

6

Promise and Paradox: Why Improved Knowledge of Plate Tectonics Has Not Yielded Correspondingly Better Earthquake Hazard Maps

Seth Stein,¹ Mian Liu,² Bruce D. Spencer,³ and Edward M. Brooks⁴

ABSTRACT

The discovery of plate tectonics in the 1960s offered the promise of a physical foundation for earthquake hazard assessment. The relative motion of tectonic plates concentrates stress along plate boundaries, where most earthquakes occur. Steady plate motions load plate boundary faults at constant rates, leading to cycles of stress buildup and release. Hence it seemed reasonable to expect quasi-periodicity in the recurrence time of large earthquakes, so if some segments of the plate boundary faults have not produced a large earthquake in the recent past, they are “seismic gaps” where a quake is due or overdue. In subsequent years, knowledge of the geometry and rates of plate motions has improved dramatically, in part due to the advent of space-based geodesy and recognition that many plate boundaries are diffuse zones. Similarly, advances in seismological data and methods provide much more information about what happens in large earthquakes. Paradoxically, even given all this knowledge, reliably assessing earthquake hazards remains difficult. Although earthquake hazard maps often do a good job of describing what occurs, in other cases large earthquakes occur in unexpected places and/or produce greater-than-expected shaking. The locations, times, and magnitude of large earthquakes turn out to be highly variable. Some of the variability can be addressed by using longer time series and knowledge of plate motions, but some reflects not-yet-understood and likely chaotic behavior. As a result, some key parameters required for earthquake hazard maps are poorly known, unknown, or unknowable. Although maps may be improved by better estimating some parameters, the fact that others cannot be much better estimated limits how good maps can be. Hence, in addition to trying to better assess hazards with new data and models, we can do better by recognizing and communicating the uncertainties involved. Agreed methods can be developed to assess how well a map performed, whether one map performed better than another, and when and how to update maps. Mitigation policies can be developed by considering the costs and benefits of various strategies, to yield sensible policies given the unavoidable uncertainties in hazard estimates. Thus, although from a scientific standpoint hazard maps can be viewed as half-empty glasses that we hope to fill somewhat further, from a societal view these maps can be viewed as glasses already half full.

*“With firm geological foundations and major earthquakes rare,
Fukushima is a safe and secure place to do business.”
Fukushima Prefecture Website [Cyranoski, 2011]*

¹Department of Earth and Planetary Sciences and Institute for Policy Research, Northwestern University, Evanston, Illinois, USA

²Department of Geological Sciences, University of Missouri, Columbia, Missouri, USA

³Department of Statistics and Institute for Policy Research, Northwestern University, Evanston, Illinois, USA

⁴Department of Earth and Planetary Sciences, Northwestern University, Evanston, Illinois, USA

6.1. INTRODUCTION

On 26 January 2015, the National Weather Service (NWS) and television weather forecasts warned residents of the northeastern United States of a “life-threatening,” “extremely dangerous,” and “potentially historic” snowstorm. New York mayor Bill de Blasio warned the city’s residents to “prepare for something worse than we have seen before” and banned all street travel, including (to residents’ horror) take-out food delivery. The satirical Onion website reported, “NYC Mayor: ‘Reconcile Yourselves With Your God, for All Will Perish in the Tempest.’” New York governor Andrew Cuomo ordered the city’s subway system shut down in advance of the storm, the first such shutdown in its 111-year history. Residents emptied stores of milk, bread, eggs, and toilet paper. Amorous New Yorkers placed online advertisements for “blizzard buddies.”

In reality, the threatened “snowmageddon” produced snowfall that ranked 36th in the city’s 125-year record. Streets and subway quickly reopened. Weather forecasters quickly started explaining their missed forecasts, admitting uncertainties that had not been mentioned. The meteorologist in charge of the NWS station in Mount Holly, NJ, acknowledged, “You made a lot of tough decisions expecting us to get it right, and we didn’t. Once again, I’m sorry.” The NWS said “rapidly deepening winter storms are very challenging to predict, specifically their track and how far west the heaviest bands will move. These bands are nearly impossible to predict until they develop. Our science has come a long way, but there are still many moving parts in the atmosphere, which creates quite the forecast challenge” [*Santora and Fitzsimmons*, 2015]. NWS admitted choosing weather models that predicted major snowfall in the city over their newest one that didn’t. However, the Associated Press reported that the NWS director “wouldn’t say his agency’s forecast was off. Instead, he blamed the way meteorologists communicated and said the weather service needs to do a better job addressing uncertainty” [*Bornstein and Mulvihill*, 2015].

To seismologists interested in forecasting earthquake hazards, this sequence seemed familiar. Despite advances in science driven by improved data and models, a forecast proved wrong due to larger-than-admitted uncertainties, causing problems for the impacted communities [*Barro*, 2015; *Flegenheimer*, 2015]. This paper is an overview of why, despite advances in understanding plate motions, crustal deformation, and earthquakes, seismic-hazard assessments often do not do as well as we would like. It also suggests some ways to do better.

6.2. PROMISE

A major goal of the commission that studied the 1906 San Francisco earthquake [*Lawson and Reid*, 1908] was to find “evidence upon which a judgment might be based as to the probability of recurrence of the earthquake in the future.” They identified the San Andreas fault as the source of the earthquake, and developed the idea of how it happened using a model of elastic rebound, in which motion on opposite sides of a locked fault produces strain that accumulates and is eventually released in earthquakes (Fig. 6.1). Based on this idea, they recommended that geodetic measurements be made to measure the motion across the fault over time as strain built up before a future earthquake. They suggested that “we should build a line of piers, say a kilometer apart” across the fault so “careful determination from time to time of the directions of the lines joining successive piers ... would reveal any strains which might be developing.”

This pioneering and insightful analysis left key questions unresolved. Why was the crust moving on opposite sides of the fault? Why was the fault oriented essentially parallel to the motion? If strain were accumulating, how would one know when it would give rise to an earthquake?

As a result, this promising approach to assessing earthquake hazards remained an abstract idea until the discovery of plate tectonics in the 1960s. Plate tectonics showed that most large earthquakes occur at plate boundaries as a result of motion between plates. Hence the geometry of faults reflects plate motions, and the long-term rate of motion across plate boundary faults is given by plate motions. In subsequent years, the advent of space-based geodesy permits determination both of motions between plates and of the motion giving rise to strain accumulation across individual fault segments. The resulting rates are consistent with those inferred from geological plate motion data (Fig. 6.1). In addition, space-based geodesy confirmed that many plate boundaries are diffuse zones, as suggested by the distribution of seismicity, topography, and active faulting [*Gordon and Stein*, 1992]. In these cases, the motion taken up on individual faults is not constrained by the plate motion, but can be measured geodetically, whereas net motion across the zone sums to the plate motion [*Kreemer et al.*, 2012].

A parallel important development has been new seismological data and methods for studying earthquake rupture. As a result, we now know a great deal about the fault and slip geometry in large plate boundaries. For example, combination of seismological, geodetic, and tsunami data make the giant 2011 Tohoku earthquake possibly the best-studied to date [*Lay et al.*, 2011; *Simons et al.*, 2011].

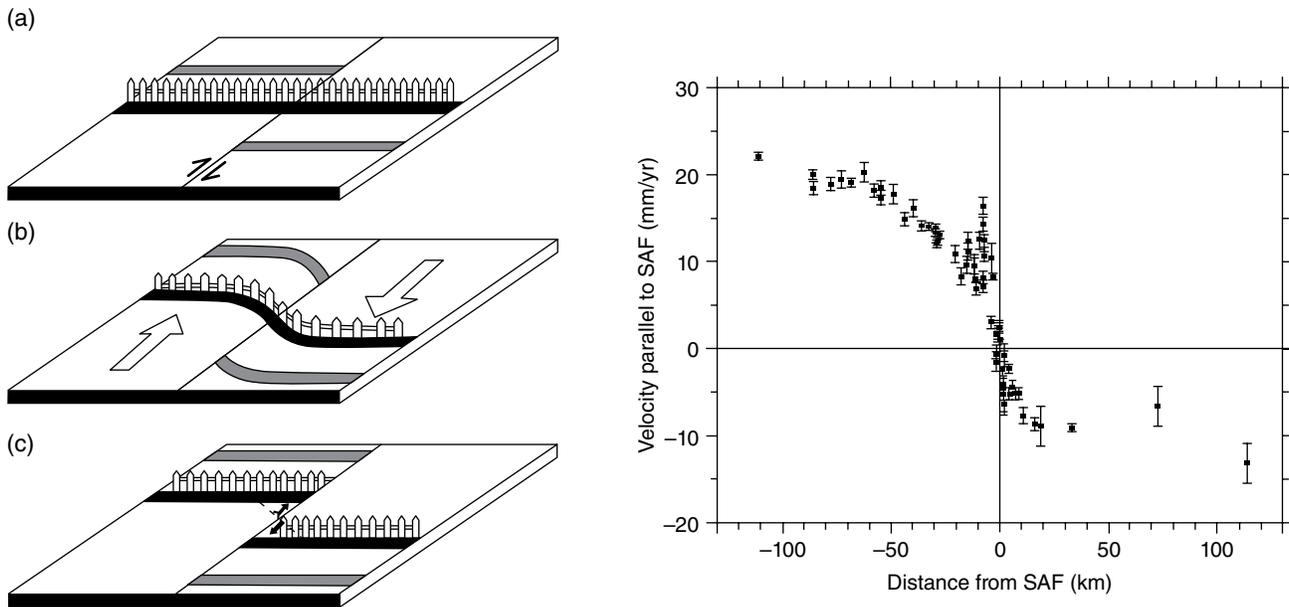


Figure 6.1 (Left) How elastic rebound works is shown by the history of a fence across a fault. (Right) GPS data showing strain accumulation across the San Andreas fault [Stein and Wysession, 2003].

6.3. PARADOX

Paradoxically, using the rapidly growing knowledge of plate boundaries and the motion on them in efforts to assess seismic hazards has yielded results that are less satisfactory than might be hoped. This paper reviews some of the recent work on the topic to explore why.

Initial efforts used the seismic gap concept, in which a gap exists on a plate boundary fault segment when it has been quiescent long enough since it produced the last major earthquake that another is overdue, more likely to occur there than elsewhere on the boundary fault zone. Although the model appeals to seismological instincts, the data do not show that gap models predict where large earthquakes will occur significantly better than a time-independent Poisson model in which the time since the last earthquake has no effect. A global test of the seismic gap hypothesis, which examined how well a gap map [McCann *et al.*, 1979] forecast the locations of major earthquakes, found that the map did no better than random guessing [Kagan and Jackson, 1991]. In fact, many more large earthquakes occurred in areas identified as low risk than in the presumed higher-risk gaps. Hence, despite its attraction, the cycle-gap model often does not yield useful forecasts [Kagan *et al.*, 2012].

Current hazard assessments rely on seismic hazard maps, which are used to develop codes for earthquake-resistant construction. Most maps are produced with the

probabilistic seismic hazard assessment (PSHA) algorithm, which uses estimates of the probability of different future earthquakes and the resulting shaking to predict the maximum ground shaking expected with a certain probability over a given time [Cornell, 1968; Field, 2010]. Larger expected shaking corresponds to higher predicted hazard. These maps are derived by estimating a variety of parameters for models that are chosen and used to forecast future seismicity and the resulting shaking. Estimates of many of these parameters for earthquakes in plate boundary zones explicitly or implicitly use the available knowledge about plate geometry and plate motions.

Because in many places the resulting maps seem sensible in general terms, with highest predicted hazard usually on recognized plate boundary faults where large earthquakes have occurred or are expected, they are widely accepted and used to make costly policy decisions. However, their predictions have never been objectively tested, primarily because on any given fault segment large earthquakes are infrequent. We have no real idea of their uncertainties or how well they predict what actually happens. Hence, the fact that they sometimes do poorly (as is becoming clear in recent years) is not surprising. In some cases, hazards are underestimated, and in others, they are overestimated. Either is undesirable; one exposes communities to undue risk, whereas the other diverts resources that could do more good if used otherwise.

The problem is illustrated by images of the tsunami from the giant 2011 Tohoku earthquake pouring over 10-meter seawalls. The Japanese hazard mappers used the historic earthquake record to divide the trench, along which the Pacific plate subducts beneath Japan, into segments about 150 km long and infer the largest earthquake to expect on each in the next 30 years. For example, off Fukushima at most magnitude 7.4 was expected. The resulting map (Fig. 6.2) predicted less than 0.1% probability of shaking with intensity “6-lower” on the Japan Meteorological Agency scale in the next 30 yr off Tohoku. Thus, such shaking was expected on average only once in the next 30/0.001 or 30,000 yr. However, within 2 yr, such shaking occurred. On 11 March 2011 five segments broke together, causing a magnitude 9.1 earthquake much larger than expected (Fig. 6.3) and a tsunami larger than anticipated. The mapping process significantly underpredicted what happened [Cyranoski, 2011; Geller, 2011; Sagiya, 2011]. Geller [2011] noted that the Tohoku area was shown as having significantly lower hazard than other parts of Japan, notably the Tokai, Tonankai, and Nankai districts to the south. Moreover, the hazard map showed low hazard in the areas of all earthquakes in the previous years that caused ten or more fatalities. He thus argued that “all of Japan is at risk from earthquakes, and the present state of seismological science does not allow us to reliably differentiate the risk level in particular geographic areas,” so a map showing uniform hazard would be preferable to the existing map.

Highly destructive earthquakes have occurred in other areas that hazard maps predict to be relatively safe. The 2008 M 7.9 Wenchuan, China, and 2010 M 7.1 Haiti earthquakes occurred on faults mapped as giving rise to low hazard [Stein *et al.*, 2012]. As *Science* magazine [Kerr, 2011] explained, “the seismic crystal ball is proving mostly cloudy around the world.”

These events stimulated discussions among seismologists and earthquake engineers about earthquake hazard mapping practices [Gulkan, 2013; Stein *et al.*, 2012; Stirling, 2012]. The underlying question is the extent to which the occurrence of low probability shaking indicates problems with the maps or chance occurrences. One explanation [Frankel, 2013; Hanks *et al.*, 2012] is that these earthquakes are low-probability events allowed by probabilistic seismic hazard maps. Some such events are expected, just as although the chance that a given lottery ticket is a winner is low, the probability that some ticket wins is high. However, the common practice of extensively remaking maps to show increased hazards after “unexpected” events or shaking (Fig. 6.4) is inconsistent with the interpretation that these were simply low-probability events consistent with the map [Stein *et al.*, 2015a]. In a lottery, the odds of winning are only reassigned after a winning ticket is picked when the operators think their prior model was wrong.

6.4. THE MAPPING CHALLENGE

The challenge for hazard map making is choosing hundreds or thousands of parameters to predict the answers to four questions over periods of 500–2500 years: Where will large earthquakes occur? When will they occur? How large will they be? How strong will their shaking be?

All of these are difficult to reliably estimate. The first is largely a tectonic question. On plate boundaries we expect large earthquakes to occur eventually along most of the boundary, whose geometry is reasonably well known. However, in diffuse plate boundary zones and even more so in plate interiors, earthquakes can occur in unexpected places. When earthquakes will occur and how large they will be are fundamental questions of earthquake physics that are not understood, namely how plate motion is released in earthquakes. Earthquakes are often larger (Tohoku) or smaller (Nepal 2015) than expected, and we have no useful ideas about when earthquakes will recur beyond estimates of long-term average recurrence intervals. Ground motion for large earthquakes is somewhat predictable in areas that are active enough that modern seismological records exist, and hard to reliably infer elsewhere.

Thus, although general statements can be made (e.g., large earthquakes are most common at subduction zones), plate boundaries are generally more active than plate interiors, Los Angeles is more hazardous than Chicago, making detailed forecasts for specific areas (e.g., hazard maps) challenging. The more specific a forecast we want, the more challenging it is.

