Jointly supported by the U.S. Geological Survey and the U.S. Nuclear Regulatory Commission

Sizes of the Largest Possible Earthquakes in the Central and Eastern United States—Summary of a Workshop, September 8–9, 2008, Golden, Colorado

By Russell L. Wheeler

Open-File Report 2009–1263

U.S. Department of the Interior
U.S. Geological Survey

U.S. Department of the Interior
KEN SALAZAR, Secretary

U.S. Geological Survey
Marcia K. McNutt, Director

U.S. Geological Survey, Reston, Virginia: 2009

For product and ordering information:
World Wide Web: http://www.usgs.gov/pubprod
Telephone: 1-888-ASK-USGS

For more information on the USGS—the Federal source for science about the Earth, its natural and living resources, natural hazards, and the environment:
World Wide Web: http://www.usgs.gov
Telephone: 1-888-ASK-USGS

Suggested citation:
Wheeler, R.L., 2009, Sizes of the largest possible earthquakes in the Central and Eastern United States—Summary of a workshop, September 8–9, 2008, Golden, Colorado : U.S. Geological Survey Open-File Report 2009–1263, 308 p.

Contents

Tables

Conversion Factors

SI to Inch/Pound

Multiply	By	To obtain
	Length	
centimeter (cm)	0.3937	inch (in.)
gram per cubic centimeter (gm/cm^3)	62.4220	pound per cubic foot (lb/ft^3)
millimeter (mm)	0.03937	inch (in.)
millimeter per year (mm/yr)	0.03937	inch per year (in/yr)
meter (m)	3.281	foot (ft)
kilometer (km)	0.6214	mile (mi)

Sizes of the Largest Possible Earthquakes in the Central and Eastern United States: Summary of a Workshop, September 8–9, 2008, Golden, Colorado

By Russell L. Wheeler

Introduction

Most probabilistic seismic-hazard assessments require an estimate of Mmax, the magnitude M of the largest earthquake that is thought possible within a specified area. In seismically active areas such as some plate boundaries, large earthquakes occur frequently enough that Mmax might have been observed directly during the historical period. In less active regions like most of the Central and Eastern United States [(CEUS; table 2)] and adjacent Canada, large earthquakes are much less frequent and generally Mmax must be estimated indirectly. The indirect-estimation methods are many, their results vary widely, and opinions differ as to which methods are valid. This lack of consensus about Mmax estimation increases the uncertainty of hazard assessments for planned nuclear power reactors and increases design and construction costs.

Accordingly, the U.S. Geological Survey (USGS) and the U.S. Nuclear Regulatory Commission (NRC) held an open workshop on Mmax estimation in the CEUS and adjacent Canada. The workshop occurred on Monday and Tuesday, September 8 and 9, 2008, in the USGS offices in Golden, Colorado. Thirty-five people attended (table 1). The workshop goals were to reach consensus on one or more of:

(1) the relative merits of the various methods of Mmax estimation, (2) which methods are invalid, (3) which methods are promising but not yet ready for use, and (4) what research is needed to reach consensus on the values and relative importance of the individual estimation methods.

Drafts of two documents were distributed to the meeting invitees about one month before the workshop. Their purpose was to provide a common basis for discussions during the workshop. C.S. Mueller (written commun., 2008; *see* summary later of his talk on Monday morning) described the effect of changes in Mmax on computed CEUS hazard. Wheeler (2009) summarized the methods that have been used to estimate CEUS Mmax in the past, listed pros and cons of each method, and tested several methods against published Mmax estimates that were derived from paleoseismic mapping and trenching of fault scarps and liquefaction features.

Agenda

Monday, September 8, 2008

Morning

Annie Kammerer (NRC): Goals of the workshop

Rus Wheeler (USGS): Logistics, list of methods of estimating Mmax

Chuck Mueller (USGS): Sensitivity of hazard to different values of Mmax

Jon Ake (NRC): Uncertainties

Rus Wheeler (USGS): Overview of methods of estimating Mmax

Martin Chapman (Virginia Polytechnic Institute and State University): Mmax and earthquake catalogs

Kevin Coppersmith (Coppersmith Consulting, Inc.): Bayesian estimates of Mmax

Afternoon

Rus Wheeler (USGS): USGS estimates of Mmax

Walter Mooney (USGS): Mmax in cratons

Steve Harmsen (USGS): Uncertainties in parameters related to Mmax

Jon Ake (NRC): Should some of the Mmax methods no longer be used?

Tuesday, September 9, 2008

Morning

John Adams (Geological Survey of Canada): GPS constraints on hazard in eastern Canada

John Adams (Geological Survey of Canada): Mmax in Canada's fourth-generation hazard model

Trevor Allen (Geoscience Australia): Neotectonics in Australia

Rus Wheeler (USGS): What is a "tectonic analog"?

John Ebel (Boston College): Magnitude of the 1663 Charlevoix earthquake

Afternoon

Jon Ake (NRC): What are our research needs?

Annie Kammerer (NRC): What are our conclusions about methods of estimating Mmax?

Minutes of Discussions

The entire workshop was recorded and three attendees took notes at my request. The informal records allowed these edited minutes to be more complete and accurate than would have been possible otherwise. The purpose of the editing was to enhance readability, clarity, and conciseness beyond what an exact transcript could provide, while preserving speakers' meanings and speaking style, as well as most details of the discussions. Inaudible comments, repetitions, and short digressions that I deleted made up approximately 10 percent of the total material. Square brackets in the minutes enclose my and reviewers' explanatory insertions, including citations of a few pertinent reports. [During the formal review process] all speakers and discussers had the chance to correct any errors or confusion in remarks attributed to them. Their corrections are identified by double-square brackets.

Table 2 defines acronyms. Table 3 indexes the speaker of each talk and comment. Speakers mentioned numerous notable earthquakes by name or year only. Table 4 gives their dates and locations [for reference]. Speakers' Power Point slides (here called figures) are in the appendixes, and are cited in the texts of the authors' remarks. The text is mostly paraphrases and should not be treated as quotations from the speakers. If you need a quotation, please obtain it directly from the speaker.

Monday morning, September 8

Annie Kammerer (NRC): Goals of the Workshop (Appendix 1)

This workshop is important to the NRC because the ground motions of large CEUS earthquakes have long return periods, which can affect sites of interest to the NRC. At about the same time as preparations for this workshop began, another project for nuclear facilities was also being planned: the Central and Eastern U.S. Seismic Source Characterization [[(SSC; table 2)]] Project. SSC results will replace the existing hazard methodology of EPRI–SOG [(table 2)] [[(Risk Engineering, Inc., and others, 1986; Johnston and others, 1994c)]]. Results of this workshop will [be used in] the SSC project and will help it to follow the guidelines recommended by SSHAC [(Senior Seismic Hazard Analysis Committee; table 2)].

The SSHAC guidelines describe procedures for capturing the [[center, body, and]] range of legitimate opinions in the informed [[technical]] community [(Budnitz and others, 1997)]. Accordingly, NRC needs us to consider all possible methods of estimating Mmax, instead of trying to find the single best method. NRC hopes that we will be able to agree on the relative weights that should be given to the various methods of estimating Mmax (fig. A1–1). Some methods might have shown themselves to be the best ones to use now. Others might be somewhat useful, but [[are]] as mature as they are likely to get. Still others might have looked promising once, but may have fallen out of favor or may now be seen as not technically justifiable. Yet other methods might look promising now, but require additional research to develop and test them. [[Additionally,]] workshop recommendations will help NRC to decide what future [[research]] on Mmax estimation they might want to fund.

Chuck Mueller (USGS): Sensitivity of Hazard to Different Values of Mmax (Appendix 2)

The USGS has been helping the NRC to review applications for new power-plant reactors. During the reviews we compare the applicants' PSHA [(probabilistic seismic hazard assessment; table 2)] results to ours. So far all the applicants have used the EPRI–SOG methodology [dating] from the middle 1980s [(Risk Engineering Inc. and others, 1986)]. For some applications the USGS [[probabilistic ground motions are as much 2–3 times greater than the applicant's. This sort of difference leads us to ask whether the USGS results are conservative and, if so, why]].

The focus here is on Mmax, although there are other differences between the EPRI–SOG and USGS [[hazard]] models (fig. A2–1). The EPRI–SOG ESTs used several diverse methods to estimate Mmax. The diversity generated large uncertainties and broad probability distributions. At quiet sites, the centers of mass of the EPRI–SOG distributions were typically around m_b 5.5, whereas more active sites had Mmax estimates that clustered around m_b 6.5.

In contrast, the USGS model used the global earthquake catalog of Arch Johnston and colleagues, who examined worldwide tectonic analogs of the CEUS. In the CEUS craton, we set Mmax at M 7.0, and in the surrounding extended margin we took Mmax to be M 7.5. [(Frankel and others, 1996; Wheeler and Frankel, 2000; Frankel and others, 2002)]. In the 2002 national hazard maps we used only the two values. In the 2008 maps, we added probability distributions around M 7.0 and M 7.5 [(Petersen and others, 2008)]. The blue lines in figure A2–2 are Wheeler's delineation of the most northwestern Iapetan extensional faults that separate the craton from the extended margin. In the black text, parentheses indicate the weights assigned to the different magnitudes of the probability distributions.

The next two figures show two typical comparisons between the EPRI–SOG and USGS Mmax values. The [South Texas] site is seismically quiet and is in the USGS extended margin (fig. A2–3). The histogram shows that the Mmax distribution for the site from the six EPRI–SOG teams peaks in m_b 5.0–5.5. [[NRC required the applicant to update Mmax to account for recent seismicity in the Gulf of Mexico. Here the original, not the updated, EPRI–SOG Mmax distribution is used because it is typical of many quiet CEUS sites]]. The orange arrows show the four points of the USGS Mmax distribution (fig. A2–2). There are eight orange arrows because M was converted to m_b using two relations: the 1996 equation of Johnston [(Johnston, 1996)] and the 1995 equation of Atkinson and Boore [(Atkinson and Boore, 1995)]. The histogram and arrows show that the difference between the EPRI–SOG and USGS distributions is large.

The North Anna, Virginia site also is in the extended margin but it is seismically more active than the [[South Texas]] site (fig. A2–4). The EPRI–SOG distribution is centered in the magnitude range that is typical of similarly active CEUS sites. The difference between the EPRI–SOG and USGS distributions is smaller.

[[The sensitivity analysis that forms the rest of this talk examines the results of changing Mmax from the preferred USGS values of M 7.0 in the craton and M 7.5 in the extended margin. Results are shown as seismic-hazard ratio maps of the CEUS (fig. A2–5). We consider seismicity-based sources only, not faults, for the reasons given in figure A2–5. For each ratio, the denominator (“reference”) map is computed using the 2008 USGS Mmax model with distributions added to the preferred Mmax values, and the numerator map is computed using an alternative Mmax model. We consider the results of decreasing Mmax by 0.0 unit, 1.0 unit and 2.0 units from the preferred values (fig. A2–6). All other parameters except Mmax are held fixed in the hazard calculations.]]

[[Figure A2–7 shows examples of decreasing Mmax for two sites. The figure compares the alternative USGS Mmax values (arrows) to the EPRI–SOG Mmax distributions (histograms) for the North Anna and South Texas sites.]] Black boxes identify the Mmax values that we decreased, from left to right, by 2.0, 1.0, and 0.0 magnitude units. Yellow arrows show the resulting [USGS] Mmax values for the craton and orange arrows show the extended-margin values. [[Again]], each arrow is doubled because of the two magnitude-conversion relations that we used.

Note that we [don't] decrease our magnitudes more than 2.0 units, to M 5.0 in the craton and M 5.5 in the extended margin. Thus, we [don't] capture the low-magnitude tails of [[some of the EPRI–SOG Mmax distributions. We stopped at 2.0 units because the USGS hazard calculations]] use a lower cutoff value of m_b 5.0. Thus, as Mmax decreases toward the cutoff value, hazard becomes very small and a comparison to our unchanged Mmax values loses meaning.

[[Before testing Mmax sensitivity, [note that] seismicity-based sources can control the mid- to high-frequency seismic hazard in much of the CEUS (fig. A2–8). For reference, the map at the top shows seismic hazard with only contributions from seismicity sources, that is, from historical seismicity and no faults (fig. A2–5). The bottom of figure A2–8 shows a seismic hazard ratio map with only seismicity sources divided by a map with all sources, including the Meers, Cheraw, New Madrid, and Charleston faults. Both maps use the reference USGS Mmax model; this is not a test of Mmax sensitivity. Warm colors in the bottom part of the figure—hazard ratios greater than 0.65—show areas where seismicity sources dominate the hazard. These maps are for peak ground acceleration for a 2-percent-in-50-years exceedance probability. Similar maps for periods of 0.2 second and 1 s look much the same, although at 1 s the faults tend to control hazard out to greater distances than shown in figure A2–8.]]

[[Figures A2–9, –10, and –11 show results of the Mmax sensitivity analysis. [The figures show] seismic hazard ratio maps using the reference Mmax model for the denominators and an alternative Mmax model (fig. A2–6) for the numerators. Each calculation uses seismicity sources only, and each figure presents results for 0.2 s (top) and 1.0 s (bottom) period spectral acceleration. Figure A2–9 shows that decreasing Mmax by 2.0 units, which is a strong Mmax truncation, decreases hazard significantly.]] At 0.2 s, the decrease is by factors that range from roughly two, in blue areas like the southernmost CEUS, to about three, in grey and lavender areas like most of the craton (fig. A2–2). In terms of percentages, the decrease is to about 60 percent in the blue areas and to approximately 30–40 percent in the small grey and large lavender areas. At 1.0 s the hazard decreases more [[because large earthquakes tend to dominate longer-period hazard. The decrease is by factors of two or three [to 30–40 percent] in the southernmost U.S. to approximately five [to about 20 percent] in most of the craton. Complexities in the spatial patterns are caused by complex interactions between Mmax and hazard-controlling earthquakes in the PSHA.]] The table of figure A2–12 shows the decreases as percentages for several CEUS cities.

Figure A2–10 shows smaller hazard reductions when Mmax is decreased by only 1.0 unit. At 0.2 s, the decrease is to about 80 percent in most of the extended margin and to approximately 70 percent in most of the craton (fig. A2–12). At 1 s, the decrease is greater, to approximately 60 percent in the extended margin and about 50 percent in the craton (fig. A2–12).

Figure A2–11 shows the effect of applying [[the Mmax distributions of the 2008 update. The effect is small]], as also shown in the right-hand one-third of figure A2–12.

Figure A2–13 is the same as figure A2–9, with Mmax lowered by 2.0 units [[for seismicity-based sources, except that now both the numerator and the denominator maps include all the sources.]] Hazard from the Meers and Cheraw faults and the large, repeating New Madrid and Charleston earthquakes are added back in. These maps may be a more realistic depiction of the effect of Mmax than figure A2–9 because near the New Madrid and Charleston seismic zones the hazard is dominated by the large repeating earthquakes[[, which are modeled as specific sources by both the NRC applicants and the USGS]].

Jon Ake (NRC): Uncertainties (Appendix 3)

Probably everyone in the room is aware of the need for a PSHA to include uncertainty and describe it numerically (fig. A3–1). Not only is addressing uncertainty good practice, but also regulations have required it since the middle to late 1990s (fig. A3–2).

As most of us are aware, there are two main types of uncertainty that have different properties and are treated in different ways (fig. A3–3). Aleatory variability is random and irreducible.
For example, in ground-motion studies we might ask, if a magnitude six and one-half earthquake occurs at ten kilometers' distance, what is the natural variability of the observed ground motion at a site?

The second source of uncertainty is what we call knowledge-based or epistemic uncertainty. It stems from incomplete knowledge, and it is potentially reducible. We often have unknown models and unknown parameters within those models, or competing models. For example, as Chuck [Mueller] mentioned, there are two different relations between m_{bLg} and moment magnitude. We generally incorporate epistemic uncertainty with logic trees. One of the things that Annie [Kammerer] and I would like to [see] come out of this workshop is identification of some of the epistemic uncertainties in estimating CEUS Mmax. We'd like to identify research that could potentially reduce the uncertainty in Mmax, or reduce the field of methods for estimating Mmax. We'd like to get a hierarchical listing of what research might be done in the next ten years or so. As we incorporate epistemic uncertainty into our models as logic trees, we are forced to wrestle with the problem that we don't know what we don't know.

Here are a couple of examples of how we think about epistemic uncertainty in the problem of estimating Mmax (fig. A3–4). If an earthquake produces a surface rupture, the relations of Wells and Coppersmith [(Wells and Coppersmith, 1994)] let us estimate the magnitude either from the length of the resulting scarp or from the size of the resulting displacement. These two competing models represent an epistemic uncertainty. However, each model has its own estimate of variability that needs to be carried through the analysis. Additionally, Kevin [Coppersmith] and Don [Wells] were trading space for time. For example, in estimating the amount of displacement on a fault that would result from an M 6.5, we don't usually have multiple occurrences of an M 6.5 on that fault. We are forced to draw in information from other M 6.5 earthquakes elsewhere. Additionally, there could be other relations than those of Wells and Coppersmith that could be included in the analysis, which indicates that the branches of the epistemic logic tree could increase rapidly. More commonly in the CEUS, we try to estimate Mmax from paleoliquefaction, or from either maximum intensity or the area within one or more isoseismals. The methods of estimating magnitude from paleoliquefaction or intensity involve a number of uncertainties that need to be taken into account.

These considerations have several implications (fig. A3–5). Pragmatically, the USGS and people who are designing power plants, liquefied natural-gas facilities, and the like must make hazard assessments by fixed deadlines while knowledge is still incomplete. This problem is especially acute in stable continental regions like the CEUS because they are data deficient. This means that substituting space for time in some way is required. To try to estimate the necessary parameters from the geological and geophysical properties of a small source is, in my view, impractical. The uncertainties would be almost insurmountable.

To summarize, we need a way to formally include uncertainty and variability in our assessments if we are going to make our results useful to the community at large, especially for assessments for critical facilities (fig. A3–6). Additionally, we need a well-defined process for including and tracking aleatory and epistemic uncertainties without double-counting any uncertainty. Finally, to keep the problem tractable we need to define a minimum subset of credible alternative models. There must be a sound technical basis for each model. Otherwise, the epistemic logic tree will quickly add branches [that might not improve the assessments].

Discussion

Salomone: There are two groups of methods, the more physical and the more mathematical. Physical evidence of Mmax in a source zone includes seismicity, rates, and geologic features. Other solutions could be more mathematical and could derive Mmax. From the viewpoint of uncertainty, which group is likely to have the greater promise of capturing uncertainty without producing a large logic tree [[that is not practical]]?

Ake: At present I'm not sure I could answer that. Perhaps the next year and one-half will tell us.

Mooney: How do you formally assess uncertainty when you substitute space for time?

Ake: You have to decide that the substitution is the appropriate one. For the example of ground motion, when we think that we are using the appropriate magnitudes, distances, site conditions, tectonic settings and so forth to select earthquakes to group together, we are doing a similar thing to what Kevin [Coppersmith] and Don [Wells] did in pooling information to set up their regression equations.

Mooney: If I substitute information from India or Australia, which have a lot of earthquakes, for [missing] North American information, it may not be a very good substitution because India and

Australia are fast-moving plates that are both subducting at their northern ends. There may be systematic differences that are difficult to include in the uncertainty.

Ake: That's where we get into the subjective evaluation of the applicability of different domains. In the work that Arch [Johnston] did, a lot of time was spent examining the applicability of tectonic regions of the world as analogs of the CEUS [(Johnston and others, 1994c)].

Johnston: Walter [Mooney] has a good point. You certainly introduce more uncertainty by including global areas. But you gain more than you lose. That's where the judgment comes in.

Coppersmith: The ergodic assumption, which this really is, is controversial in some places. At Yucca Mountain, for instance, some would argue that you shouldn't import large ground motions into their regressions because they have seen precarious rocks that have not been toppled.

Ake: Sometimes this leads to more branches on the epistemic logic tree, to incorporate including those data instead of excluding them.

Coppersmith: Some in the uncertainty community maintain that logic trees must be collectively exhaustive and mutually exclusive. Sometimes we have models that overlap and are not mutually exclusive. How can we claim that we have all the models? Some purists prefer that we call them weighted alternatives, but not logic trees.

Frankel: Part of the problem might be sloppy terminology. The ergodic assumption is that the spatial variability of some quantity is the same as the temporal variability of that quantity. That's not what we are talking about. We are talking about estimating Mmax for a region, and we are looking at [estimates of] Mmax everyplace else. We are not trying to characterize the variability of Mmax in time in a given location. It is not analogous to the ground-motion example.

Rus Wheeler (USGS): Overview of Methods of Estimating Mmax (Appendix 4)

I'm going to give you an overview of the various methods of estimating Mmax. If anyone wants to describe any method in more detail than I will, please come up and do so. This summary is largely of section 3 of the foundation document that I circulated to you before the workshop [(Wheeler, 2009)]. Section 3 is summarized itself in the document's table 5, and table 5 is boiled down here to figure A4–1. The figure lists 13 methods or groups of methods, some of which Chuck Mueller also listed in his talk.

The figure groups documents by common disadvantages, with one disadvantage per column. Focusing on the disadvantages might be a more effective way of eliminating some methods than focusing on the advantages. Additionally, all methods have advantages and disadvantages, but it is the disadvantages that we are trying to overcome.

During the early and middle 1980s, when everything was expressed in m_b, two of the EPRI–SOG teams estimated Mmax as 7.5 if they couldn't use any other method, because that is where m_b saturates. That was reasonable at the time. Now the method is irrelevant because we all use moment [magnitude] for large and moderate earthquakes. If we need m_b for computations, we convert moment into m_b. I won't mention this method again.

Several of the methods suffer from a small-sample problem. The table identifies some of them with the notation "SS(src zone)" in the second column. Jon [Ake] referred to the small-sample problem when he concluded that it forces us to use global analogs. The small-sample problem arises in small source-zones east of the Rocky Mountains because, with a few exceptions, small source-zones contain too few earthquakes to usefully constrain Mmax. Also, most small source-zones do not contain enough different geologic settings, or enough examples of any single geologic setting, to show which settings are associated with large earthquakes. The small-sample problem, I think, strongly undercuts some of the methods. In the first line of the table (fig. A4–1), Mobs is the largest historical magnitude in the source zone being considered [(table 2)]. Chuck Mueller mentioned that some workers add a constant to Mobs, presumably the same constant over a group of source zones or perhaps over the entire study area. The M-f method [(table 2)] consists of making a magnitude-frequency graph and extrapolating it out to any desired recurrence interval. Mmax is estimated as the magnitude that has this recurrence interval. However, determining which recurrence interval to use requires additional information.

One EPRI–SOG [Earth-science] team and one Lawrence Livermore expert took the first step in trading space for time to mitigate the small-sample problem. They examined the North American SCR [(stable continental region; table 2)], thereby avoiding possible differences between SCRs. It's a large SCR, and one of the most seismically active. It might contain enough different geologic settings and enough examples of each setting to indicate which settings tend to have larger earthquakes. The next step is to go global, with global analogs or Bayesian analysis. Both methods use the Earth's entire historical sample of SCR earthquakes.

Toward the bottom of the table, I grouped together all the physics-based methods, following Kevin [Coppersmith] and Chinnery [(Chinnery, 1979; Coppersmith and others, 1987)]. Similarly, all statistical methods are grouped together, although I chose to break out pattern recognition [and Bayesian analysis]. Finally, I recently found a 1998 paper about China that reported an inverse relation between the value of Q_0 and the magnitude of the largest historical earthquake.

The foundation document [(Wheeler, 2009)] explains how the small-sample problem afflicts a lot of these methods. The problem is not necessarily fatal to a method, but it must be addressed by anyone who wants to use one of the methods that have the problem.

The third column in the table points out an unanticipated problem that the foundation document describes at length. Estimating Mmax from seismicity rates turns out to be insensitive to Mmax, [when the rate estimate is tested against] Mmax [as] estimated from paleoseismic results. Paleoseismic results give us the fourth column. These results are a valuable kind of information that we did not have a decade or two ago. The first reports I've been able to find on CEUS paleoseismic results are Steve Obermeier's on the area around Charleston, South Carolina [(Obermeier and others, 1985a; Obermeier and others, 1985b)]. He didn't have enough data for an authoritative paper for another half decade or so. This kind of work goes slowly. The EPRI–SOG and Lawrence Livermore personnel did not have access to these paleoseismic results and that severely limited the methods that they could rule out. The methods of Mobs, Mobs plus a constant, M-f extrapolation, and Q_0 give results that, however reasonable the methods might seem, are inconsistent with the paleoseismic results that we now have. The results are few. Five CEUS areas or faults have been studied enough that each of them has a chronology of prehistoric earthquakes of at least approximately M 7. The magnitudes have the kinds of uncertainties that Jon [Ake] was discussing. The chronologies are in hand for New Madrid, Charleston, the Wabash Valley [the southern half of Illinois and the southwest quarter of Indiana], and the Meers and Cheraw faults. The tests of estimation methods against paleoseismic findings are described in detail in the foundation document [(Wheeler, 2009)]. Each of the three tested methods near the top of the table under estimates the sizes of the largest earthquakes in the five historical and paleoseismic records by 1–3 magnitude units. Thus, for at least some of these 13 methods, I think that the test against paleoseismic methods is strong. The uncertainties in the paleoseismic magnitudes can be large, but paleoseismologists have developed a rough feeling for how large these uncertainties are, even if they can't put a number on them yet.

The disadvantage of the methods that use continental or global analogs is our poor understanding of what constitutes an analog. Some of the discussion this morning has already touched on this problem.

Figure A4–2 summarizes what I've said. Probably most of us would agree that our ultimate goal is a physical analysis that starts with the measured properties of whatever source zone we specify. Then, from that information alone, with no assumptions that most of us would not accept, the analysis would produce a value of Mmax and its uncertainty. We are not close to that goal, but we need to keep focused on it. I'm not sure how to solve the small-sample problem. The obvious answer, "wait a thousand years", is not helpful. I've already described the valuable tool that paleoseismic results give us. One problem with this tool is that we have good paleoseismic results from only five areas. That number will increase, but slowly. Paleoseismic work that produces results that we can use is labor-intensive. Thus, the work is funding-limited. In the next 5–10 years we might expect to add a few areas to the present five. We might add another 5–10 areas, but not another 50. The sample of paleoseismically-characterized areas will probably increase too slowly to be of use to us in the next few years.

In contrast, the problem of defining analog areas may be more tractable. This may be largely an intellectual problem, which we might be able to solve by thinking logically about what we mean by analog, and, just as importantly, what signal would tell us that an area is not an analog to the one we are considering. Tomorrow morning I'll suggest a straw man definition of analog.

Discussion

Hough: I have two comments instead of questions. One is that we need to explain what we mean by Mmax. Are we are talking about the maximum observed magnitude, or something else? The other comment is that it is sobering to keep in mind that, although we talk about physics-based models, a lot of these debates are still raging in California. The historical record is two orders of magnitude longer than the repeat times, and we still can't answer these questions.

Kimball: I think you're falling into a trap that my 30 years of experience warns me we have to be careful of. That's concluding that the paleoseismic estimates of Mmax are uncertain, but not that much. Paleoseismic information is relatively new. It's only been in the last 10–15 years that we've been getting a lot of information. It's not a simple job to unravel Mmax from paleoseismic data. It takes an integrated approach by geologists, seismologists, and geotechs. I claim that it's one of the most uncertain subsets of our information. Time and time again in PSHA we find that we under appreciate [the uncertainties in] the things we think we know the most about.

<u>Wheeler:</u> We tend to consider New Madrid 1811–1812, whatever their magnitudes, and Charleston 1886, whatever its magnitude, as Mmax because the prehistoric earthquakes appear to have caused similar abundances, sizes, and geographic distributions of liquefaction features. What that says about magnitude, I'm not sure. A still newer method is geotechnical measurements that estimate strength of shaking and, from that, the magnitude of the causal earthquake. It's only been applied in a couple of places. Again, it's slow work. But it's a completely independent approach from any that we've used before. I'd very much like to see this applied to New Madrid, for instance.

<u>Kimball:</u> There's site response, too. We've seen how the site-response estimates for New Madrid have changed in the past 15 years as people accumulated more and more data. Site response data have been very important to that increase in understanding, and I think they are going to be just as important to backing out Mmax.

<u>Salomone:</u> On the uncertainty with respect to back-calculating, and from a geotechnical perspective, the cyclic strength that [[caused]] liquefaction is a key value. Also, care [[must]] be taken to include factors like the aging effect of soils.

<u>Wheeler:</u> What Larry [Salomone] is referring to is even more general. In the Wabash Valley, a series of four [or more] papers were published over five to ten years by [Steve] Obermeier, Russell Green, Scott Olson, and [Eric Pond]. Each paper adds quantitative consideration of another

complicating factor, like cyclic stress [and aging]. Eventually they will come to the end of the factors that are worth including, but it's not clear that they have yet.

Johnston: One thing that we didn't have until recently is the 2001 Bhuj earthquake. It gives us calibration of a large upper magnitude-seven earthquake with the distribution and size of liquefaction features that we never had before. We can improve on the crude estimates that we had.

Lettis: In the future we may be able to use the absence of paleoliquefaction as a ceiling on Mmax in a particular area, especially if the Mmax that we would be estimating has a recurrence of 10,000 years or less. We should see it [the liquefaction record of a large earthquake] in susceptible deposits. If enough work is done and we don't see it, we could be able to bring down the threshold of Mmax values that could be hiding there.

Wheeler: What Bill [Lettis] is referring to is a technique that two people have used in a handful of places. Steve Obermeier has done it in a couple of areas; Tish Tuttle has taken this approach in at least one area, in western New York State; and I'm probably not remembering one or two more. The idea is that you survey an area, so this is time intensive. It's probably more time intensive than just looking for liquefaction because you have to guarantee that you have sufficient coverage. You canoe streams and otherwise search the area for liquefaction features. The idea is to get a dense enough coverage of exposures, and a long enough age range of deposits in those exposures, that you can say this: if a magnitude, say, seven [earthquake] had occurred here within the time spanned by the liquefiable exposures, its liquefaction field should have been so wide that it could not be hidden between the streams that I surveyed. Therefore, an earthquake of M7 or larger did not occur within the time span sampled. That's how Obermeier was able to say that the Central Virginia seismic zone has not had an earthquake of M 7 or larger within the last 2,000 years in one part of the zone, and within the last 5,000 years in the rest of the zone.

Hough: Geologists looked at the Wabash Valley and said there was no liquefaction there in 1811–1812, and yet there is a historical account that says there were wagonloads of sand blows. You can argue about whether that was a New Madrid event or a local event, but it's a pretty clear account that there was significant liquefaction.

Murphy: We've been talking about techniques like paleoseismology. I'm not sure how a technique falls into the categorization of 13 methods. Also, we have techniques that we are using with a method, say global analogs. How do we associate the technique with the quality of the analog?

Wheeler: I view paleoseismology as a way to collect or build data. Using paleoseismology to estimate Mmax is an application of the technique.

Adams: When we are looking at global analogs, we shouldn't be dealing just with historical earthquakes, but also with prehistoric ones, or paleoseismological results. I'm looking forward to the talk about Dan Clark's Australian work, if Trevor [Allen] wants to give it. The second thing is that you probably fell into the trap of putting up methods but leaving off other ones. The big one that I think is missing is using geodetic strain rate to constrain Mmax in eastern North America. Mazzotti of the Geological Survey of Canada has done something like that, and I'll show a couple of [his] slides tomorrow. It's very uncertain, but in some sense the top end of the uncertainty puts an upper limit on Mmax, as long as you believe that the [geodetic] strain rate is actually the long-term rate.

Break

Annie Kammerer (NRC): An Invitation

I'd like to reiterate something that was said earlier. As we developed the agenda we were partly influenced by the needs of the CEUS SSC project that I mentioned at the start of this morning. The SSC project is going on right now under the SSHAC guidelines and the SSC will replace the EPRI–SOG database. We were thinking of this part of the workshop as a chance for proponents of each of the methods of estimating Mmax to discuss the strengths, weaknesses, and uncertainties of the different methods. We have fewer talks by proponents of individual methods than we wanted. We'd like to encourage any of you who think that some methods have been missed or incorrectly characterized to please let Rus [Wheeler] know that you'd like to talk about them, however informally. It would be best to do this as early as possible so that we have a strong basis for later discussions of all the methods.

Martin Chapman (Virginia Polytechnic Institute and State University): Mmax and Earthquake Catalogs (Appendix 5)

I'd like to look at the maximum catalog magnitude and what it may or may not tell us about Mmax (fig. A5–1). Let's look at the whole Earth to start with (fig. A5–2). The red dots are from the NEIC PDE [(National Earthquake Information Center's Preliminary Determination of Epicenters catalog; table 2).] The black dots from magnitude eight and one-half up to nine and one-half are from the NEIC's listing of significant earthquakes. The [regression] line I've put through the red dots has a well established b-value of –1, and an a-value of 8.13. The black dots for the bigger earthquakes trail off downward. The one at the end at about magnitude nine and one-half is the Chilean earthquake of 1960 [table 4]. We don't know how big it really was because we didn't have the instruments at the time to see the low-frequency component. It may have had a huge amount of aseismic slip, even preceding what was recorded at the long-period stations that were in operation at the time.

We have a candidate for a maximum earthquake because it is so outstandingly large compared to the other earthquakes that we've seen since 1900. The question I want to address here is whether the absence of events larger than nine and one-half in the global catalog is significant (fig. A5–3). In our case, the question is whether the maximum catalog earthquake is a useful estimate of Mmax. The context we'll put it in is a Poisson process. That's the standard model for a probabilistic hazard analysis. We'll assume a Gutenberg-Richter relationship.

The mean rate of earthquakes greater than magnitude nine and one-half is 0.0427 per year. That amounts to a return period for magnitude nine and one-half or greater of about 23 years. The Poisson process allows us to calculate the probability of not seeing any earthquakes [larger than M 9.5] in that time interval. Multiply the rate by the exposure period and exponentiate [the negative of] it. When we do that we find that we have a 1 percent probability of not seeing any earthquakes bigger than nine and one-half in the catalog [if bigger earthquakes are possible].

Now, it may be that a bigger earthquake than the Chilean earthquake may have slipped into the catalog. We didn't have seismographs operational prior to the middle 1950s that could have told us just how big an earthquake it could have been. If we decrease our time interval to 50 years, we see that we've got an 11 percent probability of not seeing any earthquake bigger than magnitude nine and one-half. I will argue here that both these numbers are small. A probability of 1 percent and a probability of about 10 percent would suggest that there's a pretty good possibility that the Chilean earthquake is the biggest earthquake that we'll ever see. In this case the maximum magnitude in our catalog may be giving us a useful estimate of Mmax.

Can we do the same sort of thing for a [smaller] region (fig. A5–4)? That's going to depend on the rate of the earthquakes and the length of time over which we have a complete catalog. I took a look at the southeastern U.S. to see if we could do something like that (fig. A5–5). We'll look at [the large cluster of earthquakes in and around] eastern Tennessee. Here's the catalog (fig. A5–6). I'd say that we're complete from magnitude 2 up to about 3.4 back to about 1980. Magnitudes 4 to about 4.4 are complete back to about the 1870s, and for the larger magnitudes we might be complete back to the start of the catalog, which is about 1840. Using those completeness times and rates, I did a regression model and got [and a-value of] 3.23 with a b-value of [-1.07] (fig. A5–7).

We'll do an analogy to what we did for the whole Earth (fig. A5–8). For an exposure period of 138 years, back to 1870, there's a 35 percent probability of seeing what we've actually seen, that is, no earthquakes greater than magnitude 5.0. If we are conservative and say that our catalog is complete back to 1840, we'll still have about a 30 percent probability of seeing what we've seen. How long would the catalog have to be for eastern Tennessee before we'd get down to a probability of about 0.1, which we found for the whole Earth (fig. A5–3)? We'd have to have a catalog that goes back about 300 years. The earliest earthquake that we've got in the whole region was in 1777.

I'm not saying that the maximum earthquake in the eastern Tennessee catalog tells us nothing about Mmax. In fact, it does tell us something about Mmax (fig. A5–9). If it's all you've got to work with, it's your best estimate. But unfortunately, unless the rate of the maximum catalog earthquakes, multiplied by the time over which the catalog is complete, is significantly large, such that its [negative] exponentiation is on the order of 0.1 or less, then the estimate may not be very reliable.

Discussion

Perkins: Magnitude 5 is the minimum magnitude for the hazard calculations. If you can't say anything definite about whether Mmax exceeds 5.0, then that might be information, but it's useless.

Chapman: The fact that you don't see anything bigger than a certain magnitude in your catalog does provide a little bit of constraint on your effort to estimate what Mmax is.

Perkins: If we can't say anything about any magnitude larger than 5 then I don't see how we have a constraint.

Chapman: For the whole Earth, if the probability is very small, then that gives you some suggestion that Mmax is not much bigger than nine and one-half. But you're correct about eastern

Tennessee, where there's a 30 percent probability that what we've seen historically is what we'd expect [even if larger earthquakes than we've seen are possible]. You have to look at it like this—the catalog can give you some information if you're in an active area, or if you have a very long catalog.

Hough: Does the global calculation assume a b-value of one?

Chapman: It assumes it, but as you saw, one fits the data.

Hough: So if you instead assumed a truncated version, then presumably the probability of a magnitude nine and one-half being the largest would decrease.

Chapman: The truncation is what we're trying to get a handle on, but where would we truncate? It's unlikely that the real Earth [follows a b-value of one] and then just stops. It probably has some kind of a functional truncation, but I didn't try to model that. I just assumed that the line went to infinity.

Adams: I may have missed something here. You say one magnitude nine and one-half after every 23 years, but the probability of one larger in 100 years is only 1 percent. Those seem contradictory. [Chapman said that the probability of no earthquakes larger than nine and one-half in 100 years is only 1 percent if bigger earthquakes are possible.]

Chapman: Twenty three years is the return period. That would be a 63 percent probability.

Adams: But we've had quite a few 23–year periods and we haven't seen anything bigger. Is that the basis for saying we can't have anything bigger? But we haven't had many earthquakes even close to nine and one-half, right? So the prediction that we'll have a nine and one-half every 23 years seems to be contradicted by the historical record.

Chapman: Yes, that's right. It [23 years] is underpredicted. We've seen fewer magnitude nine and one-half [earthquakes] than the data between magnitudes six and eight would predict if you assume a linear recurrence relation and an infinite maximum magnitude.

Adams: And that leads you to the truncation.

Chapman: Yes, that leads me to the conclusion that we might be seeing something like a real Mmax for the globe somewhere in the range of nine and one-half to nine and three-quarters, or maybe ten.

Kimball: Shouldn't we find the longest catalog that we can, and look at fifty-year windows in that catalog to apply this kind of thought process? Perhaps [the catalog] from China?

Chapman: Certainly we need to make some use of the fact that the maximum catalog magnitude does provide some information on Mmax. It will only be useful where we have a long catalog and/or a high activity rate. For the whole Earth, we have both.

Adams: A prime data set, then, would be the western Australian one, which is a long one.

Chapman: Even there you may run into problems. The longer the catalog the better.

Harmsen: I did a study similar to yours. I looked at the majority of the Central and Eastern United States and considered magnitudes greater than or equal to 6. We might like to talk about it at some point and get a little more information about this methodology.

Chapman: Absolutely. If I did include the whole southeast, the conclusions might have been different. If you did that you'd run into the dilemma of the Charleston earthquake [table 4]. The historical record has a big gap between the maximum [observed] earthquakes in all the other seismic zones in the southeast and in the Charleston area. The fact that we don't have any earthquakes in the southeast between magnitude 5.8 and 6.7 could also be because of chance and a short sample duration.

Ake: I have a short question in the context of looking at a more extensive historical catalog. Isn't that one of the things that was done in the Swiss study? [Didn't they] look at catalogs that in parts of central Europe go back to the [A.D.] 1400s or 1300s? The report explains that some of the experts went through an exercise similar to what Martin was talking about.

Coppersmith: Particularly [that's true of the catalog of] the Swiss Seismological Service. The catalog is complete for large magnitudes back 700 years or so. For moderate earthquakes [it is complete back about] 300 years. By the way, the technique that I'll talk about next was used, but they [the Swiss] are enamored with paleoseismic information. They are looking at delicate cave structures, lake landslides, liquefaction, [and] all types of things, which gives them a longer catalog still. If you deal with hundreds of years, it's still not very long relative to rare events. If you could make it thousands of years with the paleoseismic record, that would be useful.

Kevin Coppersmith (Coppersmith Consulting, Inc.): Bayesian Estimates of Mmax (Appendix 6)

Bob Youngs is sick but I'll give his talk (fig. A6–1). This work on understanding maximum earthquakes really began as the EPRI–SOG teams were developing their assessments in the early 1980s. The problem of estimating Mmax is very difficult. You deliberately limit yourself to a small source zone, where you've seen no moderate events or only a few, and you're trying to assess how large they can get. It became clear that we needed to do something to help the teams. In June, 1985 a seminar reviewed methods that were used then for estimating intraplate Mmax on a regionally varying basis and evaluated each method's strengths and weaknesses [(Coppersmith and others, 1987)]. How far have we actually come from the time when this was envisioned?

The approach that I'll be talking about is based on a notion that was almost anecdotal at the time. Arch [Johnston] and some of his colleagues looked at the largest earthquakes that had occurred within regions that he considered to be analogous [to the CEUS], and came up with the definition of stable continental regions. The biggest earthquakes seemed to be occurring in [certain] types of environments, and that's what was discussed with the EPRI–SOG teams. They seemed to be in extended crust, they seemed to be in old rifts of certain ages, and so on. It turns out that the anecdotal part is very difficult to prove statistically. It turns out that there are many rifts or extended margins that haven't had large (M 7 or larger) earthquakes. Ultimately we agreed that we can't wait around a thousand years, and we need to look spatially somewhere else.

You do that by first defining where you are going to look. We asked Arch to help us on a study of maximum earthquakes that we expected to last about six months. You [Arch] were funded on the basis of a year and did about five years of work. First, it was a huge amount of work to define what is meant by a stable continental region [(Johnston, 1989, 1994b; Kanter, 1994b)]. Second, you had to pull together a catalog of magnitude four and one-half or larger earthquakes from all over the world and look systematically at moment magnitudes and the associations of the earthquakes with certain types of geologic features [(Johnston, 1994a; Johnston, 1994c)].

The goal was to go beyond the anecdotal evidence and, first, to be able to say that we can uniquely associate large-magnitude earthquakes with certain types of geologic features. Second, if we can make those association, then can we say if we capture the maximum earthquake that can occur globally in a certain tectonic setting, then [the earthquake] represents a maximum earthquake for a source zone, say in [a similar tectonic setting in] Nebraska? A lot of assumptions go into that [reasoning]. We also said that the longest record that exists is the global record because at that time the paleoseismic record didn't exist. That global record may be close to giving us the maximum earthquake, but we needed to estimate the maximum earthquake in a way that would capture its uncertainty. So we enlisted Bob [Youngs] and Allin Cornell and his student Steve Winterstein to do the calculations for a quantitative assessment of the anecdotal notion. The bottom line is that the approach I'll outline still has a lot of uncertainty, and still has a lot of assumptions, but provides a more formal way of informing an assessment [(Coppersmith, 1994a; Cornell, 1994)].

The basic notion in all this is that earthquake-size distribution is a truncated exponential distribution between the minimum magnitude and the upper-bound magnitude (fig. A6–2). The likelihood of seeing this maximum earthquake goes up with the number of observations, [which we call] N. If we have a large number of [observed] earthquakes in our source zone, then the likelihood of having captured the maximum earthquake goes up. When we deal with the Central and Eastern U.S. the number of large earthquakes stays very low. Suppose we have a source zone with a maximum observed earthquake of M 6 (fig. A6–3). The larger the N, the more the sample likelihood function favors the maximum observed earthquake as Mmax. If you have only one earthquake, the likelihood function is almost flat above 6. If you have ten, it begins to peak [at M 6]. If you have 100, it gets even more peaked. However, the likelihood function has no upper bound (fig. A6–4). It integrates to infinity and doesn't uniquely define an upper bound. Because of that, we have to have some type of prior distribution.

Arch had broken out all these earthquakes into many tectonic categories—age of rifting, orientation relative to local stress regime, type of stress, evidence of local reactivation, how many reactivations—all the things that you would think physically are probably important to the problem. That only cut the sample size down smaller and smaller and made it more difficult to prove anything statistically. We ended up with a subdivision that was very fine-grained, with very poor statistics. Then we pooled earthquakes back together to improve the statistics a bit.

How did we do this (fig. A6–5)? The first step was to define stable continental regions and break them into domains based on things that may be important to the maximum earthquake problem [(Kanter, 1994b)]. We feel in fact that some things that really control the maximum earthquake have to do with the development and rupture of faults. Earthquakes occur on faults, so if we can limit the dimensions of the ruptures, we can limit Mmax. But where are the faults in the Central and Eastern U.S.? If you think we don't have enough data now, go back 25 years. The Cheraw [fault] wasn't studied. The Meers [fault] wasn't studied. We didn't feel that we had a basis for limiting Mmax [with] the things that you would think would be the best. So the main idea was to look at the maximum observed earthquakes that occurred around the world in stable continental regions, and to try to subdivide the SCRs in some way.

The element we want is the maximum magnitude (fig. A6–6). We have a maximum observed magnitude for a domain, and that becomes a median estimate. Then we need to adjust [correct] that on the basis of the number of samples [earthquakes] that we have, [whether that] be 1, 10, or 100, to assess how well that median gives an assessment of the true upper bound. This is the type of correction that is carried out (fig. A6–7) [(Cornell, 1994)]. We can make an estimate, by using an exponential distribution, of the median value. If the maximum observed [magnitude] is 5.7 ["Example" in the figure] and we have 10 earthquakes, the median [vertical axis] needs to be corrected to get to the true upper bound magnitude [horizontal axis]. You can see that the curves flatten downward. With many earthquakes [upper curves], the median estimate is close to the maximum. With few earthquakes, the curves become almost horizontal [lower curves], to the point where you have a very large bias correction to move from the maximum observed to the true upper bound magnitude.

How can we get a better prior distribution (fig. A6–8)? After several attempts to use the subdivisions of tectonic features, it became clear that only through domain pooling would you get any statistical significance at all. We went first to what Arch said, to extended crust, or Mesozoic rifting, or rifts that had multiple reactivations, to see whether those areas have a significantly different maximum observed earthquake [than the non-extended crust]. In fact, at the mean level of the distributions, they [the two distributions for the extended and non-extended crust] were essentially the same. The only thing that was really different was a much broader range of variability in the distribution [for the extended crust]. Ultimately, pooling the domains into just these [two] categories was the only way we could show any statistical significance.

This is the process of developing a prior distribution (fig. A6–9). We first compute the statistics for extended crust and we get a mean of the distribution (actually it's a median) of about 6 and a large standard deviation. Then we do the sample size adjustment [correction], which allows us to get a better handle on the upper bound magnitude, and we are dealing with a mean upper-bound magnitude for extended crust of M 6.4. These are all moment magnitudes, which is one of the great contributions of the work that Arch did. But look at the standard deviation for extended crust. It's very large. The mean for non-extended crust is very similar but with a smaller standard deviation. This becomes the prior. A number of dummy variables were used (fig. A6–10) to include variables that Arch and everyone else thought to be important. None of those helped. We ended up with low r-squared values, [which are] not very effective in explaining the variations [(Cornell, 1994)].

The way [this approach that I've been describing] is used in an application for seismic hazard analysis is to look first at the prior distribution [upper graph on the left of fig. A6–11] that comes from a particular type of SCR crust. This example is for extended crust. For that particular [type of] source, we develop a [likelihood] function [middle graph on the left] that takes into account the number of earthquakes [larger than the minimum or cutoff magnitude and] smaller than the maximum observed earthquake, in this case an M 5. The more earthquakes you have, the more peaked the likelihood function becomes. Then you multiply the prior by the likelihood function to get the posterior function [lower graph on the left], and that becomes the maximum magnitude distribution [histogram on the right]. It takes into account the prior information from the global data, and source-specific information that's been actually observed in the source zone.

When this came out after EPRI–SOG was finished, distributions like these, which went from magnitude 5 to magnitude 8.5, were met with shock and awe in the PSHA community. The argument was made, and will continue to be made [for the next] 25 years, as to whether or not it is appropriate to have this kind of space for time substitution. I would argue that the Mmax distribution for a source zone in Nebraska is very uncertain. We might have been kidding ourselves 25 years ago to say that it [Mmax] was five to five and one-half. In fact, we would need a basis [to justify] excluding similar types of distributions from our source zone of interest.

One of the things that Bob [Youngs] mentioned in his wrap up is that the Bayesian approach is not limited to priors that result from the global data set (fig. A6–12). For example, in the report [(Coppersmith, 1994a)] we used Charleston [South Carolina] as an example of this approach, starting with a prior that comes directly from the global data set, and then modify[ing] that prior in the light of other information. We used the example at that time of the length of rupture that might have been associated with the Ashley River fault and some other features. If those provide a basis for modifying our prior, then we can expect a magnitude 8 earthquake simply because of the dimensions associated with those features. We can modify the prior distribution up front so it doesn't have a massive standard deviation, particularly in the upper tail. That [is an example of developing the prior by] incorporating more source-specific information [instead of] relying only on imported earthquakes from the rest of the world. That is an area for future work. A second area, and this may be going on for the SSC project, is to update the information included in the regression to see whether we would have some better statistics associated with the regression and a better basis for the prior distribution.

Allin [Cornell] went to great lengths in section five of [Johnston and others (1994c)] to point out that, even though the statistics are poor, we can disprove something [(Cornell, 1994)]. [Even so,] the first reaction of some people when this [concept of using global analogs] first came out was to say, "if that's the case, then the maximum earthquake everywhere is [magnitude] 8.3. It's the largest observed anywhere in an SCR area." However, he [Cornell] points out that you can disprove that statistically, certainly within non-extended crust. In fact, those sorts of [objections] were brought up when this [discussion on the use of global analogs] first came out. Now these are distributions, and we can make arguments that the tails of these distributions can be truncated meaningfully. The arguments that Arch made in the report [(Johnston and others (1994c)] that the seven or eight or nine largest earthquakes all have occurred in extended crust, they've all been compressional, and a number of other conclusions—none of those were violated by the arguments that were made with statistics. It's just that we can [also] turn around and ask, given the extended crust, and stress regime, and everything else, what is the maximum earthquake? That argument is a different one than the one that we were making in the report [Johnston and others (1994c)].

Discussion

Johnston: I might add that, of the 15 magnitude 7.0 or greater [earthquakes] that we'd had, the most recent one was 1951 [South of Tasmania; table 4]. It preceded the WWSSN [(Worldwide Standard Seismograph Network)]. Even though [that] was an instrumental earthquake, it was a poor one. And that held true up until Bhuj in 2001.

Lettis: In this report, one uncertainty that would be nice to clear up is the definition of extended crust. Is it Mesozoic and younger extended crust, or does it extend back to Paleozoic extended crust when we are looking at these larger events?

Wheeler: Arch [Johnston] has one paragraph, noting that none of those 15 [earthquakes of M 7.0 or larger] occurred in Paleozoic extended crust, but he was sure it had the potential [to make big earthquakes] [(Johnston, 1994b, p. 4–8)]. Probably we should not expect [to have seen] a [magnitude larger than] 7 in Paleozoic extended crust, because there is so little area of it [, even if an M larger than 7.0 were possible there].

Coppersmith: According to this [(Johnston and others, 1994)], it [the prior distribution] does include Cenozoic, Mesozoic, Paleozoic, and Precambrian. One of the things that was looked at is reactivation, and that's one of the things that was thought to be very important. It was used as a descriptor variable and couldn't be proven to be significant. So we ended up putting in all ages and all numbers of reactivations to get a sample size that we could use.

Mooney: Arch, what would you think of this? You had two types of crust, but I think it's reasonable to have four. Two would be continental margins, or rifted margins, which border oceans as extended terranes, and continental interior rifts, like [those under] New Madrid and [the] Wabash [Valley]. That distinguishes geologically Charleston [South Carolina, in a continental margin] from the continental interior rifts. The third category that I think we should consider is sutures, where continents have been joined together. The Appalachian [Mountains] are about 400 million years old and they are overthrust on top of the Grenville Province, [which is] 1.2 billion years old. That is as significant a geologic contact as a rift, and there's seismicity like the Ausable Falls earthquake that occurred at the boundary between the Appalachian [Mountains] and the Grenville [Province]. I would say that we can now have four continental crustal types [with the first three being] continental margins, interior continental rifts, and sutures. The fourth would be [named] "unexplained."

Wheeler: At least in North America, the Appalachian and the Ouachita Mountains override, by some unknown distance [that may be] tens of kilometers, the Paleozoic-Precambrian passive margin. Unless we have a depth for the earthquake, which we do for the Ausable [Falls earthquake], we don't know whether it occurred in the orogenic crust or in the underlying passive margin.

Coppersmith: This type of brainstorming, which occurred over a two-year period, probably started the whole process. What type of things do you think are important [in controlling Mmax]? Things that are truly crustal in dimensions, for example? In a failed rift or a successful rift, a passive margin, the faults have got to be deep seated. They've got to be truly crustal in extent. [The same holds true for] large sutures and things like that. The problem is that they are so few in number in terms of having magnitude 5 or larger earthquakes that you can't make the statistical argument. What we need to brainstorm, if we're going to use this approach at all, is how to develop priors and modify them in light of these types of observations. [In contrast,] we'll never get to a sample size that's large enough [for statistical analysis]. I'm sure that the past 23 years haven't added a lot.

Ake: The more you break that [global sample] down, the smaller N becomes. You are getting to almost noninformative priors where you start to see sigmas of 0.84.

Adams: Let's suppose that we use this tool to fix zones across the eastern U.S., and we come up with a prediction. Five years elapse. We have three big [earthquakes]. Each of those events, for the zone in which it happens, would probably cause you to recalculate the Mmax, [that is,] the Mu [(upper bound magnitude; table 2)]. What I would say is that any one of those should invalidate the Mu's for all of the zones. The reason I say that is Pete Basham, who was involved in the EPRI–SOG process around 1983, and who had done the 1985 seismic hazard map of Canada, chose Mmax on the basis of historical [earthquakes] and a few other things, but we had geological source zones in mind at the same time. In three of the zones, within a matter of eight years, events [occurred that equaled or] significantly exceeded the upper-bound magnitudes. They were Miramichi in the Appalachian [Mountains] [table 4], where the largest one was [M] 5.7. There were actually four events and added together they would have been 5.9, which is very close to the 6.0 that we'd chosen; Nahanni [table 4], which was a 6.6 followed by a 6.9 where the upper bound had been chosen at 6.0; and Saguenay [table 4], where basically in a background zone we had an upper bound magnitude of 5.0 and we had a 5.9. So those three examples caused us to reevaluate the Mmax's in all of the zones. It wasn't just in the zones in which they happened. Clearly we'd done something wrong, and I worry that the statistical argument here might be revised but only in the places where an event happened.

Coppersmith: The only thing that would change would be the likelihood function.

Adams: But would you change it just in the one zone where it happened?

Coppersmith: Would you change the prior (fig. A6–11)? You can see that this posterior [lower left graph] is dominated by the prior [upper left graph]. That's because we had five events, and now we have six. The likelihood function would get a little more peaked at the observed [largest magnitude], and if it was a 6.0, then this whole [likelihood function] would slide over [to the right] because the largest observed [earthquake] always defines the minimum of the likelihood function. So it would change this [the likelihood function] for this zone, but it's [the posterior distribution] driven by the prior.

Adams: It would be useful to see this in a case where a 6.0 [earthquake] happens in that source zone, because then we would see the actual changes. But a more difficult question is, should you be changing the outcomes in every other zone that uses that method, or only in the one zone?

Coppersmith: The only place you would need to change it would be where it [the new earthquake] affected your prior distribution. In other words, where it added to the global knowledge sufficiently.

Hanson: We've applied this approach in doing a large regional study of the southeastern United States for the TVA [(Tennessee Valley Authority; table 2)]. In doing that, you can start with the large domain[s], which is essentially the USGS model [of] two domains, and in which case you [would] truncate the prior up at a New Madrid or Charleston [magnitude]. To apply [this approach] to come up with alternative source zones, [as] for instance [in] the Canadian model where you [Geological Survey of Canada (GSC); table 2] have partitioned the crust into a number of specific zones, we considered alternatives that included partitions from very large regions to smaller ones. We updated the prior on the basis of for instance, the Fort Payne earthquake [table 4], which would [affect only] those zones that included that earthquake. That was how we tried to capture the uncertainty. I think the challenge in applying this technique is, how do you define those zones? Do you go with "we don't know anything" and use very large zones, so that a large Mmax applies all the way across, which is essentially the USGS model, or do you come up with the best geologic rationale for coming up with different domains?

Coppersmith: Eastern Tennessee is a good example. Arguments have been made to use the maximum lengths of the zone as a coherent rupture, or individual small segments might be used. Those would all be used to modify the prior distribution that comes from the global data set, to say that eastern Tennessee is different in such and such ways. The occurrence of events like John [Adams] is talking about will have very little effect. In eastern Tennessee, the number of events larger than M 5.0 continues to be zero. A more fruitful approach would be to try to modify the prior to incorporate what makes eastern Tennessee different from comparable regions around the world.

One of the things we are talking about is the SSHAC implementation guidance project and how we do these things in the future. How do we set up a system so that we can update things? Will the occurrence of some earthquake cause us to have to recalculate a national hazard map?

Hough: Am I correct that your prior distribution is assuming Poissonian [behavior]?

Coppersmith: Yes, we are assuming a Poissonian exponential.

Hough: What is the 8.3, the maximum SCR [magnitude]?

Johnston: That would be the 1994 estimate of New Madrid [the February 7, 1812 earthquake].

Frankel: I don't quite understand the top part of the prior there (fig. A6–11). You're basically taking an average of these magnitudes, the largest magnitude that's been observed in an SCR and its scatter. But that's not the distribution of Mmax. You don't know that the earthquakes that you've seen are the Mmax of those regions. You're just taking the largest events you've observed in those regions.

Coppersmith: You take the largest events you've observed and then you make that bias correction, which actually goes from the median of the largest [earthquakes] you've observed to the true Mmax. That takes into account the N, the number you've observed.

Frankel: Where does the bias correction come from?

Coppersmith: You have an exponential distribution, you assume that you have the median maximum observed event, the median of that distribution, and you know that this median estimate lies somewhere between the smallest event that you're looking at and the true upper bound. The more events you have, the better predictor that median estimate is of the upper bound. It's very similar to the approach of Kijko and Graham [(1998)]. If you have an area with hundreds of earthquakes and assume an exponential distribution, and if you have a mean or a median of that distribution, then you can estimate with some confidence the upper bound of that distribution. At 100 or 200 earthquakes you start to get a pretty good handle on that distribution. I don't know of anywhere in an SCR where you have that many earthquakes, but it's the same sort of concept.

Monday afternoon, September 8

Rus Wheeler (USGS): USGS Estimates of Mmax (Appendix 7)

[After the workshop, attendees told me that I had created the mistaken impression that the Mmax estimates in the USGS national seismic-hazard maps come from a Bayesian analysis. Recordings and notes showed related errors. One of the purposes of the workshop was to describe and critique the USGS Mmax methodology. Therefore, this insertion corrects those misunderstandings. Other speakers were offered the opportunity to submit similar insertions.]

[The idea of geologic and tectonic analogs is used in two ways. Some of our uncertainty about what constitutes an analog may be cleared up by distinguishing the two usages. The first usage looks outward from North America to the world, whereas the second usage looks back inward from the world to North America.]

[The first usage seeks worldwide analogs that share some of the geologic and tectonic characteristics of the CEUS and adjacent Canada (CEUSAC; [table 2]). The need to do this led to the definition of SCRs in terms of CEUSAC tectonic elements (Johnston and others, 1994). Kanter (1994b; 1994a) applied the definition globally and described the resulting analog SCRs. Kanter divided the SCRs into domains according to geologic and geophysical variables that were chosen "without consideration of seismicity" (Kanter, 1994b, p. 2–6). The talk that follows this insertion described another example of the first usage by blocking out the main tectonic elements of CEUSAC and seeking their analogs worldwide (for example, figs. A7–2, A7–3).]

[Later the domains were assigned values of additional variables that "might be related to seismicity" (Kanter, 1994b, p. 2–8) and combined into superdomains on the basis of variables that might control Mmax (Cornell, 1994). My talk on Tuesday morning ("What is a tectonic analog?") suggested an alternative way to divide the SCRs into domains.]

[The second usage of analogs is the more common. Arch Johnston examined the global SCR catalog and Kanter's domains. Arch made the seminal observation that SCR earthquakes of M 7.0 and larger occurred historically in extended terranes whose most recent extension was of Mesozoic and Paleogene age ("young"). I verified the observation from Arch's digital global catalog and probably other workers have done the same. Thus, the observation is reproducible. I applied the observation to North America to identify CEUSAC analogs of the globally observed young extended terranes (USGS's "extended margin"). The rest of the North American SCR is commonly referred to as non-extended SCR crust (USGS's "craton"), although it includes Precambrian extended terranes.]

[The distinction between outward analogs and inward analogs simplifies the problem of defining analogs. Outward analogs could be those that follow the definition of an SCR as given by Johnston and others (1994). Inward analogs could be those that follow Johnston's observation that large SCR earthquakes tend to occur in young extended terranes. The most straightforward way to identify inward analogs may be by sorting Kanter's domains according to Johnston's observation.]

[The justification for my application of the observation is that the global SCR catalog was compiled with the aim of collecting all historical SCR earthquakes of M 5 or larger regardless of their geologic settings (Coppersmith, 1994b). Each earthquake occurred within a domain, but as noted earlier, the domains were defined without considering seismicity. Thus, the SCR catalog is a single unit whose compilation was not strongly affected by ideas about possible geologic controls on Mmax. Our problem is how to divide the catalog into parts having different Mmax values.]

[The division of the North American SCR into young extended regions and all non-extended regions is common to the USGS and Bayesian methods. I know of no direct observations of the internal structure of the SCR database that support further division of the SCR crust. Other than direct observations, additional defensible divisions seem to me to require two things that we don't yet have: (1) physical understanding of how specific geologic properties control the propagation of seismic ruptures in SCR crust, and (2) the geophysical ability to detect the presence or absence of those geologic properties at hypocentral depths. With those two things absent, I can't see any justification for the USGS to divide the North American SCR into more than the craton and extended margin. This decision not to divide further is neither conservative nor non-conservative in terms of hazard in itself. Neither does the decision involve any assumption about whether Mmax is high or low. The decision can be changed easily whenever someone identifies additional defensible divisions of SCR crust in terms of Mmax.]

[Once SCR crust is separated into cratons and extended margins, it is straightforward to select the world's historical SCR main shocks above any particular M cutoff and to make histograms of the magnitudes in each of the two kinds of crust. Domains are unnecessary to the histograms beyond telling us whether an earthquake occurred in a craton or in an extended margin. To this point the USGS methodology appears to me to be simple, objective, and reproducible. Interpretation of the histograms is judgmental, as Kevin Coppersmith and Bill Lettis correctly pointed out at the workshop, but then so was the choice of six of the eight geologic variables that Cornell used to pool domains into superdomains—only crustal type and tectonic (geologic) age are based on Johnston's observation (Cornell, 1994, table 5–2).]

[Thus, the USGS method includes no Bayesian analysis, prior distribution or likelihood function. Because it uses no prior distribution, it does not use most of the geologic and geophysical variables that Kanter and Cornell used to define the individual domains and superdomains. The only two of these variables that are used are non-extended vs. extended crust and age of latest extension, both of which Johnston's observation identified as spatially related to Mmax. Because the USGS method uses no likelihood function, there is no need for a bias correction. There are no duplicate earthquakes to remove from the histograms—each earthquake in a histogram is needed to help define the histogram. The USGS estimates magnitude from scarp lengths and displacements by individual earthquakes, but otherwise we don't use physical principles because we don't know which ones are pertinent to propagation of SCR ruptures. For the same reason, we don't use most local geologic information beyond what is necessary to define the extended margin and craton. The second most important kind of local geologic information used is paleoseismic results. The USGS uses all chronologies of prehistoric earthquakes that include estimated ages, magnitudes, and locations of individual earthquakes (Wheeler and Frankel, 2000). Finally, we don't use analogs beyond identifying SCRs, dividing them into young extended terranes and non-extended terranes, and identifying their CEUSAC representatives. Most of the problems that trouble the use of analogs don't affect the USGS Mmax methodology.]

To assign Mmax values to the CEUS part of the 1996 USGS national seismic-hazard maps, I started with Arch's [Johnston] observation that M 7 and larger SCR earthquakes occur preferentially in young extended terranes (fig. A7–1).

We can apply the observation to the CEUSAC. We're going to explicitly set aside New Madrid and Charleston as special cases because those areas don't use our regional Mmax values. Arch [Johnston] recognized four kinds of SCR terranes globally. Each of these four kinds of extended and non-extended terranes—rifts, passive margins, Paleozoic orogens, and cratons—has North American analogs, and it's important to recognize that. We have a lot of rifts. The Reelfoot rift is probably the best known, the South Georgia rift [is] south of Charleston, [and there are many more] of them, mostly of Eocambrian or Cambrian age and of Mesozoic age. Passive margins [include] the Mesozoic Atlantic and Gulf Coast seaboards, and then inboard of that, we've got a Cambrian passive margin. They [the Mesozoic and Cambrian passive margins] are separated by the two orogens, the Appalachian and the Ouachita Mountains. Northwest of all of that is the craton. Our local craton is the Canadian Shield, where the cratonic basement is exposed, and its southward extension into the States, where the cratonic basement is covered by largely flat lying, thick to thin Paleozoic sedimentary rocks: sandstone, shale, and a lot of limestone.

Thus, we have North American examples of each of these four kinds of terranes, and here they are (fig. A7–2). Let me run through them briefly to point out one thing. All of these terranes are shingled to the southeast, in the sense that when any one of them forms, it overrides or faults or otherwise imprints itself on what is now the southeast boundary of the next older terrane. So we have the craton, and mainly in the Cambrian—starting a little earlier than that in the Precambrian and extending a little later than that—the blue part formed [early Paleozoic passive margin and rifts]. What is now the southeastern edge of the craton was the interior of a supercontinent. A [number] of more or less Cambrian rifts started. Some of them linked up into passive margins, [although] others stayed as failed rifts, like the type aulacogen, the Southern Oklahoma aulacogen. There were similar margins that surrounded CEUSAC on all sides including the west. That [rifting] broke up the supercontinent Rodinia and the craton's other half went somewhere else, and the Iapetan and its later successor oceans began to open. Almost immediately, starting in the Ordovician and continuing through the rest of the Paleozoic, those oceans started to close through subduction. Eventually, by the end of the Paleozoic they had all closed and the result was the late Paleozoic Appalachian and Ouachita orogens along the southern edge of the craton. That [ocean closing] produced the next supercontinent, Pangea, which lasted for about [fifty to] a hundred million years, and then it started to split up in the Mesozoic rifts that form today's oceans for the most part.

The orogens overrode the Cambrian passive margins: the red areas [in the figure] overrode the blue areas, mainly on thrust faults. The green [Mesozoic passive margins and rifts] overprinted the red by imposing a system of extensional faults on the red. The last thing to happen is that a [great deal of sediment was] dropped by large rivers [to form the Coastal Plain]. The tectonic elements are best exposed in the north. Farther south, there are more of these Coastal Plain sediments covering more of the tectonic elements.

One thing to note is that we've got rifts, passive margins, orogens, and a craton in the North American SCR. If we're going to look for analogs, we want to look for analogs to them, not to anything else that's not pertinent [to North America]. I'll come back to [this point] in another talk tomorrow. Here's a cross section through those tectonic elements, showing the overprinting relationships that I talked about (fig. A7–3). Here's the central craton [the non-extended terrane]. Everything from this line [between the craton and the Paleozoic extended terrane] southeastward is grouped into our extended margin. Here's the result (fig. A7–4). We drew six small polygons to handle special cases like New Madrid and Charleston. Why [these two are] special is not just the sizes of their earthquakes, but their frequency, every 500 years or so. The two exceptions are the big background zones, the craton and the extended margin.

Now I'd like to digress. I like the Bayesian analysis and I think I'm going to like what Walter [Mooney] has [to say] because they [both] use the global catalog. As John Ake mentioned, I can't see any reason to use methods that do not use the global catalog. There will be [some uses but those methods] just don't have enough earthquakes [to have stable statistics]. But what [use you make of] the global catalog varies. The Bayesian analysis seems to start with small areas and think about how to amalgamate or, in Allin Cornell's word, pool them into a smaller number of bigger areas that [each] have more earthquakes. Our [USGS's] approach is to take the whole global catalog as a given, and to start to think about how we can justify breaking it into parts. And so, of course, we're going to wind up with bigger Mmax [estimates] than the Bayesian analysis does because the bigger the area is, the more likely it is to include a big earthquake. But for a variety of reasons I suspect that the two methods are going to get closer [in their Mmax estimates]. Some fundamental differences [may remain, but that situation is an improvement over] "wait a thousand years".

This (fig. A7–4) is what the 2008 USGS national maps use. Here is how the different tectonic elements of the North American SCR show up globally (fig. A7–5). [On figure A7–5, the green area represents] the Mesozoic and [early] Cenozoic extended margins. Here [in eastern North America] is the one I've just been talking about, and as you can see, it has analogs worldwide. The only SCR that doesn't have any green is a small one, Arabia. The blue [area represents] the Paleozoic extended crust. There are several [Paleozoic extended terranes]—one in Australia, [and] three small ones in India—but there are not that many. So we have analogs of the extended margins. The orogens are not [shown] on this map but the Ouachita and Appalachian orogens are shown [along the green-blue boundary in eastern North America]. There are other orogens, but I don't know enough [about them] to generalize. The heavy black lines are the boundaries of the SCRs and the gray areas within the black lines are cratons.

We have cratons, orogens, passive margins, and rifts. We have analogs of the tectonic elements that make up the CEUS, and they're present globally. The red and purple dots are from Arch's [Johnston] global catalog. They are all [of the] M 6.5 and larger SCR earthquakes. The red ones are extended margin earthquakes and the purple ones are cratonic earthquakes. I put the earthquakes into those [two] bins. This is the version of the catalog that we used in the 2008 national maps.

I made these histograms (fig. A7–6) of the magnitudes. [The upper histogram is of the 17 cratonic earthquakes. The lower histogram is] of the extended margin earthquakes, which are about twice as abundant. Both are still small samples. This drop off [in the number of cratonic earthquakes smaller than M 6.7] I suspect shows incompleteness. I only sampled down to M 6.4, with the idea that any new conversions from these older magnitudes to moment magnitude would not increase the resulting magnitudes more than about 0.1 unit. By collecting earthquakes down to 6.4, I thought I'd be assured of capturing everything down to 6.5. I'm going to go back and pick up everything down to M 6.0. If the drop off disappears I'll be happy, otherwise it's probably not a completeness problem [attributable to sampling]. We don't have that problem as strongly, although it's here, for the extended margin earthquakes.

These histograms show the basis for the distributions that the 2008 national maps use for Mmax. The preferred magnitudes, which [have the largest weights] in the distributions, predate the distributions. I came up with [the highest-weighted points] in 1996 [and 2002]. Before we get into the values, look at the shapes of the histograms. These two histograms of different and non-overlapping types of SCR crust share three features.

First, they both have a tall peak in the middle to high 6's. The peaks are made of earthquakes in many continents and SCRs and a wide variety of plate tectonic settings, and just about everywhere. They say to me that the true value of Mmax probably can't be below the peaks anywhere.

Second, each of [the histograms] has one high outlier that is separated from everything else by a gap. I'm not comfortable giving much weight to a single anomalous earthquake [as the basis for] a conclusion. The cratonic outlier is from eastern China in 1917 [table 4]. It may well be an M_S from Gutenberg and Richter. The extended margin outlier is Kutch, India, 1819 [table 4], the Allah Bund earthquake. We can argue about whether Kutch was an SCR earthquake. That might be an explanation for this gap [that separates the Kutch magnitude from the rest of the histogram]. Whatever the reason, I'd want to know more than I know now before I throw this outlier into the distribution, because it's a single outlier from an old pre-instrumental earthquake.

Third, both histograms have a long tail running from the peaks almost to the outliers, and of more or less constant height. For instance, if this were a lognormal distribution you'd expect to see [the smaller-M part of the tail rising more, just above the peak] and I wouldn't expect to see [the slightly higher part of the extended-margin tail at M 7.3–7.6].

Thus, the problem reduces itself to picking a preferred Mmax value within the long tail, above the peak, and not as far out as the outlier. What actually happened was that I went to Arch's [Johnston] 1994 publications, because his 1996 papers were not yet published, and they showed that the largest cratonic earthquake was M 6.8. [For the 1996 maps, we used M 6.5 for cratonic Mmax. The consensus of a workshop in 2000 was that we should raise the value to M 7.0, and we did so for the 2002 maps.] What I called the extended margin similarly gave an M 7.5. That's what we used in 1996 and 2002.

During the Spring 2006 CEUS workshop, which was part of the preparation for the 2008 maps, Ivan Wong raised the question: why are you still using a single value? Why not put a distribution on it? That was the consensus of the workshop and after thinking about it a little, we did it. The distribution we picked is pretty straightforward. Whatever the preferred value is, the distribution would have four points and extend 0.2 up and 0.4 down, with a fourth point at –0.2. If you do that, you get the black brackets on the histograms. The large decimal numbers are the weights we put on the points—one half on the preferred value, and 0.2 on the top and next lower values. To that extent they are symmetrical distributions. There's a weight of 0.1 down here, on the preferred minus 0.4. To my eye, these distributions fit the histograms fairly well. Both of them exclude the high outliers. Both of them honor the high peaks. This one [the craton] goes into the peak, but it only has a weight of 0.1. This one [the extended margin] is a little too shy of the peak. If we did it again we might extend [the extended margin distribution] down to here [M 6.7 or 6.8], but again the weight is only 0.1. And both of [the distributions] fit values somewhere on the long flat tail. That's how we got to where we are [today].

The biggest advantage of this approach, and this applies to anything that uses the global catalog, is that it does just that (fig. A7–7). Thereby, it avoids the small-sample problem that is insurmountable for a lot of the older methods that used small regions. The disadvantage is what I referred to in my earlier summary, and like I mentioned then, I have a couple of suggestions about how to start getting around this disadvantage. This list of uncertainties could be longer. The first one of course is ["what's an analog?"]. The second is that the individual magnitudes are uncertain. The most important magnitudes are the biggest. They are likely to be the oldest so they are likely to be the most uncertain. But that's tractable. You can put pluses and minuses on each of the earthquakes [magnitudes] in those histograms and figure out how that should affect the distributions. The third thing, which I haven't done yet, is to test the robustness of the histograms. There are one or two earthquakes or groups of earthquakes that I included which perhaps I shouldn't have included. The outer parts of the passive margins, the more seaward parts of the Mesozoic passive margins, might [contain] such a group of earthquakes. The reason is that faults mature as they accumulate slip. The more the cumulative slip, the more short faults link into longer ones, and the more faults of any size tend to smooth out their geometric and strength irregularities. That makes a long, smooth homogeneous fault with few barriers to a propagating rupture, and we'd expect it to allow larger Mmax. I need to see how those histograms change if I take out the outer continental margin earthquakes. There are one or two other sets of earthquakes that should be taken out [to see how the histograms change]. There are also earthquakes that I excluded that perhaps should be in there. So I need to add them [in] to see what they do to the histogram. I'm testing the robustness of the histograms under uncertainty about exactly which earthquakes should be included.

Discussion

Frankel: We used Mmax of 6.5 in the craton for the 1996 maps. For the 2002 maps, most people at the CEUS workshop thought we should bring that up to M 7.0 for the craton.

Adams: The values in the histograms (fig. A7–6) should be bias-corrected by the Geomatrix method. It won't take the distribution down, it will take it up. Maybe not very much, but it might explain, for example, why you've got 6.6 for one craton and 7.1 for another. If you've chosen your analogs correctly, and if the correction is correct, then they'll all fit on the same value.

Lettis: I have comments along the lines of John's [Adams]. First, make the bias correction on all of them. Second, if there is a region where there is more than one earthquake, the Mmax should not be less than the largest observed magnitude. You could eliminate the smaller earthquake if there were two or three close to each other. That could be a way of getting rid of some of the smaller earthquakes in the histograms. Third, after making the bias corrections and eliminating the smaller earthquakes, why wouldn't the histograms represent both the most likely Mmax and its uncertainty? You only represent the upper tail. Why wouldn't the preferred [value] be 6.7 for the craton, with a lower uncertainty going out to capture the tail? Why are we being conservative in picking the upper tail of the distribution?

Wheeler: I don't think it's conservative. The way I interpreted these peaks is that's the lower limit of Mmax. It can be argued that this 0.1 can be moved down [to M 6.7 in the extended-margin histogram of figure A7–6], that is, to extend the lower tail [of the distribution] a bit.

Lettis: But why do you consider that to be the lower limit of Mmax? That's where the majority [of the values] are falling and we want to pick the most likely value of Mmax, and then capture the uncertainty on each side.

Ake: What would happen if you adjusted Rus's choice of 6.5 [as the lower limit of compiled earthquakes]? What if you made it 6.2? The way we're doing this right now, it's somewhat dependent on what Rus defines as his first bin.

Wheeler: I guess I can't answer Bill's question until I rework the catalog [by compiling down to M 6.0] and see whether this gap [between M 6.5 and the tallest peak in each histogram] fills in. What we would expect to see is this [the histogram] continuing on up [past the peak to still lower magnitudes], because these [smaller] earthquakes are more frequent than these [larger] ones.

Lettis: There are a lot [of these earthquakes] in South China.

Wheeler: Yes, there are. I've come across a few papers recently that suggest that North China might have had a lot of Mesozoic tectonic activity and it does have a lot of Mesozoic intrusives. I want to look into whether it's worthwhile to just remove that block and no longer call it an SCR. But I don't want to do it without talking with Arch and thinking a lot about those papers.

Johnston: I don't remember all the arguments, but I do know that South China was the most problematic of all the SCRs that Hsü defined [cited and summarized in Kanter (1994)]. That cluster of four earthquakes, shown in red on figure A7–5, is just [onshore from] Taiwan, where Taiwan is colliding with China.

Adams: We came to the same conclusion to define South China as a much smaller block than that. It would basically cut off all of those earthquakes. Why don't you include New Madrid?

Wheeler: The earthquakes are so darned frequent there [at New Madrid]. The only other places where big earthquakes occur at 500–year intervals, and this goes for Charleston [South Carolina] too, are at plate boundaries. I'm not saying that's a plate boundary, but something is happening very rapidly there, and it's got to be geologically short lived because there are no young mountains there.

Adams: Maybe the same goes for Charleston, or for other places. It's just that we don't have information there like we do at New Madrid. I think it's a jump in logic to leave [[out]] New Madrid and keep Kutch. I think you have to be consistent.

Wheeler: That's right. Now we know that Kutch has a short recurrence interval [too]. We knew that by 1956 [when the Anjar earthquake occurred near Kutch: table 4]. We certainly knew it by 2001 [when the Bhuj earthquake occurred near Kutch].

Hough: And if New Madrid is a seven and one-half instead of eight, then it's much more logical [to include it].

Wheeler: I will look into what happens when I keep New Madrid and Charleston. With New Madrid I need to look into what happens with each of the three groups of magnitudes that have been suggested, and I should do the same for the three published magnitudes for Charleston. Thank you, that's an excellent suggestion.

Bilham: Going back to the slide with the histograms [fig. A7–6], how well do you know these magnitudes? The Allah Bund event [Kutch, India, 1819] I happen to know is plus or minus 0.2

magnitude units. You've got an awful lot of faith in plopping these down. In the construction of that histogram, where are the uncertainties?

Wheeler: They're not in there now. That's the second uncertainty that I looked at in the slide after this one [fig. A7–7]. Actually I'm delighted that you put a plus or minus as small as 0.2 on the Allah Bund earthquake. There's been a lot of work done on it, and that may be why.

Bilham: The instrumental ones to your left [in fig. A7–6] are probably pretty good but as soon as you get into the ones that are [calculated from] intensities, then the uncertainty increases. The long tails might look very different if you actually knew what the magnitudes were. I think that most seismologists would guess better than 0.2 on most earthquakes, but these large earthquakes are very different.

Wheeler: Arch [Johnston] did publish something that I'm probably going to use for this, at least to start with. Each of his earthquakes had one of about ten levels of uncertainty that ranged from very small, for teleseismic moment magnitudes, up to plus or minus one, for some of the older earthquakes with sparse intensities.

Adams: You don't like the 7.5 [in the craton histogram], and maybe I accept your lapse for one aberrant analysis. Tell me why the next one down shouldn't be Mmax for the whole analogous region. I put it to you that the reason you are unhappy with doing that is because you are trying to do two things. You have a very large area you think is analogous, and you want to choose the biggest [earthquake]. But you also worry that the analogy is not correct. So your uncertainty range up there is really trying to sample the epistemic uncertainty in the analogs. You're trying to do two things, in other words.

Wheeler: You're probably right. I need to justify throwing the [aberrant] one out, and I need to justify keeping it if I don't throw it out.

Adams: Throwing up a histogram like this is quite valuable. Let's suppose we take magnitude into account. That 7.5 might be over a certain range. What I suggest you do is to give it ten building blocks which you distribute over that range. Randomize them. Add all of them together. Count down from the upper block ten blocks and choose that as your Mmax.

Wheeler: How should we randomize the blocks? Should we use a uniform distribution or something else? It's somewhat like a Bayesian distribution in that the shape of the randomization will dominate the result. If you use a uniform distribution you'll get something that looks like a ragged long tail. If you use a normal distribution you'll get something that's peaked. If you use a log normal distribution, you'll get something different.

Adams: Do it all three ways and see if the results differ.

Lindvall: On the extended margin distribution of Mmax, in picking a lower bound you picked a spot on the upper end of the frequent earthquakes. That's not exactly what you did for the craton. Why did you go down to [a lower bound of] 6.6 instead of staying up at 6.8 on the upper end of the [peak]?

Petersen: At the [CEUS] workshop [in 2006, in preparation for the 2008 national maps,] there was discussion about Mmax in the craton and the extended margin. Like Rus said, people felt like we needed to put a distribution on it. You can see that the distribution on the extended margin is the same as the cratonic one: [Mmax] plus 0.2 and we wanted to go to minus 0.4. It's not quantitative, but we felt like there could be an asymmetric distribution.

Lindvall: It may have been more driven by wanting a consistent distribution on both.

Frankel: I think I'm hearing that there are two different ways of looking at Mmax. One of them says there is spatial variability of Mmax. When you do a Bayesian analysis, you do an average and we get a distribution, so that when you see a New Madrid earthquake it's only part of the distribution of Mmax that you can use in any given spot. The other view of it is, I've got a bunch of earthquakes, I have a seven and one-half at New Madrid, plus or minus whatever the uncertainty is, that's my Mmax for everywhere. It seems there are two different ways of looking at the problem. There's nothing wrong with either of them. That's why we sometimes talk past each other.

Coppersmith: These were developed for the forward application, to say here's the Mmax for all cratonic areas, and here's the Mmax for all extended areas. Art [Frankel] was saying that a typical PSHA for a [power plant] site would have a different Mmax for each source. That's where the source-specific information would come into play. For example, as I showed you, the likelihood function certainly would define the minimum size of the maximum earthquake. The magnitudes that are below that are effectively taken out of the Mmax distribution anyway. What you've done in the grouping is very similar to what was done to develop the prior distributions for these two, cratonic and otherwise, except that we integrated everything from four and one-half and larger. Potentially those also defined maximum earthquakes, but then in application almost every zone we've looked at has had at least a 5.0 and in some cases larger. So that part of the distribution is removed in the second step. The issue here of trying to decide what can be generally applied to a region as opposed to what could be applied to an individual source zone is what causes us, as Art said, to talk past each other. They are different applications.

Wheeler: It's not as hopeless as talking past each other. I think there's a more fundamental similarity between the two methods than just both using the global catalog. The Bayesian analysis starts with domains. Lisa Kanter [(Kanter, 1994b)] divided the world's SCRs up into things she called domains, places [areas or regions] that were [each] tectonically more or less the same [homogeneous]. The one I know best is the central and southern Appalachian Mountains, which she put into a single domain. They [domains] are typically bigger than traditional source zones, which are the size of one or two states. Some domains are small. I don't remember for sure, but probably the Reelfoot rift is a single domain. Others are bigger. She had about 200 of them worldwide. Some of them were set aside because they didn't have any earthquakes above the cutoff magnitude, which was four and one-half or five. They were set aside in Allin Cornell's analysis.