Some of the parameters required are reasonably well known, some are somewhat known, some are essentially unknown, and some may be unknowable [Stein *et al.*, 2012; Stein and Friedrich, 2014]. As a result, mappers combine data and models with their sense of how the Earth works. Such models, which of necessity require subjective assessments and choices among many poorly known or unknown parameters, are examples of what risk analysts term BOGSATs, from “Bunch Of Guys Sitting Around a Table” [Kurowicka and Cooke, 2006]. In Freedman and Stark’s [2003] words, this involves “geological mapping, geodetic mapping, viscoelastic loading calculations, paleoseismic observations, extrapolating rules of thumb across geography and magnitude, simulation, and many appeals to expert opinion. Philosophical difficulties aside, the numerical probability values seem rather arbitrary.”

The resulting maps thus have large uncertainties, in that different plausible assumptions about key parameters yield quite different hazard maps (Fig. 6.5). Not surprisingly, sometimes maps do well at predicting what occurs in future earthquakes, and sometimes they do poorly. Hence

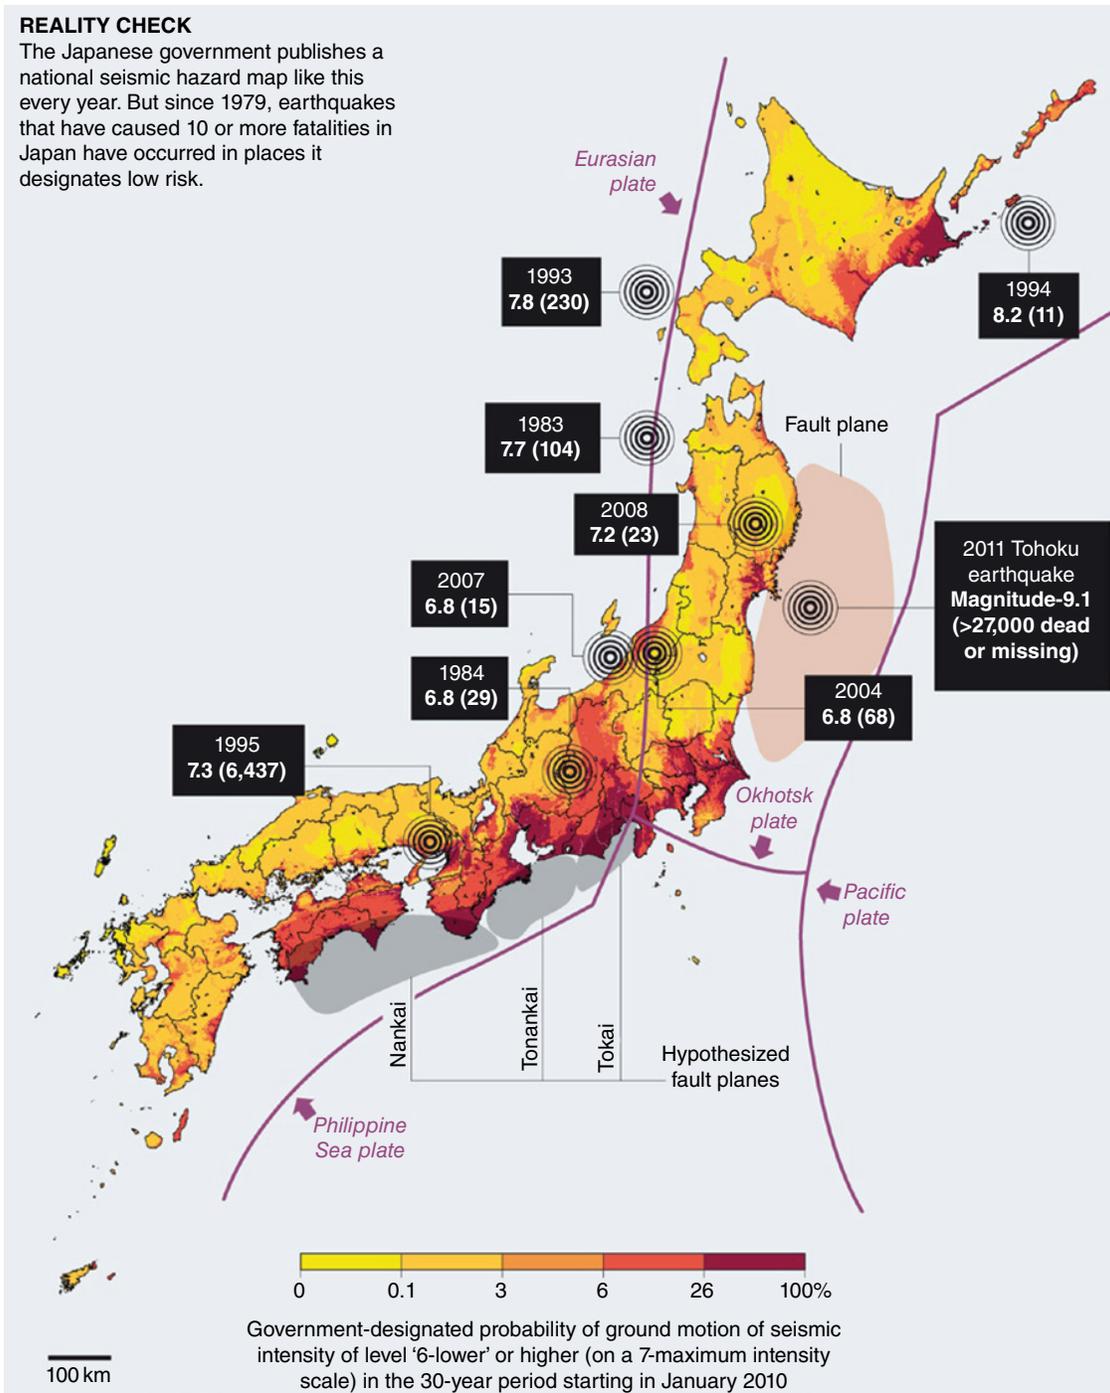


Figure 6.2 Comparison of Japanese government hazard map to the locations of earthquakes since 1979 that caused 10 or more fatalities. Hazard is shown as probability that the maximum ground acceleration (shaking) in any area would exceed a particular value during the next 30 yr. Larger expected shaking corresponds to higher predicted hazard. The Tohoku area is shown as having significantly lower hazard than other parts of Japan, notably areas to the south. Since 1979, earthquakes that caused 10 or more fatalities occurred in places assigned a relatively low hazard [Geller, 2011; Reproduced with permission of *Nature*].

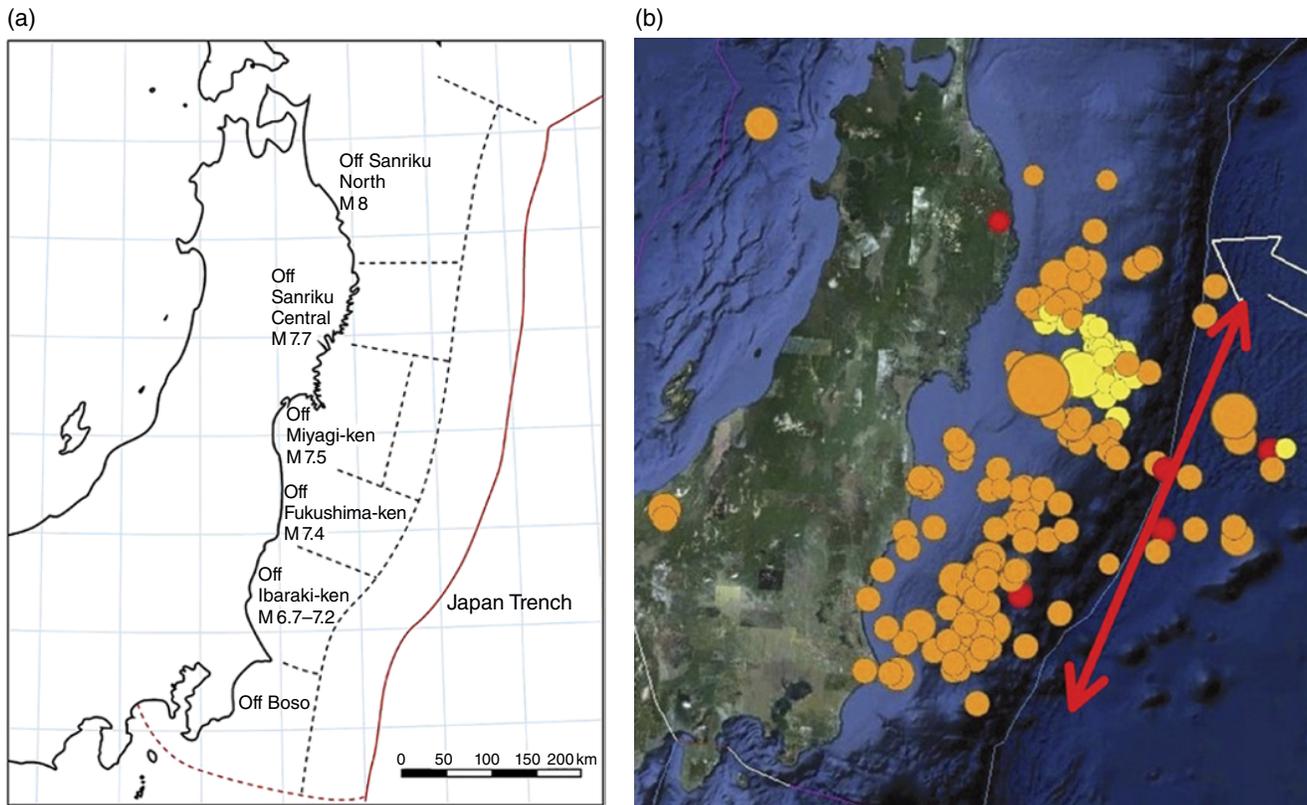


Figure 6.3 Comparison of the trench segments and corresponding maximum earthquake magnitudes assumed in the Japanese hazard map (left) to the aftershock zone of the 11 March 2011 earthquake (right), which broke five segments [Stein and Stein, 2014].

questions arise as to how to do better and how much better is practical.

In some cases, as discussed shortly, maps do poorly because they do not yet incorporate data or concepts that are available. In other cases, the problem is that where and when large earthquakes will happen and how big they will be are more variable than assumed in the hazard maps. The conceptual model used in hazard mapping comes from studies of plate boundaries, where steady motion between plates loads a plate boundary fault rapidly at constant rate. In this case, the largest earthquakes should ideally be spatially focused on the fault, temporally quasi-periodic, and have similar magnitudes (Fig. 6.6). However, these assumptions are challenged by the growing body of evidence showing that fault ruptures are highly nonlinear, spatially and temporally. As a result, the concepts of recurrence time, seismic gaps, and characteristic earthquakes, while somewhat useful in general terms for thinking about basic earthquake mechanics, are inadequate for reliable hazard analysis. Some earthquakes appear where and when they were not expected and others are much larger than expected.

Some of the space-time variability of large earthquakes can be addressed by using longer time series and knowledge of plate motions, but some results from not-yet-understood and likely chaotic behavior. Table 6.1 gives an assessment of how much better various parameters could be estimated as new data and models become available. As discussed next, some uncertainties are reducible, whereas others seem likely to remain [Stein and Friedrich, 2014].

6.5. UNCERTAINTIES: SHALLOW VERSUS DEEP

Because hazard maps seek to describe unknown future events, characterizing the sources of uncertainty is crucial. Various formulations are available. Seismic hazard analysis follows engineering literature in distinguishing uncertainties by their sources. Aleatory uncertainties are due to irreducible physical variability of a system. Epistemic uncertainties are due to lack of knowledge of the system, and so can be reduced by more knowledge.

Uncertainties can also be described as being shallow or deep [Cox, 2012; Hallegatte et al., 2012]. Typically,

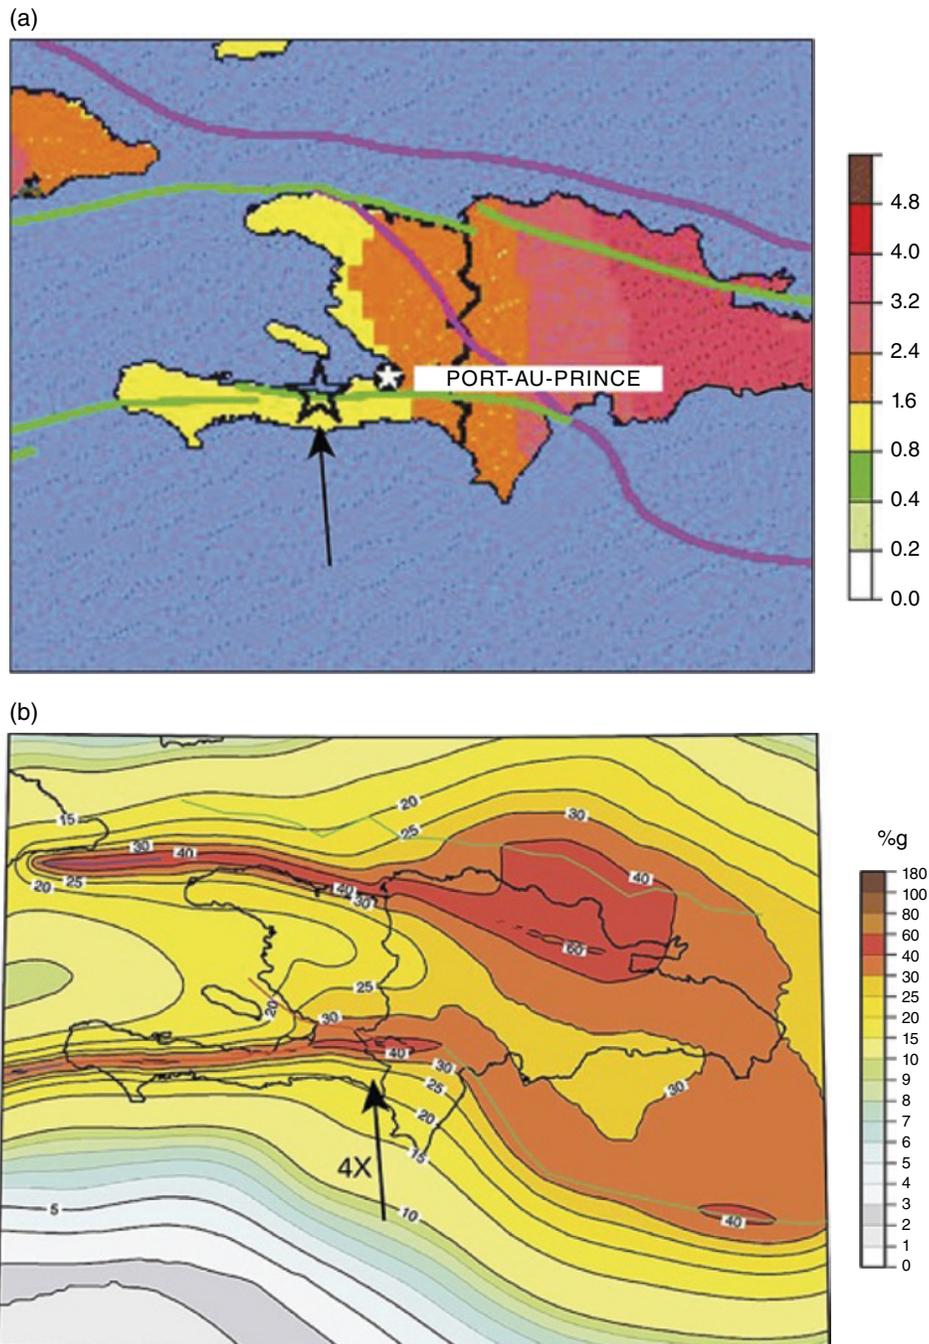


Figure 6.4 Comparison of seismic hazard maps for Haiti made before [GSHAP, Global Seismic Hazard Assessment Program, 1999] and after [Frankel et al., 2010] the 2010 M 7.1 earthquake. The revised map shows a factor of four higher hazard on the fault that had recently broken in the earthquake.

scientists consider shallow uncertainty, recognizing they do not know the outcomes, but assuming they know a probability density function describing them. In this case, models based on a system's past are good predictors of the future. The alternative is deep uncertainty in which the probability density function is unknown, so models

based on a system's past are likely to be poor predictors of the future [Stein and Stein, 2013a].

In sports terms, a baseball player's batting average describes shallow uncertainty, because it is a good predictor of the chance he will get a hit. Deep uncertainty is involved in trying to predict the winner of next year's

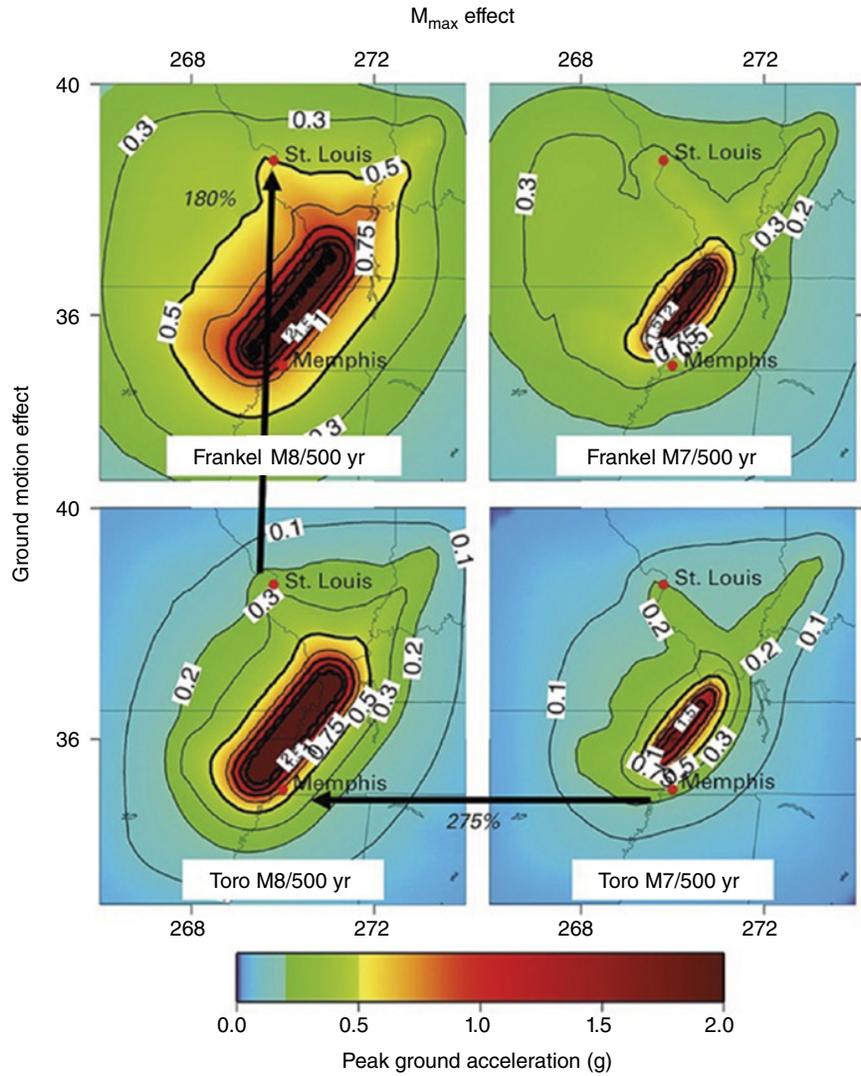


Figure 6.5 Comparison of the different hazard maps (2% probability in 50 yr) showing the effect of assuming different ground motion prediction equations (columns) and maximum magnitudes (rows) of the main New Madrid fault source [Newman *et al.*, 2001].

World Series. The teams’ past performance provides only limited insight into the future of a complicated process. We could develop various models based on the past performance, but would place little confidence in them.

6.6. SPACE-TIME VARIABILITY

Because large earthquakes on a given fault segment occur hundreds or thousands of years apart on average, the short records from seismology (about a hundred years) and historical accounts (hundreds to thousands of years) are often inadequate to show where to expect the next large earthquake.

To see this, consider the coast of North Africa (Fig. 6.7), part of the slow-converging boundary between west

Africa (Nubia) and Eurasia. During the time period over which we have good seismological data, roughly the past century, only parts of the boundary have had magnitude 7 earthquakes. However, the convergence rate is similar along the boundary. A simulation assuming that these occur randomly along the margin at their recent rate yields apparent concentrations of large earthquakes and seismic gaps for records up to thousands of years long. These artifacts can bias hazard assessment either to assume that areas with recent large events are more dangerous, or conversely that areas without recent large events are dangerous “gaps” “overdue” for earthquakes. Both biases are common (e.g., Fig. 6.2).

In the simulation, approximately 8000 years of record is needed to show that the seismicity is uniform. Any

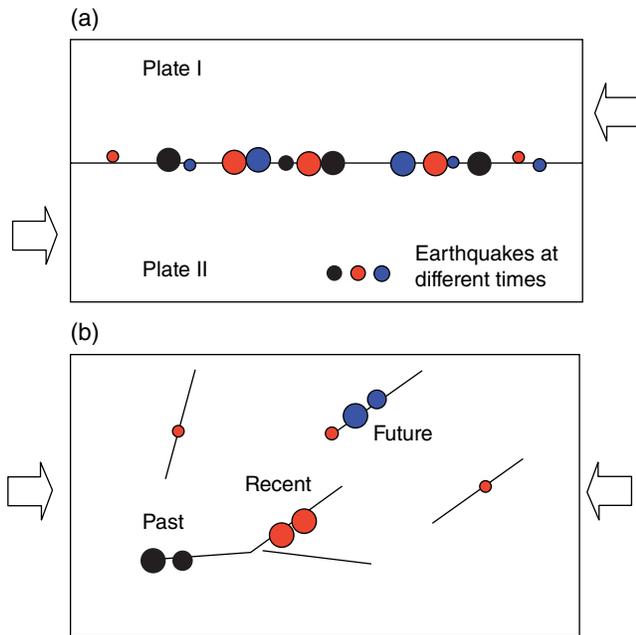


Figure 6.6 Cartoon showing the difference between earthquakes (a) at plate boundary faults and (b) in midcontinents. The plate boundary fault is loaded at a constant rate by the steady relative plate motion, causing quasi-periodic earthquakes to concentrate along the plate boundary. The midcontinent is loaded from the far field, and the loading is shared by a complex system of interacting faults. Hence, on each fault, the loading rate may be variable, and earthquakes may shut off on one fault and migrate to another [Liu *et al.*, 2011].

Table 6.1 Earthquake Hazard Assessment Uncertainties

Cause of uncertainty	How much can the uncertainty be reduced?
Where will large earthquakes occur?	Significantly on plate boundaries, somewhat in interiors
When will large earthquakes occur?	Little if at all
How large will they be?	Significantly for lower bound (paleoseismology), not for upper (short sample)
How strong will the shaking be?	Significantly in seismically active areas, less so in less active one

shorter sample (like that available today) would give a distorted view. If the seismicity and thus hazard are uniform, a hazard map produced using the seismic record alone will overestimate the hazard where previous large earthquakes occurred and underestimate it elsewhere. Hence, the 1999 Global Seismic Hazard Map, showing shaking expected at 10% probability in 50 years, features a prominent “bull’s-eye” at the site of the 1980 M 7.3 El Asnam earthquake. The largest subsequent earthquakes to date, the 2003 M 6.8 Algeria and 2004 M 6.4 Morocco

events, did not occur in the regions designated as high hazard (Fig. 6.8).

Moreover, the plate motion could be released in earthquakes or occur in part aseismically. This issue arose for Tohoku, where the presumed absence of giant earthquakes on the trench was implicitly interpreted as indicating that much of the subduction occurred aseismically. The Kurile trench, just to the north, seemed to show this discrepancy. The largest seismologically recorded earthquakes there are magnitude 8, which only accounts for about one third of the plate motion. Hence it had been assumed that most of the subduction occurred aseismically [Kanamori, 1977]. However, more recently discovered deposits from ancient tsunamis show that much larger earthquakes had happened in the past [Nanayama *et al.*, 2003], accounting for much of the subduction that had been thought to occur aseismically. In hindsight, the same applied off Tohoku.

Geodetic data can help address this problem by showing whether strain is accumulating and how fast. For Tohoku, GPS data were recognized as showing a much higher rate of strain accumulation on the plate interface than would be expected if a large fraction of the subduction occurred aseismically [Loveless and Meade, 2010]. Including these data in the hazard analysis would have strengthened the case for considering the possibility of large earthquakes.

In such cases, using a longer record from historical accounts and paleoseismology can give a better view of where to expect earthquakes. In the decade prior to the 2011 Tohoku earthquake, increasing attention was being paid to data showing that large tsunamis had struck the area in 869 [Minoura *et al.*, 2001], 1896, and 1933. Some villages had stone tablets marking the heights reached by previous tsunamis and warning “Do not build your homes below this point” [Fackler, 2011]. However, these data were not yet incorporated into the hazard map.

Similarly, the 2010 Haiti earthquake occurred on a fault mapped in 2001 as having low hazard because the map was based on recent seismicity (Fig. 6.4). Much higher hazard would have been predicted by considering the long-term earthquake history of faults in the area and GPS data showing strain accumulating across them [Manaker *et al.*, 2008].

The challenge is even greater within plates [Leonard *et al.*, 2007, 2014; Stein *et al.*, 2009]. In these situations, tectonic loading is collectively accommodated by a complex system of interacting faults, so the loading rate on a given fault is slow and may not be constant. As a result, earthquakes can cluster on a fault for a while and then shift elsewhere (Fig. 6.6) [Camelbeek *et al.*, 2007; Li *et al.*, 2009].

A striking example is a 2000-year record from North China showing roaming of large earthquakes between

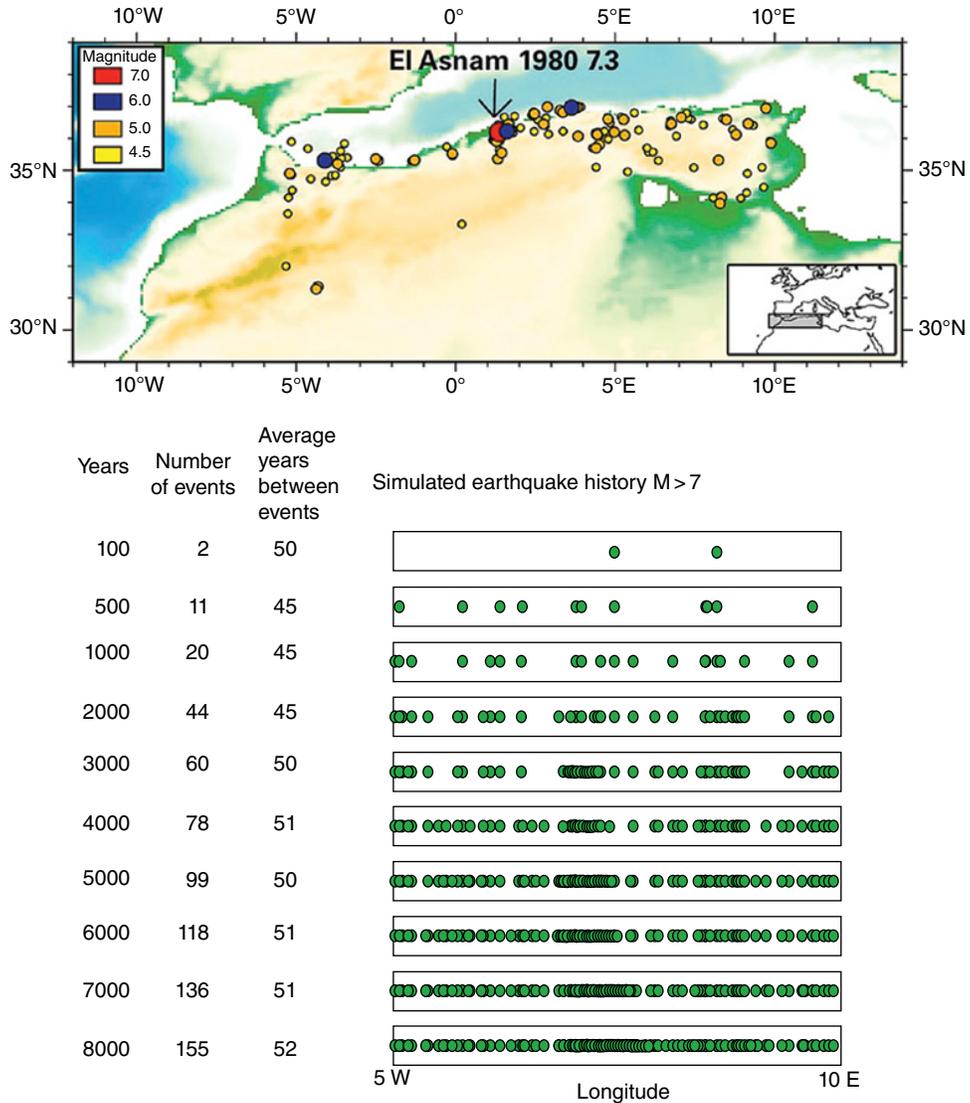


Figure 6.7 (Top) Seismicity along the North Africa plate boundary for 1963–2004. (Bottom) Simulations using a frequency-magnitude relation derived from these data predict that if seismicity is uniform in the zone, about 8000 yr of record is needed to avoid apparent concentrations and gaps [Swafford and Stein, 2007].

fault systems over a broad region, such that no large earthquake ruptured the same fault segment twice during the past 2000 years [Liu *et al.* 2011]. A map made from any short subset of the record would be biased. For example, one using the 1900 years prior to 1950 would miss recent activity including the 1976 Tangshan earthquake, which occurred on a previously unknown fault and killed nearly 240,000 people. Moreover, even a long past record may not indicate what will happen.

In such situations, we cannot use plate motion data to predict where strain will accumulate. Geodetic data can show strain accumulating. However, how to interpret faults on which little or no strain is accumulating is unclear, because faults may release strain that

accumulated over very long periods of time. An example is the current seismicity in the New Madrid seismic zone (NMSZ) in the central United States, that appears to be aftershocks of a cluster of $M \sim 7.0$ events in 1811–1812 [Stein and Liu, 2009]. These large events and similar events in the past millennium release strain much faster than GPS shows strain accumulates today [Calais and Stein, 2009; Craig and Calais, 2014], suggesting that they result from recent fault activation that releases prestored strain energy in the crust. If so, this earthquake sequence is similar to aftershocks in that the rates of energy release should decay with time and the sequence of earthquakes will eventually end (Fig. 6.9). Estimation of the duration of large earthquakes from this transient energy release

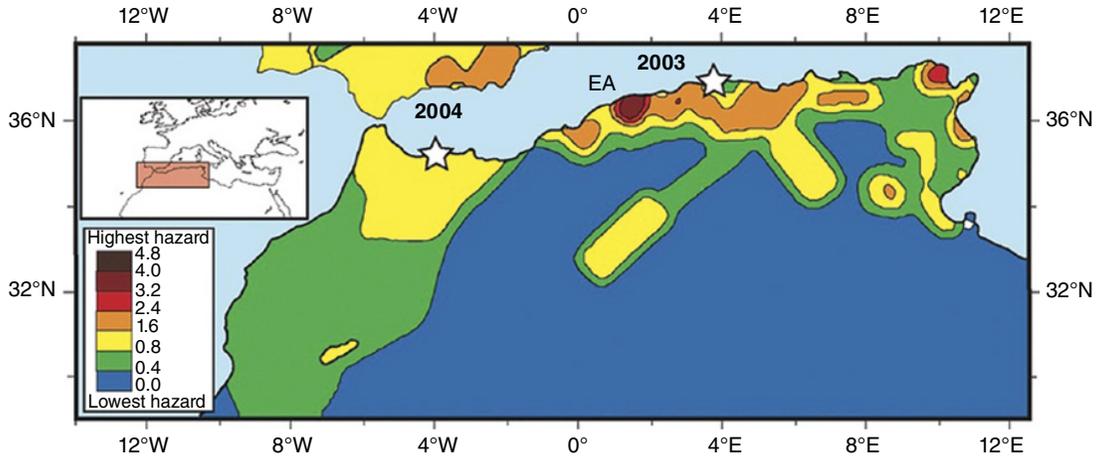


Figure 6.8 Portion of Global Seismic Hazard Map (1999) for North Africa, showing peak ground acceleration in m/s^2 expected at 10% probability in 50 yr. Note prominent “bull’s-eye” at site of the 1980 M_s 7.3 El Asnam earthquake (EA). The largest subsequent earthquakes to date, the May 2003 M_s 6.8 Algeria and February 2004 M_s 6.4 Morocco events (stars), did not occur in the predicted high-hazard regions [Swafford and Stein, 2007].