Then you and Allin [Cornell] started looking for ways you could amalgamate [pool] like domains into superdomains. The larger you expand a single domain or source zone, the more likely it becomes that you will include a larger earthquake than any you started with. In fact, Mobs can't go down. It can only go up or stay the same. So the more you pool small domains into a smaller number of larger superdomains, the more the average Mmax of all the superdomains will go up. Cornell's chapter 5 [(Cornell, 1994)] describes a particular level of pooling. There was a particular set of Mobs, one for each superdomain, and that [set] was the source of the prior.

I did some experiments. If you use no pooling, so that you're dealing only with Mobs from all the individual domains, then you can make a prior and you get a prior distribution that is centered at comparatively low values. If you do the level of pooling that Allin [Cornell] and Kevin [Coppersmith] did, you get a prior but it is centered at a slightly higher value. I think you got something like 70 or 90 superdomains. I did another level of pooling using fewer variables, got just 12 superdomains, and the center of the prior went up still more. If you do a lot of pooling, you wind up with the central craton and the extended margin of the national maps, and Mmax goes up to these values, 7.0 and 7.5 (fig. A7–6).

The Bayesian analysis and the national maps can be thought of as two points on a continuous distribution of level of pooling. I think that Allin [Cornell] chose too low a degree of pooling because he thought that too many of the descriptor variables should have been included. Fundamentally, they started with a bunch of domains and tried to figure out how much they should pool. We go the other way. We start with a given global catalog and try to figure out how we can subdivide it. So far we've only been able to divide it down into two [parts], using Arch's [Johnston] observation. It wouldn't surprise me at all if someday somebody comes up with a reasonable way to subdivide it even more, in which case our Mmax's would come down. So I think there's promise for shrinking that gap, but I don't know whether it will ever disappear.

Coppersmith: You started with one and came down to two, we started with about 80 and came up to two. Because N is small, the bias correction is very small, and you end up with what's basically the same distribution [the same prior distribution that you had before the bias correction]. The only assumption that plays into this is that there is an exponential distribution.

Wheeler: For each of those two "superduperdomains" [extended and non-extended crust, which together make up the world's SCRs], you make a prior from the domains [superdomains] in that "superduperdomain". You make one prior for all the extended domains [superdomains] and a second prior for all the non-extended domains [superdomains]. And I'm not sure you should do that. I'm not sure how I would change it, but I'm not quite comfortable with it.

Coppersmith: The one that has the low number [of superdomains], the extended [superdomains], has N of 30, and [it] goes through a bias correction that raises [the estimated Mmax] up about 0.3 magnitude units. They [the global-analog Mmax estimate and the Bayesian estimate] should be identical by the time it's done. The only difference might be in the catalog that was used in this study and the updated catalog you might use now. Conceptually they arrive at the same thing. Then the next step is different. The next step in this [Bayesian] approach is source-specific. The likelihood function truncates [the prior] at the largest observed magnitude within our source. If you're in New Madrid, then that's the source, and [the likelihood function] will truncate [the prior] at that part of the distribution [at a high magnitude]. [The posterior] will be a fairly narrow distribution that starts at whatever the New Madrid size is and extends upward. Its application is for that particular seismic source.

Walter Mooney (USGS): Mmax in Cratons (Appendix 8)

I'm going to follow in a similar vein and address the issue that Art Frankel brought up, which is that Mmax can vary within the craton (fig. A8–1). The point I'm going to focus on is that we know quite a bit about this craton (fig. A7–2) and it's not just a homogeneous white mass. The result is that we have ways that we can use to estimate Mmax and its spatial variability within the white area (fig. A8–2). I'm going to talk about how cratons evolve, the correlation of earthquakes with their tectonic settings, and then my main point about Mmax. I'm going to conclude that Mmax for cratons is 6.8.

North American geology is familiar to you (fig. A8–3). The exposed craton is to the north, in Canada [reds and oranges]. The unexposed craton is to the south, in the U.S. [blues and purples], where it's covered with [sedimentary rocks].

How does this [craton] evolve? The white area that Rus [Wheeler] was showing is actually very complicated. The way you form a craton is through the amalgamation, starting [billions of] years ago, of cratonic blocks (fig. A8–4). When they come together they form sutures (fig. A8–5). This would be the Superior Province [yellow and orange on the figure]. You do have zones of weakness even within cratons. As we evolve further (fig. A8–6) you bring in other accreted terranes from the sides, [and] suture them [to the preexisting craton]. Then you have things like New Madrid that occur, shown here with a plume-generated rift. So that white cratonic area is very complicated. It has a lot of geology and it's just as complicated as Paleozoic geology. We know that there are earthquakes associated with incipient rifts, but there can also be earthquakes associated with suture zones and other [features].

<u>Frankel:</u> What's the largest earthquake associated with a suture zone?

<u>Mooney:</u> The Ausable Forks earthquake [2002, northern New York State] was only a 5.4. I can't [think of an example at the moment] of a larger earthquake like a 6 on a suture zone.

Johnston: The Narmada-Son lineament [is] a suture zone in India. It's had a magnitude [in the] middle sixes.

Frankel: I remember looking at all these [geologic and tectonic] maps with earthquakes all over them. I don't see how you can correlate them with geology.

Mooney: The next three slides are taken from this paper in Geophysical Journal International (fig. A8–7). This is something very similar to what Arch [Johnston], Lisa Kanter, and Rus [Wheeler] did (fig. A8–8). This is all the margin events. We updated Arch's catalog from 942 events to 1,483 and you can see that a lot of intraplate seismicity is associated with the yellow regions [orange on some printers or screens] and you can see a lot of events that are on margins. These are all the rift events (fig. A8–9). A lot of the correlations with rifts are quite clear, New Madrid being only one of the many examples. Our rift catalog is taken from a very complete summary. It's about a 180–page article by Celal Şengör. These are the events that there's no explanation for (fig. A8–10), and a lot of them I bet are on sutures. They are about 36 percent of the total. Those are the events that are non-rift and non-margins. When we looked at the seismicity of North America after this earlier study, it began to look like as you go into the exposed craton, you have fewer events (fig. A8–11). We wanted to look at the magnitudes. There's a geological correlation [between fewer earthquakes and greater distance into the craton].

Now, what is a craton (fig. A8–12)? A geophysicist thinks of the lithosphere as having two parts, the crust and the mantle lithosphere. The thicker lithosphere is Archean, in the center of the craton. The point I'm going to make is that in areas of the craton where you have this thick lithospheric column, there's little deformation. Deformation is concentrated at the margins [of the thick lithosphere]. [To determine whether or not it's] a craton, we have to go to petrology (fig. A8–13). This [diagram] shows that a geochemist has no problem telling you whether you're on a craton or not. North China is not a craton. You are correct in removing that. When a kimberlite comes through a true craton, it brings with it samples of the rocks that are beneath the crust. These will be found to be depleted in iron. They are very different from the convecting mantle below. The convecting mantle is enriched in a huge number of elements. A geochemist has no trouble at all telling whether you are on an undeformed stable craton. The geochemistry of the xenoliths coming through that lithospheric column looks completely different depending on whether you are on the thick Archean, the thinner Proterozoic, or the younger material. In China there's abundant evidence that you are not on the original cratonic lithosphere.

There's a geochemical definition [of craton] as well as a surface geology definition, and now I'm going to add a third definition from seismic tomography (fig. A8–14). The theory behind what I'm saying is that the deformation in the lithosphere follows a power law and the power law has temperature [T] as a primary quantity. For the slow deformation that occurs in the lithospheric column, those areas that have a cold geotherm are going to deform much slower than those areas that have a warm geotherm (fig. A8–13). This area [Archean crust with cold geotherms] is going to deform at its boundaries. You really can't load it very much because it keeps slipping at its boundaries. There's a limit to the maximum size earthquake because in this equation (fig. A8–14), the deformation rate and the loading of the faults on the boundaries are going to be much faster because of the higher temperature regime.

Seismic tomography (fig. A8–15) can tell us where we have these thick lithospheric roots. Where are these cratons? Hopefully the seismologist and the geochemist will agree on where these cratons are located, how thick they are, and how big the earthquakes are. This is a diagram prepared by an unusually careful seismologist, Steve Grand at the University of Texas, Austin (fig. A8–16). This is an S-wave model from double bounces. It shows the speed of the lithosphere at depths between 100 and 175 [km depth]. Where it's blue it's cold and fast, and that's the lithospheric roots where the cratons extend down into the mantle. Where it's red it's hot, with convecting asthenosphere at those depths.

Where it's yellow, it's thinner lithosphere on the order of about 120 km. So the blue areas are the real cratons, like West Africa, the Congo, and South Africa. Australia can be divided into eastern Australia, which is thin lithosphere, and central and western Australia, which is thicker lithosphere. And of course [there's] the famous Baltic shield and then our favorite, which is the Superior Province.

What we're going to do is correlate the thickness of the lithosphere with seismicity (fig. A8–17). This map is similar to the previous slide. Blue is cold, thick regions, red is warm, thin regions, and black dots are global earthquakes from the NEIC catalog. The next slides show how magnitude correlates with thickness of the lithosphere (fig. A8–18). We are measuring the S-wave velocity down to depths of 175 km. Thick lithosphere is on the right of the slide, corresponding to a positive S-wave velocity anomaly of 6 percent. Very thin lithosphere is on the left. There is not much of a correlation for M 4–5. For magnitudes 5–6, the number of earthquakes begins to drop where S-wave anomalies are about 4 percent or larger (fig. A8–19). For magnitudes between 6 and 7, I was hoping the correlation would be a little bit stronger (fig. A8–20). For magnitudes greater than 7, the global catalog contains no earthquakes in areas with S-wave anomalies greater than about 3 percent (fig. A8–21). This is where the lithosphere is thick and cold, the mantle has a different chemistry (fig. A8–13), and the global catalog contains no earthquakes greater than M 7. Now here's everything (fig. A8–22). You can see that all the earthquakes larger than M 7 occurred with S-wave anomalies less than about 3–4 percent.

Let's look at just North America (fig. A8–23). I would want to put a transition zone, maybe 300–400 km wide, around this [blue to green] boundary. North of that boundary I would estimate Mmax to be on the order of M 6.8, and higher within the transition zone. I don't have a number for you for Mmax south of that transition zone.

Our goal is to correlate deep geophysics with the maximum size of the earthquakes so we can quantify the spatial variability of Mmax within cratons. I think it's a pretty robust method because with better seismic coverage we are getting better maps of the lithosphere. This is a ten-year old map of the lithosphere. We've got new ones coming out now that do a better job. We'll be able to take what Rus [Wheeler] has been doing one step further to add spatial variability within cratonic regions.

Discussion

Bilham: This is an empirical thing. What if you ran the clock forward, say, ten thousand years. Would you expect something to pop up to the right [on your histograms], or would that change your opinion about the largest [earthquake] you could get?

Mooney: Your point is well taken. It's entirely empirical and it's entirely based on the analogs that you use.

Bilham: What you've found is obviously very interesting. What's the physics behind it? You were implying that you've got a nice thick lithosphere and therefore low strain-rate. But if you have low strain that doesn't necessarily mean you can't have a very large earthquake. Is this limited by the size of the fault, or by the strength conditions before they get relieved?

Mooney: I agree with you again. This only means a low strain-rate. My thinking is that the strain is coming from ridge push or basal tractions in mantle convection, so the strain will be relieved by faulting in the weaker zones and a continuing cycle back to a low strain-rate in the interior. But I agree with your second point. We can only say that it has a low strain-rate. We don't have any evidence that it can't have a large earthquake every five thousand years. Sichuan reminded us of that, that low strain-rates can correlate with a 7.9 earthquake [2008, southern China; table 4].

Johnston: It could be that you just raise the strength [for an old, cold, thick craton]. You could double or more the thickness of the brittle strong layer in the crust of the craton. For example, the Bhuj earthquake, although it happened in a rift that propagated into the craton, ruptured perhaps all the way through the crust, but certainly down to 35–36 km. Other evidence is that the Fennoscandian ice sheet was on the Baltic shield [dark blue in fig. A8–17]. There was post-glacial faulting on the Pärvie fault and they've got microearthquake activity down around 40 km. The earthquake that produced the Pärvie fault was certainly middle 7 to 8 in magnitude. Granted, those were very special conditions, where you removed the ice sheet suddenly. Maybe it takes special conditions to get that large an earthquake.

Mooney: Before I take the next question, there's a point I forgot to make (fig. A8–13). This lithospheric column [of Archean age] has a different chemistry than elsewhere and one of the qualities is that it's much stronger. It's stronger because it is dewatered. The cratons have lasted 3.5 billion years even though the continents have made their ways around the world so many times, because they [cratons] are cold and strong. [We know that] they are cold and strong because [strengths of samples of their rocks have] been measured in the laboratory and because they've lasted 3.5 billion years.

Bilham: That's sort of the same argument. You could say that cold and strong means much bigger earthquakes when they go [occur].

Hough: In [fig. A8–16], Bangladesh and parts of the Himalayas are awfully blue.

Mooney: There are both artifacts and cratons there. You're seeing a blurry image of the slab hanging into the mantle. He [Grand] doesn't have the resolution to get the Indian craton because it's too small a target.

Allen: My colleague back at Geoscience Australia, Dan Clark, found evidence of very large paleoseismic earthquakes in the Yilgarn craton in the southwest of Western Australia, upwards of magnitude 7.4 with a slip of about 4 m per event. [But see remarks in Allen's talk tomorrow morning.]

Mooney: That's an extremely important observation. The Yilgarn has gold and diamonds, so this (fig. A8–13) is a good model for the Yilgarn.

Allen: But there may be some question about the magnitudes because a lot of the earthquakes are very shallow. They nucleate in the top 5 km or so and there's a mid-crustal detachment. So the faults may spread out further instead of jumping down into the crust.

Mooney: There are other regions that are anomalous. Because the continents are under compression, most continental rift breaks are either strike-slip or reverse faults. But in South Africa's cratonic lithosphere, because of uplift you have normal faulting.

Steve Harmsen (USGS): Uncertainties in Parameters Related to Mmax (Appendix 9)

Mmax is one of a collection of uncertain parameters in PSHA (figs. A9–1, A9–2). The first point I'd like to make is that an effort to improve a PSHA by improving Mmax isn't guaranteed to improve the result if other linked parameters aren't improved correspondingly. For example, if an Mmax is for a million-year long recurrence in a source, it's unlikely to have a material [effect] on the PSHA. What we're really talking about is: how do Mmax and associated parameters [affect] the hazard curve?

In the USGS PSHA we divide sources into those for which we have a significant amount of local information, and those that are based more on concepts (fig. A9–3). The faults in the CEUS that we have good knowledge of, and which go into the PSHA, are the Cheraw Fault in Colorado, the Meers Fault in Oklahoma, and [faults] in the New Madrid seismic zone. The earthquake size that is related to Mmax, which we call "Mchar" [(table 2)] with an uncertainty distribution, is estimated from local evidence. Similarly, Charleston, South Carolina is a characteristic source zone whose size [Mchar] can be estimated from local evidence. Background sources, as Chuck [Mueller] mentioned, are an important part of CEUS hazard. They are often at locations where there is negligible evidence for estimating fault size, so the methodology that we use at the USGS assumes that long, capable faults are essentially everywhere.

What I’ve done is to find some regions that cover a big part of the CEUS (fig. A9–5), and compute the expected rates of exceedance of different magnitudes (fig. A9–4). Then I compare that information with available information about historical earthquakes, and I also talk a little bit about paleoearthquakes. The regions I chose are a central U.S. region (CUS; [table 2]), which has a radius of 970 km and has both margin and craton sources in it, and a Northeast U.S. region (NEUS[; table 2]), which has a radius of 700 km and also has craton and margin sources. For computational purposes, the only important logic-tree branch on Mmax in the USGS model says that in the craton Mmax is seven and in the margin it’s seven and one-half. Given these values, the fault length within the margin is 90 km and within the craton it’s 41 km. Those are not verified numbers, but they are [from] the Wells and Coppersmith regressions of length on moment magnitude for unspecified slip type.

Here are my results (fig. A9–6). The red curve shows the USGS model for the rate of earthquakes exceeding different moment magnitudes from four and one-half up to seven and one-half in the CUS region. The blue curve is the same rate except for the NEUS region, and the black curve is the sum of the two curves. I focused on the rate of exceedance of M 6, which is about one every 43 years [in the CUS and NEUS areas combined], and the rate of exceedance of M 7, which is about [one in] 480 years [in both areas combined]. These are just for the background sources. They don’t include faults that have local information, such as Charleston, New Madrid, Meers, and Cheraw.

I tried to compare those numbers with the actual historical record (fig. A9–7). The result depends on whom you ask. There is a variety of magnitude estimates for earthquakes, and USGS chose the NCEER magnitude estimate as preferred if it is available [Mueller and others, 1997]. There were no earthquakes with magnitude greater than 6 in the CUS and NEUS circles, according to the declustered USGS catalog. As I mentioned before, I omitted the New Madrid main shocks, foreshocks, and aftershocks because they are not in the background and are handled separately. I omitted the 1843

[Marked Tree; table 4] and 1895 [Charleston, Missouri; table 4] earthquakes because they have NCEER magnitude less than 5.5, but several other sources give these magnitudes greater than 6. I omitted Cape Ann, November 1755 [table 4], because it has an NCEER magnitude of 5.8, although others such as John Ebel give this a higher magnitude. As I've said, I omitted Charleston, not only because it is handled separately but because it is outside the CUS circle. I omitted the Canadian sources—St. Lawrence River, Charlevoix, Timiskaming, and Grand Banks [table 4]—because they are outside the study area. Finally, I've omitted Giles County [table 4] because its inferred magnitude also was less than [magnitude] 6.

I also needed some kind of estimate of completeness of magnitude 6 and greater (fig. A9–8). This is a pilot study, so I said there's a 300–year completeness in the NEUS and a 200–year completeness in the CUS. There might be some question about Oklahoma, which was a frontier [until the late nineteenth century], but certainly [many people were living] along the Mississippi and Missouri Rivers by 1808.

We can make a probability statement, again using the Poisson distribution, and proceed exactly as Martin Chapman did in his study (fig. A9–9). Let mu1 be the expected number of earthquakes with magnitude greater than 6 in a 300–year period in the NEUS. From the recurrence curve I showed earlier (fig. A9–6) that value is 2.85. Let mu2 be the expected number in a 200–year period in the CUS, and [that value is] 2.72. The sum is 5.57 and the Poisson probability of observing no earthquakes, using the recurrence model we used, is 0.0038. If I missed one, the probability is 0.025, so it's a pretty significant result[, implying that M 6 might be about the largest we'll ever see in the two circles].

Is the significance test any good (fig. A9–10)? Only if the observed number of earthquakes is less than two. If it's two or more, that number [the Poisson probability of the last figure] rises to more than 0.05. Complaints that I received [include] that I omitted [the] Cape Ann, Timiskaming, and

Charlevoix [earthquakes]. I [do not agree with] that argument [for Timiskaming or Charlevoix], but [for] Cape Ann I have to agree with [it]. As far as the New Madrid “aftershocks” go, that’s an open issue. Maybe NCEER did underestimate those magnitudes. If that’s the case, then this historical comparison is not so significant. Also, completeness issues can be argued. These regions were [unsettled] frontier for substantial parts of the time. Also, you could argue that the [seismicity of the] last 200 or 300 years may not be typical. For example, New Madrid may have put much of the CEUS into a shadow zone. That’s not in our model, but maybe it should be.

Even if you could overcome all the objections and say the test is good, you wouldn’t necessarily have to reduce Mmax to get a non-significant test result (fig. A9–11). For example, you could reduce the predicted earthquake rate in the M 6 to 7.5 range by a factor of two or so. In fact, that’s precisely what was done in California. They reduced the background-source rates to one-third of what the predicted rates from the smaller events were to get a better match with the historical rates.

The paleoseismic test is a little more interesting (fig. A9–12). There’s just one event in those regions outside the areas that I’ve omitted. That’s the Vincennes [western Indiana] event, which occurred about 6,100 years ago. [Evidence of this earthquake] was discovered in the 1985–1990 period. We’ve found one significant [larger than M 7] paleoseismic earthquake in the last 20 years outside New Madrid and Charleston. Mu is 12.5 and the probability of one or fewer M 7 or greater, given that mu, is zero. Furthermore, [as far as I know,] there are no CEUS nuclear power plant applications that indicate evidence for Holocene earthquakes near the sites.

Discussion

Lettis: What about the Saline River liquefaction features [in southeastern Arkansas]?

Harmsen: The Saline River liquefaction features haven’t been well enough associated with a source, at least for the USGS to include in their model.

Also, you might want to make circles big enough to include the 1663 Charlevoix earthquake [table 4], which was approximately M 7 (fig. A9–12). [If you enlarged the circles] the probability of two or fewer M 7 is 0.0003. How many paleoearthquakes do we need to increase the probability of getting that many or fewer when the expected number is 12? We need six paleoearthquakes to give support to our model, and five [paleoearthquakes] only give a 0.015 probability. Not that it makes a strong statement about Mmax. It makes a statement about the magnitude-frequency distribution that we're using up to Mmax (fig. A9–13).

Johnston: I know the NCEER catalog is in m_b or m_{bLg}. Is that one of the reasons it was used in the hazard analysis?

Harmsen: NCEER is used for the background sources, and that is because most catalog magnitudes in the CEUS are given in m_b units. Magnitude 6 m_b and moment magnitude generally are about the same. Some regressions say that moment magnitude is a little bit less at M 6.

Johnston: It gets more divergent the larger you get. I thought the rest of the hazard mapping was done in moment magnitude.

Harmsen: All sources are done in moment magnitude and western U.S. background sources are done in moment magnitude. My graph was done by converting body-wave magnitude to moment magnitude, using both Atkinson-Boore [(Atkinson and Boore, 1995)] and Johnston [(1996)].

Ake: How sensitive would your probabilities be if we were investigating, say, [M] 5.5?

Harmsen: I believe there's a good collection of [magnitude] 5.5s. My gut reaction is that it wouldn't be significant at 5.5. Six is an important cutoff, but I'm not certain because I didn't go [to a magnitude that low].

Ake: Did you look at sensitivity to completeness in the CUS? In parts of western Arkansas and Oklahoma, there weren't a lot of people living in the area until maybe the middle 1800s, if then.

Harmsen: I think that's a valid concern. Has anybody worked on this? What are your feelings, Sue, about the 200–year completeness for M 6?

Hough: I sort of sympathize with Jon's [Ake] point. There was one New Madrid event on December 16 [1812] that was widely felt but wasn't strongly felt anywhere. One way to do that is if it were west of the New Madrid seismic zone. That would be consistent, but we just don't know.

Ake: We have to go back and look at census data, and when the first newspapers appeared in various communities. It's just grunt work that has to be done.

Ebel: In your northeast zone up there you have 300–year completeness. That would take the catalog into the 1600s. But when I looked at the earthquakes in the 1600s in the Northeast, I was looking at reports [only] from basically Trois-Rivières in Quebec and from eastern Massachusetts. There were several earthquakes that were felt noticeably in both areas. To be felt in both of those areas they have to be at least above magnitude five. Of course, they could have been anywhere to the west, and the farther you put them from those two regions, the larger the magnitudes. We don't really have the data, but if we were able to go back a little bit further in time, there may have been some [magnitude] 6s centered somewhere within this region that would bring the statistics up. And on that note, I think the Cape Ann earthquake was above m_{Lg} 6. I think it was about a six and one-quarter. [As a] moment magnitude, that probably puts it around 5.8 or 5.9.

Harmsen: I was using moment magnitude (fig. A9–6), so Cape Ann according to your estimates still wouldn't fit the criterion of six or greater.

Murphy: We've had an ongoing discussion, with the French in particular, about PSHA and its comparison with the historical record. We have been adamant that PSHA has been giving us a good look at what [is important for] nuclear power-plant design purposes. What you're showing would seem to imply that the PSHA process is in some sense overestimating the number of M 6 or M 7 earthquakes that would be expected. Would you comment on that?

Harmsen: NRC has a set of regulations, as Jon [Ake] explained in his talk, that requires that a broad distribution of uncertainty be explicitly included in a PSHA. The USGS doesn't have any such requirement. We just use [a model developed by] two people—Rus [Wheeler] and Art [Frankel]. You don't have the broad distribution of Mmax [that's] in the current literature. Some people say that our model is an extreme model because our distribution is a lot higher than others' [models], say, [that of] EPRI–SOG. So that's part of the explanation for why you see the differences in these magnitude-frequency curves.

Frankel: These were magnitudes that the [personnel of the national hazard mapping] project came up with. They were clarified at the workshops and we had a lot of interaction with people at the workshops. I wouldn't say it's one person's thoughts.

Harmsen: It may be a model agreed on by dozens, but it's only one model, that's the point here.

Frankel: Sure there's one model, but [it has] some distribution around the values [of Mmax].

Harmsen: But even with that, Chuck [Mueller] showed us this morning that it gives you a 1 or 2 percent change in the hazard curve. One or 2 percent really isn't a significant amount, so I still think that there is just one model of Mmax in the USGS hazard calculations.

Hough: I just want to be sure I understand how those curves are generated. Do you take the observed seismicity with the Mmax and extrapolate?

Harmsen: The curves are generated from the a-value and the b-value at each point in the analysis. The USGS models have their sources at every one-tenth of a degree in latitude and longitude. The source defines a rate of magnitude zeros, which is the a-value, and assigns a b-value to allow you to extrapolate up to higher values.

Hough: And then you extrapolate up to Mmax.

Harmsen: You extrapolate up using a Gutenberg-Richter relation.

I forgot to mention that Martin's [Chapman] early slide, where he showed the NEIC catalog frequency-magnitude distribution (fig. A5–2), is important because it shows that, at least using that sample, the [straight] Gutenberg-Richter relation fits the data up to a certain value. But it didn't go all the way up to the Mmax of 9.5. There was a huge overestimation, not only of the 9.5s, but of the [smaller M] two points. The basic assumption of the Gutenberg-Richter relation up to Mmax, [as] Allin Cornell said, is as uncertain as Mmax itself. That uncertainty is an important feature of PSHA.

Chapman: I haven't done a lot of testing on this, but I think it's pretty difficult to fit those data I showed (fig. A5–2). If you simply truncate it [the Gutenberg-Richter relation] at various values, you have a much sharper fall off than the data actually suggest. So it may be a much more complex behavior as you approach the maximum magnitude, whether it be for the whole Earth or a subregion. I don't think that a truncated model is really what's going on.

Harmsen: That is a sensitive issue for PSHA. We're talking about these larger-magnitude events and their frequency, not just one earthquake at the maximum magnitude but a distribution. The Mmax source is less important if the distribution is Gutenberg-Richter, because Mmax is the least frequent. In a sense, from a PSHA point of view the focus should be on the distribution up to Mmax, not just on Mmax.

Hough: I wonder if the b-value isn't the real problem. I think most people would argue that b should be 1 for Mw.

Harmsen: We've been requesting a moment-magnitude catalog to do the analysis on for quite a long time. I think that would be a nice thing to have, a good constraint.

Salomone: An output from the CEUS SSC project will be a moment-magnitude [[seismicity]] catalog.

Wheeler: Certainly for M 7—I don't know if it holds for M 6—I don't know what the variation of recurrence interval is for sevens in the CEUS. I'm not sure what it means to say we haven't seen one or we've seen one or two in 300 years. I don't know whether the Poisson distribution includes that kind of variability. You've got an annual return for exceedances, and you invert it to get a roughly 500–year interval for return of M 7. We've seen maybe one, at Charlevoix [1663]. I'm not sure we should be surprised if in 1663 we'd just finished a 50–year recurrence interval, and now we're in a 2,000–year recurrence interval.

Harmsen: Are we talking about time-dependent hazard analysis rather than Poissonian? It has been dealt with in Charleston, New Madrid, and perhaps in Canada.

Wheeler: We're not talking about an area with paleoseismic control, like Charleston or New Madrid. And even in New Madrid, the radiocarbon dates are uncertain enough that they would allow for recurrence intervals as short as 200 years and as long as 800 years. We use 500 as an average, but even in the three or four recurrence intervals that we have there is a substantial percentage of variation.

Kammerer: We don't know whether these are Poissonian or not, but even if they were, 300 years is a short sample. It's hard to look at these things and say we should have seen something in 300 years that we haven't seen.

Adams: Perhaps one could generate a synthetic catalog according to some rules, and then ask how many intervals look like the present. Is the present an abnormal sample of this population?

Coppersmith: What is the variation of recurrences at well-studied locations?

[Cramer provided a clarification of his answer to Coppersmith's question. The clarification comes next in double square brackets, followed by the original paragraph.]

[[Cramer: The higher the number of intervals you have, the more time you have to sample the tails of the recurrence interval distribution and hence the higher the estimate of the lognormal standard deviation (lnsd). Here are four examples from the San Andreas [fault] in California and the Cascadia subduction zone in California and Washington that illustrate this point: (1) A sequence of five recurrence interval observations from Parkfield, California, yields a lnsd of 0.35 (Savage, 1991, p. 871). (2) A sequence of tsunami deposits at Eureka, California, provide nine recurrence intervals with a lnsd of 0.43 for the rupture of the southernmost Cascadia subduction zone (Cramer and others, 2000, p. 6). (3) A sequence of similar deposits at Willapa Bay, Washington, provide seven recurrence intervals with a lnsd of 0.58 for the rupture of the northern Cascadia subduction zone (Petersen and others, 2002, p. 2,154). (4) In Southern California at Pallet Creek on the San Andreas fault, a sequence of 24 recurrence intervals indicates a natural lognormal standard deviation of 0.77 (Cramer and others, 2000, p. 6). This illustrates why two or three sample intervals in the eastern U.S. don't tell you anything about variability. You have to import that from elsewhere (world-wide).]]

Cramer: [At Parkfield there are] five [recurrence intervals], and the longer the interval you have, [the more] time you have to sample the tails, and you actually get higher lognormal values [lognormal standard deviation] near 0.6 or 0.8. That's why two or three sample intervals in the eastern U.S. don't tell you anything about variability. You have to import that from elsewhere. I've looked at a few cases, 7, 10, and 25 [sample intervals] on the southern San Andreas. The sevens and tens are [have lognormal standard deviations] around 0.5. If you get to 25 maybe you get 0.7 or 0.8.

Ebel: If you take the curves and, instead of evaluating them at M 6, evaluate them at M 5.5, then at least for the Northeast and perhaps for the central U.S. circle also, I think that you'd probably be getting pretty close to the numbers you're expecting up there. Your problem is that you picked your cutoff right at M 6, and for whatever the reason, the earthquakes are getting up close to that [magnitude] but not over M 6.

Kammerer: One thing that I think we could all agree on is that if we did more fieldwork on paleoliquefaction we'd have more data. We have one new source in 20 years. I would argue that's not because they [[additional unknown sources that left paleoseismic evidence]] are not there, but because we don't have enough people in the field looking for them.

Ebel: Tish Tuttle would claim that she's got some features in the eastern Tennessee seismic zone that are worth studying, and she's also got some features in [western] Tennessee that are not part of the New Madrid zone.

[NRC is currently funding a multi-year paleoliquefaction search in the eastern Tennessee seismic zone by Bob Hatcher (University of Tennessee, Knoxville) and colleagues.]

Kammerer: Right. We can't sit around waiting for big instrumental earthquakes. The one potential database that we can create now is paleoliquefaction to constrain magnitudes better, and to identify sources. We need to look at more funding in that area. [[(After the workshop, the NRC funded work by Tish Tuttle (USGS) to develop an updated paleoliquefaction database and to document field techniques used to study paleoliquefaction, along with estimates of uncertainties.)]]

Ebel: Just remember that the glaciation in the northern part of the country hid a lot of the evidence and made it very difficult to find what would be easier to find in other parts of the country.

Jon Ake (NRC): Should Some of the Mmax Methods No Longer be Used?

One of the things we were hoping to do today is a brief recap of the various methods that Rus [Wheeler] put together in a slide earlier today and in the framework document [(Wheeler, 2009)] (fig. A4–1). Are there methods [of estimating Mmax] that we can take off the table and don't want to discuss anymore? We'd like to focus our discussions on those methods that as a group we think are the most promising in the short term. That is, in the next few years, during which NRC, DOE and others will be producing an updated hazard evaluation for the CEUS and the USGS will be moving forward to the next set of maps.

Discussion

Hough: I would take off number 13 [Q_0]. I would take off numbers one and two [Mobs and Mobs + c], because I think we can agree that the statistical assumptions are very tenuous. Q_0 is just such a huge can of worms.

Cramer: It changes over time.

Kammerer: So it's just not viable and it's never going to be viable? Okay.

Lettis: You can't have an Mmax less than the maximum observed. The maximum observed is an important data point. And what about the 1811–1812 sequence? If we pretended that we didn't have the maximum observed, if we didn't use that as a method for New Madrid, what Mmax value would you use? [One calculated from] the length of the fault?

Lindvall: We can't do away with that method, and certainly it will be more valid for some zones than others. If you don't have a large earthquake it's not going to be very useful.

Lettis: In the Central and Eastern U.S., we need an estimated Mmax for New Madrid. What method are we going to use? Isn't one of the methods that you would consider [using] the maximum observed earthquake, just to inform your judgment?

Kimball: I think you are going to go down the path of diminishing returns. We have at least one approach, the Bayesian approach, that is a tool that can incorporate a number of the other things up there (fig. A4–1). I don't know if you want to start with that one and ask if you use this one, what others come along with it? A number of the methods could come along with the Bayesian approach, and in that way they have some viability. We could continue to discuss New Madrid and Charleston, but the most important areas to discuss are away from New Madrid and Charleston. Maybe we could rapidly come to the conclusion that the Bayesian approach is the most user-friendly [framework] to take into account the most information. Then we could discuss which methods should go with it.

Ake: It sounds like this [the Bayesian approach] subsumes this [the Mobs method], as opposed to this [Mobs] being an independent observation.

Kimball: That's another way of saying it.

Arabasz: I think we need to distinguish between assessing a distribution for the maximum magnitude as opposed to assessing the number. [Two things] seem to be mixed here: settling on a most likely value, and how to get to the tails [of the distribution]. There are methods here that overlap, but [because] you are estimating a true upper bound value, I think we need to distinguish some of these terms.

Ake: That's a very good point. In my mind, even this [the Mobs method] is something that has a distribution. Very few of these large older events have unambiguous magnitudes. That uncertainty has to be carried forward.

Hanson: I agree with Jeff [Kimball]. It seems to me this [Bayesian approach] would be a broad approach that gives what Walt [Arabasz] was talking about, a distribution, and then you have observed seismicity and paleoseismicity. You have good, well documented paleoseismic earthquakes like you have at New Madrid and Charleston with distributions and ages, and then you have random events elsewhere for which you don't have anything to constrain magnitudes, like in western and southern Illinois and Indiana. I think you do start with [[a Bayesian approach]] and you have these different data sets [that] you use to truncate the tails.

Salomone: One reason there may be some confusion is that in the first column (fig. A4–1) you may be mixing methods and inputs. One way to get a final answer might be to sort that column by what is really an input and what is really a method.

Coppersmith: Rus [Wheeler] did a good job of describing each of those in the foundation paper [(Wheeler, 2009)]. There are a couple up there that deal with rate. He makes the comparison between rate and the largest observed earthquake [in the foundation paper]. The study for EPRI did the same thing worldwide. Arch [Johnston] showed that the strain rates, when compared to [graphed against] the largest observed [magnitude], [showed no relation]. I think that the things related to rates—earthquake rate, magnitude-frequency extrapolation—are things that you might not want to call valid [for estimation of Mmax]. They are probably not promising areas for future work. Otto [Nuttli], in the 1985 workshop, argued that the 1,000–year event over some normalized area would give you the maximum earthquake. That's very much a function of the size of the normalized area. It seems like no matter [what] you do, the maximum seems to be decoupled from rate. We haven't seen anything that tells you that's a promising approach.

Ake: It seems to fly in the face of ultimately wanting to do something on the basis of the physics of how earthquakes rupture and of what makes a big earthquake instead of a small earthquake.

Coppersmith: We might like it because we know we have to calculate rates anyway for hazard analysis, but it never seems to pan out as a predictor of Mmax.

Ake: If you look at what's been done in many of the source characterizations in the less seismic areas, you find that there's almost a one to one correlation between [estimated] Mmax and the rate within a particular zone. If the rate is low, there's always a low Mmax.

Coppersmith: Let me finish the definitions. I still don't know what we call a pure statistical method. Is that a [method like the one of] Kijko and Graham [(1998)]?

Wheeler: Yes, I would put their [method] in that category.

Coppersmith: And then what's pattern recognition?

Ake: This usually implies that you've got a lot of data that you can learn with. [However,] that's the whole source of our problem: a lack of data.

Wheeler: I'm thinking of two examples for the CEUS. They were done [around] 1980. One was done by Meridee Jones-Cecil here, and the other was done by Noel Barstow and coauthors, working for Rondout. In both cases, they looked at the spatial distributions of a lot of geological and geophysical kinds of data. They looked to see whether there were particular combinations of different kinds of data that were systematically spatially associated with big earthquakes. It was worth doing [but] they didn't find anything. Elsewhere in the world, where you have more earthquakes, it might be feasible to make something out of that [approach]. It's also possible that they had too many variables to examine. [Almost thirty] years later, we still don't know which of those variables we ought to be looking at.

Lettis: Was that for Mmax, Rus?

Wheeler: No, it wasn't, but it could be. If you find that there is a characteristic combination of gravity and aeromagnetic anomalies, fault length, and age of the youngest recognized tectonic deformation that is systematically associated with six and greater, and then you see that combination in an aseismic area, maybe Mmax is six or larger there. However, it never got that far.

Lettis: Can we eliminate the magnitude-frequency extrapolation? Doesn't that depend on Mmax? You're taking [extrapolating] that down to a preselected M. Isn't that kind of circular?

Johnston: We've seen examples of it overestimating Mmax, where earthquakes are characteristic, and it frequently underestimates it [as well].

Coppersmith: I think earthquake rate and M-f extrapolation should be put on the back burner. We can't say they're invalid, but they're not promising areas for future work.

Hough: Getting back to the Mobs as a method, even for New Madrid all we really know is that there were three big earthquakes. But do we know that there couldn't be something bigger than 1811–1812? I'll put it to Arch [Johnston], could it rupture the Cottonwood Grove fault and up through the Bootheel lineament and on north? It seems to me that it would be interesting to look at that with a Bayesian approach and see what you come up with.

Coppersmith: The way to do it would be to modify the prior, right?

Lettis: Rus [Wheeler], about the distribution that you use to pick your definition of Mmax that you described to us earlier: where would you say that fits in here (fig. A4–1)?

Wheeler: It's the global analog method.

Lettis: So it's a combination of global analog and Mobs. But then it was primarily judgmental, just looking at a suite of earthquakes and judgmentally picking Mmax, and you're assigning weights to some distribution at the upper end of the tail. So the global analog is kind of a judgmental approach.

Wheeler: Yes. As a geologist, you can't get away from judgment completely. You can certainly minimize it.

Coppersmith: This is what the EPRI–SOG teams had available. [During the EPRI–SOG study,] Arch [Johnston] got up and showed us, this is where the big earthquakes are occurring: they're in extended [domains], they're in rifts, they're in multiply-reactivated rifts, and so on. It's just a judgment.

Wheeler: It's an observation also. You could call it a judgment, but it's pretty soundly based on what data we have.

Coppersmith: Yes, but there is a bunch of rifts that don't have big earthquakes. That's why I call it anecdotal. It gave the side of the story that had to do with where big earthquakes are occurring, and that's still true. But when you look at the statistics, there are lots of similar features that haven't had big earthquakes.

Wheeler: Did that statistical examination include the effect of short historical records?

Ake: They are observations that will hopefully roll up into some methodology.

Salomone: Rus [Wheeler] made an important point about North American and global [methods]. He said there would be trade offs between the advantages of getting an increased sample with the global [approach] and the negatives that might [[be introduced]]. [This would] make you want to double check against the small sample from North America. I think that's important to study further.

Coppersmith: The people in Switzerland are working with New Madrid earthquakes for Mmax, so it works both ways.

Chapman: I think with the global approach, you have to use the Bayesian approach. The issues are what data [do you need and] what work do you do to establish the prior distribution and the likelihood function for your source. The likelihood function for an individual source will depend, at least to some degree, on the maximum catalog earthquake. Your prior distribution, which you use to obtain your Bayesian estimates, can involve anything that you think is relevant. Nearly everything you've got on that list (fig. A4–1) could potentially be incorporated into a Bayesian estimation.

Ake: I can amplify that. In my view, Bayesian analysis is probably the place to go. Martin's [Chapman] quite right. What you get for your likelihood function will depend entirely on how you divide up your [study area into] source zones. Given the sparse seismicity in the East, most of the likelihood functions are going to be step functions, as Kevin [Coppersmith] showed. Reject everything below Mobs in the source zone, with unity beyond that. So the question is, how many of these other things [methods] do we want to roll up into producing the prior? It's really going to boil down to what we do for the prior. Which of these [methods] are the most promising for coming up with the prior for reasonable tectonic analogs of the CEUS?

Johnston: If you take the global analog and combine it with Mobs, then you're going for the largest sample and you'd choose your Mobs from the largest observed worldwide. That's essentially what's been done to date to get the prior distribution.

Coppersmith: The question is, how are we going to modify the prior? The example we gave in 1994 was for Charleston [South Carolina], where we used physics-based approaches to modify the prior. At that time there were some faults that had been proposed, and their lengths would imply certain magnitudes. We used that to modify the prior. There could be other information. Paleoseismic information could come into play. Do we see the distributions of liquefaction features that would indicate such and such magnitudes? That could modify the prior.

Johnston: But those are all ways to estimate what Mobs really is. Most of your maximum earthquakes, even going global, will still be small.

Coppersmith: Yes. Let's assume we've made the bias corrections to truly estimate the upper bound magnitudes. Then that becomes the prior and how do we modify it? That's where you can incorporate other information. You could incorporate physics, or paleoseismic [results]. That's how the Swiss did the problem. They took Allin's [Cornell] distribution for non-extended crust (all of Switzerland is non-extended). They had broad distributions and modified them on the basis of things like the length of a zone of seismicity and its depth distribution [of earthquakes].

I would add paleoseismicity as a method on that (fig. A4–1). Right now it's used as an absolute basis for comparison.

Hanson: It's just [a way to determine the] maximum observed [magnitude], except instead of instrumental [data] it's paleoseismic [data].

Ake: I would argue [the Mobs method] has to include paleoseismic [magnitude estimates].

Lettis: What about a discrete fault source like [the] Meers [fault]? Would you still go through the prior and likelihood function, or would you just say that we have a fault scarp, trench data, displacement information, and fault length? I would just use that.

Ake: I would fall back on those physical properties.

Lettis: That would be a non-Bayesian approach. So, not all sources in the CEUS would use a Bayesian approach. The discrete sources would not [not be incorporated into a Bayesian framework] but there are only a couple [of discrete sources].

Chapman: I think that the Bayesian concept of doing the estimations is the framework that you want to refer to.

Coppersmith: You'd be setting up a prior that is so strong that the likelihood function [has little effect].

Lettis: So it could still be done under a Bayesian approach.

Ake: What we've been talking about mostly is estimating Mmax in source zones, as opposed to discrete faults that we can identify. Would we approach the problem differently [if instead we used] a large number of discrete source zones that make up the eastern U.S., instead of the approach that Rus [Wheeler] and Martin [Chapman] are taking that uses just two large source zones in the CEUS?

Coppersmith: They are two different activities. The stated goal of the EPRI–SOG study was to understand and estimate spatially varying Mmax. If in fact it [the goal] is to develop a distribution that can be used for an entire region or continent, that's a different activity, and the way it's estimated would probably be different as well.

Hanson: The people in this room need both.

Wheeler: For different purposes.

Ake: That's why I'm asking you to think about this overnight, so the outcome of this meeting will be useful to anyone who's trying to do this and to answer these questions.

Tuesday morning, September 9

John Adams (Geological Survey of Canada): GPS Constraints on Hazard in Eastern Canada (Appendix 10)

This is using geodetic information in the East to constrain Mmax (fig. A10–1). It's a published paper in the Journal of Geophysical Research. It's related to using logic trees to analyze geodetic data. We are here in the lower St. Lawrence (fig. A10–2), with the Charlevoix active zone [lower red arrows and underlying grey dots] and the Lower St. Lawrence cluster [upper red arrows and underlying grey dots]. These [sets of red arrows] are the strain results that Stéphane [Mazzotti] has derived from ten years of data. Within the large polygon [green dotted line] the strain vector is quite small [dark green arrows in center of figure]. In the two active zones it's quite large. Although the magnitudes are not large, the fact that the orientations come out quite similar gives a little bit of hope that there's true information there. One way of interpreting this is that there's shortening across the St. Lawrence Valley on the order of 1 mm/yr.

With that geodetic signature, and making some assumptions about the strain model, the earthquake recurrence statistics, and the seismogenic thickness, on which we've got some handle from earthquake magnitudes, you can talk about the implications for the magnitude and recurrence of large earthquakes (fig. A10–3). If we assume a Gutenberg-Richter relationship, here's the Charlevoix data (fig. A10–4, upper graph, green triangles), including the big earthquake in 1663 [right-most green triangle at M 7.0; table 4]. In order to accommodate that geodetic shortening and its uncertainty, you need Mmax in this range [shown in yellow]. Here's the best estimate of Mmax [7.8] and there's the uncertainty [7.2–8.5]. The lower end of the uncertainty probably doesn't have a lot of information, because we probably wouldn't go much lower than that anyway, having had an M 7.0. The upper end is actually horrendously large. It is an upper bound, and it's a valid way of estimating how much bigger Mmax can't be. It's just not very useful. For the Lower St. Lawrence [lower graph], the uncertainties are larger. You can see at the bottom here that the best estimate is more or less the same number [7.3], but the range is very large. Mmax could be quite small or it could be quite large.

That's if you have a Gutenberg-Richter model, and it actually makes sense. Stéphane asked what happens if you use a characteristic model (fig. A10–5). With a characteristic model [for Charlevoix], we have a Gutenberg-Richter relationship coming down to about here [upper graph, downward-concave dashed black line] and then stopping [at M 6.0] with characteristic earthquakes being thrown in [yellow ellipses with red centers]. The red dots are the median rates of M 6.0 [upper left red dot], M 7.0 [middle red dot], and M 8.0. Any one of those is a valid model. The yellow ellipses give you some idea of the uncertainty. From the historical record we can remove the M 6 model; we're not having M 6 earthquakes every decade. We can't quite rule the M 7 model out, because if the 1663 earthquake were M 7 and occurred every two hundred years or so as a characteristic earthquake, it would more or less accommodate the geodetic shortening. The other alternative is that 1663 was much larger and much rarer.

If we move to the Lower St. Lawrence, which is a much more poorly known region, it has [[had]] a largest earthquake of [only] M 5 [lower graph]. The prediction is that an M 6.0 should occur about every 10 years [upper left yellow ellipse and red dot]. We haven't seen that [one or more M 6.0 earthquakes]; [therefore, the hypothesis of M 6.0 characteristic earthquakes] should be eliminated. This one here [middle ellipse and dot], an M 7.0 every 500 years, we can't eliminate.

So that's the method. The conclusions are very inexact at the moment (fig. A10–6). Presumably they will get better with time, but I think they do provide another way of estimating Mmax.

Discussion

Wheeler: What excites me about this is that it's a possible way to subdivide a very large region into areas that have different Mmax values.

Adams: Of course, one of the assumptions is that what you measure over a decade is what you would get for the long-term seismic hazard. If we look at the data here (fig. A10–2), this very large area [green polygon] has a very low strain rate, so in fact you'd end up with infrequent large earthquakes. The area between [the red arrows] is presumably just as loaded, and it has the same Paleozoic faults. We don't have enough data yet to analyze whether small subareas are just the same. Stéphane's working on it. He's putting in about a dozen stations every year and visiting them every two years, with the idea of getting a geodetic strength signature across the St. Lawrence.

Wheeler: Are there enough liquefiable post-glacial sediments [to search for evidence of large prehistoric earthquakes]?

Adams: Yes. Tish Tuttle worked there last summer and the conclusions were [[these]]: she saw liquefaction structures in Charlevoix representing paleoearthquakes, and she looked very hard here [upstream, or southwest, from Charlevoix] and didn't find anything. If there had been a very clear one [liquefaction structure], or a lot of them, she probably would have seen them.

Ebel: I'd like to take a few minutes this morning to talk about the 1663 earthquake [table 4]. As you know, I think it was not just M 7, but perhaps significantly above M 7. My question here is about the seismicity rate as opposed to the strain rate. One of the things that I know for the northeastern U.S. is that the strain rate, as inferred by the people who are doing the geodetics throughout eastern North America, is probably no more than about 10^{-11} or so per year. In California it's on the order of 10^{-7} in the active zones. And yet our seismicity rate, the rate at which small earthquakes are occurring, is only two orders of magnitude less in the northeastern U.S. than it is in California. So we have this discrepancy between seismicity rate variations between East and West and estimated strain-rate variations between East and West. Can you comment on that?

Adams: Part of the conclusions of Stéphane is that the geodetic and seismic strain rates agree quite well in these two zones (fig. A10–6). The data he was analyzing is fairly tightly spaced around those active zones. Maybe it's because the strain signature is actually very localized. In a profile across the St. Lawrence, going from northwest to southeast, you have a number of data points [along the profile], and the null hypothesis that they change smoothly is okay. You could also accommodate that with a very large strain rate in a smaller area. The data points are not sufficient to choose between the two hypotheses [of uniform or localized strains]. If, in fact, the strain rate is low in the large area, which is what the large polygon looks like, and strong in the small areas, you have to have quite a dense [Global Positioning System (GPS; table 2)] network to find those areas, and I'm not sure that's been done other than in New Madrid, and that's controversial. If you want more details on this, you really ought to talk to Stéphane and read the paper.

Li: How many years of data are there?

Adams: I think it's on the order of ten now. It's a little variable. These are campaign GPS [data point that are collected periodically], not continuous [records]. Stéphane is putting in continuous GPS measurement [stations], which give you a lot more redundancy and a better signal to noise ratio.

John Adams (Geological Survey of Canada): Mmax in Canada's Fourth-Generation Hazard Model (Appendix 11)

I'd like to talk about the Mmax experience in Canada's fourth-generation hazard maps (fig. A11–1). In a lot of ways it's a counterpoint to the USGS maps. Rus [Wheeler] and I have been throwing ideas back and forth for about 20 years about this. I won't say there's full agreement, but in general the GSC approach and the USGS approach are fairly similar for Mmax. We're dealing with our fourth-generation model. It was developed in 1994–1997 and finalized in 2003. The documentation is in this reference at the bottom [[of fig. A11–1]]. It was only implemented in the 2005 [building] code, so it's quite an old model in terms of the ideas that went into it.

We have a lot of big earthquakes in eastern Canada (fig. A11–2), a lot more than in the eastern U.S. In the U.S. you're focused on Charleston and New Madrid, and then you cut them out and worry about the rest (there isn't much of the rest), whereas we throw ours into a large box and ask the question: how big? We've had big earthquakes on the eastern [continental] margin [the two largest red dots offshore], so we throw all of those inputs into a [frequency-]magnitude curve [red, curved line on the graph]. The best estimate [of Mmax] is M 7.5. It could be M 8.0, or perhaps as low as M 7.3, which is the Mobs [[maximum observed earthquake magnitude]]. Most of those earthquakes are actually extended margin ones. You'll see that even back in 1994 we put uncertainty on Mmax into the model, and it was a bit of a shock that the USGS hadn't. I'm not sure it makes a lot of difference but there should be some sort of uncertainty in there.

One of the other things we did back in 1997, although only published in 2006 in a peer reviewed journal, is look at the stable cratonic cores worldwide (fig. A11–3). These are more or less the same as the non-extended SCR regions that we've talked about. They are quite a bit more rigorous than, for example, one that was shown [(the Chinese craton in fig. A7–5)]. For example, here's our [version of the] Chinese craton: that little polygon there. We cut a lot of things out. We stayed 200 km away from the nearest things that we were worried about [not being a craton] on the basis of young activity. We [didn't do it] perfectly because our knowledge is not perfect. For example, the Timiskaming earthquake [red dot east of the Great Lakes in southern Canada; table 4] is in the Canadian Shield, although I would argue that it is actually not a stable craton earthquake.

There were 67 earthquakes that passed [a] completeness [test]. They gave a really well-defined magnitude-recurrence curve that we took down to an Mmax of M 7 (fig. A11–3). It's hard to go much smaller than M 7 because Tennant Creek [table 4] was M 6.7–6.8. There were three events. If you add them together you get about M 6.9. So it's not unreasonable to [assume] that you could get up to magnitude 7. Our conclusion was that anywhere on the continents you could see [a] magnitude 7.

This [[slide]] is a bit of a digression from Mmax. For the hazard estimate we took rates of activity from the world, North America, and the part of Canada that's not in a source zone, weighted them, and then came up with a hazard calculation (fig. A11–4). The uniform hazard spectrum for our floor level, which for Winnipeg is here (lowest curve in fig. A11–5), is more or less what we'd used for, say, Minneapolis. A deaggregation plot shows that a lot of the contribution to that calculation comes from magnitude 6 events [near the site of interest] (fig. A11–6). [That is, the highest parts of the third and fourth histograms in the figure are at magnitudes between 6 and 7 and at distances around 50 km.] There is truncation because our upper-bound magnitude is 7. If you were to make the upper-bound magnitude seven and one-half, there would be extra contributions along here [there would be a sixth histogram in the back of the graph] that would increase the hazard slightly. The [[deaggregation shown]] was done for 2 percent in 50 years of hazard. The floor model of seismic hazard (fig. A11–7) was applied to a lot of areas, obviously shield areas, but not [[to]] any areas where we had clusters of earthquakes [blue area in southernmost Saskatchewan]. It was applied to areas in western Canada where there was very low seismicity in the interior plateau of B.C. [pale grey area in central British Columbia]. We came to the conclusion that nowhere should have a value [of hazard] lower than the floor value.

Let me take you through a few of the source zones that go into the [[full]] model (fig. A11–8). This is a Mesozoic extended-margin source. It's the whole eastern continental margin. It goes roughly from Boston up to Greenland, 3,500 km. We've got two big earthquakes in it, Grand Banks and Baffin Bay [table 4; respectively, the southern and northern largest red dots offshore in figure A11–2]. The biggest magnitude observed is about 7.4. We gave three branches to this: M 7.5 with a very heavy weight, 8.0 with a low weight, and 7.3 (which has actually been observed) with a low weight. What I would say is that there are plenty of potential large faults out there that could give you a magnitude 8. We just don't know whether they are going to rupture like that, but it's not implausible.

Moving to the Paleozoic rifted margin, which we think is active, this is the Charlevoix source zone (fig. A11–9). We argue that we had a magnitude 7 [1663], and that's what John [Ebel] will talk about later. So the observed magnitude is 7. We put in branches for M 7.5, with the highest weight, M 7.7, and M 7.2. Unlike the eastern margin, those [magnitudes] are m_{bLg}, so we're playing a bit loose with the magnitudes at this point. It would be more sensible to make those magnitudes on the M_S scale. I would say that there are enough potentially large faults to do this. We have lots of rift faults along the St. Lawrence, and the only real question is whether they are segmented. But even if they are segmented, more than one segment could go. The problem there is that we might have a large rupture. That doesn't seem implausible.

Let's move into the interior (fig. A11–10). This is the slightly extended crust and it happens to be the area that includes the southern two Great Lakes. The [largest] observed magnitude there is roughly 5, depending on whether you include the Attica earthquake [table 4], which was five and one-half or five and one-quarter. We took, [[for consistency]] with the stable craton, magnitude 7.0 with heavy weight, 7.2 with low weight, and 6.8 with low weight. We don't really know whether there are potential faults that are large enough [to generate earthquakes of these sizes], but the stable craton says that you can get earthquakes this large even if you don't identify faults that are large enough. The hazard happens to be pretty insensitive to Mmax at the national building code's [annual] probabilities, so to some degree we don't care so much about Mmax here. At least we've got a consistent hypothesis: the Mmax's here are consistent with our model of the stable craton core.

Here's one place where we weren't consistent (fig. A11–11). This is the lower St. Lawrence region, which is the other area I talked about with respect to geodetic strain (fig. A10–2). The [largest] observed magnitude here is around M 5.0. So the [data constrain the] magnitude-recurrence curve down to here [less than M 5.0 in figure A11–11], and then there's a large extrapolation. We took the same best and upper values as at Charlevoix, but for the lower value we chose M 6.0, representing the fact that the biggest historical [earthquake] was magnitude five. We've got those same potential large faults there, the hazard is not insensitive to Mmax, unfortunately, and that [M 6.0] is inconsistent with our stable craton model. We would probably change that the next time around. It'll go up to M 6.8 to be consistent with the stable craton Mmax.

I'm not convinced that M 7 is actually the right value for Mmax in the stable craton, but Trevor [Allen] will be talking about that.

Discussion

Arabasz: You concluded that you can get a magnitude 7 anywhere in the craton from the recurrence period. Does that apply only in regions where there is seismicity?

Adams: The assumption is not really from the recurrence period but from the global analogs. One of the things I would have said, when I put up that magnitude-recurrence curve for the whole of eastern North America, is that we tailed it off around seven and one-half, but you could take it down to eight or nine. I don't believe you can extrapolate the magnitude-recurrence curve [[much further. But]] you don't know where to stop.

Arabasz: But my point is that you're analyzing a collection of seismicity in this aggregate compilation. Let me ask again, although I think you've answered me by saying global analogs. [Are you saying that] anywhere in the stable cratonic core, even in areas where there's no seismicity, you'd still put an Mmax of 7?

Adams: Yes. Of course, is that the correct value or not? I was really taken by Walter Mooney's stuff because I've been involved with some Ontario seismicity that seems strange. But it may be typical shield seismicity: it's very shallow. We have a paper in press in the Bulletin of the Seismological Society of America [showing that] there does seem to be a relationship between the activity rate of small earthquakes and the seismic velocity down at, say, 150 km [[in the mantle]]. We strongly support what Walter was saying. It looks like an interesting thing to do. We see it the other way around, if the velocity is slow and [the rocks are] strong, we see higher levels of small earthquakes. Where the [deep crustal] roots are fairly warm we see low rates. In a way we disagree with him but we think the approach is good.

Wheeler: We need to keep in mind the distinction between the value of Mmax and the frequency of Mmax. We asked, what's Mmax in an area where there's no seismicity? It could still be high, but with a long recurrence interval. We can predict one [value or frequency] from the other with a magnitude-frequency diagram, but [[we need to know one of them to determine the other]].

Adams: The other thing I would say [to follow that point] is that if the activity rate is low, then it doesn't matter very much which Mmax you choose. In a way, you can choose a big number and it doesn't affect what you're doing [[that is, the hazard values at an annual probability of about 0.0004]]. That way you can have a consistent hypothesis for everything you're doing, rather than do it case-by-case.

Coppersmith: That's only in the context of building codes[, which consider a higher probability of exceedance than is used for critical structures].

Kimball: We have the same controversy in the Central and Eastern U.S. Are people using the building-code hazard curves for more critical applications [in Canada]?

Adams: What we try to do is dissuade them from using this model for one-in-ten-thousand-year calculations. The reason we don't like it is that we don't think the fault [[or other source]] characterization is robust enough. In other words, choices are being made which are good enough for a building code at 2 percent [probability of being exceeded] in 50 years, but not necessarily good enough for a site-specific assessment of hazard. I'm wiggling a little bit here but that's what we tell the people [who ask about annual probabilities of] one in ten thousand years. For the nuclear plants it's lower than that now.

The message we give is that we're doing national maps. We can't possibly do all of Canada at the quality you need for your critical plant. You should be doing a good site-specific assessment. If [your] numbers come out lower than [ours] at this probability, then we'll look at it and if it looks okay we'd accept it. It's not rigid at all.

Kimball: Mmax could become more of a sensitive issue.

Adams: I don't think it has been in the past. I don't think it's been the driving issue. A lot of it has been the source-zone boundaries. If you cross a source-zone boundary and you're looking at high frequencies, the gradient could be so steep that it could change the hazard a lot. The justifications for the boundaries in Canada are judgmental.

Ake: That's sort of like in the Central and Eastern U.S.

Adams: By the way, we are still using Cornell-McGuire source zones, even for relatively small clusters [of seismicity], whereas the USGS is using smoothed seismicity. Steve Halchuk, several years ago, ran the smoothed seismicity and showed that our H model, which is the small clusters, and smoothed seismicity come out with very similar seismic hazard. We could have saved ourselves the trouble of doing that source model, basically.

Wheeler: You'd get away from the steep gradients.

Adams: Yes, and I think we will implement some sort of smoothed seismicity [[in the future]]. We also want to look into fuzzy boundaries. One of the problems is that you wind up with subjective decisions about parameters, where in the end you can come up with any steepness of gradient that you want by choosing different values. Unless there is a clear consensus that, for instance, you should use 50 km [[as a smoothing kernel]], and I'm not sure that there is, then we will probably do what seems right.

Coppersmith: As you just said, you can escape source-zone boundaries by smoothing, but in fact there's a subjectivity related to smoothing that will lead to gradients, too. There can be adaptive kernels that take into account uncertainty in the locations of epicenters. There can be elliptical kernels that take into account some preference for orientation. [There are] all types of tools in the toolbox.

Wheeler: But still, it's [smoothing is] a more clearly defined, tightly focused question, for which it's more possible to explain with a greater level of detail what you've done.

Coppersmith: They're all tools to express your degree of belief in things that they [source-zone boundaries, earthquakes, and faults] share. If you believe in your locations of past earthquakes or the locations of features, then you can [express] that with tight source-zone boundaries, or with small bandwidth.

Wheeler: A lot of source-zone boundaries are likely to be harder to explain than the mathematics that you choose to represent smoothing.

Coppersmith: It depends on who you are. To some people, source-zone boundaries are much more intuitive and easier to explain than smoothing. What are you smoothing: a[-values], b[-values], or both? Volcanologists, for example, are used to doing spatial smoothing of events of various ages. They have very complicated smoothing operators that are adaptive, but it can be very complicated to explain.

Adams: If there's anything that I would leave you with, it's that we would choose a large Mmax in Canada, based mostly on the stable craton arguments. It's hard for us to see that it would be smaller than M 7 anywhere. We're sort of like the USGS, coming from the top down [to divide the global earthquake catalog]. The segue into Trevor [Allen's presentation] is that the Australian paleoseismic catalog suggests that, in fact, that [[Mmax]] may not be large enough.

Coppersmith: But you do have a distribution?

Adams: Yes. By the way, we use a three-point distribution for many of our parameters. For some we've decided that's not enough. We'll use about five, in order to get smoothing of the percentile curves. I think on the whole we won't have as large a range of Mmax next time around, because I think that 6.0 (fig. A11–11) was left over from the way we were thinking in 1992 and it wasn't cleaned up at the beginning of 2000.

Lettis: Is the M 7 for the stable craton based on the global analog data?

Adams: That's right. The individual earthquakes that really drive that, I think you could say, are from Tennant Creek, but there are other ones. But at Tennant Creek, the magnitudes were good, and there were three of them that, summed up, got you close to M 7. You could find excuses for other earthquakes which might otherwise turn out to be the biggest, but I'm not sure that their data are as good as for Tennant Creek.

We have a habit of drawing boxes around things that we think are different. For the Canadian activity rates, we've given weights of 0.2 where we've taken things away [[Figure A11–4, model on the right]. I think that's a bad philosophy. We take things away, and then we assume that what's left is actually a valid representation. But the very act of taking things [[earthquakes]]away drives the rate down. So I don't think that's a very sound way of doing it. You have to be very careful to not subtract from the large picture by pretending that you understand the small picture. The discussions we had yesterday about New Madrid and Charleston are exactly that extraction process. I'm rather unhappy with extracting New Madrid because I'm not so sure it's different from Charlevoix. In fact, there's probably a continuum. Maybe New Madrid is a little more special than Charlevoix, but maybe Charlevoix is a little more special than the lower St. Lawrence. Where in that continuum do you draw a line?

Trevor Allen (Geoscience Australia): Neotectonics in Australia (Appendix 12)

I'm giving this talk for Dan Clark of Geoscience Australia (fig. A12–1). Dan has more than doubled the number of known neotectonic features on the Australian continent. This provides a test bed for other stable continental regions and shows that we can have big earthquakes in stable continental regions and [that] they may be possibly used as analogs.

The Australian plate is moving northward at 60–70 mm/yr (fig. A12–2). It is the fastest moving plate in the world, and it is also connected to India, where there's a continent-to-continent collision. The stress regime is fairly complex. Unlike most stable continental regions of the world, our stress regime doesn't necessarily parallel the absolute plate motion. In the northern regions, the stress direction is pretty close to the direction the plate is traveling; however, in the southern part of the country it becomes more complex. We have northwest-southeast trending stress in the eastern part of the continent, which [changes] to east-west compression in the western part of the continent.

By zooming in on the stress regime a bit more and overlaying the seismicity (fig. A3–3), we see here [yellow dots in the north between the north-south oriented blue arrows] the Tennant Creek sequence that John [Adams] mentioned earlier. Here [clusters of yellow dots in southwest Australia] is the southwest seismic zone where the majority of our seismicity has been over the last 50 or so years.

Here [abundant onshore yellow dots in southeastern Australia and extending northward along the coast] are the eastern highlands where we have quite a bit of seismicity although there hasn't been any earthquake above five and one-half there since European settlement, but there is certainly evidence for larger earthquakes occurring in this area [green triangles in figure A12–6]. Also in the Adelaide fold belt, Flinders ranges, and Mount Lofty ranges [northward trending, elongated cluster of yellow dots at the blue arrows that trend east-west] is a lot of concentrated seismicity. This region is where the stress regime changes [trend] from northwest-southeast to east-west. [[Published hypotheses suggest that high heat production in the middle crust acts in concert with the favorably oriented structure of an inverted rift to make the Flinders and Mount Lofty ranges particularly susceptible to strain localization.]]

It has been hypothesized by Mike Sandiford at Melbourne University that the whole Australian continent is tilting to the north as it subducts to the north (fig. A12–4). The bottom part of the continent is actually popping up out of the water. Some evidence of this is across the Nullarbor Plain in the Great Australian Bight region (fig. A12–5). This section shows [three sets of sea cliffs that have been successively cut and uplifted through time], getting older as they get farther into the continent. This is an extremely flat part of the continent as you can see in this image. In one part you can drive for 400 km without turning or going up any significant elevation.

Here is the historical seismicity (fig. A12–6). Overlaying that are Dan's neotectonic features. These are the green triangles that show the areas he's identified as having some neotectonic feature, mostly from digital elevation models. All across the continent he's identified more than 200 of these features. In some places the neotectonic features agree well with the seismicity [red rectangle in the Flinders and Mount Lofty ranges]. However, in some places like southwest Western Australia we see fairly obvious features in the landscape that don't coincide at present with the seismicity [red rectangle]. This suggests to us that in some parts of the continent the seismicity is migratory. One part of a region might experience very high levels of seismicity for a short time, and then that seismicity will migrate to other parts of the region.

These are the surface expressions of some of the features that Dan [Clark] has identified (fig. A12–7). The Narrabeen scarp ranges up to 3 m in height. The Dumbleyung scarp was trenched just before I left Australia [in 2006]. At that time they thought it was [formed by] a single rupture, which would have led to 4 m of uplift. I think they've revised that now to a couple of ruptures, so I apologize to Walter [Mooney] for suggesting that some of these earthquakes may have been up to M 7.3–7.4. These magnitude estimates now have been revised to perhaps M 7.0–7.1, based on the Wells and Coppersmith relation [(Wells and Coppersmith, 1994)]. We see some very subtle features here [on the Merredin scarp]: an uplift of about 1.5 m over a [[100]]–m range [horizontal distance].

Dan has discovered many of the scarps in southwest Australia from digital elevation models (fig. A12–8). We can clearly see that there is some sort of feature running generally north-south [identified by red arrows] [[in both panes and these are]] perpendicular to the current stress regime. They've gone [into the field] and verified that many of these features are actually fault scarps. Unfortunately he hasn't been able to trench a lot of these features to find out how many ruptures might have gone into generating the scarps, but many of the scarps he's identified are probably [results of] multiple ruptures.

<u>Stirewalt:</u> Is the material at all susceptible to liquefaction?

<u>Allen:</u> [It's] quite a dried out area. The reason Dan is able to go out and find these features is that there is very little rainfall. They are preserved in the landscape a lot better than they possibly could be in the eastern U.S.

Moving across to the Mount Lofty ranges, in central South Australia, we see that the expression of the neotectonic faulting is quite different than in Western Australia (fig. A12–9). In Western Australia, we tend not to see a great buildup in topographic relief, whereas in South Australia the seismicity tends to be concentrated on just a few [[proximal]] faults and the buildup in topographic relief is quite high. Dan mentioned that on the Milendella fault, which is in the bottom photograph, he's found evidence of a single rupture that may have ruptured 7 m. At the moment he's not quite sure whether to believe that but he can't find any evidence that would suggest that it's not a single rupture. If that's the case, then it could be up to a magnitude seven and one-half earthquake. Here's the regional topography (fig. A12–10). The Mount Lofty ranges are bounded by these very large faults, which may be capable of producing very large earthquakes.

Based on this neotectonic information and seismicity, Dan [Clark] developed these neotectonic domains (fig. A12–11). We have the western and central domain[s], and the Nullarbor domain—which has essentially no seismicity at the moment—the Flinders/Mount Lofty ranges, and Eastern Australia. You can see that he's also added some extended margins around the edge of the continent. Historically we've had quite a few earthquakes in Western Australia, up in the northwest shelf (fig. A12–3). There may have been an M 6 or so up in that area, but nothing around M 7.

<u>Stirewalt:</u> There's a hinge line that crosses the continent and you have the south coming up and the north going down (fig. A12–4). How do you get extended terrane in the north when that's being subducted?

<u>Adams:</u> That's an old passive margin.

<u>Allen:</u> These extensional margins [offshore from the south coast] are from the breakup from Antarctica about 60 million years ago. [[To the best of our knowledge, the entire Australian plate is in a state of compression. The extended terranes are Meozoic in age, and currently are being reactivated in compression. In the case of northwestern Australia, the compression is associated with a large component of strike slip.]]

This is an example from the eastern domain that Dan is highlighting (fig. A12–12). It's the Cadell fault scarp, which is an 80–km long feature. About 50,000 years ago, the Murray River, Australia's largest river, flowed [westward] across [where the scarp is now]. This is the paleochannel [brown curved line crossing the gray zone of highest uplift, about 40 percent of the way up the digital elevation model]. The river then flowed to the north. After repeated episodes of seismicity, it now flows south down [along the east foot of] the fault and then cuts across the fault [[immediately south of the bottom of the elevation model]] and the fault continues southward on the other side of the river. This is a seismic line that was shot along here [red line near the south end of the elevation model]. There appears to be 100 m of displacement across the fault, which is not evident at the surface, and there appears to have been two episodes of seismicity; there's been an initial buildup of relief, and then a period of quiescence and erosion, and then there's been additional uplift.

This is an interesting area of Australia (fig. A12–13). I pointed out that historically there hasn't been a lot of seismicity in the area. It's the Nullarbor plain. This terrace [midway along the coastline] is Pliocene in age, less than two-million years old. We're going to look at features in here [the red box]. You have the Miocene terrace up here [gray area] and [topographic] profiles 1–3 across the [scarp] (fig. A12–14). The terrace is less than 15 million-years old and you can see that there's evidence to suggest more than 25 m of [relative] uplift [across the scarp]. If we go down to the younger, Pliocene material, you see that there's also been activity across here [the scarp on the Pliocene surface] with about 10 m of uplift. This is currently an area where there is essentially no seismicity. The length of this feature is on the order of 100 km. If it were to go all at once, that could result in a very large earthquake.

Here in the western and central domains, the Meckering earthquake [table 4] had a surface-wave magnitude of 6.8 (fig. A12–15). There's a fault scarp related to it. This [short black line about 100 km north of the Meckering scarp] is the Cadoux earthquake's surface rupture [table 4] and Calingiri [table

4] [is] right here [black dot about 100 km northwest of the Meckering scarp]. All the red features are those that Dan has identified with the digital elevation models. He's significantly built on the known neotectonic database in southwest Western Australia. Many of these features are on the order of 45 km long. There's conjecture about whether any of them could rupture entirely in one earthquake. If one did, that would be equivalent to a low M 7, using the Wells and Coppersmith relations. This is in the Yilgarn craton, which is some of the oldest crust in the world.

If we go to the northwest corner and look at some of the extended margin of western Australia, we see some evidence of neotectonic activity (figs. A12–16, A12–17). [[These terraces are being uplifted on the back of an anticline. The anticline is in a compressional jog of a larger strike-slip fault system. The Carnarvon Basin formed as the result of Paleozoic and Mesozoic extension, and the extension lasted into the Miocene. However, since the middle Miocene the extensional architecture has been inverting in the current compressive, strike-slip stress regime.]]

In eastern North America, among faults that have been trenched, many show a series of large earthquakes within a relatively short time span (fig. A12–18). Going further back in time, you can identify additional large earthquakes. Here's the database that Dan has compiled for the Australian faults (fig. A12–19). The Tennant Creek earthquakes are [in the sixth line down], and from there up are all the western and central Australia scarps [that have been trenched]. On the Hyden fault, on which Tony Crone [USGS] worked in the 1990s, there was a series of earthquakes here [upper red ellipse] and another perhaps more than 250,000 years ago. The events that cluster on a fault occurred [three red ellipses] during a relatively short time period, and there's evidence to suggest that these faults have been reactivated from longer ago.

This is something I grabbed from a paper by Mark Leonard [Geoscience Australia] and Dan [Clark] that uses Dan's neotectonic database, which is shown here in red (fig. A12–20). The black dots are from the historical seismicity of southwest Western Australia. What we see is that the rate of historical seismicity is about an order of magnitude higher than what the neotectonic catalog suggests. Dan is giving a completeness of M 7 and greater in this region over a period of about 100,000 years. This (fig. A12–21) is the comparison [of the Australian neotectonic catalog] with the [Australian Shield part of the] SCC [stable continental crust; table 2] catalog that John [Adams] and his colleagues at the GSC put together [CSS in the figure caption apparently should be SCC]. We can see that there is some relatively good continuation [of the black dots into the red dots], but it still suggests that the current seismicity of the Australian shield is about a factor of two higher than what we can see in the neotectonic record. [Now] we are looking at the recurrence rates for the entire Australian catalog, including the Flinders ranges and eastern Australia (fig. A12–22). The rate of neotectonic activity appears to be fairly well correlated with the current Australian catalog. Bear in mind that these tectonic regimes are very different [[in crustal age, structures, and stress orientation, although they are all in compression]].

Wheeler: Are the red dots just from southwest Australia?

Allen: Yes, they are. And finally, this is a comparison with the worldwide stable continental crust database (fig. A12–23). The black dots are the worldwide [historical] database and the red dots are the southwest Australian neotectonic record. This suggests that, in the prehistoric record, we see higher seismicity rates in Australia relative to [the historical record of] other stable continental regions in the world (fig. A12–24).

Discussion

Ebel: One of the things that's said about earthquakes in eastern North America is that they occur on preexisting zones of weakness. Is the same statement said for Australia?

Allen: I guess we'd like to think so, but unfortunately we don't know where those preexisting zones of weakness are. For example, with the Tennant Creek earthquakes, we had no idea there was a fault there until the earthquakes eventually occurred.

Lettis: That's an excellent question, though. Are the neotectonic features, the scarps, reactivating old faults, or are they new faults cutting across the craton?

Wheeler: If you look at geologic maps of exposed basement, like the Adirondacks as shown on the New York State geologic map, the thing that's always struck me is that, constrained by the map scale, there are faults of all lengths and all orientations everywhere. If that's actually true, then it really doesn't matter whether we're reactivating or making new ones.

Adams: As you fly over the Canadian Shield where there's not much vegetation, you see lineations all over the place. So there are lots of old faults, and there may well be active faults hiding in that; we don't know.

Coppersmith: Here's a place for detailed calibration. Look at the trends and see whether the scarp goes into a preexisting bedrock fault.

Lettis: I am not a believer that all faults are the same everywhere. There are different types of faults, some of which may be a deeper crustal structure, others may be a shallow folded feature, or they might be more aligned, or have a displacement history that [makes them] more likely to reactivate. They may be healed or unhealed, they may have been uplifted after ductile slip at depth and now they still look like a fault but are not easily reactivated, or they may always have been in the brittle regime. This is the first time I've seen the potential for a great database to test what types of faults are being reactivated in a cratonic environment. We can't do in the U.S. or Canada what you can do in Australia because of the [exposure] conditions.

Mooney: Rus [Wheeler], you seem to be contradicting what you said yesterday. Yesterday you showed us a map of the Paleozoic rifted margins with their big normal faults with a particular orientation, and then the accretion of terranes like the Appalachians makes thrust faults with a particular orientation. Now you're saying that there are faults of all lengths and all orientations, the exact opposite of what you told us yesterday.

Wheeler: Part of the wiggle out is what Bill [Lettis] just said. The other part is that many of the exposures of basement in the maps that I've looked at have a diameter of about the rupture length of what we'd consider to be an Mmax earthquake, say seven or seven and a half. But you can see that there are lots of small ones and medium-sized ones as well. So I'm not sure that the two generalizations get close enough together to compare and contrast them.

Mooney: Maybe the difference is that I look at deep reflection profiles that image the entire crust, and you look at geologic maps. We need to look at these data together. On a deep reflection profile, you see a particular fabric of the crust with definite orientations to deep faults. The crust is not randomly faulted. It matches the tectonic evolution very well.

Adams: It's not random, but there are lots of them.

Stirewalt: Even on a geologic map, they aren't random and certain things are preferential to being rejuvenated and reactivated. No matter how many lineations you have, some are going to be reactivated, and some are not because they are not positioned correctly in the stress field.

Wheeler: I did make some measurements about a decade ago. I looked at just the exposed Mesozoic basins along the eastern seaboard of the U.S. and Canada. I measured length and orientation. There's no surprise that the longest ones [faults] are oriented parallel to the long axes of the individual basins. They have to be that way or they couldn't fit into their basins. But as you go down in fault length, you find all orientations. That's a better expression of what I just said. I think it's consistent with what Bill [Lettis] was pointing out and what Walter [Mooney] was saying.

Li: When you have different kinds of stress fields, you can have conjugate sets of faults that can have different properties. [This can occur with] normal faults and with strike-slip faults.

Lettis: There's a lot that can be done with this database. I'm really fascinated. We have grouped, because of lack of knowledge, all global [stable continental region] earthquakes into cratonic environments or extended margins. They're all lumped together. [However,] not all earthquakes are the same, and not all faults are the same. Because these are neotectonic features [with surface expressions], all the earthquakes [that formed them] are ones that have ruptured to or near the Earth's surface, and deformed the Earth's surface. In the eastern U.S. we're facing this uncertainty -- do the large earthquakes occur below a detachment and never reach the surface, or are they occurring above the detachment and should they rupture the Earth's surface? Are these [Australian] kinds of earthquakes, which are essentially driving Mmax, the same kind of earthquakes that we'd expect in the Central and Eastern U.S.? We need to interrogate these kinds of data and try to learn more about the causes of earthquakes. What kinds of structures are earthquakes reactivating? Are they reactivating detachment features, or high-angle features? All these features that you [Allen] just showed have vertical displacement and they're all reverse faults. [In contrast,] we believe that where we are [in California] they are strike-slip faults. That's a fundamental difference in the way earthquakes are breaking. [Perhaps we can use this database to] learn to differentiate cratonic regimes that are breaking by strike-slip versus by reverse faults. There's so much that can be done with this database.

Allen: For a lot of the western Australian, earthquakes we know that there's a mid-crustal detachment. Most of the earthquakes in that region are occurring above the detachment. They rupture to the surface but they're probably not rupturing much below 10–12 km depth.

Johnston: The Meckering earthquake would provide a nice comparison to a modern fault scarp.

Petersen: When I talked to Mark Leonard [of Geoscience Australia] this summer [of 2008], he said that they had developed distributions of the magnitudes associated with the lengths of those ruptures in both the margin and the craton, using Wells and Coppersmith. You showed some of that, I think, when you were showing those plots of the magnitude frequencies. Does that paper have those distributions?

Adams: That paper's available online. It's the Australian Earthquake Engineering Society, 2006, and that's enough to find it. [http://www.aees.org.au, accessed June 17, 2009]

Ake: You showed the shallow seismic-reflection profile across the Cadell fault (fig. A12–12). It showed on the order of 100–200 m of total throw on the fault. Has that kind of work been done on any of the other scarps in the area?

Allen: No, none of them. They did try to trench the Cadell fault, but it was unsuccessful because it's right next to the Murray River. The fault could be projected up to the surface to try to locate the scarp on the ground. But it was pretty difficult to identify because of the fluvial environment. In Western Australia we don't need to do that because we pretty much know where the scarp is.

Ake: Do you have any estimate of the total throw on any of the faults? That actually gets at what Bill [Lettis] was asking about. We are potentially having to search for active faults that have total throws of only 100 or 200 m. Those are hard to find. Is that a number that would apply to any of the other faults?

Allen: This is characteristic of the eastern domain (fig. A12–11). A couple of examples of that are the [[Lapstone]] monocline, which is just to the west of Sydney, and the Lake George fault, which you drive over along the road from Sydney to Canberra. They all have very similar topographic relief. There was a significant buildup of relief, and then there's really not much else around.

Lindvall: The sheer number of fault scarps is impressive, and clearly because of the ability to preserve those things. Do you have any thoughts on whether Australia is a special continent because of its high rate, or whether we would we expect that other continents would have this many features buried beneath the soil or regolith?

Wheeler: Lee-Ann Bradley, in [the USGS here in Golden], may know in two or three years. We're starting here, using Dan's [Clark] techniques.

Adams: Arch [Johnston] showed that the current seismicity in the stable craton is much higher in Australia than anywhere else [in a craton]. The fact that the number of scarps is so large is perhaps just because Australia is more active. Another thing I would say about the Australian seismicity is that a lot of [it] is shallow, really quite shallow, and we know that because we see the earthquakes being shallow and we also see the faults [scarps]. One earthquake, the Ayers Rock earthquake, was down at 26 km, but we don't know that others can't be deep. So it may be that this rupture population only is part of the earthquakes that happen.

Allen: There's some evidence that suggests we can get deep earthquakes in the Flinders ranges and Mount Lofty regions. Depths are down to about 20 km.

Salomone: Do you see the same issue about climate change that we see in the U.S.? Do you see changes in the climate or changes in the water table [[in the geologic record]]?

Allen: Australia didn't experience the glaciations. A lot of the features may have been preserved because of that. But a lot of the features may have been covered by sediment. Most of central Australia is just covered with dirt, so we have no idea what's beneath it.

Rus Wheeler (USGS): What is a "Tectonic Analog"? (Appendix 13)

This is a straw man, but it may give us some directions to pursue (fig. A13–1). When we talk about Mmax we're talking about large amounts of [seismic] moment [release], and we're also talking about faults. The moment is released when the fault slips. It may be useful to talk about Mmax in terms of the components of moment. We'll set aside shear modulus and treat it as more or less constant, varying just a few percent. That leaves us with slip, length, and width. We can't do much with slip because that refers to single earthquakes and we'll be talking about [cumulative slip over large] groups of earthquakes. So for Mmax we're talking about large widths and large lengths of rupture zones.

The new Geologic Map of North America has a scale of 1:5,000,000 [Reed and others, compilers, 2005, Geologic Map of North America: Boulder, Colorado, Geological Society of America, 2 sheets, 28 p. pamphlet] and its predecessor for the U.S. was at 1:2,500,000. These maps could be made at those scales because there were lots of larger-scale geologic maps that could be compiled. Probably that's also the case in Australia and Europe. In a lot of other SCRs we are likely to have to get along with smaller scale maps. For example, in 1985 Exxon published a tectonic map of the world that Walter [Mooney] has used. Its scale is 1:10,000,000. We're going to have to work on those scales if we want to consider tectonism and tectonic analogs [worldwide]. That means we are going to have to look at large tectonic elements. It will [be] hard to consider individual faults. Good examples of the tectonic elements are what Arch [Johnston] broke out in 1994 -- rifts, passive margins, orogens, and cratons. They are big, so we can recognize them on whatever geologic and tectonic maps are available for a given SCR. Those four kinds of tectonic elements are represented in the CEUS, as I explained in more detail yesterday. If they weren't, then they might be tectonic analogs of something but they would not be helpful in estimating CEUS Mmax.