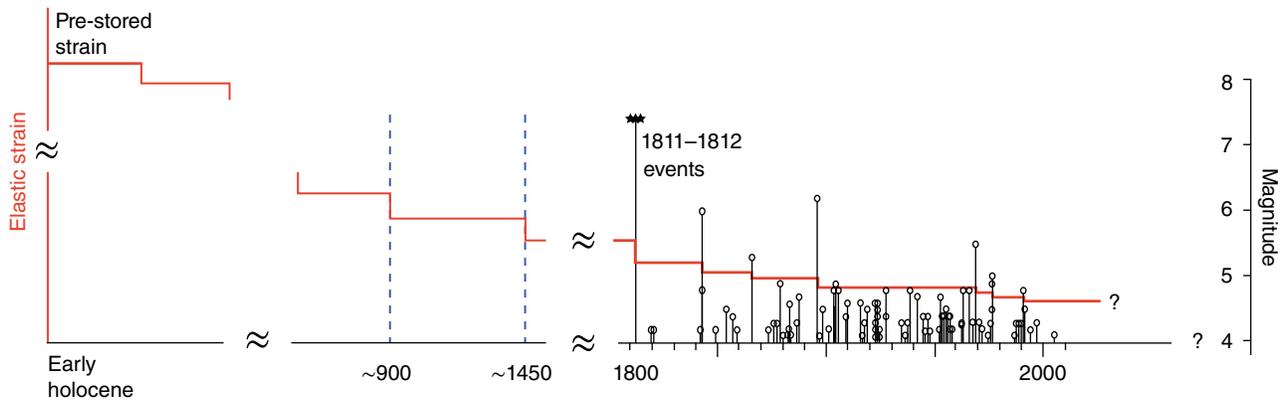


Figure 6.9 Conceptual model of a decaying NMSZ earthquake sequence, showing large earthquakes during 1811–1812 and similar events around 900 and 1450, and smaller events since 1812 [Liu et al., 2014].

shows that within uncertainties of model parameters, it is plausible that the NMSZ’s large earthquakes are ending now [Liu et al., 2014].

6.7. MAXIMUM MAGNITUDE

The Tohoku, Haiti, and Wenchuan earthquakes were surprising because they were much larger than the assumed magnitude of the largest future earthquakes expected on a fault or in an area, termed M_{max} . Unfortunately, knowledge of plate motions offers little insight into M_{max} , because even where we know the long-term rate of motion across a plate boundary fault or the deformation rate across an intraplate zone, neither predict how strain will be released. As a result, quite different estimates can be made [Kagan and Jackson, 2013].

Attempts to relate maximum magnitude to plate tectonics have not been successful. An analysis in 1980 of the largest known earthquakes at different subduction zones (Fig. 6.10) showed a striking pattern: magnitude 9 earthquakes occurred only where lithosphere younger than 80 million years was subducting rapidly, faster than 50 mm/yr [Ruff and Kanamori, 1980]. This result made intuitive sense, because both young age and high speed could favor strong mechanical coupling at the interface between the two plates. Because oceanic lithosphere cools as it moves away from a ridge and ages, young lithosphere is less dense and thus more buoyant. Similarly, faster subduction should increase frictional stress at the interface. The stronger coupling was, in turn, assumed to give rise to larger earthquakes when the interface eventually slips in a great thrust-fault

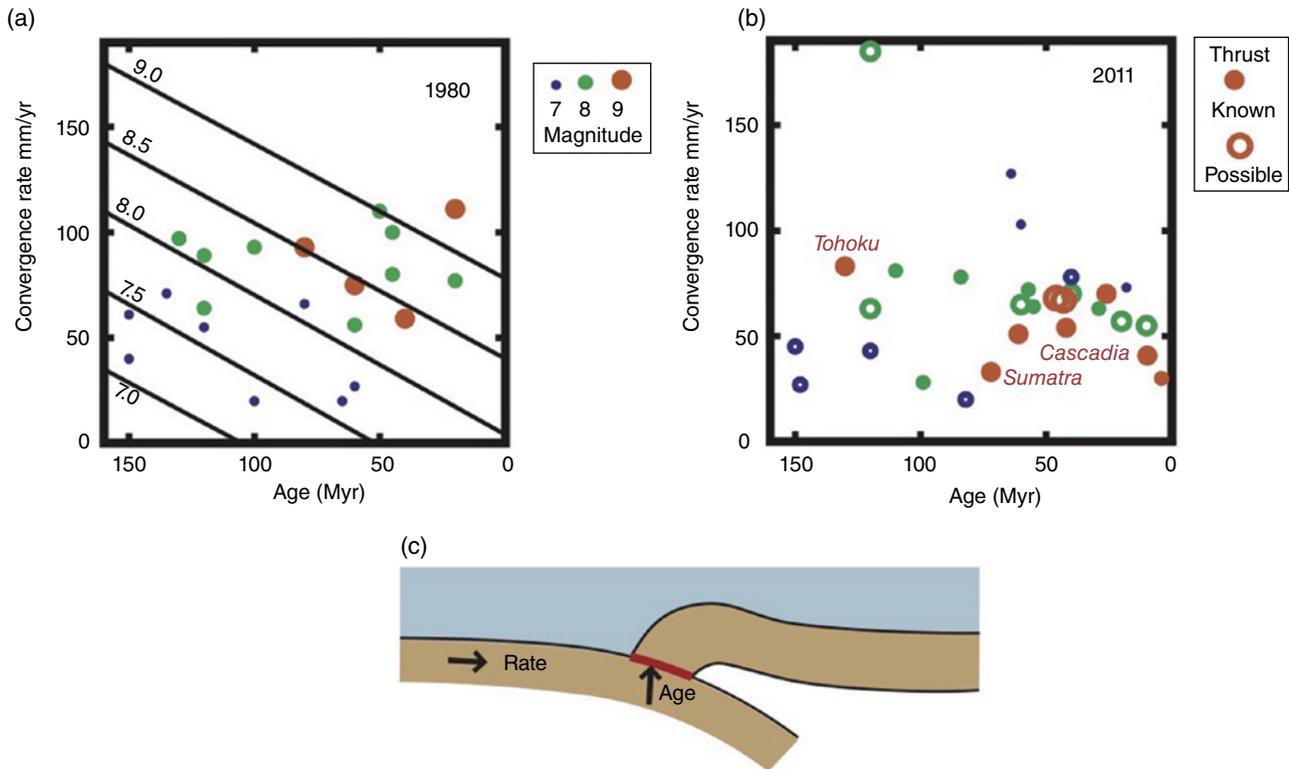


Figure 6.10 (a) Data available in 1980, showing the largest earthquake known at various subduction zones. Magnitude 9 earthquakes occurred only where young lithosphere subducts rapidly. Diagonal lines show predicted maximum earthquake magnitude [Ruff and Kanamori, 1980]. (b) Data available today, updated from Stein and Okal [2007] by including 2011 Tohoku earthquake [Stein and Okal, 2011]. (c) Physical interpretation of older result (a) in terms of strong mechanical coupling and thus large earthquakes at the trench interface.

earthquake. Using the model, the maximum expected earthquake size could be predicted.

The model was widely accepted until the 2004 M_w 9.3 Sumatra earthquake that generated the giant Indian Ocean tsunami. According to the model, this trench should have generated at most a low-magnitude 8 earthquake. However, reanalysis found a quite different picture [Stein and Okal, 2007]. With the newer data, the proposed correlation vanished, as the 2011 Tohoku earthquake subsequently confirmed. Thus, instead of only some subduction zones being able to generate magnitude 9s, it now looks like many or all can [McCaffrey, 2007]. The apparent pattern resulted from the fact that magnitude 9s are so rare, on average there is less than one per decade [Stein and Wysession, 2003]. These are about 10 times rarer than magnitude 8s. Thus, the short seismological record (the seismometer was invented in the 1880s) misled seismologists into assuming that the largest earthquakes known on a particular subduction zone were the largest that would happen. Unfortunately, the accumulating paleotsunami data (e.g., Minoura et al., [2001]) that spanned a

longer time and showed larger earthquakes had not yet been assimilated into these ideas.

M_{max} estimates can be made by assuming that a fault of a certain length will rupture. However, as the Tohoku example (Fig. 6.3) showed, the Earth may not rupture as expected.

Thus, all one can say with certainty is that M_{max} in an area is at least as large as the largest earthquake in the available record. However, numerical simulations show that the maximum magnitude appearing in a catalog tends to be that of earthquakes with mean recurrence time equal to the catalog length. Because catalogs are often short relative to the average recurrence time of large earthquakes, larger earthquakes than anticipated can occur (Fig. 6.11). Estimating M_{max} is particularly challenging within plates, where large earthquakes are infrequent compared to the length of the available earthquake history, vary in space and time, and sometimes occur on previously unrecognized faults.

Longer records can help address this issue. Paleoseismic and paleotsunami studies can improve estimates of the

lower bound on M_{max} by finding past earthquakes larger than those previously known. However, they cannot resolve the upper bound issue, because we have no way of knowing whether a bigger earthquake will occur.

6.8. CHAOS: A WEATHER ANALOGY

The Tohoku example illustrates that, although over many years the total slip in earthquakes on a plate boundary should be determined by the plate motion, the detailed history of where and when large earthquakes

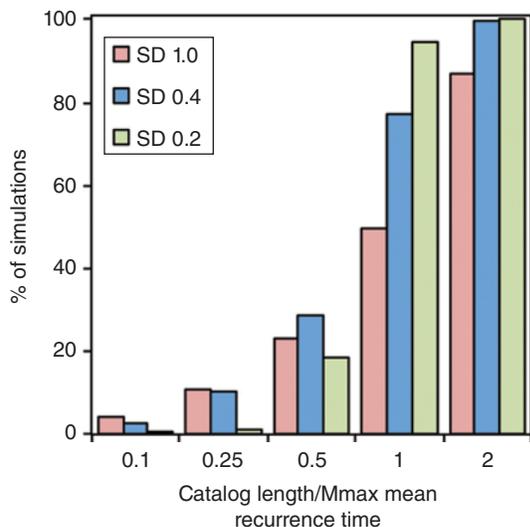


Figure 6.11 Numerical simulations assessing how well earthquake catalogs of different lengths recover the actual maximum magnitude M_{max} of earthquakes in an area. Catalog lengths are given as a fraction of the mean recurrence time for earthquakes with magnitude M_{max} . Colors show results for Gaussian distributions of recurrence times with standard deviation equal to the indicated fraction of the mean. The largest earthquake observed likely reflects the length of the history used, even if larger earthquakes occur, so a catalog shorter than an earthquake’s mean recurrence time is likely to not contain an event of that size [Stein and Friedrich, 2014].

occur and how big they are can be very complicated. Hence it is likely that these key parameters for hazard modeling are difficult to know or unknowable, placing fundamental limits on how well we can assess hazards.

A useful analogy is the challenge of forecasting weather. Lorenz [1995] explains that the overall frequency of storms depends on the energetics of the atmosphere, but minuscule disturbances can modify when they occur. This effect was suggested by Lorenz’s surprising observation in 1963 that small changes in simple computer models of the weather could give very different results. Assuming that the real atmosphere acts in this chaotic way, small perturbations could grow to have large effects. Lorenz described this effect using the famous analogy that the flap of a butterfly’s wings in Brazil might set off a tornado in Texas. This concept reached the public in the movie *Jurassic Park*, where small problems grew into big ones that made the dinosaur park collapse. A simple illustration of this idea comes from considering a system whose evolution in time is described by the difference equation $x(t+1) = ax(t)^2 - 1$ (Fig. 6.12). For $a=2$, two runs starting off at time $t=0$ with slightly different values, $x(0) = 0.750$ and $x(0) = 0.749$, yield curves that differ significantly within a short time.

The fact that small differences grow is part of the reason why weather forecasts get less accurate as they project farther into the future: tomorrow’s forecast is much better than one for the next five days. By about two weeks, the uncertainties are so large that forecasts are not useful.

An interesting thought experiment is to ask what the weather would be like if it were not chaotic [Lorenz, 1995]. In this case, weather would be controlled only by the seasons, so year after year storms would follow the same tracks, making planning to avoid storm damage easy. In reality, storm tracks differ greatly from year to year (Fig. 6.13). Thus, in Lorenz’s words, “the greater difficulty in planning things in the real world, and the occasional disastrous effects of hurricanes and other storms, must therefore be attributed to chaos.”

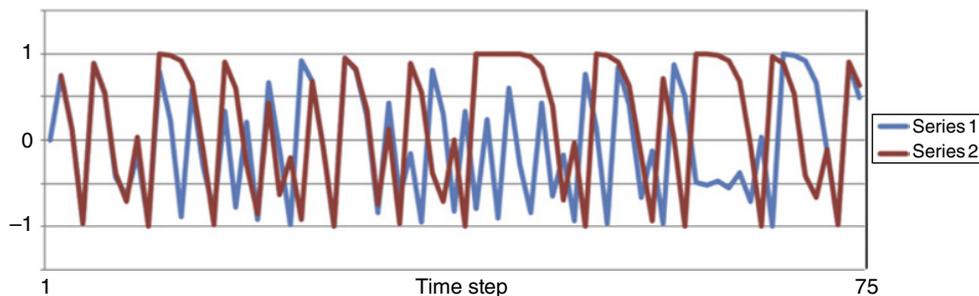


Figure 6.12 Comparison of two time series generated by the same equation with slightly different initial conditions, which quickly lead to quite different values. (For color detail, please see color plate section).

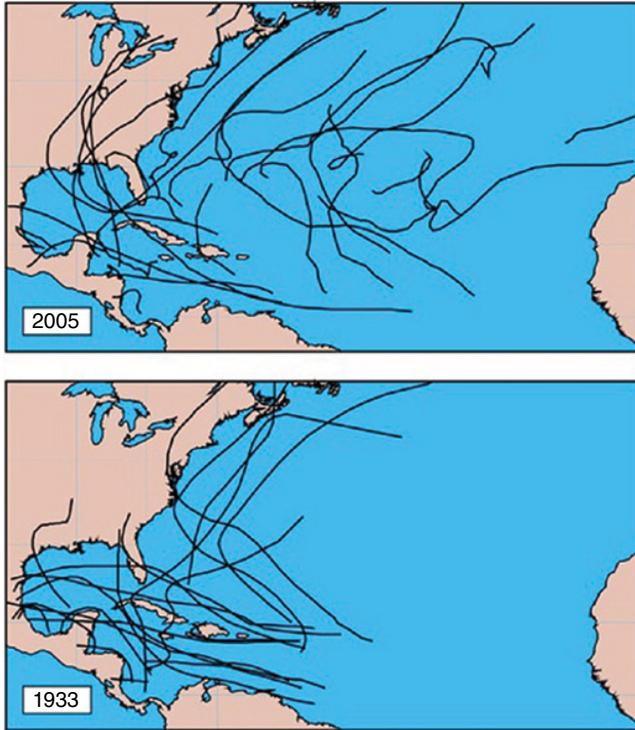


Figure 6.13 Tracks of North Atlantic hurricanes, tropical storms, and depressions for two very active hurricane seasons [Ebeling and Stein, 2011].

Thus we can say something about space-time variability, but only in general terms. Although we know that on average some areas (e.g., Florida) will be hit by hurricanes more than others (e.g., New York), we cannot predict what will happen next year, beyond the general fact that hurricanes will occur during the hurricane season, when the ocean water is warm enough. Similarly, although we can say that over time some areas have higher earthquake hazard than others, but the more detailed a forecast we try to make, the more uncertain it will be.

6.9. WHEN IS THE NEXT EARTHQUAKE?

By analogy to weather, without chaos, steady motion between plates should produce a pattern of earthquakes that repeats in space and time. In contrast, the chaos view predicts that the locations of big earthquakes on a plate boundary and interval between them should be highly variable. As Fig. 6.14 shows for the Nankai Trough, this is what the geological record shows. This variability is why using seismic gaps to predict earthquake locations often fails.

Paleoseismic records show time-variability of large earthquakes. Fig. 6.15 (left) shows a recurrence history at Pallett Creek, on the segment of the San Andreas fault that broke in the 1857 Fort Tejon earthquake [Sieh et al.,

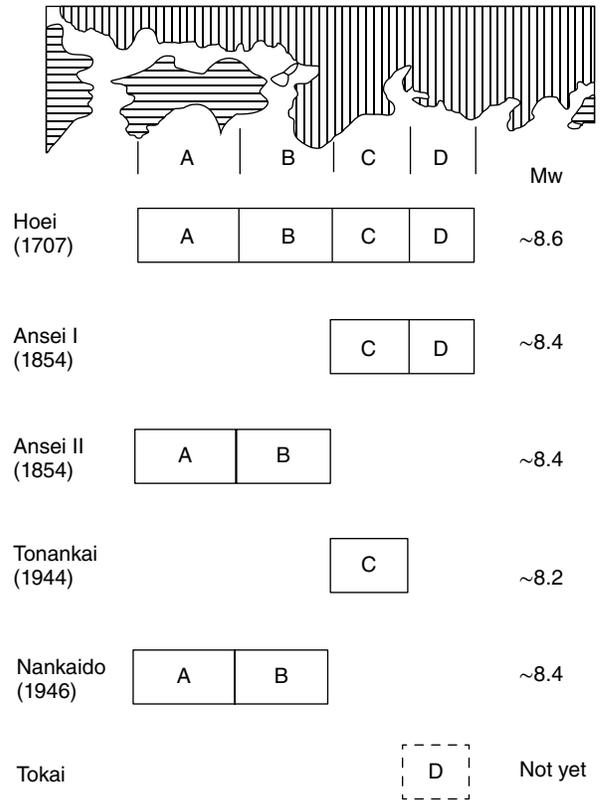


Figure 6.14 Earthquake history for the Nankai trough area [Ando, 1975] illustrating how different segments rupturing cause earthquakes of different magnitudes. [Stein and Okal, 2007].

1989]. The nine recurrence intervals have a mean of 132 years and a standard deviation of 105 years. This large variability results from the presence of several clusters of large earthquakes, which, together with the observational uncertainties, make it difficult to characterize the sequence and estimate earthquake probabilities. Hence, Sieh et al.'s [1989] estimates of the probability of a similar earthquake before 2019 ranged from 7% to 51%. Moreover, using different subsets of the series will yield different results [Stein and Newman, 2004]. As a result, the actual variability is greater than inferred from studies that use short earthquake sequences, typically 2–4 recurrences [Nishenko and Buland, 1987].

Hazard mapping requires assuming some probability density function that describes the distribution of future earthquake recurrence intervals. The traditional choice among the possible probability density functions is between ones representing two models of earthquake recurrence [Stein and Wysession, 2003]. In one, the recurrence of large earthquakes is described by a time-independent Poisson process that has no “memory.” Thus, a future earthquake is equally likely immediately after the past one and much later, so earthquakes often cluster in time. Under this

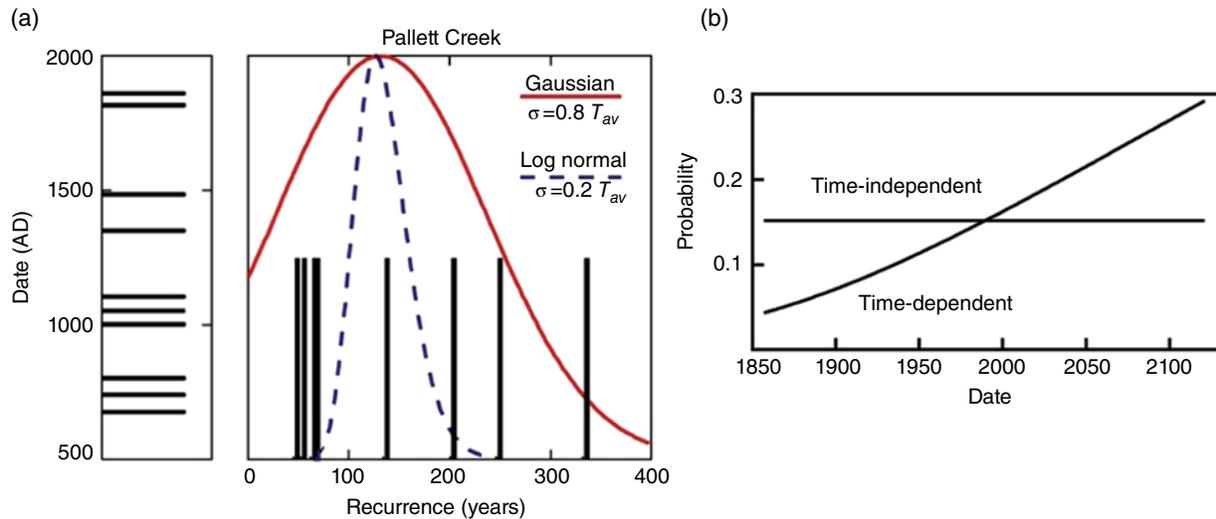


Figure 6.15 (a) Variability in recurrence intervals for the long earthquake sequence at Pallett Creek [Stein and Newman, 2004]. (b) Conditional probability that a large earthquake will occur in the following 20 yr for two different models [Stein and Stein, 2014].

assumption, the probability that an earthquake will occur in the next t years is approximately t/T , where T is the assumed mean recurrence time, and an earthquake cannot be “overdue.”

The common alternative is to use some time-dependent recurrence model in which a probability distribution describes the time between earthquakes. In this model, earthquakes are quasi-periodic, with the standard deviation of recurrence times small compared to their mean. In such models, the conditional probability of the next large earthquake, given that it has not yet happened, varies with time. The probability is small shortly after the past one, and then increases with time. Eventually, if a large earthquake has not occurred by this time, the earthquake is overdue in the sense that time-dependent models predict higher probabilities (Fig. 6.15, right).

Neither of these common model types does a good job of describing the earthquake clusters. A possible alternative is a different time-variable model in which earthquake probability increases with time between earthquakes and decreases after an earthquake, but does not reset, allowing for clusters with long gaps. The model is derived via the classic formulation of drawing balls from an urn containing balls labeled “E” for event and “N” for no event. E-balls are added after a draw when an event does not occur, and removed when an event occurs. This makes the probability of an event increase with time until one happens, after which it decreases and then grows again. Events are not independent, because one happening changes the probability of another. The resulting time histories (Fig. 6.16) are suggestive of the clusters, also termed supercycles, that also appear in long records of subduction zone earthquakes [Goldfinger *et al.*, 2013].

The sequences of earthquakes result from both the model parameters and chance, so two runs with the same parameters look different. The model parameters control the average time between events and the variation of the actual times around this average, so models can be strongly or weakly time dependent.

Fig. 6.16 illustrates that although generating synthetic earthquakes from a probability model is easy, using an earthquake history to infer the probability model is very difficult. From the earthquake sequences, it would be difficult to tell what process generated them. Moreover, we can easily convince ourselves that we see all sorts of patterns, many of which would look different in another run, as illustrated by the differences between the lower two series that are due only to randomness. Some parts of the sequences look pretty regular, so if we had only these samples, we might decide the system was periodic, and then would be surprised when the next events did not fit that apparent pattern. We might then decide that something in the system gave clusters of periodic events separated by longer intervals, and would be disappointed when that pattern also broke down.

Similarly, even long paleoseismic and historic earthquake records often cannot resolve the probability density function very well [Biasi *et al.*, 2002; Parsons, 2008; Parsons and Giest, 2009; Savage, 1991, 1992, 1994]. Freedman and Stark [2003] concluded that estimates of earthquake probabilities and their uncertainties are “shaky.” In their view, “the interpretation that probability is a property of a model and has meaning for the world only by analogy seems the most appropriate. ... The problem in earthquake forecasts is that the models, unlike the models for coin-tossing, have not been tested against

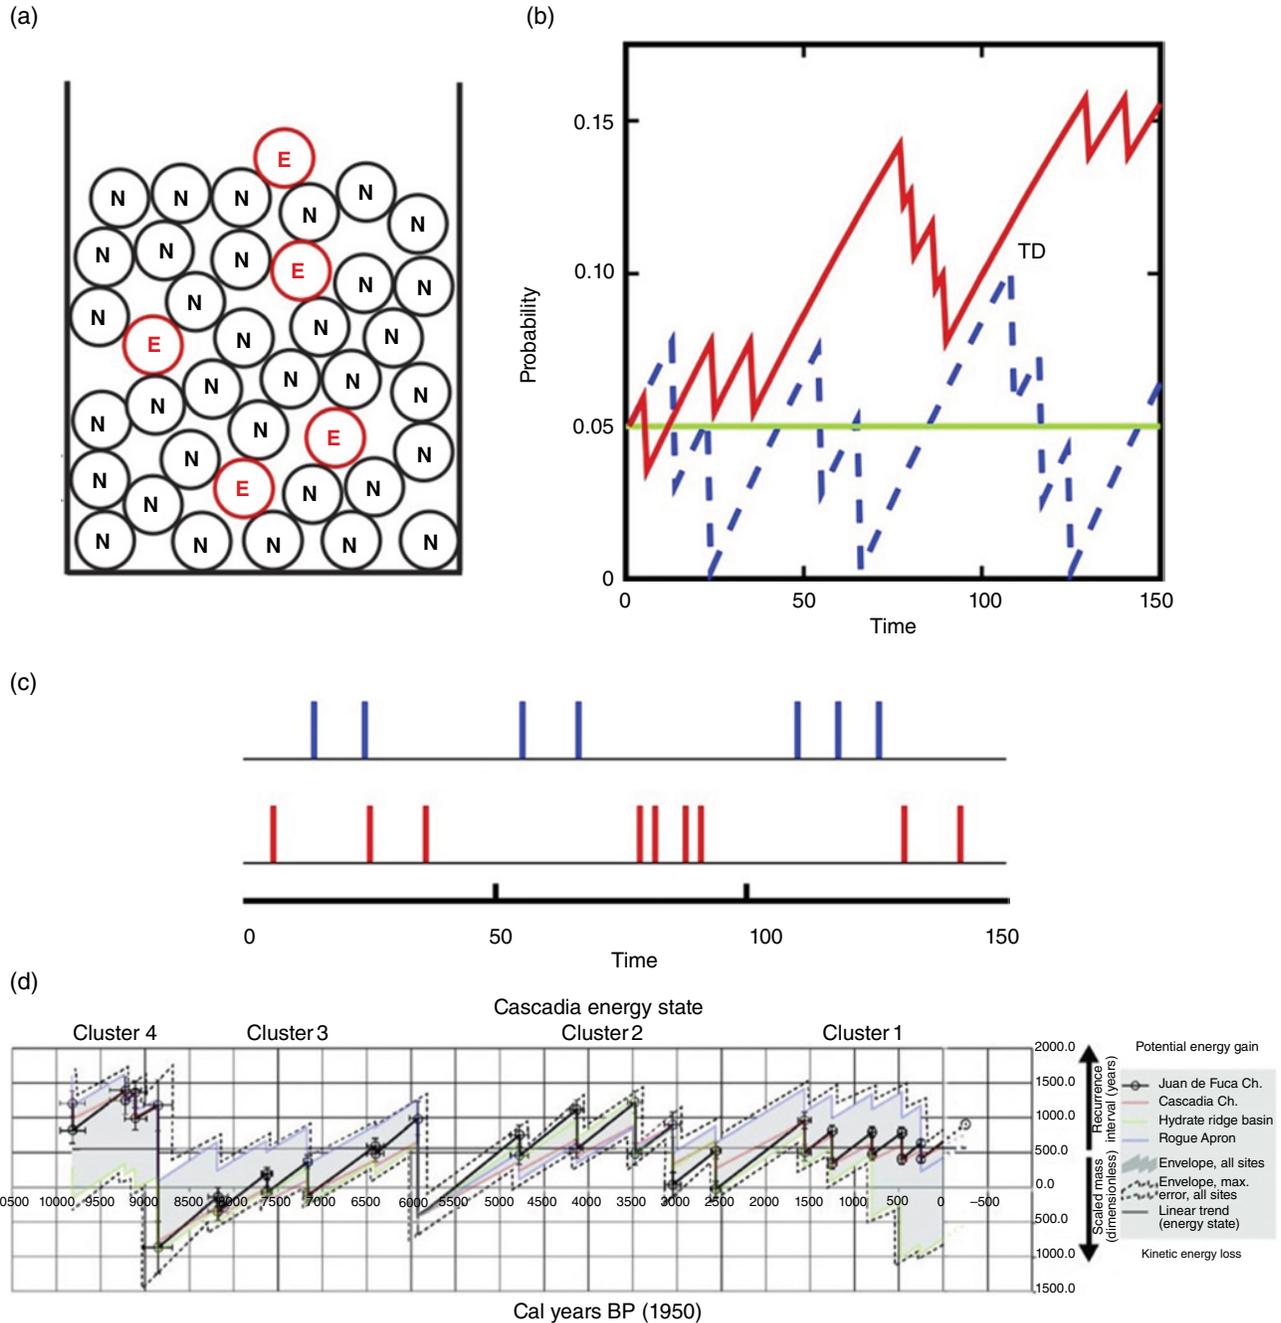


Figure 6.16 (a) Model for the probability of an event as drawing from an urn with balls labeled “E” for event and “N” for no event. (b) Comparison of the probability of an event as a function of time for time-independent (green line) and two runs (red and blue lines) of a time-dependent model. (c) Sequence of events as a function of time for the two runs in (b) [Stein and Stein, 2013a]. (d) Supercycle model for large Cascadia earthquakes [Goldfinger et al., 2013]. (For color detail, please see color plate section).

relevant data. Indeed, the models cannot be tested on a human time scale, so there is little reason to believe the probability estimates.” Savage [1991] similarly concluded that earthquake probability estimates for California are “virtually meaningless” and that it would be meaningful

only to quote broad ranges, such as low (<10%), intermediate (10–90%), or high (>90%). In other words, it seems reasonable to say that earthquakes of a given size are more likely on some faults than others, but quantifying this involves large uncertainty.