Now let's bring faults back in. Do any of these [kinds of tectonic elements] have distinctive fault styles that have aspects that might lead us to suspect that the fault styles might favor larger rupture zone lengths or widths than you might find in other tectonic elements that have different fault styles? Let's try to make that more specific (fig. A13–2). These are the tectonic elements that I described yesterday -- extended terranes, orogens and cratons.

Now let's look at plate tectonics (fig. A13–3). We're trying to get to statements like -- in such-and-such a plate tectonic setting, we would expect such-and-such [kinds of faults to form]. The most straightforward are extensional plate motions. They make rifts and passive margins. These are linear features so faults parallel to rifts can be many tens of kilometers long and faults parallel to passive margins can be several times that long. In principle, we could find any length of fault we want in extensional terranes. Faults in rifts and passive margins dip steeply, at least in the upper few kilometers. Also, faults in rifts and passive margins penetrate deeply. What tells us that is partly seismic reflection profiles, but also the fact that the incipient stages of rifting worldwide are characterized by alkaline igneous rocks. Laboratory melting experiments have shown that when you melt bulk compositions characteristic of the crust, under small amounts of horizontal extension that relieve the confining pressure, the first melts that come off are unusually rich in potassium. The implication is that when you first start to extend the crust, you make small amounts of alkaline igneous rocks. As the extension proceeds, the early alkaline rocks are swamped by larger volumes of igneous rocks of more common compositions. The presence of the alkaline rocks tells us that the extension goes deep, to melting depths, below the brittle-ductile transition, below the base of the seismogenic zone. So the extensional faults can have large lengths and large widths, so we should not be surprised to see large Mmax in extensional terranes. We've been hearing many times, yesterday and today, that that's exactly what we see.

Now let's look at contractional plate motion. That makes orogens when the contraction is finished. The brittle parts of orogens, not the hot ductile cores, but the outer parts like the western Appalachian Mountains and the northern Ouachita Mountains, do their contraction mostly on thrust faults. Thrust faults have low dips, and they have very great lengths and very great widths. Thrust faults can carry the rocks above them tens of kilometers up shallow dips, over the rocks beneath the faults. Thus, contractional plate motions can produce faults with large lengths and large widths, but they are shallow. In the Appalachian Mountains seismic-reflection profiles can detect these thrust faults extending down to depths of roughly 5 km. In 1992, Wheeler and Johnston [(Wheeler and Johnston, 1992)] published histograms of reasonably well-constrained depths from local networks and the hypocentral depths of 15–20 larger earthquakes that had been the subjects of special studies. With local variations, depths in the CEUS and adjacent Canada have a median of something like 10 km. Most of the seismicity is within plus or minus 5 [km of the median]. It's shallower in New England, and elsewhere it may be deeper. So the thrust faults, although they are large, probably are too shallow to support much of a rupture. They can support small to moderate seismicity, but maybe not something we'd consider for an Mmax earthquake. Underneath [the thrust faults] you can have bigger faults. Thrust faults form in an orogen when an ocean closes, and the thrust faults override a passive margin that formed in the opening of that same ocean. The result is that we might expect smaller Mmax in orogens than in extensional terranes.

Now let's look at cratons. They are distinguished from other terranes because they are old [as Walter Mooney described yesterday]. That means they've been through several different episodes of extension and contraction, normal faulting and thrusting, with their fault systems overprinted onto each other so that the fault system we might expect in a craton would combine all of these together. We'd expect to see lots of different lengths and lots of different orientations. Another thing enters into the fault systems of cratons. Remember that yesterday I brought up the topic of fault maturation and what happens to fault systems as slip accumulates. That's counteracted by faults healing, and I have no idea how fast healing works. I've seen a couple of seismological papers that argue for healing starting within a couple of years after a fault rupture in California. Probably fault healing is not going to be uniform because when rocks are hot they form their own mineral assemblages, but as they cool, they go through retrograde metamorphism. Retrograde metamorphism uses up water so a lot of cratonic crust has been dewatered. Some of it will [still] be wet, but the water will be patchy [in its distribution]. The reason it matters is that fault healing involves moving dissolved chemical species through cracks. You can dissolve small grains, or grains that are not optimally oriented in the ambient stress field. The water moves that [dissolved material] around to places where it grows new grains that are optimally oriented, and larger than the original faulting-comminuted grains. The result is that we shouldn't expect fault healing to be ubiquitous but we can expect that it will take place. Cratons have been around for a long time. There have been perhaps billions of years in many cases for healing to decrease the [effective] length and the width of [old] rupture zones that can be readily reactivated in the modern stress field. So we shouldn't be surprised to see smaller Mmax in the cratons.

You'll recall that I've argued that extended terranes might be expected to have larger Mmax than orogens or cratons. However, I haven't been able to develop arguments that would suggest whether orogens should have larger Mmax than cratons or vice versa.

Here are some of the large tectonic elements like the ones that we find in the CEUS and adjacent Canada (fig. A13–4). Yesterday, I explained that the green is Mesozoic [and Paleogene] extended terranes, blue is Paleozoic extended terranes, and gray is cratons inside the heavy black lines that mark the SCR boundaries. These are the global tectonic analogs to parts of the CEUS. They are the places in which we'd expect to find fault systems like the ones we find in [different parts of] the CEUS.

Let me pull some of these ideas together (fig. A13–5). Extended terranes have distinctive fault styles, so do orogens, and so do cratons. Fault styles, and perhaps some of the associated rock types, let us recognize a rift or passive margin anywhere in the world, whatever SCR it is in. We can say the same thing for orogens, and we can say the same thing for cratons. The last slide [figure A13–4] demonstrated that it can be done. I didn't show the orogens but I know where to go to get them. We can recognize examples of each of these tectonic elements wherever they occur, and they are different enough that we can readily distinguish examples of each element from examples of the other elements in any single SCR, like this one. What that suggests is that the differences that go into defining tectonic analogs are much greater within a single SCR than between SCRs. That's the justification for defining a tectonic analog that looks like something we have here in the CEUS, where "looks like" means in terms of some criterion or set of criteria related to faults and to the length and width of expected rupture zones.

That's a suggestion as to what we might mean by "tectonic analog" and how we might go about identifying analogs. It should come as no surprise that we end up with Arch's [Johnston] results and with the USGS Mmax zones, but there are similarities to what's done for the Bayesian analyses.

Discussion

Lettis: This discussion was based on the assumption that SCR earthquakes are occurring on preexisting faults. It may be that SCR earthquakes are producing new faults and that they have nothing to do with what you described in terms of looking at characteristic fault styles. Going back to the talk on Australia, what type of basement faults are they occurring on, and what type of fault that you talked about are being reactivated? I suggest that a lot of those earthquakes are breaking new faults and not reactivating old ones.

Kimball: I think we know little about the source size that results in a given magnitude, particularly in an SCR. It's something that we should try to get a better understanding of. Wells and Coppersmith [(1994)] is our [[standard reference]], but that doesn't mean it's right. We need to get a better appreciation for the size fault that's needed to generate a moment magnitude six and a half or seven. If we underappreciate that and it's a much smaller rupture than we think, particularly given the variability of the stress, then it would change our way of thinking about [the problem].

Johnston: For the Bhuj earthquake, the fault size as defined by aftershocks is about two and a half to three times smaller than you would expect for a magnitude 7.6–7.7. It raises the question in my mind, was the Bhuj earthquake a new fault? If we can nail down the difference in source scaling between SCRs and active tectonic regions, it might be large enough that it could actually be used to define SCRs. We're not there yet, but at least the data on a large earthquake are good now.

Li: Are analogs based on tectonic maps around the world?

Wheeler: I started with the CEUS and Canada.

Li: Should we have a three-dimensional description attached to the sources to give a better idea of where a hypocenter is [which source contains the hypocenter]?

Wheeler: The vertical dimension of the seismicity is much, much smaller than the horizontal dimension of these tectonic elements. If you look at the scales of the maps I showed, the depth [of the hypocenter] is irrelevant. It may well have something to tell us about the mechanism and about fault size; that was part of Arch's [Johnston] point.

Frankel: Rus [Wheeler], you seem to dismiss these thrust faults as sources of earthquakes, but we know worldwide that thrusts make earthquakes, for example at El-Asnam [table 4]. Why couldn't we have a big earthquake on the Appalachian detachment?

Wheeler: How comfortable are [all of] you as seismologists with reactivating a subhorizontal thrust fault at, say, 5 km depth? If it's steeper [in dip then it might be reactivated].

Adams: It's [El-Asnam is] not analogous to the Appalachians.

Wheeler: There's another example -- Coalinga [table 4]. El-Asnam was like Coalinga, except that Coalinga is a feature within a strike-slip boundary, and I think it's compressional at El-Asnam. It's an active anticline with a thrust fault underneath.

Ebel: To what degree is the idea of fault reactivation supported by geologic evidence? The one thing I've never seen is where the geologist does a systematic study to find out which faults have been reactivated in the past. If they have been, there must have been earthquakes. Has anyone gone out and concluded that these kinds of faults consistently show reactivation, whereas these other faults do not show reactivation?

Wheeler: I don't know, but for quite a few years Rob Holdsworth at the University of Durham has headed the Center for Fault Reactivation. They do this sort of thing. Also, last fall I was on a field trip to central Colorado to see pseudotachylites. [The trip leaders] found pseudotachylites that formed in a ductile fault. That's a brittle reactivation of a ductile fault. But I don't know of anybody who has looked at it from exactly that viewpoint [that you asked about].

Lettis: One example that we know almost for certain is, in many areas, Triassic basin normal faults that have been reactivated with a reverse or oblique sense of movement.

Wheeler: And there are thrust faults that have sometimes been shown to have reactivated Cambrian extensional faults.

Arabasz: In the whole eastern Basin and Range [Province] the thrust-belt structures were reactivated under extension. Typically the steeper ramps become the normal faults.

Lettis: In the Laramide fold-and-thrust belt in Wyoming and elsewhere, it's argued that those thrust faults have been reactivated.

Wheeler: This sounds like what John [Ebel] was suggesting, to compile faults that have and have not been reactivated. The problem that would have to be overcome, and it's certainly not intractable, is that usually a fault that [is known to have] been reactivated, almost of necessity, has been studied in unusual detail to detect at least two different episodes of movement and date them. It's a lot easier just to find a fault. You'd have to make sure that you were collecting only faults that had been studied well enough to detect reactivation above a certain size, if it were there. In principle, it's possible. I don't know of anyone who's tried, partly because the folks who generate and publish and use detailed evidence on faults are either university people who are interested in petrology or structural geology, or explorationists for minerals or petroleum. None of those folks care too much about earthquakes, and we don't know much about their specialties. Again, it's not a problem that can't be overcome.

Coppersmith: The issue is how to go farther. We could list all the cases with clear reactivation. The big normal faults in the Norwegian North Sea have clear Cenozoic reactivation. You can see digital tracings [of reflections on seismic-reflection profiles] of big normal faults with nearly a kilometer of normal slip, and in the upper part [you can see] reverse reactivation and local folds. In the foothills of the Sierra Nevada are places that are locally reactivated, but that only answers the first part of the question. The second part of the question -- is there something distinctive about all that, which allows you to predict in advance that this fault is likely to be reactivated, [but] that one is not. What are the criteria for it? What are the types of things that you would be looking for?

Wheeler: That's a geologic variant of the question that has plagued us in the East from a seismological perspective. Suppose you can link an earthquake to, say, a Cambrian extensional fault. Why that fault and not any [others of the same age and faulting style]?

Hanson: I think Kevin's [Coppersmith] comment about the Foothills Fault System in California is a perfect example of a long, continuous shear zone. Portions of that have evidence for reactivation that generated historical earthquakes. There's [been] a lot of trenching and a lot of study along that fault system. I still think there's a lot of uncertainty about what the maximum magnitude along that fault system could be because you have extensive faults that could be reactivated. Why certain pieces have been [reactivated] I don't think you could pin down.

Ake: To get to the nature of this problem in general, that's a very good analog because the rate appears to be quite low. It is right at the threshold of whether the evidence is being removed in the landscape about as fast as it is accumulating in the landscape. Of those [fault] sections [of the Foothills Fault System that] are similarly aligned in the same stress field, how many of them are reactivating and how many are not?

Lettis: The Foothills Fault System has all the issues that you listed. It's a former crustal suture. It's been [altered] by hydrothermal activity. It's a very long fault, which has been intruded by plutons and locally deformed. And yet certain segments, which are hard to predict in advance, break. We don't think they break in magnitude seven to seven and one-half. They're breaking in roughly magnitude six and one-half, but we allow for sevens now in some of our source zones. It has a low slip rate, maybe a millimeter per year. It's a good example and it's hard to say that even a healed suture in a contractional environment can't be reactivated if certain conditions are met, and I don't know what those conditions are.

Wheeler: That's encouragingly consistent with the eastern paradigm of reactivation.

Lettis: I hate to cast guilt on every fault in the crust.

Wheeler: We don't have to. The problem is [that] we don't know which ones are guilty.

Chapman: One of the things that I've noticed is that a lot of the Mesozoic basins in the Appalachian Piedmont [Province], or at least their northwestern boundaries, appear to be correlated with the Paleozoic thrust faults. I believe in some cases we are seeing that earthquakes, at least in the Appalachian Piedmont [Province], may be associated with these Mesozoic features. They've gone through several generations of reactivation. It's very difficult to say exactly what's happening there, except that the most recent major reactivation that we've seen might have something to do with the seismicity.

Ake: Rus [Wheeler], one of the things that you had in your slides is that healing happens, and the age of reactivation could be important. What do we mean by that? [Do we mean --] is the last major activity [on some structure that's pertinent to seismicity] Paleozoic [? Do we mean --] is there anything more recent in terms of documented activity on these structures, or does it have to be more recent?

Wheeler: I don't know enough to go beyond the observations. In other words, I don't understand the processes well, especially their kinematics. The observations are Arch's [Johnston], that Mesozoic and [early] Cenozoic extended terranes have higher Mmax [observed] than other places. That doesn't say whether it's the faults that are weaker, or the crust and maybe the upper mantle that's weaker. Paleozoic extensional terranes, at least the one example that we know best, do have earthquakes, but not as many, but Precambrian extended terranes are as dead as the non-extended parts around them. This suggests a healing half-life, or a healing rate, such that the healing [of] faults, perhaps to the strength of the unfaulted rock around them, takes place over hundreds of millions of years. I'm not sure that's a time scale that's particularly useful to us. I'm not sure what light it could shed on Mmax.

Coppersmith: It could eliminate structures that are [too old for reactivation because they've been healed].

John Ebel (Boston College): Magnitude of the 1663 Charlevoix earthquake (Appendix 14)

I'd like to talk with you about the 1663 earthquake, which we think was in the Charlevoix seismic zone [table 4] (fig. A14–1). Some of my thinking is biased by this paper that we published in 2000 (fig. A14–2)[[Ebel and others, 2000]], where we speculated that some of these areas with clusters of earthquakes really are very late aftershocks, hundreds or thousands of years after the main shocks. They'd have to have been quite large main shocks in order to have aftershock activity that would last that long.

The Charlevoix seismic zone is up here [the large cluster of numerous earthquakes northwest of Maine, along the St. Lawrence River]. Here's the report of the Charlevoix earthquake at Roxbury, Massachusetts [[Danforth, 1880]], which is now part of the city of Boston and at a distance of almost 600 km (fig. A14–3). The [reported damage] is about intensity VI shaking at a distance of 600 km from the earthquake.

What kind of argument can we make for the size of this earthquake (fig. A14–4)? One [argument] goes back to that paleoseismicity idea, that this zone is aftershocks of 1663. The zone is about 70–75 km long, so it's not an insignificant zone [[Ebel, 1996]]. [It's] comparable in size to the cross-zone in the New Madrid seismic zone. This length and the Wells and Coppersmith relation [(1994)] give magnitudes from 7.1–7.2 to approaching 7.5, depending on whether you use the subsurface or surface rupture length, or multiply the length times the width to get the rupture area.

Ake: How deep do the earthquakes go?

Ebel: Those earthquakes go down to below 20 km, and there are a few hypocenters that go down to 30 km. It's a relatively high-angle zone. Doesn't it dip about 60 degrees?

There are a couple of ways to estimate the magnitude from the ground-motion information at Boston (fig. A14–5). One is to assume intensity V, VI, and VII [and] take an intensity attenuation relation. I used one by George Klimkiewicz (1980), which included intensity attenuation information from the 1925 Charlevoix earthquake. [I] get [these Lg magnitudes].

Another way to do it is to consider the chimney damage itself. Here's some work by Bob Whitman, who's a retired geotechnical engineer from MIT [Massachusetts Institute of Technology] (fig. A14–6). The threshold for chimney damage if you use PGA [(peak ground acceleration; table 2)] is about 0.01 g. For a typical two story chimney in Massachusetts, the natural period is somewhere close to 0.3 seconds, so you might look at [the acceleration threshold for] 0.3 seconds. 0.01–0.03 g is where you would expect chimney damage to start, and you're probably right at or just above that threshold. I took two attenuation relations [by] Atkinson and Boore (1995) and Toro and others (1997). For a distance of 600 km and a spectral acceleration of 0.3 seconds, Atkinson and Boore could explain the ground motions by a magnitude seven and a half earthquake. Toro's ground motions, both for PGA and for 0.3 second spectral acceleration, are predicted to be less than the threshold for chimney damage. This analysis suggests magnitude 7.5 or perhaps even a little bit larger. For both of these I believe I used the Lg attenuation relationship. I think that the damage at Boston is consistent with an Lg magnitude around 7.5, which is moment magnitude 7.2–7.5. Sue [Hough], how far was the chimney damage from [the New Madrid earthquakes of] 1811–1812?

<u>Hough:</u> There was chimney damage reported in Cincinnati, which is about 600 km from New Madrid. Site effects would affect the reported intensity.

<u>Ebel:</u> Right, but we're talking about 1663 and I'm not sure you can argue for any significant soil amplification at Roxbury. Roxbury goes up into Roxbury Heights. I don't know exactly where the chimneys were, because you can go down to the water, so you probably have some range of soil conditions in Roxbury. It's not clear in 1663 that the structures would have been on much soil. For instance, in South Boston, which is not far from Roxbury, the soil is only 5–10 m thick.

<u>Kimball:</u> Two or three damaged chimneys is not much damage for a large population. You really don't have enough evidence to support [your] conclusions.

<u>Ebel:</u> There are several issues. In 1663 you did not have a large population. You also have to worry about the type of construction. I used Whitman's fragility curves, which are for an older kind of mortar, somewhat less strong mortar than would be used today. I'm trying to take into account some of these things. This was only forty years after the area was settled so probably you're only talking about a few thousand people in the whole area.

Frankel: This doesn't amount to much evidence, in view of the uncertainties in magnitude, ground motions at the site, and vulnerability of the structures. All of the evidence could be off by a factor of two or three.

Wheeler: Were Montreal and Quebec settled then?

Ebel: I don't really know what the settlement patterns were. There were people up along the St. Lawrence River in 1663, because a lot of what we know about this earthquake came from Jesuit reports that were sent back to Europe and published. There may be more information out there. I don't know whether it would have been felt in New York City, or up into the maritimes of Canada.

Johnston: If you accept the argument that the tabular seismicity is an aftershock zone, which I think it probably is, then a 70 km length and a 20 km width corresponds to a fault area of 1,400 square kilometers. The Bhuj fault area was about 1,200–1,300 square kilometers. That gives you a modern-day comparison.

Ebel: So [the 1663 earthquake could have had] about the same size rupture area as Bhuj.

Kim: Most of the focal mechanisms in the area trend north-south, which implies north-south striking faults, but the Charlevoix seismicity trends northeast.

Ebel: The work I've done and the work Allison Bent has done suggests that the 1925 earthquake here [in the Charlevoix seismic zone] had a P–axis like this [trending northwest] and was a thrust earthquake on a fault that strikes in this direction [northeast].

Adams: This may have been a strike-slip earthquake, which would make the stress field complex. The main shock could have distorted the stress field quite a lot. [[In that context subsequent earthquakes (even ones 350 years later) might be aftershocks reacting to the distorted stress field. Thus the thrust mechanism documented by Allison Bent for the 1925 Charlevoix M 6.2 event need not represent the mechanism of 1663. Furthermore, it is easier to conceal the geomorphic evidence of a large strike-slip earthquake (especially if it occurs underwater) than of a large thrust event, in which case one might see meter-scale shoreline elevation changes for a single event and many events would accumulate into hundreds of meters or a kilometer of deformation (which is not seen).]]

Frankel: If the small earthquakes are aftershocks they should be decreasing with time. Are they?

Ebel: The rates are pretty constant over a few tens of years. I don't know if we can say anything about the rates over hundreds of years.

Adams: We can't tell whether the aftershock rate is dying down. [[If the aftershock sequences lasts about 1,000 years]] we have only a 10 percent time slice and that's not enough to draw any conclusion.

Frankel: Steve Wesnousky recently looked at a survey of surface ruptures in active tectonic regions to see what the maximum distance [is that] you can get between faults that jump between surface ruptures. He found that 5 km is about the maximum jump. It's his view that this could have applications to Mmax in the East. In eastern Tennessee you can see segments. The question is: can they link up in big earthquakes? How big are the faults out there? Can we use seismicity or seismic reflection [profiles to determine whether they are segmented]? We should think about earthquakes in terms of sources.

Ebel: Let me get on to my thinking (figs. A14–7, A14–8). What got me started is the paper that was published by Stu Nishenko and Gil Bollinger (1990), on recurrence relations in the CEUS and eastern North America. They were trying to take very large regions and extrapolate the seismicity to larger magnitudes and estimate how often [large] earthquakes might recur. They worked with M 5 and greater and M 6 and greater. They used two different relations, as shown by the red and green lines in figure A14–8. I extrapolated their relations up to M 7 [and greater] and asked, how many earthquakes did they predict, and how many have been observed? In [the case of eastern North America] from 1760 to 2003, and for the CEUS going back to 1620, at M 7 and greater, I noted that there have been more earthquakes observed than predicted by these recurrence relations. We are talking about very, very small numbers of earthquakes. We can't go back in time to add to the statistics, so I took the map that Art Frankel published in 1995 as an early version of the a-value map for the CEUS [[Frankel, 1995]]. He had contoured the seismicity. This (fig. A14–9) is a map of the number of M 0 earthquakes per 60 years. Where you have darker colors, that's just reflecting higher rates of seismicity and higher a-values. I counted all areas that are within the eight[-earthquake] contours, and then all areas that are within the 16[–earthquake] contours. Those are all indicated by these [red] arrows. Let's assume that all of those places had M 7 or greater earthquakes, and that's why they are so active today.

Then I took the Nishenko-Bollinger relations, went back 1,100 years, and assumed that the areas that had 16 or more M0 earthquakes in the last 60 years are areas where there were M 7s. I have eight of those, so I said that there should have been eight earthquakes of M 7 or greater in the last 1,100 years or so [upper table in figure A14–10]. Nishenko and Bollinger predict less than half that number. Next I looked at the areas with eight or more M0 earthquakes, and assumed that they had M 7 or greater earthquakes in the last 2,100 years. The recurrence relations predict the number of arrows that I had on the map, and I'm a factor of three greater. For this calculation I assumed that everything between M 7.0 and M 7.5 has equal rate.

I did another calculation using the Gutenberg-Richter distribution [lower table in the figure]. This was to get the factor of three down, which happens because there are fewer M 7.5s than M 7.0s, which means fewer earthquakes in total. But I only got them down to 6, 4, 13, and 7, which are still a factor of two higher than [predicted by] the Nishenko-Bollinger relations. If these clusters really are reflecting large earthquakes that have happened in the past, I would argue that perhaps there have been more M 7s than the extrapolation of the modern seismicity tells us we should have had in the CEUS or in eastern North America. When we are talking about Mmax, rates of those large earthquakes are also something we need to know for seismic hazard.

I was using their [Nishenko and Bollinger's] b-value. I did not fit one. I was using their recurrence relations. [I was just] distributing numbers of earthquakes between magnitudes seven and seven and one-half. I needed a b-value to do that. It doesn't have anything to do with the a-value.

Discussion

Lettis: Are these aftershocks? What's the seismologists' opinion? Are these aftershocks, or are these areas of greater seismicity that might have a use for identifying where the next big earthquake might occur?

Ebel: As I study more and more specific areas in greater detail, I'm leaning more and more toward thinking that a lot of the earthquakes are aftershocks. I'm amazed, when I look at California seismicity, how much of it you can explain as aftershocks.

Lettis: Do these rates fit Omori's law?

Ebel: The problem is that Omori's decay function initially is very steep, and then it flattens out. Where you have constant seismicity, you can argue that it is approximately flat at the tail of the 1/t of Omori's law. Once you get out on the tail, you can't really tell unless you can sample for hundreds of years and see the rate decreasing.

Frankel: We see that the average stress drop of eastern U.S. earthquakes is about 150–200 bars. I'd be a little surprised if aftershocks have that kind of stress drop. Also, is the Grand Banks earthquake still having aftershocks?

Adams: Yes.

Ebel: Art [Frankel], I just completed a study and sent a report in to the [USGS] external grants Web page. I looked at intraplate and cratonic aftershocks, and looked at the statistics and a version of

Omori's law, the same one that Reasenberg and Lucy Jones used [(Reasenberg and Jones, 1989)]. I see the same distribution of parameters for the intraplate aftershocks as for California aftershocks. At least statistically, whatever we see in California we also see for intraplate earthquakes. I have some from Australia and Europe and North America. They statistically seem to follow the same distribution. Now of course you have active aftershock sequences and inactive aftershock sequences, and you have somehow to take that into account if you want to do this more rigorously.

Tuesday afternoon, September 9

Annie Kammerer, Jon Ake (NRC): What are our Conclusions About Methods of Estimating Mmax? (Appendix 15)

Ake: After the presentations and discussions we've had thus far, I'd like to go back to a couple of key things. One of them is to try to capture what we've heard about our research needs for both the short term and the long term. We have two discrete user groups for the Mmax work. One is for more site-specific studies and the other would be, for example, the Canadian survey and the USGS who make hazard maps.

Kammerer: Several people have pointed out things that we need to look into more. I've pulled them together here (fig. A15–1). [The version of the figure shown during the workshop did not include several items that were suggested during the following discussion. This version of the figure contains the suggested items.] One thing I heard a lot is that we can't wait another thousand years for instrumental records [of large earthquakes]. One of the best approaches we have for getting more data is paleoseismology. Liquefaction was discussed a lot and that's something we are really interested in. [Kevin Coppersmith] also mentioned lake deposits and fragile features. Currently we are starting some proof of concept work in caves. Historical data is another area that came up quite a bit. At the first

[SSC] workshop we heard of a compilation of a lot of newspaper accounts just in the last three years [[by Jeff Munsey of the TVA]]. We also need to look at census data to address the questions: how complete are some of these [earthquake] catalogs? Were people living in the area? When did the first newspapers [start publishing] in different areas? [Answering these questions can lead to] a better understanding of the completeness of the information that we do have. We heard today discussion of site effects, as a way to evaluate shaking in different areas. Paleoseismology and historical data seem like things that we should be mining for new information.

Coppersmith: Arch [Johnston] looked at all of the historical SCR earthquakes. We need to go back to them globally and apply paleoseismology to them. Is there any evidence of past occurrences of similar earthquakes? We could go around the world and try to get a better handle on rates of recurrence. Historical accounts would be important, but so would these geologic effects. Results could then be associated with particular types of faults.

Wheeler: A lot of the earthquakes are pre-instrumental. Ideally we might want to do the work on intensity data and on site effects first.

Salomone: A lot of this work on intensity and site effects requires a lot of experience. The people who have that experience need to be identified. What they've found needs to be preserved. Also, the expertise in identification of specific paleoliquefaction [[features, that]] Steve Obermeier and Tish Tuttle have [[, for example]], needs to be [shared with] younger [, less experienced] people [who will carry on the work]. Documenting those two areas could be started quickly.

Kammerer: Many of the areas of research include knowledge transfer and training the next generation. We're trying to support doctoral students and things like that as much as possible. I've added [in the figure] field studies to [[enhance]] interpretation (fig. A15–1).

About global analogs, specific items came up: to define what is meant by analog, to remove multiple earthquakes, to perform bias corrections, and to incorporate magnitude uncertainty.

There was a lot of discussion about [fault] reactivation and interest in it. There was a lot of interest in the Australian neotectonics database, which is a fantastic resource. There also was discussion of studying reactivated faults to try to determine whether certain types tend to be reactivated.

[We] discussed the applicability of the Wells and Coppersmith results to other regions [like the CEUS]. We may be looking at smaller rupture areas, so rupture widths are important to constrain.

Johnston: For Wells and Coppersmith, an SCR-specific source scaling has been developed by Saikia and Somerville. I'm not sure if that's been published other than in a NEHRP report to USGS. Does anybody know?

Kammerer: The national LiDAR project is something of interest. LiDAR has been very useful in the Pacific Northwest. I'd be interested in whether people think that would be of any use in the Central and Eastern U.S. I'd also like to hear if there is anything missing from this list (fig. A15–1).

Hough: Analysis of strain budgets from GPS data is a promising development.

Kammerer: We should put that on the list. At a workshop on the CEUS there was a presentation on GPS data. It was inconclusive, but it is worth looking at to see what opportunities there might be.

Coppersmith: Permanent GPS stations should be installed to map lower-crustal strain anomalies.

Adams: Geodetic strain can be used as a constraint on the rate of large earthquakes, although there is little information with which to do this in the East.

Kammerer: With regard to the NRC program, it is a lot easier to support things like instrumentation rather than the ongoing maintenance cost. Maintenance is seen as the USGS's job, or as industry's job.

Ake: An earlier comment recommended looking at offsets in faults. What is the likelihood of segments linking up to make a very large event? What sort of research do we need to do to help answer those questions? We have some idea of the inventory [of faults], but are they active or not?

Frankel: To amplify on that, the standard method for this kind of work is dynamic rupture modeling. That might be an interesting way to see how a higher level of stress changes some of these kinds of things. I don't think anybody really works on the stable continental region problem, with its different stress drops and different complexities of faulting.

Adams: In the relatively short term, LiDAR is not a great way [to look for scarps][[because of the large area and the high cost of flying it]]. [A] better [approach] might be visual examination of digital elevation models (DEMs; table 2), for example in Google Earth. It might be worth getting someone to do it systematically.

Kammerer: We're doing that, using Google Earth to identify features.

Hanson: It's a good application for research. You can put the viewer right where you want, so you know where you're flying. I agree with John.

Wheeler: Last September and early October we had Dan Clark [of Geoscience Australia] up here for five weeks. We took him out in the field about half the time to show him paleoseismological features. The other half of the time he spent here in the office, demonstrating to us what he does [with DEMs to identify the neotectonic features that he's found in Australia] -- what criteria he uses, what software he uses, and how he manipulates it. He also examined some targets that we already knew about in the midcontinent -- things that have been shown to be tectonic [in origin], and things that have been shown not to be tectonic, like the Anton escarpment in Colorado and the Ord escarpment in Nebraska. He took a look at New Madrid because he was going to Memphis for a few days, and convinced himself and all of us that his approach does not work where there are large, active streams reshaping the landscape rapidly. We now have part of one geologist dedicated to establishing a system like Dan uses and applying it to the semi-arid western half of the Great Plains. We know that none of us are going to be around long enough to finish it, but we're going to start in Colorado and add adjacent states once we know what we're doing. Eventually parts or all of ten states will be covered in the western Great Plains. If that works, then we intend to go a bit farther to the east. I don't know how much farther east it's worth using DEMs. It could well be that by then there will be enough LiDAR that [it can be used to greater advantage]. For many years I've been looking for a way to prospect enormous regions quickly and cheaply to find targets that are worth investigating on the ground. So Dan's work in Australia is an example of the direction we're hoping to take in the next few years in the western Great Plains.

Hanson: What is the resolution of the DEMs that you are using?

Allen: He used SRTM [(Shuttle Radar Topography Mission; table 2)] data. [[Across the vast majority of the Australian continent, the 3–second SRTM DEM is the only regional DEM available. In contrast, the detailed scarp map in the southwest of western Australia (fig. A12–15) was compiled using an orthophoto-derived DEM that has 10–m horizontal resolution. Features become discernible in the orthophoto dataset above about 1–1.5 m in height in rolling countryside, so we can see scarps of most earthquakes above M 6.5. I [Dan Clark] could see fewer than half of the features in the same area using the 3–second SRTM data. In the SRTM dataset features emerge from the noise if they are higher than 3 m. If the southwest is anything to go by, we might expect to be missing half the features or more compared to what we'd see with a 10–m orthophoto dataset. A postscript is that, for this work, the 1–second SRTM data are not a significant improvement over the 3–second data because of high noise content. The resolution is closer to 70 m than the state 30 m.]]

Wheeler: That means that they couldn't see any of the active Holocene scarps in west Texas or a lot of the stuff in the Basin and Range [Province]. It tells us that they're probably missing a lot of features [in Australia]. For the ones they do see, they probably are underestimating the lengths, which could be corrected once they get there [on the ground]. But [in Australia] there is one paleoseismologist plus a couple who are coming along, as coauthors, for an entire continent. My guess is that Dan's [Clark] estimates of frequencies, sizes, and numbers probably are underestimates. Of course, there's only one way to find out. [In the western Great Plains of the central U.S., we'll be using DEMs that started out as 1:24,000–scale topographic maps, so probably they're going to have better resolution [than the SRTM data]. But still, we'll miss some things.

Nelson: Using Dan's methods you'd miss maybe 80 percent of the scarps that are seen with LiDAR, even if the leaves were off. LiDAR is expensive but it's getting cheaper and there are many more LiDAR data sets than there were a year ago. Different groups are acquiring them for different reasons. If you can process them, you're very likely to find a lot more [scarps]. Steve [Personius] and the rest of us have been trenching one-half meter high scarps in areas that were clear cut 20 years ago and are a total jungle to walk through now. With GPS and LiDAR you can go right up to a half-meter high scarp and trench it.

Wheeler: If the DEMs highlight an area that looks like it has a lot of marginal features, or if there is a particular feature whose length you really want to know, and if LiDAR data are there, they would help to answer those questions even faster and cheaper than going there on the ground, which of course would be the next step [to collect field data].

Nelson: [LiDAR data would also be valuable] if you have a forested region and you think it is quite possible that there might be surface deformation or fairly narrow folds, or [if the locale of interest is] in an urban area.

Kammerer: Is there anything more to add [to this figure]?

Stirewalt: Is the concept of triggered seismicity pertinent here?

Hough: It’s hard to use historical data [for this], but New Madrid, I’m convinced, has triggered earthquakes. If you don’t recognize that you could inflate estimates of [the size of the triggering earthquake].

Lettis: In terms of capturing research needs, there are a number of approaches that we’ve identified as needing more work. Perhaps [triggered historical seismicity is] a focus area of research.

Kammerer: Yesterday we discussed various methods [of estimating Mmax]. For some methods, someone said we should throw it out and no one objected. For other methods, someone pointed out reasons to keep the method. This is where my notes tell me we ended up (figs. A15–2, A15–3).

Ake: Some of these methods can be used directly, whereas others can be used as inputs to a Bayesian analysis. Some, like Mobs and global analogs, can be used in both ways.

Lindvall: Would a bias correction get any closer to Mmax than Mobs plus a constant would?

Kammerer: Other people might disagree, but I feel that while some sort of a bias correction is appropriate, a correction that is constant needs a strong basis.

Li: Is paleoseismic part of local geology or part of Mobs?

Kammerer: Local geology. As a liquefaction engineer, I think there are a lot of problems with saying that something didn’t occur because you don’t find liquefaction. It could be a drought year. It could be that you weren’t looking in the right place. Paleoseismic fits into local geology (fig. A15–2), and needed research (fig. A15–1). Here’s where we put it (fig. A15–4) [in three places].

Lettis: Local geology, which gives information like rupture dimensions, is both a method and an input to other methods. It could be a direct input to Mmax. Where in the table (figs. A15–2, A15–3) would we put using fault length and fault width for a direct estimate of Mmax?

Coppersmith: We could put it into “Local geology/rupture dimensions.”

Kimball: We're thinking of local geology in terms of how it was used in the past. We need to move beyond that, and including paleoseismic information in local geology does that.

Arabasz: M-f extrapolation seems to be part of statistical analysis. I think that the technical community might support asking the question of whether an estimate of Mmax is consistent with an M-f extrapolation of the historical seismicity.

Kammerer: So M-f extrapolation is not an estimation method itself. It could be applied to an Mmax estimated with other methods, as a sanity check. We could put it here [in statistics] (fig. A15–3).

Murphy: Going back to the differentiation between North American analogs and global analogs, they seem to form a continuum with North American analogs being part of global analogs.

Kammerer: In the broader technical community there are people who would prefer to rely more on North American analogs than on global analogs [[and so would possibly apply different weights to the data]].

Hanson: Both methods have the same goal. Both can be used on their own and as inputs to a Bayesian approach.

Wheeler: Under global analogs, "does it agree with North America" is a question that one should always ask. One reason I've always been so comfortable with those two subdivisions of the global catalog [into young extended areas and all other areas] is that once I knew they were [different] globally, I could look back at the CEUS and Canadian seismicity and see the [same] pattern. I could not have argued from this one continent's pattern that there is this division. There just aren't enough earthquakes in this one SCR. If you find a global pattern that doesn't show up in North America, then [the global pattern will be of little use in North America].

Adams: Does that mean that the Australian pattern of many large surface ruptures doesn't apply to the CEUS?

Kammerer: So are we good with leaving it this way [with North American analogs and global analogs listed separately but with the same applications]? [No one objected.]

We didn't have a lot of discussion on statistics or pattern recognition. That's why [the] statistics [cell in the right side of figure A15–3] was blank [before insertion of the note about M-f extrapolation]. I think that pattern recognition means different things to different people as well.

Perkins: One of the things you might consider under statistics is the statistics of the different regions or analogs -- are the extended margins quite as active as the cratons? What kind of statistical tests can you do to show whether they [cratons or individual SCRs] are deficient in their maximum magnitude behaviors?

Coppersmith: At the time we looked at statistical approaches in 1985, there were two. One was the Kijko and Graham [(1998)] approach, which was to take a catalog and look at the frequency statistics, assuming an exponential distribution, to see whether you had sufficient numbers to define an upper bound. The second was [from] the Soviet Union, that had statistical approaches [in which] they did multivariate regression of everything you can imagine into one model, and came out with a maximum magnitude. They put in topography, frequency of stream intersections, crustal structure, [and] a whole manner of things, and did a multiple regression based on observed earthquakes in their country and arrived at a maximum earthquake assessment. That's where it [the statistical-methods category in fig. A15–3] came from.

Arabasz: Someone could check that. It gives fairly extreme values that are still used in other disciplines.

Wheeler: Kijko and Graham [(1998)] is one of the papers I included under that [in the foundation document (Wheeler, 2009)].

Ake: The way Dave Perkins was talking about it [statistical methods] was really as a basis for statistical tests on the validity. Is that right?

Perkins: It would be a way of dividing [the global set of SCRs and their earthquakes] into different categories. It could define the analogs.

Coppersmith: Allin [Cornell] did this and couldn't find a statistical difference.

Frankel: You might have to do statistical tests to see if there's any difference in Mmax any place in the world, between active tectonic regions and [others].

Ake: You may end up deciding that outside the subduction zones there isn't any difference.

Salomone: We might be able to look at this table [figs. A15–2 and A15–3] in this way: we've got the Bayesian method as the solid method, we've identified inputs, [and] we [[have]] one item as a sanity check. You might look for other items that could be used as sanity checks to [[confirm]] the answers that are developed.

Ake: In a sense, a couple of those things that you [Salomone] just outlined, and that Walter [Arabasz] and Dave [Perkins] are trying to do, are things that we would be trying to do [anyway] as ways to cross check that we haven't made some grievous error.

Arabasz: Does this point speak to some project for the adoption of a standardized approach, which would have these elements or steps, and which allows some transparency rather than an ad hoc source by source approach?

Kammerer: To me, a structured Bayesian approach does that, and you can get more information about uncertainties, which is something that we need for a lot of projects. I like that because it gives me a much better sense of how to do a check on uncertainties.

Ake: Let's look at the next slide (fig. A15–4). The way we've broken it down is into two different approaches. One is some use of global analogs similar to what the USGS has used in their current approach, which is global analogs with experienced judgment. The alternative is a Bayesian approach that we've outlined based on Kevin's [Coppersmith] discussion here, and which is the development of the prior and the likelihood function. We've broken out the different data and techniques that one would use to construct the prior and the likelihood function. Does this make sense?

Coppersmith: I don't know how you get physical principles into the likelihood function [This refers to an early version of figure A15–5, which listed "physical principles" under the likelihood function as well as under the prior.]. Physical principles is the relationship between rupture length and width and magnitude, and things like that. That definitely can be used to modify your prior. The likelihood function is really looking at what's occurred within your zone, and counting.

Ake: My feeling is that things that should be there [in the likelihood function] for sure are paleoseismic and the observed earthquakes.

Coppersmith: Yes. It's easier to remove it from the prior to make it source-specific than it is to put it into the likelihood function. Statistical analyses [are involved in] counting of the observed events for the likelihood function.

Ake: Statistical analyses should stay [under the likelihood function] because beta and N enter into the development of the likelihood function whose parameters are developed by statistical analysis.

Lindvall: Physical principles could still be there because if you find a scarp like [the ones] in Australia that's how you convert scarp length into magnitude.

Coppersmith: I agree, but I would rather use it to modify the prior.

Kimball: You could put it on both sides. We have the underlying theme that this depends on how we define seismic zones. Maybe it is a sidebar [on the right], meaning that it is not directly part of the likelihood function, but we need to recognize that we have a dependency on how we define seismic sources. It seems that geological features, statistical analyses, and physical principles may become important as this dependency is true.

Coppersmith: My interest is that I need an Mmax technique that allows spatially varying Mmax. Not the whole eastern U.S., and not two subdivisions of it. I need spatial variability, so I would rather put physical principles on the prior side. I come now to my source zones, in which I want to estimate Mmax. I take all the information that comes from global analogs, and develop a prior distribution, and say that I can modify that prior with things like paleoseismic information, rupture dimensions, and other things. Then I'll look at the number of earthquakes that have occurred within that region, paleoseismic and historical and instrumental, and build up a likelihood function.

Kimball: I want us to recognize this dependency. If you take the limited amount of correction when you start microzoning, then the likelihood function becomes more meaningless. The fact is, we could advocate that the Bayesian approach is the right way to do it if we have a prior and a likelihood function, but if we're not careful and don't recognize that we have a dependency on how seismic sources are also defined, then we're kidding ourselves that we're making a breakthrough. We'll just transfer the problem to the next person who's trying to figure out how to do this so they can microzone Oklahoma.

Adams: That's a worry I have, too. It seems that with this Bayesian approach you could wind up with different [values of] Mmax for adjacent sources for reasons that may not be valid.

Coppersmith: We have adjacent sources with different Mmax's all the time.

Adams: Then I wonder what the basis is for your assumption that it should be the case.

Coppersmith: When you say "different" I picture two distributions that overlap. They won't be single-valued. I can't imagine a case where adjacent sources have no overlap, except perhaps at [the] New Madrid [seismic zone].

Ake: Unless you are starting with a radically different prior [in one of the adjacent zones].

Hanson: I think Jeff's [Kimball] concern is that understanding what Mmax method you're ultimately going to use affects how you subdivide or create a zonation model. But that's part of the first step, which is developing the zones. I'm not sure that's part of the Mmax method. I agree with Kevin [Coppersmith]: once you have your zones, then you have your prior and you modify it.

Ake: Let me make sure I understand what you are saying, Kevin [Coppersmith]. Suppose that we were doing zones that were relatively small, and for one of my zones I had my global prior that had a tail that extended to eight and one-half, but I was utterly convinced that based on the dimensions of my zone it just would not support an eight and one-half. I would then trim my prior based on those physical principles.

Coppersmith: That's almost an exact example from the 1994 report [(Johnston and others, 1994)] in looking at dimensions and trimming the tail. [Perhaps] I'd add more weight to places where I think the dimensions do more strongly support a large magnitude. The result would change from a smooth, flat prior to one that is more peaked and trimmed. That's reflecting additional information about the source.

Frankel: How do you apply these ideas when you use smoothed seismicity?

Ake: You'd start with a broad prior. If your zone is large so N is large, then Mobs [may be large].

Coppersmith: The problem is that Mobs changes with where you are. Can you adapt the smoothing kernel so it recognizes the size of the earthquake and truncates the distribution?

Frankel: It seems like you are putting in a correlation when you're doing this that may not be valid, to say the Mmax is correlated with the rate of earthquakes or something like that. You don't want to say that.

Ake: Look at the likelihood function. N gets bigger faster for high-rate areas. As N gets bigger, you have a higher chance of your maximum observed being close to Mmax.

Coppersmith: I'm not sure how you would gather the earthquakes to do this approach when there are no zones, other than just to take the entire zone.

Ake: I guess that was my presumption about how you would approach that problem, if you were doing it on a continental scale.

Frankel: When we originally started with the hazard maps, we had smoothed seismicity and we had these background zones. We had our workshops where we started with the whole CEUS, minus New Madrid and Charleston, and we asked: how would you break this down into source zones? The consensus was two [zones]. They didn't want anything else, and that's how we did it.

Adams: If I understand correctly what Kevin [Coppersmith] said about the sizes of ruptures within zones, no rupture can extend from one zone into another zone. That's a bit of a worry.

Coppersmith: No, you can have leaky boundaries, and you can have semi-permeable boundaries. Those are all tools that have been used now in hazard.

Adams: Does that say that you can't truncate the prior?

Coppersmith: It's just saying that there might be some characteristic lengths of this zone. It isn't so much that it is adjacent to other zones at the boundaries. It's just saying that we see zones of seismicity. For example, in Switzerland the Freiburg zone of seismicity is about 30 km long. I don't want to make an eight and one-half on this zone.

Kimball: There are two big blue boxes [in figure A15–4]. Is there clear agreement that those are the two right paths to take? [Should there be] a third or a fourth in the figure? This is a pretty big change from what PSHAs have done in the past.

Coppersmith: The box in the figure that's called Bayesian has been applied in a lot of places.

Wheeler: It's not been applied to the two big PSHAs [of EPRI–SOG and the USGS national hazard maps].

Coppersmith: But it has been applied in other settings, for example for TVA.

Lettis: I would still hold out that, if you know your fault, then you can use fault dimensions to estimate Mmax without going through global analogs or a Bayesian approach. [Is this figure] for the CEUS zones excluding [the] Meers [and] Cheraw [faults], and other [similar] faults that we might identify?

Ake, others: Yes. [At about this point figure A15–5 was added.]

Lettis: I'm trying to think ahead to when we actually try to apply this. We're looking at Triassic basins. As we go from basin to basin, one model for generating earthquakes in the CEUS might be reactivation of the basin-bounding faults. We know the dimensions of that fault. Would you use that to modify the prior of the region, as opposed to saying that this is the Mmax that we could have?

Coppersmith: From my point of view, you could do it either way and you'd get the same answer. But that might be an alternative approach to another one that says, maybe those aren't the ones that are being reactivated.

Chapman: You [Lettis] were talking about something that is almost like a deterministic estimation of magnitude based on the length of a fault. Something like that can be incorporated into the Bayesian estimation procedure. It doesn't require something external to the Bayesian approach.

Coppersmith: I'd like to go to the top box [of figure A15–4]. To me it's an informal application of exactly what we've got in the Bayesian approach. You [Wheeler] showed your plot of the numbers of earthquakes globally, with stars on the ones that were North American [fig. A7–6], and then you developed a four-point distribution of Mmax for the two types of crust. To me those are very similar to an informal application of the same concept as the Bayesian approach.

Wheeler: They're very similar. The differences are in the level of pooling. You're right, the two big superduperdomains are in the Bayesian analysis. They are in the prior, but the prior also has superdomains in it. The global analogs don't use superdomains. We start at the top and divide the global catalog into two and stop there. [The global analog and Bayesian methods] share that feature [using the two superduperdomains]. I identified the North American earthquakes with asterisks, but I don't treat them any differently. I added those asterisks so more people would pay attention to the rest of the histograms. The histograms mean the same without the asterisks.

Kimball: So the national map is not a Bayesian analysis?

Wheeler: Right.

Kimball: Then the question comes, there are two approaches, one is Bayesian and one is not. Are they viable alternatives?

Wheeler: I think that's a valid statement because it allows them to have certain features in common, which is true.

Coppersmith: It [the global analog approach] starts out with the same data, takes all the global SCR data and divides it into two regimes, then it plots them [magnitudes] on a histogram, and then there's judgment on what seems to be a reasonable maximum earthquake given those global data. But it doesn't take into account the number of earthquakes in a likelihood function. It doesn't take into account physical principles or paleoseismic because it's not source-specific, it's for a whole region. It doesn't include a lot of the local information that potentially the Bayesian could include. At a minimum it [the Bayesian approach] includes the same [information]: global analogs and their size [earthquake magnitudes].

Wheeler: I'll admit that I don't understand all parts of the Bayesian approach well enough to know whether I believe in it. But it's certainly encouraging, not least because it includes the global catalog. [For example], for the Clinton [nuclear power] site in central Illinois that several of us worked with several years ago, the paleoseismic results included an M 6.8 close to the site.

Hanson: The magnitude range was M 6.2–6.8. There was uncertainty in the size.

Wheeler: Right. The site is in our cratonic source zone, so we said [Mmax was] 7.0; at that [time] we were not putting a distribution on it. Whatever result the Bayesian analysis came up with, it was so close to seven that we decided we had no argument with it. The numbers were different, but within the uncertainty they were indistinguishable.

Hanson: We used paleoseismology, and we put in uncertainty in the number of possible paleoearthquakes that the dataset we had would allow, which was from one to as many as four or five. We used the uncertainty in the magnitude of the one paleoearthquake that was best documented from the paleoseismic record, and sort of a craton prior. We wouldn't have had a tail that went out to seven plus or eight. I don't know whether we truncated it based on the reassessment of New Madrid at that time. The likelihood function based on the paleoearthquakes moved a large part of the [low–M] tail into the central part of the distribution, so that there was a lot of mass at M 6.7 or 6.8.

Wheeler: That's one case in which the two approaches produced indistinguishable answers, so they [the two approaches] are compatible.

Petersen: Bob Youngs told me at the workshop that we held for the SSC that the likelihood function really doesn't make very much difference. All the work really goes into developing the prior.

Coppersmith: Because N is so small.

Petersen: Right. The ones [methods in fig. A15–4] that are really going to matter are the global analogs, maybe some physical principles if you can identify some faults in the area that have certain dimensions, and North American analogs. That's about what you've got.

Hanson: I think paleoliquefaction [results can change] that. When you include those events, that's the closest to having a moderate-sized earthquake [and can make the likelihood function more important].

Wheeler: Think of what drives our interest in paleoseismology in the East. We're looking for prehistoric earthquakes bigger than anything we've seen in the historical record. That means that almost inevitably the paleoseismology is going to drive the priors up.

Coppersmith: If we go [with the] Bayesian [method, then] people have got to get used to the idea of global analogs. The Mmax distribution is going to be broader than we've ever seen before. If people are used to seeing Mmax's around five or five and a half, they're going to astounded by this approach.

Kammerer: There's a lot of epistemic uncertainty here, and to me it seems like a broad distribution might be more representative of the uncertainty that we have.

Coppersmith: From a process point of view, the Bayesian prior is broad appropriately. We can [modify and] narrow the prior in various ways, but we should start broad. That isn't often the mind set. The Bayesian analysts say we shouldn't start broad, that often we should start with some best estimate and try to add uncertainty to it.

Wheeler: One partial digression: I think global analogs should include whatever Walter Mooney comes up with. [He's] looking at all the cratons and seeing what he can tell about our craton.

Kammerer: For the CEUS can we say that we're comfortable with these two approaches [fig. A15–4], and that the Bayesian approach is preferred to characterize uncertainty? Are these the two preferred approaches in general?

Ake: Let me rephrase the question. Do these two, at this point, represent the [viewpoint of the] informed technical community on how to approach this problem?

Coppersmith: Speaking for the people I know in Europe, they are really intrigued and are definitely using this approach. They like gathering information from around the world and using it more rigorously than just subjectively.

Lettis: Would you use the Bayesian approach for New Madrid?

Coppersmith: Yes. The maximum observed allows truncation of the lower tail so you have an upper tail. With uncertainty in the values of the parameters, you have a distribution.

Frankel: I'm bothered by this break down into two things because they seem like two different kinds of things. One is global analogs, which is something you look at. Bayesian is a technique. Like Bill [Lettis] says, you could have another box for, say, fault lengths. Even though it's in the Bayesian, it can be a separate technique, and you can do it without doing the Bayesian. You can do global analogs without doing the Bayesian. You have judgment everyplace here. You have judgment in the global analogs, and you have judgment in the prior. To say the Bayesian is some objective scheme is [invalid].

Kammerer: No one's saying it's objective. It's just a framework.

Coppersmith: I don't think anyone would ever say that. In fact, statistical analysts will deny the Bayesian is a proper approach.

Frankel: The bottom line is that you have the distribution of the maximum magnitude you've seen globally, and it's totally non-unique how you got that. Whether you really think that the average Mmax of anyplace is six and one-half, and seven to seven and on-half is on the tail of the distribution, or if you think the real global Mmax is seven and one-half and because you have a shorter observation period you see more of those six and one-halves clumped up on the left sides of those plots, that's the problem. What you see is totally non-unique. You don't know what the real underlying distribution is. You can parse the Bayesian in terms of these bias corrections [and other things], but all of those are assumptions that are not necessarily true.

Kammerer: That's true, except that the advantage of the likelihood function is that you take what has a lot of expert opinion in it [the prior], and you at least temper it by the [historical] record.

Frankel: The largest earthquake you've seen in these regions globally might be the only information you really have. That might be it. The rest is just a sampling issue.

Ake: In one sense I agree with you. The first time I saw these plots [fig. A7–6] from Rus [Wheeler] I thought, okay, this is very interesting, but some people would immediately gravitate to these two highest points and say: for each of these two regions, that's Mmax. But keep in mind what you touched on a minute ago. This is a sampling problem, something that doesn't happen very often, and this is a couple of hundred years worth of data. Would we have anything like this if our 200-year observation period was [a different] 200-year observation period? As infrequent as these things are, probably not. In answer to your question, there is a distribution drawn based on these observations and there is clearly some subjective judgment. The way I look at the Bayesian approach is that it's just a slightly different framework that sets out a more detailed process to go through [in order] to come up with that [distribution]. That's probably the only thing I can think about that's different about [these two methods].

Petersen: Is there a formal approach for creating a prior? I understand that there's a formal approach for linking the likelihood function with the prior, but is there a formal approach for developing a prior?

Ake: This [fig. A7–6] is a prior. [It was not developed to be a prior, but it could be used as one.]

Petersen: So we've developed a prior. You can say it's not Bayesian but in some ways they are pretty similar techniques.

Wheeler: If you use this as a prior, then what's your likelihood function? That earthquake and that earthquake [the single highest magnitude in each histogram]. Throw away everything else, and there are your Mmax's. Is that reasonable?

Coppersmith: It was made clear from the beginning of the Mmax project that there's a belief that Mmax varies spatially around the world. If you don't believe that then you'll go to the largest [earthquake] you've seen. We believe that the largest earthquake observed in each zone has the potential to be the maximum for that zone. If you believe that all SCRs have [the same] maximum earthquake, then you'll follow the approach you described. Your likelihood function will truncate at the minimum size of the maximum observed earthquake, which will be the largest we've seen anywhere.

Wheeler: Yes, and that's why I'm glad that the national maps are not Bayesian.

Kammerer: But the question you're answering with the maps is not what is the biggest earthquake around the world, because if it was, then you couldn't justify less than that.

Wheeler: We're not sure what we're seeing because the magnitudes are uncertain. That's why I left out the outliers and why I think we want to look at something in that long flat tail. Also, I think it's very unlikely that Mmax is the same everywhere in all the world's cratons, but I have no idea how to break it down.

Frankel: The Bhuj earthquake was M 7.6–7.7. The aftershock zone was maybe 30–50 km long. Are you telling me that you can't fit that in any source zone you draw in the CEUS? How small a source zone are you going to draw? Seismic reflection [data] shows us faults all over the place. They may not all be active, but there are all these faults that are that long and much longer. These things are not quite smoking guns but they are out there. How can you say that Mmax is different everywhere?

Adams: That's a hypothesis: to test whether these faults have the same Mmax or different ones.

Perkins: One of the things that often is done is to convene a bunch of experts for an expert elicitation. If you convene a bunch of experts and ask them to estimate the maximum magnitude in a zone, would that be an acceptable procedure? I don't think it's reflected in the two boxes [fig. A15–4].

Coppersmith: I view that as a process rather than a method.

Perkins: It's a process in which one of these two boxes?

Coppersmith: It's not [in either box]. They would use various methods to get there, and these [the box headings] might be the methods they would use. You'd need a tool from the toolbox to get Mmax, and the process you would use to capture Mu would be a SSHAC level three or four [process] or whatever it is. That's the process, but these [fig. A15–4] are tools to get at Mmax. We're trying to define what those tools are. And they all have assumptions. There's an assumption in the Bayesian approach and that's that Mmax varies spatially around the world.

Kimball: Probably it's fair to say that the informed technical community does include people who believe that the most likely Mmax in the near future is Mobs. The good thing about the Bayesian [approach] is that it allows that to be factored in.

Coppersmith: That's where you'd start. The distribution at the low end will be the observed [magnitudes]. That's not the expected [magnitude], you're right. I'm sure there are those in the community who will say "I've seen an M 5.8, and [so] Mmax is M 5.8."

Petersen: Jeff [Kimball], you referred to a time frame ["in the near future"]. We've never referred to a six-year period as being [appropriate for Mmax]. [The USGS national hazard maps are updated approximately every six years.]

Coppersmith: There are some people who would say, "forever". There are some people who would say, "I told you five to five and one-half, and in 25 years it hasn't been exceeded; I don't think it ever will be."

Petersen: There are some zones where they've done that, so they had to go back and explain them, even in the last 20 years.

Coppersmith: That's why I think this will start at the observed and go up. All of these distributions will be higher than they were 25 years ago. That's it. You're looking more at a rare earthquake and not thinking about an expected event.

Adams: I am thinking that this Bayesian method that you're talking about does put a low limit on Mmax, as you just said. That's a good thing. I just don't know whether it gives you the right number. That's what I'm not happy about. Even if we came up with a number we'd never know if it's exactly the right number.

Petersen: We have to define Mmax. Is it [something of short applicability], or this really something longer lasting? There's a difference.

Ake: Are we ready to wrap this up?

Wheeler: I think we've come a long way. We've knocked out a couple of methods that used to be used a lot. The two methods that we favor most, if we don't have an actual fault to study, use global analogs, which was considered in both Livermore and EPRI–SOG but not given too much weight by most people. From a biased viewpoint, I'm glad to hear what Kevin [Coppersmith] forecast for the future of Mmax [in studies done outside the USGS], which is that it will gradually go up. I'd also like to point out that whatever we decide here, no agency, no company and no person can be obliged to follow it. That doesn't detract from the substantial progress that we've made here.

Kammerer: On behalf of the NRC I'd like to thank you all very much for your time and your thoughtful comments. I, and probably all of us from the NRC, feel like this was really useful.

Kimball: A SSHAC level 4 process involves a group of experts who have come together by themselves to select the best methods to use. The conclusion of this group is that the best methods to use are these two [fig. A15–4]. If someone uses a different method, it would be regarded as not the best method to use. The implication that we all have to be comfortable with is that, if someone were to have a level 4 [process], there would not be an alternate method that could be defended as legitimate.

Ake: Only if the method they were using, especially if it were the sole method they were using, were essentially one of the methods that Rus [Wheeler] pulled together for us and which we looked at. If it were a completely new and different approach, and if they could provide a sufficiently good technical basis for it, why wouldn't they be able to use it?

Kimball: I was thinking about [fig. A4–1]. What could happen is that you would get somebody who is supposedly part of the informed community, and they would pluck one of the methods out of the table. We are saying that the community does not view that as an appropriate method.

Kammerer: There would be a feedback loop with the informed community. If [someone did that] they would be asked why they did it, and they would have to justify it. As long as they understand the implications of everything and can justify what they want to do, they could do it. Also, the SSHAC process is supposed to capture both the center and the range of informed opinion.

Lettis: It seems like the Bayesian approach can factor in about every method that I can think of into either modifying the prior or developing a likelihood function.

Kammerer: Class dismissed.

Opinion Polls and Summary

Attendees were asked to fill out two opinion polls, one at the start of the workshop and the other at the end. The first poll characterized the opinions with which the attendees arrived, and the difference between the two polls summarized how the opinions changed during the workshop. The polls asked the attendees to evaluate each of 13 methods of estimating Mmax that were to be considered during the workshop (fig. A4–1; Wheeler, 2009). The polls asked three questions about each method: (1) is the method invalid, (2) is the method unready for use, that is, does it require more work, and (3) if the method is valid and ready for use, what weight should it be assigned when combined with other valid, ready methods?

Of the 35 attendees, 25 turned in the first poll but only 15 turned in the second one. Accordingly, the poll taken at the start of the workshop probably represents the center and range of opinions of the attendees. However, it is unclear whether the second poll is similarly representative because of its small response. Large differences between the two polls are more likely to represent changes in the collective opinions of the attendees than are small differences. Tables 5–7 summarize the two polls and their differences.

Table 7 shows the main result of the workshop. Early in the workshop a consensus began to emerge that Mmax estimation should be based on the global catalog of SCR earthquakes, which was compiled by Johnston and others (1994). The two main methods that are based on the global catalog are that of global tectonic analogs, which the USGS uses for its national hazard maps, and the Bayesian analysis, which is used in industry PSHAs that follow SSHAC processes. The polls showed that both methods were among the most favored methods at the start of the workshop, and that by the end [of the workshop] both of these methods and the use of North American analogs were the favorites (table 7). The second poll showed that the Bayesian analysis was assigned approximately twice the median weight of the methods of global and North American analogs (table 6). The preferences for methods using the global catalog, and for the Bayesian analysis among those methods, could be seen developing during the course of the workshop. Accordingly, both preferences appear to represent majority opinions despite the small sample size of the second poll.

A second result has implications for both past and current PSHAs. Methods 1–6 formed the core of Mmax estimation in many industry PSHAs in recent decades. Nonetheless, at the start of the workshop methods 1–5 were among the least favored and by the end method 6 had joined them. I do not know whether attendees had formed unfavorable opinions about these methods long before the workshop, or whether their opinions had been affected by the analyses in the foundation document that they received a month before the workshop (Wheeler, 2009). If the former is the case, then it is possible that some past PSHAs used estimation methods that did not represent the center of informed opinion at the time that the PSHA computations were done.

A third result is that paleoseismology is now seen as the most feasible way to acquire additional direct observations of large SCR earthquakes within a reasonable time. Figure A15–5 illustrates the pervasive importance of paleoseismology. The figure lists paleoseismology as one method among several that could be used to develop a Bayesian prior distribution, as a tool for estimating Mobs for a Bayesian likelihood function, and as a specific kind of local geologic information. Wheeler's insert at the start of the Monday afternoon minutes notes that the USGS national seismic-hazard maps make heavy use of paleoseismic chronologies of large earthquakes. Where an active fault has been studied paleoseismically, its scarp length and single-earthquake displacements can provide two independent estimates of Mmax for the fault. Early in the workshop Ake noted that the importance of paleoseismic information lends urgency to improving the characterization of its uncertainties.

The workshop produced other results. First, Mobs was discarded as a valid estimator of Mmax, except where a large historical or prehistoric earthquake has occurred, as at the New Madrid, Charleston, Charlevoix, and Wabash Valley seismic zones and the Cheraw and Meers faults. Second, elaborate physical and statistical models lost favor as being untestable in SCRs, at least with present information. Third, geodetic measurements of strain may soon be able to provide upper limits on Mmax.

Acknowledgements

The U.S. Nuclear Regulatory Commission provided funding for the workshop. I thank the attendees for making the workshop a success. J. Ake and A. Kammerer provided frequent questions, answers, suggestions, and patience. C. Mueller, A. Nelson, and D. Perkins took detailed notes, and R. Smith taped the full two days of discussions. Without these informal records, no useful summary could have been made. All speakers generously allowed reproduction of their slides as figures in the appendixes. Figures of non-USGS speakers are unchanged except that I added figure numbers and deleted slides that were not used. USGS slides benefitted from USGS editing. The clarity of the manuscript was improved by reviews of C. Mueller and A. Nelson, and by comments and corrections from J. Adams, D. Clark, C. Cramer, J. Ebel, K. Hanson, S. Harmsen, A. Kammerer, J. Kimball, C. Mueller, and L. Salomone. Any remaining errors in the minutes are mine.

References Cited

Atkinson, G.M., and Boore, D.M., 1995, Ground-motion relations for eastern North America: Bulletin of the Seismological Society of America, v. 85, no. 1, p. 17–30.

Budnitz, R.J., Apostolakis, G., Boore, D.M., Cluff, L.S., Coppersmith, K.J., Cornell, C.A., and Morris, P.A., 1997, Recommendations for probabilistic seismic hazard analysis—Guidance on uncertainty and use of experts: Washington, D.C., U.S. Nuclear Regulatory Commission NUREG/CR–6372, 2 v., 1,109 p.

Chinnery, M.A., 1979, Investigations of the seismological input to the safety design of nuclear power reactors in New England: U.S. Nuclear Regulatory Commission report NUREG/CR–0563, 72 p.

Coppersmith, K.J., 1994a, Conclusions regarding maximum earthquake assessment, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, assessment of large earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 6–1—6–24.

Coppersmith, K.J., 1994b, Introduction, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, assessment of large earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 1–1—1–10.

Coppersmith, K.J., Johnston, A.C., Metzger, A.G., and Arabasz, W.J., 1987, Methods for assessing maximum earthquakes in the Central and Eastern United States: Palo Alto, California, Electric Power Research Institute, working report EPRI RP2556–12, 312 p.

Cornell, C.A., 1994, Statistical analysis of maximum magnitudes, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, assessment of large earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 5–1—5–27.

Cramer, C.H., Petersen, M.D., Cao, T., Toppozada, T.R., and Reichle M., 2000, A time-dependent probabilistic seismic-hazard model for California: Bulletin of the Seismological Society America, v. 90, p. 1–21.

Danforth, S., 1880, Journal of Rev. S. Danforth, Roxbury, Massachusetts: New England Historical and Genealogical. Registry, v. 34, p. 84–89, 162–166, 297–301, 359–363.

Ebel, J.E., 1996, The seventeenth century seismicity of northeastern North America: Seismological Research Letters, v. 67, no. 3, p. 51–68.

Ebel, J.E., Bonjer, K.-P. and Oncescu, M.C., 2000, Paleoseismicity—seismicity evidence for past large earthquakes: Seismological Research Letters, v. 71, no. 2, p. 283–294.

Exxon Production Research Company, 1985, Tectonic map of the World: 21 sheets, scale 1:10,000,000 at the equator (copyrighted and publicly distributed by American Association of Petroleum Geologists Foundation, Tulsa, Oklahoma, 1994).

Frankel, A., 1995, Mapping seismic hazard in the Central and Eastern United States: Seismological Research Letters, v. 66, p. 8–21.

Frankel, A., Mueller, C., Barnhard, T., Perkins, D., Leyendecker, E.V., Dickman, N., Hanson, S., and Hopper, M., 1996, National seismic-hazard maps—documentation June 1996: U.S. Geological Survey Open-File Report 1996–532, 70 p., available online at *http://pubs.er.usgs.gov/usgspubs/ofr/ofr96532*. (Accessed June 19, 2009.)

Frankel, A.D., Petersen, M.D., Mueller, C.S., Haller, K.M., Wheeler, R.L., Leyendecker, E.V., Wesson, R.L., Harmsen, S.C., Cramer, C.H., Perkins, D.M., and Rukstales, K.S., 2002, Documentation for the

2002 update of the national seismic hazard maps: U.S. Geological Survey Open-File Report 2002–0420, 39 p., available online at http://pubs.er.usgs.gov/usgspubs/ofr/ofr02420. (Accessed June 19, 2009.)

Frankel, A.D., Petersen, M.D., Mueller, C.S., Haller, K.M., Wheeler, R.L., Leyendecker, E.V., Wesson, R.L., Harmsen, S.C., Cramer, C.H., Perkins, D.M., and Rukstales, K.S., 2005, Seismic-hazard maps for the conterminous United States: U.S. Geological Survey Scientific Investigations Map 2883, scale 1:7,000,000, 6 sheets, available online at *http://pubs.er.usgs.gov/usgspubs/sim/sim2883*. (Accessed June 19, 2009.)

Johnston, A.C., 1989, The seismicity of 'stable continental interiors,' *in* Gregersen, S., and Basham, P.W., eds., Earthquakes at North-Atlantic passive margins—neotectonics and postglacial rebound: Dordrecht, The Netherlands, Kluwer Academic Publishers, p. 299–327.

Johnston, A.C., 1994a, Appendix C—summary tables, SCR seismicity database, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 2: Palo Alto, California, Electric Power Research Institute, p. C–1—C–46.

Johnston, A.C., 1994b, Seismotectonic interpretations and conclusions from the stable continental region seismicity database, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, Assessment of large earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 4–1—4–103.

Johnston, A.C., 1994c, The stable continental region database, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, Assessment of large-earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 3–1—3–80.

Johnston, A.C., 1996, Seismic moment assessment of earthquakes in stable continental regions—I. Instrumental seismicity: Geophysical Journal International, v. 124, p. 381–414.

Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., 1994, The earthquakes of stable continental regions: Palo Alto, California, Electric Power Research Institute, 5 v., 2,519 p., 16 folded plates, 1 diskette.

Kanter, L.R., 1994a, Appendix B—tectonic domain data sheets, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions, v. 2: Palo Alto, California, Electric Power Research Institute, p. B–1—B–196.

Kanter, L.R., 1994b, Tectonic interpretation of stable continental crust, *in* Johnston, A.C., Coppersmith, K.J., Kanter, L.R., and Cornell, C.A., eds., The earthquakes of stable continental regions—v. 1, Assessment of large earthquake potential: Palo Alto, California, Electric Power Research Institute, p. 2–1—2–98.

Kijko, A., and Graham, G., 1998, Parametric-historic procedure for probabilistic seismic hazard analysis, Part I—Estimation of maximum regional magnitude m_{max}: Pure and Applied Geophysics, v. 152, p. 413–442.

Klimkiewicz, G.C., 1980, Earthquake ground motion attenuation models for northeastern United States Earthquakes: Chestnut Hill, Massachusetts, Boston College, M.S. thesis, 150 p.

Mueller, C., Hopper, M., and Frankel, A. 1997, Preparation of earthquake catalogs for the national seismic-hazard maps—contiguous 48 states: U.S. Geological Survey Open-File Report 1997–464, 36 p.

Nishenko, S.P. and Bollinger, G.A., 1990, Forecasting damaging earthquakes in the Central and Eastern United States: Science, v. 249, p. 1,412–1,415.

Obermeier, S.F., Gohn, G.S., Weems, R.E., Gelinas, R.L., and Rubin, M., 1985a, Geologic evidence for recurrent moderate to large earthquakes near Charleston, South Carolina: Science, v. 227, p. 408–410.

Obermeier, S.F., Weems, R.E., Gohn, G.S., and Rubin, M., 1985b, Distribution and recurrence of prehistoric earthquakes near Charleston, South Carolina [abs.]: Earthquake Notes, v. 55, no. 1, p. 25.

Petersen, M.D., Cramer, C.H., and Frankel, A.D., 2002, Simulations of seismic hazard for the Pacific Northwest of the United States from earthquakes associated with the Cascadia subduction zone: Pure and Applied Geophysics, v. 159, p. 2,147–2,168.

Petersen, M.D., Frankel, A.D., Harmsen, S.C., Mueller, C.S., Haller, K.M., Wheeler, R.L., Wesson, R.L., Zeng, Y., Boyd, O.S., Perkins, D.M., Luco, N., Field, E.H., Wills, C.J., and Rukstales, K.S., 2008, Documentation for the 2008 update of the United States national seismic-hazard maps: U.S. Geological Survey Open-File Report 2008–1128, 119 p., Available online at http://pubs.er.usgs.gov/usgspubs/ofr/ofr20081128. (accessed June 19, 2009.)

Reasenberg, P.A., and Jones, L.M., 1989, Earthquake hazard after a mainshock in California: Science, v. 243, p. 1173–1176.

Risk Engineering, Inc., Geomatrix Consultants, Inc., Woodward-Clyde Consultants, and Cygna Corporation, 1986, Seismic hazard methodology for the Central and Eastern United States: Palo Alto, California, Seismicity Owners Group and Electric Power Research Institute report NP–4726, 11 vol.

Savage, J.C., 1991, Criticism of some forecasts of the National Earthquake Prediction Evaluation Council: Bulletin of the Seismological Society of America, v. 81, p. 862–881.

Toro, G.R., Abrahamson, N.A., and Schneider, J.F., 1997, Model of strong ground motions from earthquakes in Central and Eastern North America—Best estimates and uncertainties: Seismological Research Letters, v. 68, p. 41–57.

Van der Lee, S., and Nolet, G., 1997, Upper mantle S velocity structure of North America: Journal of Geophysical Research, v. 102, no. B10, p. 22,815–22,838.

Wells, D.L., and Coppersmith, K.J., 1994, New empirical relationships among magnitude, rupture length, rupture width, rupture area, and surface displacement: Bulletin of the Seismological Society of America, v. 84, p. 974–1,002.

Wheeler, R.L., 2009, Methods of Mmax estimation east of the Rocky Mountains: U.S. Geological Survey Open-File Report 2009–1018, 43 p. (Available online at http://pubs.er.usgs.gov/usgspubs/ofr/ofr20091018).

Wheeler, R.L., and Frankel, A., 2000, Geology in the 1996 USGS seismic-hazard maps, Central and Eastern United States: Seismological Research Letters, v. 71, no. 2, p. 273–282.

Wheeler, R.L., and Johnston, A.C., 1992, Geologic implications of earthquake source parameters in Central and Eastern North America: Seismological Research Letters, v. 63, no. 4, p. 491–514.

Table 1. Attendees and affiliations.

Name	Affiliation
Adams, John	Geological Survey of Canada
Ake, Jon	U.S. Nuclear Regulatory Commission
Allen, Trevor	U.S. Geological Survey/Golden, and Geoscience Australia
Arabasz, Walter	University of Utah Seismograph Stations, University of Utah
Chapman, Martin	Virginia Tech Seismological Observatory, Virginia Tech
Coppersmith, Kevin	Coppersmith Consulting, Inc.
Cramer, Chris H.	Center for Earthquake Research and Development, University of Memphis
Ebel, John	Weston Observatory, Boston College
Frankel, Art	U.S. Geological Survey/Golden
Fuller, Chris	William Lettis and Associates, Inc.
Halchuk, Stephen	Geological Survey of Canada
Haller, Kathy	U.S. Geological Survey/Golden
Hanson, Kathryn	AMEC Geomatrix Consultants, Inc.
Harmsen, Steve	U.S. Geological Survey/Golden
Hough, Susan	U.S. Geological Survey/Pasadena
Jarernprasert, Sittipong	Paul C. Rizzo Associates, Inc.
Johnston, Arch	Center for Earthquake Research and Development, University of Memphis
Kammerer, Annie	U.S. Nuclear Regulatory Commission
Kim, Won-Young	Lamont-Doherty Earth Observatory, Columbia University
Kimball, Jeff	Defense Nuclear Facilities Safety Board
Lettis, William	William Lettis and Associates, Inc.
Li, Yong	U.S. Nuclear Regulatory Commission
Lindvall, Scott	William Lettis and Associates, Inc.
Mooney, Walter	U.S. Geological Survey/Menlo Park
Mueller, Chuck	U.S. Geological Survey/Golden
Murphy, Andrew	U.S. Nuclear Regulatory Commission
Nelson, Alan	U.S. Geological Survey/Golden
Perkins, David	U.S. Geological Survey/Golden
Personius, Steve	U.S. Geological Survey/Golden
Petersen, Mark	U.S. Geological Survey/Golden
Salomone, Lawrence	U.S. Department of Energy Savannah River Site
Stirewalt, Gerry L.	U.S. Nuclear Regulatory Commission
Wheeler, Russell L.	U.S. Geological Survey/Golden
Williams, Robert	U.S. Geological Survey/Golden
Withers, Mitch	Center for Earthquake Research and Development, University of Memphis

Table 2. Acronyms used in the text.

Acronym	Definition
CEUS	Central and eastern United States
CEUSAC	Central and eastern United States and adjacent Canada
CUS	Central United States
DEM	Digital elevation model
EPRI–SOG	Electric Power Research Institute-Seismicity Owners Group
GPS	Global Positioning System
GSC	Geological Survey of Canada
Hz	Hertz (cycles per second)
LiDAR	Light Detection and Ranging
M	Magnitude
m_b, m_{bLg}	Body wave magnitude
Mchar	A characteristic magnitude used in USGS PSHA computations
M-f	Making a graph of frequency versus M for historical seismicity
Mmax	The largest magnitude thought to be possible within a specified area
Mobs	The largest magnitude observed within a specified area
M_S	Surface wave magnitude
Mu	Upper bound magnitude (Mmax)
NEIC	USGS National Earthquake Information Center
NEUS	Northeastern United States
NRC	United States Nuclear Regulatory Commission
PDE	USGS NEIC Preliminary Determinations of Epicenters
PGA	Peak ground acceleration, in percent g
PSHA	Probabilistic seismic hazard assessment
SCC	Stable continental crust earthquake catalog of the GSC
SCR	Stable continental region
SRTM	Shuttle Radar Topography Mission
SSC	Central and Eastern U.S. Seismic Source Characterization Project
SSHAC	Senior Seismic Hazard Analysis Committee
STP	South Texas Project
TVA	Tennessee Valley Authority
USGS	U.S. Geological Survey
WWSSN	Worldwide Standard Seismograph Network

Table 3. Speaker index.

[Bold and underlined page numbers: scheduled talk. Double square brackets: contains numerous or substantial corrections submitted by speaker (*see* Minutes of Discussions in text). Normal page numbers: comment or question. A speaker may be listed more than once on a single page]

Name	Pages Containing Talks or Comments
Adams	19, 23, 24, 33, 34, 46, 48–50, 68, **79–81**, 81–82, **82–86**, 86–90, 94, 98–99, 101–102, 108, [[115–116]], 118, 121–122, 126, 130, 132, 139, 141
Ake	**9–11**, 11–12, 24, 33, 47, 64–65, **70–71**, 72–74, 76–79, 88, 101, 111–113, **119**, 122, 125, 128–129, 131–133, 136, 138, 141–142
Allen	59–60, **91–97**, 97–98, 100–102, [[123]]
Arabasz	72, 86–87, 109, 126, 127–128
Chapman	**20–22**, 22–24, 67, 76, 78, 111, 133
Coppersmith	12, 25, **25–31**, 31–36, 51, 53, 68, 73–78, 88–90, 98, 110, 112, 120–121, 125, 127–137, 139–141
Cramer	[[69]], 69, 71
Ebel	65, 69–70, 81, 97, 109, **112–118**, 118
Frankel	12, 35–36, 46, 50, 54–55, 66, 108, 115, 116, 118, 122, 128, 131–132, 137–138, 139
Hanson	34–35, 72, 77–78, 110, 122–123, 126, 131, 135–136
Harmsen	24, **60–63**, 63–68
Hough	16, 19, 23, 35, 48, 59, 65–67, 71, 75, 114, 121, 125
Johnston	12, 18, 31, 35, 47, 55, 58, 64, 74, 77, 101, 108, 115, 121
Kammerer	**4–5**, **20–21**, 68, 70–71, **119–120**, 120–122, 124–128, 136–139, 141–142
Kim	114
Kimball	16–17, 24, 71–72, 88, 107, 114, 126, 129–130, 132, 134, 140–142
Lettis	18, 31, 46–47, 63, 71, 74–75, 78, 90, 98, 100, 107, 109, 111, 118, 125, 133, 137, 142
Li	82, 100, 108, 125
Lindvall	50, 71, 102, 125, 129
Mooney	11–12, 32, **54–58**, 58–60, 99
Mueller	**[[5–9]]**
Murphy	19, 65–66, 126
Nelson	124
Perkins	22, 127–128, 140
Petersen	50, 101, 135, 138, 140-141
Salomone	11, 17, 68, 73, 76, 102, 120, 128
Stirewalt	93, 94, 99, 124
Wheeler	**13–16**, 17–19, 31, 32, **[[36–40]], 40–46**, 47–51, 53, 68, 74–76, 78, 81, 87, 89, 97–99, 102, **102–107,** 108–112, 115, 120, 122–124, 126, 127, 133–136, 138–139, 141

Table 4. Earthquakes mentioned in the text.

[Years are listed separately because an earthquake is often referred to only by its year of occurrence]

Name	Year	Date	Location
Alaska	1964	March 28	United States, southern Alaska, Prince William Sound
Anjar	1956	July 21	Northwestern India, Gujarat Province
Attica	1929	August 12	United States, western New York State
Baffin Bay	1933	November 20	Arctic Canada, northeast of Baffin Island
Bhuj	2001	January 26	Northwestern India, Gujarat Province
Cadoux	1979	June 2	Southwestern Australia
Calingiri	1970	March 10	Southwestern Australia
Cape Ann	1755	November 18	United States, offshore northeastern Massachusetts
Charleston	1886	September 1	United States, central-coastal South Carolina
Charleston	1895	October 31	United States, southeastern Missouri
Charlevoix	1663	February 5	Canada, southeastern Quebec, along St. Lawrence River
Charlevoix	1925	March 1	Canada, southeastern Quebec, along St. Lawrence River
Chile	1960	May 22	Central Chile
Coalinga	1983	May 2	United States, central California
China	1917	July 30	Eastern China
China	2008	May 12	Southern China, eastern Sichuan Province
El-Asnam	1980	October 10	Northern Algeria
Fort Payne	2003	April 29	United States, northeastern Alabama
Giles County	1897	May 31	United States, southwestern Virginia
Grand Banks	1929	November 18	Atlantic Ocean, south of Newfoundland, Canada
Kutch	1819	June 16	Northwestern India, Gujarat Province
Marked Tree	1843	January 5	United States, northeastern Arkansas
Meckering	1968	October 14	Southwestern Australia
Miramichi	1982	January 9	Canada, central New Brunswick
Nahanni	1985	October 5 and December 23	Canada, Northwest Territories
New Madrid	1811–1812	December 16, January 23, February 7	United States, southeastern Missouri and adjacent Arkansas, Tennessee, and Kentucky
South Tasman Rise	1951	May 1	South of Tasmania, Australia
Saguenay	1988	November 25	Canada, southeastern Quebec, north of St. Lawrence River
Tennant Creek	1988	January 22	Northern Central Australia
Timiskaming	1935	November 1	Canada, northeastern Ontario

Table 5. Opinions of validity and readiness-for-use of Mmax estimation methods.

[Each entry is the number of attendees who turned in an opinion form and assigned a method to the indicated column. Forms were distributed and collected at the start of the workshop, and again at the end to measure opinion changes during the workshop. Some respondents left some cells blank on their forms, and others checked both cells for some methods. See table 2 for definitions of abbreviations]

Method	Start of workshop		End of workshop	
	Invalid	Not ready to use	Invalid	Not ready to use
1. Mobs	13	0	5	1
2. Mobs + constant	9	0	9	0
3. Seismicity rate	13	3	10	2
4. M-f extrapolation	11	3	12	3
5. m_b saturation at 7.5	20	1	14	1
6. Local geology	2	2	3	3
7. North American analogs	3	1	1	2
8. Global analogs	1	0	0	0
9. Bayesian analysis	2	3	1	1
10. Physical principles	2	5	3	6
11. Statistical analyses	4	4	6	6
12. Pattern recognition	9	10	9	5
13. Q_0	12	7	13	2

Table 6. Summary of weights assigned by all respondents.

[See table 2 for definitions of abbreviations. N, number of attendees who entered weights for the method; leaders (--), median and range are undefined because N=0]

Method	Start of workshop			End of workshop		
	Median	Range	N	Median	Range	N
1. Mobs	0.075	0.03–0.2	7	0.075	0.05–0.1	4
2. Mobs + constant	.1125	.1–.3	12	.05	.05–.1	3
3. Seismicity rate	.1125	.05–.2	6	.05	.05–.05	1
4. M-f extrapolation	.1	.05–.15	7	--	--	0
5. m_b saturation at 7.5	.03	.03–.03	1	--	--	0
6. Local geology	.15	.1–.3	16	.1	.1–.2	5
7. North American analogs	.1	.03–.3	18	.2	.1–.35	6
8. Global analogs	.225	.075–1.0	22	.25	.125–1.0	9
9. Bayesian analysis	.25	.025–1.0	15	.5	.2–1.0	10
10. Physical principles	.1	.025–.3	10	.1	.1–.2	5
11. Statistical analyses	.1	.025–.2	13	.125	.125–.125	1
12. Pattern recognition	.0625	.025–.1	2	--	--	0
13. Q_0	.03	.03–.03	1	--	--	0

Table 7. Summary of validity, readiness for use, and weights from tables 5 and 6.