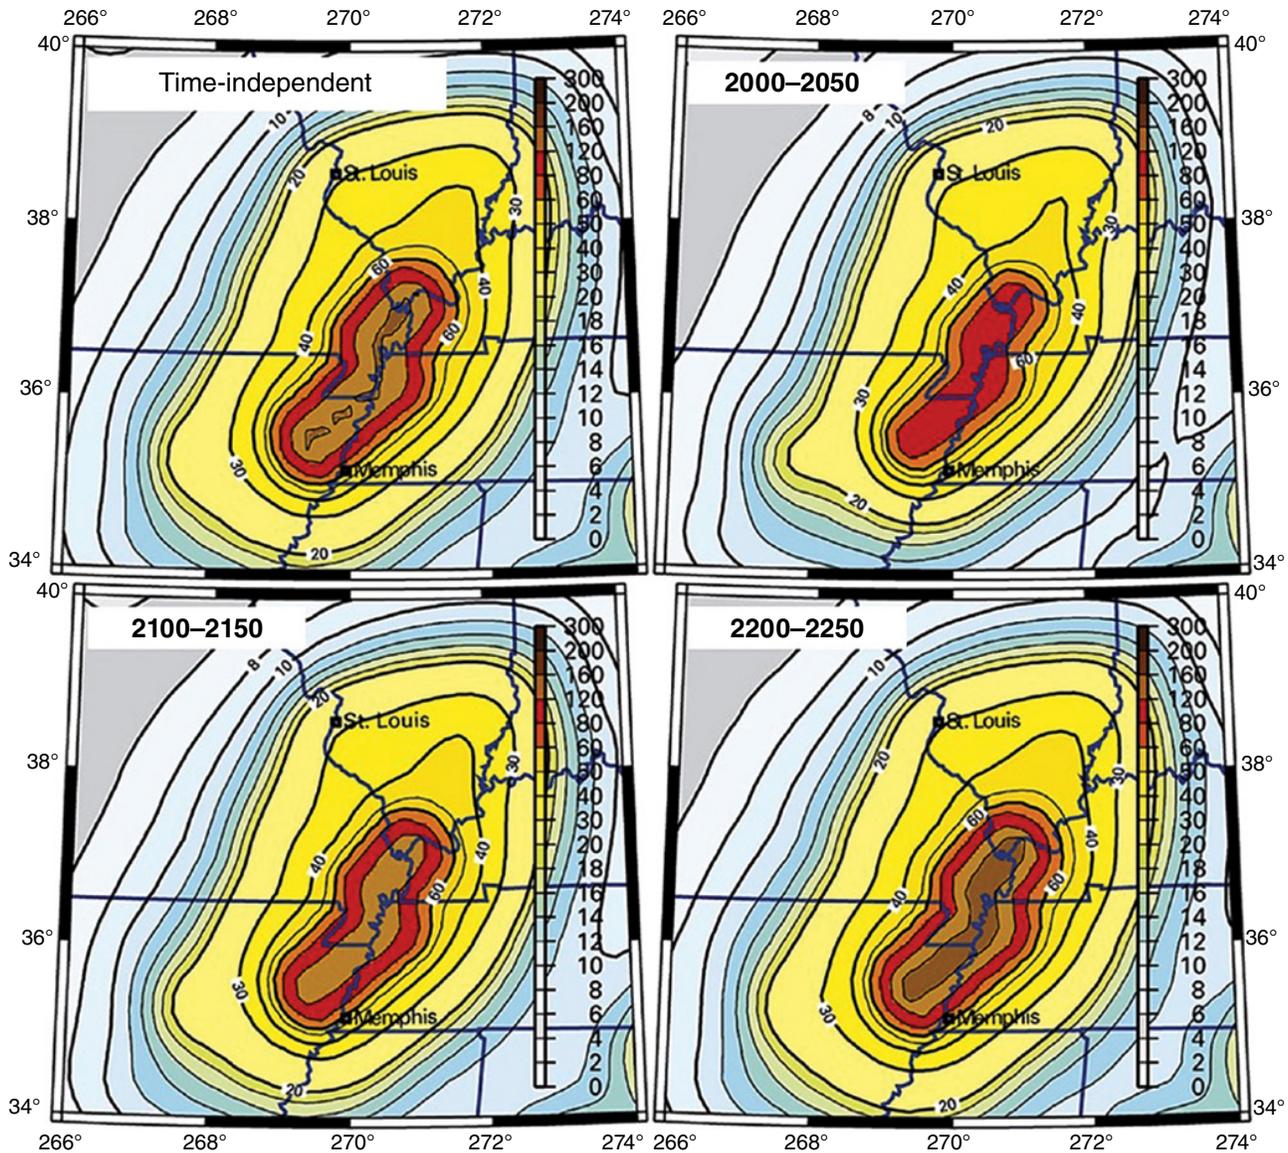


Figure 6.17 Comparison of hazard maps for the New Madrid zone. Colors show peak ground acceleration as percentages of 1 g. Compared to the hazard predicted by the time-independent model, the time-dependent model predicts noticeably lower hazard for the 50-yr periods 2000–2050 and 2100–2150, but higher hazard if a large earthquake has not occurred by 2200 [Hebden and Stein, 2009].

This situation presents a major source of uncertainty in hazard maps, as illustrated in Fig. 6.17 by alternative maps for the New Madrid zone. The biggest effect is close to the three faults used to model the jagged geometry of the earthquakes of 1811–1812, where the largest hazard is predicted. Such earthquakes are assumed to have a moment magnitude 7.3 [Hough *et al.*, 2000] and to recur every 500 years. Compared to the hazard predicted by the time-independent model, the time-dependent model predicts noticeably lower hazard for 2000–2050 and 2100–2150, because only about 200 yr have elapsed.

However, if a large earthquake has not occurred by 2200, the hazard predicted in the next 50 yr would be higher than predicted by the time-independent model.

Given the limitations of our knowledge, such uncertainties are hard to reduce and may be irreducible. Hence, at present and in the foreseeable future, choosing how to model earthquake recurrence on faults in a hazard map largely reflects the mappers' preconceptions. The Japan map (Fig. 6.2) reflected the mappers' view that a large earthquake would happen much sooner on the Nankai Trough than off Tohoku, which proved not to be the case.

6.10. WHAT TO DO?

Post-Tohoku discussions among seismologists suggest several approaches to improving the situation. The obvious approach, acquiring more and better data, is being pursued. Modeling of fault processes will also help. Because many of hazard mapping's limitations reflect the present limited knowledge about earthquakes and tectonics, we anticipate advances both from ongoing studies and new methods such as seafloor geodesy [e.g., Newman, 2011; Dixon *et al.*, 2014].

However, we still do not know how to effectively use these data for anything beyond relatively general forecasts. For example, even had the GPS data showing strain accumulation off Tohoku been appreciated, there was no good way to forecast how large an earthquake might occur or how soon. This limitation is illustrated by the fact that communities inland from the Nankai Trough are now being warned of much larger tsunamis than previously anticipated, assuming that a future earthquake could be as large as March 2011's Tohoku earthquake [Cyranski, 2012]. These communities face the challenge of deciding what to do for a possible 20-meter tsunami whose probability cannot be usefully estimated beyond assuming that it would be rare, perhaps once in a millennium.

Many challenges for hazard mapping are unlikely to be resolved soon, and some may be inherently unresolvable. Hence, in addition to research on the seismological issues, changes to current hazard mapping practices would make maps more useful.

6.10.1. Assess and Present Uncertainty Estimates

Hazard maps clearly have large uncertainties. When a map fails, it is often clear in hindsight that key parameters were poorly estimated. Sensitivity analyses like that in Figs. 6.5 and 6.17 illustrate this point: the maps are uncertain in the sense that their predictions vary significantly depending on the choice of many poorly known parameters.

Estimates of the expected uncertainty in the predicted hazard should be presented and explained. Fig. 6.18a compares the predictions of the models in Figs. 6.5 and 6.17 for the hazard at Saint Louis and Memphis, which vary by a factor of more than three. This representation shows the effects of the three factors. At Memphis, close to the main faults, the primary effect is that of magnitude, with the two M 8 models predicting the highest hazard. At Saint Louis, the ground motion model has the largest effect, so the Frankel models predict the highest hazard. Most models show hazard well below that predicted for California. The predictions for a maximum magnitude of 7 are similar to ones in which the large

earthquake sequence has ended and the hazard reflects continuing aftershocks [Stein, 2010].

Similar approaches are used to present uncertainties for analogous forecasts with significant economic and policy implications (Fig. 6.18b–d). Meteorologists [Hirschberg *et al.*, 2011] have adopted a goal of “routinely providing the nation with comprehensive, skillful, reliable, sharp, and useful information about the uncertainty of hydrometeorological forecasts.” Although, as the “snowmageddon” (section 6.1) shows, this is not yet fully the case, they increasingly present the public with uncertainty information by comparing the predictions of various models or by showing uncertainty estimates. Even though forecasts sometimes miss their targets (Fig. 6.19), uncertainty estimates are still useful.

Although seismologists have a tougher challenge and a longer way to go, we should try to do the same. Assessing and communicating their uncertainties would make hazard maps more useful. At present, most users have no way to tell which predictions of these maps are likely to be reasonably well constrained and which are not. Having this information would help users make better decisions about mitigation strategies [Stein and Stein, 2014].

6.10.2. Characterize and Assess Map Performance

Maps should specify what they seek to predict and how their performance should be measured so users can know what the mappers' goals are and be able at a later time to assess how well the map met them. One key question is, how well did a map perform compared to one that assumed a much smoother variation in the predicted hazard [Geller, 2011]? If smoother maps work better, then the more detailed maps are likely overparameterized, containing so much detail so they fit the past better but predict the future worse [Stein *et al.*, 2012]. Another is how the performance of probabilistic hazard maps compares to that of deterministic hazard maps, which use specific scenarios rather than expected values of shaking [Peresan and Panza, 2012; Wang, 2011].

Conceptually, the issue is how to compare a map of predicted shaking to the maximum shaking observed at sites within it over a suitably long period of time after the map was made. There is increasing interest in this issue, and a variety of approaches have recently been used [Albarelllo and D'Amico, 2008; Beauval *et al.*, 2008, 2010; Kossobokov and Nekrasova, 2012; Mak *et al.*, 2014; Miyazawa and Mori, 2009; Mucciarelli *et al.*, 2008; Nekrasova *et al.*, 2014; Stirling and Gerstenberger, 2010; Stirling and Peterson, 2006; Wyss *et al.*, 2012] and are being developed under auspices of the Global Earthquake Model project (<http://www.globalquakemodel.org>).

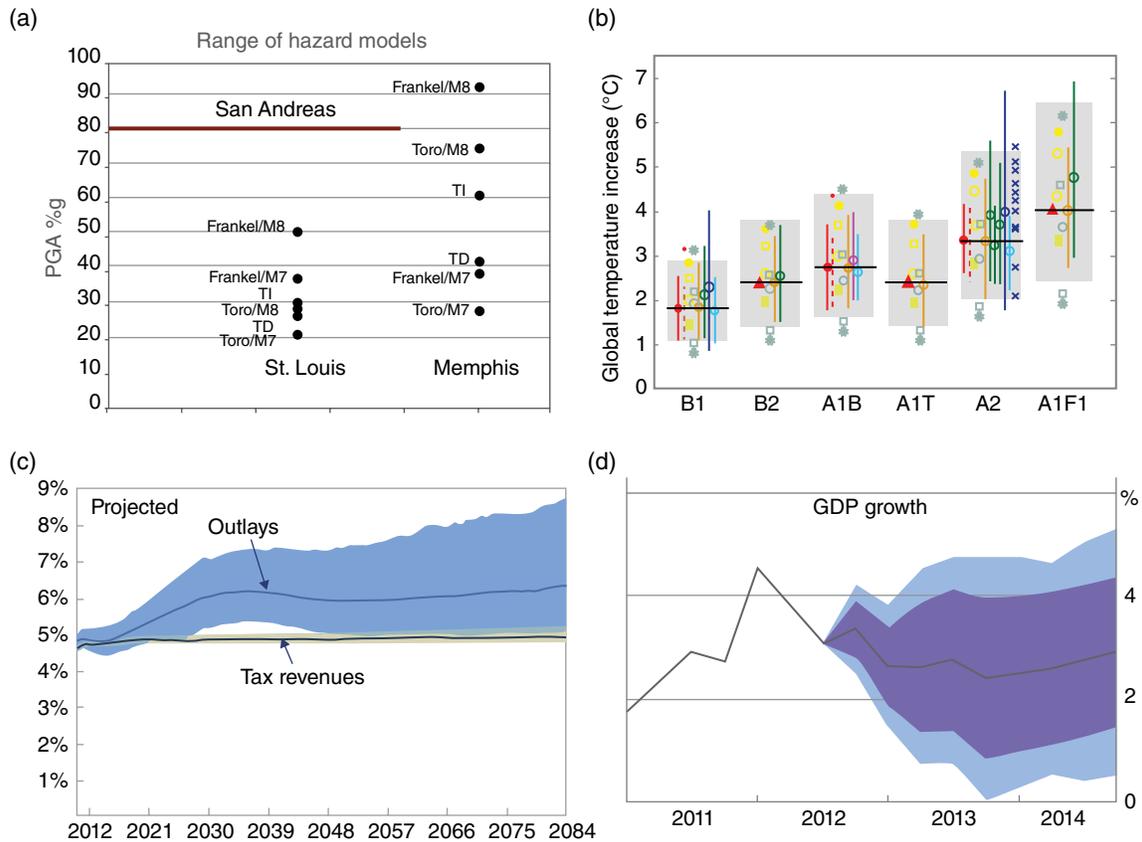


Figure 6.18 Presenting forecast uncertainties. (a) Comparison of earthquake hazard, described as peak ground acceleration (PGA) as a percentage of the acceleration of gravity expected with 2% risk in 50 yr, predicted by various assumptions for two sites in the central United States [Stein *et al.*, 2012]. (b) Comparison of the rise in global temperature by the year 2099 predicted by various climate models. For various carbon emissions scenarios, for example, B1, the vertical band shows the predicted warming [IPCC, Intergovernmental Panel on Climate Change, IPCC, 2007]. (c) Forecast of US Social Security expenditure as percentage of GDP [Congressional Budget Office, 2010]. (d) Forecast of Australian GDP growth. Uncertainty bounds are 70% and 90% [Reserve Bank of Australia, 2013]. Figure from Stein *et al.* [2015a].

At present, there is no agreed way of assessing how well a map performed and thus whether one map performed better than another. The fractional site exceedance metric implicit in probabilistic maps, that during the chosen time interval the predicted ground motion will be exceeded only at a specific fraction of the sites, is useful. However, it permits maps to be nominally successful although they significantly underpredict or overpredict shaking, or to be nominally unsuccessful but do well in terms of predicting shaking (Fig. 6.20). Although no single metric alone fully characterizes map behavior, adapting and using several metrics can provide useful insight for comparing and improving hazard maps.

For example, as measured by the exceedance metric, a 510-yr-long record of earthquake shaking in Japan is better described by a map in which the area has uniform hazard than by the actual maps [Stein *et al.*, 2015c]. However,

using the squared misfit between maximum observed shaking and that predicted as a metric, the actual maps do better than uniform or randomized maps. The observation that the actual maps do worse than uniform or randomized maps by one metric and better by another reflects the fact that a system’s performance has multiple aspects. For example, how good a baseball player Babe Ruth was depends on the metric used. In many seasons Ruth led the league in both home runs and in the number of times he struck out. By one metric he did very well, and by another, very poorly.

6.10.3. Develop Objective Methods for Updating Maps

Whether and how much to revise a map after an earthquake that is “unexpected,” bigger, or causes more shaking than expected is complicated. The question is

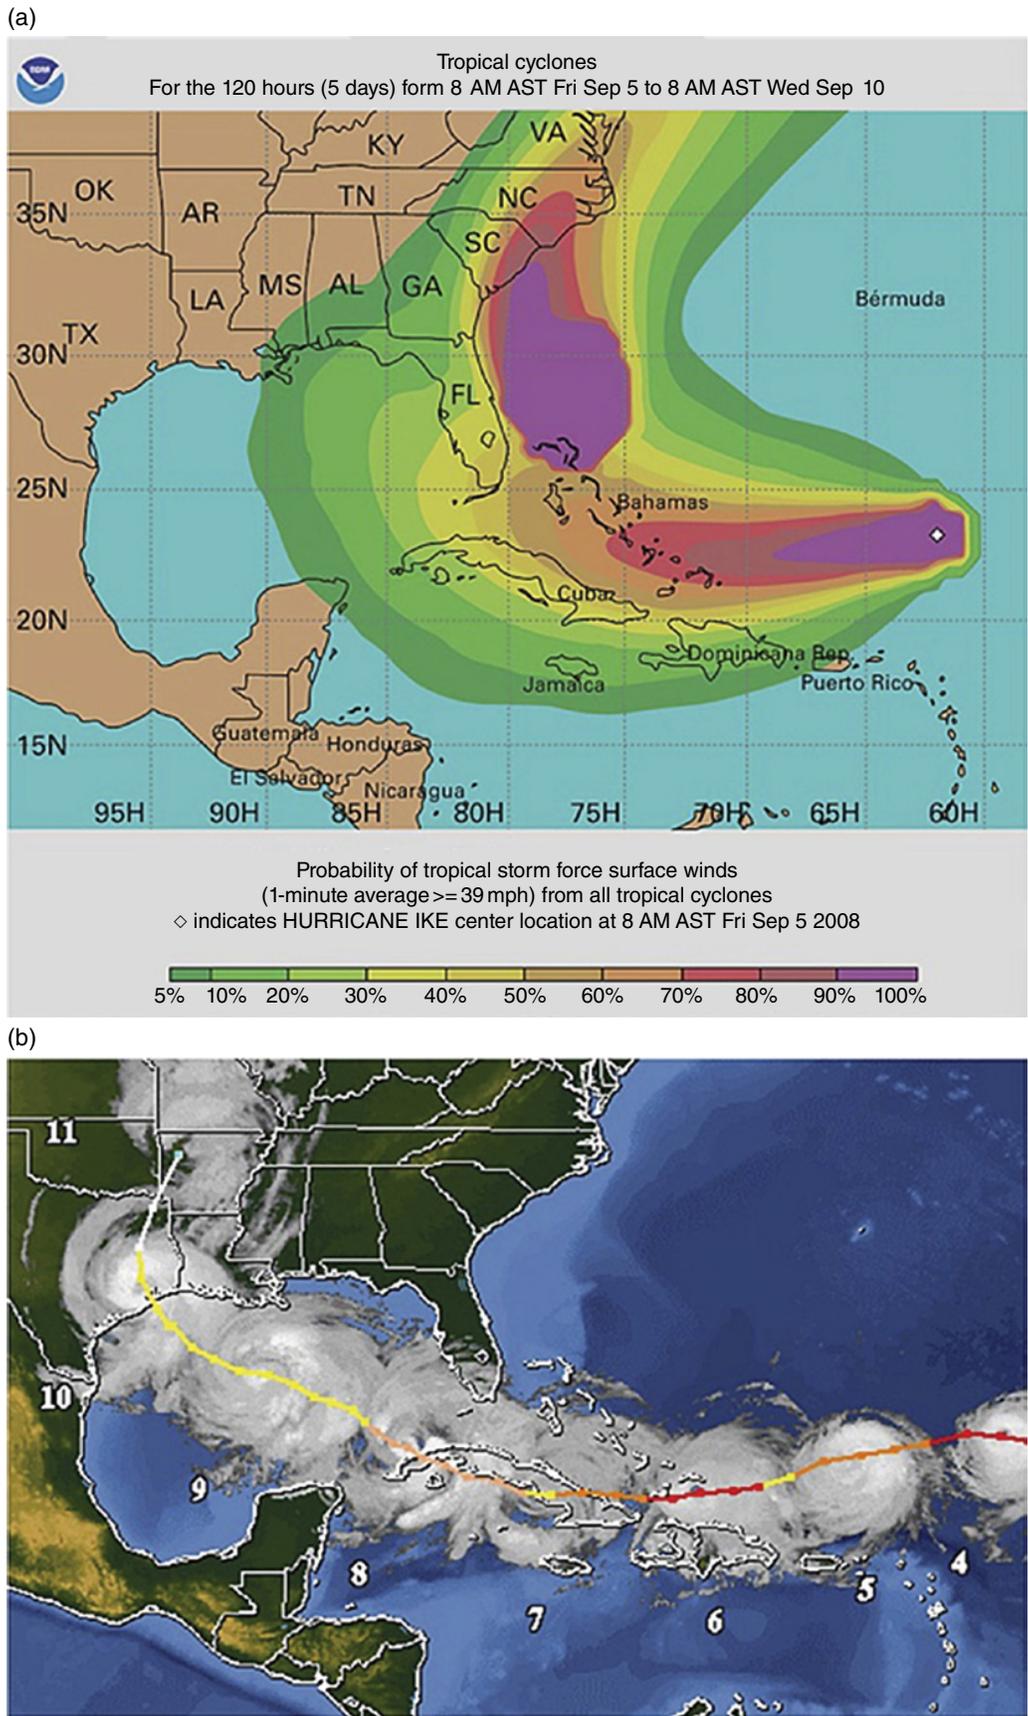


Figure 6.19 Comparison of the predicted (top) and actual (bottom) tracks of Hurricane Ike in December 2008. The storm was predicted to continue westward and then turn north along the Florida coast, but instead followed a track outside the 95% uncertainty cone that headed into the Gulf of Mexico, striking the Texas coast [Stein and Stein, 2014].

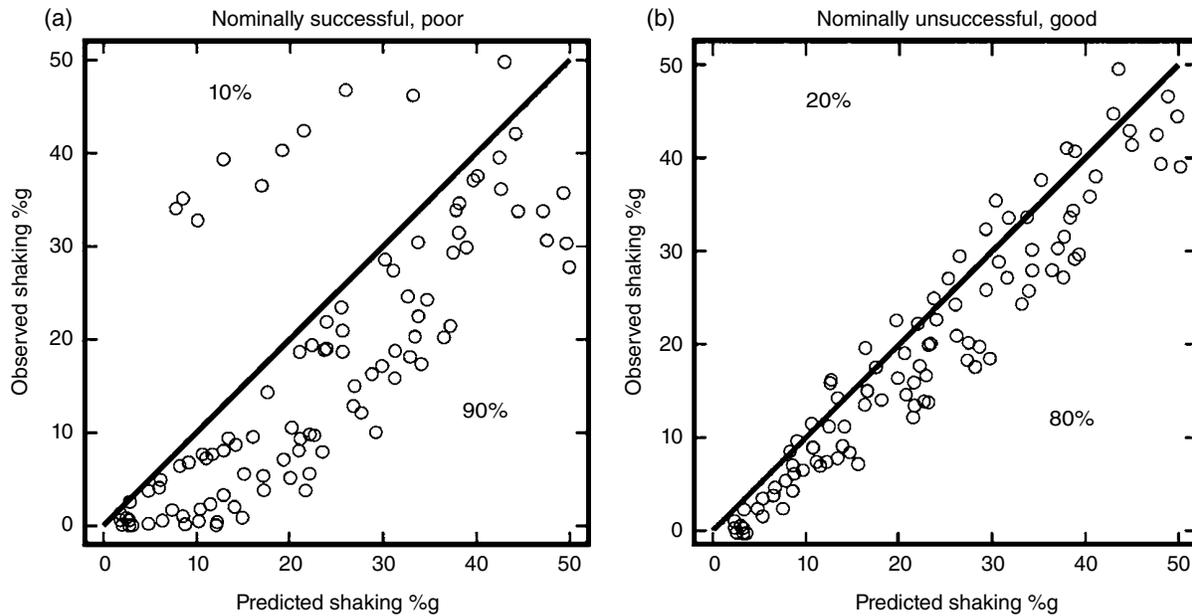


Figure 6.20 Comparison of the shaking predicted in various subregions of hazard maps to the maximum observed shaking, for two hypothetical hazard maps. That in (a) is nominally successful as measured by the fractional exceedance metric, in that the observed shaking exceeds the predicted only for the desired fraction (10%) of sites. However, it significantly underpredicts the shaking at many sites and overpredicts that at others. That in (b) is nominally unsuccessful as measured by the fractional site exceedance metric, in that the observed shaking exceeds the predicted at many more sites than desired. However, it better predicts the shaking at most sites [Stein *et al.*, 2015b].

akin to deciding whether and how much to revise your estimate of the probability that a coin will land heads after it landed heads four times in a row [Stein *et al.*, 2015a]. Either choice runs a risk. If the coin is severely biased, staying with the assumption that it is fair will continue to yield poor predictions. However, if the coin is fair and the four heads were just a low-probability event, changing to the assumption that the coin is biased does a better job of describing what happened in the past, but will make your prediction worse.

Your choice would depend on how confident you were in your assumption, prior to the tosses, that the coin was fair. If you were confident that the coin was fair, you would not change your model, and continue to assume that a head or tail is equally likely. However, if you got the coin at a magic show, your confidence that it is fair would be lower and you would be more apt to change your model to one predicting a head more likely than a tail.

Similarly, it makes sense to update maps to reflect both what occurred in earthquakes after a map was made and other information that was either unknown or not appreciated [e.g., Manaker *et al.*, 2008; Minoura *et al.*, 2001; Sagiya, 2011] when the map was made. However, a new map that better describes the past may or may not better predict the future. For example, increasing the predicted

hazard after an earthquake on a fault will make better predictions if the average recurrence time is short compared to the map's time window, but can overpredict future shaking if the average recurrence time is long and underestimated.

The most practical approach appears to develop methods based on Bayes' rule [Marzocchi and Jordan, 2014; Stein *et al.*, 2015a]. In this formulation one starts by assuming an initial or prior probability model based on information available prior to the additional observations, calculating how likely the observations were, given that model, and using the product as the revised or posterior probability model to account for the additional observations.

Fig. 6.21 shows a simple example in which an earthquake previously assumed to have a mean recurrence described by a Poisson process with parameter $\lambda = 1/50 \text{ years} = 0.02$ recurs after only one year. The updated forecast of λ , described by the posterior mean, increasingly differs from the initial forecast (prior mean) when the uncertainty in the prior distribution is larger. The less confidence we have in the prior model, the more a new datum can change it. This approach is intermediate between staying with model parameters regardless of new data, and remaking models to exactly match recent

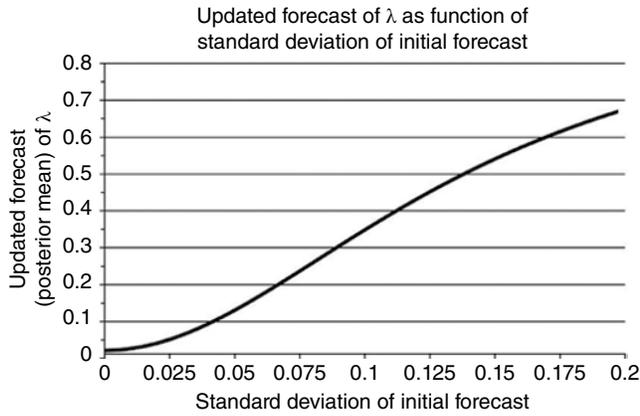


Figure 6.21 Sensitivity of updated forecast of λ , initially assumed to equal 0.02, to assumed prior uncertainty. The lower our confidence in the initial forecast, the more the new datum changes it [Stein et al., 2015a].

events. The later can be viewed as “Texas sharpshooting” in which one first shoots at the barn and then draws circles around the bullet holes.

6.10.4. Incorporate Uncertainty in Hazard Mitigation Policy

Our limited ability to assess earthquake hazards reflects the paradox that plate tectonics tells us a lot about earthquakes, but there is also much that we don’t know. The resulting uncertainties in hazard assessment affect our ability to help communities make sensible hazard-mitigation policies. The challenges illustrates Cox’s [2012] description, “Some of the most troubling risk management challenges of our time are characterized by deep uncertainties. Well-validated, trustworthy risk models giving the probabilities of future consequences for alternative present decisions are not available; the relevance of past data for predicting future outcomes is in doubt; experts disagree about the probable consequences of alternative policies.”

Society has a range of mitigation options for natural hazards, but operates under major constraints. First, we have only inadequate estimates of the hazard. Second, we have limited resources to allocate between hazard mitigation and other needs. Third, we have a wide range of societal, political, and economic considerations. Given these, we have to decide how much mitigation is appropriate, how much mitigation is enough.

Fig. 6.22 illustrates the issue. The total cost of earthquakes to society is the sum of the expected loss in future events and the cost of mitigation. This total depends on the amount of mitigation, described by the variable n . The optimum level of mitigation n^* minimizes the total

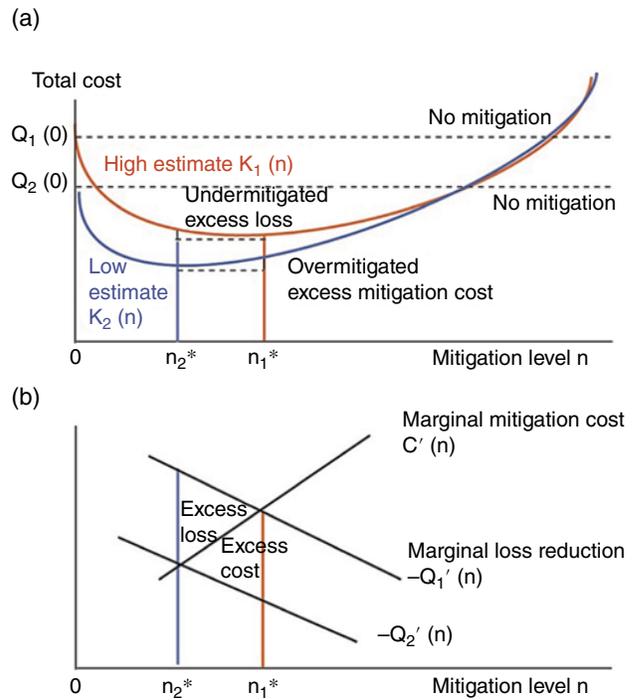


Figure 6.22 (a) Comparison of total cost curves for two estimated hazard levels. For each, the optimal mitigation level, n^* , minimizes the total cost, which equals the sum of expected loss and mitigation cost. (b) In terms of derivatives, n^* occurs when the reduced loss $-Q'(n)$ equals the incremental mitigation cost $C'(n)$. If the hazard is assumed to be described by one curve but is actually described by the other, the assumed optimal mitigation level causes nonoptimal mitigation, and thus excess expected loss or excess mitigation cost [Stein and Stein, 2013b].

cost $K(n)$, the sum of the expected loss $Q(n)$ and mitigation cost $C(n)$

$$K(n) = Q(n) + C(n)$$

The U-shaped $K(n)$ curves illustrate the trade off between mitigation and loss. For no mitigation, $n=0$, the total cost $K(0)$ equals the expected loss $Q(0)$. Initial levels of mitigation reduce the expected loss by more than their cost, so $K(n)$ decreases to a minimum at the optimum. If we undertake no mitigation, we have no mitigation costs (left side of the curve) but we expect high losses, so it makes sense to invest more in mitigation. Increased mitigation should decrease losses, so the curve goes down. Eventually, however, the cost of more mitigation exceeds the reduction in losses, and the curve rises again; the additional resources required would do more good if invested otherwise. The optimum amount of mitigation is the “sweet spot” at the bottom of the curve.

The optimum can be viewed using the derivatives of the functions, which for simplicity are shown as linear

near the optimum (Fig. 6.22b). Because increasingly high levels of mitigation cost more, the derivative, or marginal cost, $C'(n)$ increases with n . Conversely, the derivative $-Q'(n)$, the reduced loss from additional mitigation, decreases. The lines intersect at the optimum, where $C'(n^*) = -Q'(n^*)$.

Because our ability to assess hazards is limited, we need to formulate policies while accepting the uncertainties involved. To see how, consider two total cost curves corresponding to high and low estimates of the hazard, high and low estimates of the loss, or, most realistically, a combination. In these limiting cases, the hazard is assumed to be described by one curve but is actually described by the other. As a result, the optimal mitigation level chosen as the minimum of the assumed curve gives rise to nonoptimal mitigation, shown by the corresponding point on the other curve. Assuming too-low hazard causes undermitigation and excess expected loss, as shown by the height of the U-curve above the dashed line for optimum mitigation. Conversely, assuming too-high hazard causes overmitigation and excess mitigation cost. However, so long as this point is below the dashed line for the correct curve, the total cost is less than from doing no mitigation. Only even higher levels of mitigation cost more than their benefit, and thus are worse than no mitigation.

Given the range of hazard estimates, we should somehow choose an estimate between them. The resulting curve will lie between the two curves, and thus probably have a minimum between n_1^* and n_2^* . Relative to the actual but unknown optimum, this mitigation is nonoptimal, but perhaps not unduly so. So long as the total cost is below the loss for no mitigation, this nonoptimal mitigation is better than none.

This is a simple example of robust risk management, accepting the uncertainty and developing policies to give acceptable results for a range of possible hazard and loss scenarios. Such graphs are schematic guides rather than functions we can compute exactly. Given the uncertainties involved, it would be unrealistic to seek an optimum strategy. However, even simple estimates can show which strategies make more sense than others. Thus, although in real cases such approaches cannot give an optimum strategy, they can identify sensible strategies.