[The columns refer to the sum of numbers of evaluations of a method as "Invalid?" and "Not ready for use?" in table 5. The rows refer to the median weights of table 6. Leaders (--), no entries. Method numbers are those in tables 5 and 6 of this report and table 5 of Wheeler (2009). At both the start and the end of the workshop, the methods judged the best are in the upper left cells and those judged the worst are in the lower right cells. The order of methods within a cell is numerical and does not indicate relative quality. See table 2 for definitions of abbreviations]

Median weight	Number of respondent evaluations of the method as invalid or not ready for use			
	Start of workshop		End of workshop	
	Fewer than 9	9 or more	Fewer than 9	9 or more
Greater than 0.3	--	--	9. Bayesian analysis	--
0.15-0.29	6. Local geology 8. Global analogs 9. Bayesian analysis	--	7. North American analogs 8. Global analogs	--
Less than 0.15	7. North American analogs 10. Physical principles 11. Statistical analysis	1. Mobs 2. Mobs + constant 3. Seismicity rate 4. M-f extrapolation 5. m_b saturation at 7.5 12. Pattern recogition 13. Q_0	1. Mobs 6. Local geology	2. Mobs + constant 3. Seismicity rate 4. M-f extrapolation 5. m_b saturation at 7.5 10. Physical principles 11. Statistical analysis 12. Pattern recogition 13. Q_0

Appendixes: Speakers' Power Point Slides

Appendixes are separate files.

Appendix 1: Goals of the Workshop (A. Kammerer)

Figure A1–1.

Appendix 2: Sensitivity of Hazard to Hifferent Values of Mmax (C. Mueller)

Figures A2–1 to A2–13.

Appendix 3: Uncertainties (J. Ake)

Figures A3–1 to A3–6.

Appendix 4: Overview of Methods of Estimating Mmax (R. Wheeler)

Figures A4–1 to A4–2.

Appendix 5: Mmax and Earthquake Catalogs (M. Chapman)

Figures A5–1 to A5–9.

Appendix 6: Bayesian Estimates of Mmax (K. Coppersmith)

Figures A6–1 to A6–12.

Appendix 7: USGS Estimates of Mmax (R. Wheeler)

Figures A7–1 to A7–7.

Appendix 8: Mmax in Cratons (W. Mooney)

Figures A8–1 to A8–23.

Appendix 9: Uncertainties in Parameters Related to Mmax (S. Harmsen)

Figures A9–1 to A9–13.

Appendix 10: GPS Constraints on Hazard in Eastern Canada (J. Adams)

Figures A10–1 to A10–6.

Appendix 11: Mmax in Canada's Fourth-Generation Hazard Model (J. Adams)

Figures A11–1 to A11–11.

Appendix 12: Neotectonics in Australia (T. Allen)

Figures A12–1 to A12–24.

Appendix 13: What is a "Tectonic Analog"? (R. Wheeler)

Figures A13–1 to A13–5.

Appendix 14: Magnitude of the 1663 Charlevoix Earthquake (J. Ebel)

Figures A14–1 to A14–10.

Appendix 15: What Are Our Conclusions About Methods of Estimating Mmax? (A. Kammerer, J. Ake)

Figures A15–1 to A15–5.

Figure A1–1

NRC-USGS Workshop on CEUS Mmax Sept. 8-9, 2008

- **Goal**: reach **consensus** on one or more of

- (1) **Relative values** of methods of Mmax estimation
- (2) Which methods are **invalid**, or promising but **not yet ready for use**
- (3) **Research needed** to achieve (1) and (2)

USGS

Figure A2-1

- EPRI/SOG Mmax
  - Six earth-science teams, diverse methods
    - largest observed eq (+ increment)
    - catalog statistics – extreme occurrences
    - seismogenic feature: size, crustal setting, etc.
    - extrapolate frequency-magnitude curve (e.g., the 1,000-year eq)
    - global analogs
    - saturation of m_b scale
    - others …
  - Broad distributions reflect diverse methods & large uncertainties
  - Approximate center: mid-m_b-5 (quiet sites), mid-m_b-6 (active sites)
- USGS Mmax
  - CEUS global analogs: stable continental regions (AJ), Bhuj
    - Craton: **M** 7.0
    - Extended Margin: **M** 7.5
    - Mmax distribution for 2008

Figure A2-2

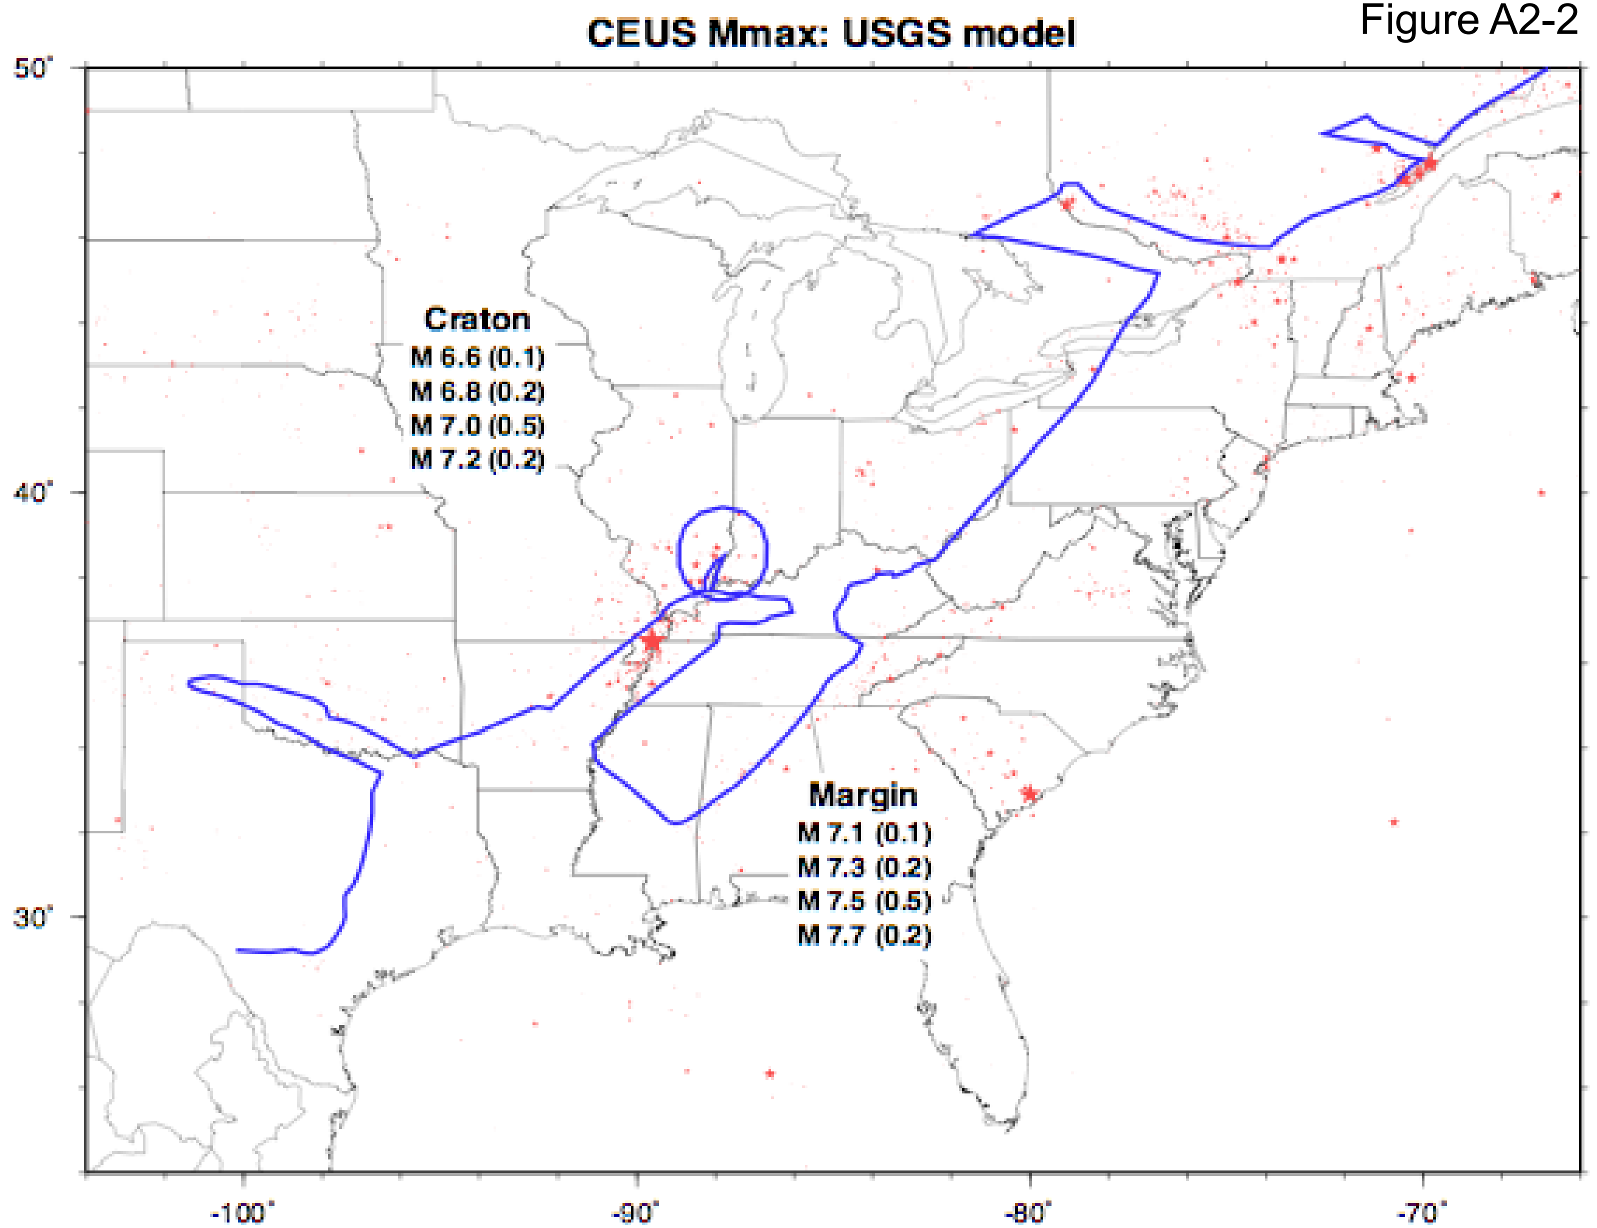

USGS
science for a changing world

Figure A2-3

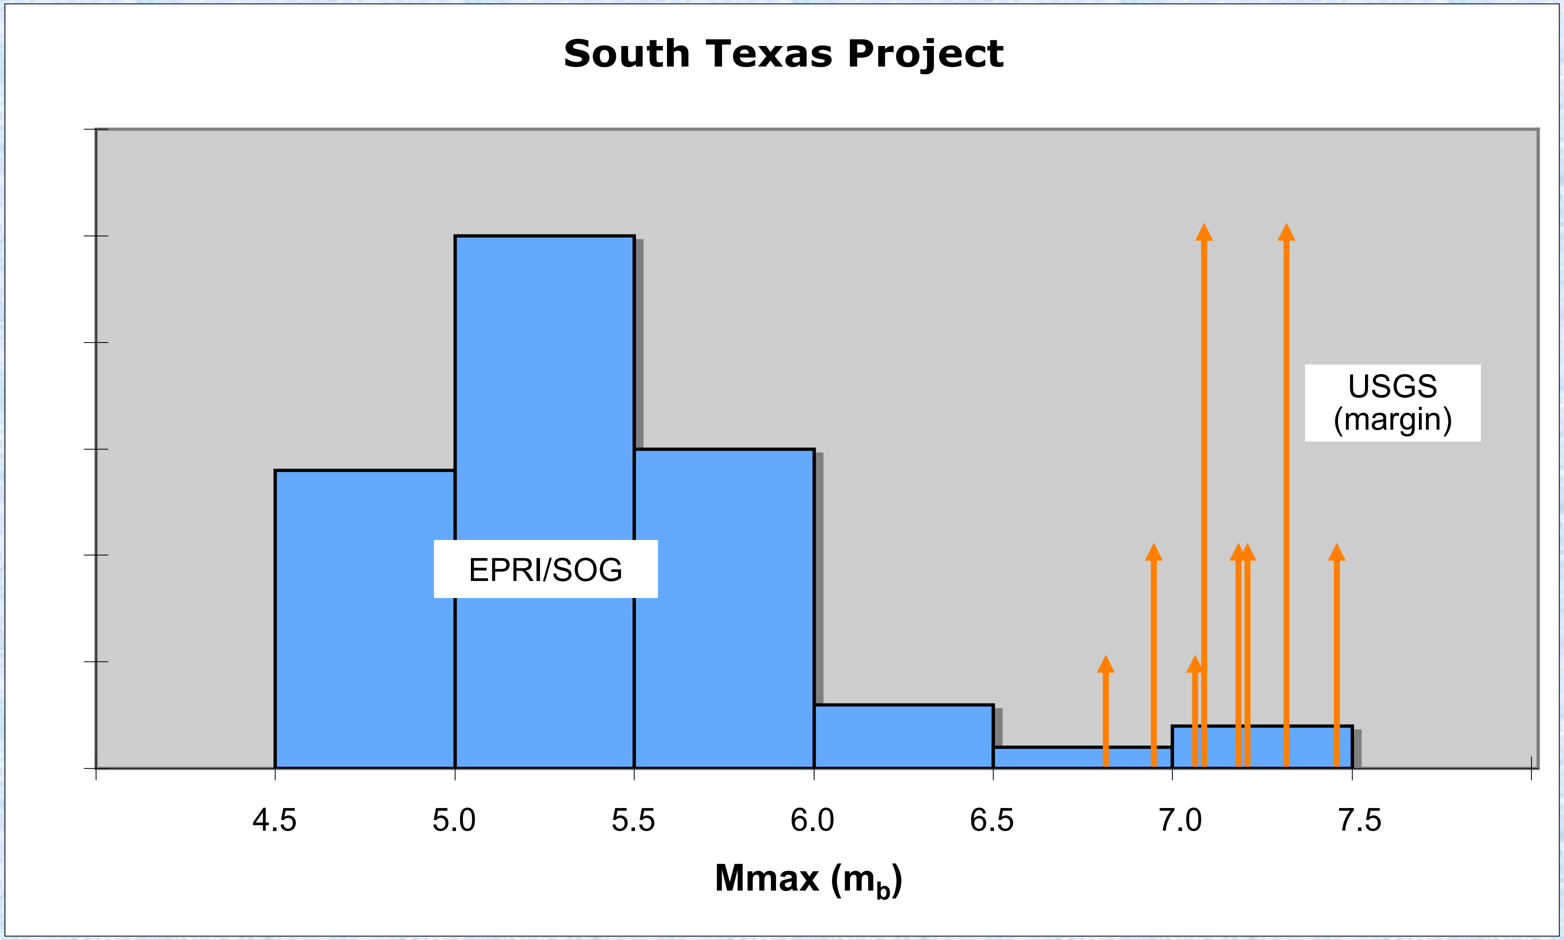

EPRI/SOG and USGS Mmax distributions for South Texas Project site's host source zones (original EPRI/SOG, not updated for 2006 Gulf of Mexico earthquakes)

Figure A2-4

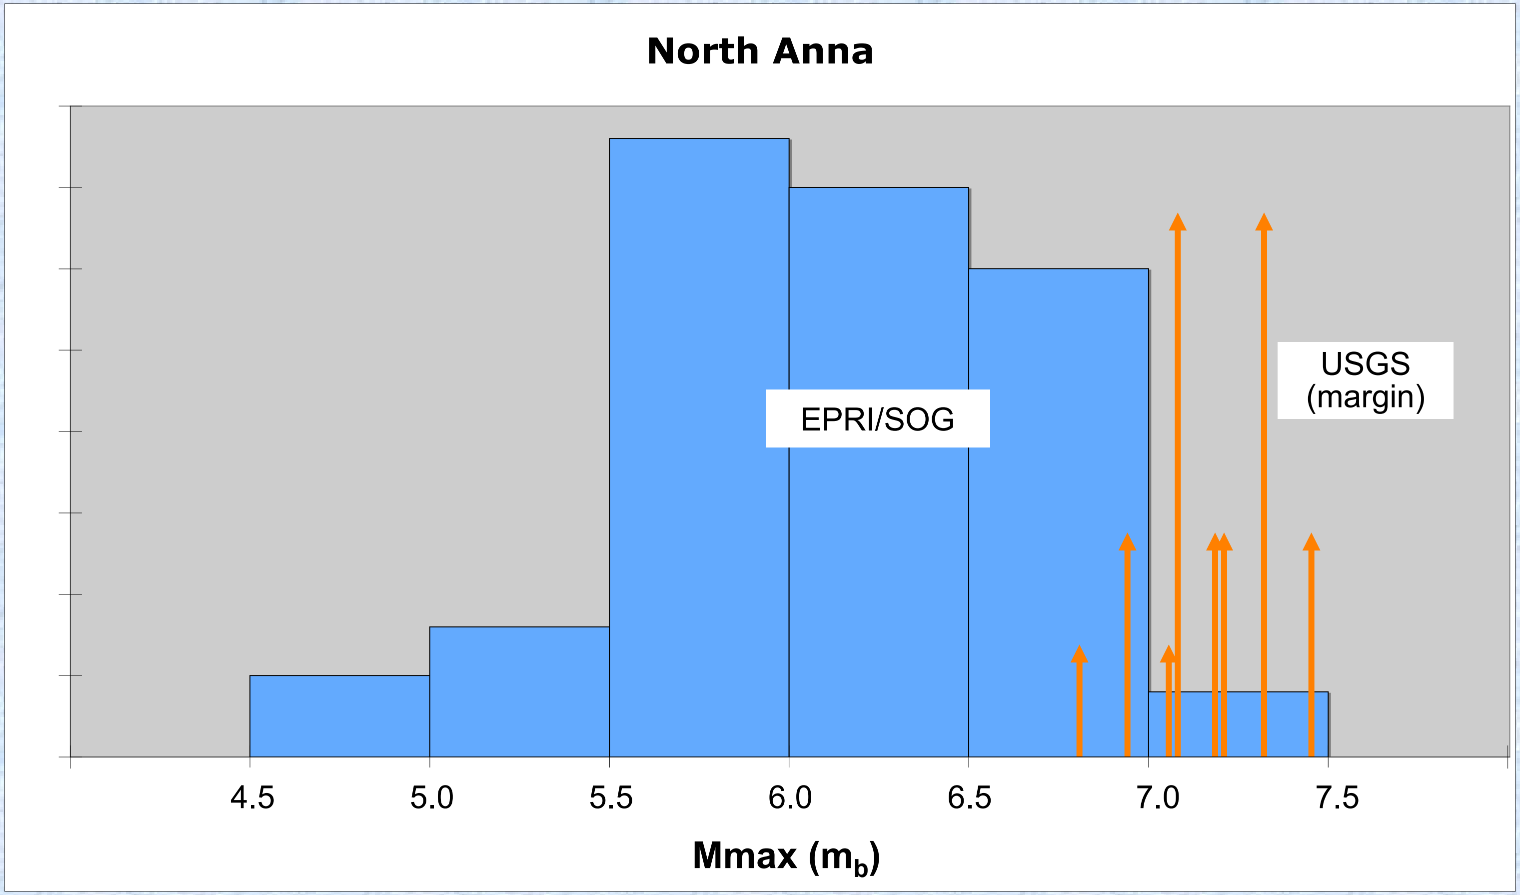

EPRI/SOG and USGS Mmax distributions for North Anna site's host source zones

Analysis

- Develop alternative Mmax models that span those used in EPRI/SOG (and current PSHA practice)
- Test Mmax models using USGS hazard model and computer codes
  - seismicity-based sources only (because they control the mid- to high-frequency hazard at many sites & Mmax is uncertain and controversial)
  - hold all parameters fixed except Mmax
- Compare hazard results with current USGS model as ratio hazard maps and lists for selected sites
  - numerator: alternative Mmax / denominator: standard USGS Mmax
  - probabilistic ground motions for PGA, 5 Hz, 1 Hz
  - 2% probability of exceedance in 50 years

Alternative Mmax Models

- M5.0c5.5m
  - M 5.0 in craton, M 5.5 in margin
  - m_b equivalent: 5.47c, 5.90m (AB95) or 5.27c, 5.67m (J96)

- M6.0c6.5m
  - M 6.0 craton, M 6.5 margin
  - m_b equivalent: 6.29c, 6.66m (AB95) or 6.04c, 6.40m (J96)

- M7.0c7.5m
  - M 7.0 craton, M 7.5 margin
  - m_b equivalent: 7.00c, 7.32m (AB95) or 6.74c, 7.07m (J96)

Figure A2-7

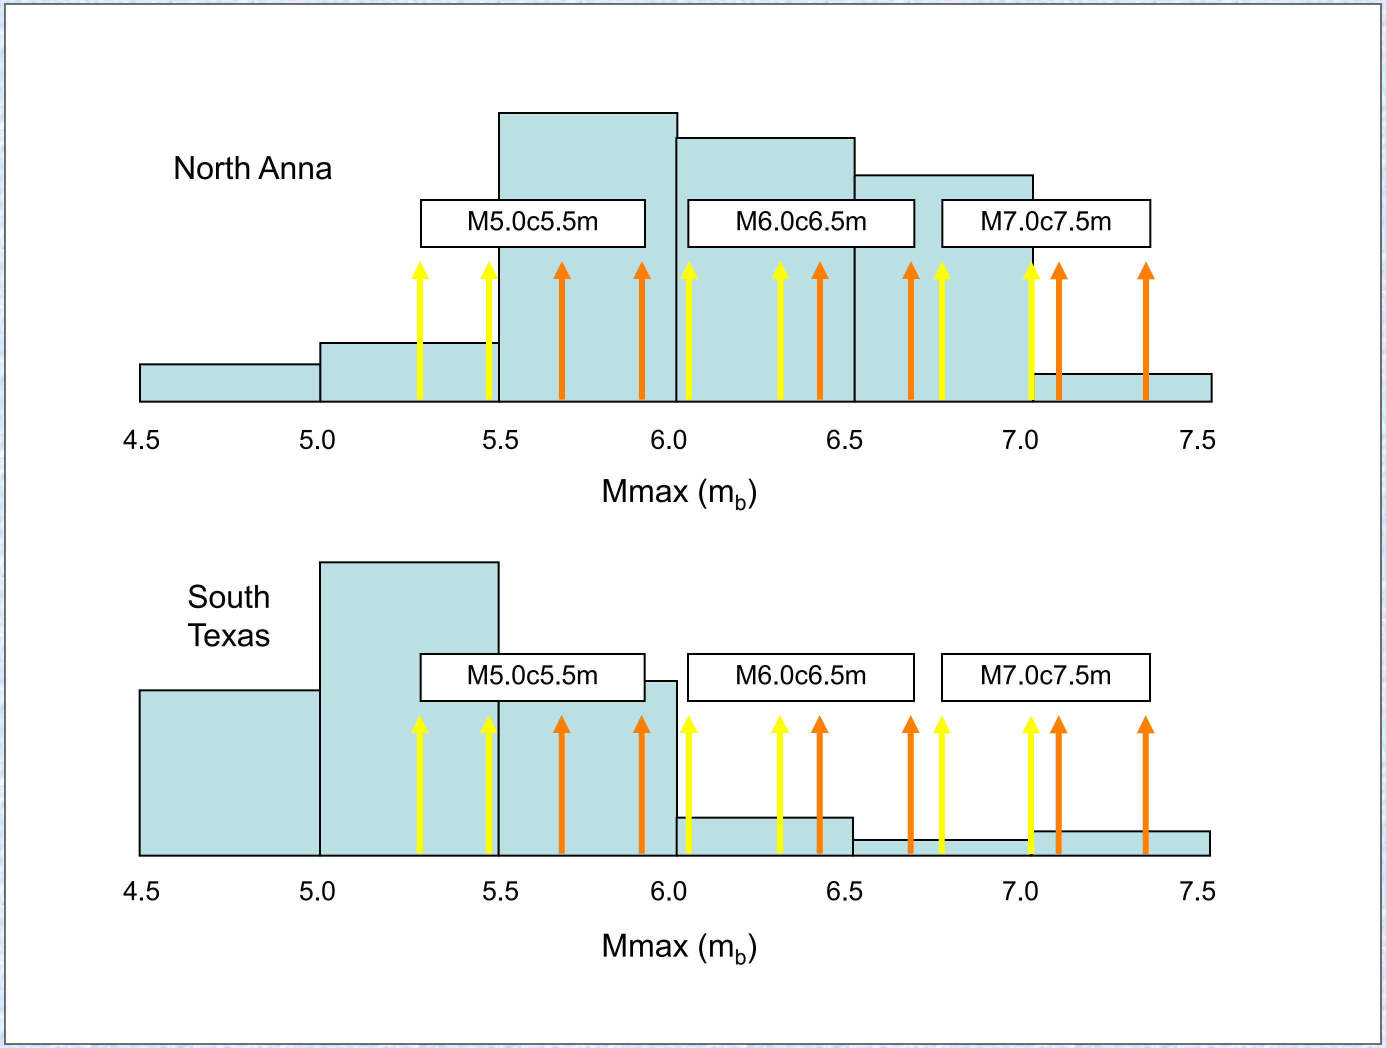
North Anna
M5.0c5.5m
M6.0c6.5m
M7.0c7.5m
4.5
5.0
5.5
6.0
6.5
7.0
7.5
Mmax (m_b)
South Texas
M5.0c5.5m
M6.0c6.5m
M7.0c7.5m
4.5
5.0
5.5
6.0
6.5
7.0
7.5
Mmax (m_b)

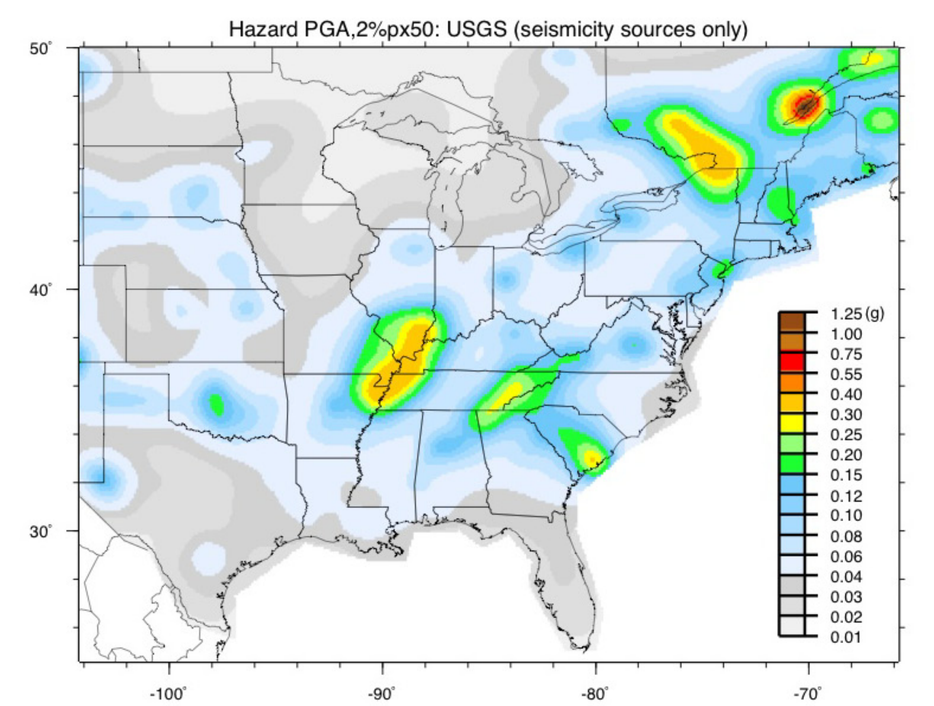

Figure A2-8

Hazard map (PGA, 2% prob of exc in 50 yrs): seismicity sources only, standard USGS Mmax

(Prob of exc, px: probability of exceedance)

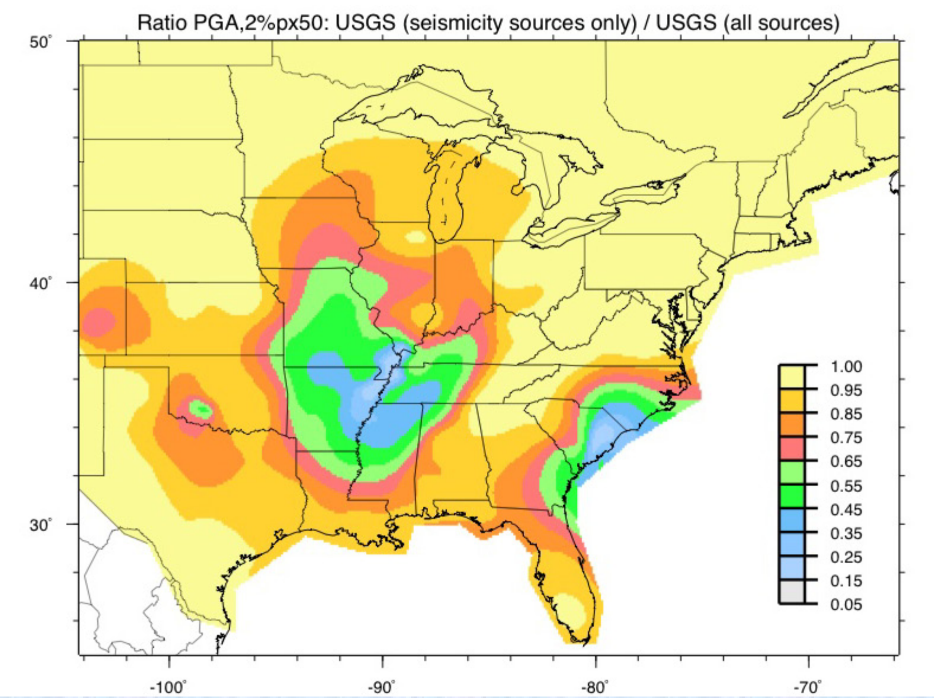

Ratio hazard map (PGA, 2% prob of exc in 50 yrs): seismicity sources only / all sources

Warm colors show where seismicity-based sources control hazard.

USGS
science for a changing world

Figure A2-9

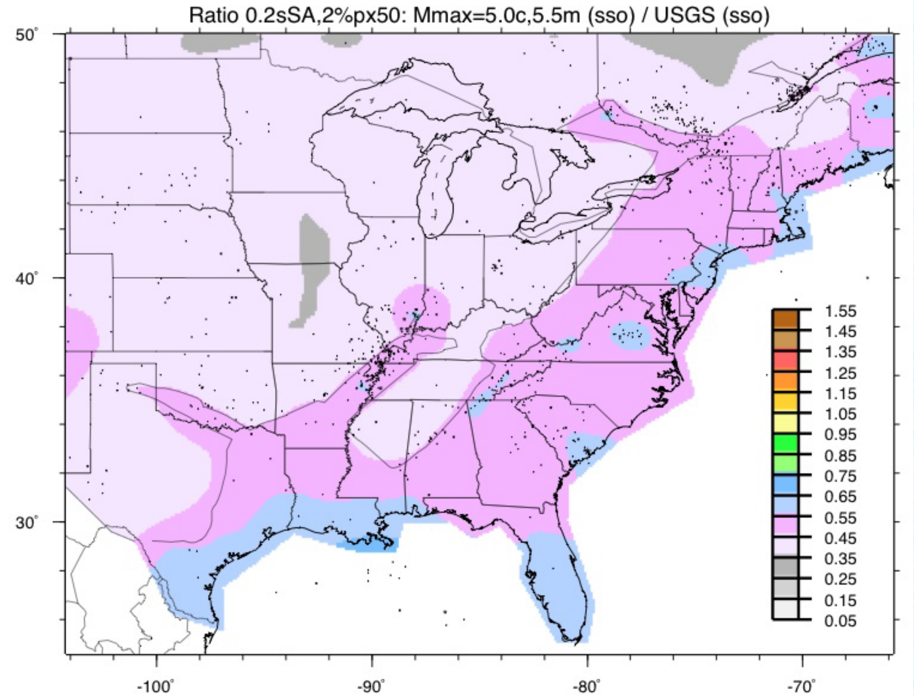

← 0.2-sec SA

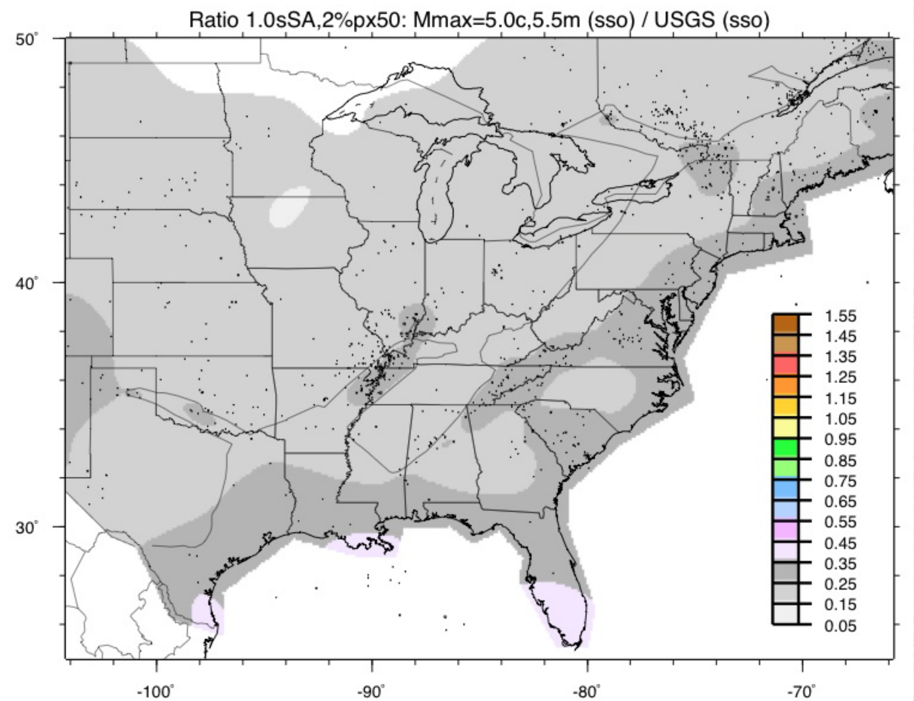

← 1.0-sec SA

Alternate Mmax model: M5.0c5.5m

Ratio hazard maps (2% prob of exc in 50 yrs)

M5.0c5.5m (sso) / standard USGS Mmax model (sso)

(Prob of exc, px: probability of exceedance)

USGS science for a changing world

Figure A2-10

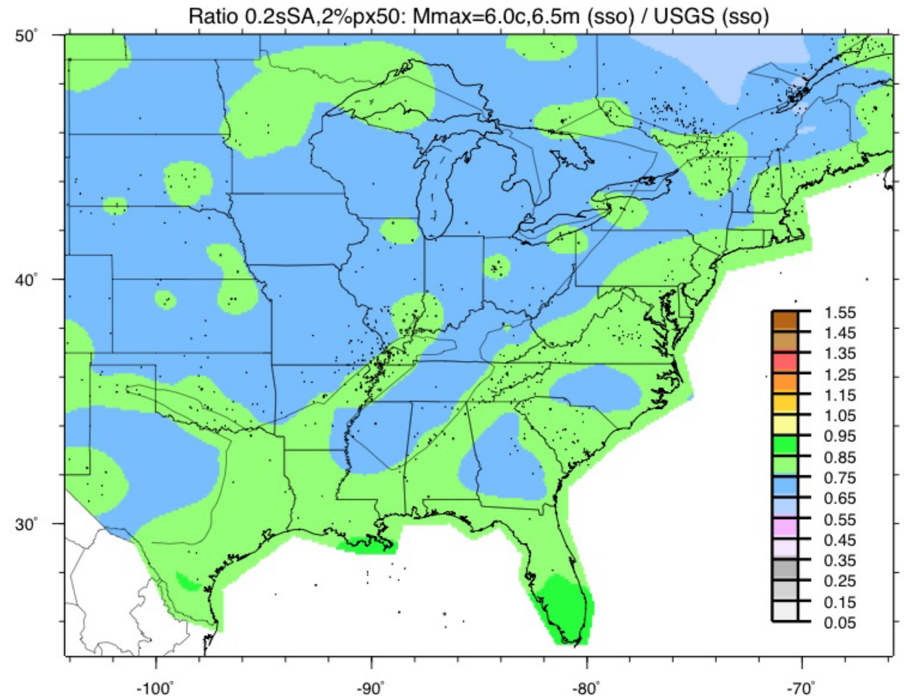

0.2-sec SA

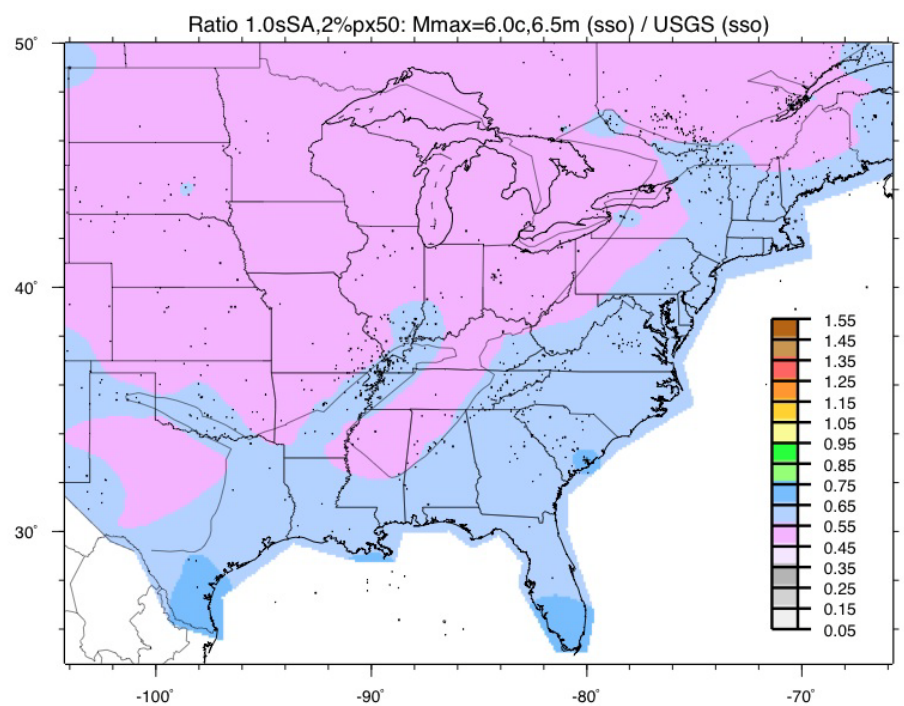

1.0-sec SA

Alternate Mmax model: M6.0c6.5m

Ratio hazard maps (2% prob of exc in 50 yrs)

M6.0c6.5m (sso) / standard USGS Mmax model (sso)

(Prob of exc, px: probability of exceedance)

Figure A2-11

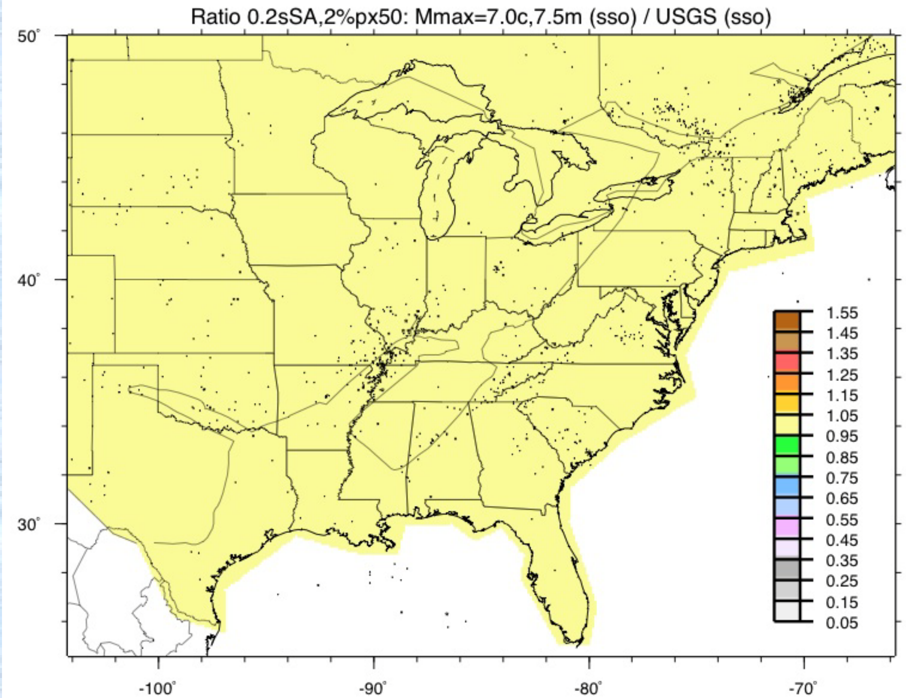

0.2-sec SA

Alternate Mmax model: M7.0c7.5m

Ratio hazard maps (2% prob of exc in 50 yrs)

M7.0c7.5m (sso) / standard USGS Mmax model (sso)

(Prob of exc, px: probability of exceedance)

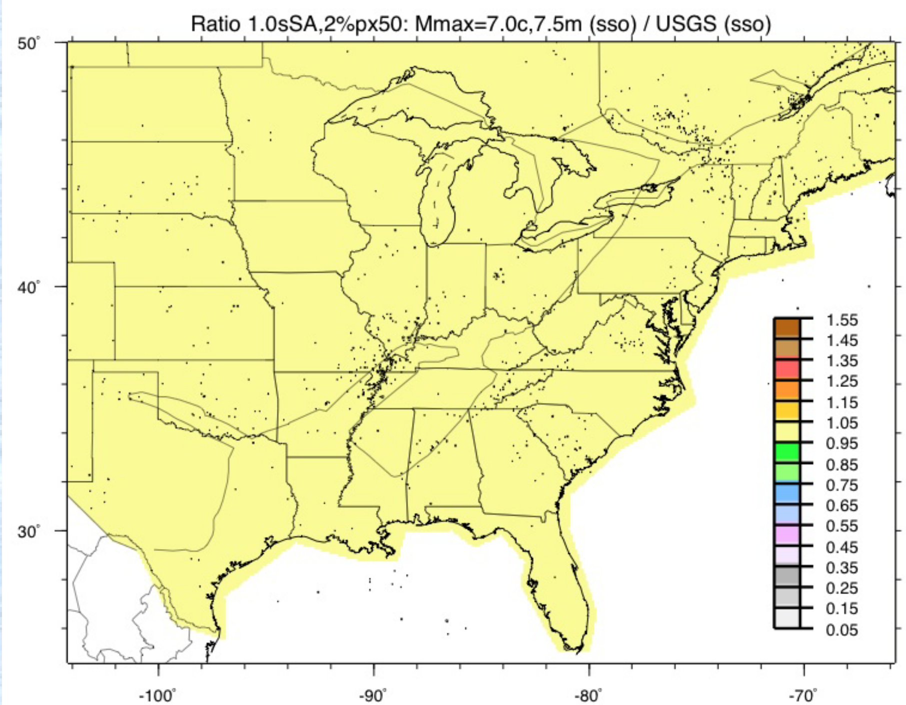

1.0-sec SA

USGS
science for a changing world

Ratios of probabilistic ground motions (2% probability of exceedance in 50 yrs): Alternate Mmax (seismicity sources only) / USGS (seismicity sources only)

	M5.0c5.5m		M6.0c6.5m		M7.0c7.5m	
	0.2 sec	1.0 sec	0.2 sec	1.0 sec	0.2 sec	1.0 sec
Boston	0.55	0.29	0.79	0.60	1.01	1.02
New York City	.57	.32	.83	.64	1.01	1.02
Washington D.C.	.51	.26	.77	.58	1.01	1.02
Pittsburgh	.45	.21	.74	.53	1.01	1.02
Charleston	.59	.34	.83	.67	1.01	1.02
Atlanta	.49	.24	.74	.56	1.02	1.03
Cincinnati	.38	.18	.70	.51	1.01	1.02
Chicago	.37	.18	.75	.52	1.01	1.02
Memphis	.52	.25	.78	.58	1.01	1.02
Baton Rouge	.57	.31	.82	.61	1.01	1.02
St. Louis	.38	.19	.72	.53	1.01	1.02
Minneapolis	.38	.16	.74	.49	1.00	1.01
Wichita	.37	.18	.72	.51	1.01	1.01
Austin	.49	.26	.78	.58	1.01	1.02
Rapid City	.37	.19	.72	.51	1.00	1.01
Denver	.41	.25	.73	.54	1.00	1.01

USGS
science for a changing world

Figure A2-13

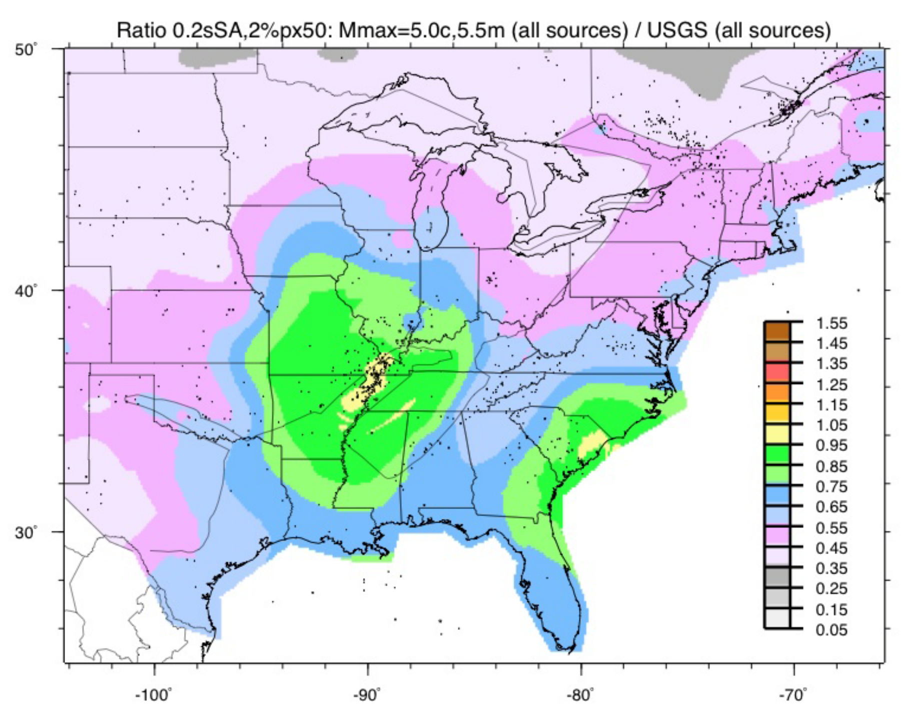

← 0.2-sec SA

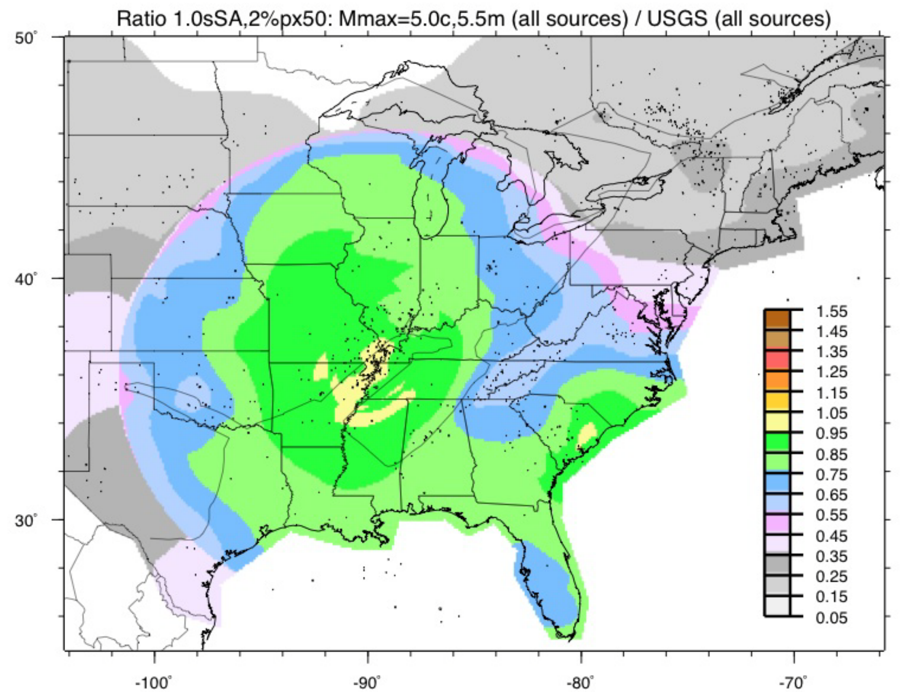

← 1.0-sec SA

Alternate Mmax model: M5.0c5.5m

Ratio hazard maps (2% prob of exc in 50 yrs)

Use M5.0c5.5m for numerator seismicity sources, but **add the faults** to both the numerator and denominator models

(Prob of exc, px: probability of exceedance)

USGS
science for a changing world

Uncertainty in Seismic Hazard Assessments (Truth in Advertising/Full Disclosure)

M_{MAX} Workshop

USGS-Golden

September 8/9, 2008

Jon Ake, U.S. NRC

Figure A3–2

Uncertainty in Seismic Hazard Assessments

- Regulatory Requirements/Guidance: 10 CFR 100.23 recognizes the nature of uncertainty in seismological and geological evaluations and the need to account for these uncertainties. Further developed in Regulatory Guides 1.165 and 1.208.
- Similar considerations are imbedded in DOE Standards, recent ANS/ANSI Standards, ICODS Guidelines
- Good Practice

Figure A3–3

Uncertainty and Variability

1. Aleatory Variability

-The natural randomness in a process

-Not predictable

-Can’t be reduced (theoretically)

2. Epistemic Uncertainty

-Knowledge-based

-Unknown models/parameters

-Competing models

-Potentially reducible

-Generally incorporated via logic trees

Figure A3–4

Examples

Epistemic: Alternative models to predict magnitude based on Intensity

- -Maximum Intensity (MMI_{MAX}) vs Area within a MMI contour ($A_{I\text{-}VI}$)
- -Each of these models has an uncertainty associated with it.

Epistemic: Alternative models to predict magnitude based on fault characteristics

- -Fault length vs displacement
- -Each of these models has an estimate of variability associated with it

Figure A3–5

Uncertainty in Seismic Hazard Assessments

- Need to take a pragmatic approach
- A snapshot in time
- Problems especially acute in SCR
- Space for time substitution is a given (both for development of models and estimates of variability)

Figure A3–6

Summary

- Necessary to formally include uncertainty/variability in our assessments
- We need to develop a well-defined process that will include and track the aleatory and epistemic uncertainties, but avoid “double-counting” of uncertainty
- Need to define <u>credible</u> alternative models

Methods Overview

Figure A4-1

Method	Small Sample	Insens. Mmax?	Paleoseismic	"Analog?"
Mobs	SS(source zone)		Inconsistent	
Mobs + c	SS(source zone)		Inconsistent	
Eq rate	SS(global sample)	Insensitive		
M-f extrapol'n	SS(source zone)		Inconsistent	
m_b 7.5				
Local geology	SS(source zone)			
NAm analogs				"Analog"?
Global analogs				"Analog"?
Bayesian				"Analog"?
Physics	SS(global sample)			
Statistics	SS(global sample)			
Pattern recog.	SS(global sample)			
Q_0	SS(global sample)		Inconsistent	

USGS

Figure A4-2

Methods Overview (cont'd.)

- Physics-based methods: ultimate goal.
- Small-sample problem: solution unclear.
  - SS(source zone): few moderate-large eqs; wait?
  - SS(global sample): few accepted Mmax's; wait?
- Inconsistent with paleoseismic results?
  - Paleoseismic M estimates are uncertain, but not that much.
  - Slow work, funding-limited, so accumulate slowly.
- What is "analog"? Perhaps solvable…lacks obvious barriers.

USGS

Figure A5–1

Mmax and the Maximum Catalog Magnitude

Martin Chapman

Department of Geosciences
Virginia Tech
Blacksburg, VA 24061

mcc@vt.edu

Mmax Workshop
Golden, Colorado
September 8-9, 2008

Figure A5–2

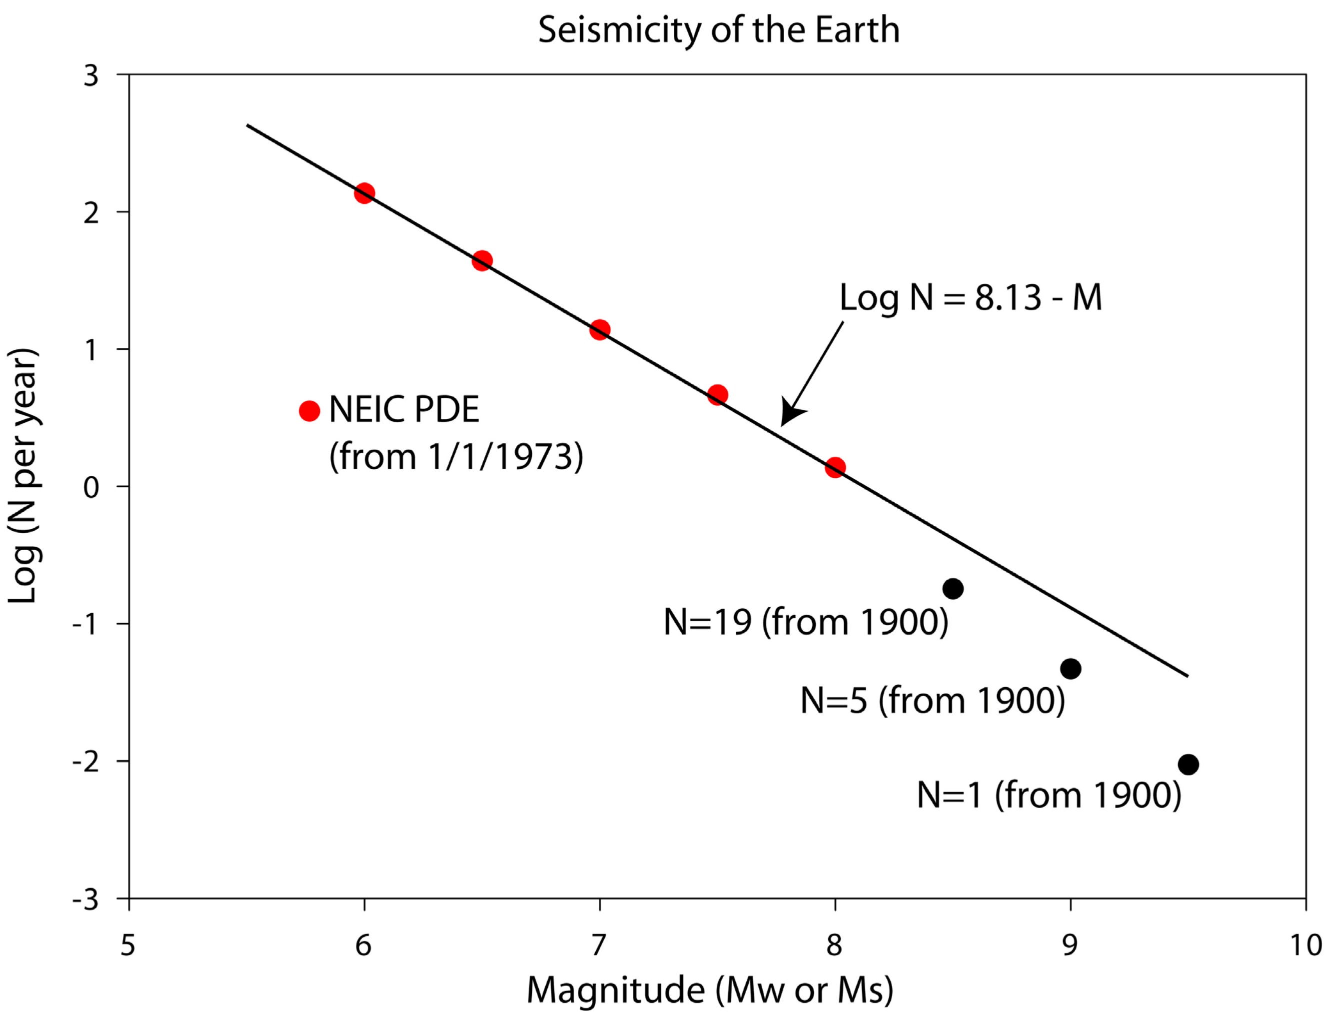

sources:

NEIC PDE catalog
M>6.0, since 1973

NEIC list of
significant earthquakes,
since 1900

Figure A5–3

Is the absence of events with M greater than 9.5 in the global catalog significant?

(i.e., is the maximum catalog magnitude Mc = 9.5 a "useful" estimate for the global Mmax, in the context of our standard PSHA model?)

Assumptions: 1) Poisson Process
2) Log N = 8.13 - M, (data for M > 6, since 1973)

Mean rate of M>9.5 = $10^{(8.13-9.5)}$ = 0.0427 yr^{-1} (return period 23.4 years)

Probability of zero M>9.5 earthquakes in 107 years

$$P(N=0 \mid M>9.5, T=107) = \exp(-0.0427 \times 107) = 0.01$$

Probability of zero M>9.5 events in 50 years

$$\exp(-0.0427 \times 50) = 0.11$$

Conclusion: These probabilities are small.
Mc = 9.5 is a useful estimate for the Global Mmax.

Figure A5–4

Can we infer useful information about Mmax from Mc
in sub-regions of the Earth?

It depends on N, and the length of time over which the catalog for a given sub-region can be considered complete.

mblg > 0.0, 1977 - 2005

Figure A5–5

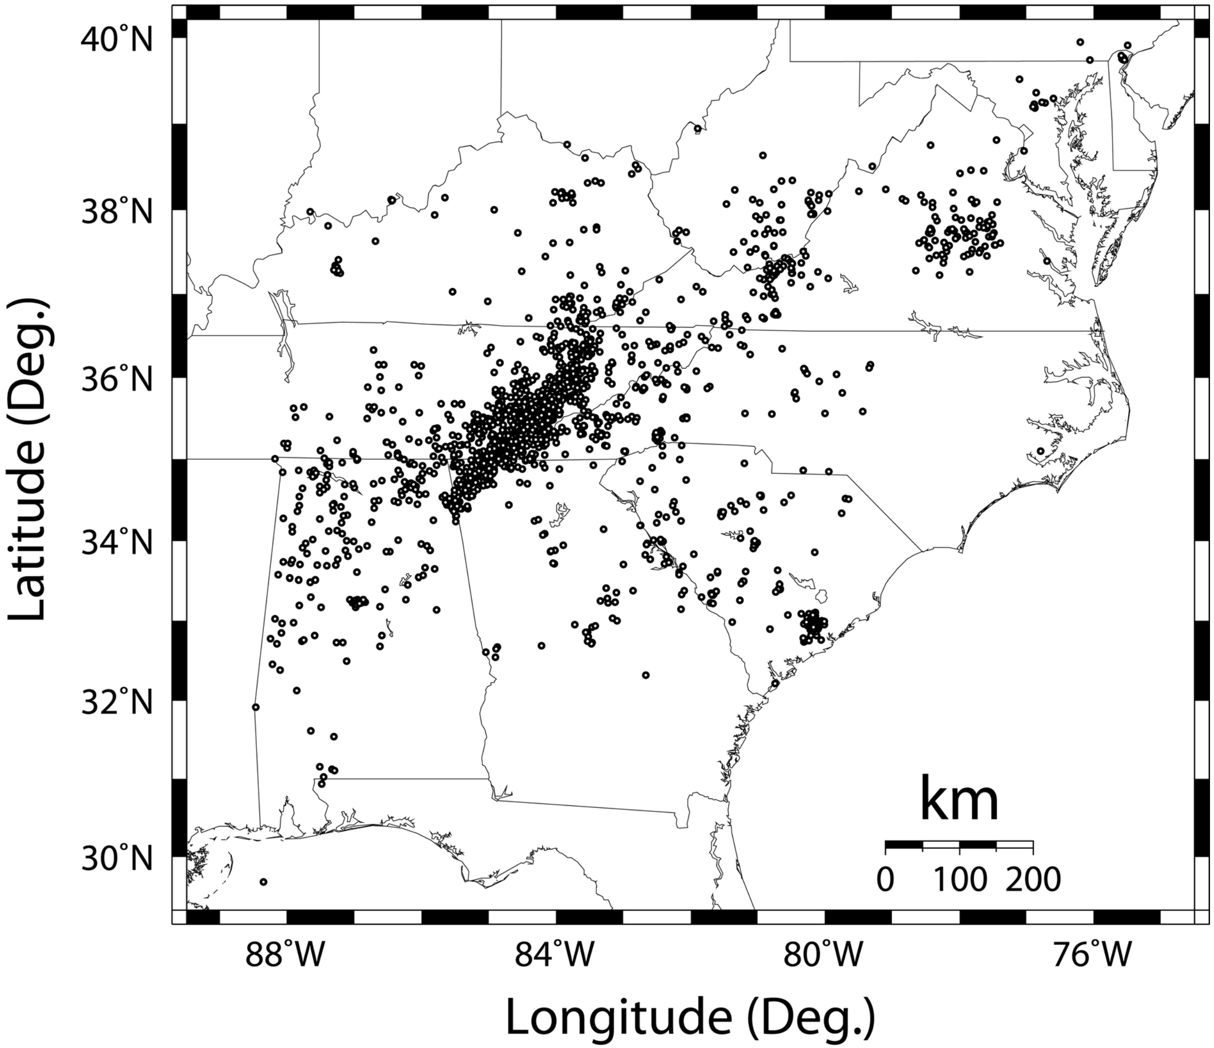

from SEUSSN Bulletin 40 (2005)

Figure A5–6

Number of Earthquakes by Decade

Eastern Tennessee Seismic Zone

Magnitude m_{blg}

DATE	2.0-2.4	2.5-2.9	3.0-3.4	3.5-3.9	4.0-4.5	4.5+
1994-90	35	19	5	0	0	0
1989-80	59	27	15	4	2	0
1979-70	1	3	2	4	0	1
1969-60	0	3	1	2	1	0
1959-50	0	0	6	4	3	0
1949-40	0	3	2	2	1	0
1939-30	0	0	1	2	0	0
1929-20	0	0	0	2	0	0
1919-10	0	2	3	3	1	0
1909-00	0	0	1	2	1	0
1899-90	0	0	0	0	0	0
1889-80	1	0	1	0	0	0
1879-70	0	2	0	0	1	0
1869-60	0	1	0	0	0	0
1859-50	0	0	0	0	0	0
1849-40	0	0	0	1	0	0
--						
1777	0	0	1	0	0	0

source: SEUSSN catalog

Figure A5–7

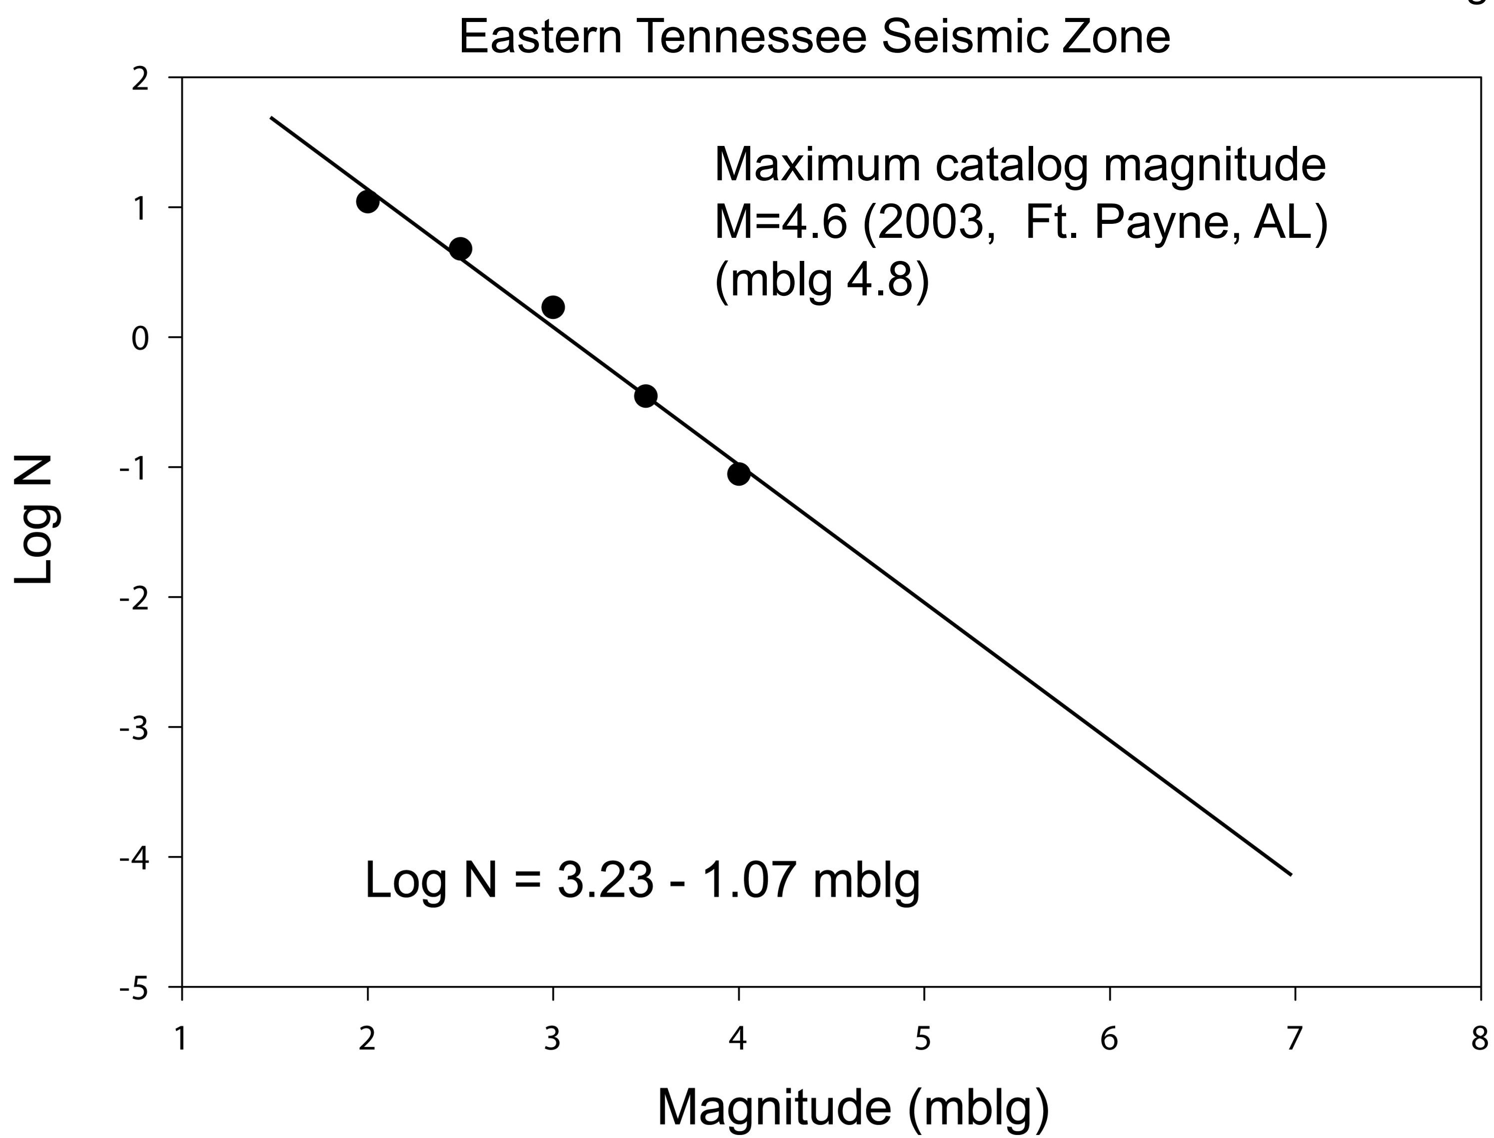

Eastern Tennessee Seismic Zone
Maximum catalog magnitude
M=4.6 (2003, Ft. Payne, AL)
(mblg 4.8)
Log N = 3.23 - 1.07 mblg
Log N
Magnitude (mblg)
2
1
0
-1
-2
-3
-4
-5
1
2
3
4
5
6
7
8

Figure A5–8

Is mblg=5 a useful estimate for Mmax in the Eastern Tennessee Seismic Zone?

Assume Log N = 3.23 - 1.07mblg:
mean rate for mblg>5.0, N(5.0) = 0.00758 yr^{-1}
return period = 132 years

Assume that the catalog is complete to 1870 (T=138 years,)

P (N=0|mblg>5.0,T=138) = exp(-0.00758 x 138) = 0.35

Complete to 1840: P = 0.28

How long would the catalog have to be complete, for the absence of mblg>5.0 to be significant: i.e., P=0.1 or less (as in the case of M 9.5 for the Earth)?

T = -ln(0.1) / 0.00758 = 303 years

Conclusion: For the maximum catalog magnitude (mblg 5.0) to be a "useful" estimate of Mmax, the catalog would have to be complete for mblg > 5.0 to 1705. The earliest recorded shock in the region was in 1777.

Figure A5–9

Summary

The maximum catalog magnitude Mc provides information on Mmax. In fact, given no other information, it provides the best estimate (under the assumptions made here).

Unfortunately, unless the product of N(Mc), and T, the time period of catalog completeness, is such that $e^{[-N(Mc) \times T]}$ is a significantly small number, the estimate will not be reliable.

The “EPRI” Bayesian M_{max} Approach for Stable Continental Regions (SCR) (Johnston et al. 1994)

Robert Youngs
AMEC Geomatrix

USGS Workshop on Maximum Magnitude Estimation
September 8, 2008

Statistical Estimation of m^u (M_{max})

- Assumption - earthquake size distribution in a source zone conforms to a truncated exponential distribution between m_0 and m^u
- Likelihood of m^u given observation of N earthquakes between m_0 and maximum observed, $m_{max\text{-}obs}$

$$L[m^u] = \begin{cases} 0 & \text{for } m^u < m_{\max-obs} \\ \left[1-\exp\left\{-b\ln(10)(m^u - m_0)\right\}\right]^{-N} & \text{for } m^u \geq m_{\max-obs} \end{cases}$$

Figure A6–3

Plots of Likelihood Function for $m_{max\text{-}obs} = 6$

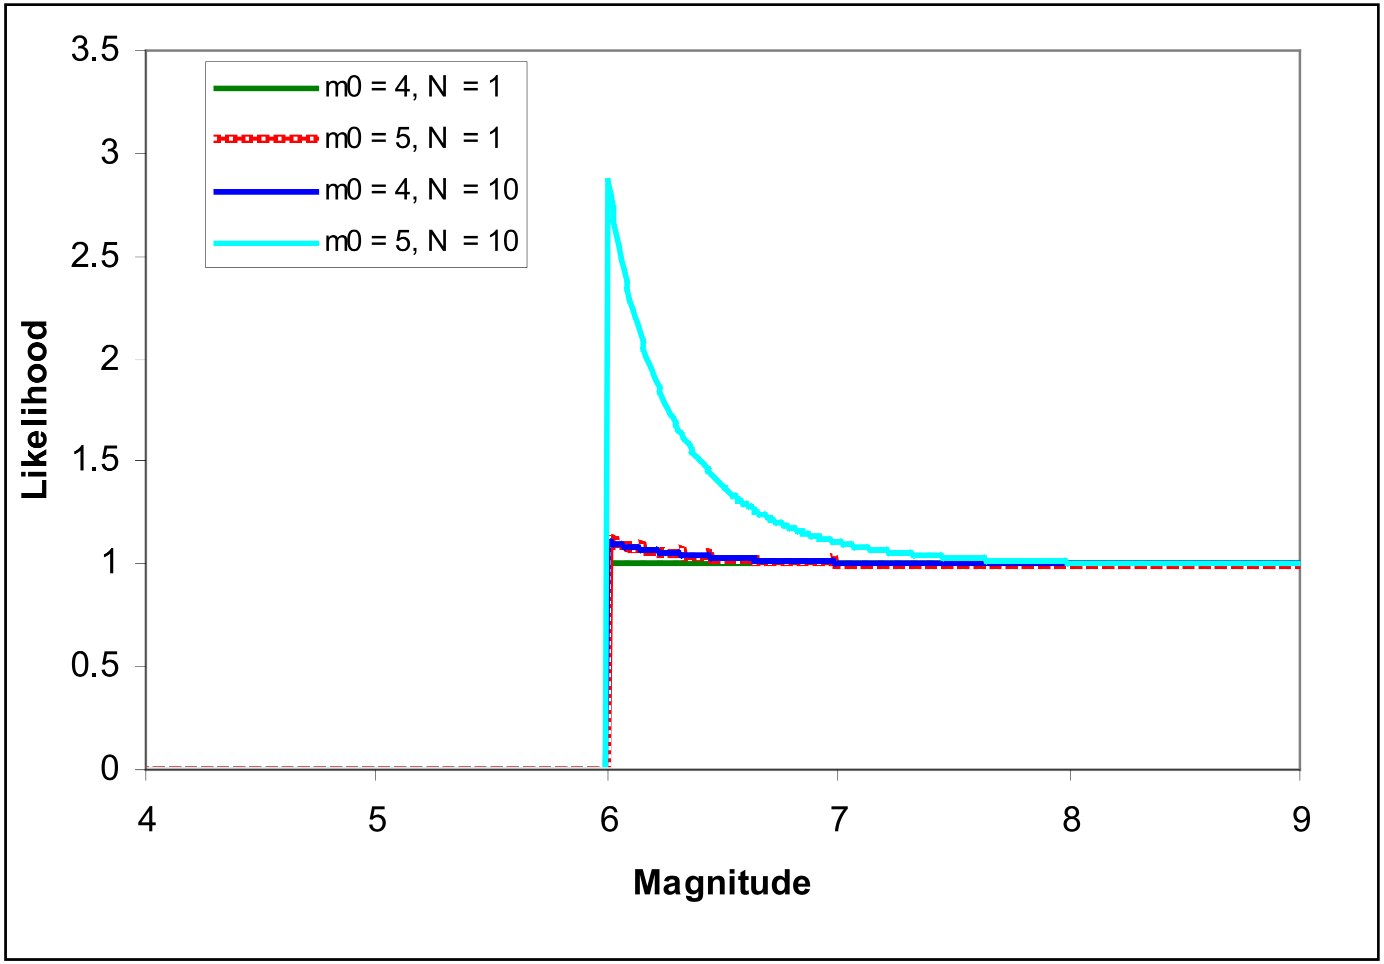

Figure A6–4

Results of Applying Likelihood Function

- $m_{max\text{-}obs}$ is the most likely value of m^u
- Relative likelihood of values larger than $m_{max\text{-}obs}$ is a strong function of sample size *and* the difference $m_{max\text{-}obs} - m_0$
- Likelihood function integrates to infinity and cannot be used to define a distribution for m^u
- Hence the need to combine likelihood with a prior to produce a posterior distribution

Approach for EPRI (1994) SCR Priors

- Divided SCR into domains based on:
  - Crustal type (extended or non-extended)
  - Geologic age
  - Stress regime
  - Stress angle with structure

- Assessed $m_{max\text{-}obs}$ for domains from catalog of SCR earthquakes

Figure A6–6

Bias Adjustment (1 of 2)

- "bias correction" from $m_{max\text{-}obs}$ to m^u based on distribution for $m_{max\text{-}obs}$ given m^u
- For a given value of m^u and N estimate the median value of $m_{max\text{-}obs}$, $\hat{m}_{\max-obs}$

$$F[m_{\max-obs}] = \left[\frac{1-\exp(-b\ln(10)(m_{\max-obs} - m_0))}{1-\exp(-b\ln(10)(m^u - m_0))} \right]^N \quad \text{for } m_0 \le m_{\max-obs} \le m^u$$

- Use $m^u - \hat{m}_{\max-obs}$ to adjust from $m_{max\text{-}obs}$ to m^u

Bias Adjustment (2 of 2)

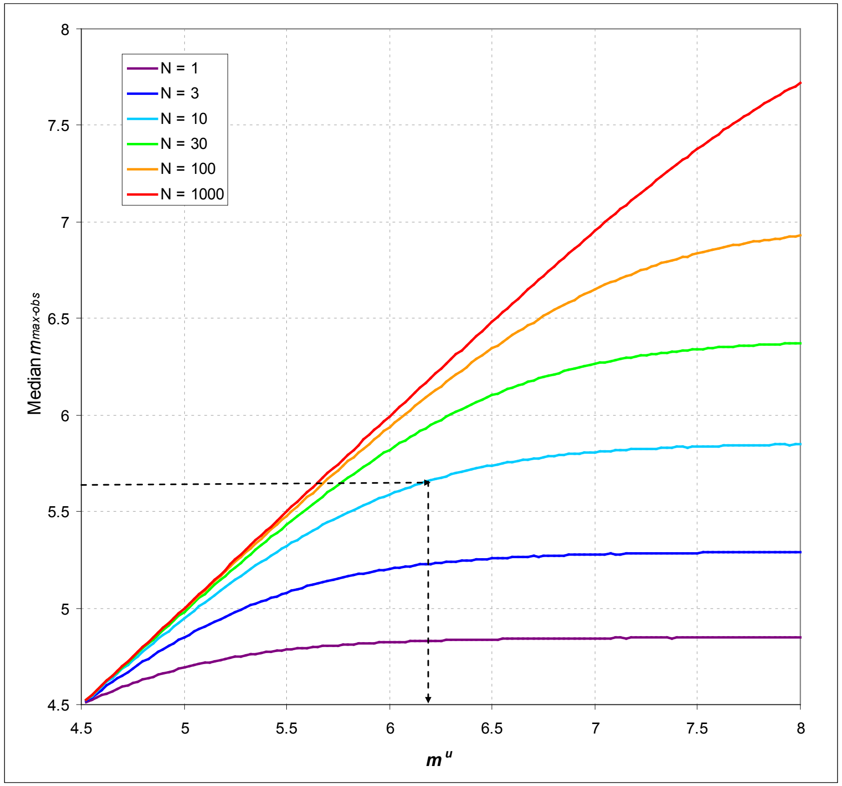

Example:

$m_{max\text{-}obs} = 5.7$

$N(m \leq 4.5) = 10$

$m^u = 6.3$ produces

$\hat{m}_{\max-obs} = 5.7$

Figure A6–8

Domain “Pooling”

- Obtaining usable estimates of bias adjustment necessitated pooling “like” domains (trading space for time)
- “Super Domains” created by combining domains with the same characteristics
  - Extended crust - 73 domains become 55 super domains, average N = 30
  - Non-extended crust – 89 domains become 15 super domains, average N = 120

EPRI (1994) Category Priors

- Compute statistics of $m_{max\text{-}obs}$ for extended and non extended crust

$$\text{for extended crust} \quad \overline{m}_{\max-obs} = 6.04 \quad \sigma_{m_{\max-obs}} = 0.84$$
$$\text{for non-extended crust} \quad \overline{m}_{\max-obs} = 6.2 \quad \sigma_{m_{\max-obs}} = 0.5$$

- Use average sample size to adjust to m^u

$$\text{for extended crust} \quad \overline{m^u} = 6.4 \quad \sigma_{m_{\max-obs}} = 0.84$$
$$\text{for non-extended crust} \quad \overline{m^u} = 6.3 \quad \sigma_{m_{\max-obs}} = 0.5$$

EPRI (1994) Regression Prior

- Regress $m_{max\text{-}obs}$ against domain characterization variables
  - Default region is non-extended Cenozoic crust
  - "Dummy" variables indicating other crustal types, ages, stress conditions, and a continuous variable for ln(activity rate) indicate departure from default.
- Model has low r^2 of 0.29 – not very effective in explaining variations

Example Application Using Category Prior

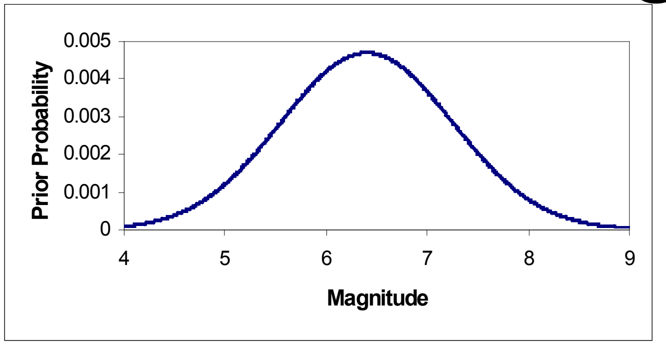

Extended crust

$\mu_{Mu} = 6.4$

$\sigma_{Mu} = 0.84$

5 events recorded

between **M** 4.5 and **M** 5

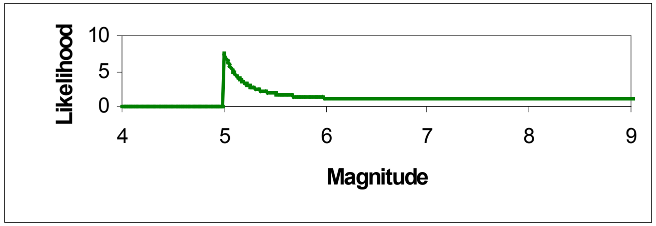

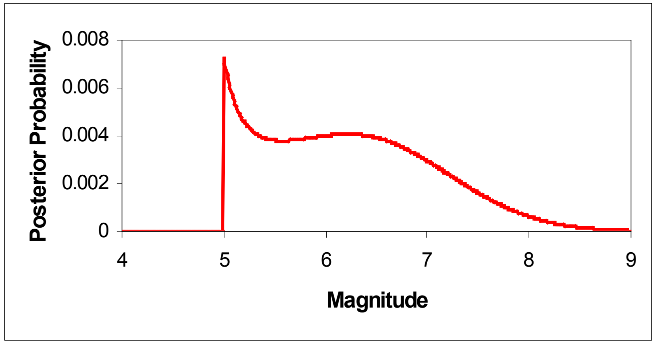

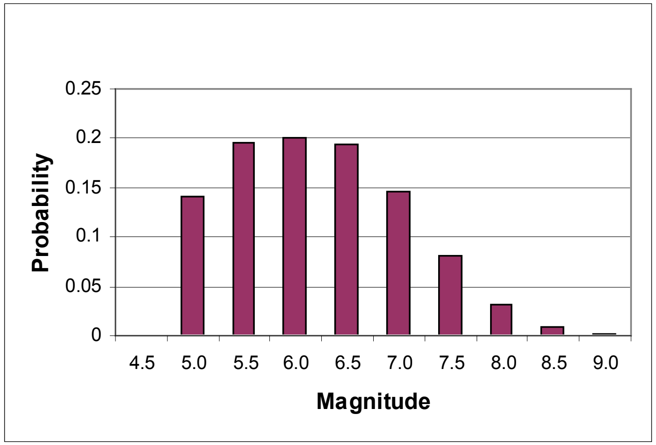

Figure A6–12

Summary

- Bayesian approach provides a means of using observed earthquakes to assess distribution for m^u
- Requires an assessment of a prior distribution for m^u
- Johnston et al. (1994) developed two types:
  - crustal type category: extended or non-extended
  - a regression model (low r^2 and high correlation between predictor variables)
- Bayesian approach is not limited to the Johnston et al. (1994) priors, any other prior may be used

Figure A7–1

Large SCR earthquakes occur preferentially in young extended terranes.

Application to CEUS & Adjacent Canada:

(New Madrid, Charleston set aside as special cases)

Extended:

- Mmax 8.3 in rifts
- Mmax 7.7 in passive margins (Atlantic and Gulf seaboards, Cambrian passive margin)

Non-extended:

- Mmax 6.4 in Paleozoic orogens (Appalachian and Ouachita Mountains)
- Mmax 6.8 in cratons (midcontinent platform, Canadian shield)

USGS

Figure A7–2

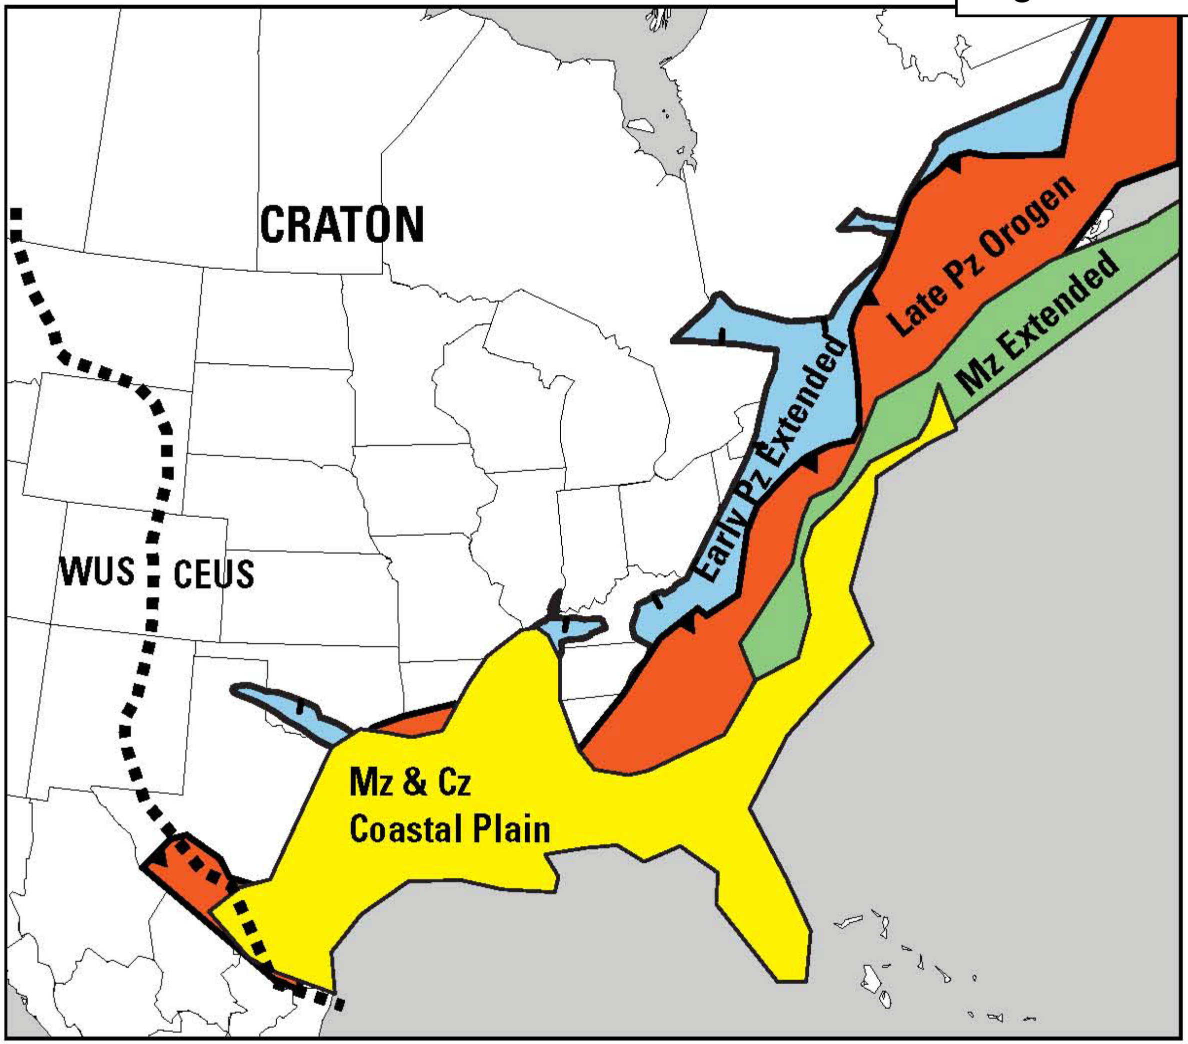

Main CEUSAC Tectonic Elements

USGS

Figure A7–3

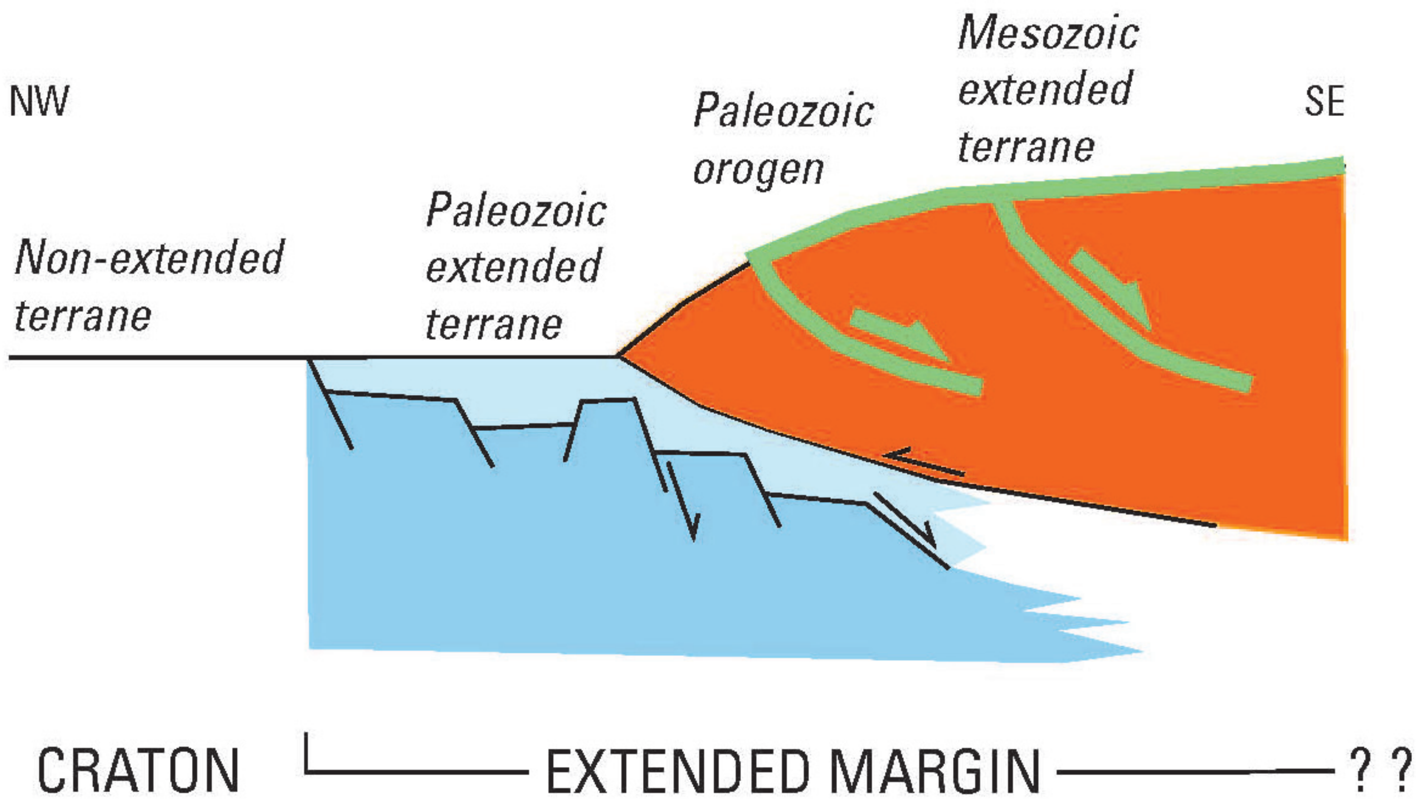
NW
SE
Mesozoic extended terrane
Paleozoic orogen
Paleozoic extended terrane
Non-extended terrane
CRATON
EXTENDED MARGIN
? ?

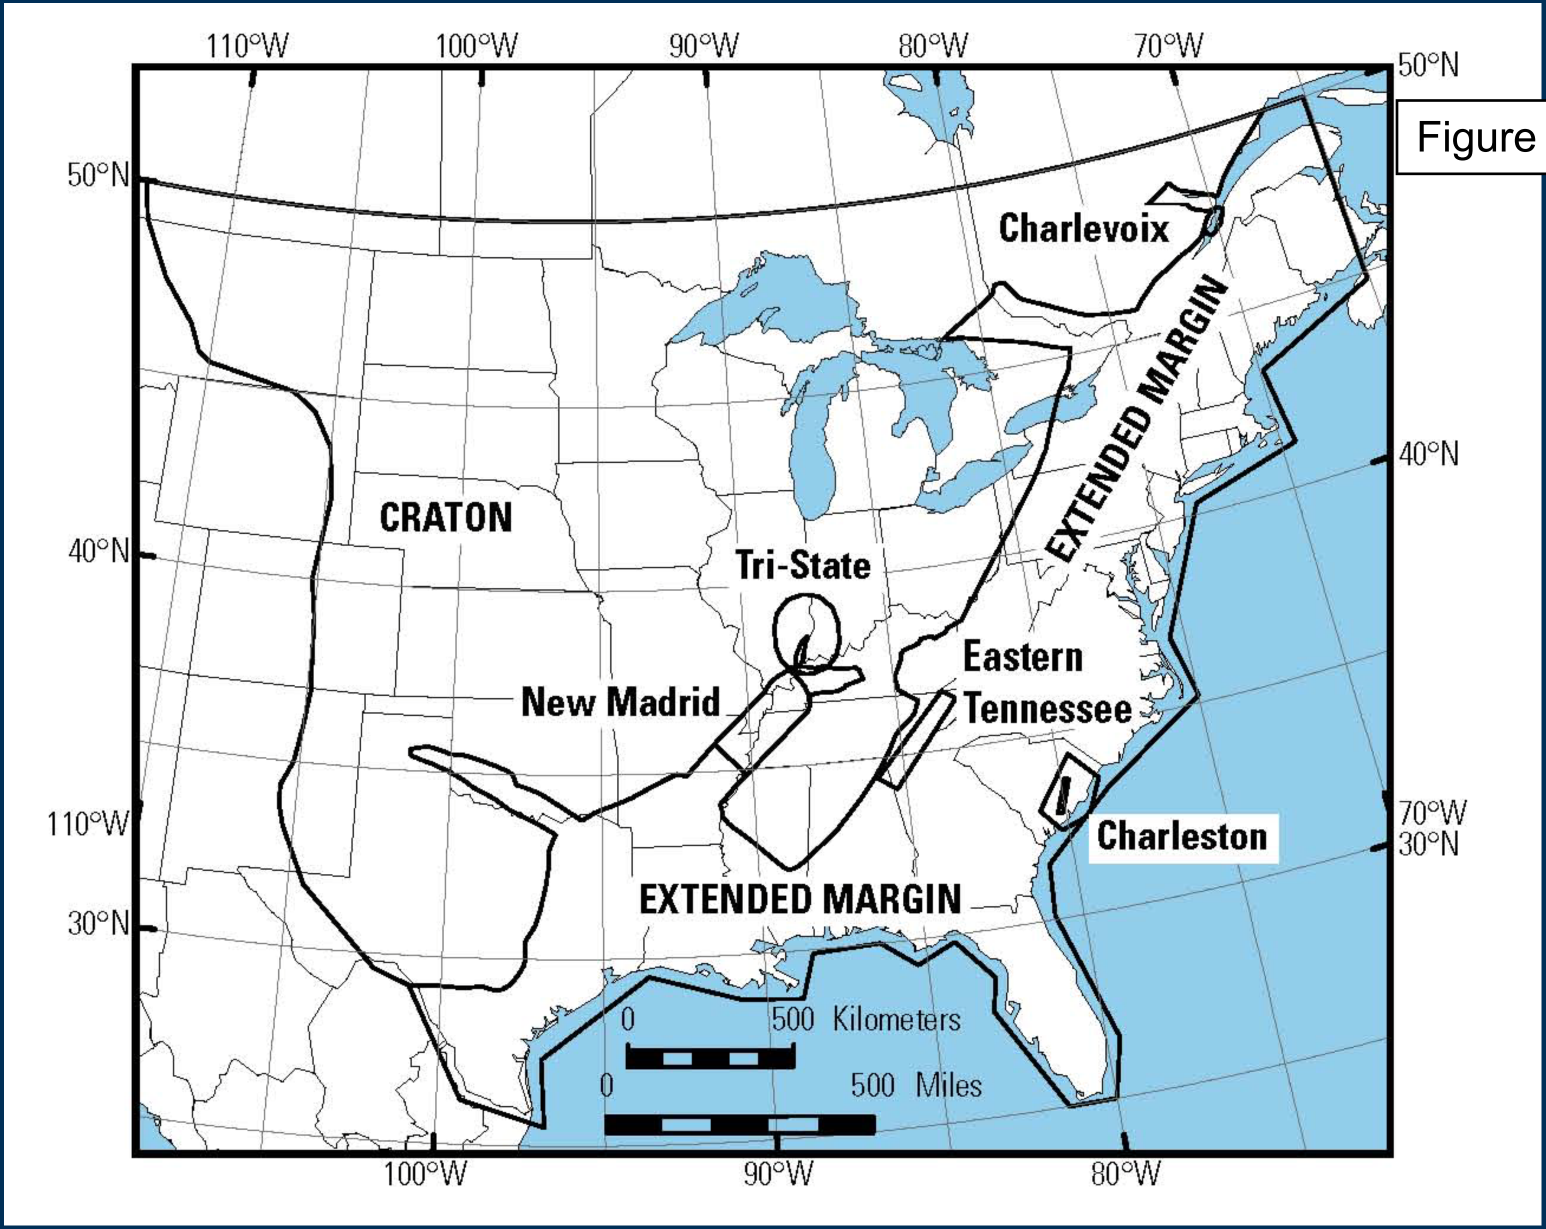

Figure A7–4
110°W
100°W
90°W
80°W
70°W
50°N
40°N
30°N
Charlevoix
EXTENDED MARGIN
CRATON
Tri-State
Eastern
Tennessee
New Madrid
Charleston
EXTENDED MARGIN
0 500 Kilometers
0 500 Miles

USGS

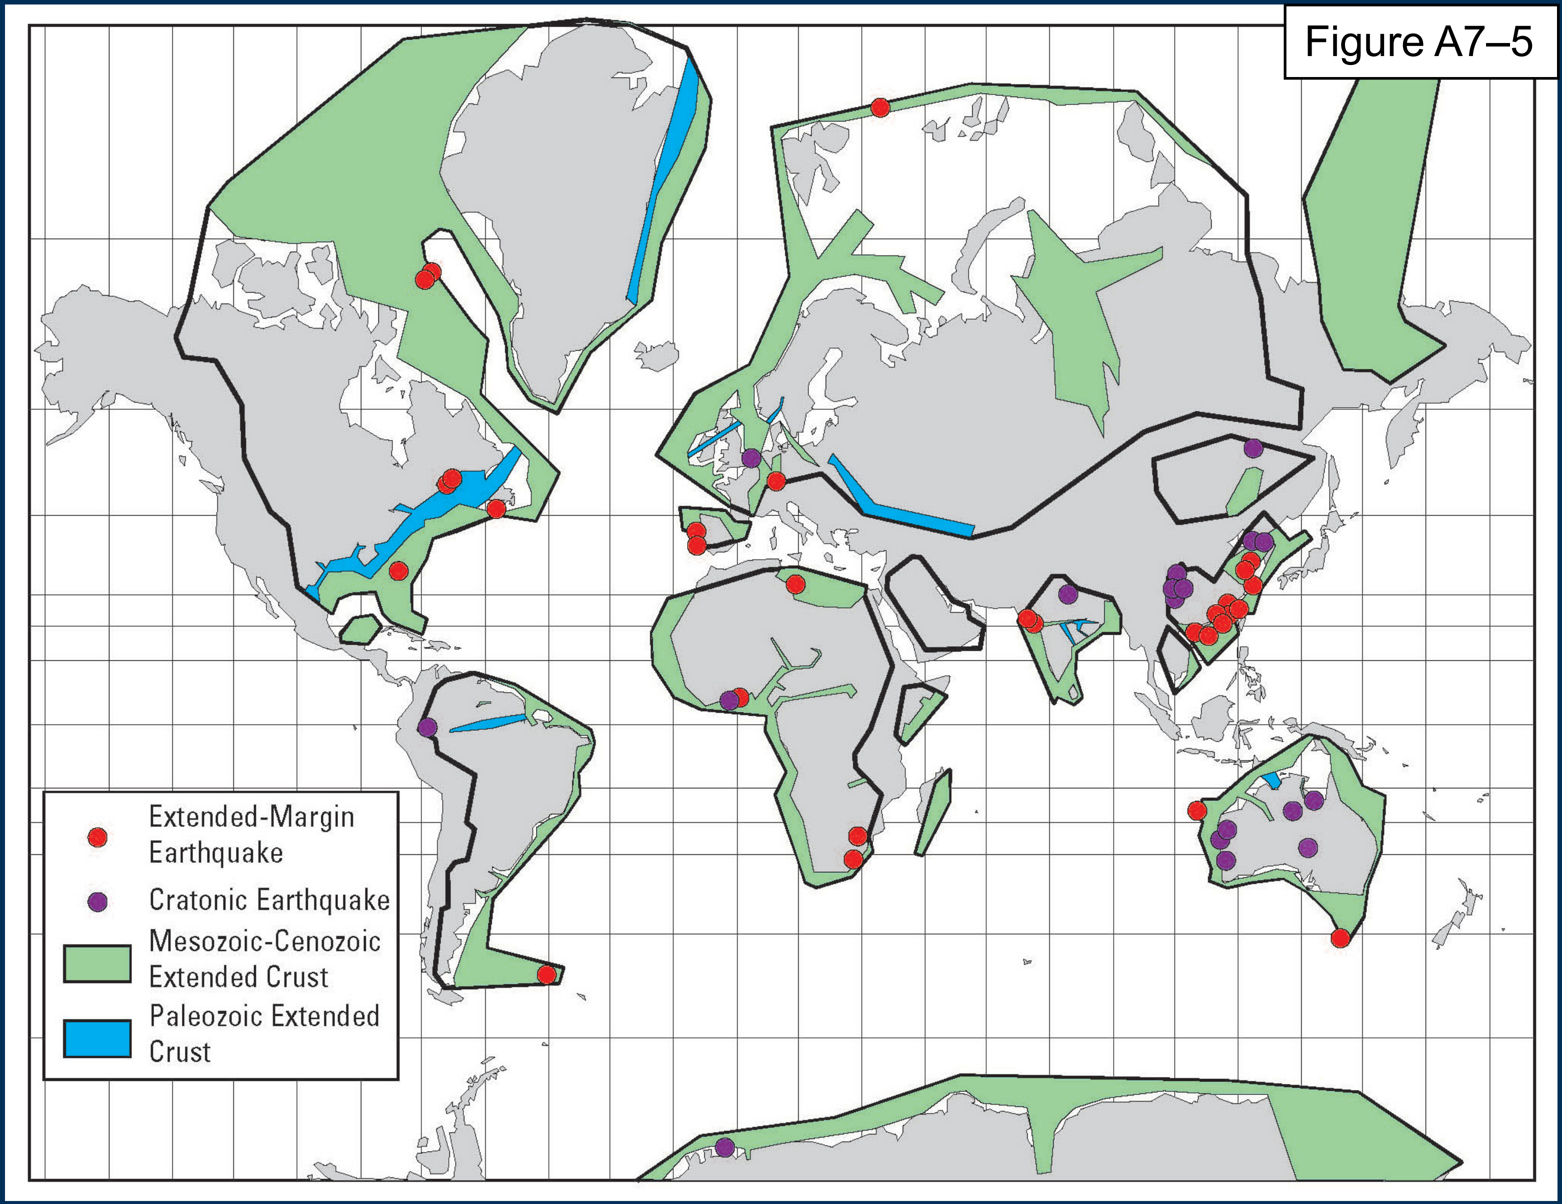
Figure A7–5
Extended-Margin Earthquake
Cratonic Earthquake
Mesozoic-Cenozoic Extended Crust
Paleozoic Extended Crust

USGS

Mw(max) for Tectonic Analogs of Central and Eastern U.S.

Figure A7–6

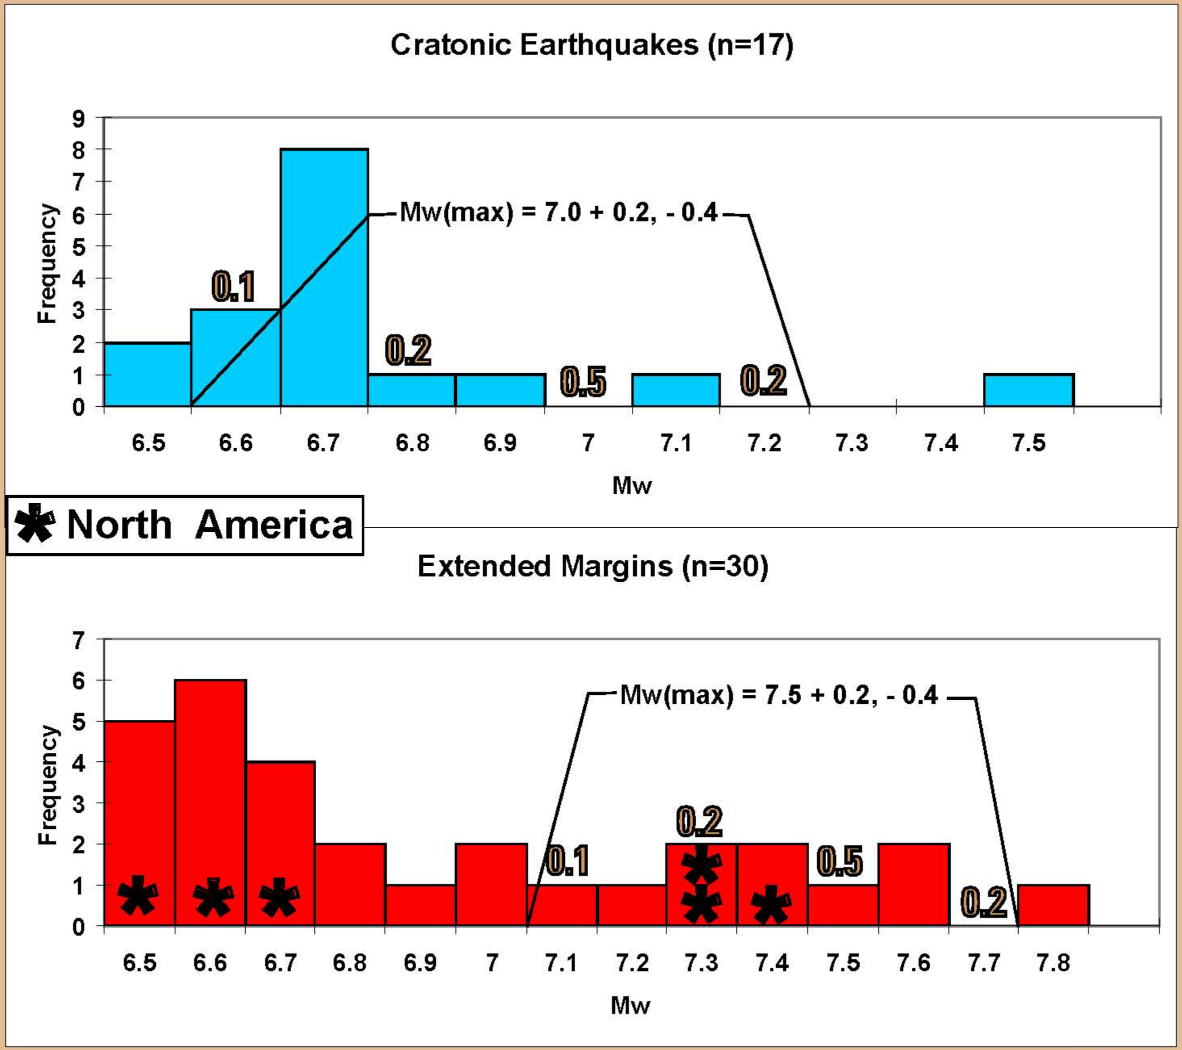

8/2/07

USGS

Figure A7–7

Pros, Cons, Uncertainties

- Pro:
  - Global catalog avoids small-sample problem
- Con:
  - What’s an “analog”?
- Uncertainties:
  - 1. Choice and justification of analogs
  - 2. M values of individual earthquakes
  - 3. Robustness of histograms

USGS

Figure A8–1

NRC Workshop on Mmax in the Central and Eastern U.S. EQ Hazards Workshop September 9-10, 2007

USGS, Golden, CO

U.S. Department of the Interior
U.S. Geological Survey

Figure A8–2

Walter D. Mooney USGS-Menlo Park

Mmax in Cratons

Outline:

1. *Introduction*
2. *Continental Evolution*
3. *Earthquakes and Tectonics*
4. *Lithospheric Structure*
5. *Lith-EQ Correlation (Mmax)*
6. *Conclusion: Mmax = 6.8 (?)*

USGS

North American Geology

Figure A8–3

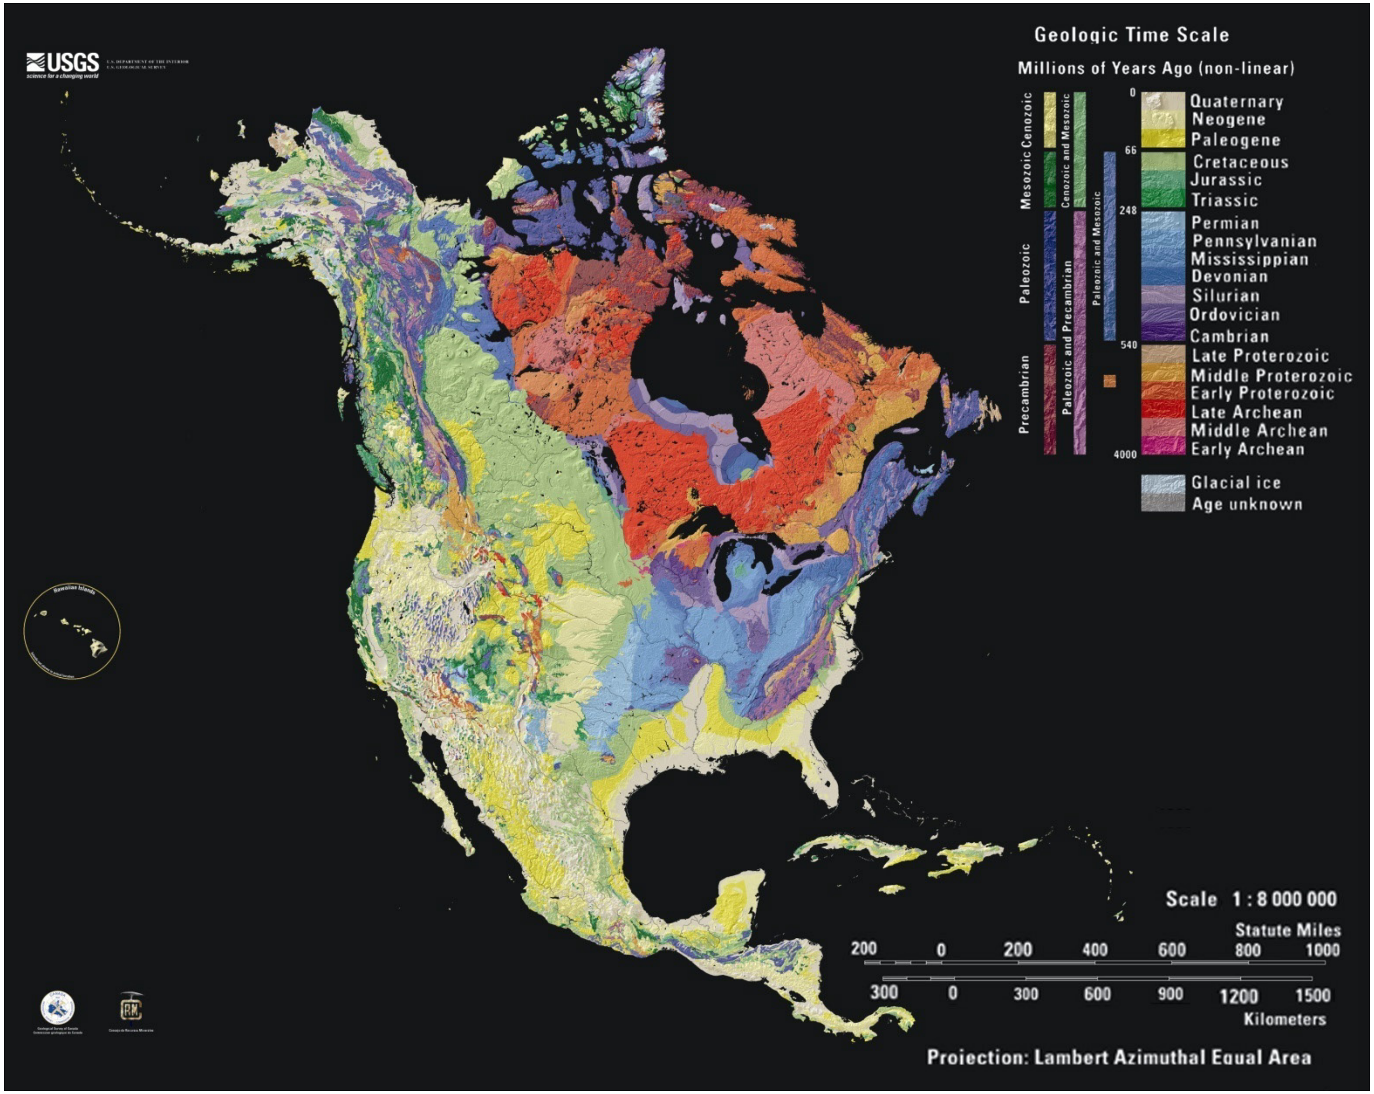

USGS

3.0 Ga (Archean): Formation of Cratons

Figure A8–4

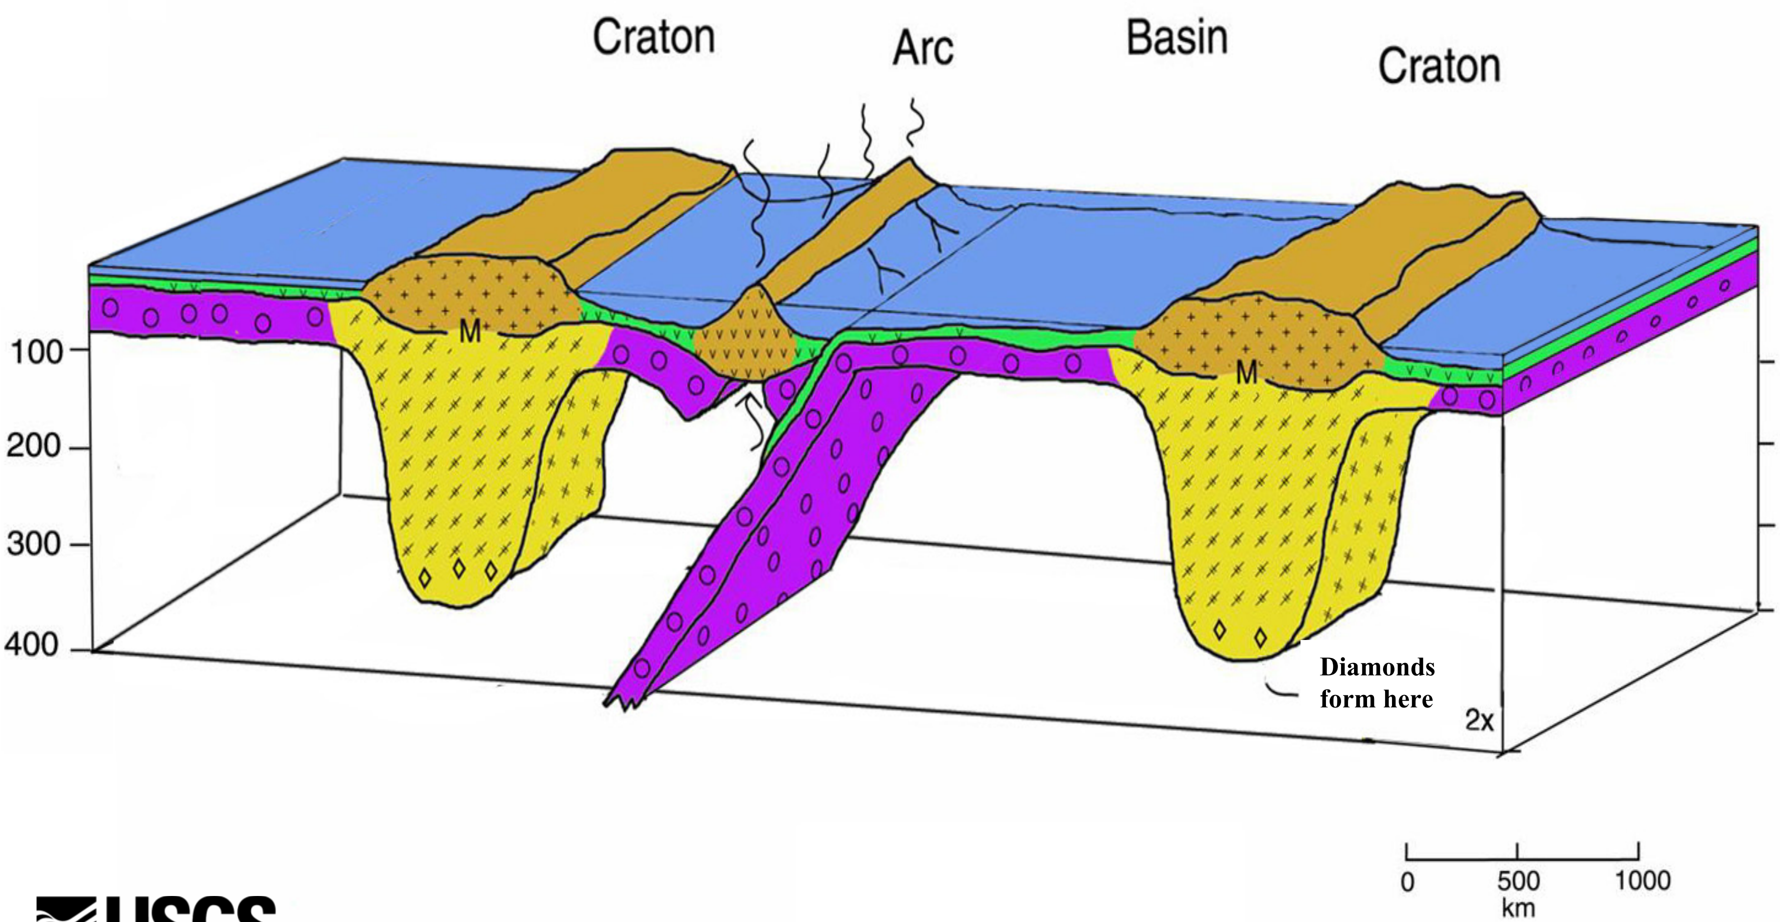

USGS

Figure A8–5

2.0 Ga (Early Proterozoic): Suturing of Cratonic Blocks

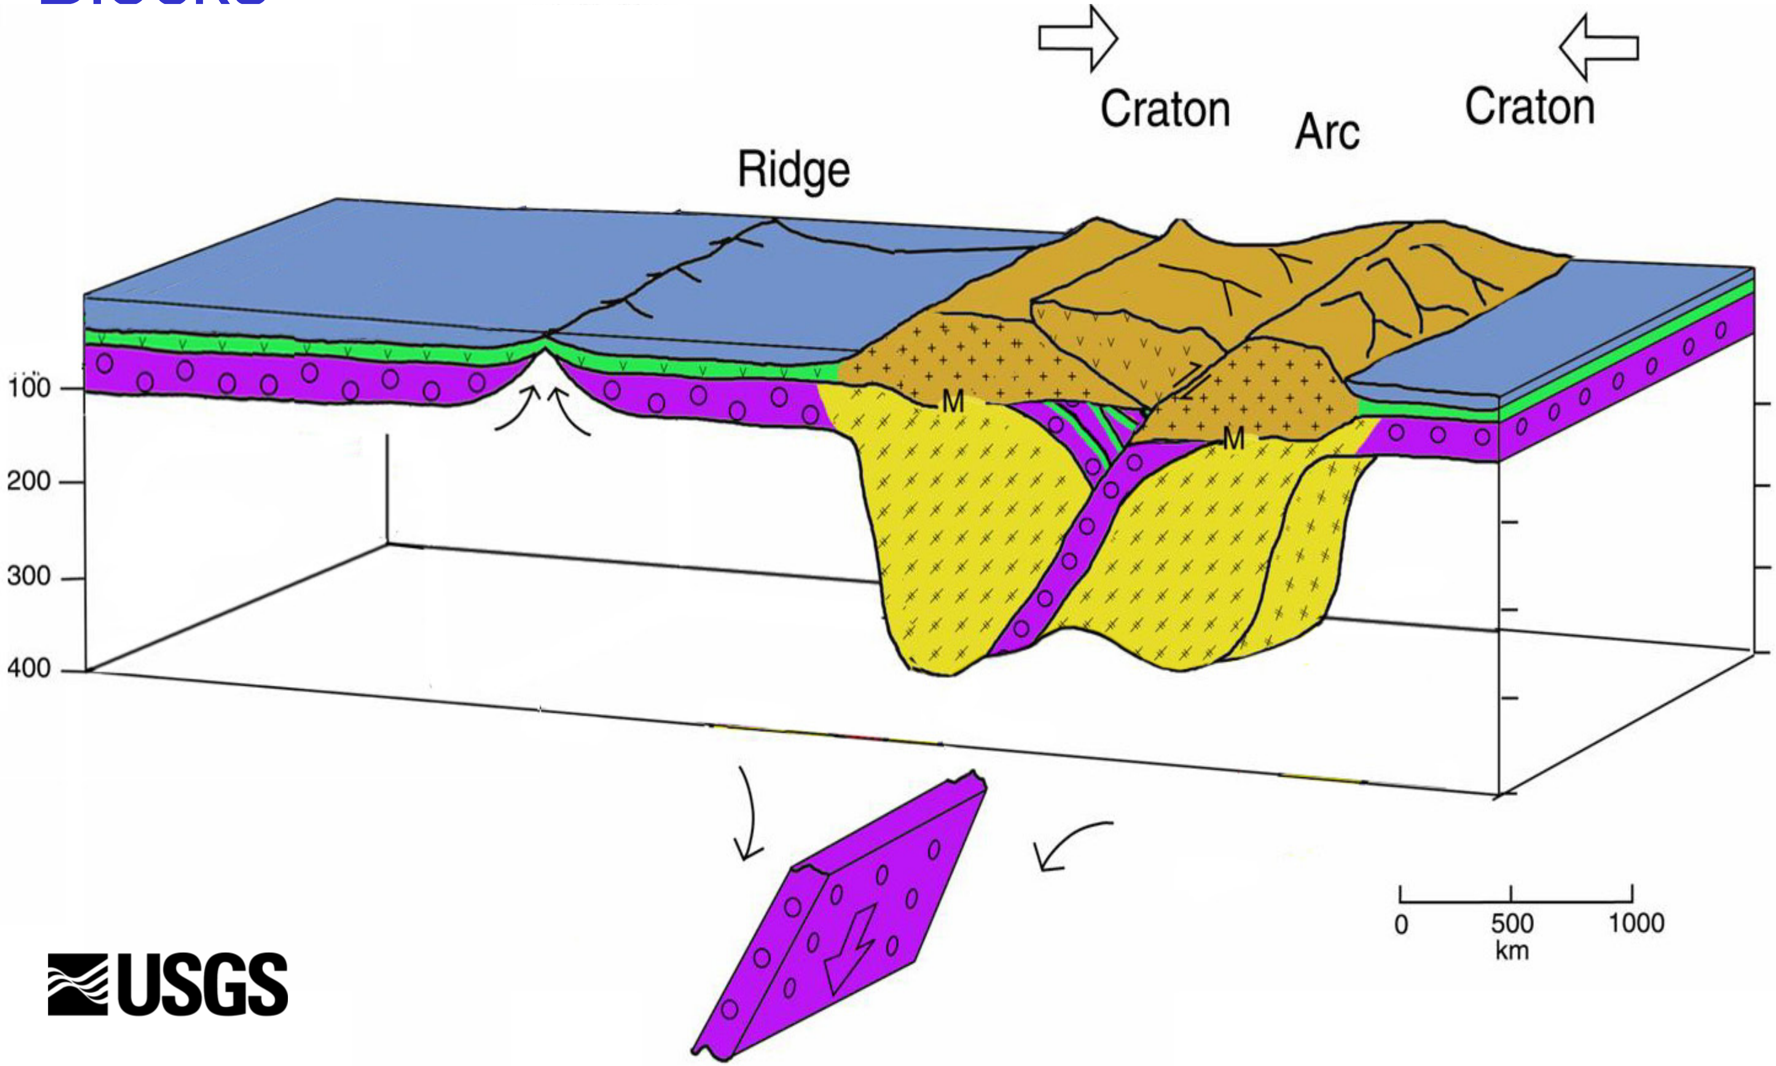

USGS

Figure A8–6

1.0 Ga (Late Proterozoic): Mantle Metasomatism at Lithospheric Suture, Intraplate Zone of Weakness

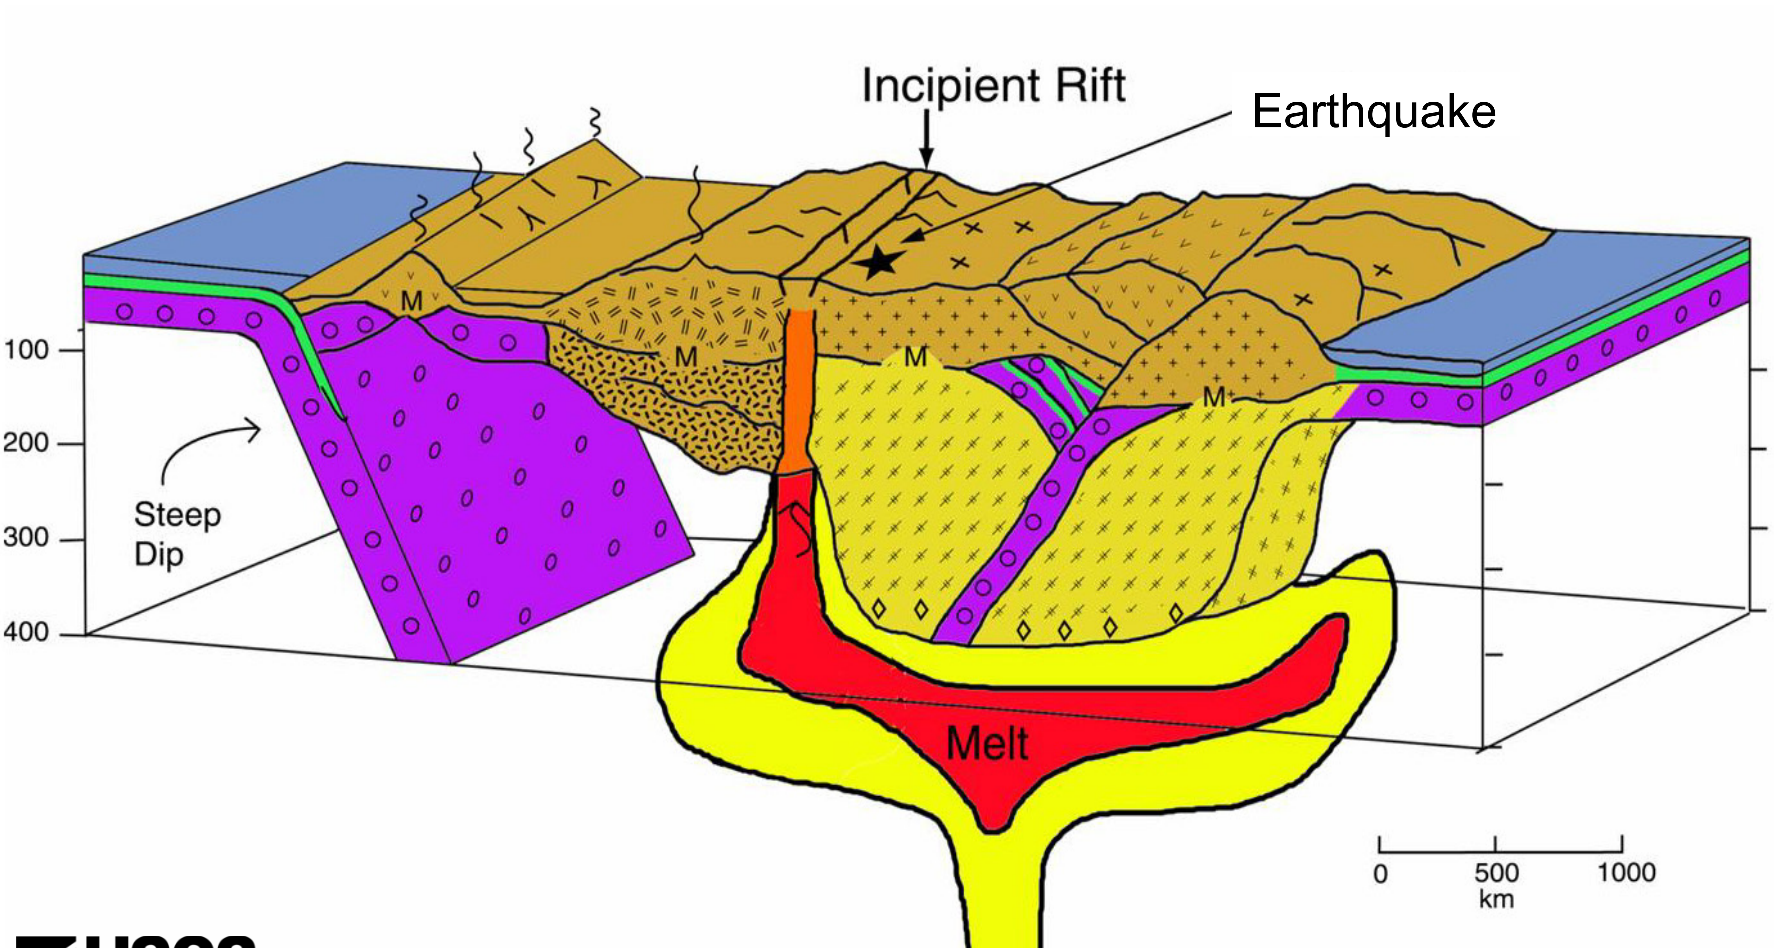

USGS

- Schulte, S. M., and Mooney, W. D., 2005, *An updated global earthquake catalogue for stable continental regions: reassessing the correlation with ancient rifts*: Geophysical Journal International, v. 161, p. 707-721.

Figure A8–8

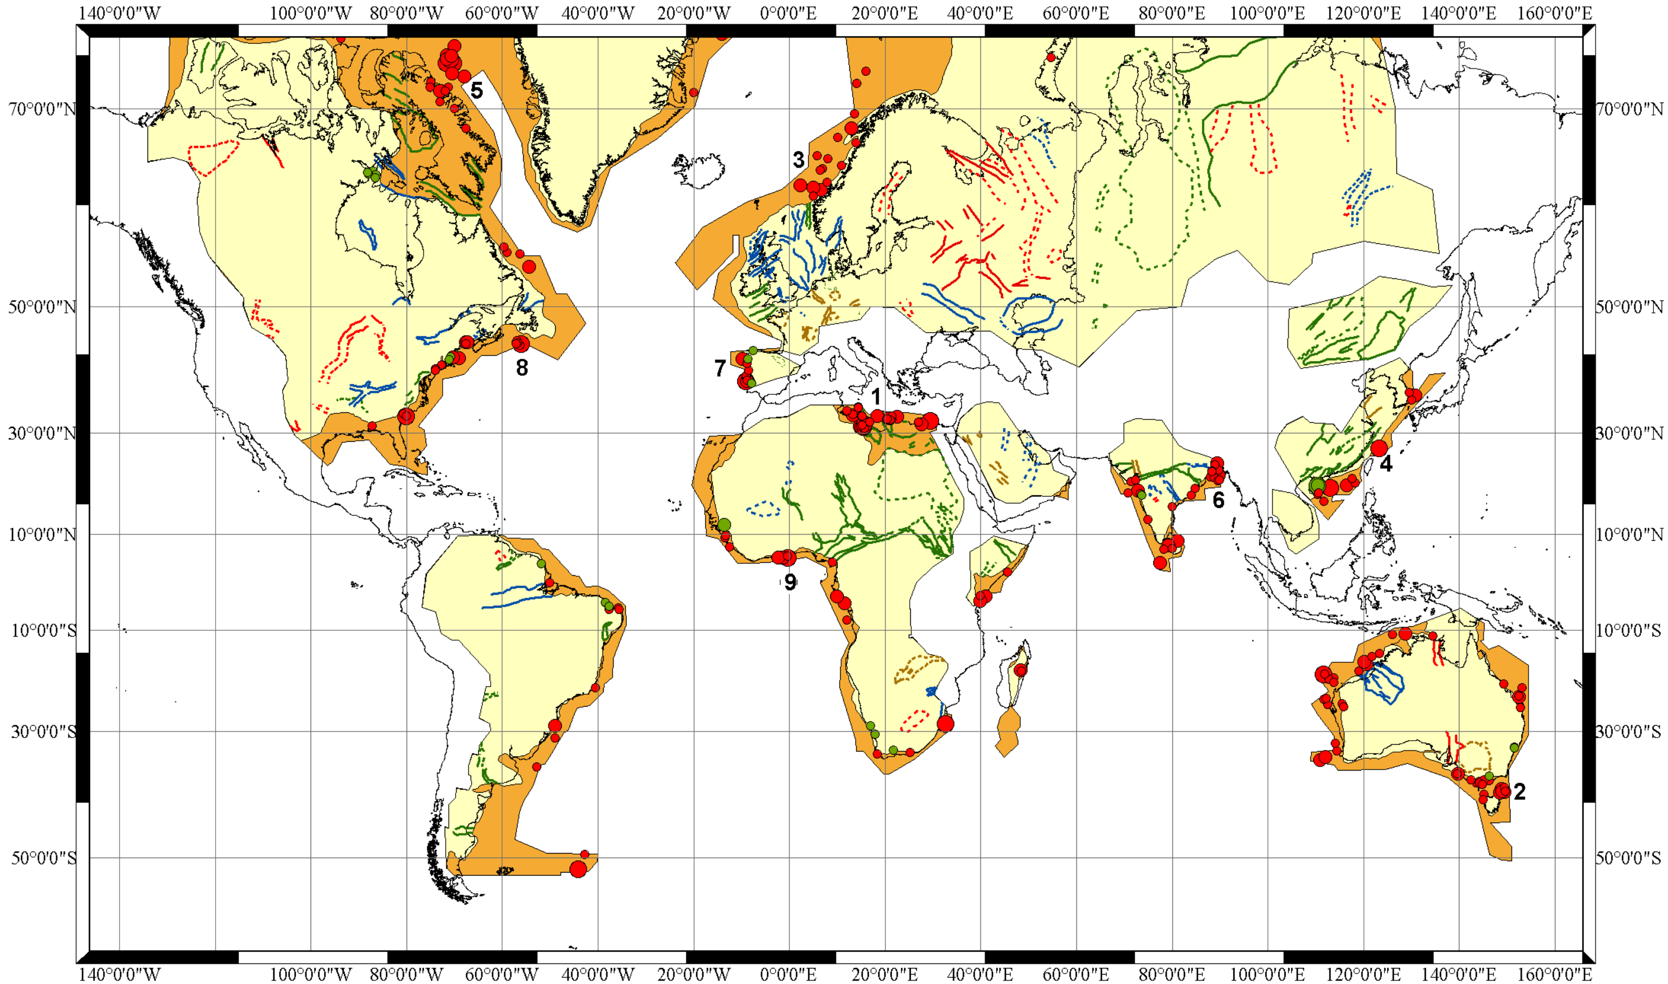

Figure 8a

Extended margin earthquakes. Red circles denote earthquakes that are associated with extended margins, green circles denote events that are possibly associated with extended margins. Regions of concentrated seismicity are numbered.

USGS

Figure A8–9

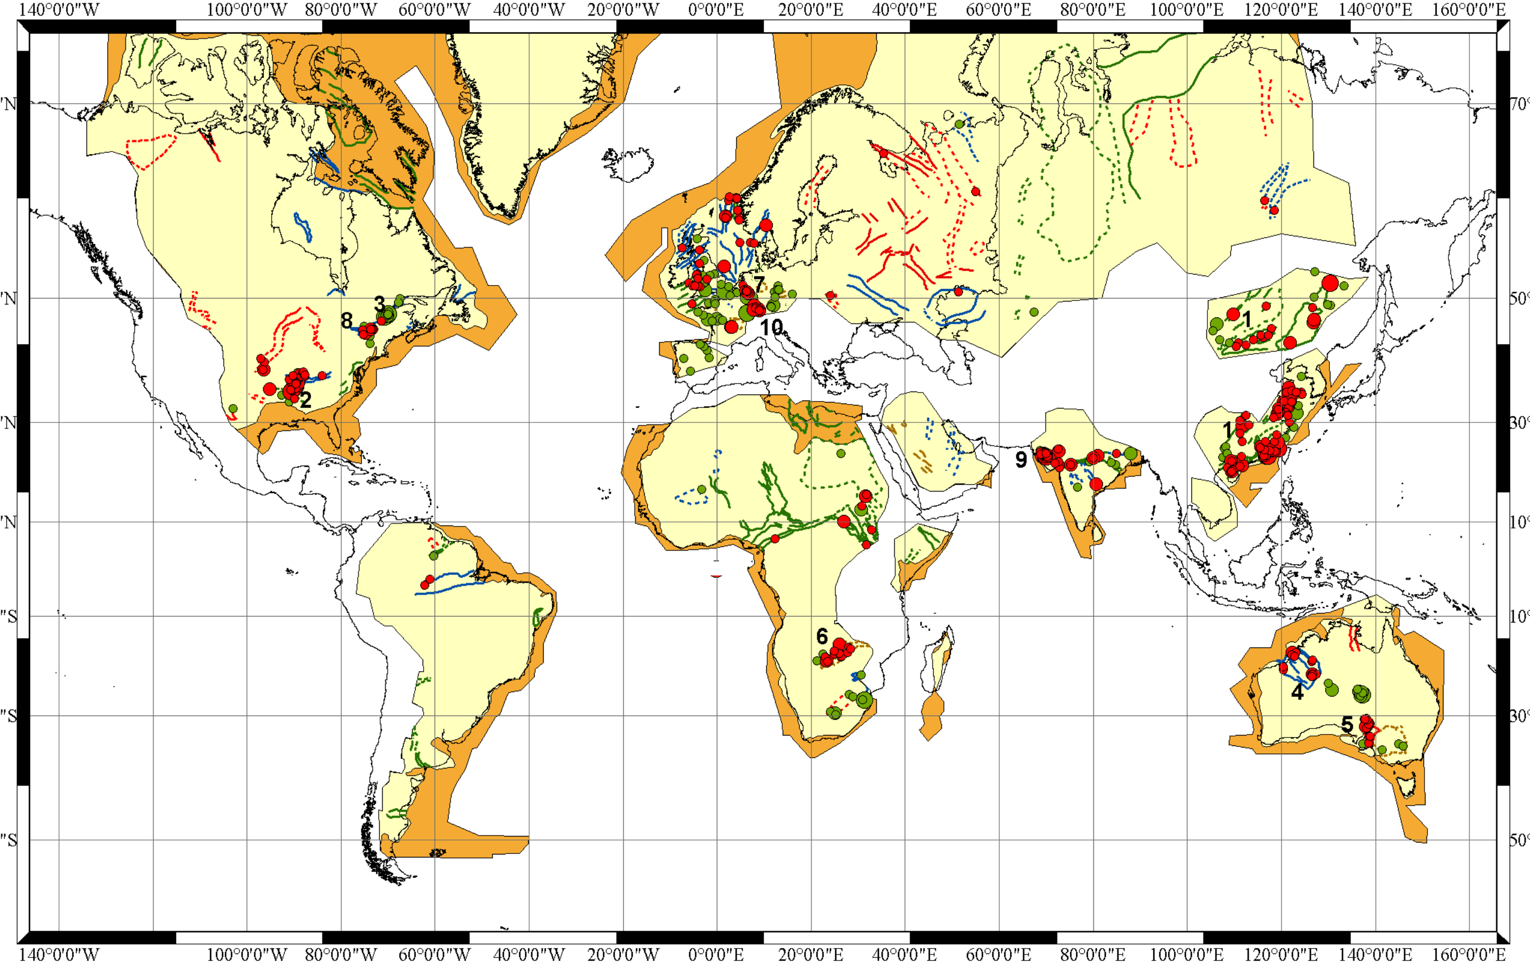

Figure 8b

Interior rift earthquakes. Red circles identify earthquakes that are associated with interior rifts, green circles identify events that are possibly associated with interior rifts. Regions of concentrated seismicity are numbered.

USGS

Figure A8–10

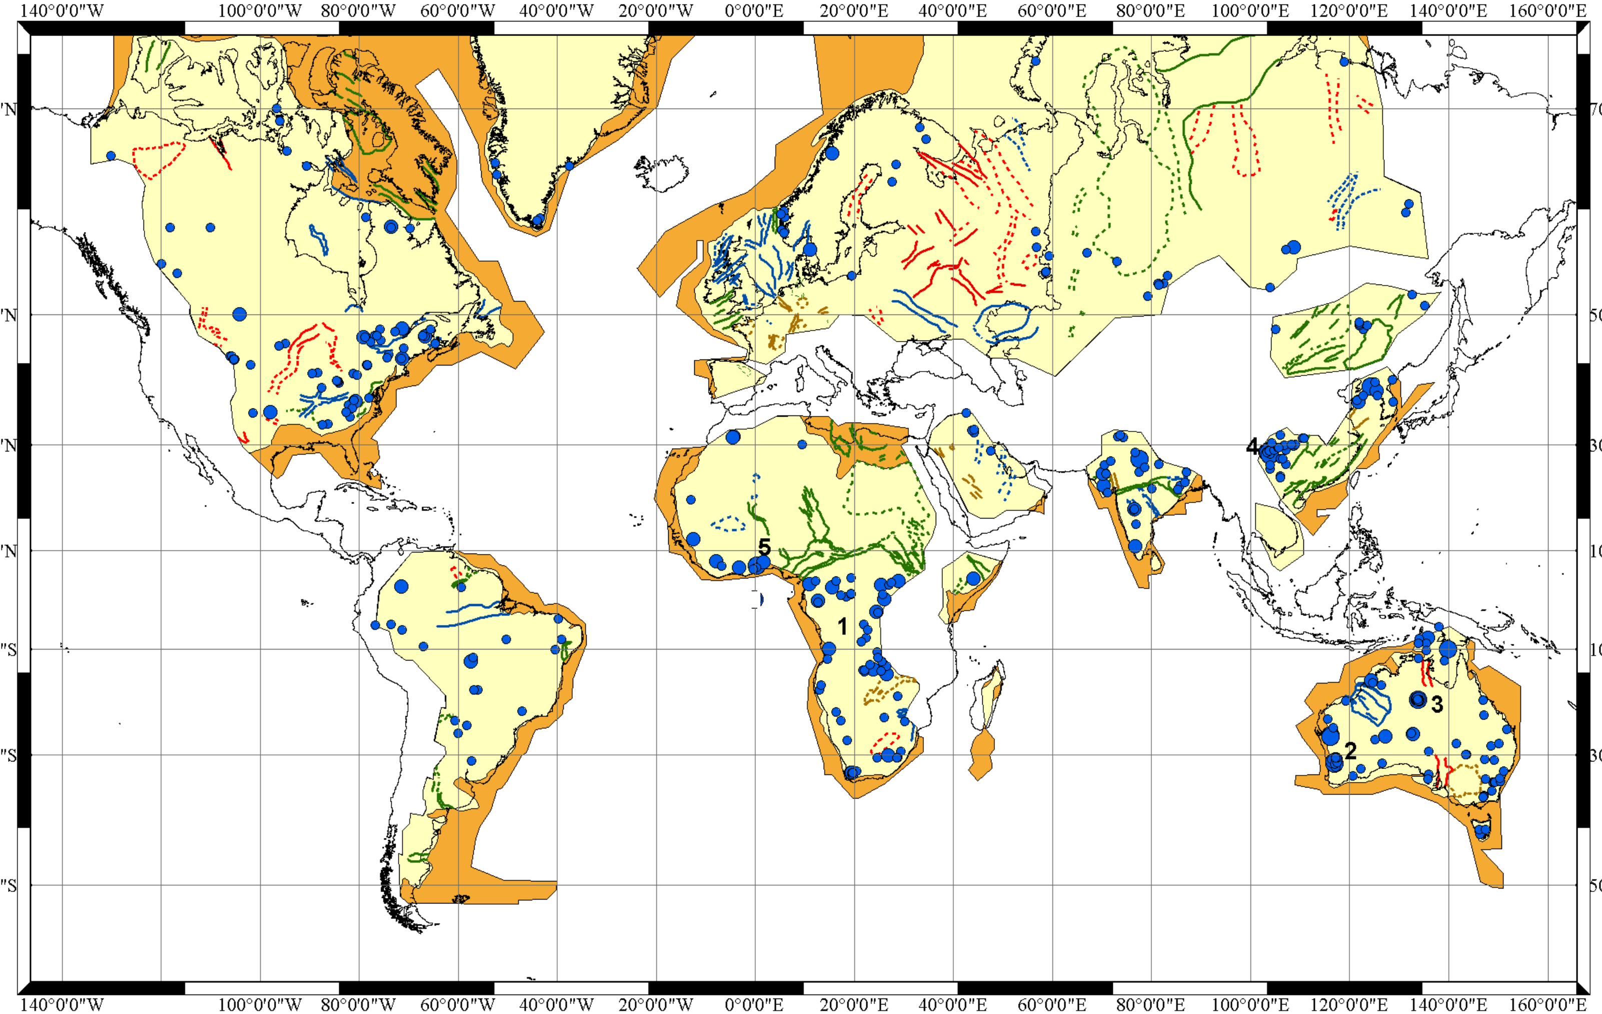

Figure 8c

Non extended crust earthquakes. Earthquakes are indicated with blue circles. Regions of concentrated seismicity are numbered.

USGS

Figure A8–11

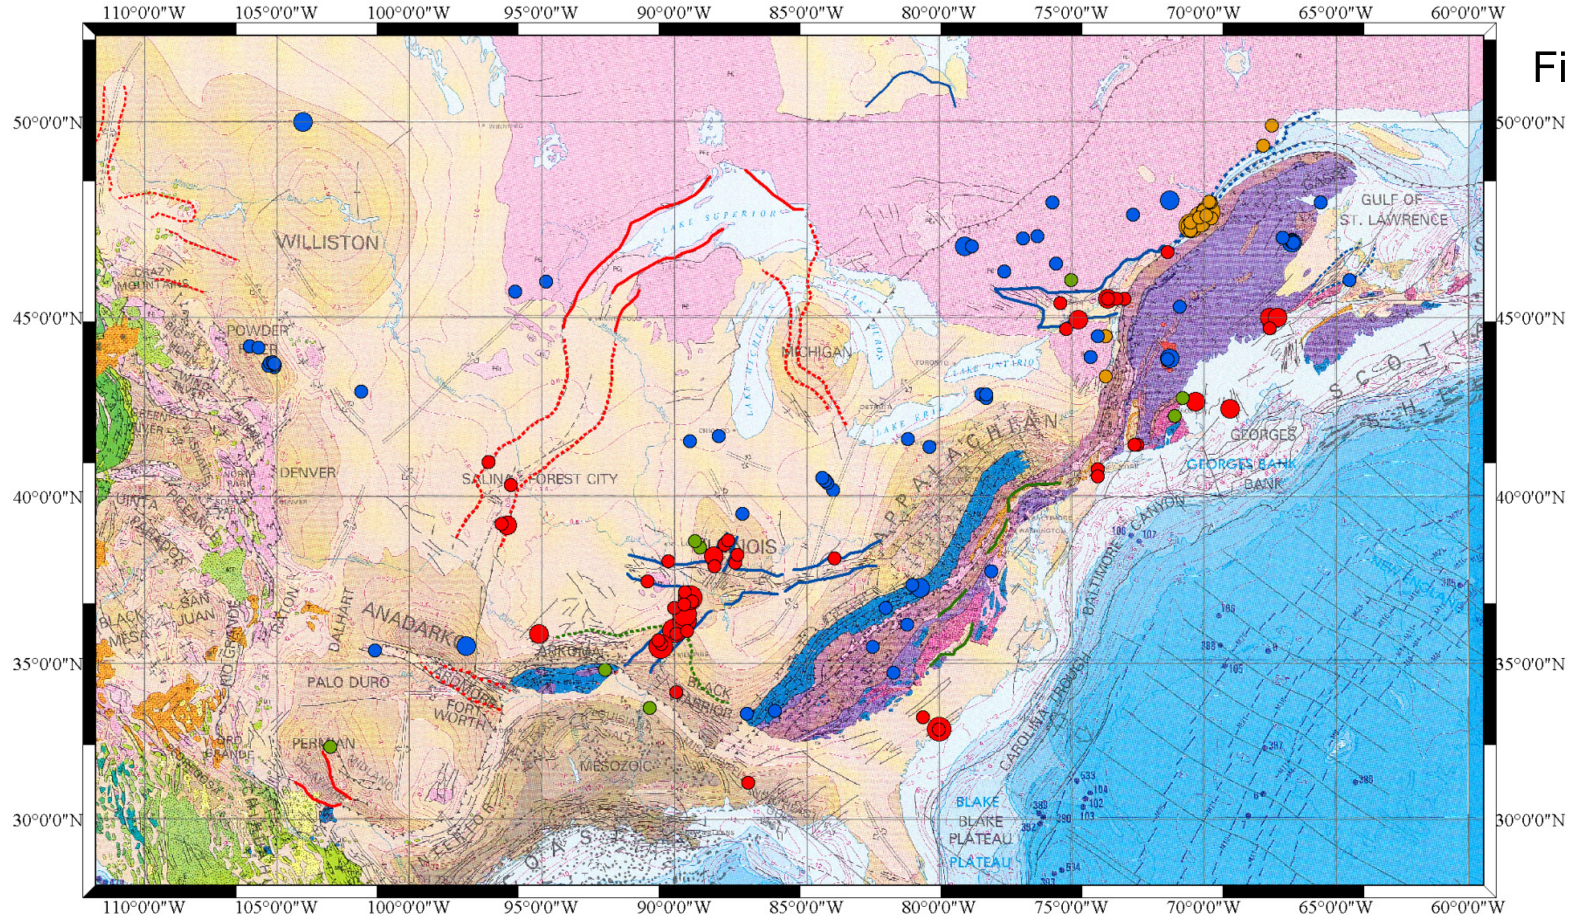

Figure 9

Figure 9. Eastern United States, taken from the Exxon tectonic map of the world [*Exxon*, 1985]. Rifts have been highlighted in red and blue. Earthquakes from the database have been plotted on top. Red circles denote events associated with extended crust, green circles denote events that might be associated with extended crust. Events that occurred within non-extended crust are indicated in blue. Earthquakes that occurred within the St. Lawrence depression are indicated in orange.

USGS

Figure A8–12

Lithospheric Root

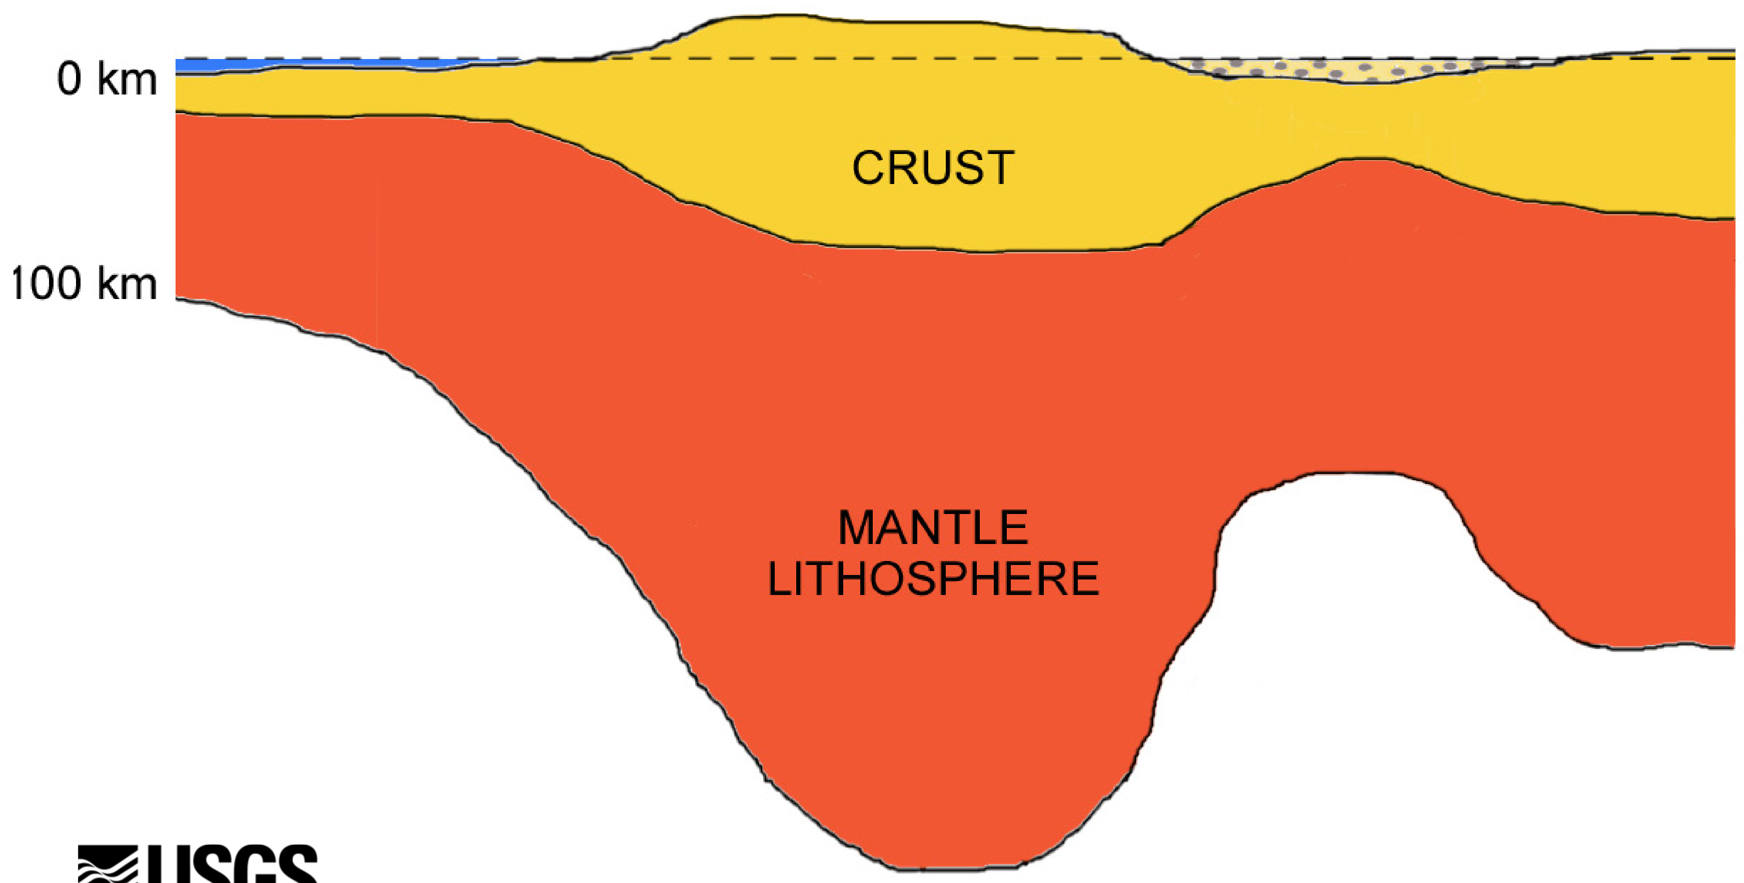

USGS

Model of Precambrian Lithosphere

Figure A8–13

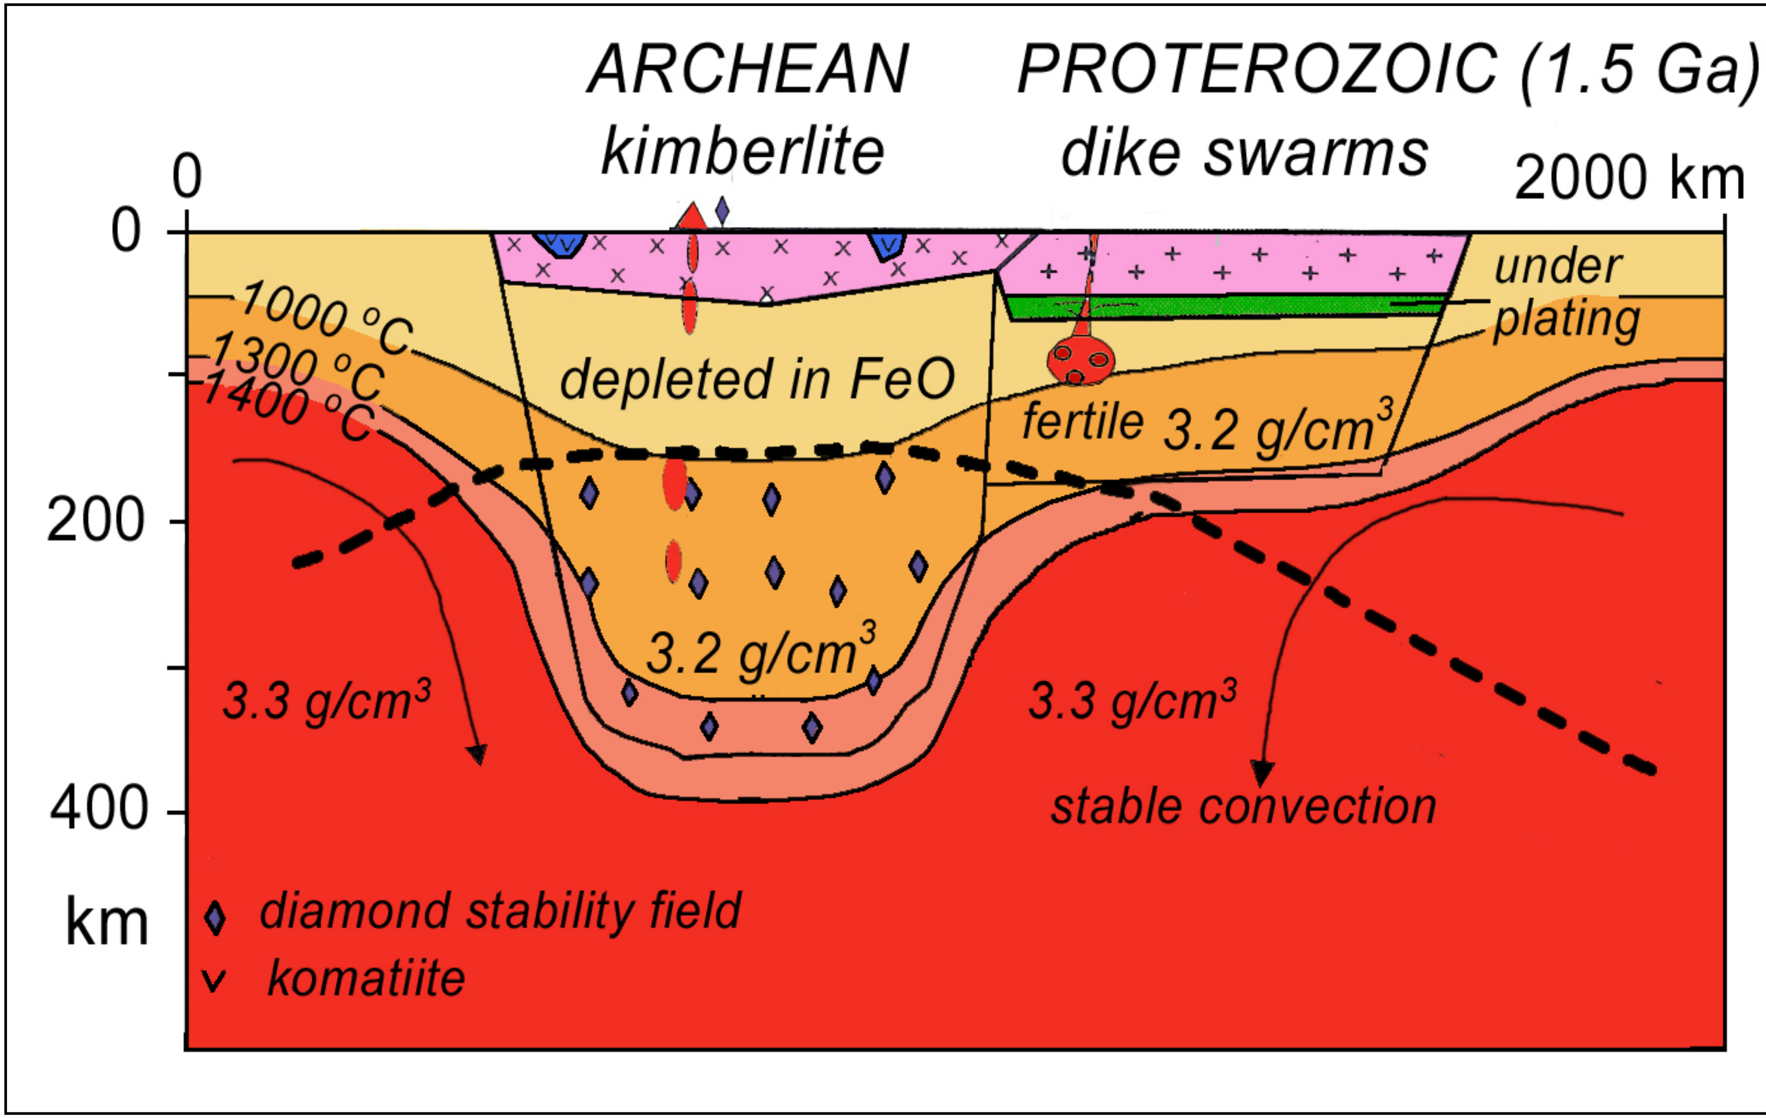

USGS

Figure A8–14

Lithosphere is Generally in a State of Failure Equilibrium

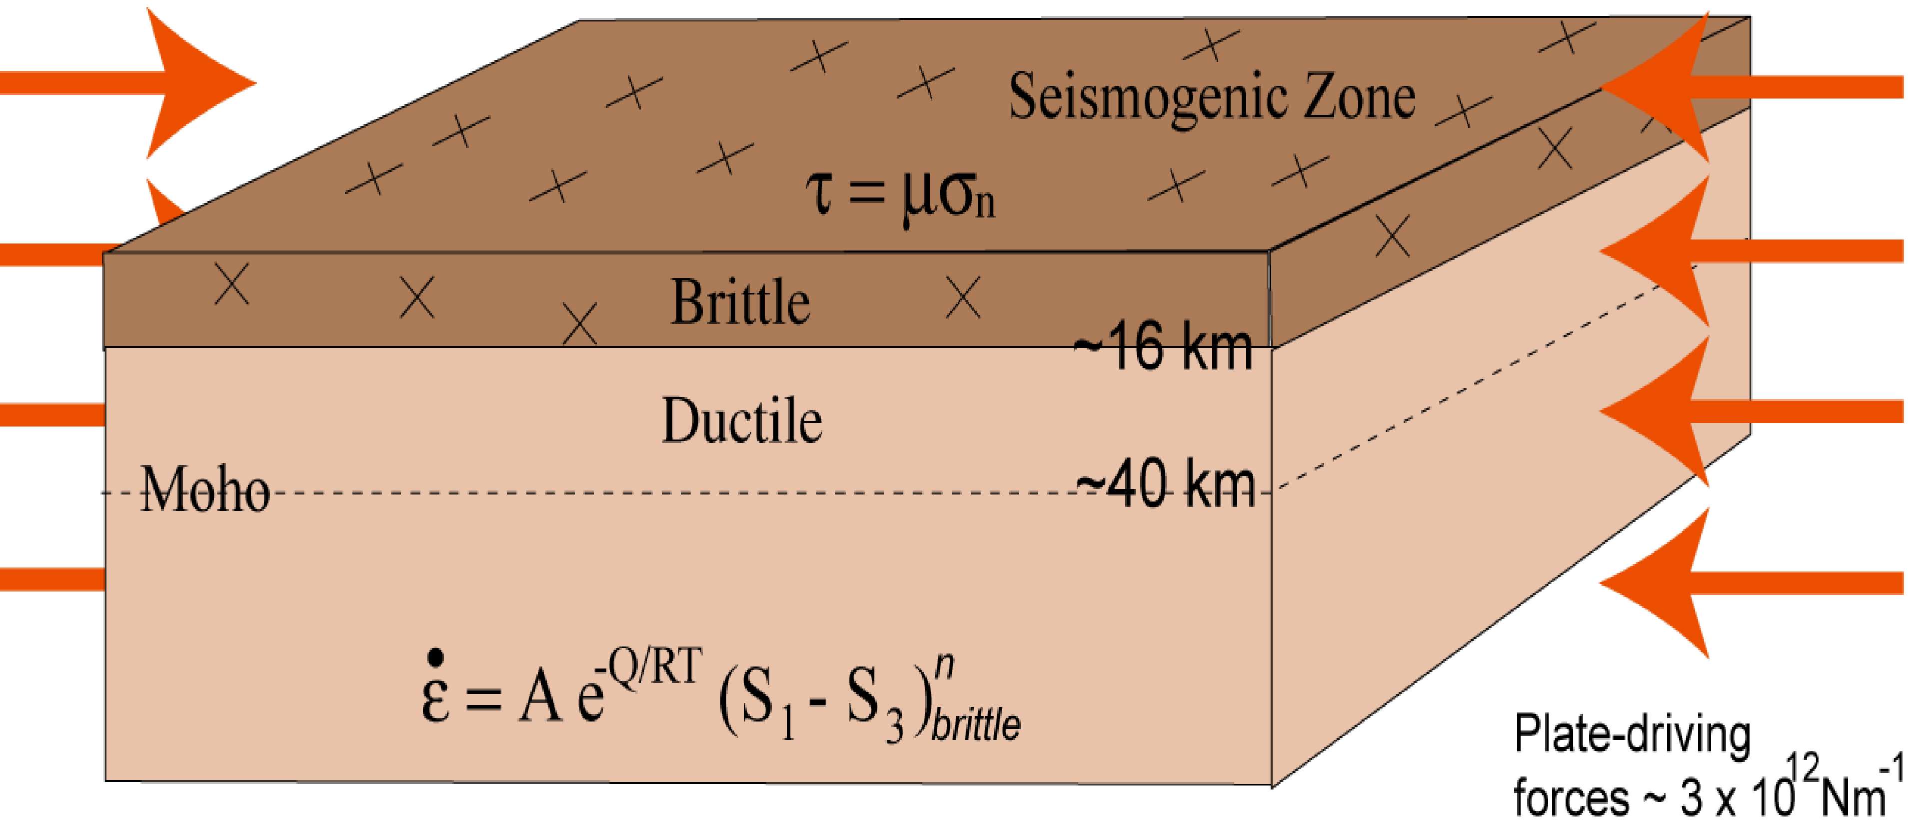

USGS

Figure A8–15

Tomographic Model

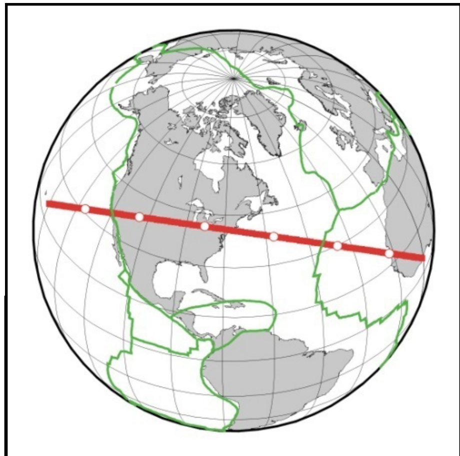

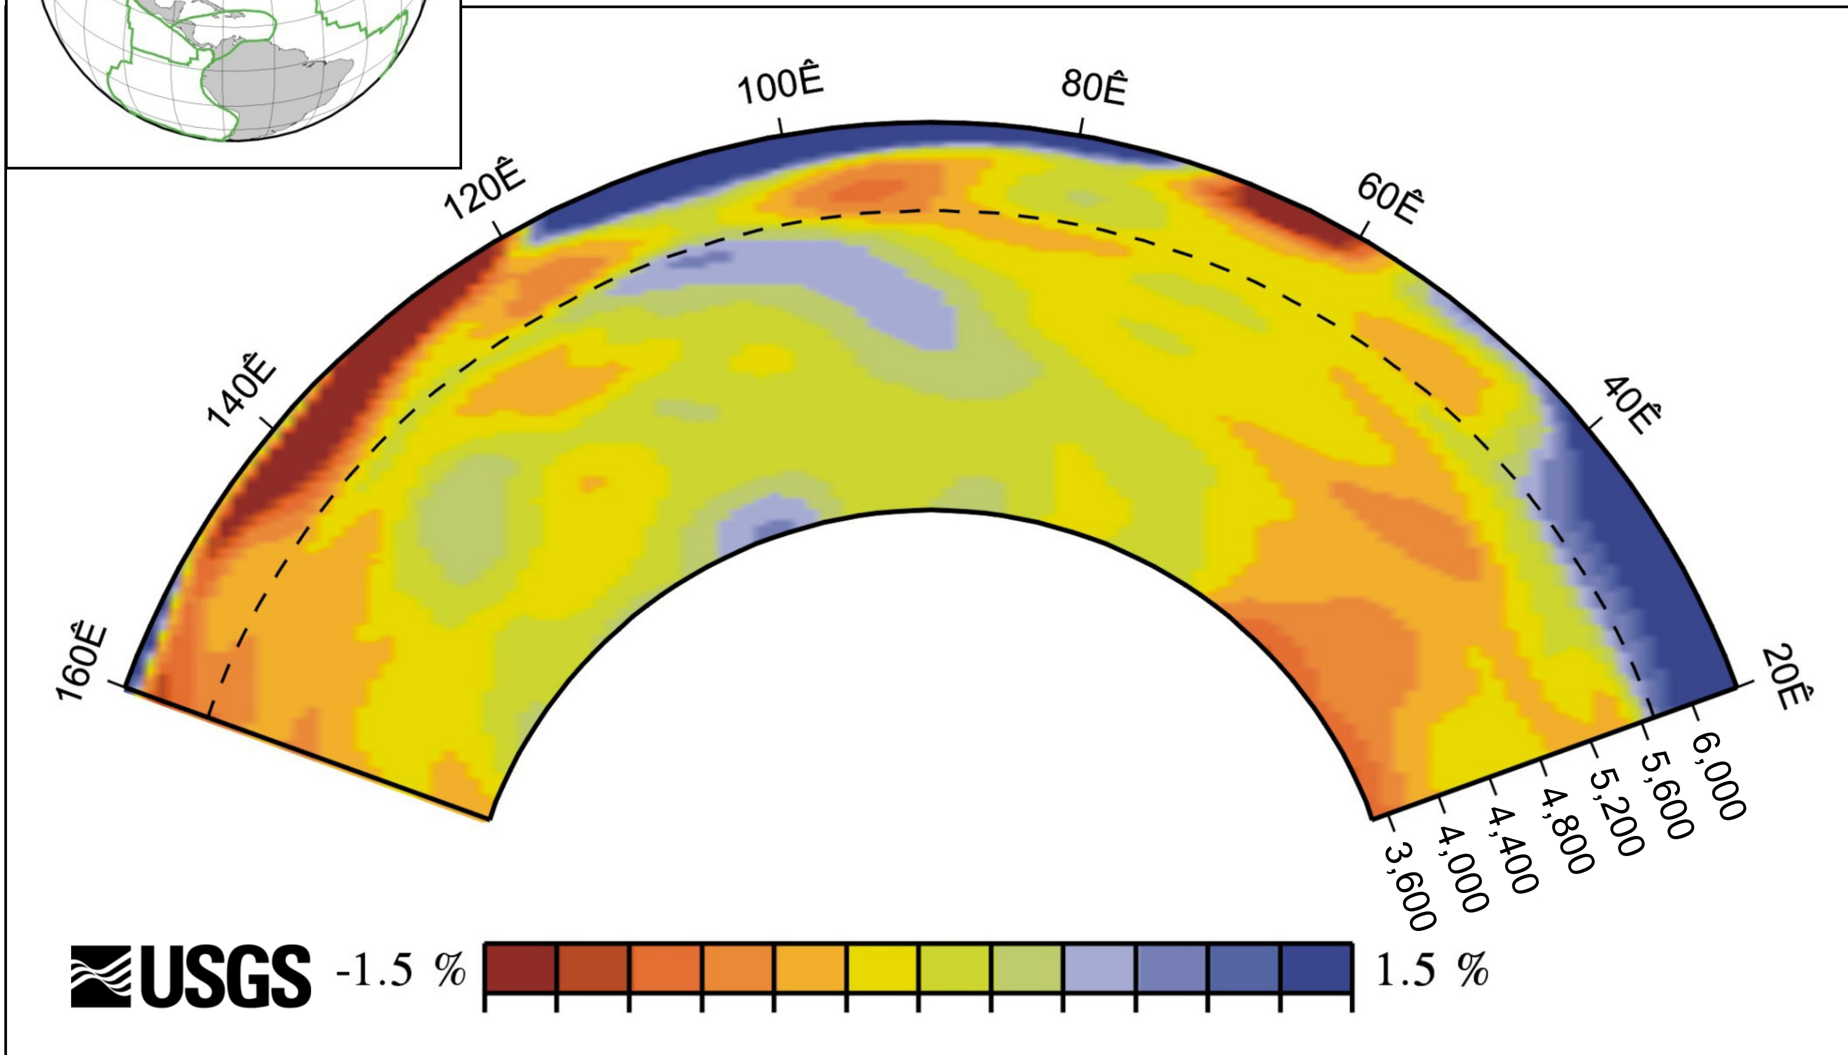

Figure A8–16

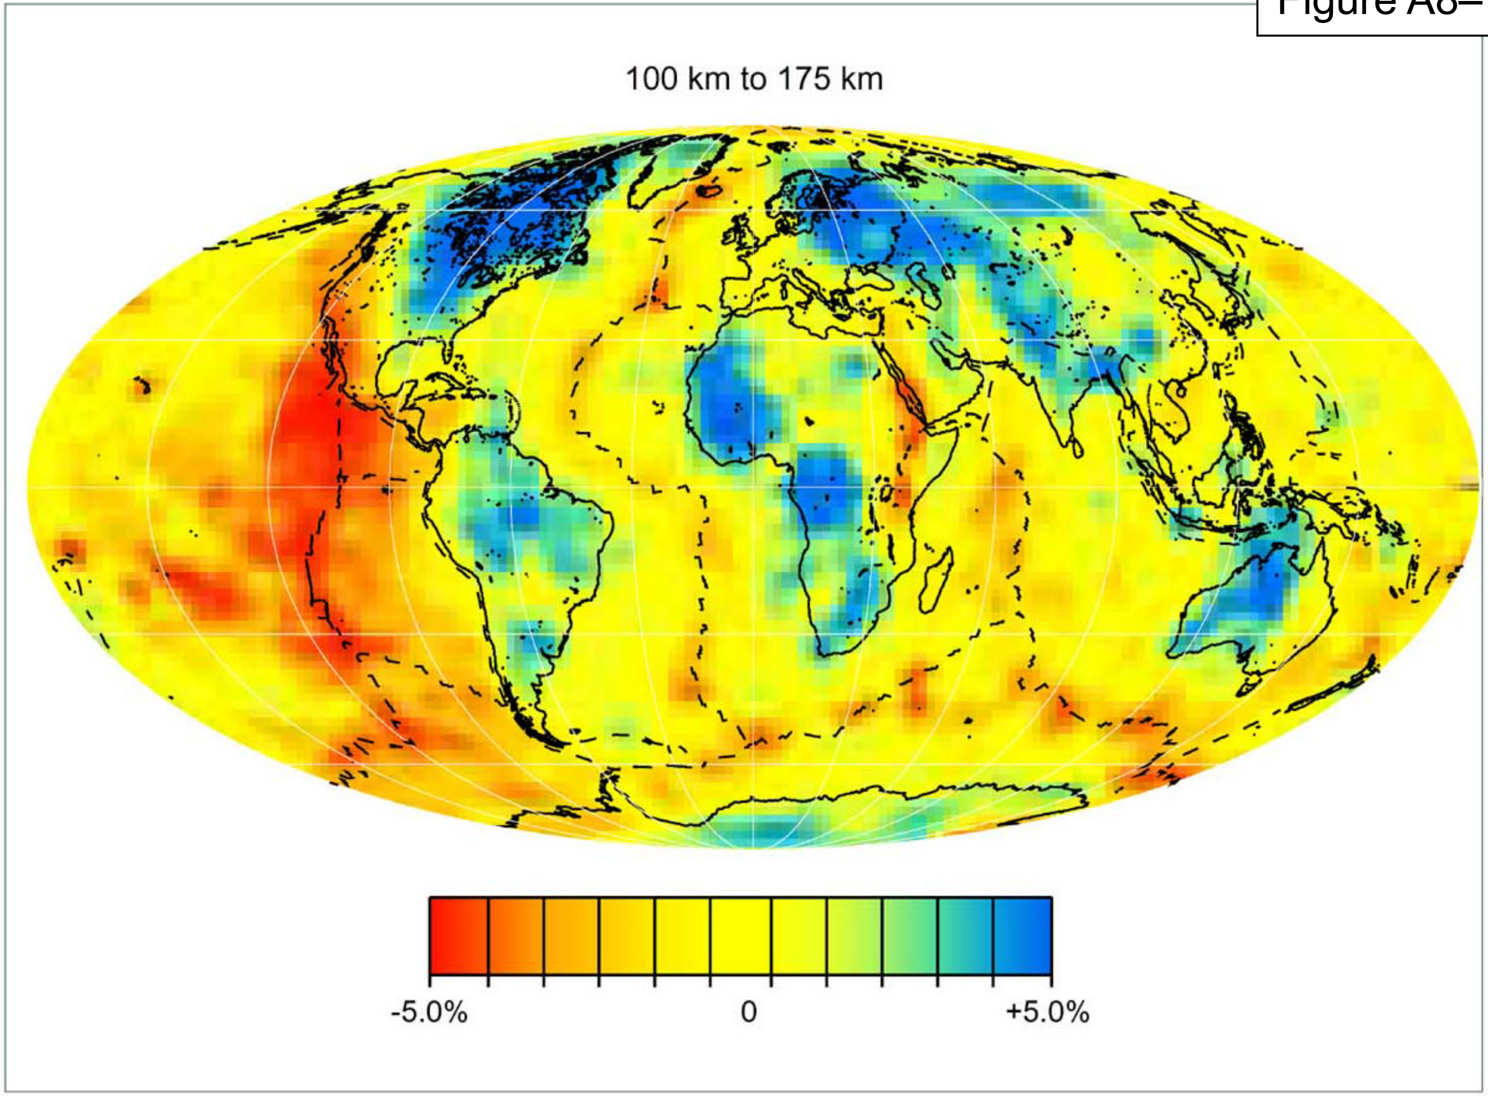

USGS

Figure A8–17

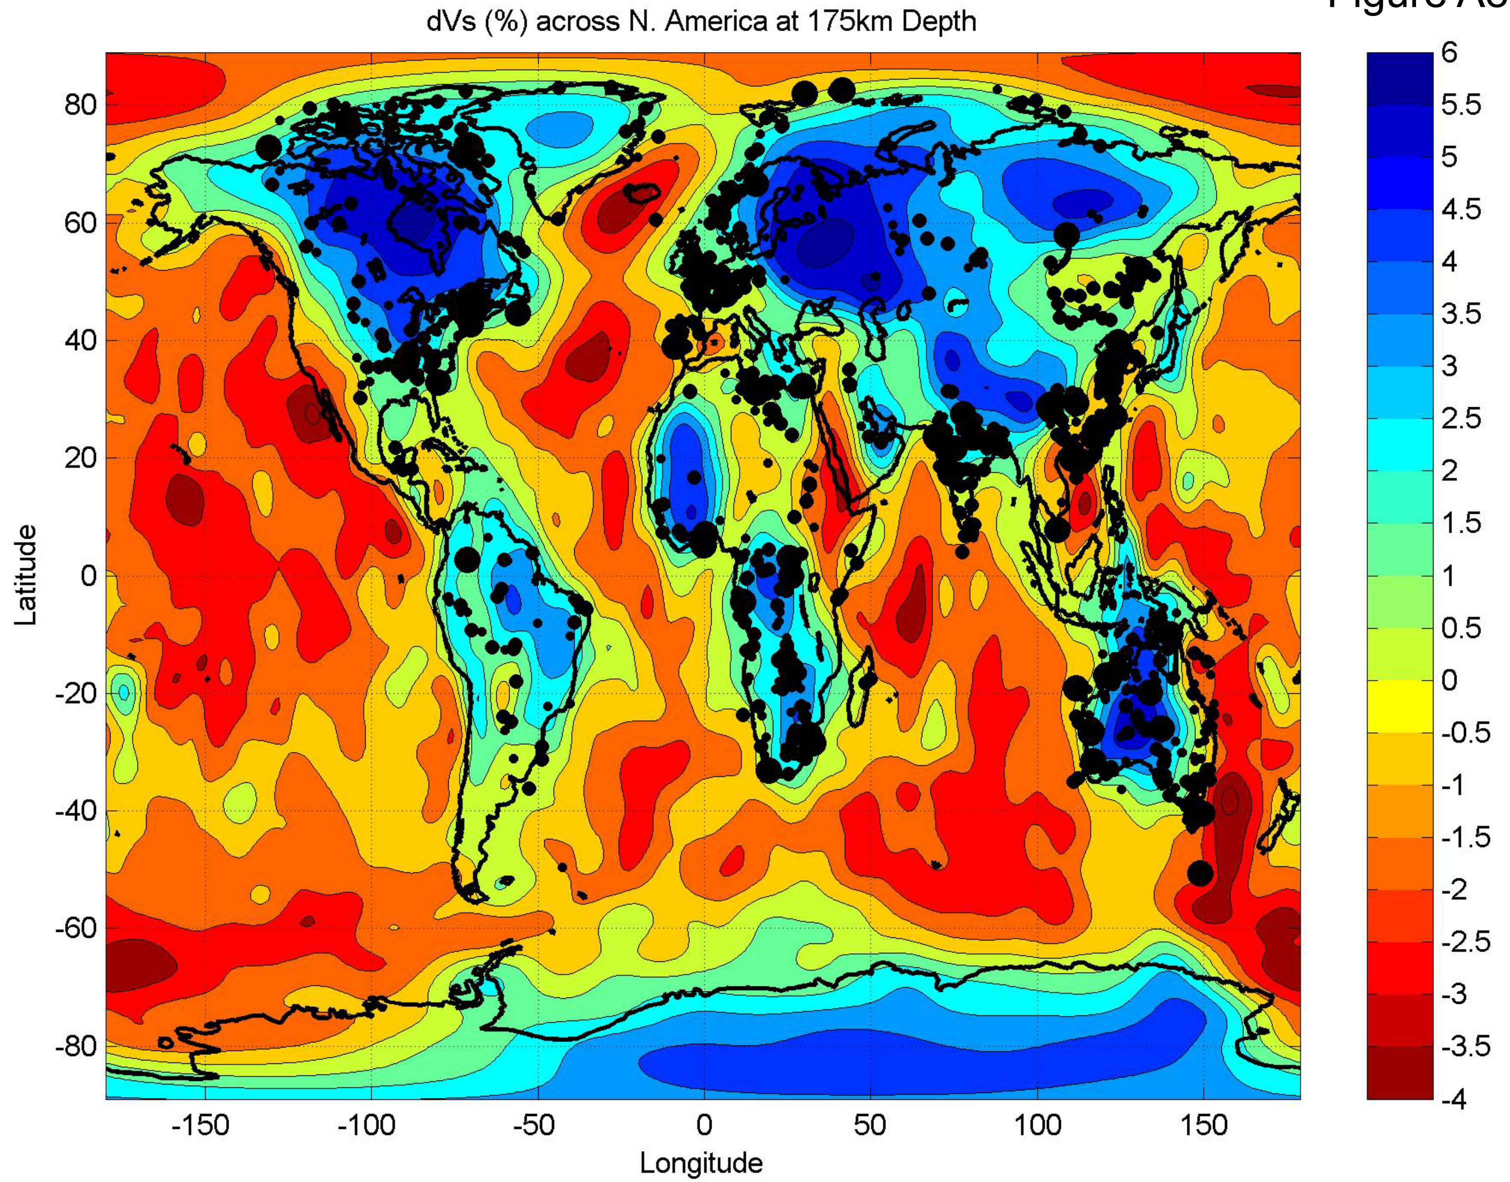
dVs (%) across N. America at 175km Depth
Latitude
Longitude
80
60
40
20
0
-20
-40
-60
-80
-150
-100
-50
0
50
100
150
6
5.5
5
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
-0.5
-1
-1.5
-2
-2.5
-3
-3.5
-4

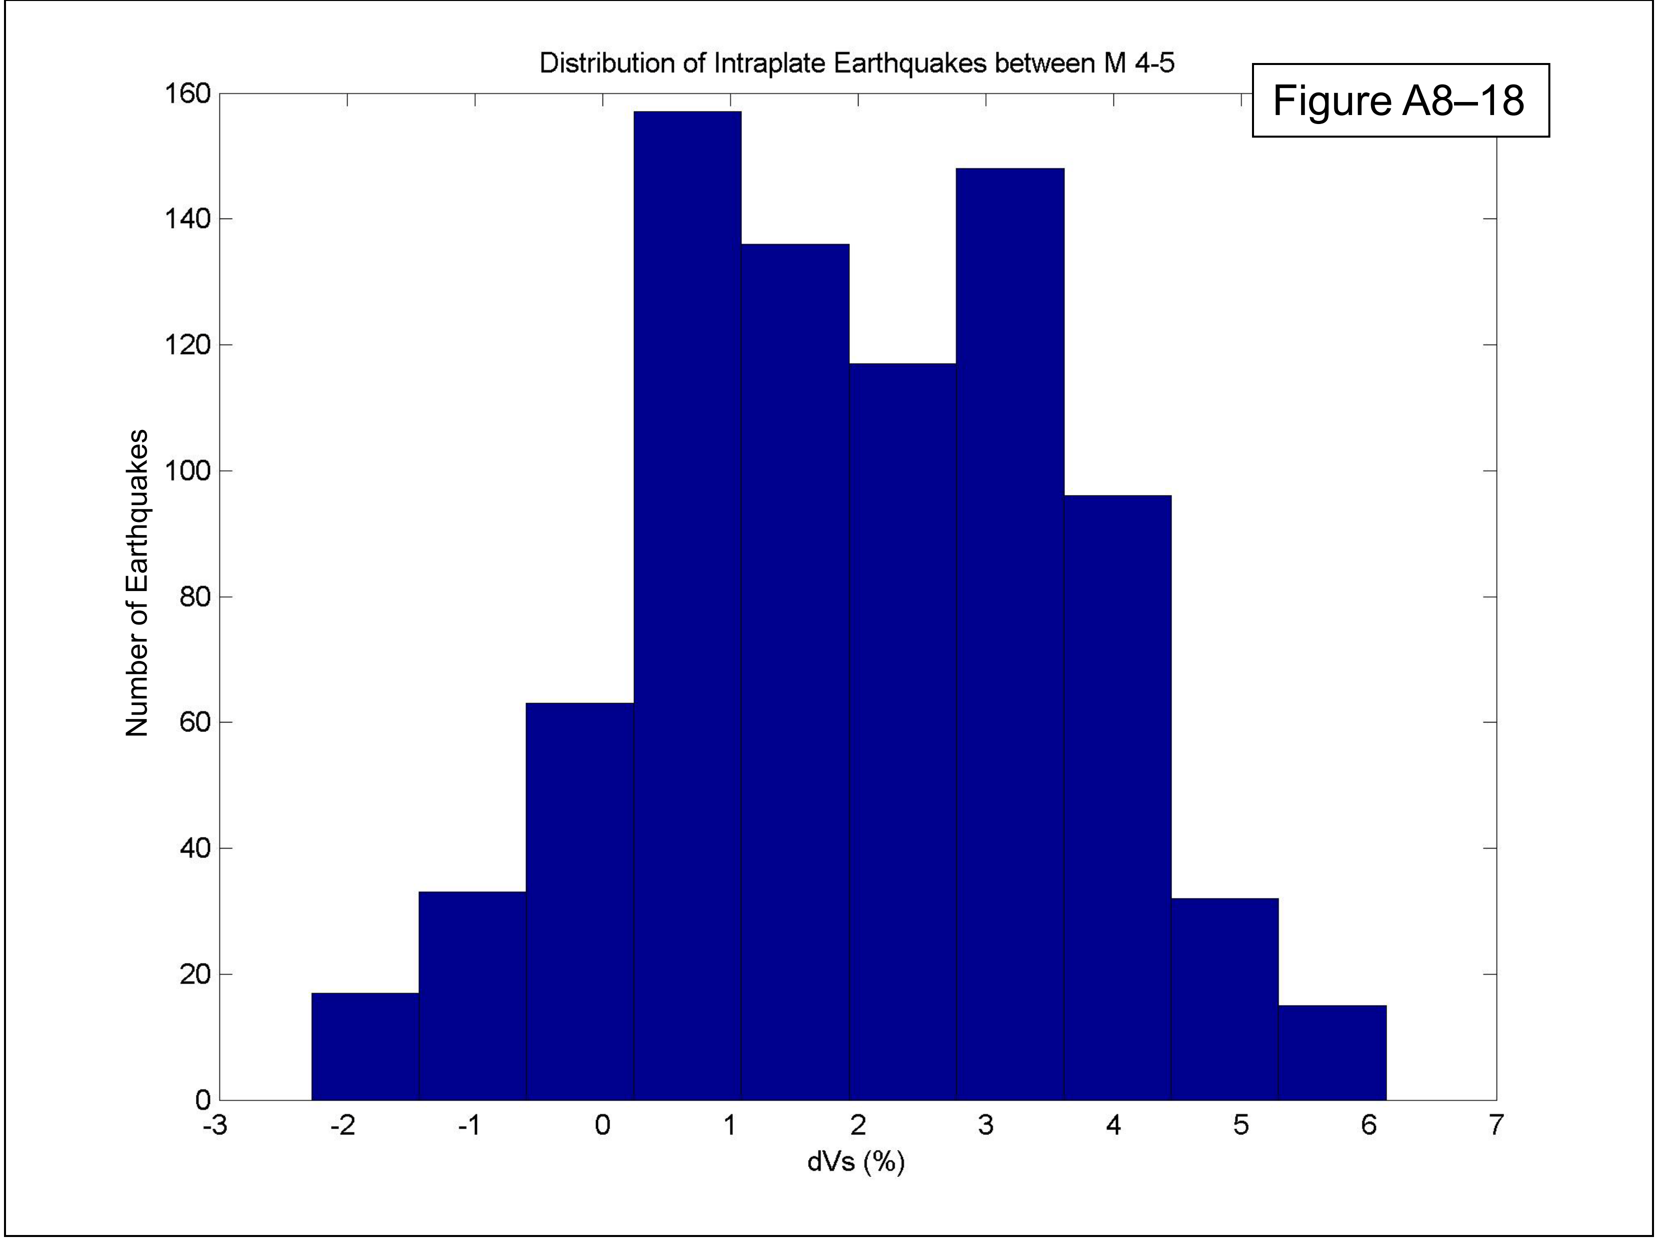
Distribution of Intraplate Earthquakes between M 4-5
Figure A8–18
Number of Earthquakes
dVs (%)
0
20
40
60
80
100
120
140
160
-3
-2
-1
0
1
2
3
4
5
6
7

USGS

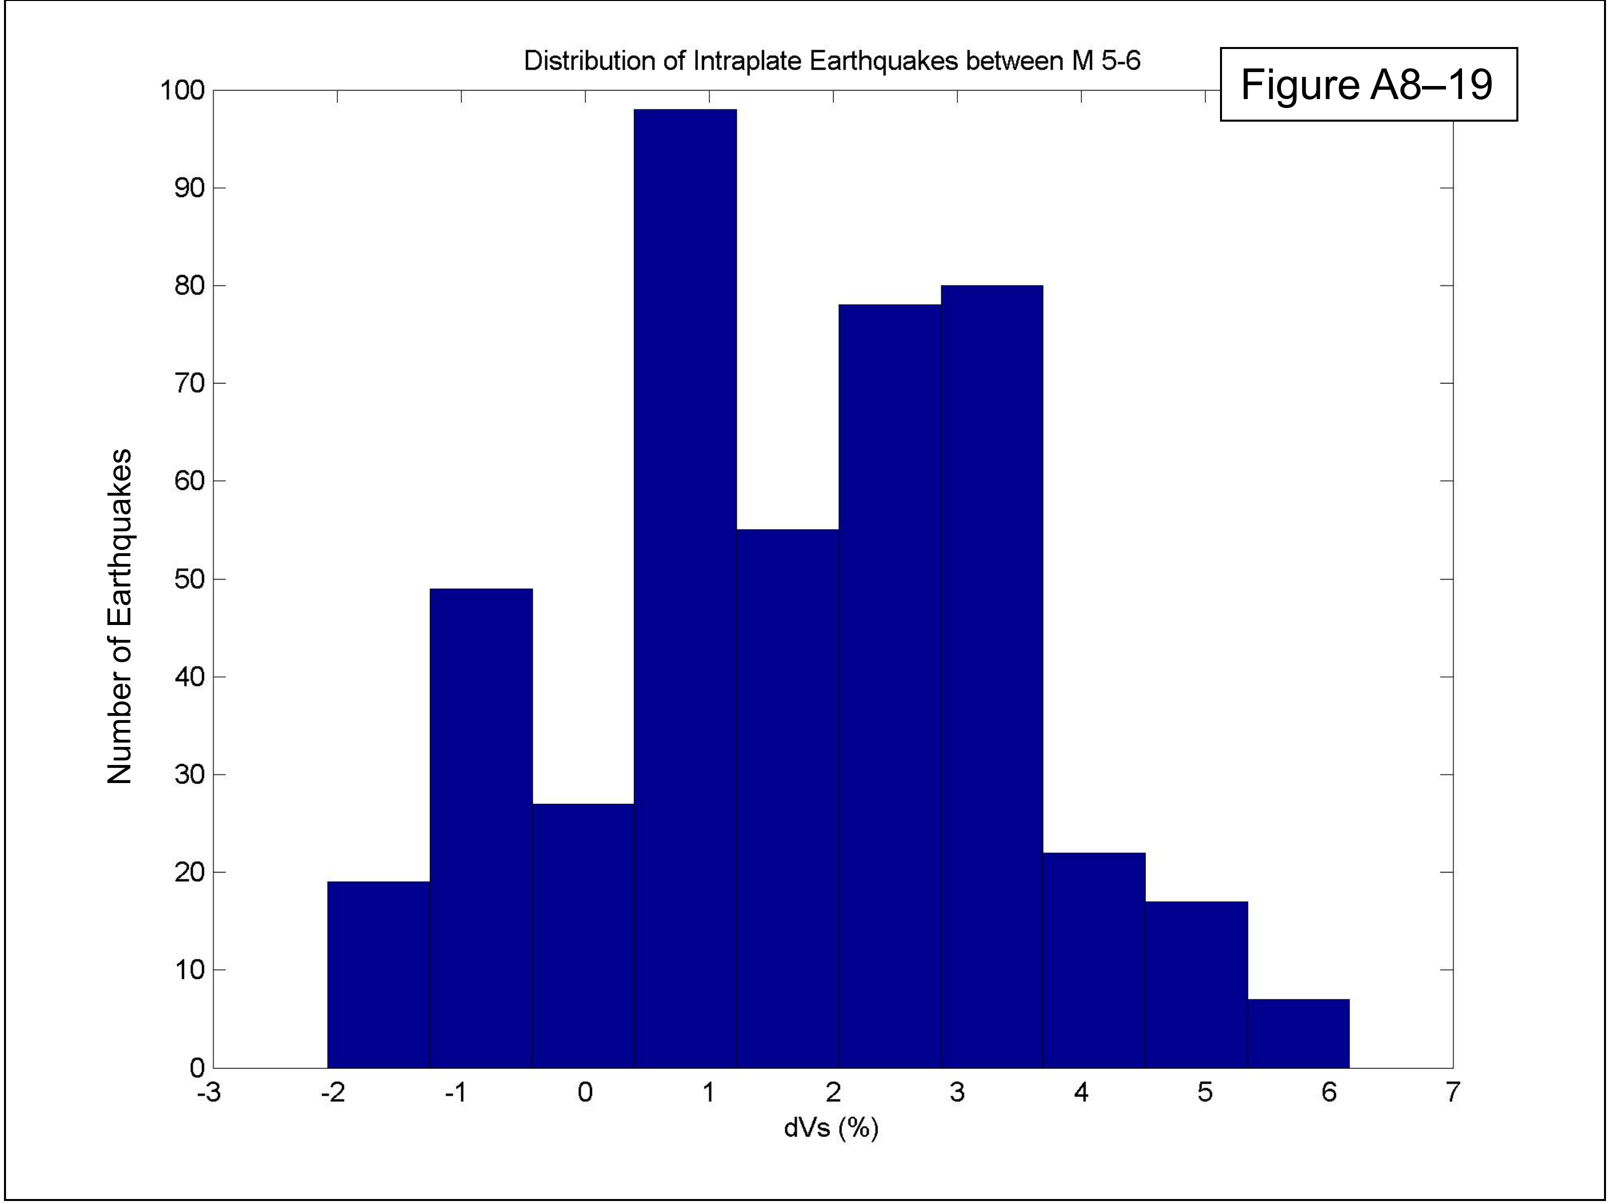
Distribution of Intraplate Earthquakes between M 5-6
Figure A8–19
Number of Earthquakes
dVs (%)
0
10
20
30
40
50
60
70
80
90
100
-3
-2
-1
0
1
2
3
4
5
6
7

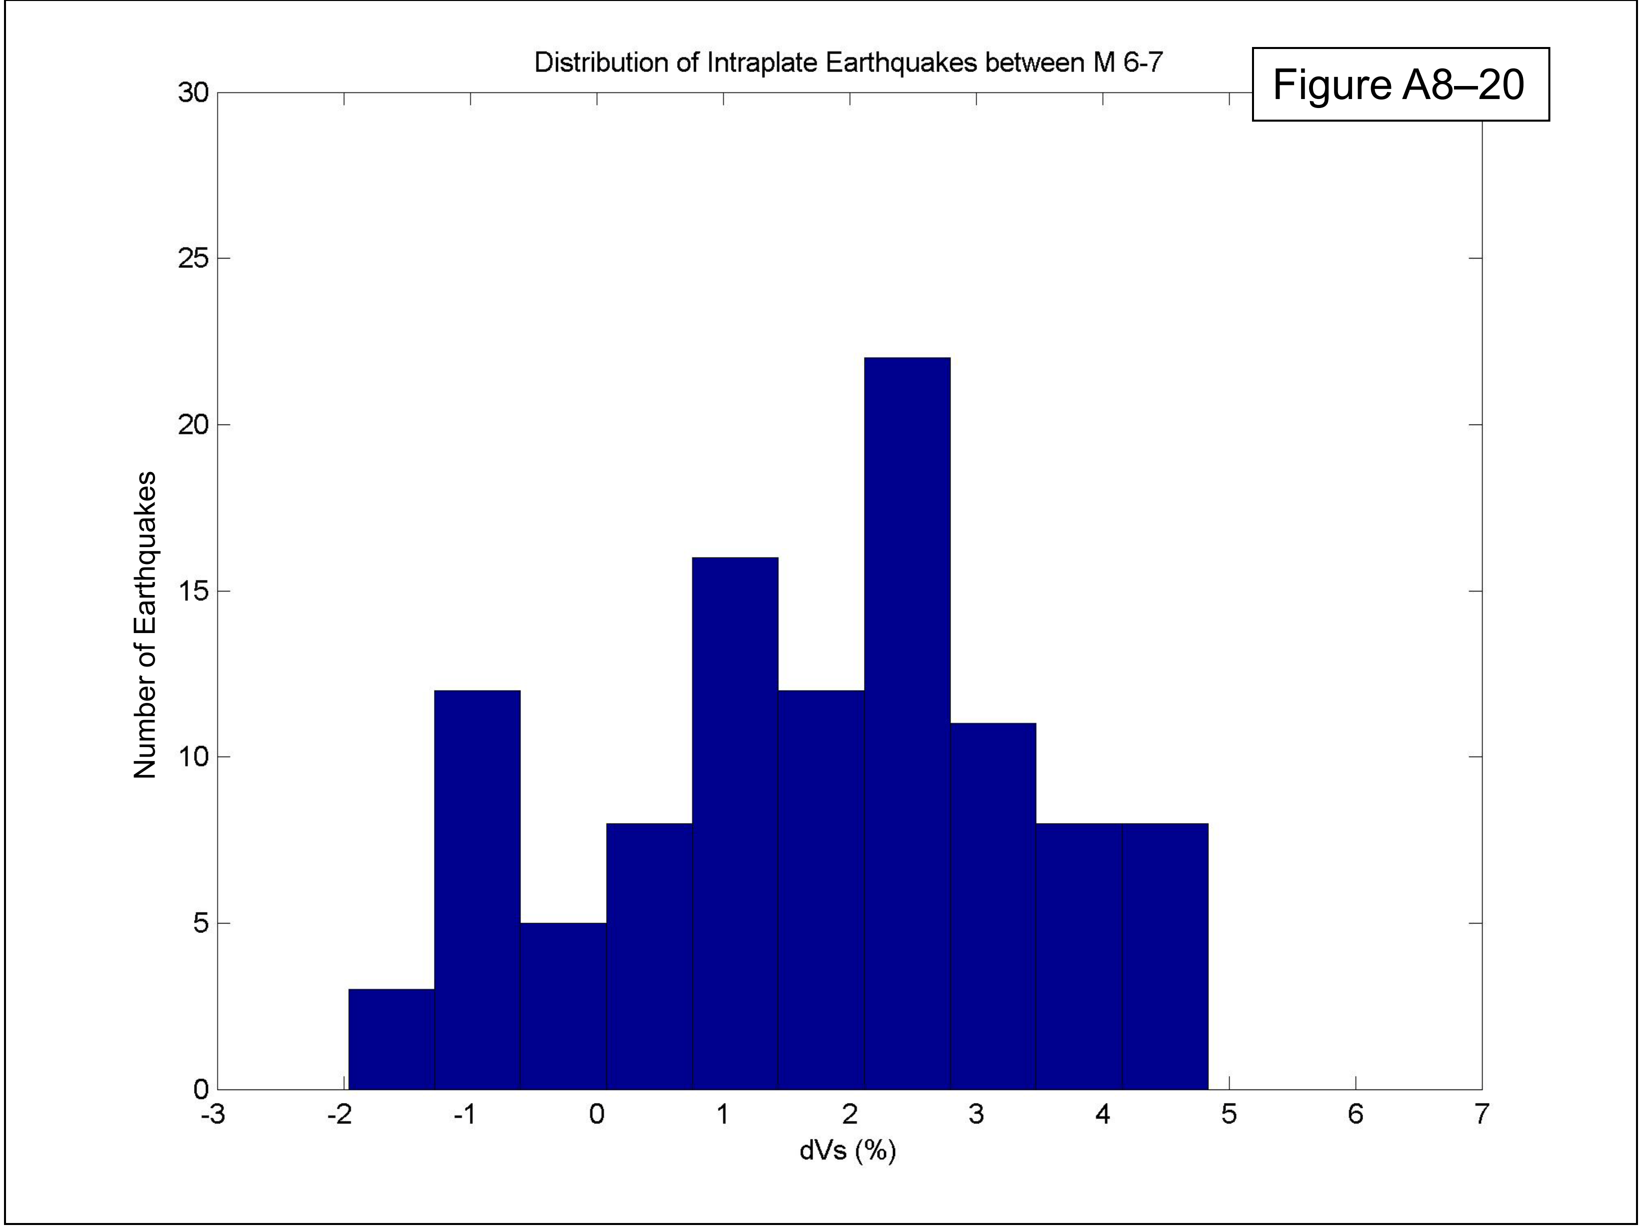
Distribution of Intraplate Earthquakes between M 6-7
Figure A8–20
Number of Earthquakes
dVs (%)
30
25
20
15
10
5
0
-3
-2
-1
0
1
2
3
4
5
6
7

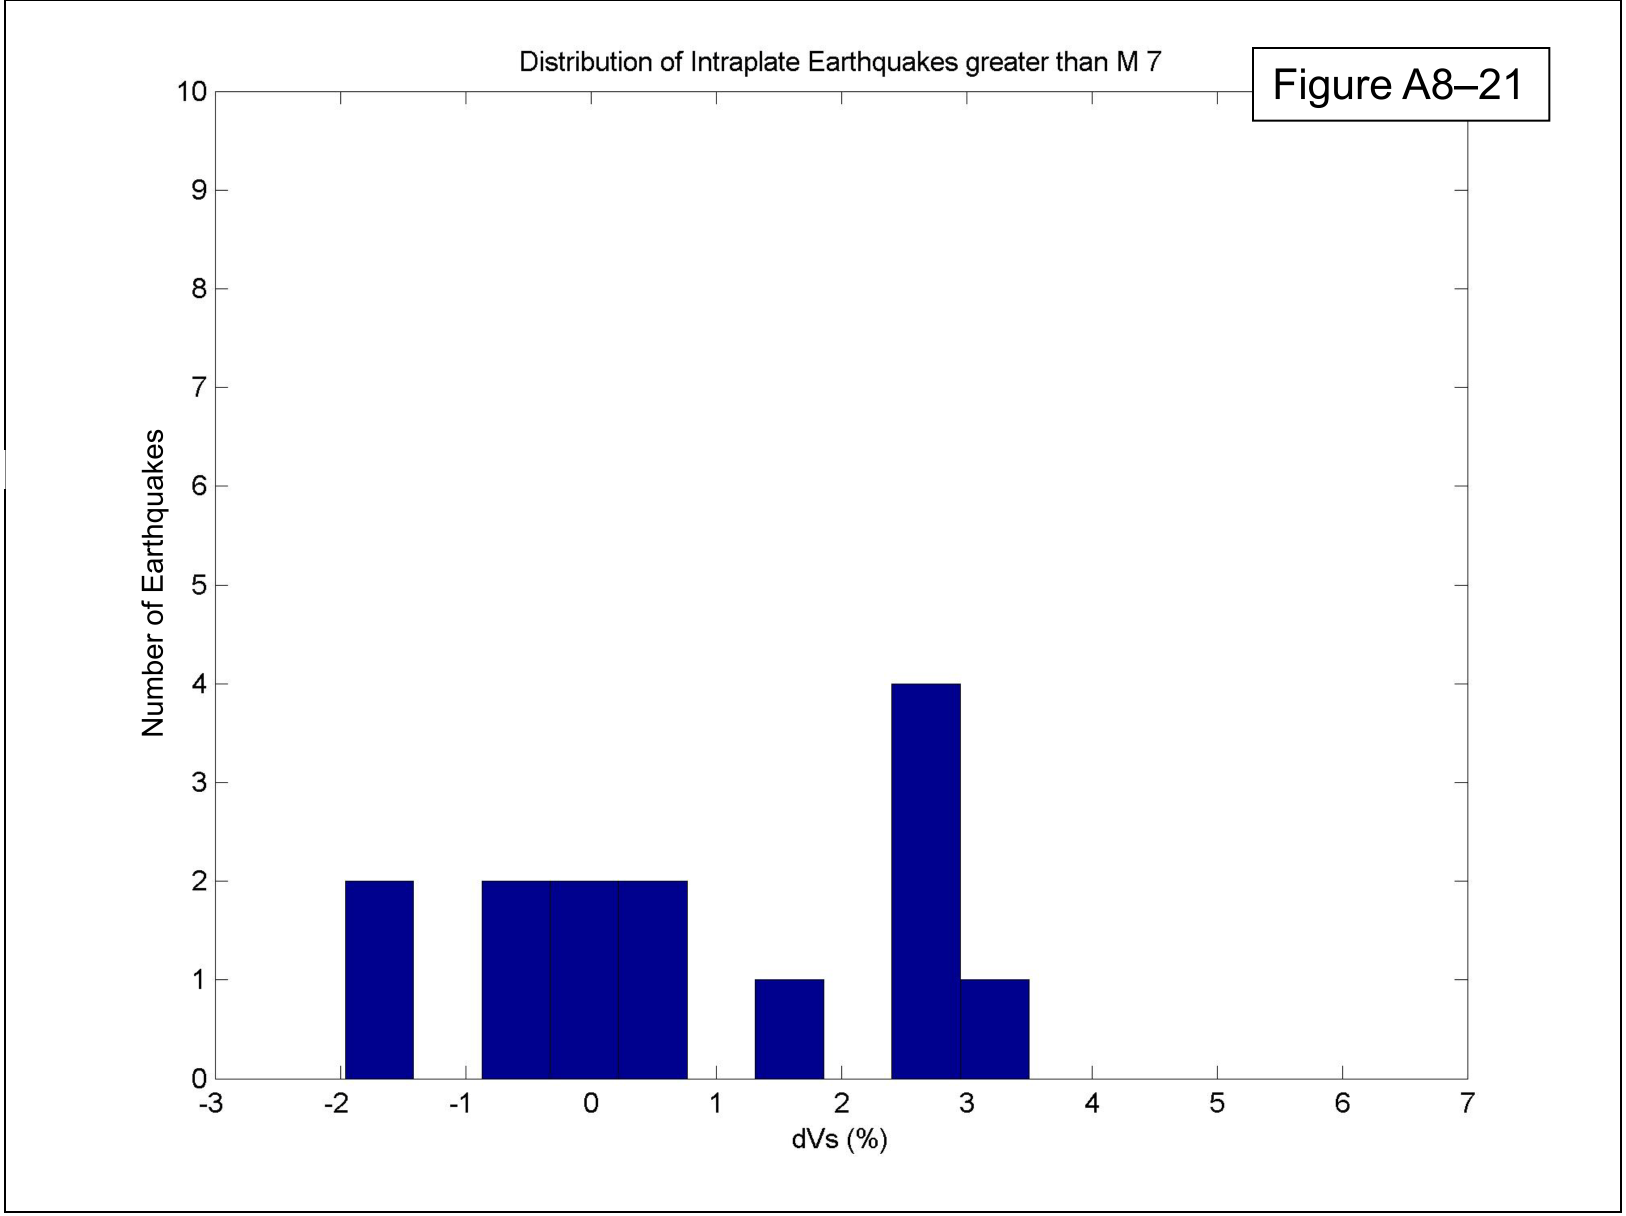
Distribution of Intraplate Earthquakes greater than M 7
Figure A8–21
Number of Earthquakes
dVs (%)
0
1
2
3
4
5
6
7
8
9
10
-3
-2
-1
0
1
2
3
4
5
6
7

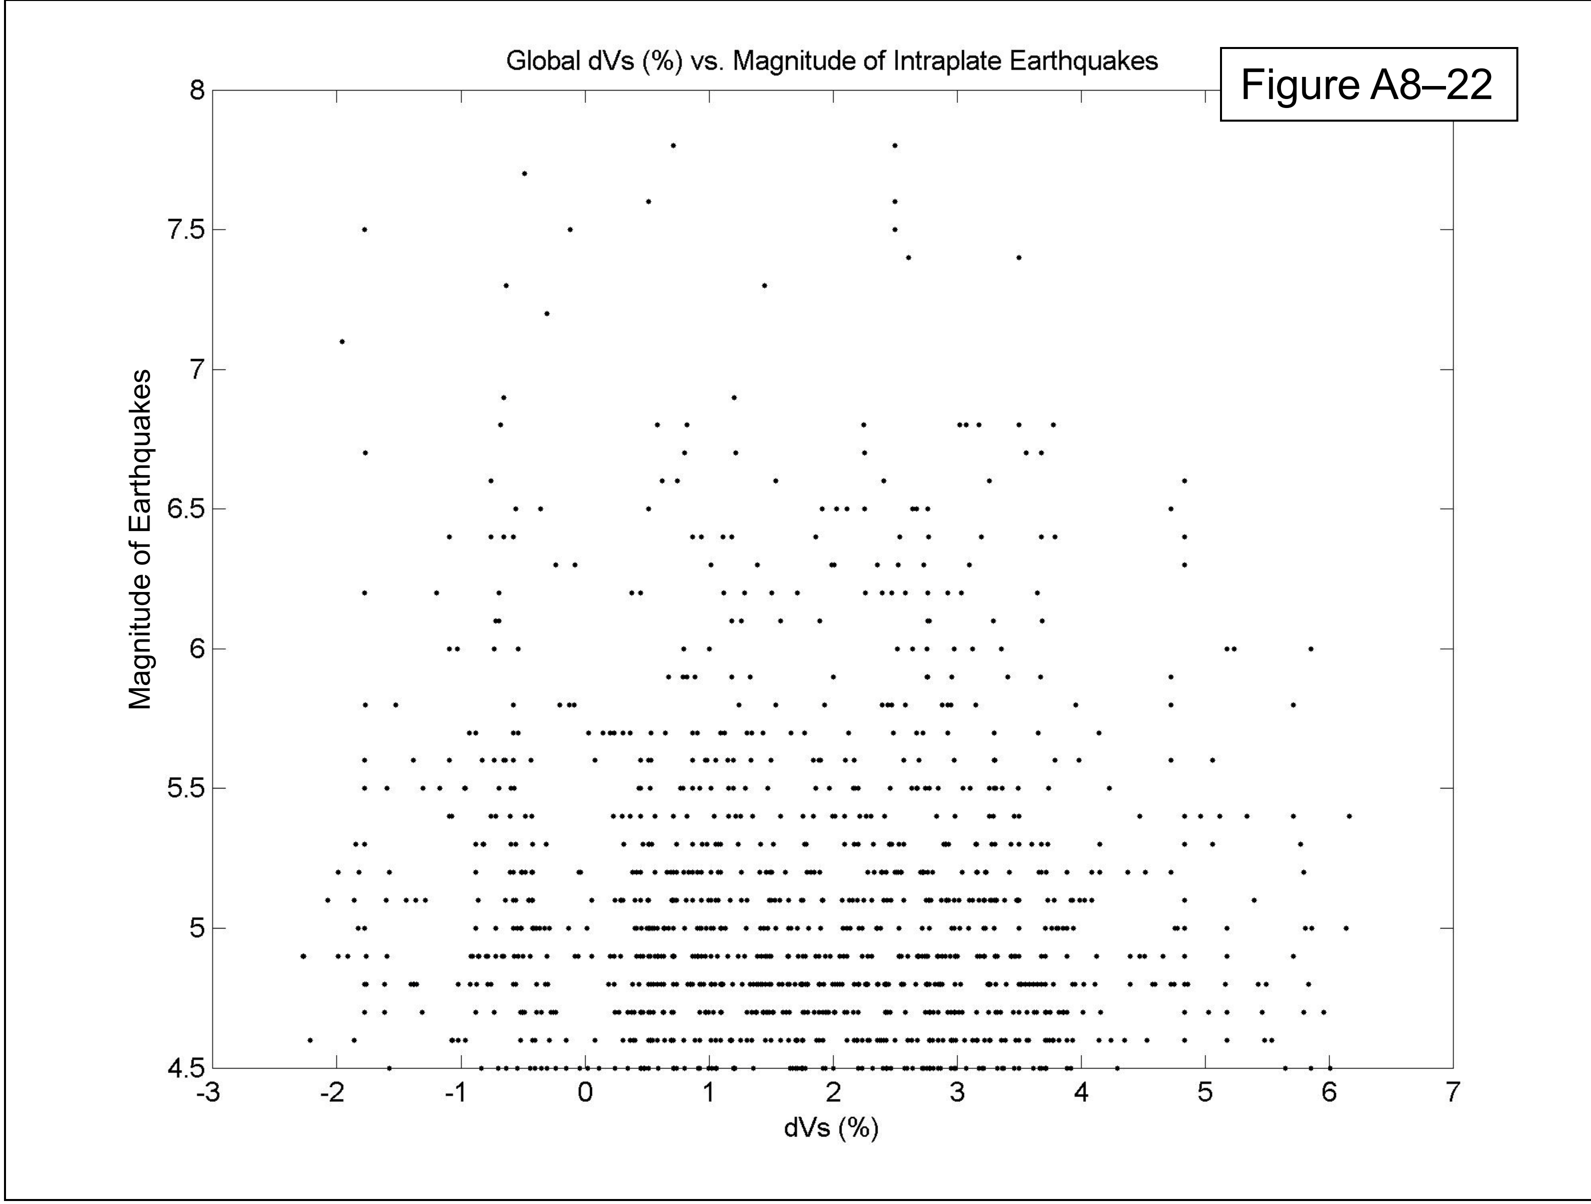
Global dVs (%) vs. Magnitude of Intraplate Earthquakes
Figure A8–22
Magnitude of Earthquakes
8
7.5
7
6.5
6
5.5
5
4.5
-3
-2
-1
0
1
2
3
4
5
6
7
dVs (%)

Figure A8–23

Vs at 100 km,

(S. Van Der Lee and G. Nolet, 1997)

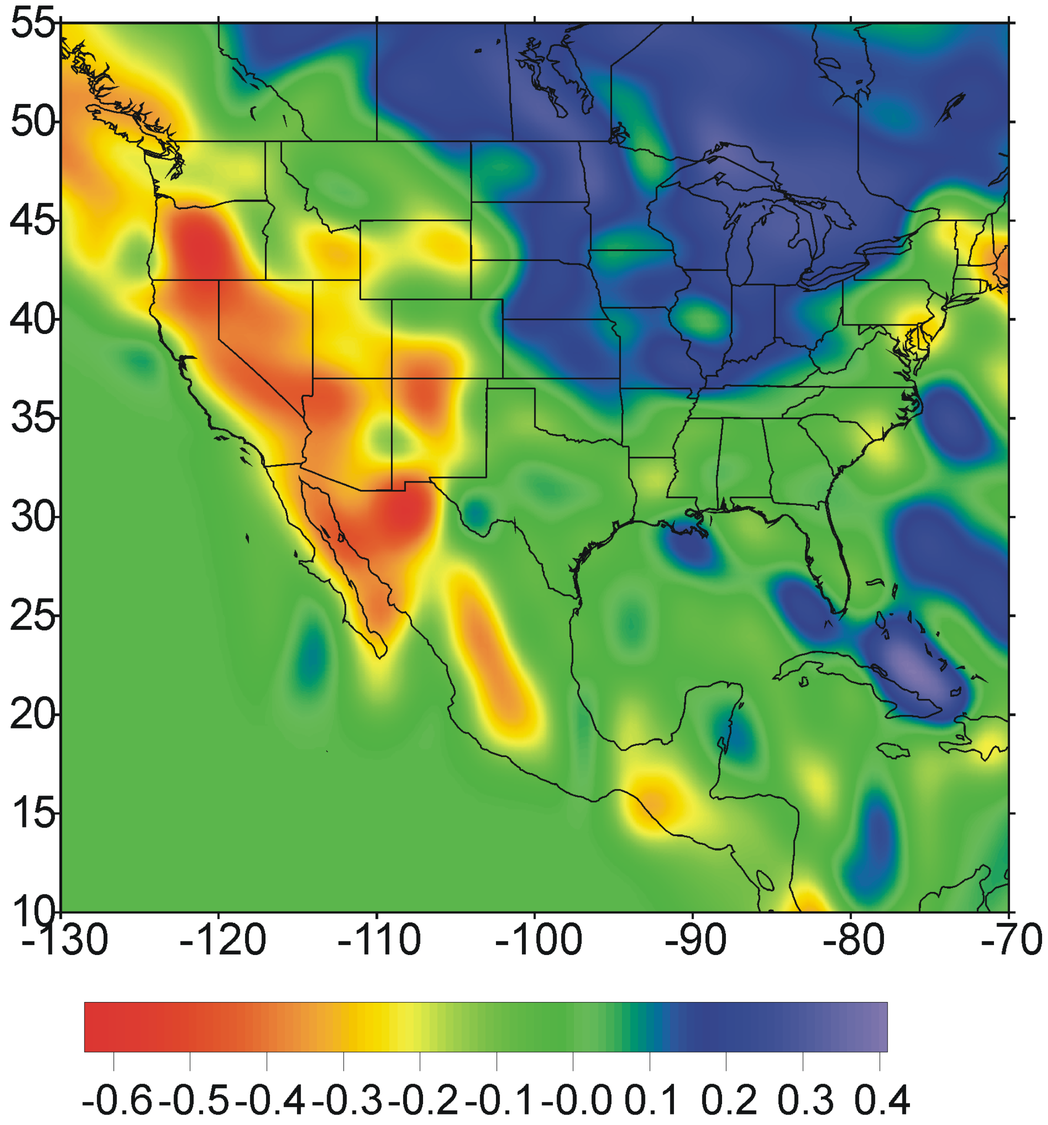

USGS

USGS
science for a changing world

Figure A9–1

CEUS Mean Earthquake Rates & Observations

Steve Harmsen, USGS

NRC-USGS Mmax Workshop

September 8-9, 2008

U.S. Department of the Interior
U.S. Geological Survey

Basic Approach

Figure A9–2

- Compare expected number of background-zone earthquakes in two sub-regions of CEUS with historical data. For expected rates, use 2008 USGS-NSHMP PSHA source model.
- Comment on expected rates of M>6 background compared to historical record.
- Comment on expected rates of M>7 background sources compared to paleoseismic evidence in those two sub-regions.
- These are initial observations. This material is not abstracted from a journal article.

USGS

Figure A9–3

What is known about the fault?

- Wide range of knowledge about faults
- Hazardous CEUS faults are Cheraw, Colorado and Meers, Oklahoma (no location uncertainty) and NMSZ (with location uncert.). Earthquake size (Mchar) can be estimated from evidence.
- Charleston, SC is a characteristic source zone. Size (Mchar) can be estimated from evidence.
- Background sources often at locations with virtually no evidence for estimating fault size. Methodology assumes that long, capable faults are everywhere (close to every location in the CEUS and WUS).

USGS

Scope of Analysis

Figure A9–4

- Background sources only
- No faults, fault zones, characteristic-earthquake source zones are included in mean rate calculations, E[N].
- Historical earthquakes used for tabulation of N on the basis of a declustered USGS catalog. Not explicitly considering range of uncertainty of historic earthquake magnitudes. Comment on the range of reported magnitudes.

USGS

Mmax; Stars are m6+ epicenters, 1708 to 2008

Figure A9–5

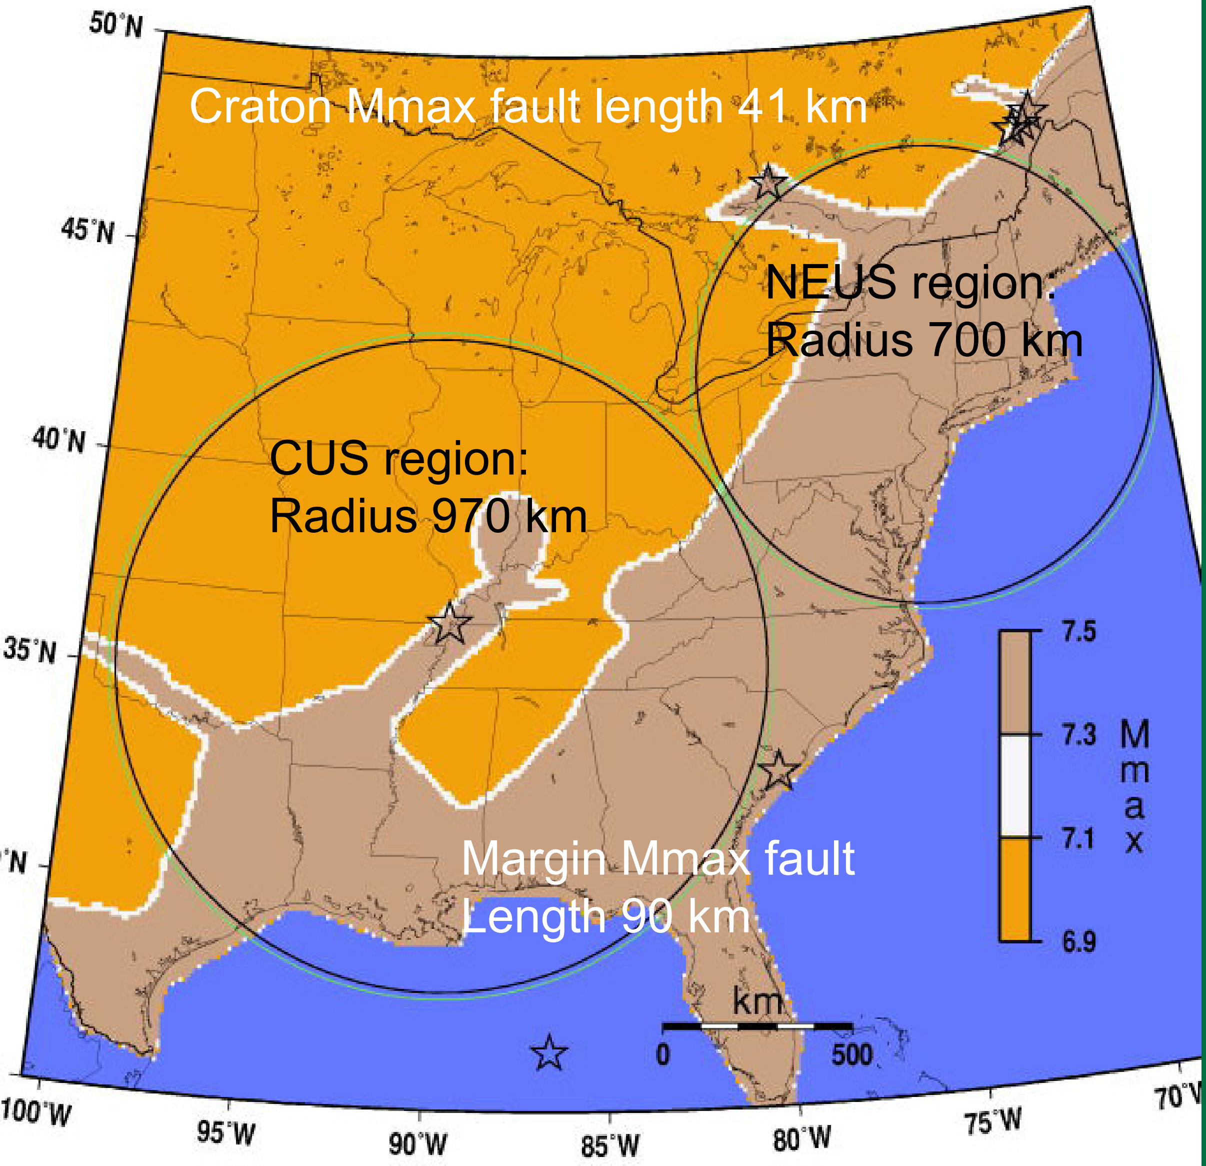

GMT Sep 3 08:19 MMax for CEUS. This is Mbmax to use with ../GR/mb_ab_7p0_7.5.asc conversion to Moment Mag

USGS

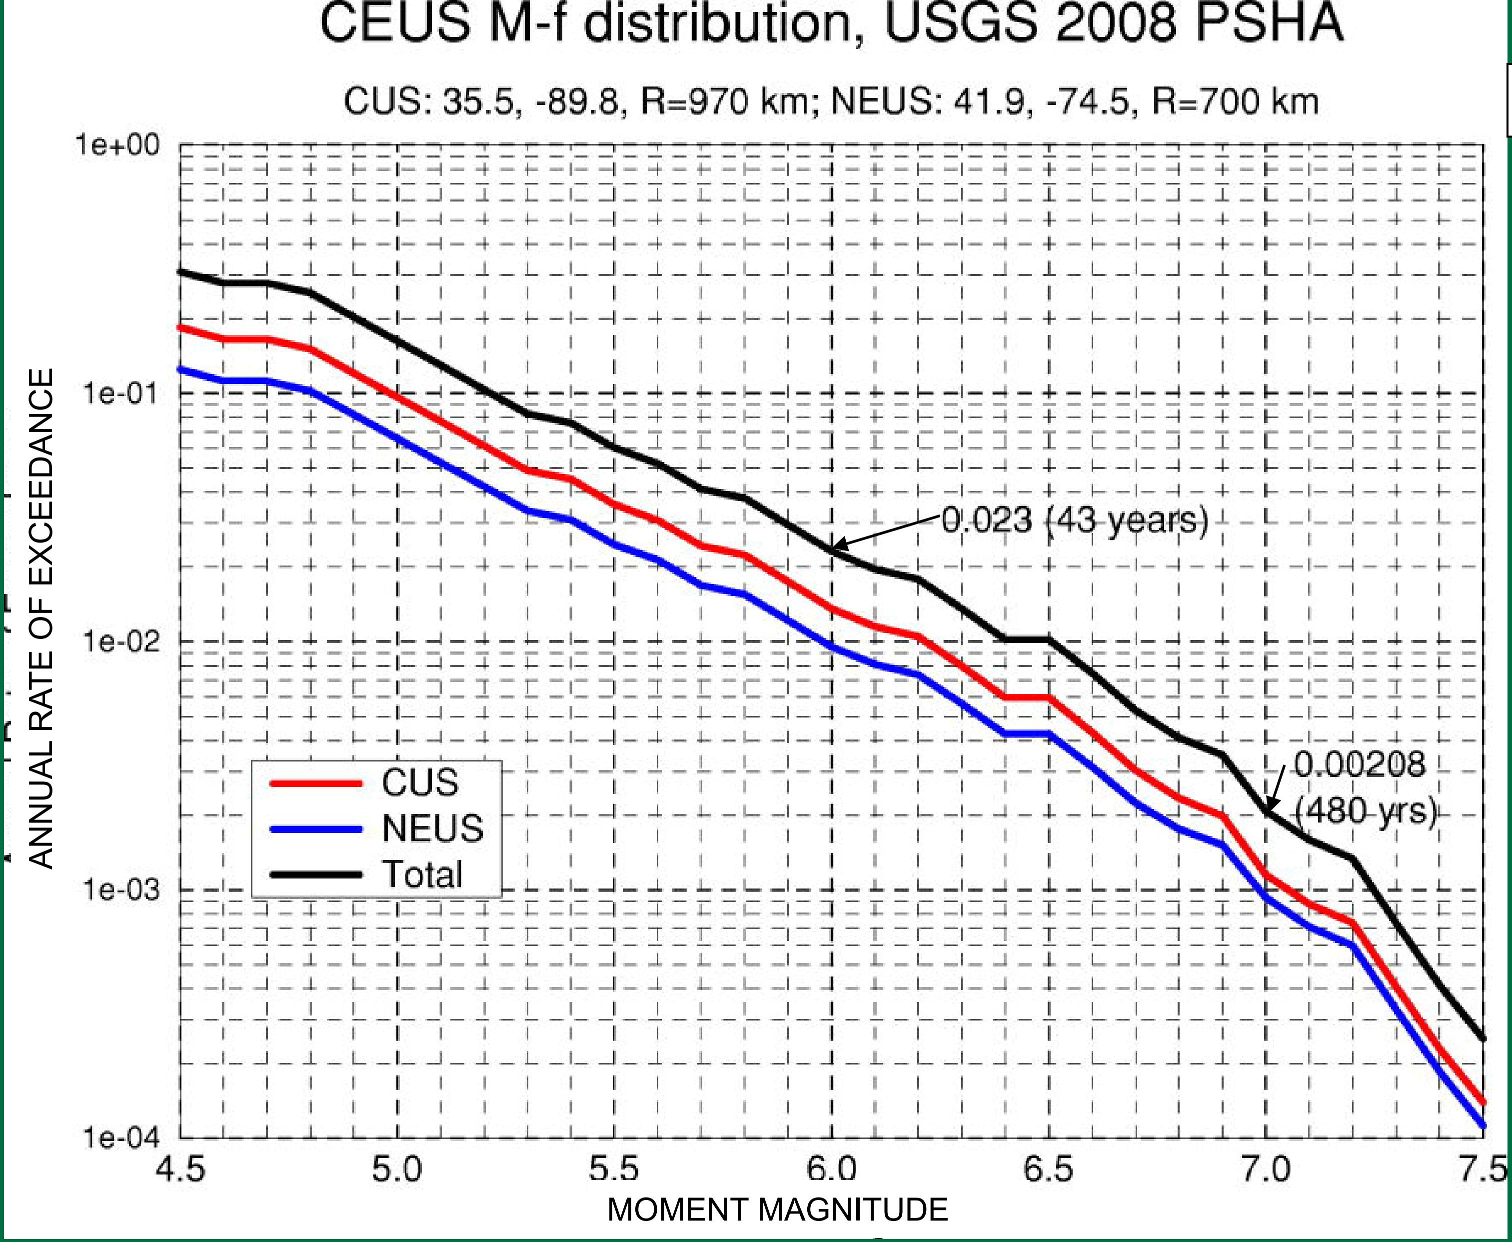
CEUS M-f distribution, USGS 2008 PSHA
CUS: 35.5, -89.8, R=970 km; NEUS: 41.9, -74.5, R=700 km
1e+00
1e-01
1e-02
1e-03
1e-04
ANNUAL RATE OF EXCEEDANCE
0.023 (43 years)
0.00208
(480 yrs)
CUS
NEUS
Total
4.5
5.0
5.5
6.0
6.5
7.0
7.5
MOMENT MAGNITUDE

Figure A9–6

Figure A9–7

How Many CEUS Eqs w/M>6 in Last 300 Years?

- Depends on who you ask. USGS ranked competing estimates, NCEER estimate came out on top of the heap.
- Zero earthquakes with m or M>6 in CUS or NEUS circles, according to declustered USGS catalog if you …
- Omit NMSZ mainshocks, foreshocks, and aftershocks of 1811-1812 (handled by fault files, not background).
- Omit 1843 and 1895 NMSZ earthquakes (all have NCEER M < 5.5, but Johnston and some others give M>6 estimates).
- Omit Cape Ann November 1755 earthquake. NCEER M is 5.8 (mb 6+, Ebel)
- Omit Charleston, SC mainshock & aftershocks (outside the study area). Also, handled by other files, not background.
- Omit St. Lawrence River-Charlevoix (M6+); 1935 Timiskaming (M6.1, Bent; MMI VII); Grand Banks (1929, M7.2) (outside study area).
- Omit 1897 Giles County, VA. Inferred M was 5.9. MMI VIII.

USGS

Mmax; Stars are m6+ epicenters, 1708 to 2008

Figure A9–8

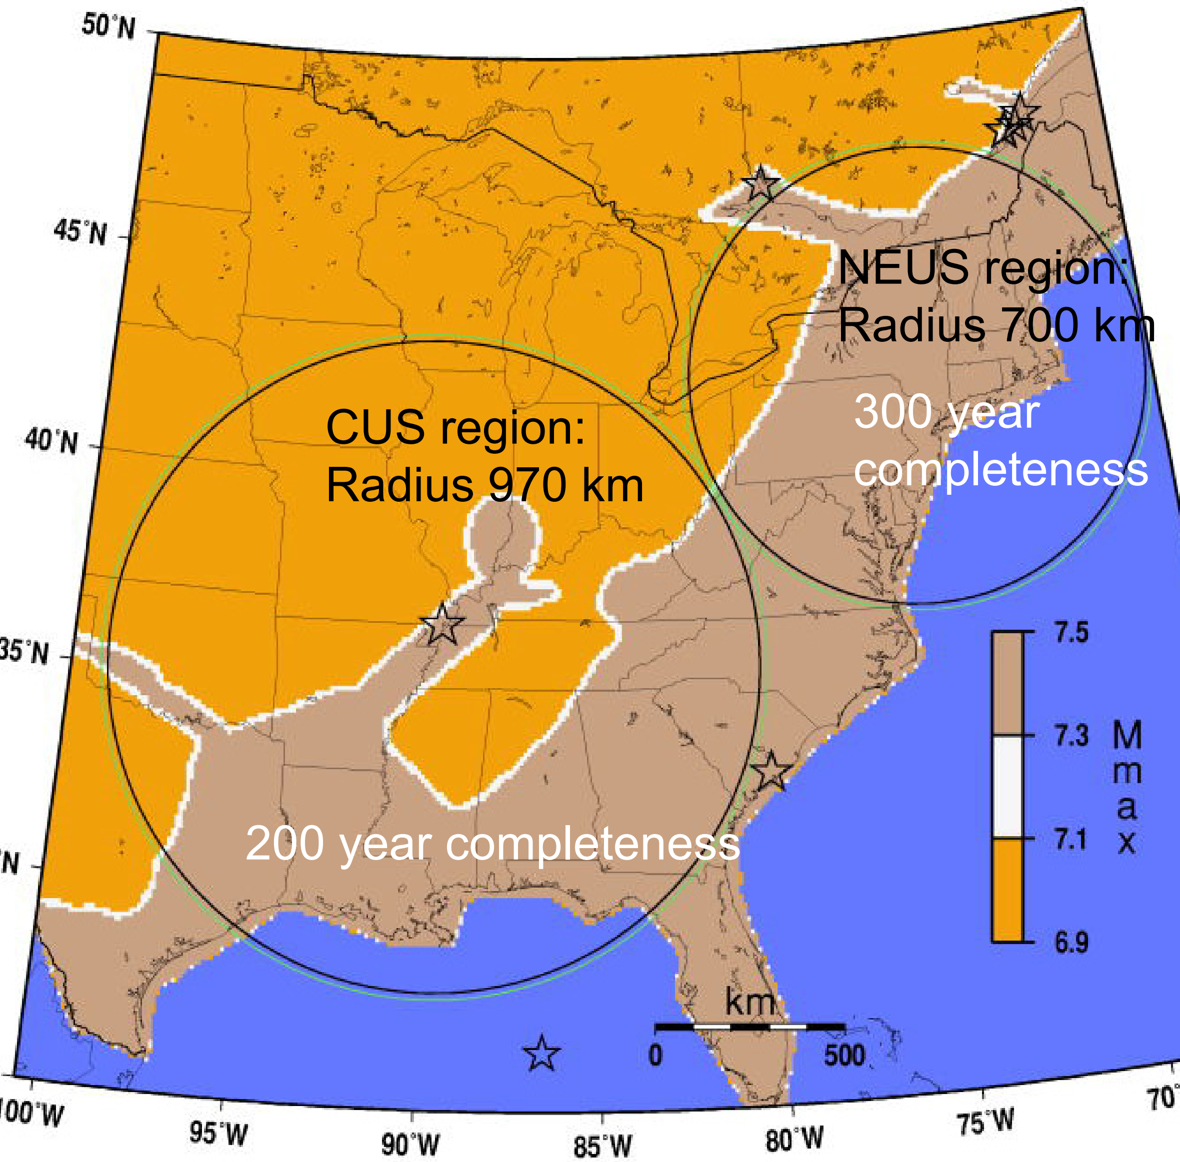

GMT Sep 3 08:19 MMax for CEUS. This is Mbmax to use with ../GR/mb_ab_7p0_7.5.asc conversion to Moment Mag

USGS

Probability Statement About M>6 Historical Earthquakes

Figure A9–9

- Let mu1 = Expected number of earthquakes having M>=6 in random 300 year period in NEUS circle. Mu1=2.85
- Let mu2=Expected number of earthquakes with M>=6 in a random 200 year period in the CEUS circle. Mu2=2.72
- Let mu = mu1 + mu2.
- Poisson Pr[0 M6+ | mu=5.57] = 0.0038
- Pr[1 or less | mu=5.57] = 0.025

USGS

Is the significance test any good?

Figure A9–10

- Only if observed N < 2.
- Some will claim that omitting Cape Ann, Timiskaming, and Charlevoix is unfair.
- Some NMSZ “aftershocks” may be independent sources. 1843 & 1895 M(NCEER)=5.4; other catalogs say M>=6 for these two earthquakes.
- Completeness issues will be argued. Large % of these regions was raw frontier for substantial part of the time intervals.
- Last 200 or 300 years may not be typical.

USGS

Even if test is good

Figure A9–11

- Do not need to reduce MMax to get a non-significant test result
- Could reduce predicted earthquake rates in the M6 to M7.5 range by a factor of 2 or so
- 2007 USGS PSHA in California did reduce background-source rates for M>6.5 sources to get better agreement with historical earthquake rate

USGS

M>7 Paleoearthquakes 6000 yrs

Figure A9–12

- 1 earthquake: Vincennes (M7+, Obermeier). 6100 years
- Let mu=expected number of M7+ background earthquakes in the last 6000 years in either CUS or NEUS. Omit NMSZ mainshocks, Meers, OK, and South Carolina coastal plain, which are handled by other files.
- From a previous slide, mu=12.5.
- Pr[1 or less M7+ source | mu=12.5] = 0.000
- No CEUS NPP site applications indicate evidence for Holocene earthquakes with M>=7 near the site.
- May want to include 1663 Charlevoix earthquake (M approx 7).
- Pr [2 or less | mu=12.5] = 0.0003
- We need 6 or more M7+ paleoearthquakes to produce a non-significant Poisson test when mu=12.5 (5 gives 0.015 p)

USGS

Summary

Figure A9–13

- Initial look at these two CEUS regions shows that expected numbers of earthquakes of M>6 or M>7 seem to exceed observations by a significant amount.

- Similar tests should be attempted after resolving some magnitude-uncertainty issues, trying better-researched completeness period estimates, and so on.

Figure A10–1

GPS Constraints on Seismic Hazard in Continental Intraplate Regions: Eastern Canada Example

S. Mazzotti, J. Henton, T. James, J. Adams

Geological Survey of Canada & Geodetic Survey Division, Natural Resources Canada

JOURNAL OF GEOPHYSICAL RESEARCH, VOL. 110, B11301, doi:10.1029/2004JB003590, 2005

(2) Crustal Strain from GPS

Figure A10–2

Regional strain =
E-W shortening
(agrees with earthquake mechanisms)

Strain rates **1-4 x 10^{-9} /yr**

<=>

Convergence across St Lawrence **0-1 mm/yr**

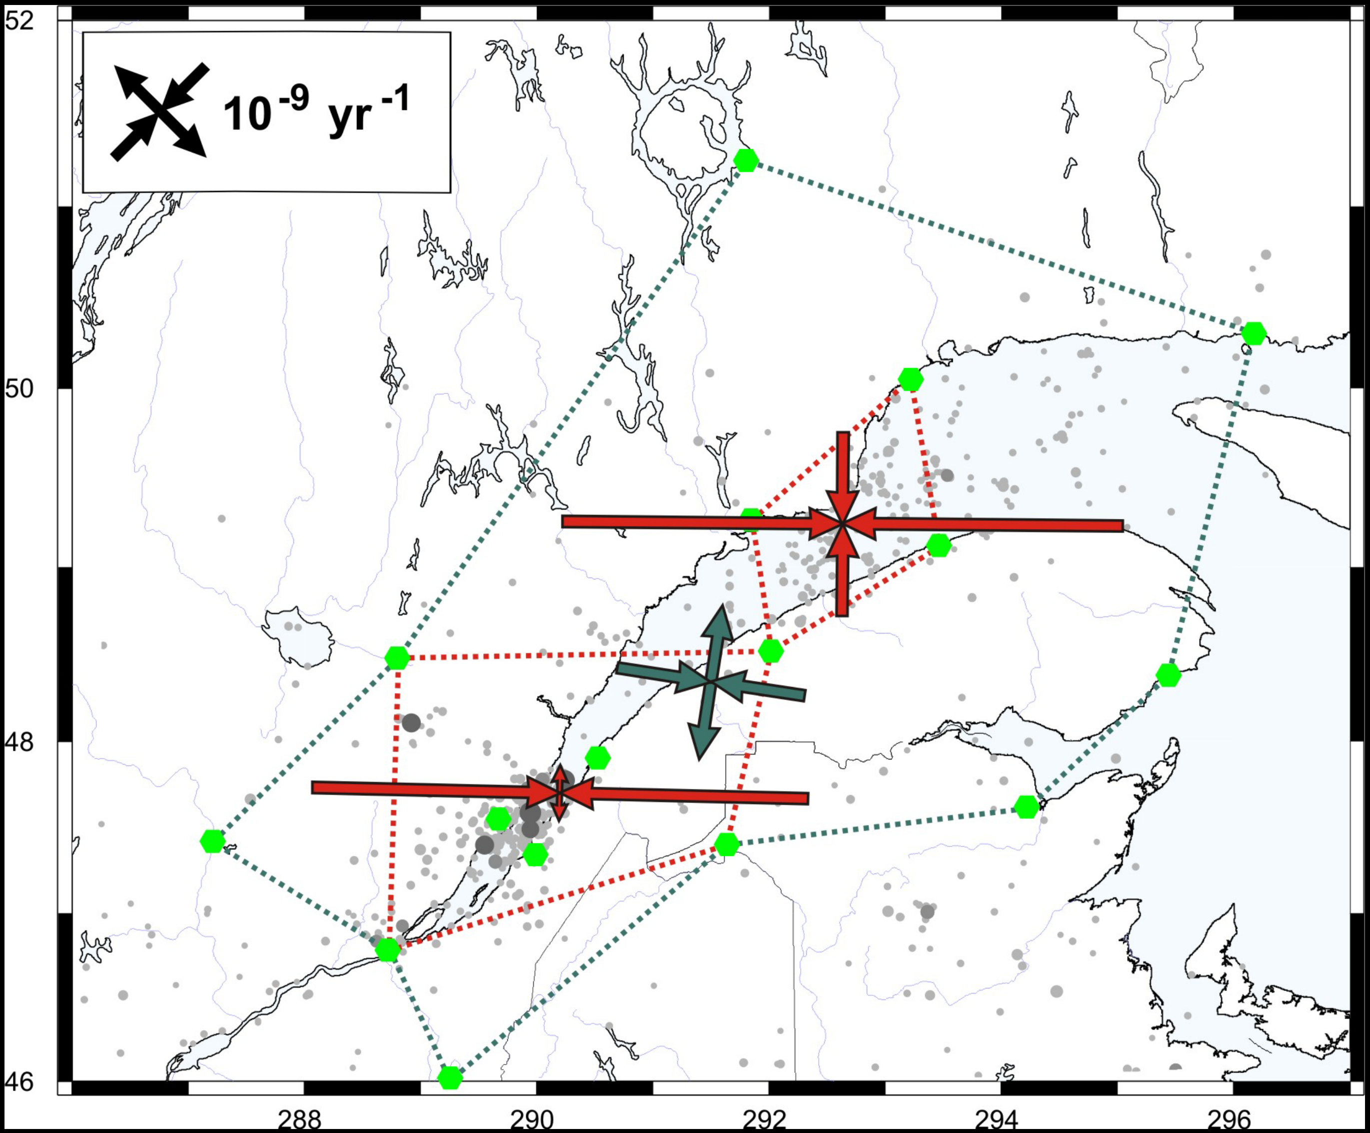

Figure A10–3

(3) GPS & Seismic Hazard

Assumptions for integration of GPS strain rates in hazard:

1) Strain model:
    GPS α Seismic strain release

2) Earthquake statistics:
    GR recurrence / characteristic

3) Seismic thickness:
    maximum depth / effective thickness

Magnitude & Recurrence of large earthquakes

(3) GPS & Seismic Hazard

Figure A10–4

GPS convergence rate:
~0.7 +/- 0.4 mm/yr

Gutenberg-Richter distribution up to Mx

Charlevoix: **Mx = 7.8**
(std. 7.2-8.5)

Lower St Lawrence: **Mx = 7.3**
(std. 5.5-8.3)

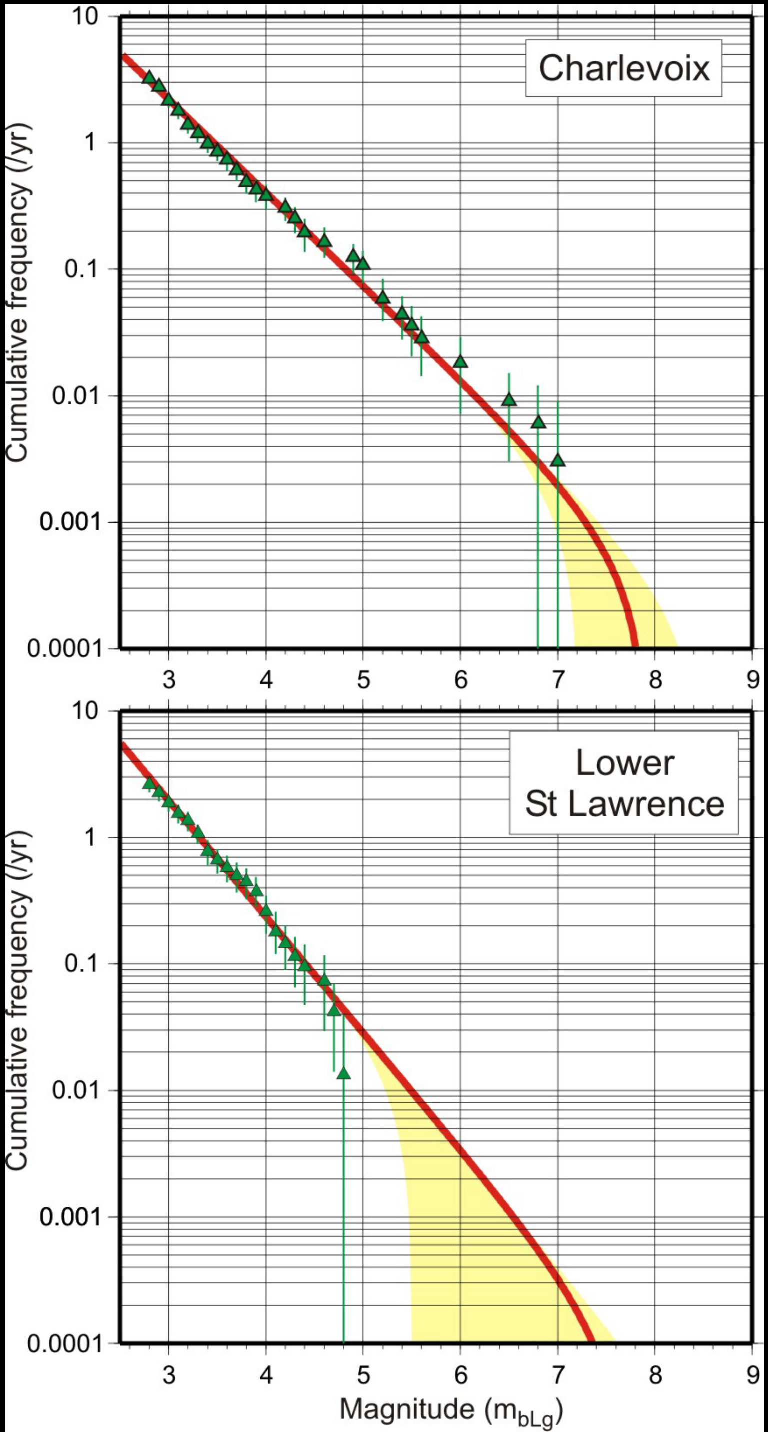

(3) GPS & Seismic Hazard

Figure A10–5

GPS convergence rate:
~0.8 +/- 0.5 mm/yr

Characteristic earthquakes

Charlevoix: **M7 = 1/200 yr**
(std. 100-400 yr)

Lower St Lawrence: **M7 = 1/1000 yr**
(std. 400-20,000 yr)

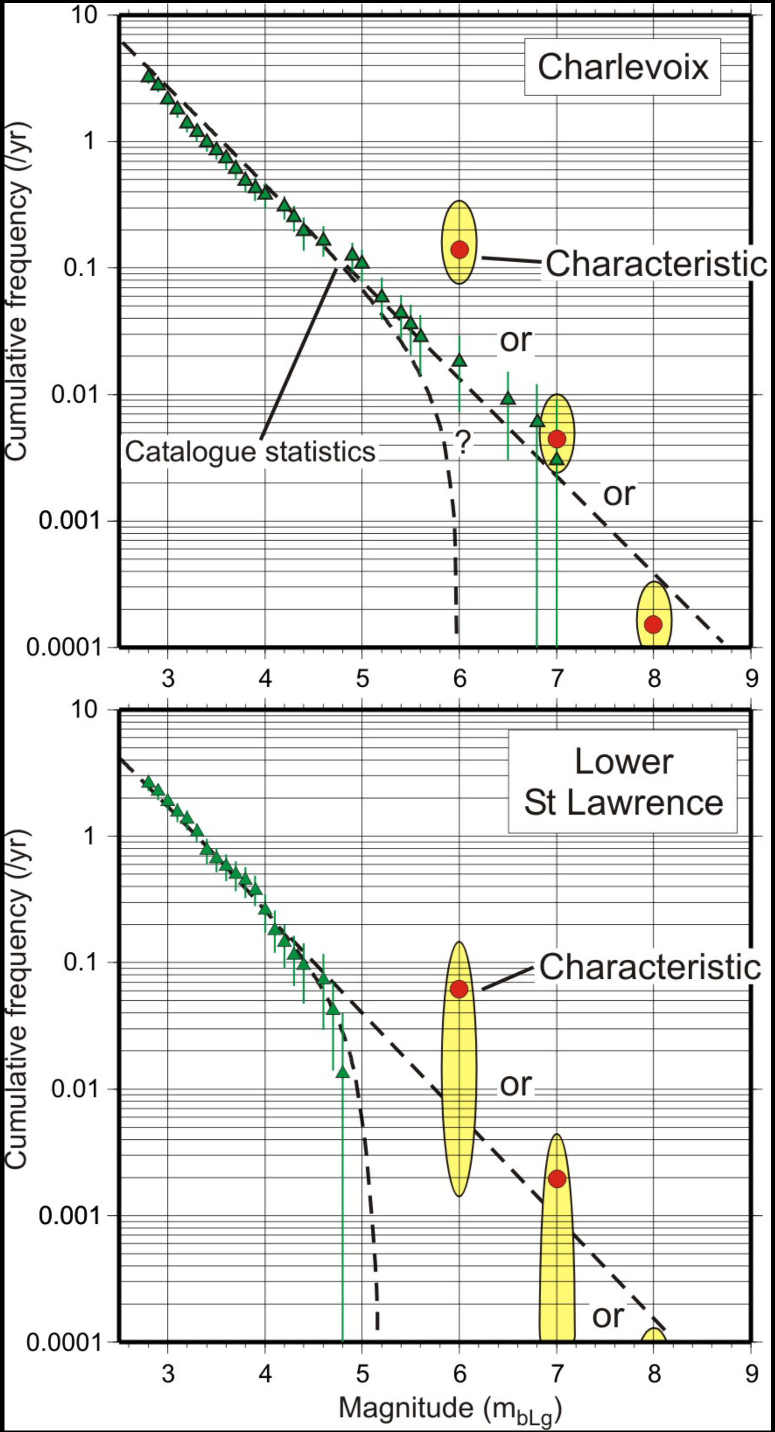

Figure A10–6

Conclusions

In Charlevoix and Lower St Lawrence,
GPS strain rates agree well with seismic strain rates

CHV: 0.7 (+/- 0.4) **mm/yr** <=> **1.0** (+/- 0.5) **mm/yr**

BSL: 0.2 (+/- 0.6) **mm/yr** <=> **0.2** (+/- 0.3) **mm/yr**

2) GPS data constrain the recurrence and magnitude of large earthquakes

CHV: Mx ~7.8 (+/- 0.6) <=> **T(M7) ~ 200 – 1000 yr**

BSL: Mx ~7.3 (+/- 1 ?) <=> **T(M7) >> 400 yr**

Figure A11–1

Mx in Canada's 4th Generation seismic hazard model

John Adams

Presentation for USGS Mmax meeting Golden 2008 09 09

4th Generation model

Developed ~1994-1997

Finalized 2003

Implemented in 2005 code

Documentation

http://earthquakescanada.nrcan.gc.ca/hazard/OF4459/index_e.php

Magnitude-recurrence for eastern Canada

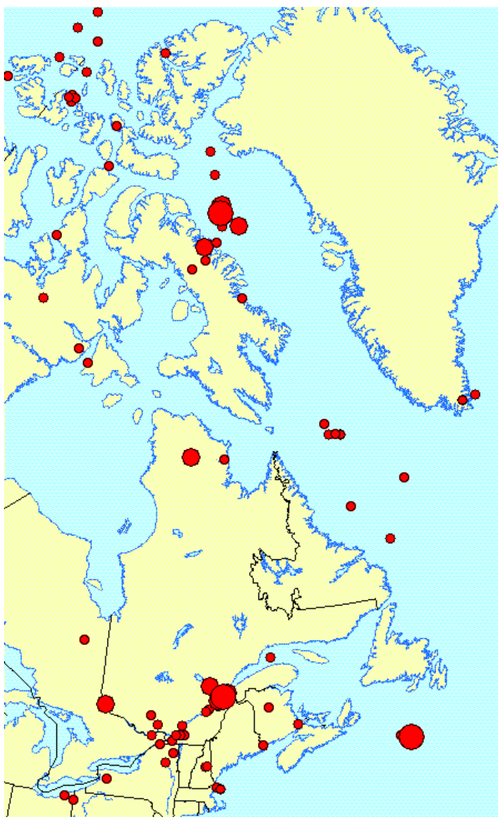

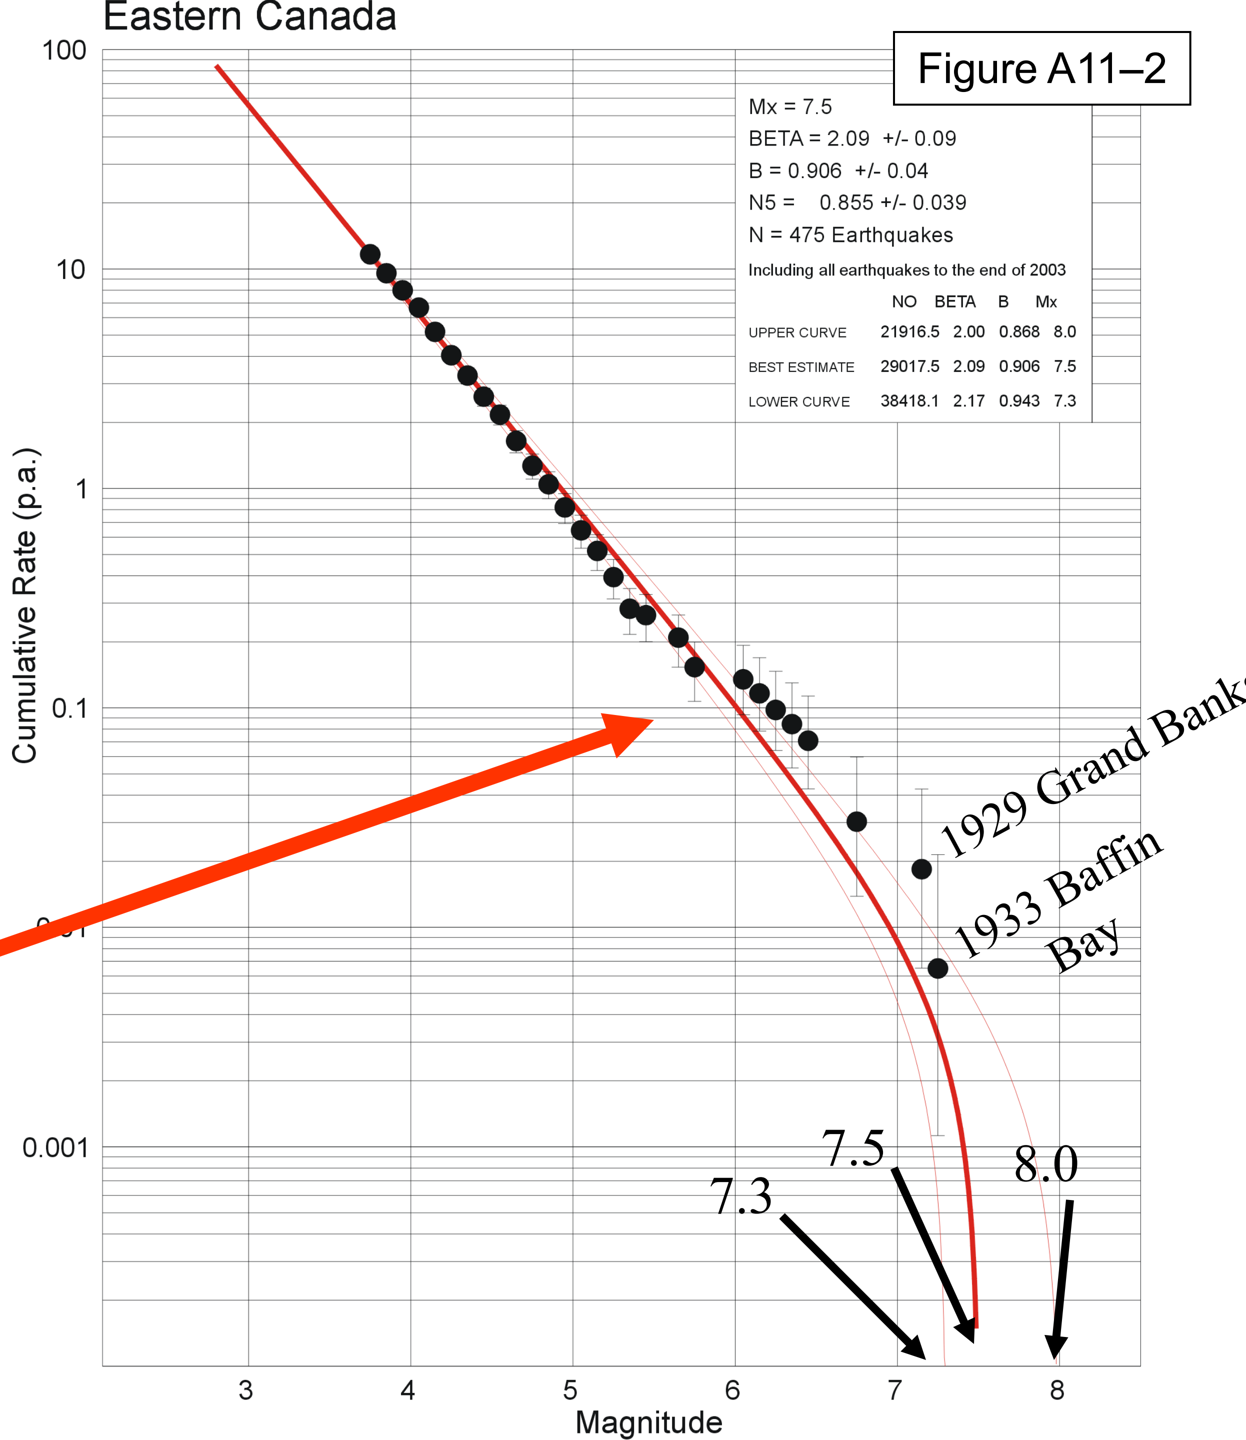

Stable Craton Core (SCC) rates and Mmax

Figure A11–3

Fenton and Adams, 1997; Fenton et al 2006

Places like Canadian Shield!

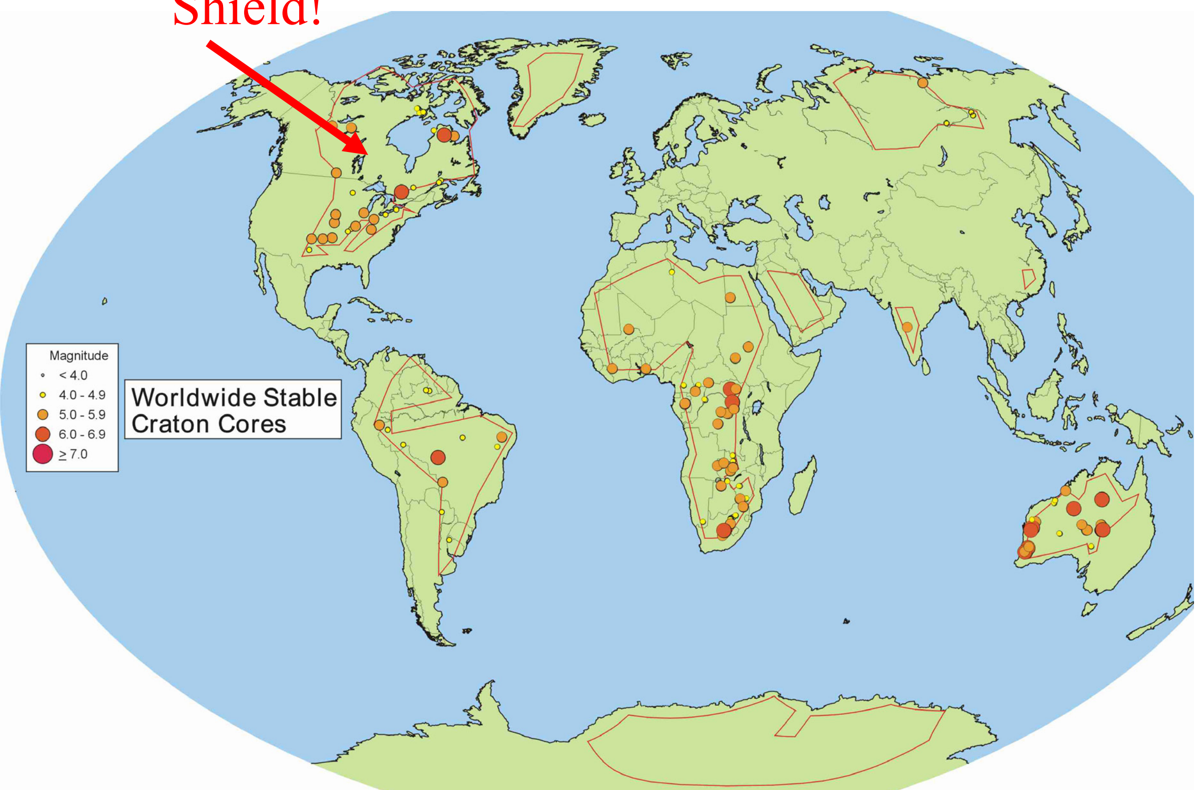

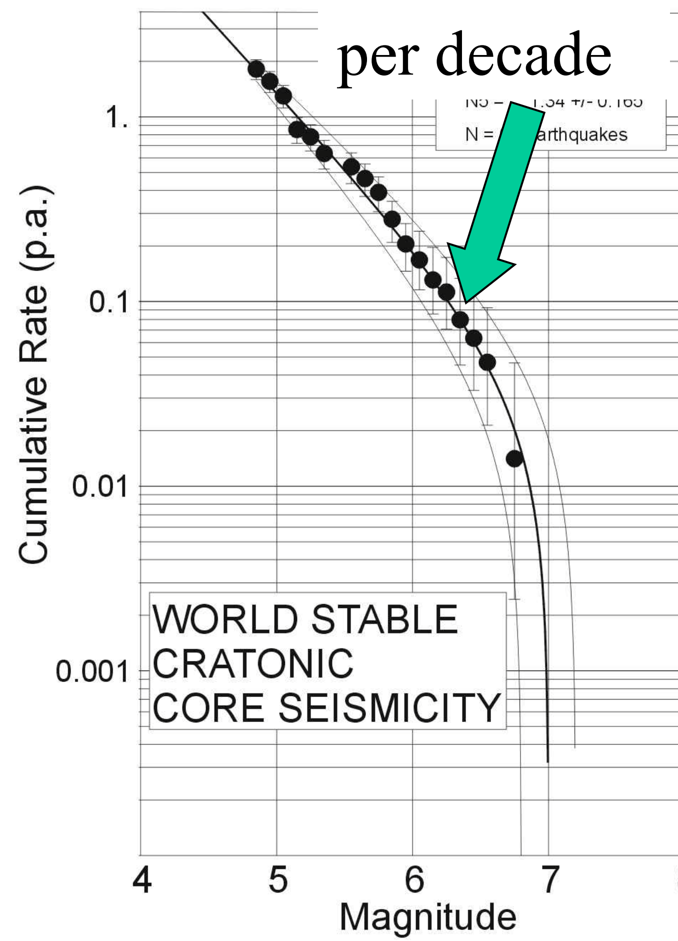

http://earthquakescanada.nrcan.gc.ca/hazard/2006/2006FentonAdamsHalchukGEGE.pdf

Figure A11–4

Floor Hazard estimates - Three rates to capture uncertainty:

A) Global earthquake activity of continental shields **Wt = 0.4**

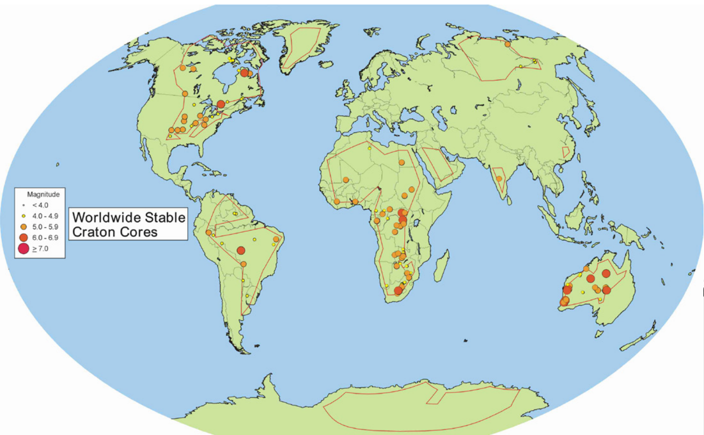

B) Observed North American shield activity rate **Wt = 0.4**

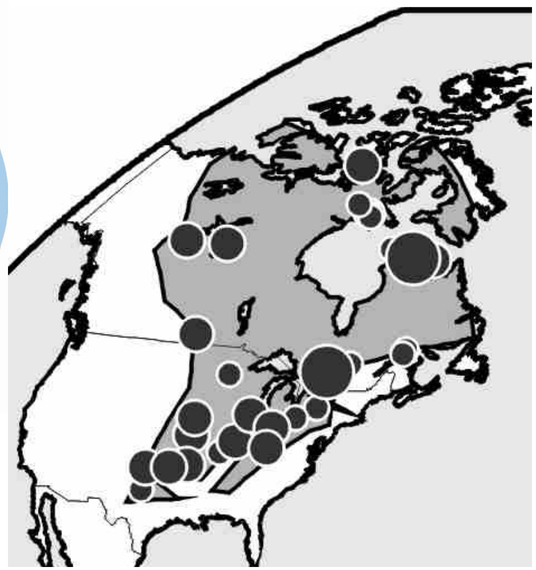

C) Rate for central Canada not in a source zone **Wt = 0.2**

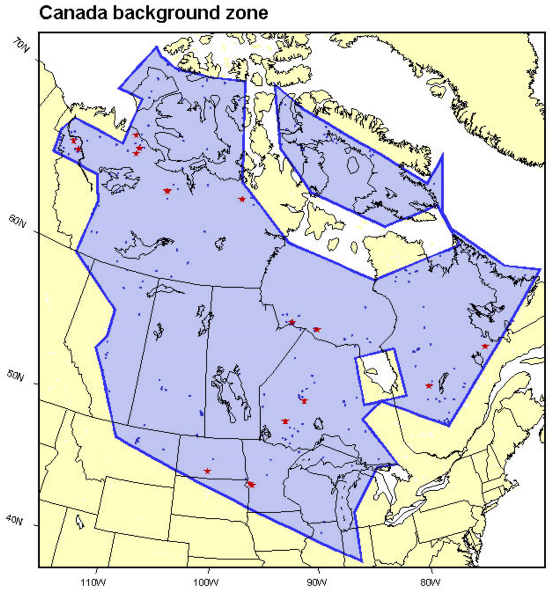

Then, seismic hazard computed for centre of large zone

Uniform Hazard Spectra

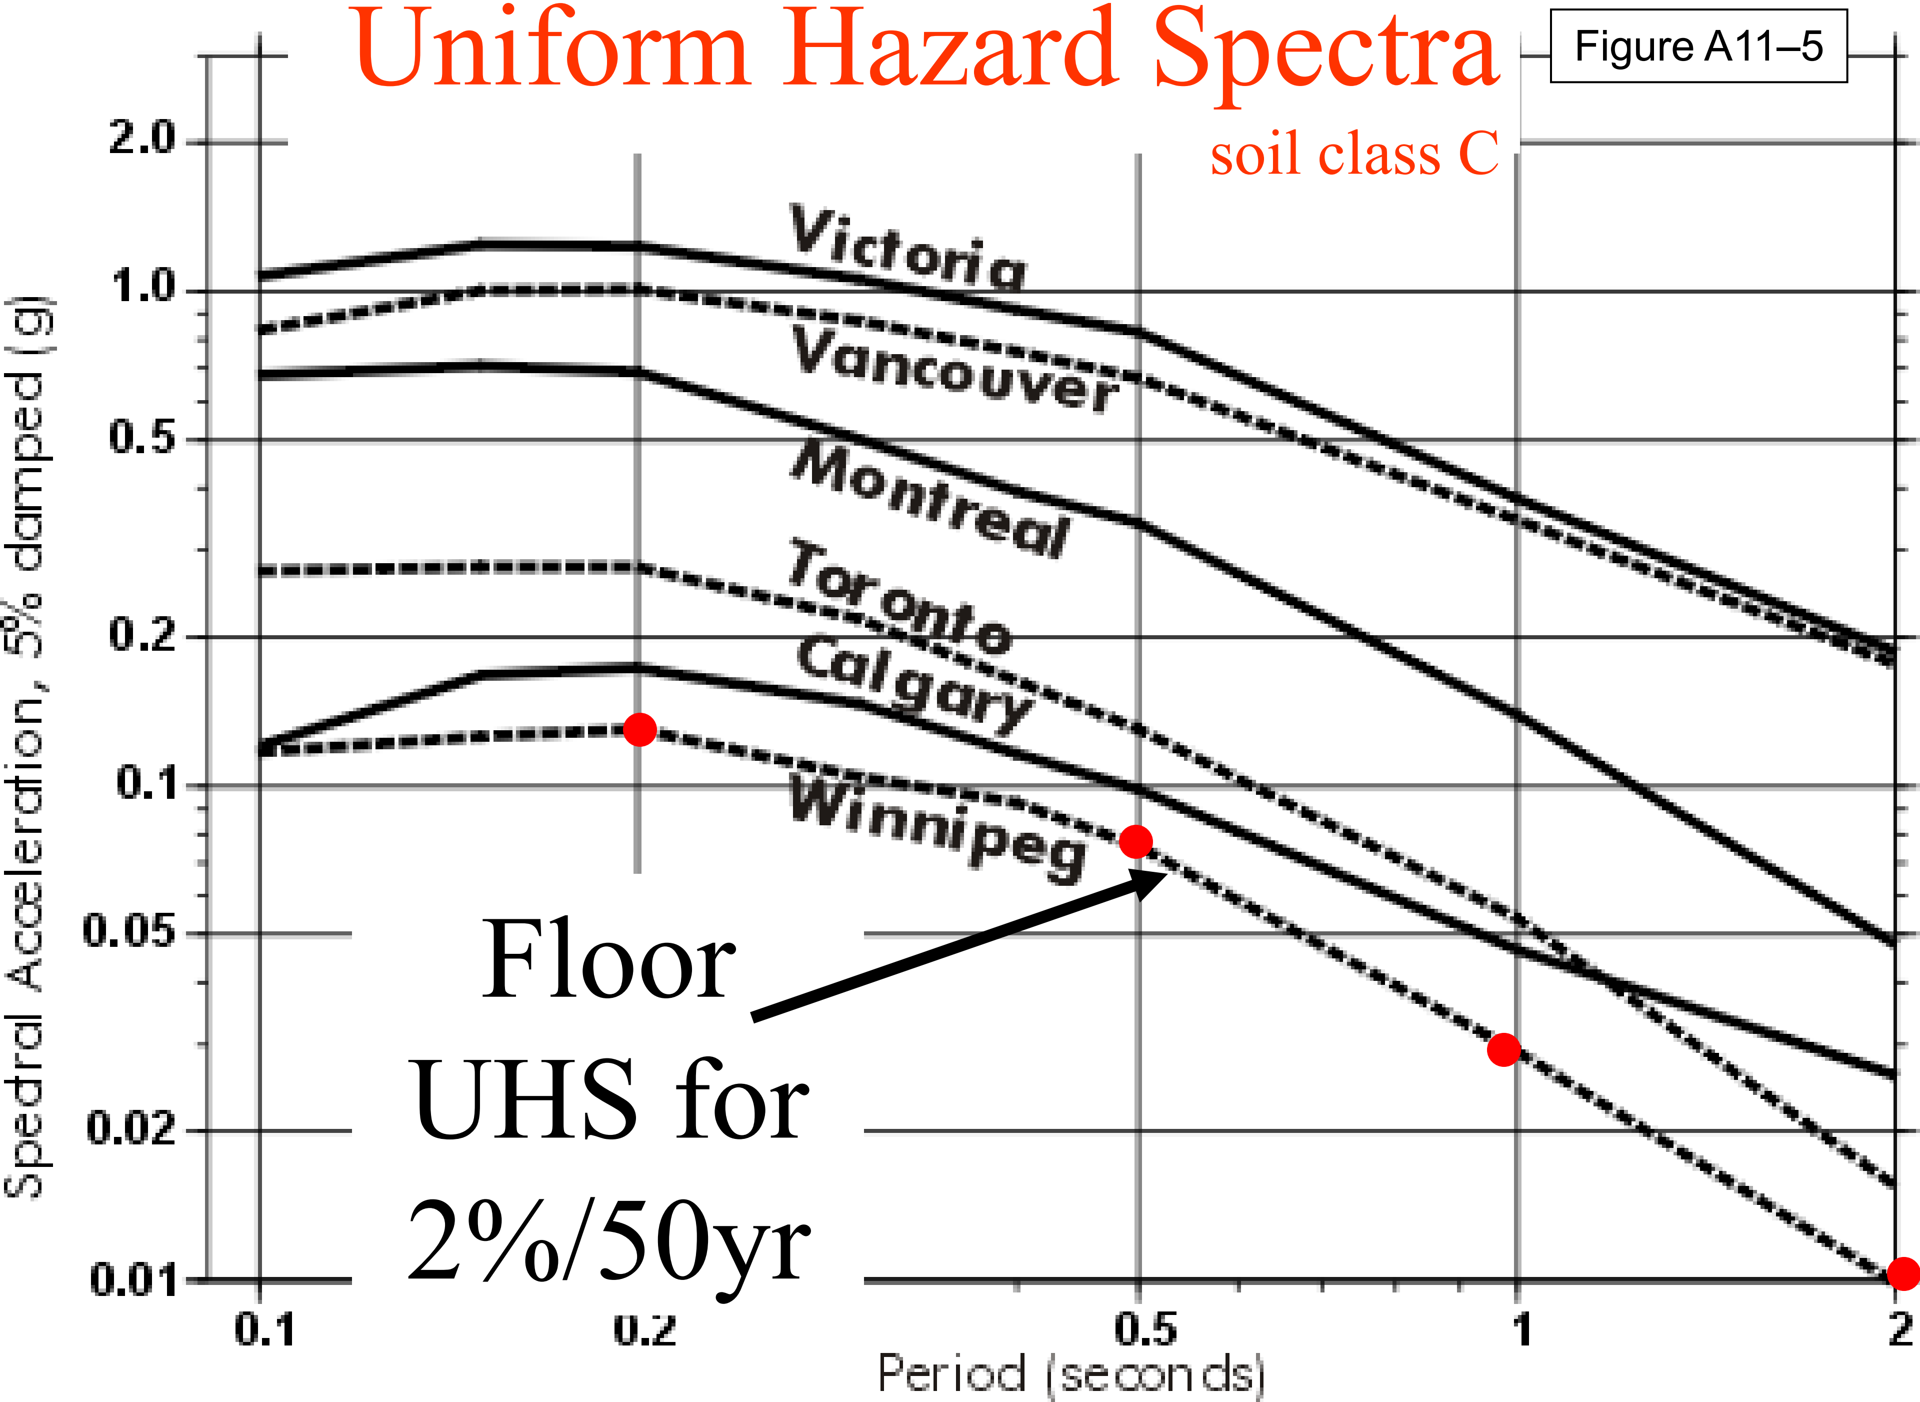

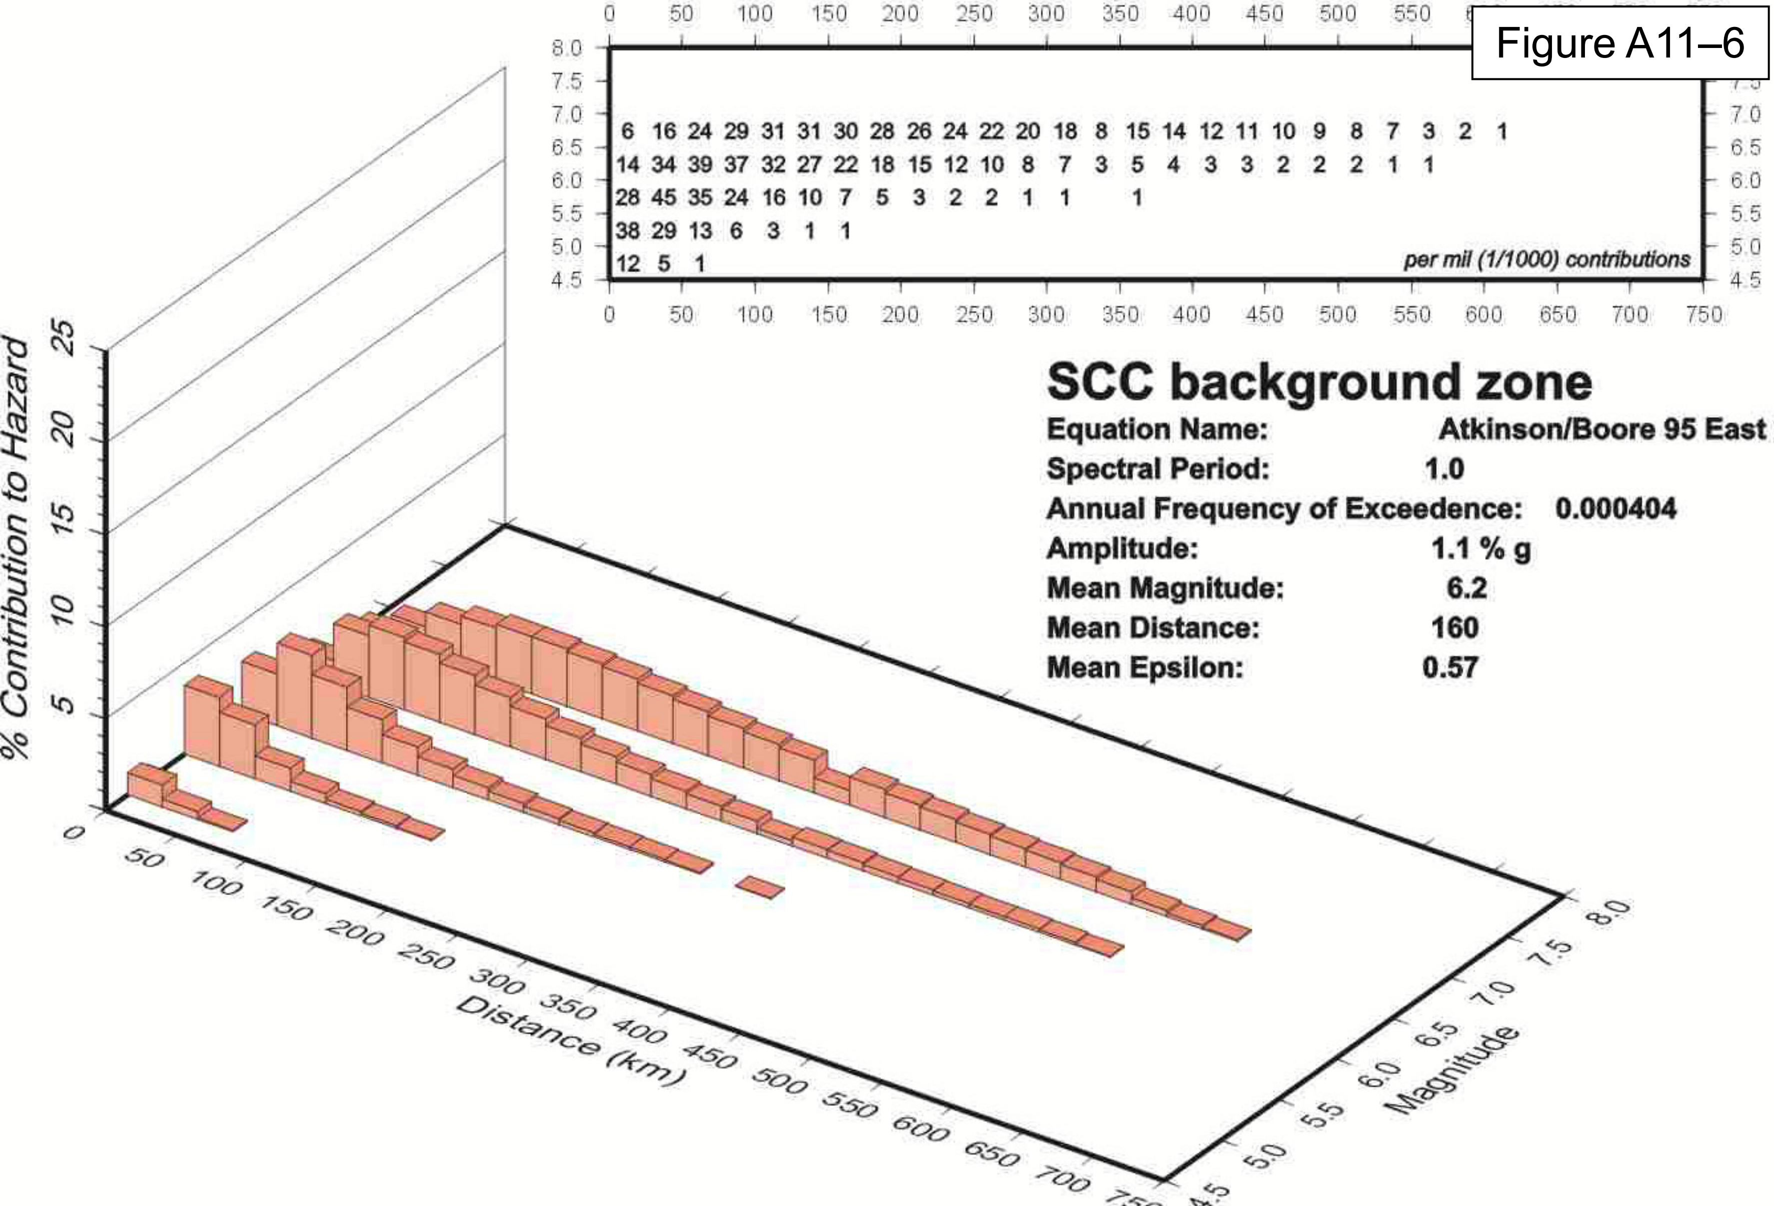

SCC background zone

Equation Name: Atkinson/Boore 95 East Hard Rock
Spectral Period: 1.0
Annual Frequency of Exceedence: 0.000404
Amplitude: 1.1 % g
Mean Magnitude: 6.2
Mean Distance: 160
Mean Epsilon: 0.57

This is a 1999 deaggregation using EZ-Frisk. Details may have changed, but pattern will be the same.

Figure A11–7

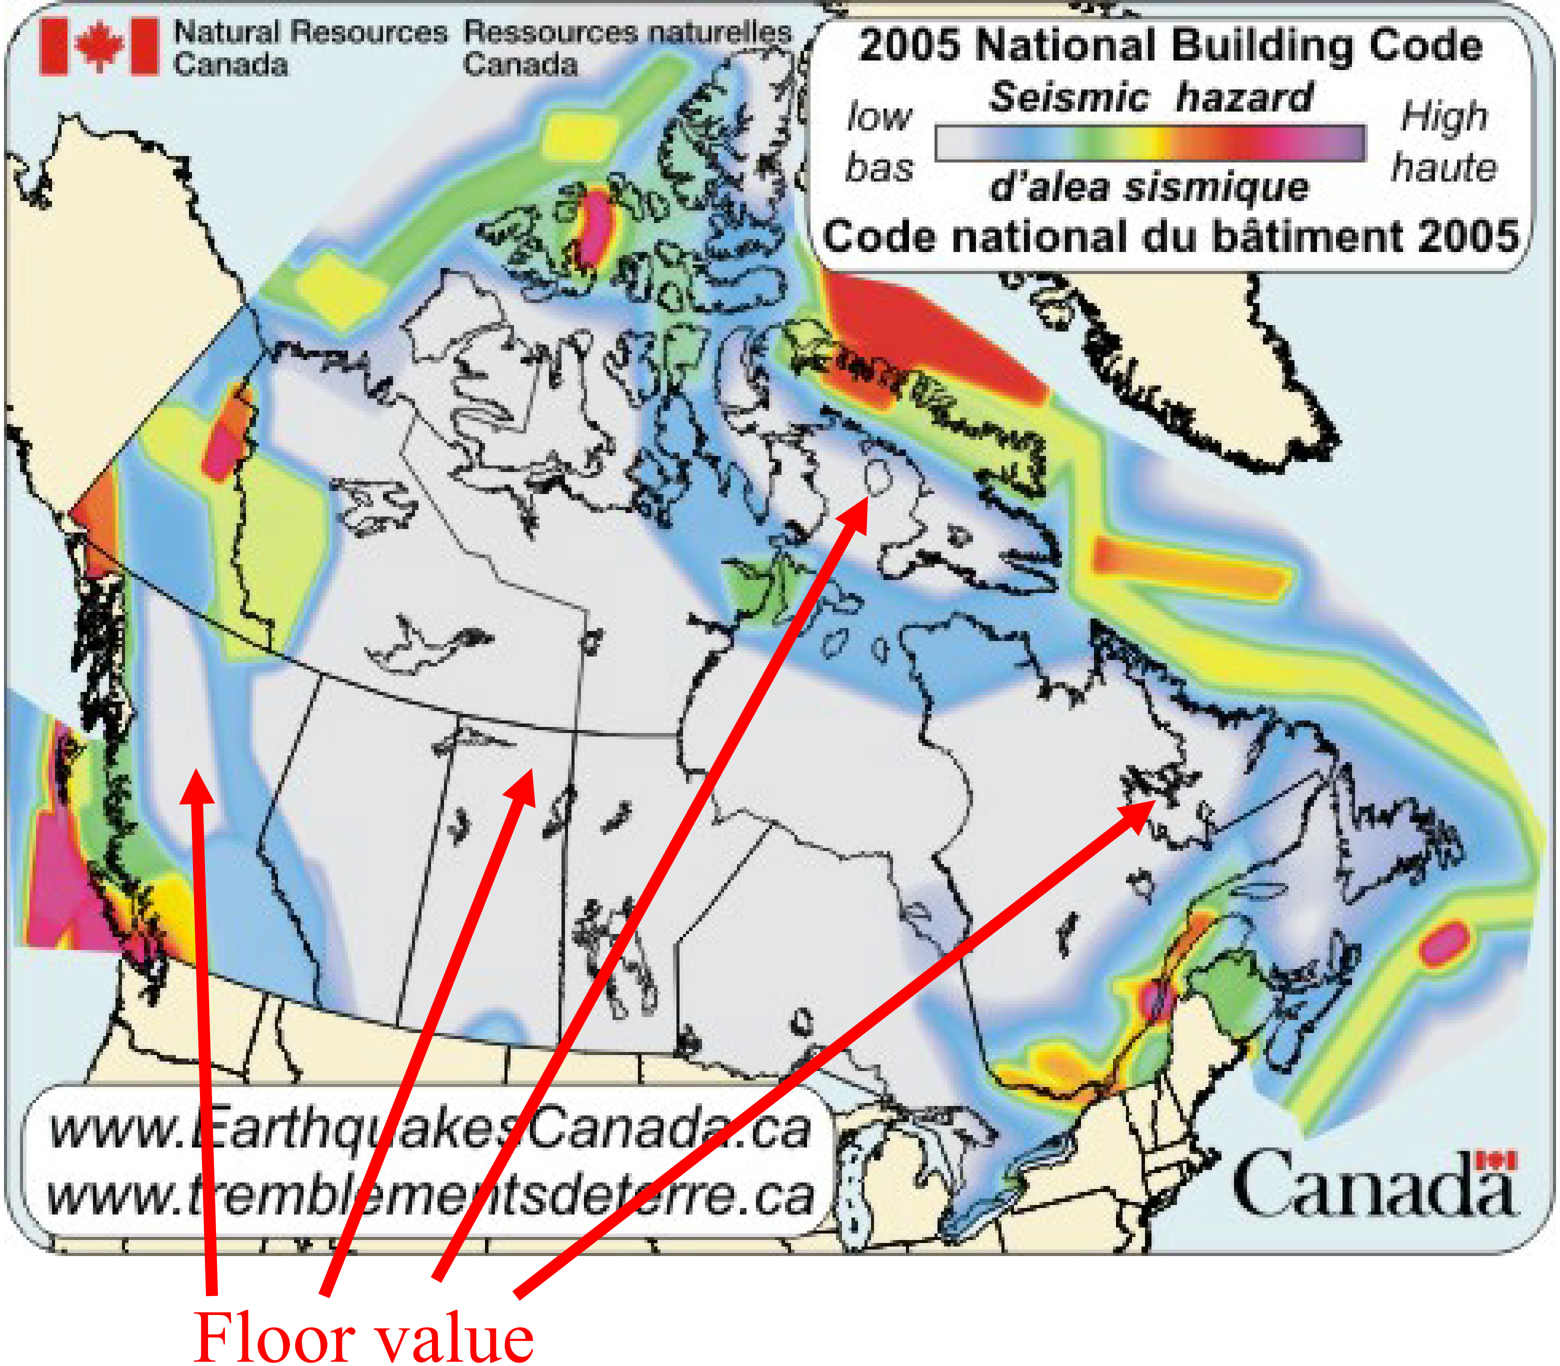
Natural Resources Canada
Ressources naturelles Canada
2005 National Building Code
Seismic hazard
low
bas
High
haute
d'alea sismique
Code national du bâtiment 2005
www.EarthquakesCanada.ca
www.tremblementsdeterre.ca
Canada
Floor value

Mesozoic rifted margin

Mobs ~7.4

Weighted branches
best, upper, lower
7.5 8.0 7.3
0.68 0.16 0.16

Plenty of potential large faults -> M8

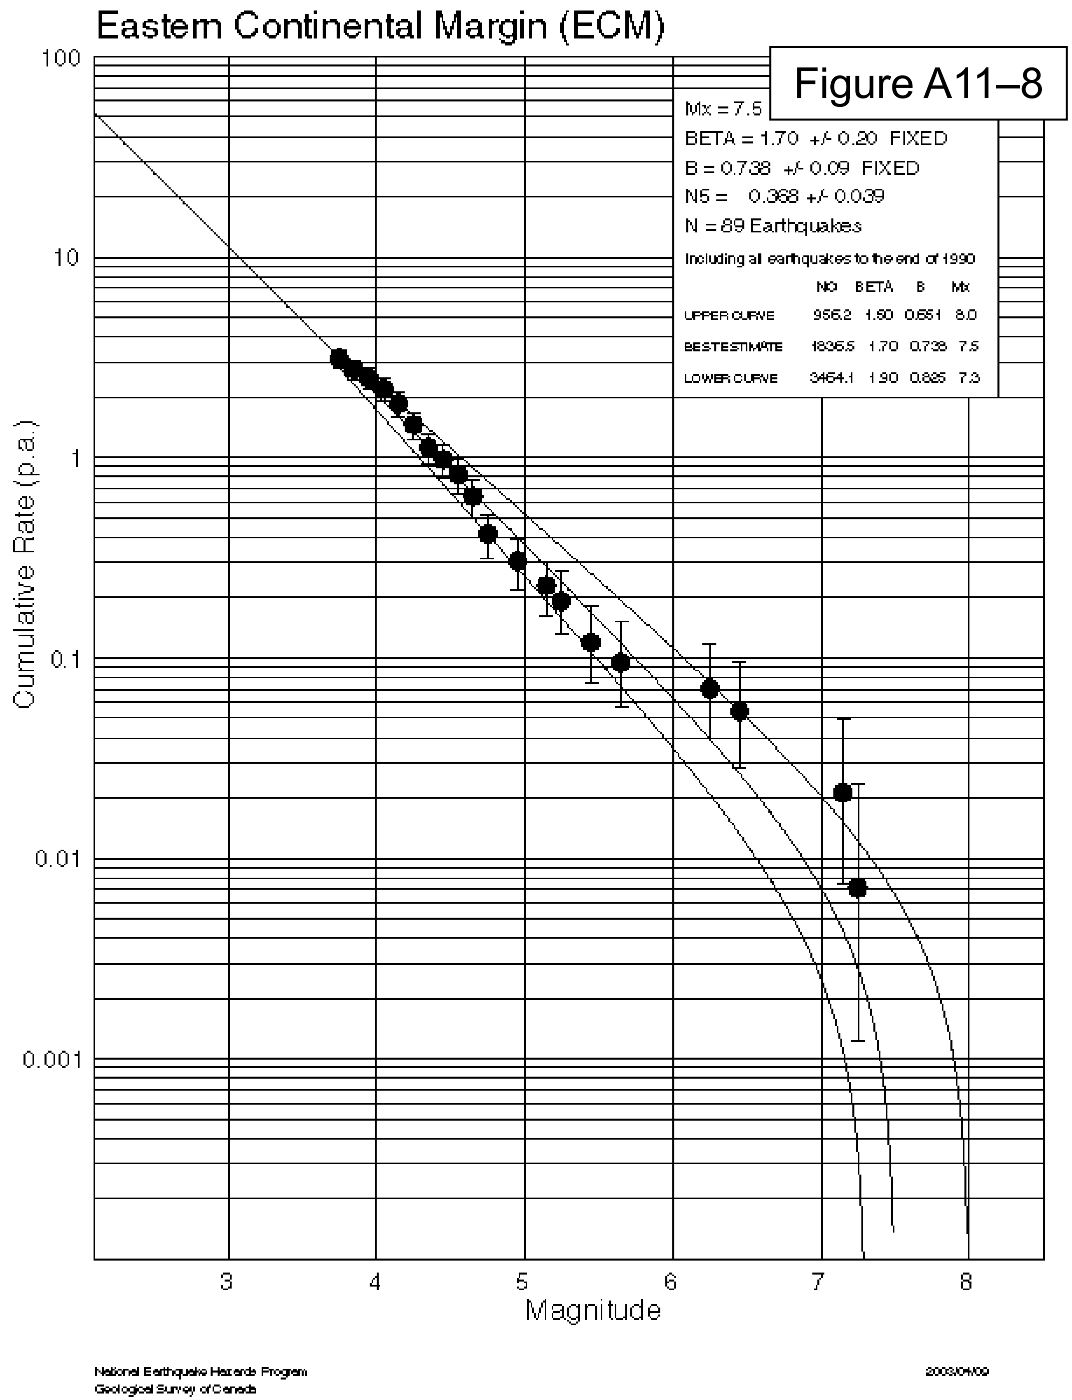

Paleozoic rifted margin

Mobs ~7.0

Weighted branches

best, upper, lower

7.5 7.7 7.2

0.68 0.16 0.16

Enough potential large faults

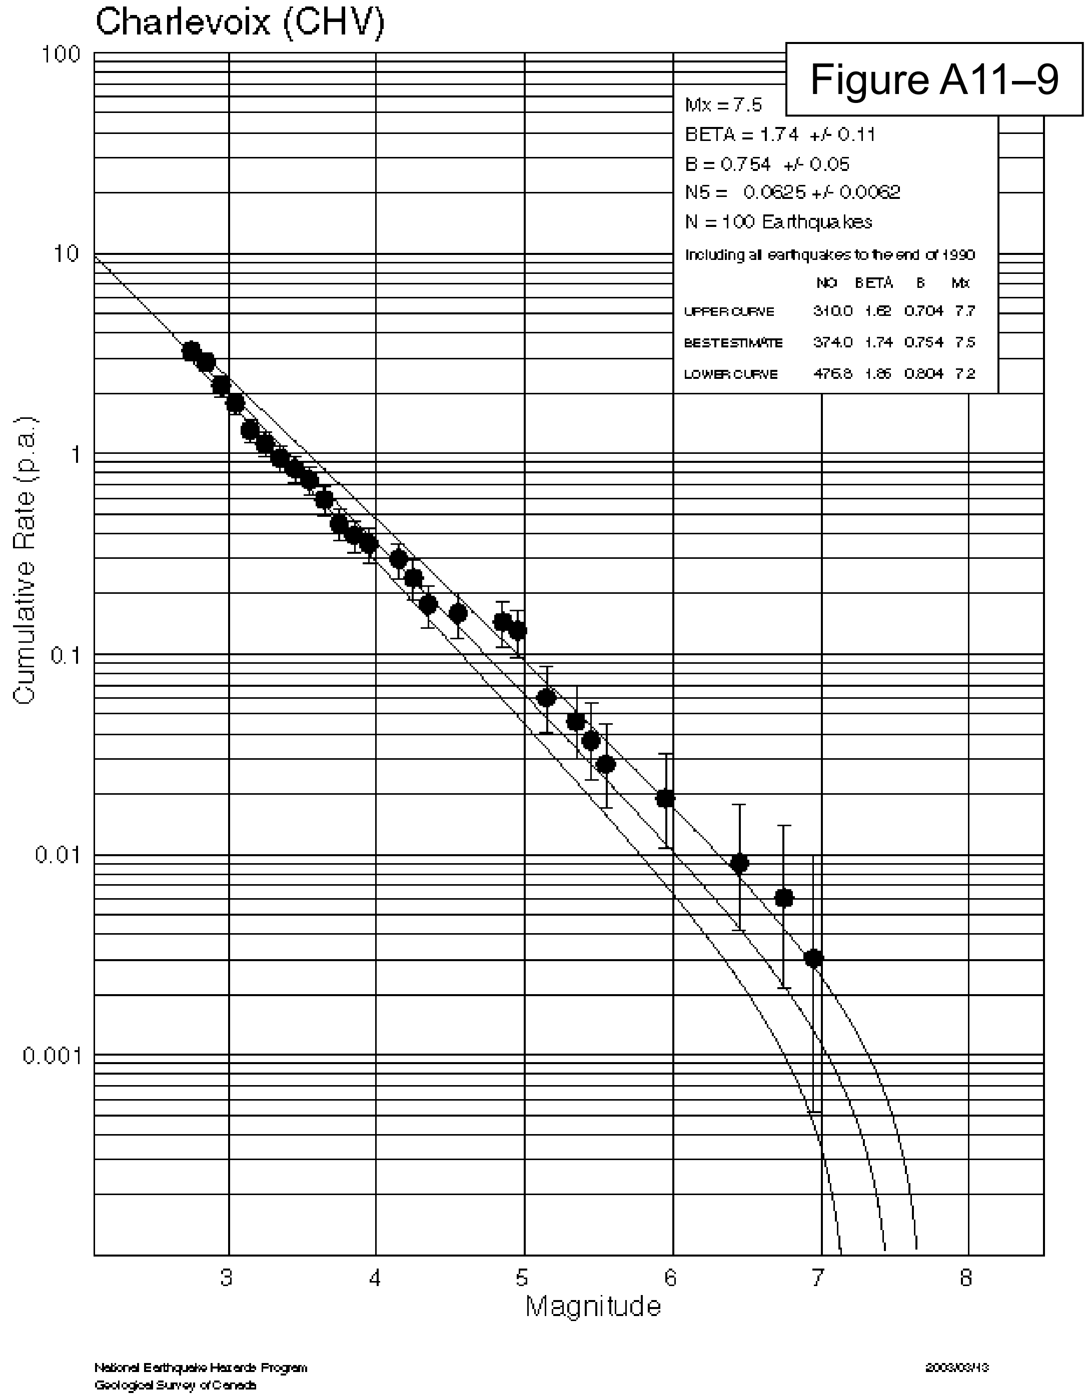

Interior

?slightly extended

Mobs ~5.0

Weighted branches
best, upper, lower
7.0 7.2 6.8
0.68 0.16 0.16

potential large faults?

insensitive to Mmax

consistent with SCC

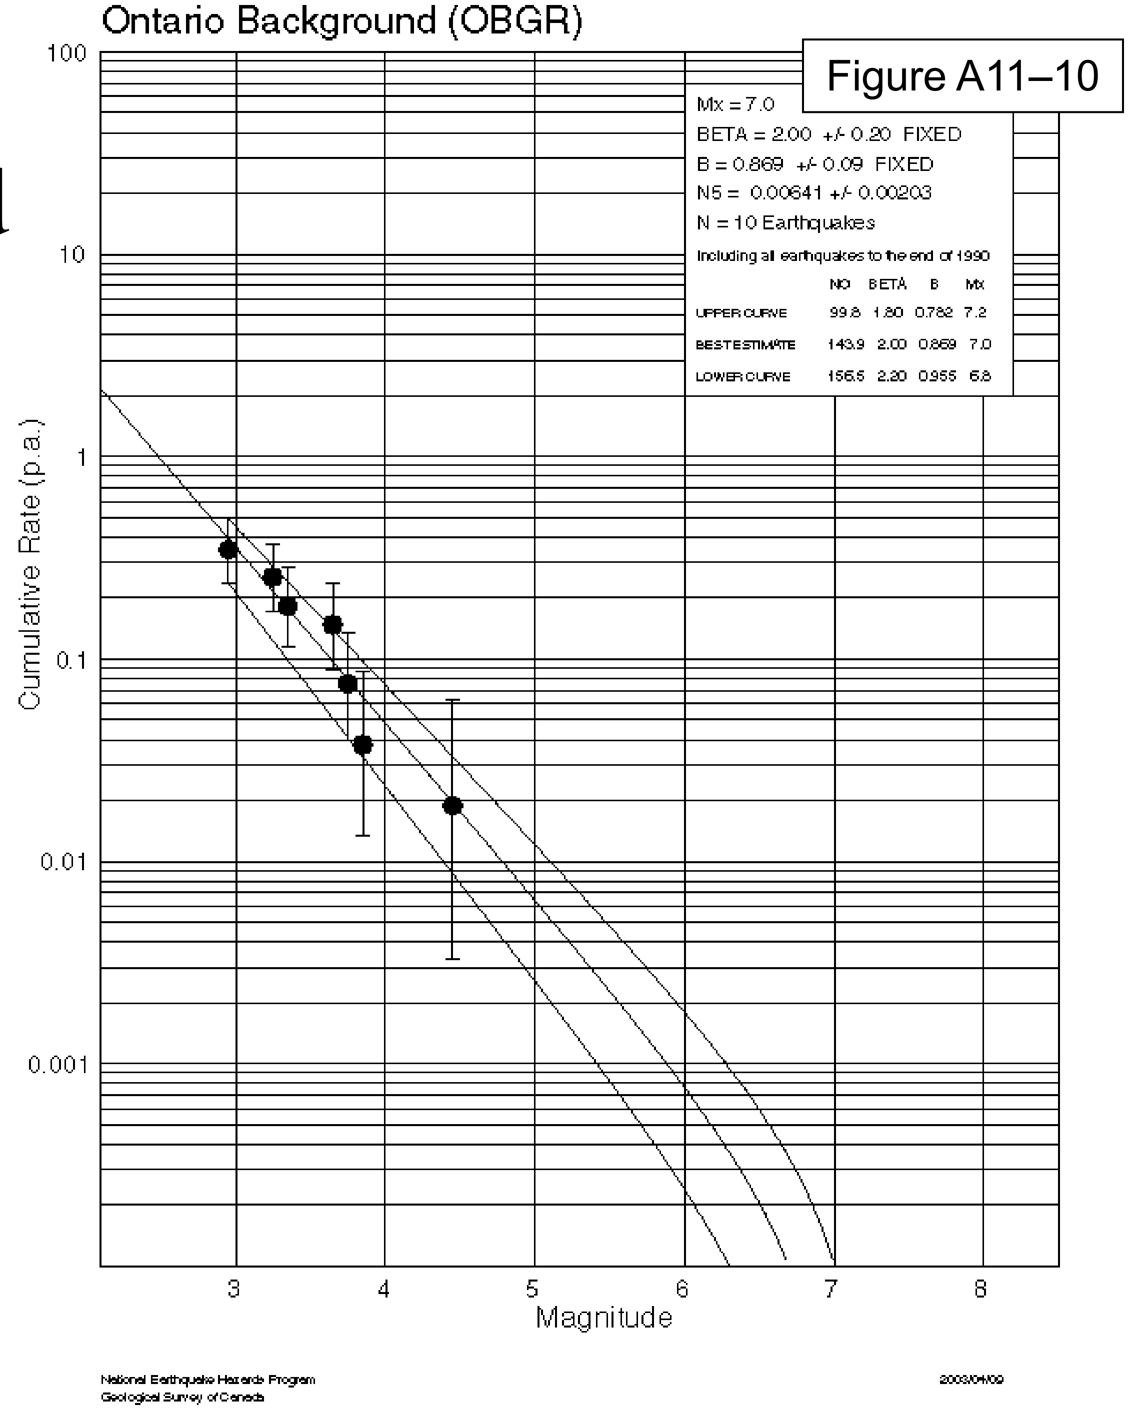

Paleozoic rifted margin

Mobs ~5.0

Weighted branches
best, upper, lower
7.5 7.7 6.0
0.68 0.16 0.16

potential large faults?

insensitive to Mmax

inconsistent with SCC!

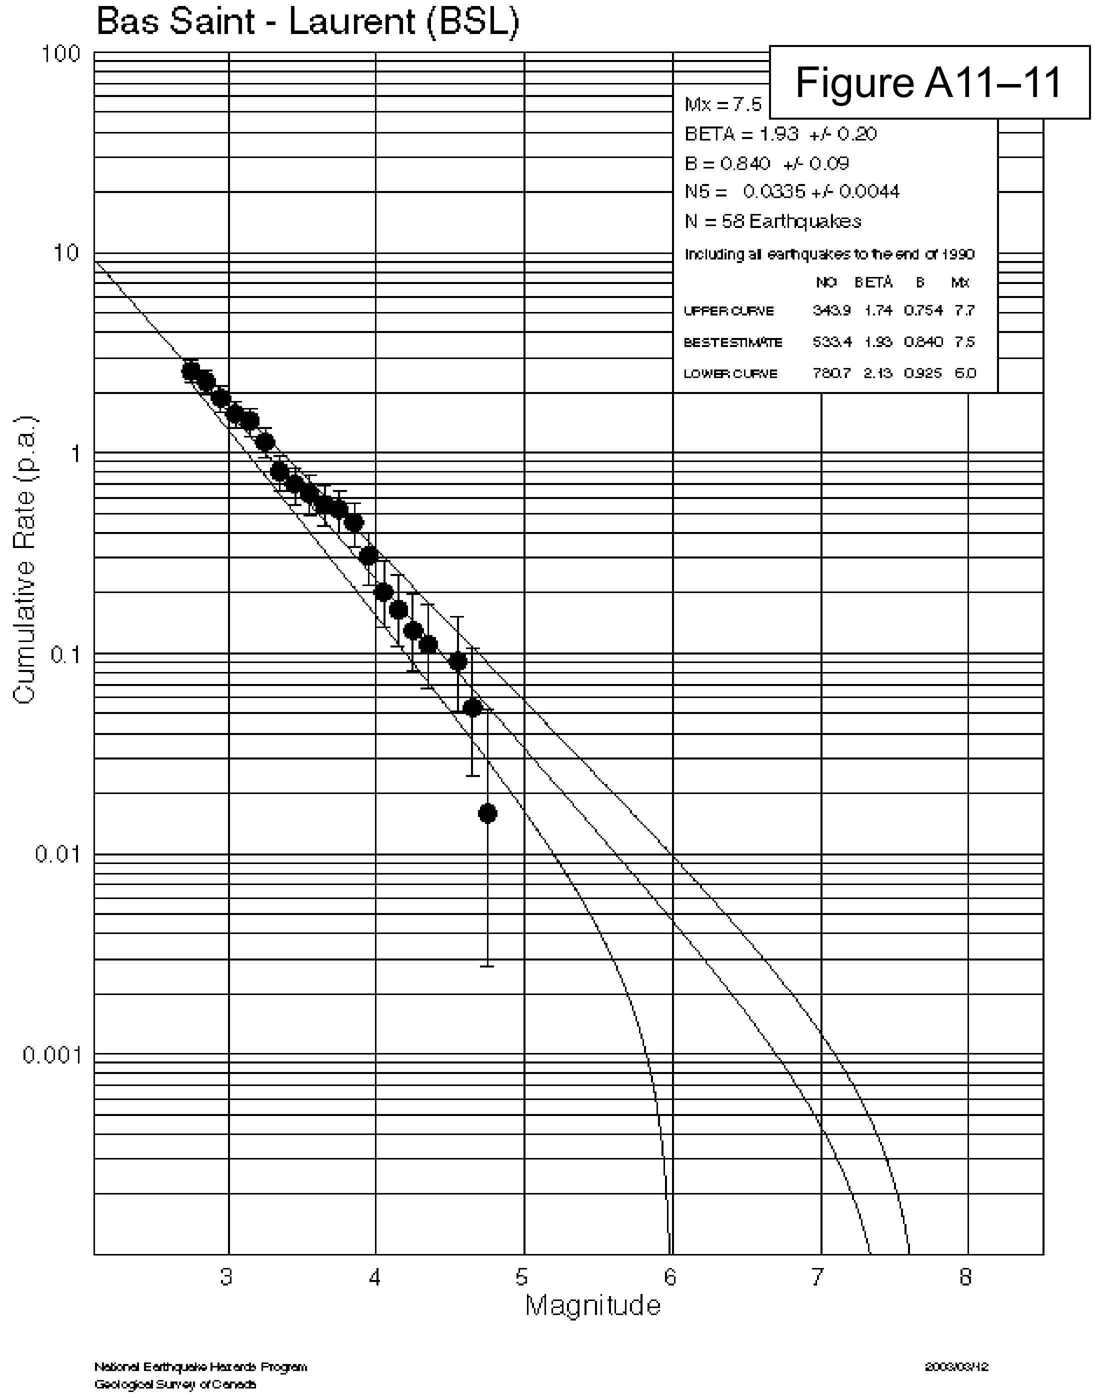

Figure A12–1

Australian Government

Geoscience Australia

Neotectonics in Australia: the not so stable continent

Dan Clark

Georisk Project

IGC Conference 2008

Figure A12–2

Stress in the Australian plate

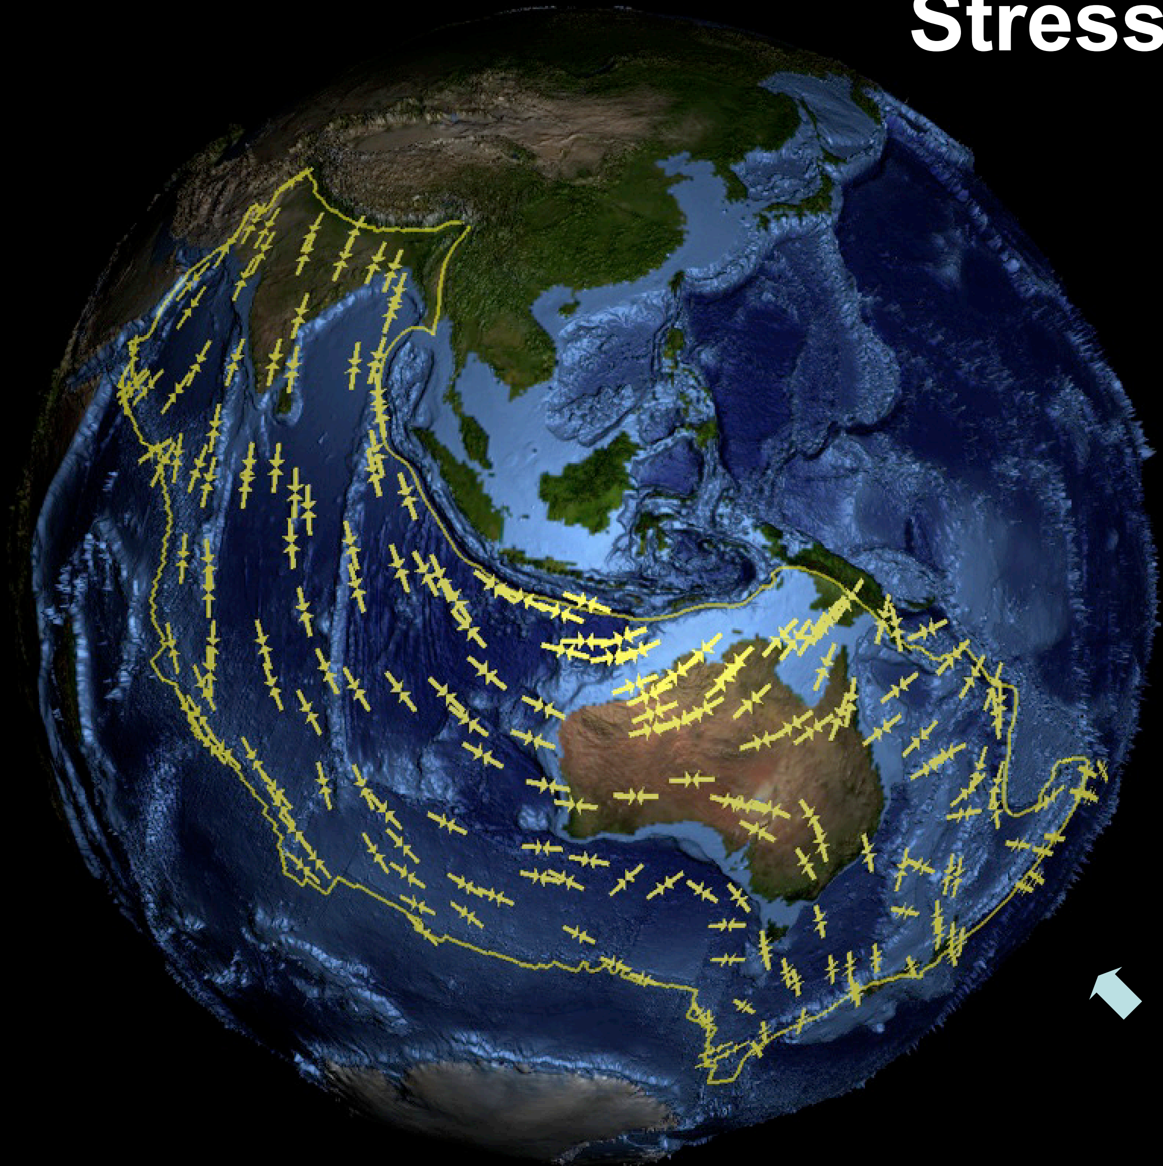

- Northerly velocity
  - 60-70 mm/yr
- Complex stress field
  - Boundary forces and basal tractions
  - E to SE in south
  - NE in north

Plate velocity vectors
(courtesy M. Sandiford)

Figure A12–3

From stress to deformation

- Mode 1 – northward tilting
- Mode 2 – long wavelength folding
- Mode 3 – faulting

SHmax and seismicity *(courtesy M. Sandiford)*

GEOSCIENCE AUSTRALIA

Figure A12–4

From stress to deformation

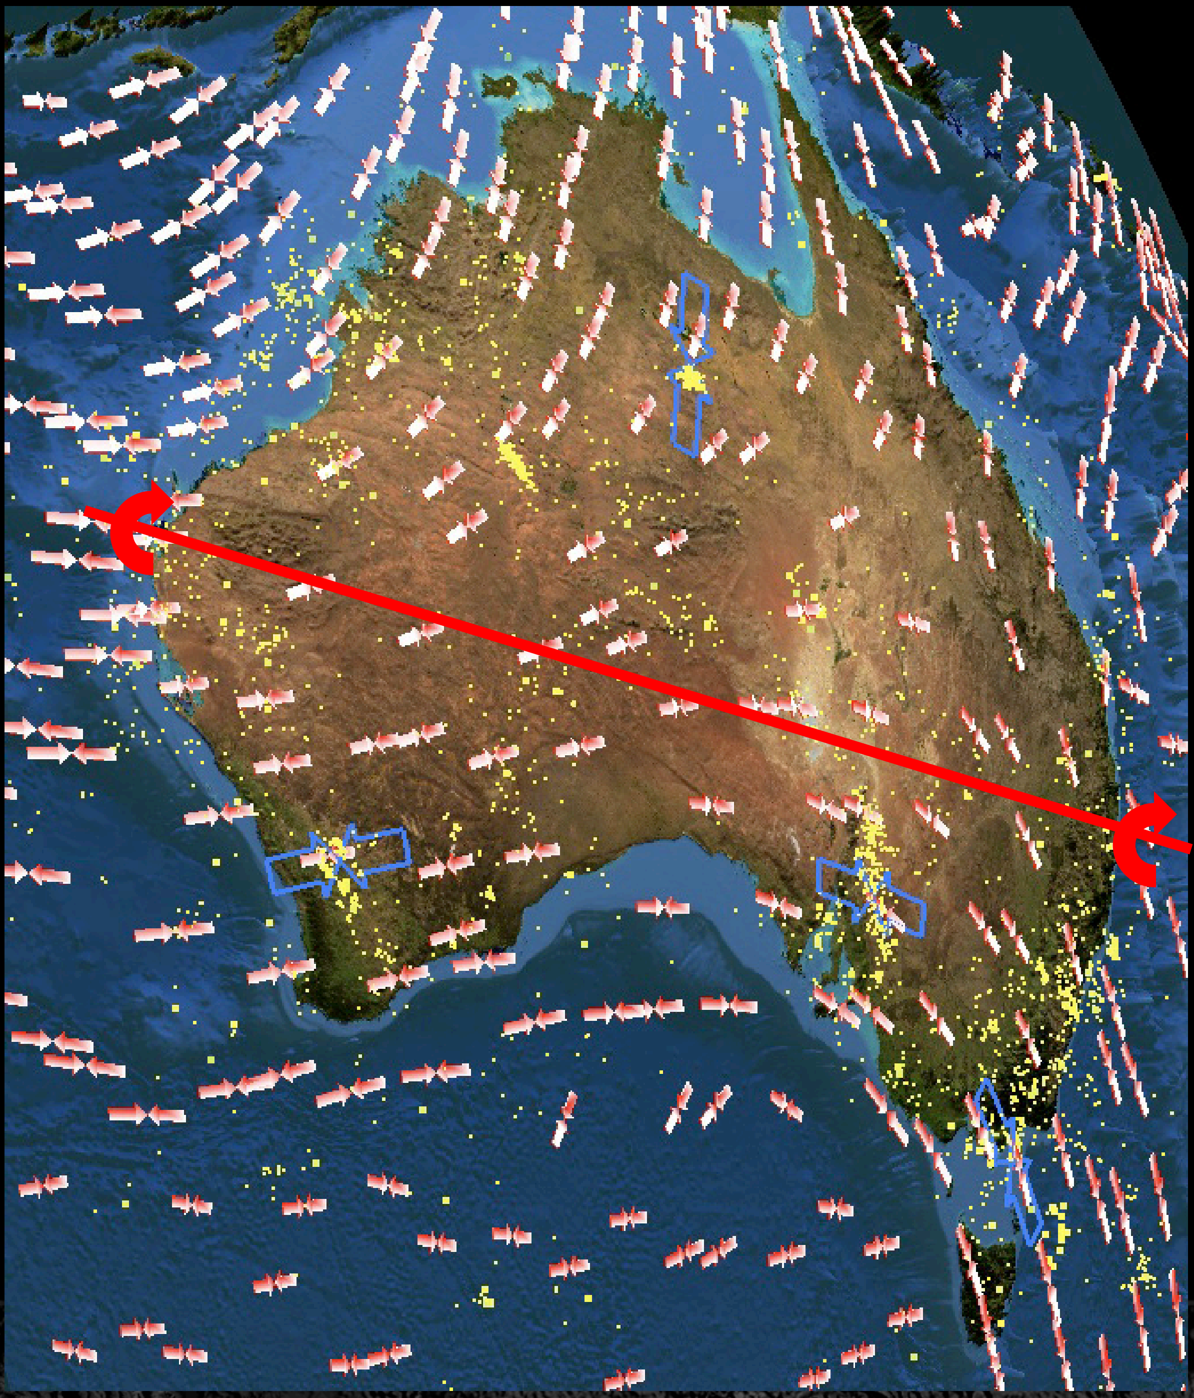

- Mode 1 – northward tilting
- Mode 2 – long wavelength folding
- Mode 3 – faulting

SHmax and seismicity *(courtesy M. Sandiford)*

GEOSCIENCE AUSTRALIA

Figure A12–5

From stress to deformation

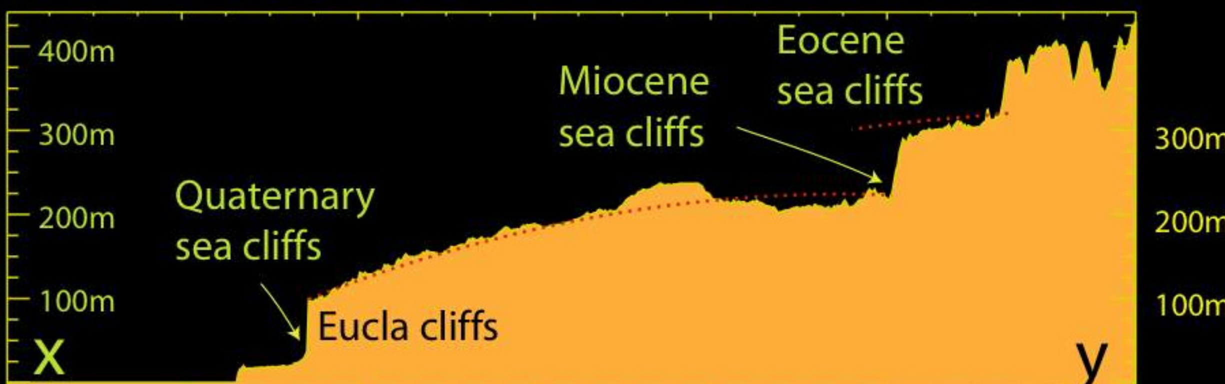

- Mode 1 – northward tilting
- Mode 2 – long wavelength folding
- Mode 3 – faulting

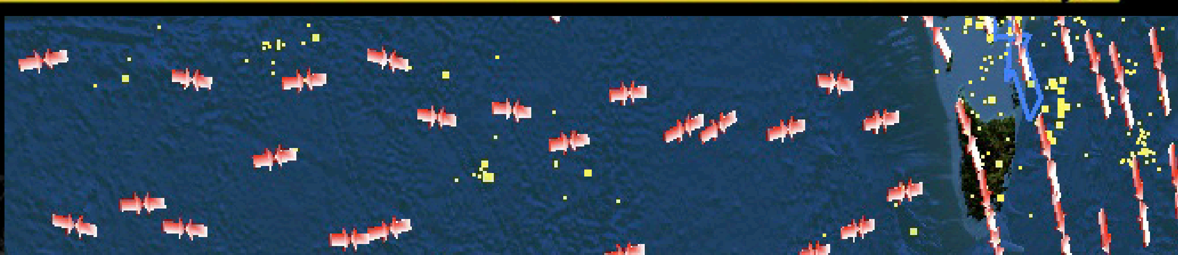

SHmax and seismicity
(courtesy M. Sandiford)

GEOSCIENCE AUSTRALIA

Figure A12–6

Australia's "neotectonic" faulting record

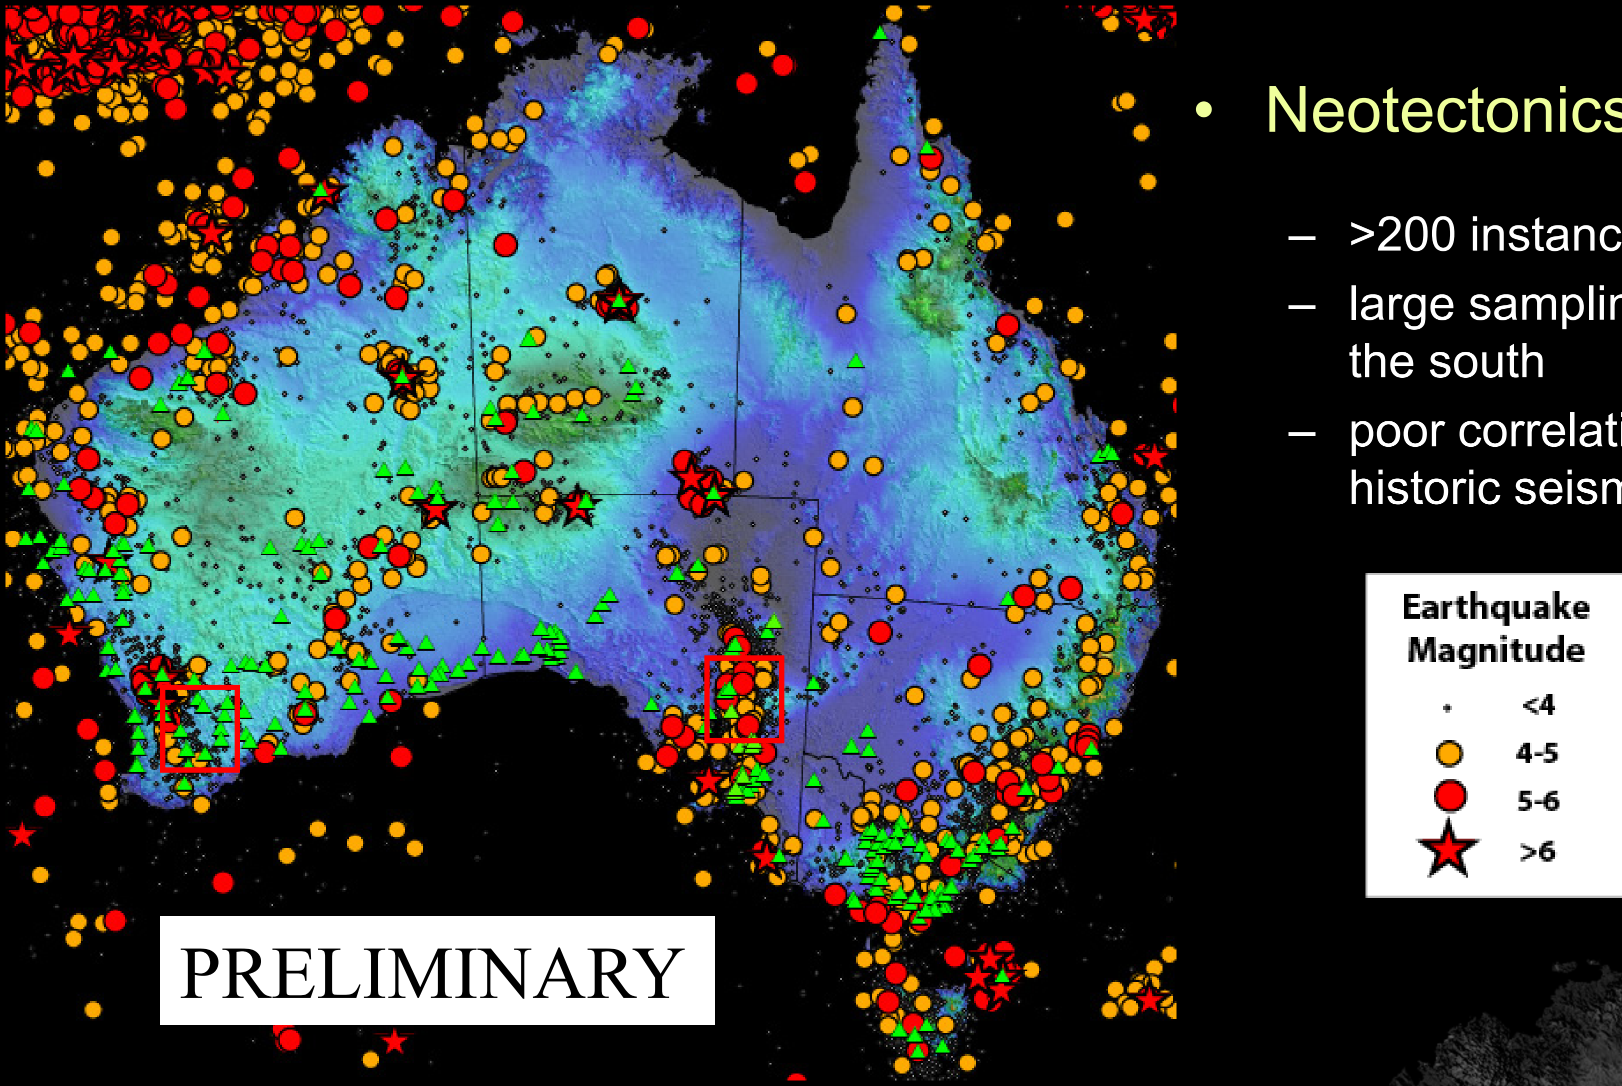

- Neotectonics data
  - >200 instances
  - large sampling bias to the south
  - poor correlation to historic seismicity

Neotectonic features from GA neotectonics database

Expression of SW WA scarps

Figure A12–7

1.5 m high Merredin scarp

3 m high Narrabeen scarp

4 m high Dumbleyung scarp

Expression of SW WA scarps

Figure A12–8

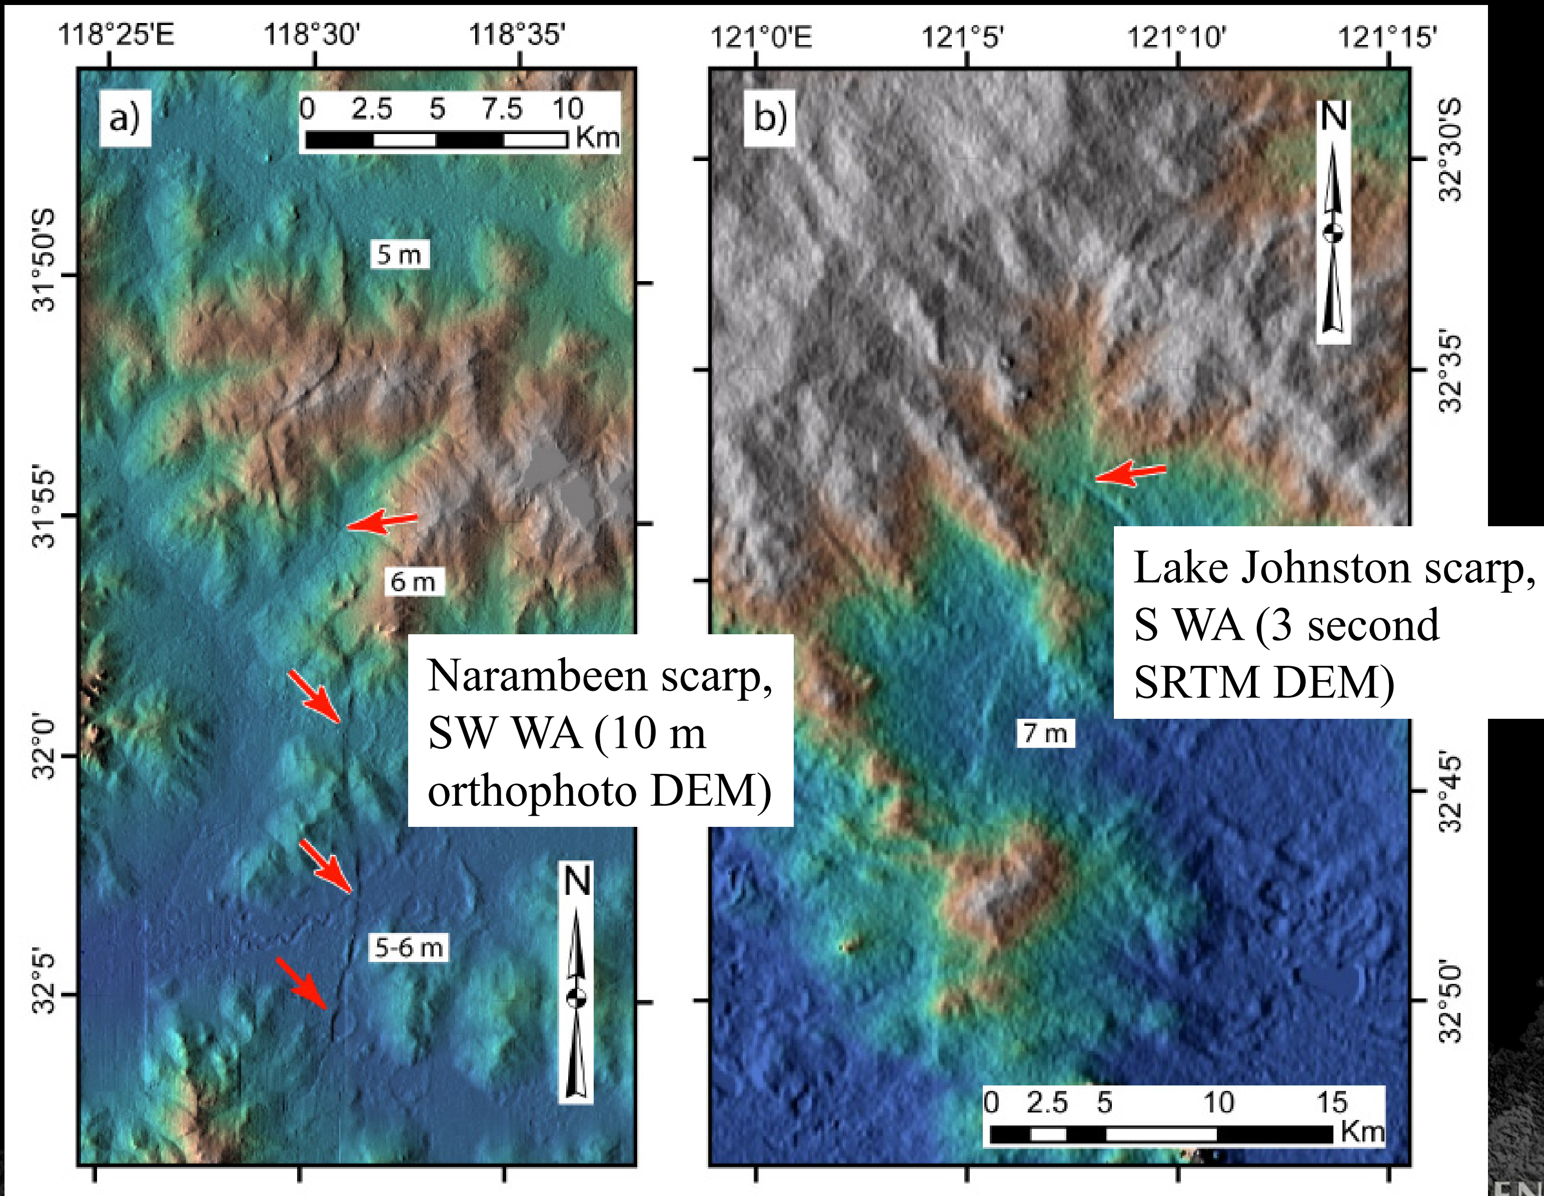

Expression of Mt Lofty Ranges scarps

300 m high Willunga scarp

200 m high Milendella scarp

Figure A12–10

Expression of Mt Lofty Ranges scarps

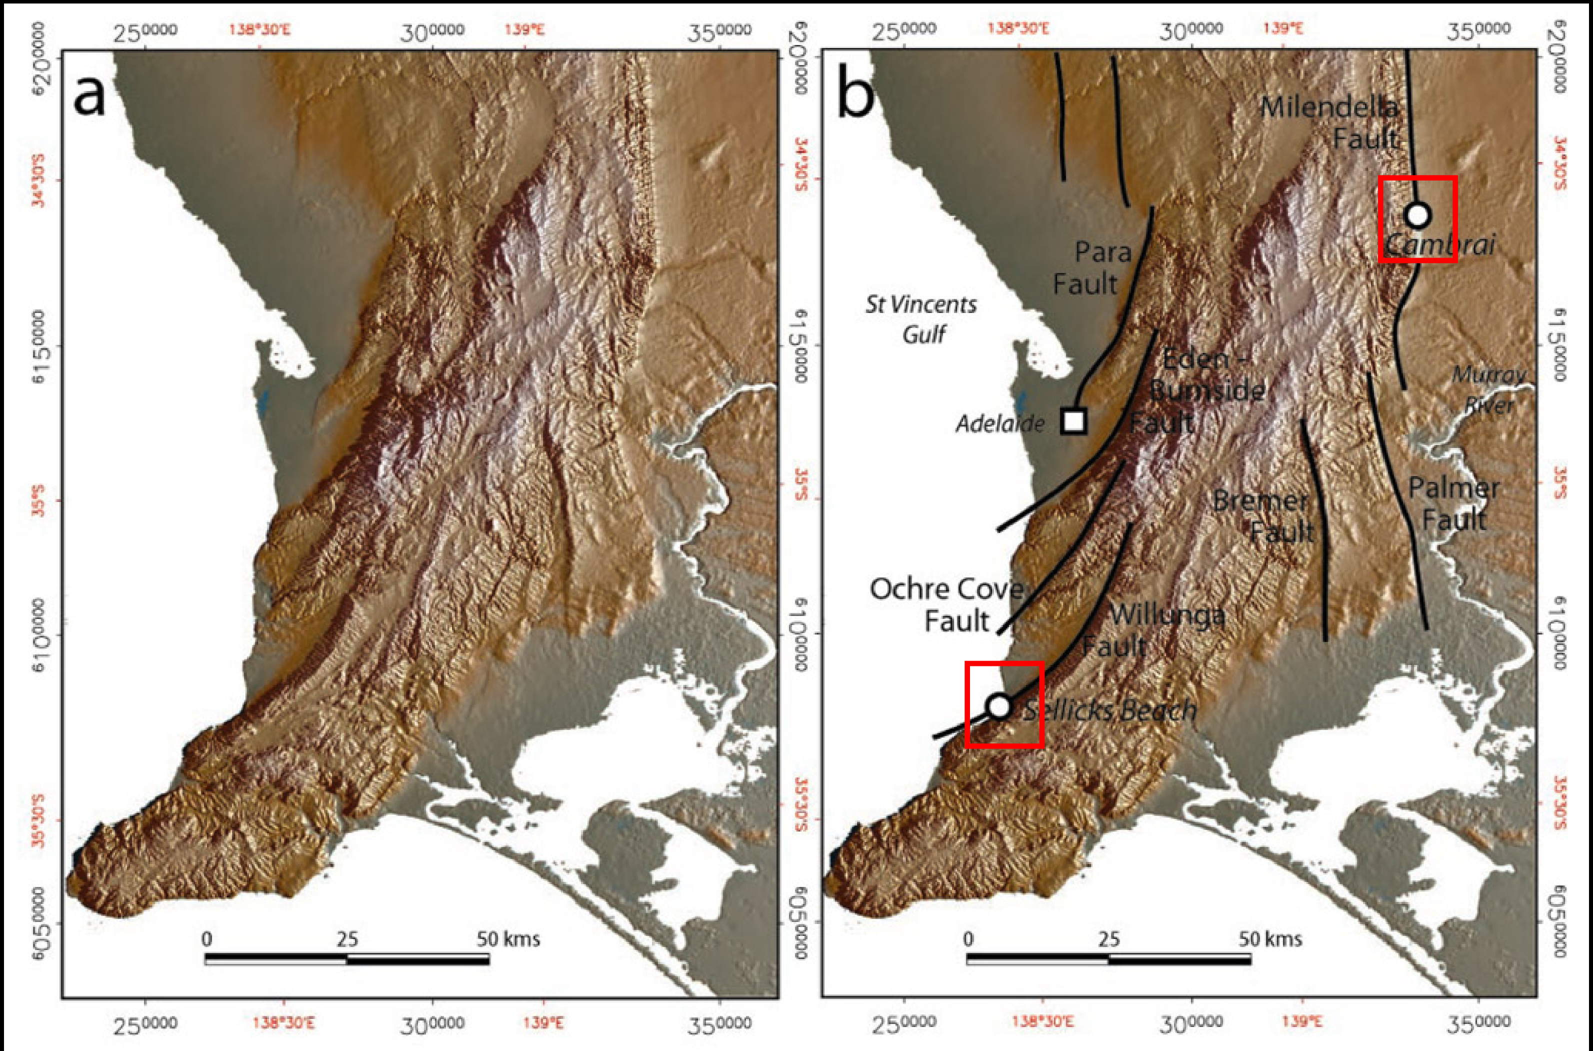

Fault bounded ranges (Sandiford 2003)

Figure A12–11

Preliminary neotectonics domains

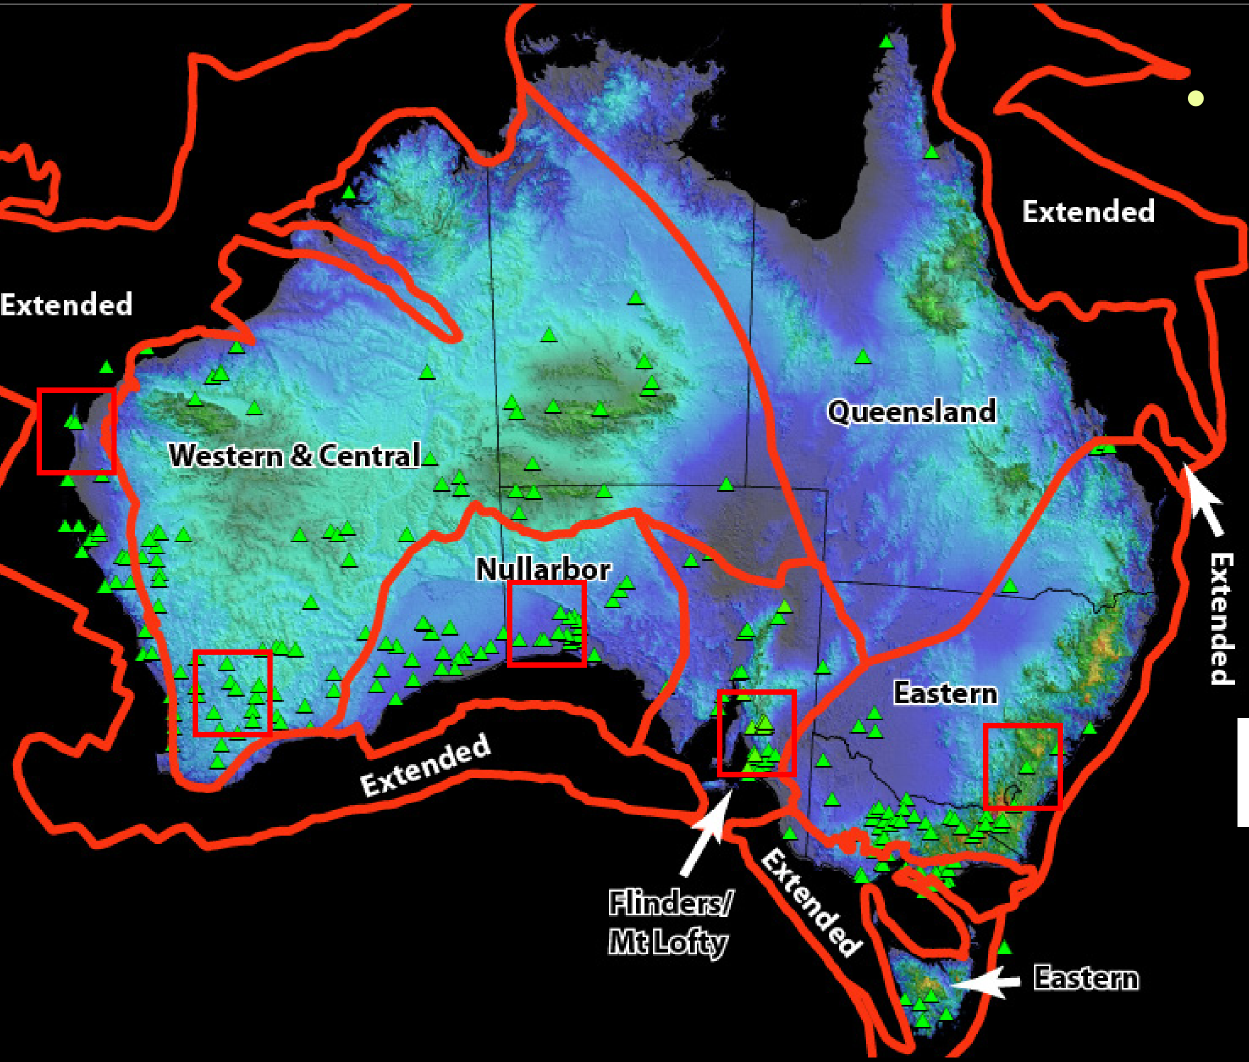

- Six onshore domains divided by:
  - total post-10 Ma offset and length
  - recurrence behaviour
  - spatial distribution / density

PRELIMINARY

Cadell Fault Structure

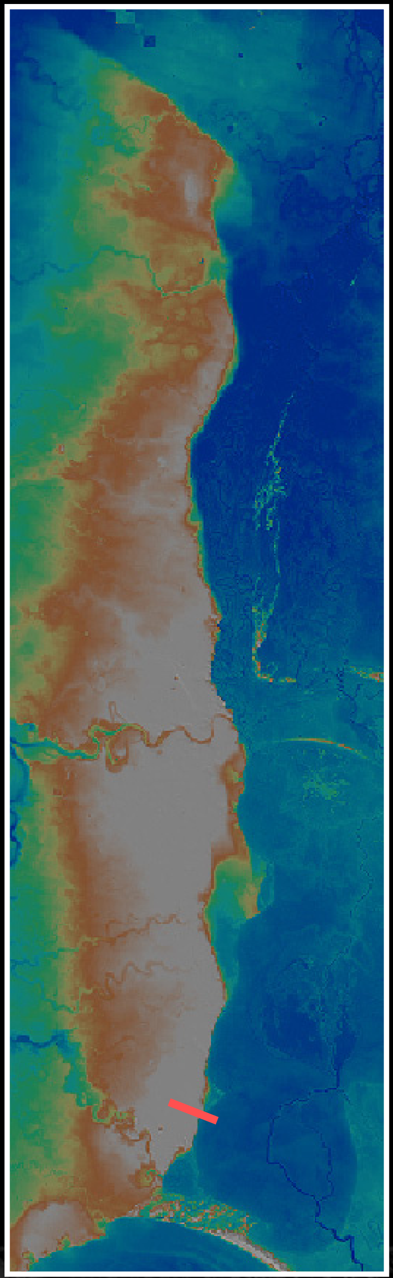

- reverse fault dipping 50 degrees west
- 100 m total displacement across bedrock
- fault propagation fold near surface

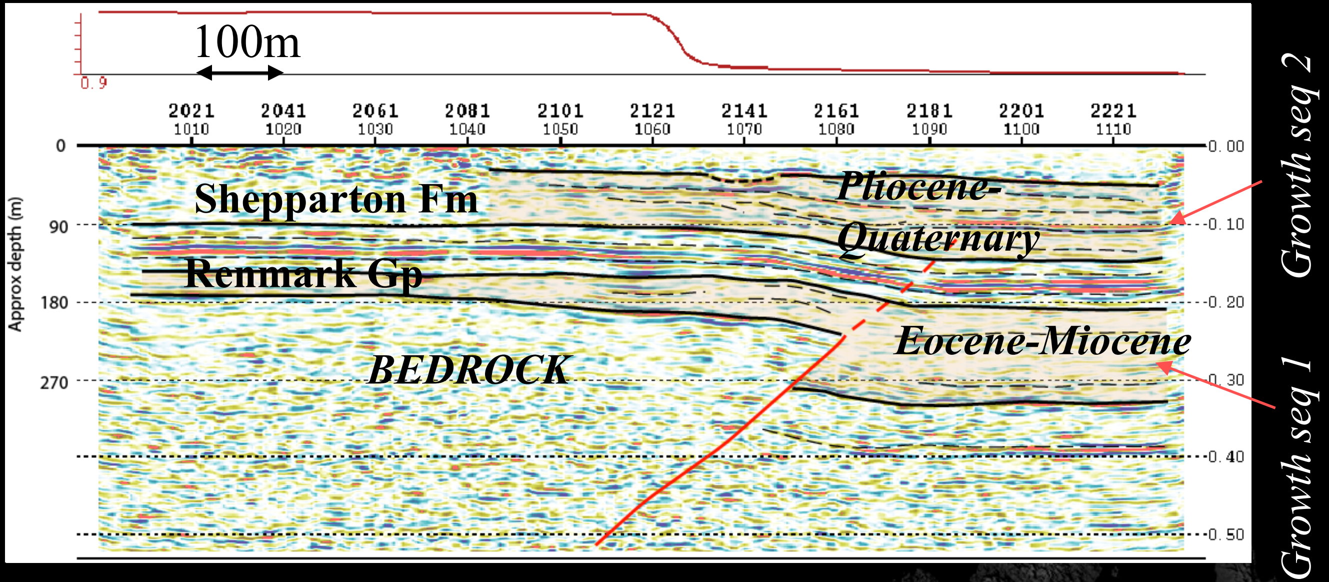

LIDAR DEM

Seismic suggests growth periods separated by inactive periods

Nullarbor domain

Figure A12–13

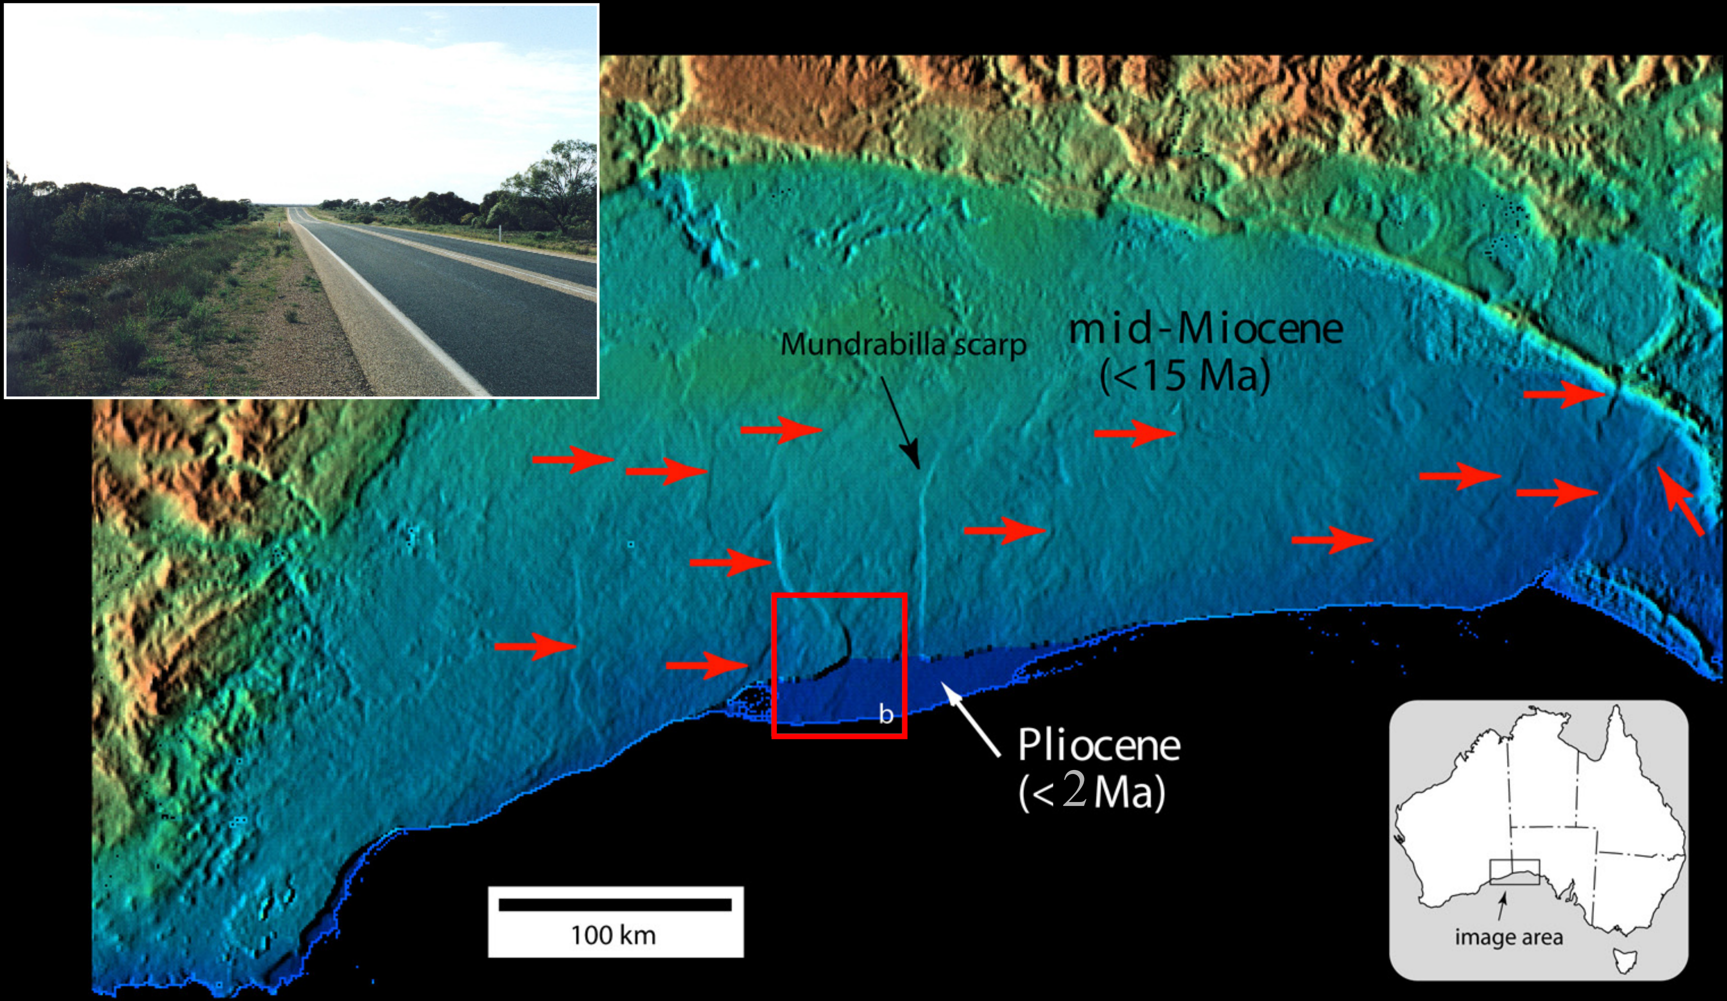

Nullarbor domain

Figure A12–14

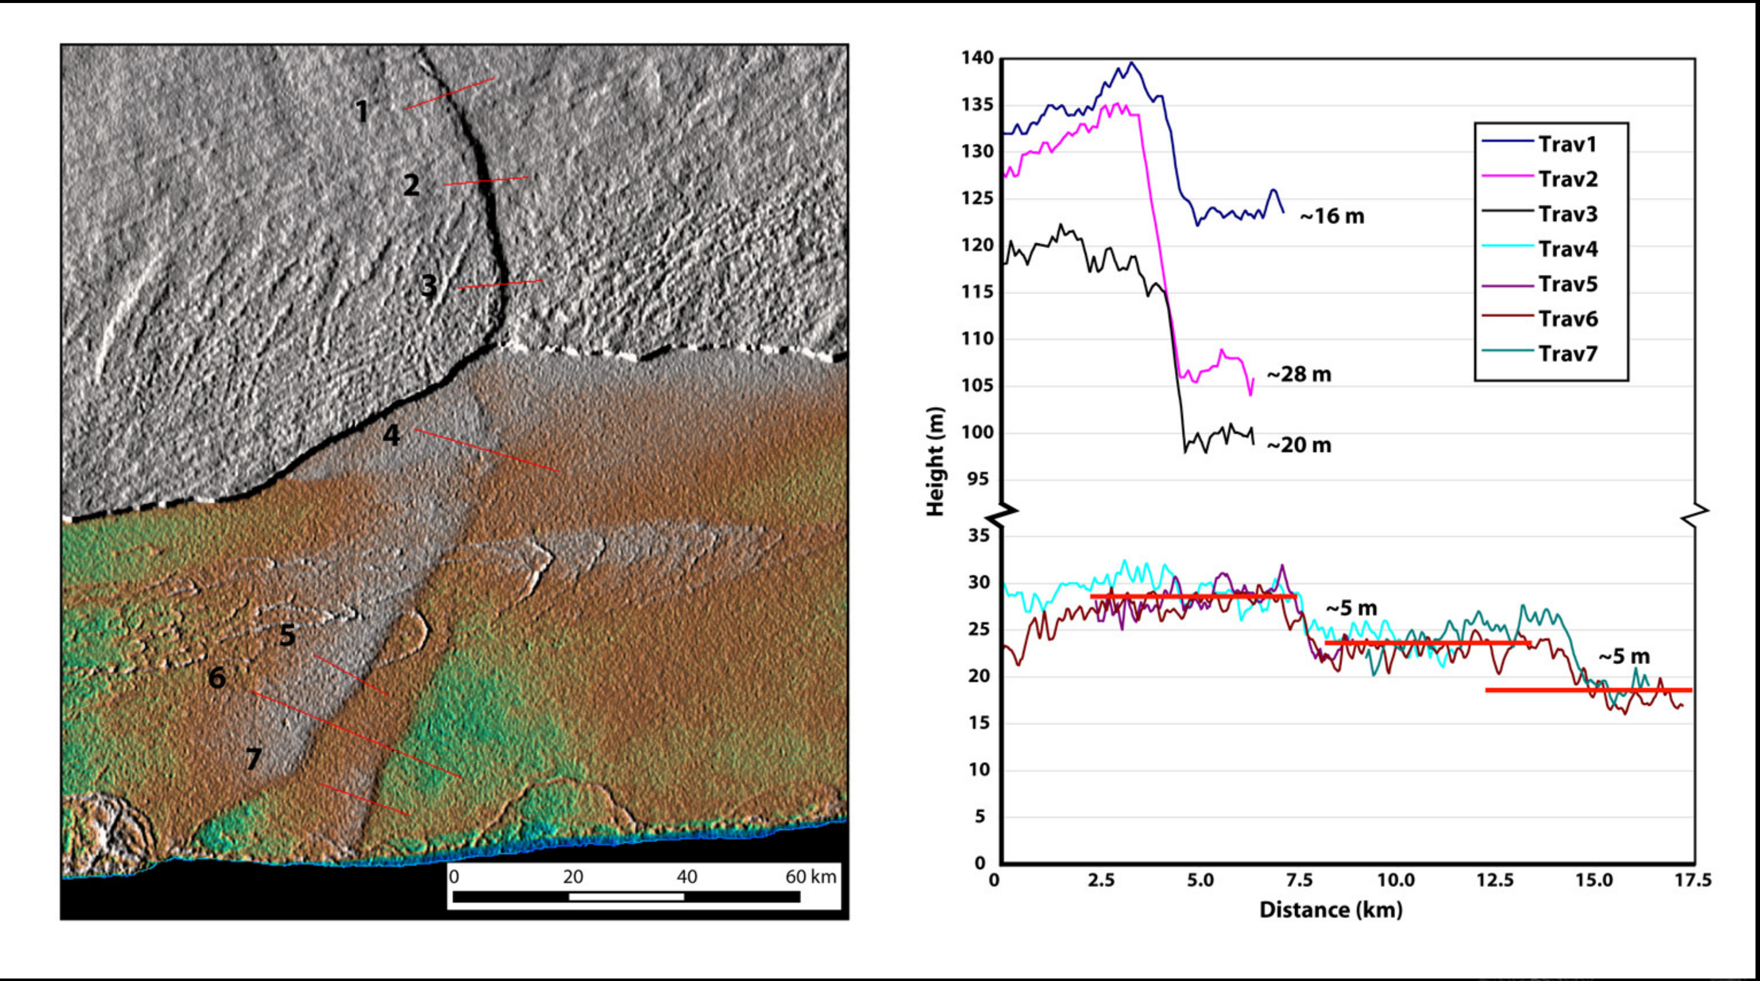

Western & Central domain

Figure A12–15

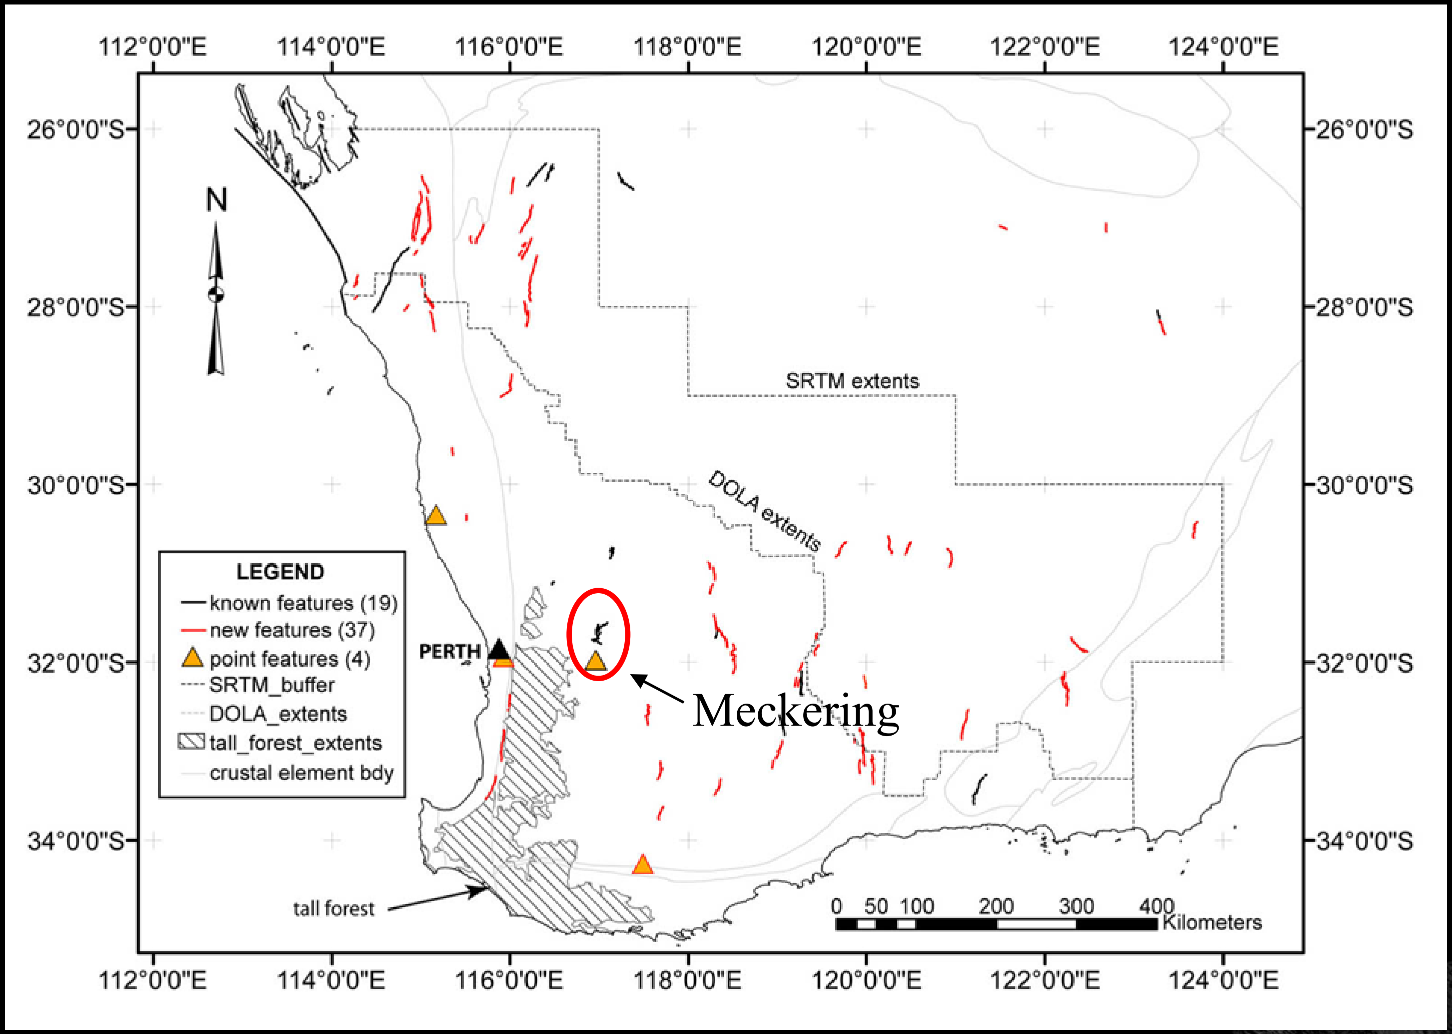

Scarps identified in hi-res DEM

Figure A12–16

Carnarvon Basin

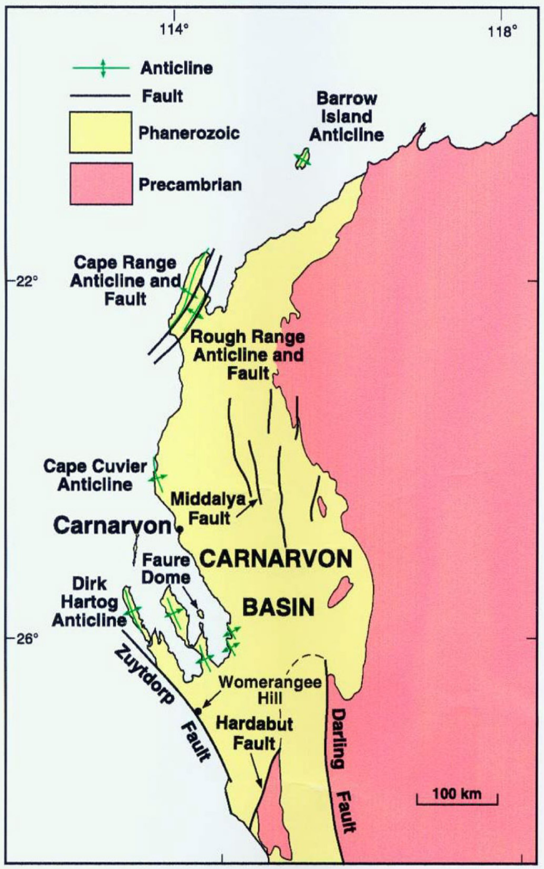

Map of the Carnarvon Basin showing important localities for study of Quaternary tectonism

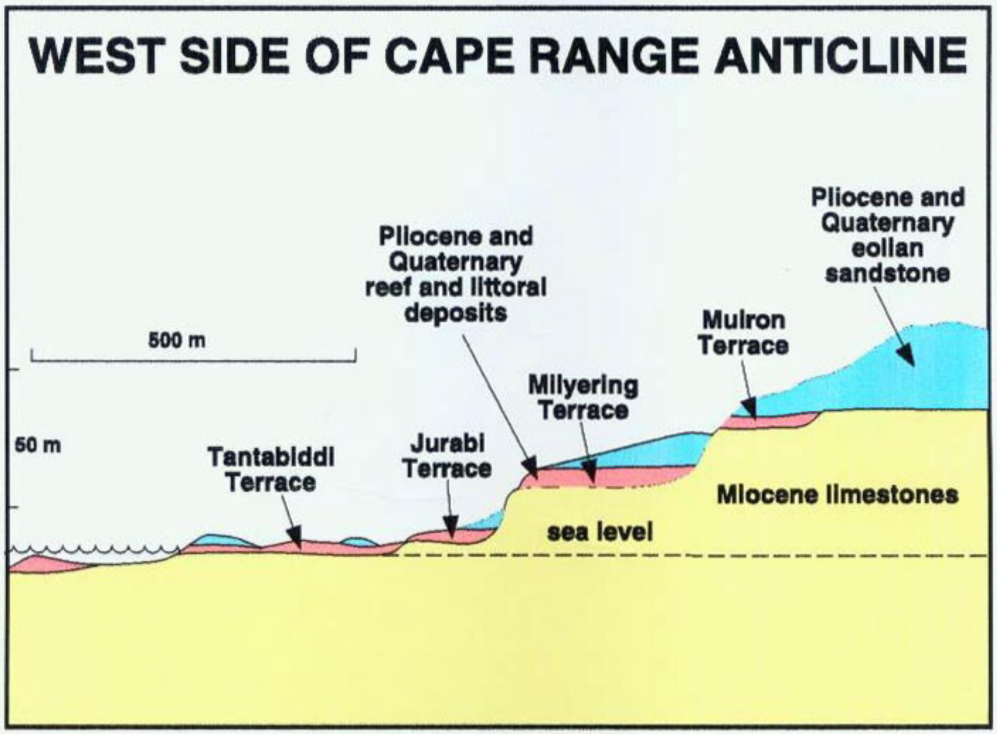

Series of uplifted Pliocene and Quaternary terraces on the west side of Cape Range Anticline

Figure A12–17

Carnarvon Basin

Tantabiddi Creek, showing a section through the lower two uplifted terraces

Intraplate fault rupture behaviour

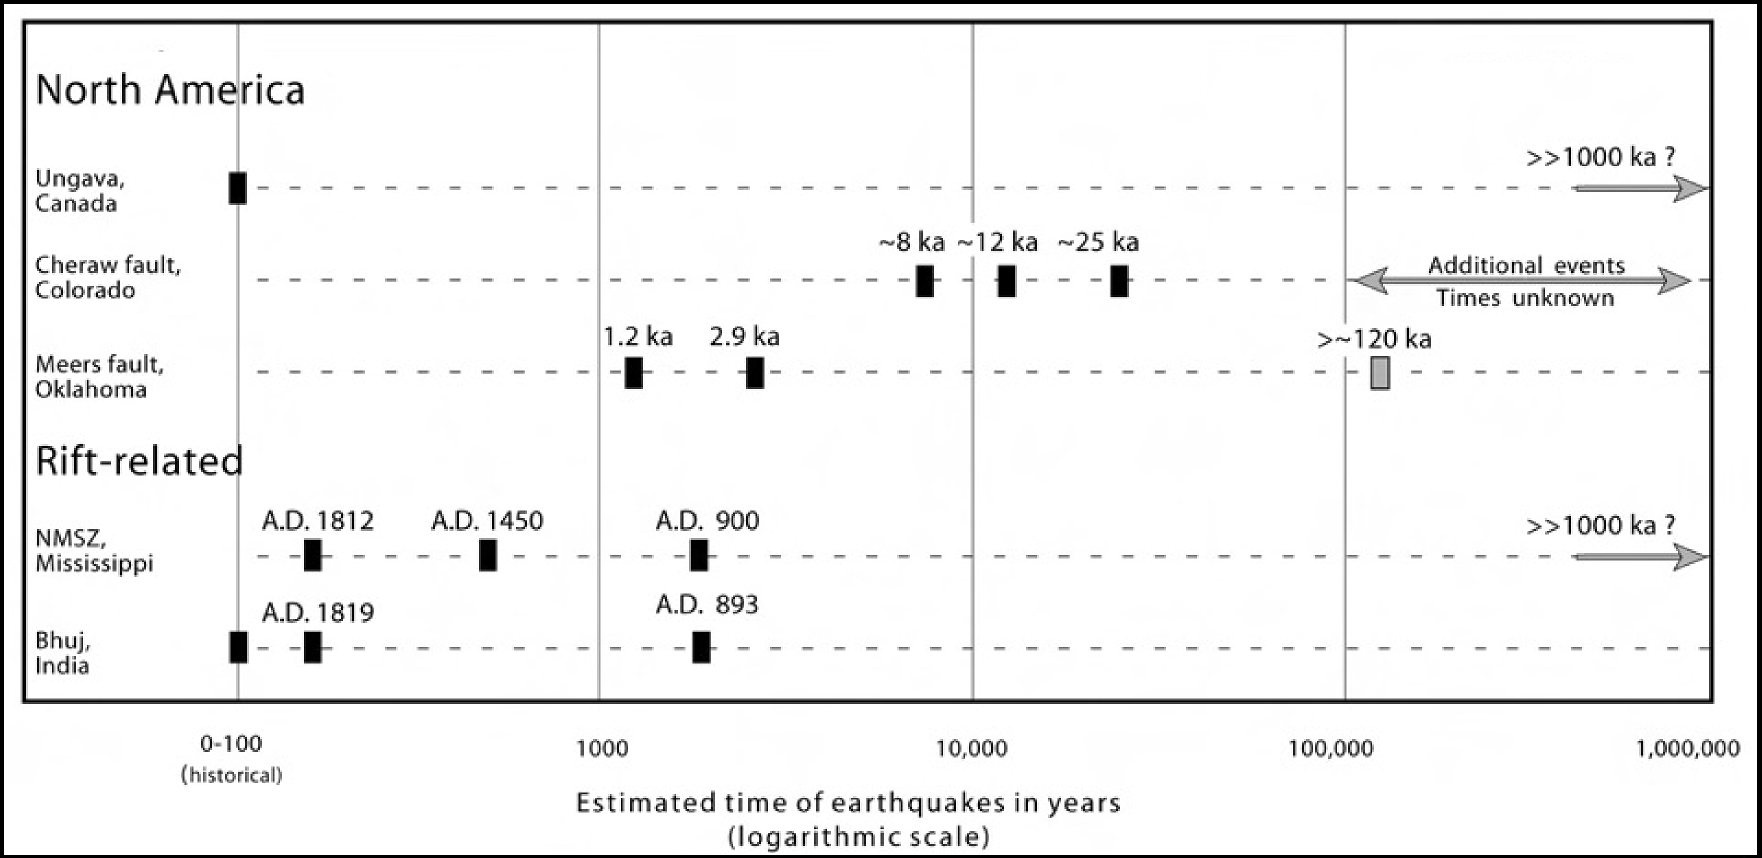

Figure A12–19

Intraplate fault rupture behaviour

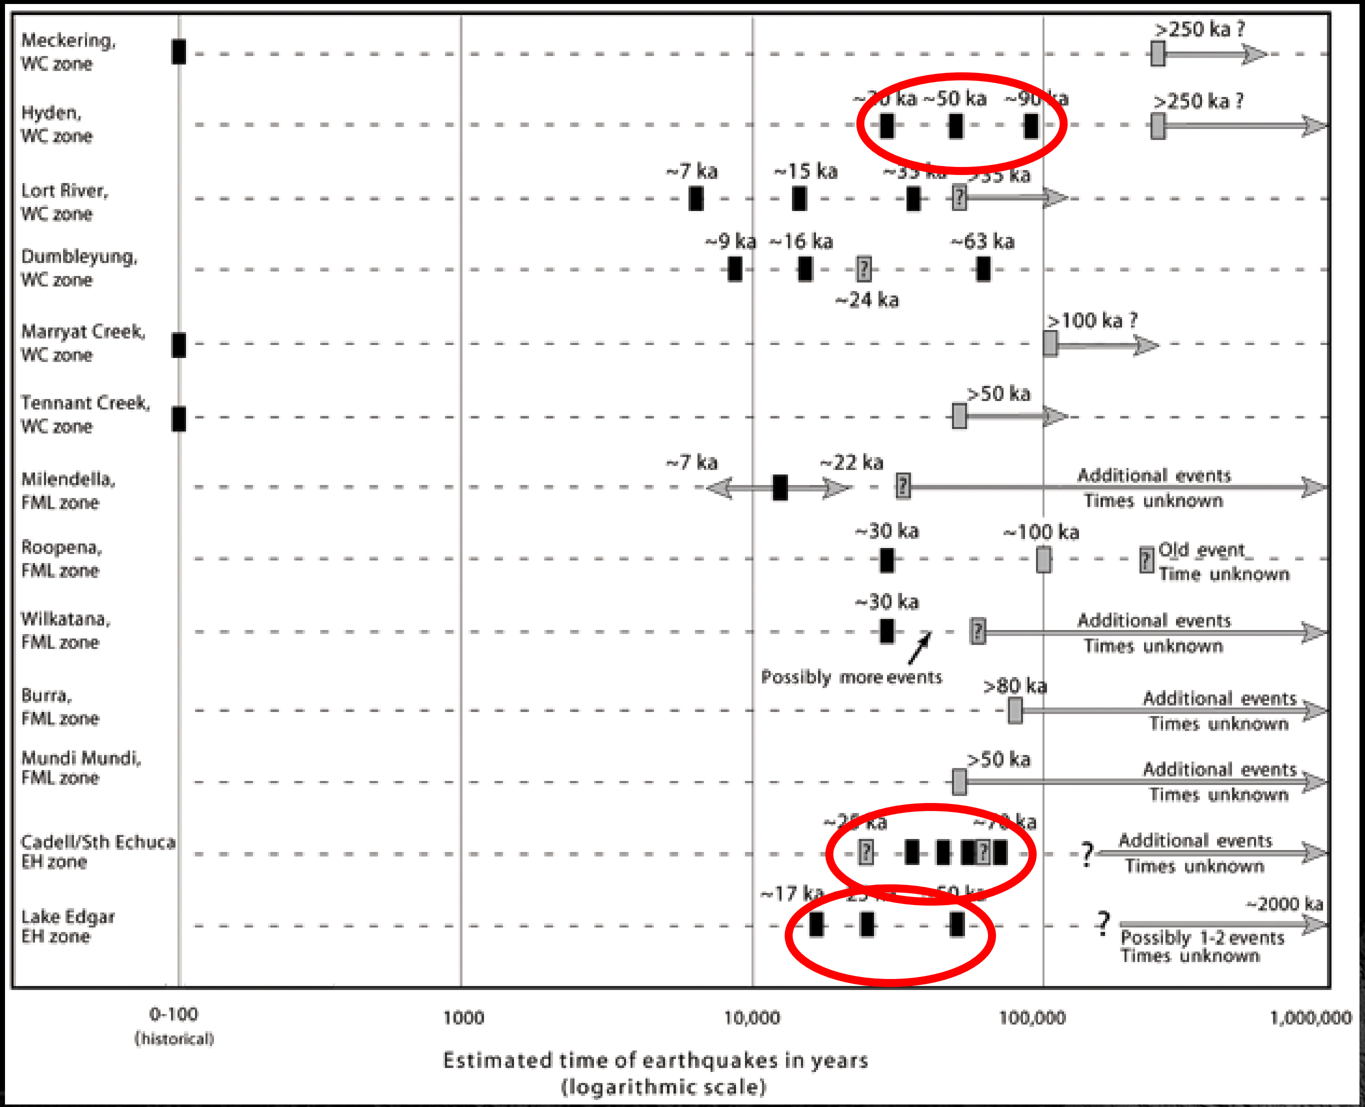

Figure A12–20

PRELIMINARY

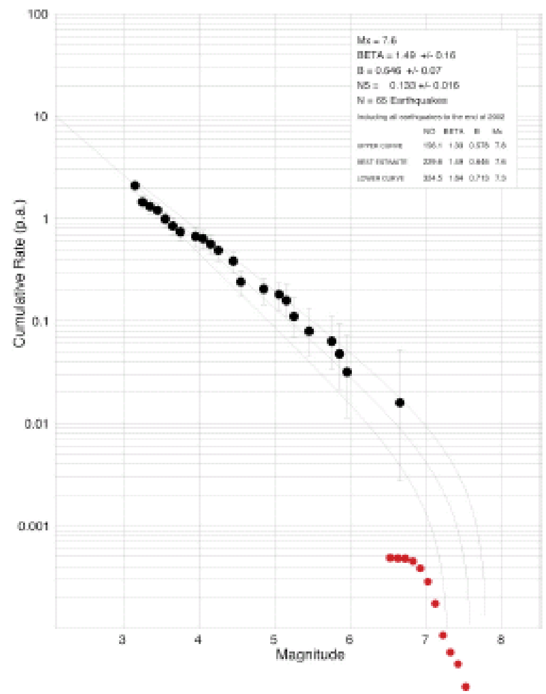

Figure 1 The cumulative recurrence rate of earthquakes in the SWSZ (black) from the GA catalogue and the SW WA neotectonics catalogue (red). The neotectonic catalogue has been scaled for time and area to match the earthquake catalogue.

Leonard & Clark, 2006

Figure A12–21

PRELIMINARY

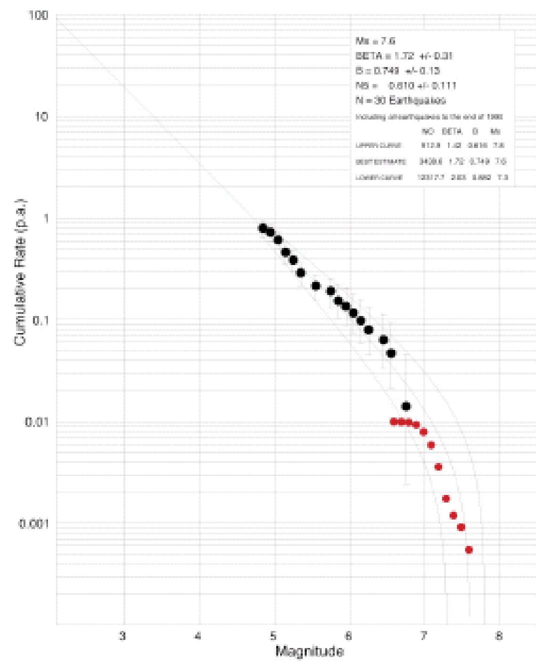

Figure 2 The recurrence rate for the Australian shield, from the world-wide CSS catalogue, and the scaled neotectonic catalogue.

Leonard & Clark, 2006

Figure A12–22

PRELIMINARY

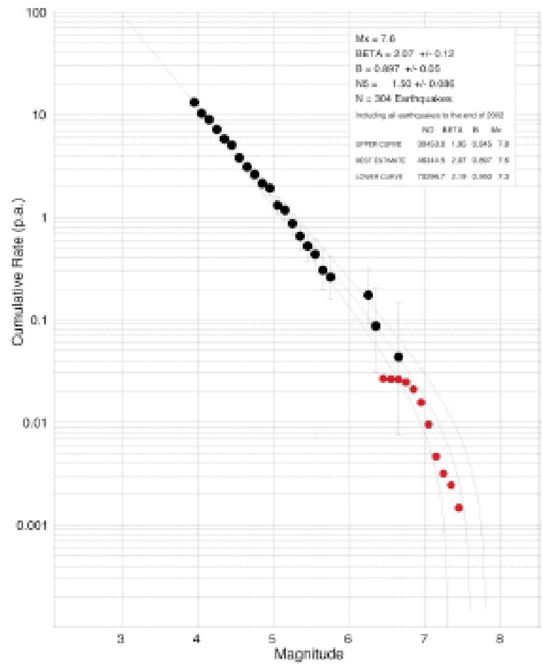

Figure 3 The recurrence rate for the Australian continent and the scaled neotectonic catalogue.

Leonard & Clark, 2006

Figure A12–23

PRELIMINARY

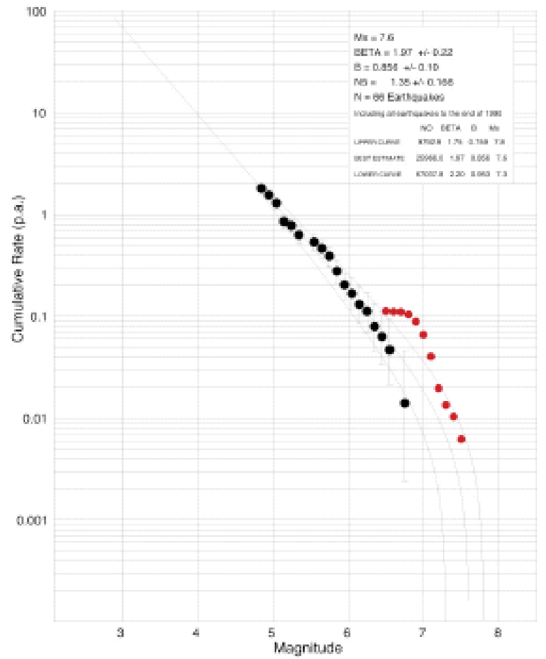

Figure 4 The recurrence rate for the world-wide SCC catalogue and the scaled neotectonic catalogue.

Leonard & Clark, 2006

Figure A12–24

Figure A13–1

What is a “tectonic analog”?

- (1) Mmax implies large L, W.
- Most SCRs are best examined with regional geologic and tectonic maps (1:2,500,000 – 1:10,000,000 scale)
- (2) Large tectonic elements (rifts, passive margins, orogens, cratons) are easiest to identify on regional maps.
- (1) + (2): Do any kinds of large tectonic elements have characteristic fault styles that favor large rupture L, W?
- (Most useful if represented in CEUS and adjacent Canada)
- Mmax = large Mo release on large faults

USGS

Figure A13–2

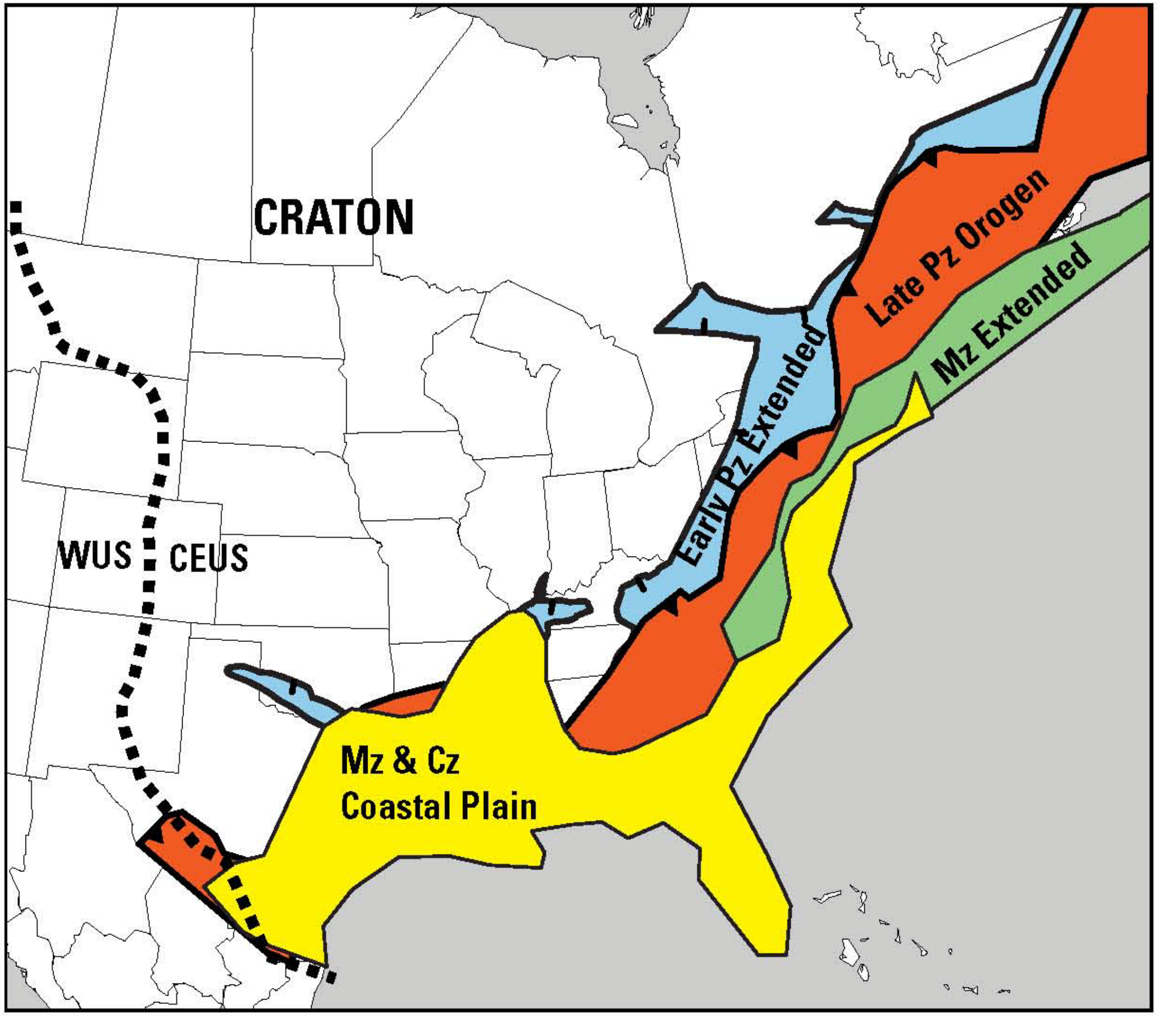

Main CEUSAC Tectonic Elements

Characteristic fault styles

Figure A13–3

- Extensional plate motions (rifts, passive margins)
  - Rift-parallel & margin-parallel faults are long
  - Steep dips, deep penetration (alkaline igneous rocks)
  - Large L, W: large Mo
  - Thick, restricted sandstones, then widespread limestones
- Contractional plate motions (orogens)
  - Low-dip thrust faults, long, wide, above most seismicity
  - Faults steepen into hot cores (healing, folding, offsetting faults)
  - Large L, W, where low dips; smaller W where steeper: smaller Mo?
  - Sandstones, shales, limestones; thicker closer to rising mountains
- Precambrian plate motions (cratons)
  - Old: rifting, folding, thrusting overprinted on each other repeatedly
  - Faults of all L and orientations, older faults can be healed and deformed; small effective L, W: smaller Mo
  - Tectonically quiet: few associated sedimentary, igneous rocks

USGS

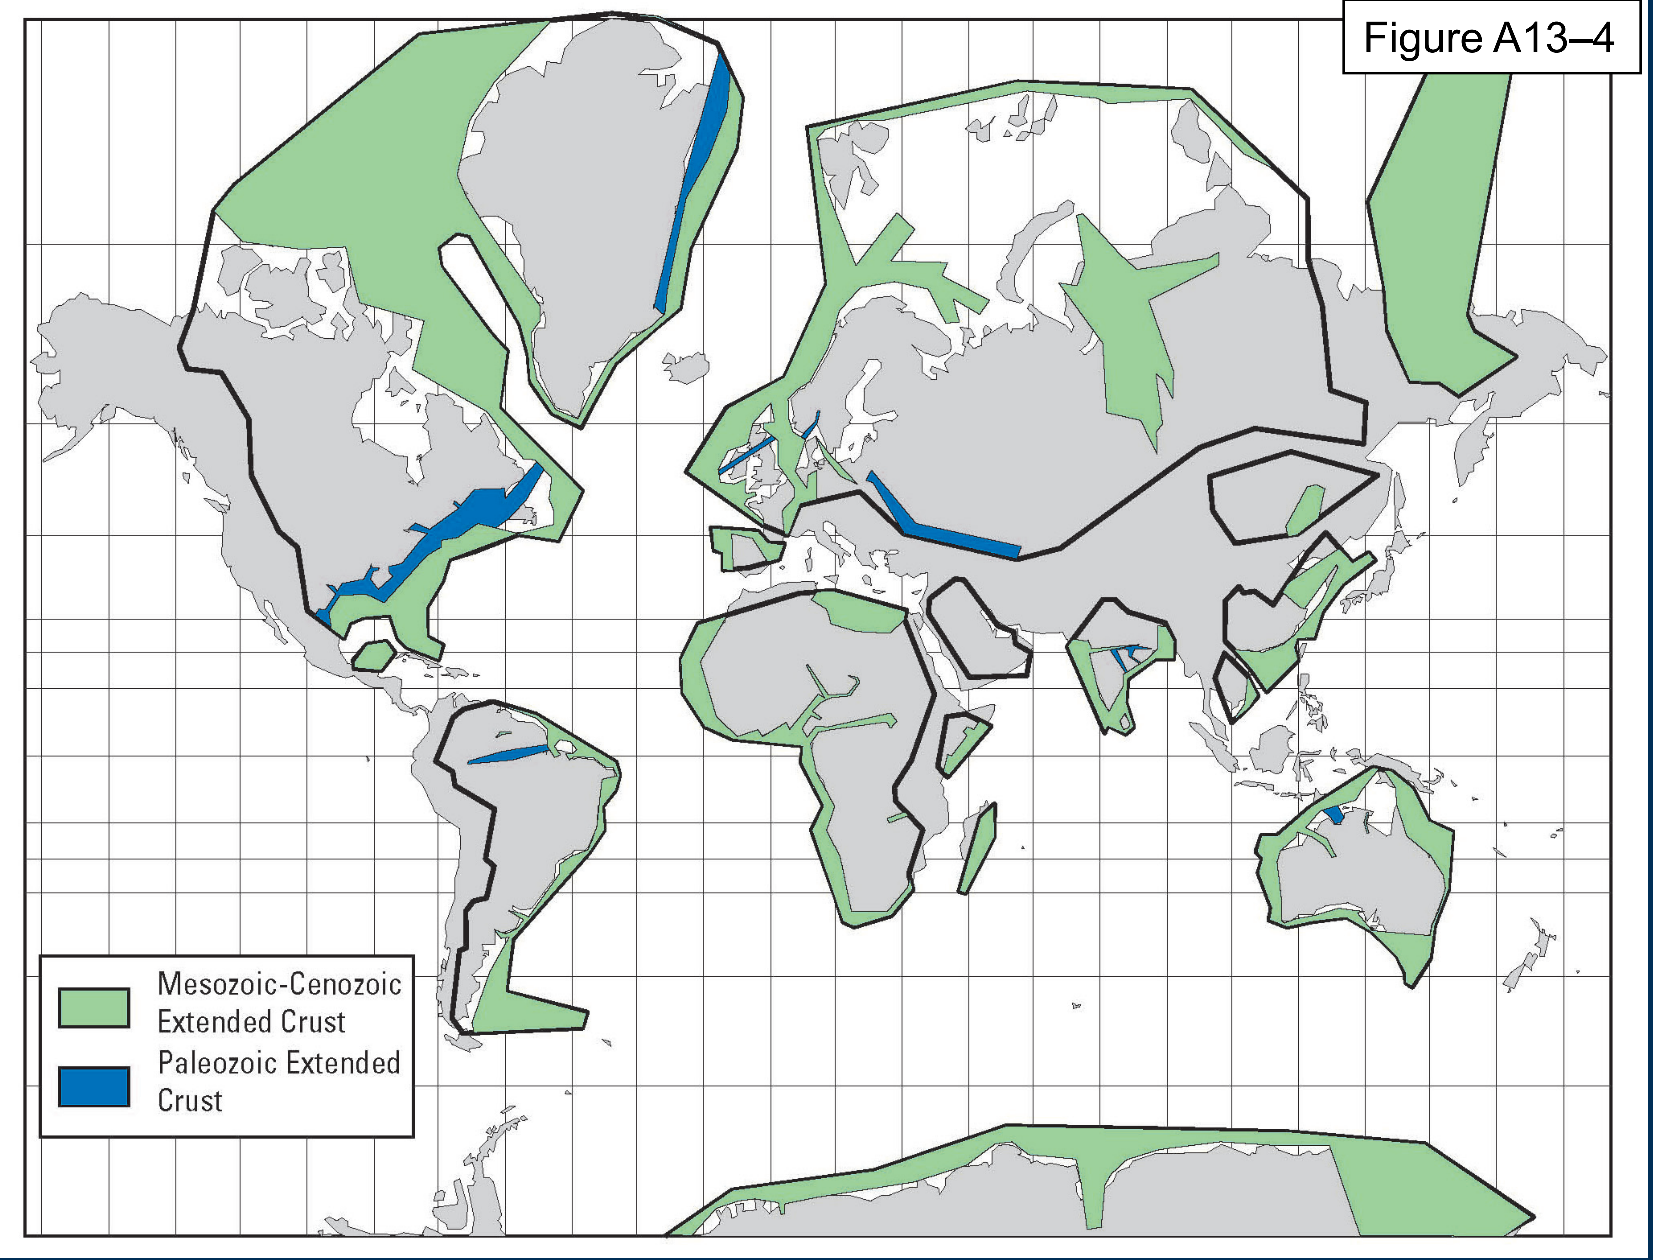

Figure A13–4
Mesozoic-Cenozoic Extended Crust
Paleozoic Extended Crust

Analogs

Figure A13–5

- Rifts & margins have distinctive fault styles and associated igneous and sedimentary rocks.
- Ditto for orogens. Ditto for cratons.
- Rifts & margins are recognizable from one SCR to another, ditto for orogens, cratons
- Rifts & margins, orogens, cratons are distinguishable from each other within same SCR
- Differences within an SCR >> differences between SCRs
- Faults heal so faulting age matters

USGS

Figure A14–1

Estimated Magnitude of the 1663 Earthquake at Charlevoix, Quebec

Figure A14–2

Many of the small earthquakes in our region may be very late aftershocks of strong earthquakes that took place hundreds or thousands of years ago. Under this "paleoseismicity" hypothesis, the spatial extents and activity rates of clusters of earthquakes can be used to estimate the magnitudes and times before present of past strong earthquakes (from *Ebel, Bonjer and Oncescu,* ***Seism. Res. Lett.,*** *2000*). Documenting persistent earthquake clusters throughout the historic record may help identify the locations of past strong earthquakes.

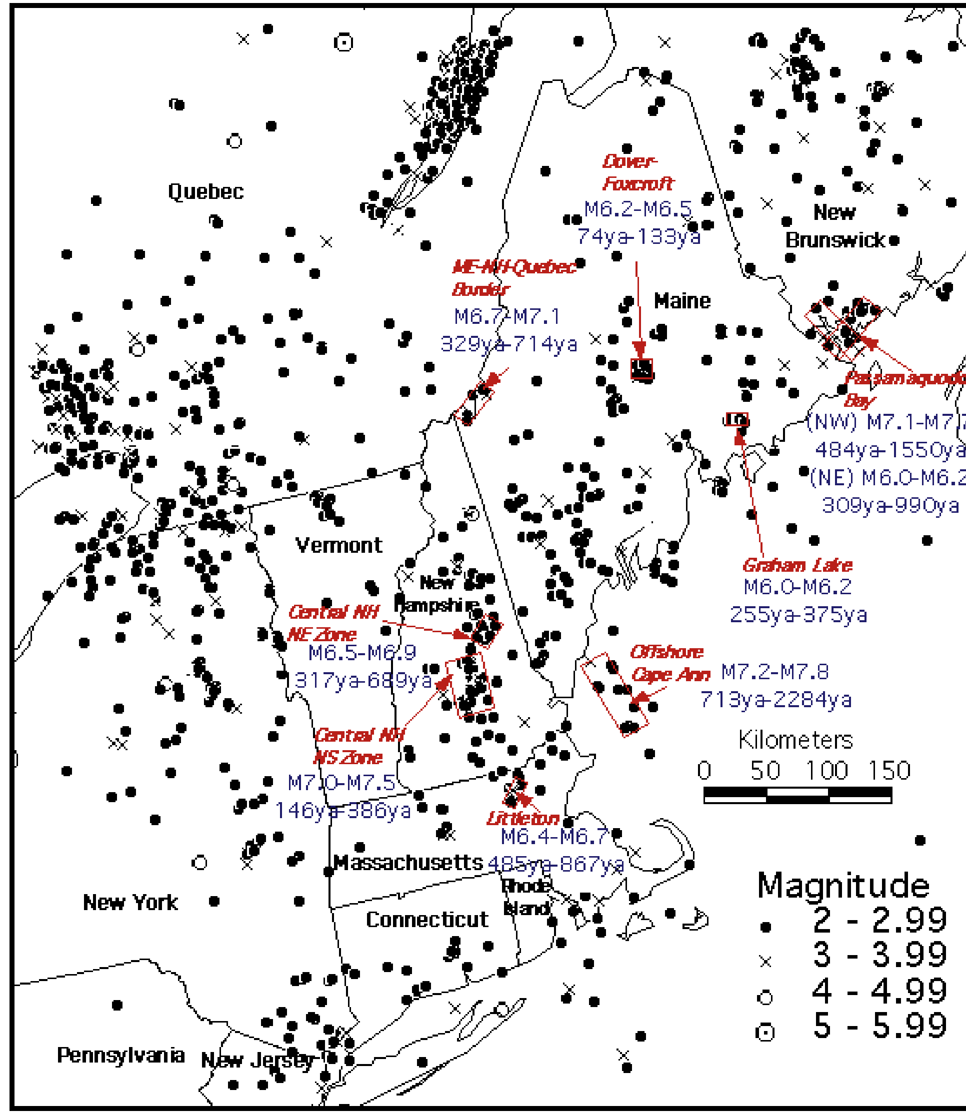

The 1663 Earthquake at Charlevoix, Quebec

This earthquake caused some minor chimney damage in Boston at a distance of almost 600 km. Several aftershocks were also were felt in Boston. Rev. Danforth in Roxbury, Massachusetts wrote:

"1662 Jan. 26 (O.S.) about 6 o'clock at night there happened an earthquake, wch shook mens houses and caused many to run out of their houses into the streets, and ye tops of 2 or 3 chimnyes fell off, or some part ym." (Danforth, 1880).

This appears to be MMI V to VI shaking (See *Ebel*, ***Seism. Res. Lett.***, *1996*).

Figure A14–4

The 1663 Earthquake at Charlevoix, Quebec

If the modern seismicity at Charlevoix is aftershocks of the 1663 event, then its rupture length must have been about 70 km. This suggests a magnitude of M7.1-7.5 based on the Wells and Coppersmith (1994) relations (i.e., a New Madrid-size event). (See *Ebel*, ***Seism. Res. Lett.***, *1996*).

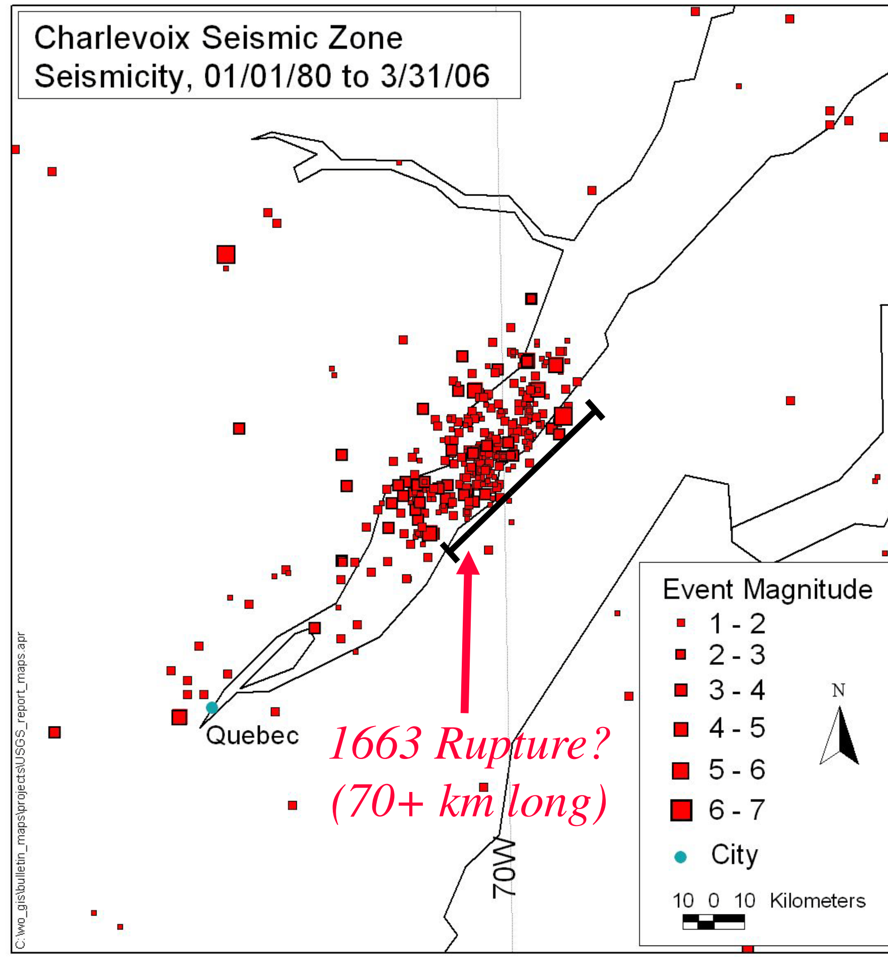

Figure A14–5

Estimated 1663 Magnitude from the MMI Estimate at Roxbury, Massachusetts

Roxbury 1663 Estimated Event Magnitude from the MMI Intensity at Roxbury, Massachusetts (Using Klimkiewicz, 1980)

If MMI = V at Roxbury, M = 7.0
If MMI = VI at Roxbury, M = 7.5
If MMI = VII at Roxbury, M = 8

The MMI estimate at Roxbury in 1663 suggests that this earthquake was M ~ 7.0 - 7.5.

Figure A14–6

Estimated 1663 Ground Motions at Roxbury, Massachusetts

Roxbury 1663 Estimated Ground Motions

	Atkinson and Boore (1995)		Toro et al. (1997)	
Mag.	Soil pga(g)	Soil SA0.3(g)	Soil pga(g)	Soil SA0.3(g)
7	0.0043	0.0222	0.0032	0.0097
7.5	0.0052	0.0311	0.0047	0.0146
8	0.0059	0.0403	0.0071	0.0213
Threshold for Chimney Damage	.01	.03	.01	.03

The chimney damage at Roxbury in 1663 suggests that this earthquake was M≥7.5.

Figure A14–7

Rates of M≥7 Earthquakes in the CEUS

Figure A14–8

For the CEUS and for ENA, the observed rate of M>~7 earthquakes is greater than expected from extrapolations of the Gutenberg-Richter curves from the smaller earthquake activity in these regions (*Nishenko and Bollinger,* ***Science****, 1990*).

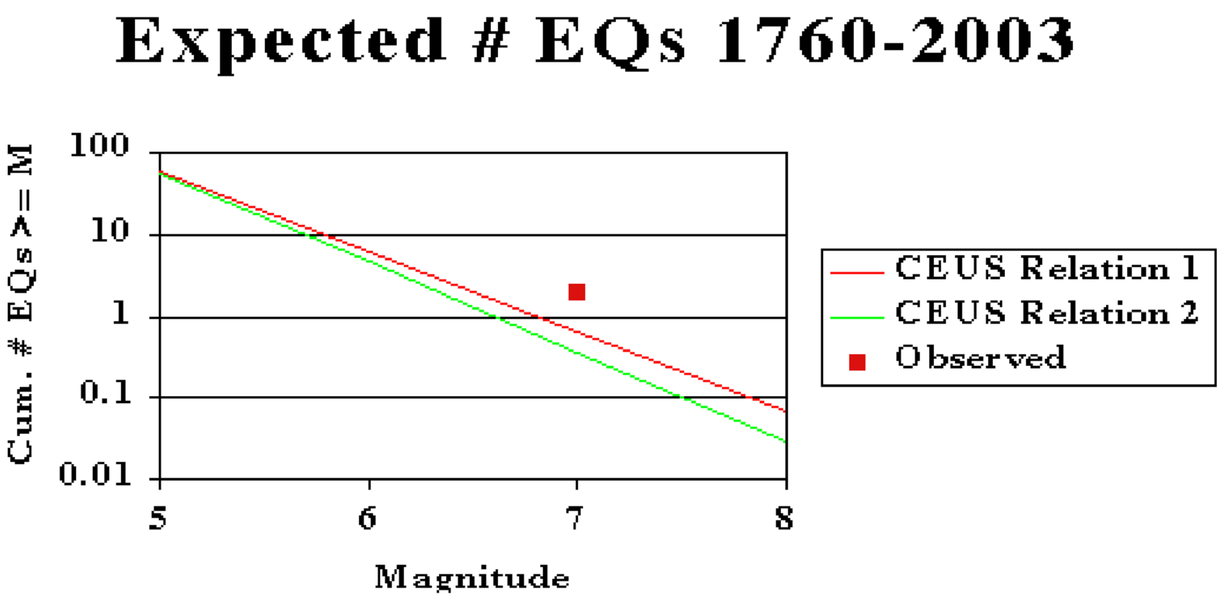

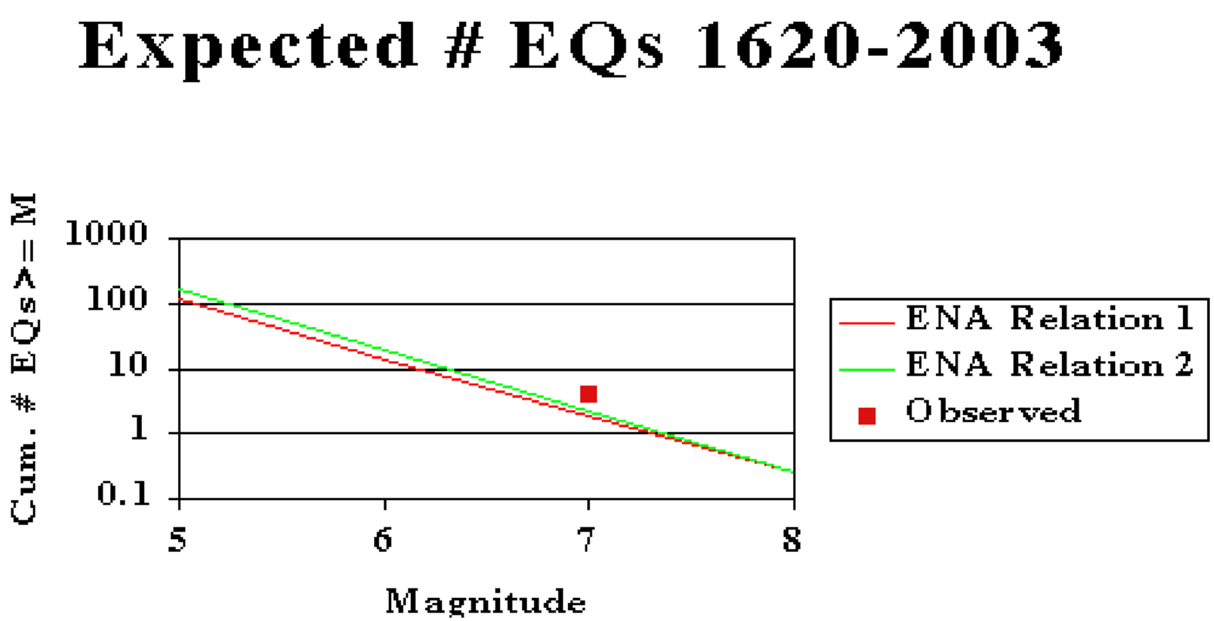

CEUS - Central and Eastern U.S.; ENA - Eastern North America

Building on the paleoseismicity idea that localized clusters of earthquakes in the CEUS delimit aftershock zones of past strong earthquakes, we can take the smaller earthquake activity and postulate locations where M>~7 earthquakes may have taken place in the past few thousand years.

Figure A14–9

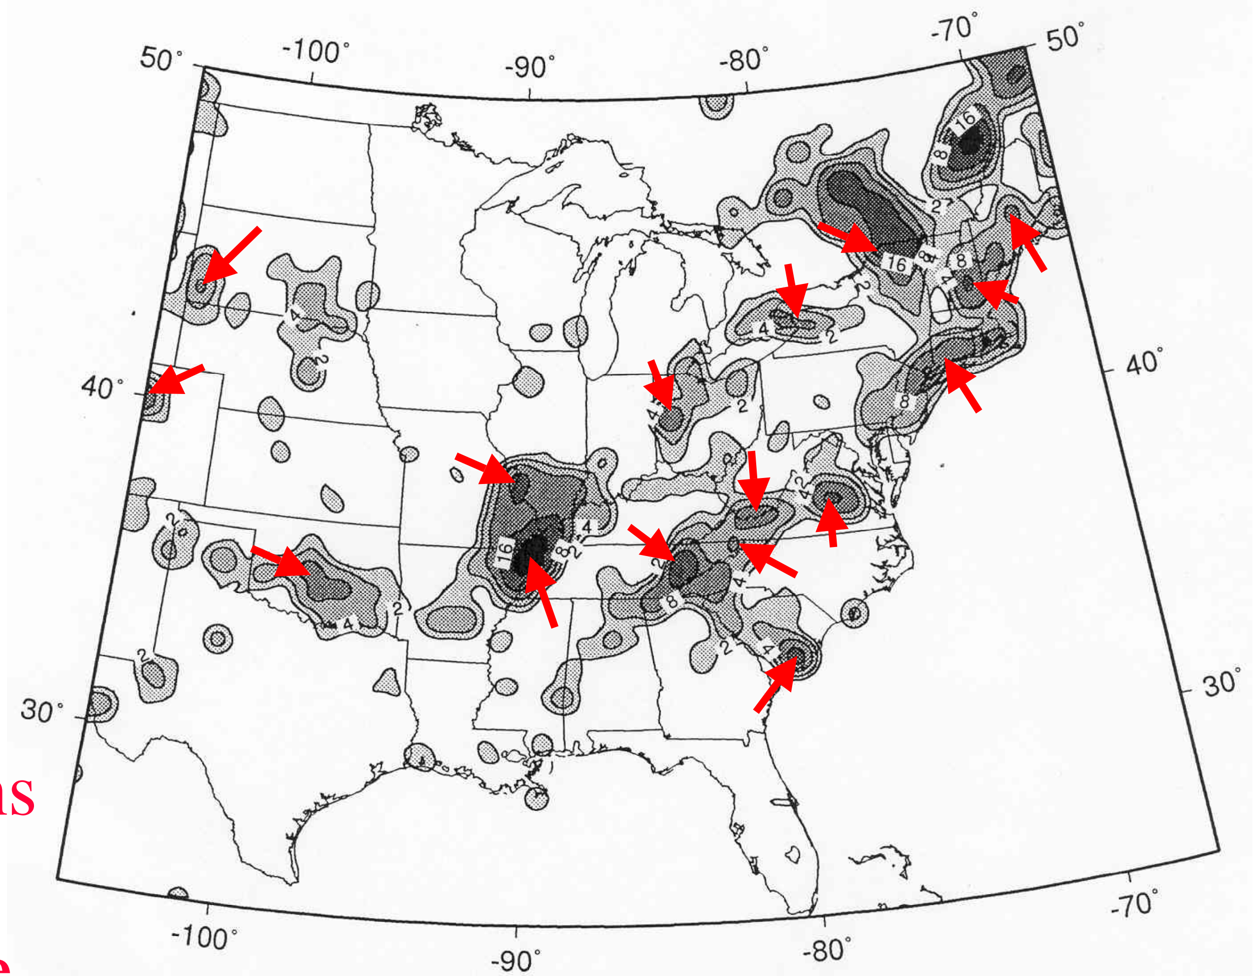

The red arrows show areas of enhanced, localized seismicity rates where the estimated rate of M=0 earthquakes per 60 years is greater than 8 (modified from *Frankel*, ***Seism. Res. Lett.***, *1995*).

M>~7 Seismicity Rates Underestimated for the CEUS?

If all of the CEUS modern seismicity clusters show locations of M>~7 during the past 2000 or so years, then the rate of M>~7 earthquakes is approximately 2 to 3 times greater than that found from extrapolations of the smaller seismicity to larger magnitudes.

Paleoseismicity Cluster Analysis Results
Main shocks between M7.0 and M7.5 Have Equal Rates

Rate of M=0 Earthquakes in 60 Years	*Time of Analysis (years)*	*Nishenko & Bollinger (1990) Relation 1 Recurrence Curve Prediction*	*Nishenko & Bollinger (1990) Relation 2 Recurrence Curve Prediction*	*Nishenko & Bollinger (1990) Relation 1 Cluster Analysis Prediction*	*Nishenko & Bollinger (1990) Relation 2 Cluster Analysis Prediction*
16 or more	**1118**	**3.0**	**1.7**	**8**	**8**
8 or more	**2124**	**5.7**	**3.1**	**15**	**15**

Table 2b
Paleoseismicity Cluster Analysis Results
Main shocks between M7.0 and M7.5 Have Gutenberg-Richter Distribution

Rate of M=0 Earthquakes in 60 Years	*Time of Analysis (years)*	*Nishenko & Bollinger (1990) Relation 1 Recurrence Curve Prediction*	*Nishenko & Bollinger (1990) Relation 2 Recurrence Curve Prediction*	*Nishenko & Bollinger (1990) Relation 1 Cluster Analysis Prediction*	*Nishenko & Bollinger (1990) Relation 2 Cluster Analysis Prediction*
16 or more	**1118**	**3.0**	**1.7**	**6**	**4**
8 or more	**2124**	**5.7**	**3.1**	**13**	**7**

Gutenberg-Richter Extrapolation

Paleoseismicity Extrapolation

Figure A15–1

Research Needs

- Paleoseismicity data
  - Paleoliquefaction (field studies, interpretation, uncertainties)
  - Lake deposits
  - Fragile features
- Historical data
  - Identification and information (newspapers, census data, etc.)
  - Analysis (site effects, intensity, etc)
  - Geologic effects
- Global analogues
  - What is an analogue? (comparison of NA & global, subdivision of global, statistical tests)
  - Indentify and incorporate uncertainty in magnitude
  - Removal of multiple events from database
  - Bias corrections
- Reactivation
  - Australian neo-tectonic features investigation
  - Study of reactivation to determine most likely fault types (CA foothills)
- Wells & Coppersmith update for SCR
- Strain rate data collection, lower crustal strain anomalies, LIDAR
- Fault segmentation, DEM/Google earth investigations, triggered seismicity

Figure A15–2

Method Resolutions

Mobs	Only useful as min M_{max} anchor or input to Bayesian method
~~Mobs±c~~	Lacks strong technical basis
~~Seismicity Rates~~	Not reliable at this time
~~M-f extrapolation~~	Not reliable at this time
Local geology/ rupture dimensions	useful as data input
NA analogues	Direct method and input to Bayesian method
Global analogues	Direct method and input to Bayesian method

Method Resolutions

Figure A15–3

Bayesian methods	Useful framework for M_{max}
Physics	Ultimate goal, not yet there as sole predictor, useful as input to data sets
GPS data	useful as data input
Statistics	M-f extrapolation useful as sanity check, min Mmax
~~Pattern recognition~~	Too little data currently, further research?
~~Qo~~	Q too unreliable a basis
~~Mb=7.5~~	No technical basis

Figure A15–4

M_{max} Methods

Global Analogues

Bayesian

Prior

- NA analogue
- Global analogue
- Physical principals
- paleoseismic

Likelihood function

- M_{obs}
  - Instrumental
  - Historical
  - paleoseismic
- Statistical analyses
- Geologic Features (paleoseismic)

Figure A15–5

Individual Methods	
Method	**Consideration for Weighting**
• Fault dimensions	• Most appropriate for well characterized faults/sources