In other words, despite their large uncertainties, earthquake hazard maps have useful information unless they are grossly inaccurate. For example, a highway department would likely use its limited funds to preferentially strengthen bridges in predicted high-hazard areas. In most cases, this approach would make sense. However, as Geller [2011] pointed out, this might not have been the case for the pre-Tohoku Japanese hazard map.

This approach amounts to accepting the uncertainty we face and working with it. We want to use what we know already about earthquakes and plate motions,

together with what we hope to learn from new data and models, to help society make sensible policies in the presence of uncertainty. Thus, even if from a scientific standpoint hazard maps may be viewed as half-empty glasses that we hope to fill further, from a societal standpoint, these maps can be viewed as already half full. Given users insight into the uncertainties involved would make them even more useful.

ACKNOWLEDGEMENTS

Stein thanks the USGS John Wesley Powell Center for Analysis and Synthesis for hosting a working group under auspices of the Global Earthquake Model project, whose stimulating discussions inspired this work. Liu acknowledges support from NSF-PIRE Grant OISE-073015 and National Natural Science Foundation of China (Grant 91214201). Spencer thanks the Institute for Policy Research for supporting his research. We thank two anonymous reviewers for helpful comments.

REFERENCES

- Albarelo, D., and V. D'Amico (2008), Testing probabilistic seismic hazard estimates by comparison with observations: an example in Italy, *Geophys. J. Int.*, 175, 1088–1094.
- Ando, M. (1975), Source mechanisms and tectonic significance of historical earthquakes along the Nankai Trough, *Japan, Tectonophysics*, 27, 119–140.
- Barro, J. (2015), Shutting down New York's subways is very expensive, *New York Times*, January 27.
- Beauval, C., P.-Y. Bard, and J. Douglas (2010), Comment on "Test of Seismic Hazard Map from 500 Years of Recorded Intensity Data in Japan" by Masatoshi Miyazawa and Jim Mori, *Bull. Seismol. Soc. Am.*, 100, 3329–3331.
- Beauval, C., P.-Y. Bard, S. Hainzl, and P. Guéguen (2008), Can strong motion observations be used to constrain probabilistic seismic hazard estimates? *Bull. Seismol. Soc. Am.*, 98, 509–520.
- Biasi, G. P., R. J. Weldon II, T. E. Fumal, and G. G. Seitz (2002), Paleoseismic event dating and the conditional probability of large earthquakes on the southern San Andreas Fault, California, *Bull. Seismol. Soc. Am.*, 92, 2761–2781.
- Bornstein, S., and G. Mulvihill (2015), Near-miss for NYC blizzard prompts backlash against forecasters, *Associated Press*, January 27.
- Calais, E., and S. Stein (2009), Time-variable deformation in the New Madrid seismic zone, *Science*, 323, 1442.
- Camelbeeck, T., K. Vanneste, P. Alexandre, K. Verbeeck, T. Petermans, P. Rosset, M. Everaerts, R. Warnant, M. Van Camp (2007), Relevance of active faulting and seismicity studies to assess long term earthquake activity in Northwest Europe, in *Continental Intraplate Earthquakes: Science, Hazard, and Policy Issues*, Special Paper 425 edited by S. Stein and S. Mazzotti, pp. 193–224, GSA, Boulder, CO.
- Congressional Budget Office (2010), *Long-Term Projections for Social Security*.

- Cornell, C. A. (1968), Engineering seismic risk analysis, *Bull. Seismol. Soc. Am.*, 58, 1583–1606.
- Cox, L. A., Jr. (2012), Confronting deep uncertainties in risk analysis, *Risk Anal.*, 32, 1607–1629.
- Craig, T. J., and E. Calais (2014), Strain accumulation in the New Madrid and Wabash Valley seismic zones from 14 years of continuous GPS observation, *J. Geophys. Res.*, 119, 9110–9129.
- Cyranoski, D. (2011), Japan faces up to failure of its earthquake preparations, *Nature*, 471, 556–557.
- Cyranoski, D. (2012), Tsunami simulations scare Japan, *Nature*, 484, 296–297.
- Dixon, T. H., Y. Jiang, R. Malservisi, R. McCaffrey, N. Voss, M. Protti, and V. Gonzalez (2014), Earthquake and tsunami forecasts: Relation of slow slip events to subsequent earthquake rupture, *PNAS*, 111(48), 17039–17044.
- Ebeling, C., and S. Stein (2011), Seismological identification and characterization of a large hurricane, *Bull. Seism. Soc. Am.*, 101, 399–403.
- Fackler, M. (2011), Tsunami warnings written in stone, *New York Times*, April 20.
- Field, E. (2010), Probabilistic Seismic Hazard Analysis: A Primer, <http://www.opensha.org/> (last accessed May 27, 2014).
- Flegenheimer, M. (2015), Leaders in New York and New Jersey defend shutdown for a blizzard that wasn't, *New York Times*, January 27.
- Frankel, A. (2013), Comment on “Why earthquake hazard maps often fail and what to do about it,” by S. Stein, R. J. Geller, and M. Liu. *Tectonophysics*. 592, 200–206.
- Frankel, A., S. Harmsen, C. Mueller, E. Calais, and J. Haase (2010), *Documentation for initial seismic hazard maps for Haiti*, Open-File Report 2010–1067, U.S. Government Printing Office, Washington, D.C.
- Freedman, D. A., and P. B. Stark (2003), What is the chance of an earthquake? in *Earthquake Science and Seismic Risk Reduction*, edited by F. Mulargia and R. J. Geller, pp. 201–213, Kluwer, Dordrecht, The Netherlands.
- Geller, R. J. (2011), Shake-up time for Japanese seismology, *Nature*, 472, 407–409.
- Goldfinger, C., Y. Ikeda, R. S. Yeats, and J. Ren (2013), Superquakes and supercycles, *Seis. Res. Lett.*, 84(1), 24–32.
- Gordon, R. G., and S. Stein (1992), Global tectonics and space geodesy, *Science*, 256(333–342).
- GSHAP (Global Seismic Hazard Assessment Program) (1999), [Available at <http://www.seismo.ethz.ch/static/GSHAP/>]
- Gulkan, P. (2013), A dispassionate view of seismic-hazard assessment, *Seism. Res. Lett.*, 84, 413–416.
- Hallegatte, S., et al. (2012), Decision making under deep uncertainty, World Bank.
- Hanks, T. C., G. C. Beroza, and S. Toda (2012), Have recent earthquakes exposed flaws in or misunderstandings of probabilistic seismic hazard analysis? *Seismol. Res. Lett.*, 83, 759–764.
- Hebden, J., and S. Stein (2009), Time-dependent seismic hazard maps for the New Madrid seismic zone and Charleston, South Carolina areas, *Seis. Res. Lett.*, 80, 10–20.
- Hirschberg, P., et al. (2011), An implementation plan for generating and communicating forecast uncertainty information, *Bull. Am. Meteorol. Soc.*, 92, 1651–1666.
- Hough, S., J. G. Armbruster, L. Seeber, and J. F. Hough (2000), On the Modified Mercalli Intensities and magnitudes of the 1811/1812 New Madrid, central United States, earthquakes, *J. Geophys. Res.*, 105, 23,839–23,864.
- Intergovernmental Panel on Climate Change (IPCC) (2007), *Climate Change*, Cambridge Univ. Press, New York.
- Kagan, Y. Y., and D. D. Jackson (1991), Seismic gap hypothesis: ten years after, *J. Geophys. Res.*, 96, 21,419–21,431.
- Kagan, Y. Y., and D. D. Jackson (2013), Tohoku earthquake: a surprise? *Bull. Seismol. Soc. Am.*, 103, 1181–1194.
- Kagan, Y. Y., D. D. Jackson, and R. J. Geller (2012), Characteristic earthquake model, 1884–2011, RIP, *Seismological Research Letters*, 83(6), 951–953.
- Kanamori, H. (1977), Seismic and aseismic slip along subduction zones and their tectonic implications, in *Island Arcs, Deep-sea Trenches and Back-arc Basins*, Maurice Ewing Ser., 1, edited by M. Talwani and W. C. Pitman, III, pp. 163–174, AGU, Washington, D.C.
- Kerr, R. A. (2011), Seismic crystal ball proving mostly cloudy around the world, *Science*, 332, 912–913.
- Kossobokov, V. G., and A. K. Nekrasova (2012), Global Seismic Hazard Assessment Program maps are erroneous, *Seismic instruments*, 48, 162–170.
- Kreemer, C., W. C. Hammond, G. Blewitt, A. A. Holland, and R. A. Bennett (2012), A geodetic strain rate model for the Pacific-North American plate boundary, western United States, *Nevada Bureau of Mines and Geology Map*, 178.
- Kurowicka, D., and R. M. Cooke (2006), *Uncertainty Analysis with High Dimensional Dependence Modeling*, Wiley.
- Lawson, A. C., and H. F. Reid (1908), *The California Earthquake of April 18, 1906: Report of the State Earthquake Investigation Commission*. (No. 87). Carnegie institution of Washington.
- Lay, T., C. J. Ammon, H. Kanamori, L. Xue, and M. J. Kim (2011), Possible large near-trench slip during the 2011 M (w) 9.0 off the Pacific coast of Tohoku Earthquake, *Earth, planets and space*, 63(7), 687–692.
- Leonard, M., D. R. Burbidge, T. Allen, D. J. Robinson, A. McPherson, A., D. Clark, and C. D. N. Collins (2014), The challenges of probabilistic seismic hazard assessment in stable continental interiors: an Australian example. *Bull. Seism. Soc.*, 104, 3008–3028.
- Leonard, M., D. Robinson, T. Allen, J. Schneider, D. Clark, T. Dhu, and D. Burbidge (2007), Toward a better model of earthquake hazard in Australia, in *Continental Intraplate Earthquakes*, Special Paper 425 edited by S. Stein and S. Mazzotti, pp. 263–283, GSA, Boulder, CO.
- Li, Q., M. Liu, and S. Stein (2009), Spatiotemporal complexity of continental intraplate seismicity: insights from geodynamic modeling and implications for seismic hazard estimation, *Bull. Seism. Soc. Amer.*, 99, 52–60.
- Liu, M., G. Luo, H. Wang, H., and S. Stein (2014), Long-aftershock sequences in North China and Central US: implications for hazard assessment in mid-continent, *Earthquake Sci.*, 27(1), 27–35.
- Liu, M., S. Stein, and H. Wang (2011), 2000 years of migrating earthquakes in North China: How earthquakes in mid-continent differ from those at plate boundaries, *Lithosphere*, 3, doi:10.1130/L129.

- Lorenz, E. (1995), *The Essence of Chaos*, University of Washington Press, Seattle.
- Loveless, J. P., and B. J. Meade (2010), Geodetic imaging of plate motions, slip rates, and partitioning of deformation in Japan, *J. Geophys. Res.*, *115*, doi:10.1029/2008JB006248.
- Mak, S., R. A. Clements, and D. Schorlemmer (2014), The statistical power of testing probabilistic seismic-hazard assessments, *Seismol. Res. Lett.*, *85*, 781–783.
- McCann, W. R., S. P. Nishenko, L. R. Sykes, and J. Krause (1979), Seismic gaps and plate tectonics: seismic potential for major boundaries. In *Earthquake Prediction and Seismicity Patterns* (pp. 1082–1147). Birkhäuser Basel.
- Manaker, D. M., E. Calais, A. M. Freed, S. T. Ali, P. Przybylski, G. Mattioli, P. Jansma, C. Prepetit, and J. B. De Chabalie (2008), Interseismic plate coupling and strain partitioning in the Northeastern Caribbean, *Geophys. J. Int.*, *174*, 889–903.
- Marzocchi, W., and T. H. Jordan (2014), Testing for ontological errors in probabilistic forecasting models of natural systems, *Proc. Natl. Acad. Sci. Unit. States Am.*, *111*, 11973–11978.
- McCaffrey, R. (2007), The next great earthquake, *Science*, *315*, 1675–1676.
- Minoura, K., F. Imamura, D. Sugawa, Y. Kono, and T. Iwashita (2001), The 869 Jogan tsunami deposit and recurrence interval of large-scale tsunami on the Pacific coast of Northeast Japan, *J. Natural Disaster Sci.*, *23*, 83–88.
- Miyazawa, M., and J. Mori (2009), Test of seismic hazard map from 500 years of recorded intensity data in Japan, *Bull. Seismol. Soc. Am.*, *99*, 3140–3149.
- Mucciarelli, M., D. Albarello, and V. D’Amico (2008), Comparison of probabilistic seismic hazard estimates in Italy, *Bull. Seismol. Soc. Am.*, *98*, 2652–2664.
- Nanayama, F., K. Satake, R. Furukawa, K. Shimokawa, B. Atwater, K. Shigeno, and S. Yamaki (2003), Unusually large earthquakes inferred from tsunami deposits along the Kuril trench, *Nature*, *424*, 660–663.
- Nekrasova, A., V. Kossobokov, A. Peresan, and A. Magrin (2014), The comparison of the NDSHA, PSHA seismic hazard maps and real seismicity for the Italian territory, *Nat. Haz.*, *70*, 629–641.
- Newman, A., S. Stein, J. Schneider, and A. Mendez (2001), Uncertainties in seismic hazard maps for the New Madrid Seismic Zone, *Seis. Res. Lett.*, *72*, 653–667.
- Newman, A. V. (2011), Hidden Depths, *Nature*, *474*, 441–443.
- Nishenko, S. P., and R. Buland (1987), A generic recurrence interval distribution for earthquake forecasting, *Bull. Seismol. Soc. Am.*, *77*, 1382–1399.
- Parsons, T. (2008), Earthquake recurrence on the south Hayward fault is most consistent with a time dependent, renewal process, *Geophys. Res. Lett.*, *35*, doi:10.1029/2008GL035887.
- Parsons, T., and E. L. Giest (2009), Is there a basis for preferring characteristic earthquakes over a Gutenberg-Richter distribution in probabilistic earthquake forecasting? *Bull. Seismol. Soc. Am.*, *99*, 2012–2019.
- Peresan, A., and G. F. Panza (2012), Improving earthquake hazard assessments in Italy: An alternative to “Texas sharpshooting,” *Eos, Transactions, American Geophysical Union*, *93*, 538.
- Reserve Bank of Australia (2013), *Statement on Monetary Policy*.
- Ruff, L., and H. Kanamori (1980), Seismicity and the subduction process, *Phys. Earth Planet. Inter.*, *23*, 240–252.
- Sagiya, T. (2011), Integrate all available data, *Nature*, *473*, 146–147.
- Santora, M., and E. Fitzsimmons (2015), New York City is spared the worst effect of snowstorm, *New York Times*, January 26.
- Savage, J. C. (1991), Criticism of some forecasts of the national earthquake prediction council, *Bull. Seismol. Soc. Am.*, *81*, 862–881.
- Savage, J. C. (1992), The uncertainty in earthquake conditional probabilities, *Geophys. Res. Lett.*, *19*, 709–712.
- Savage, J. C. (1994), Empirical earthquake probabilities from observed recurrence intervals, *Bull. Seismol. Soc. Am.*, *84*, 219–221.
- Sieh, K., M. Stuiver, and D. Brillinger (1989), A more precise chronology of earthquakes produced by the San Andreas fault in southern California, *J. Geophys. Res.*, *94*, 603–624.
- Simons, M., S. E. Minson, A. Sladen, A. F. Ortega, J. Jiang, S. E. Owen, L. Meng, J.-P. Ampuero, S. Wei, R. Chu, D. V. Helmberger, H. Kanamori, E. Hetland, A. W. Moore, and F. H. Webb (2011), The 2011 magnitude 9.0 Tohoku-Oki earthquake: mosaicking the megathrust from seconds to centuries, *Science*, *332*, 1421–1425.
- Stein, S. (2010), *Disaster Deferred: How New Science Is Changing Our View Of Earthquake Hazards in the Midwest*, Columbia University Press, New York.
- Stein, S., and A. Friedrich (2014), How much can we clear the crystal ball? *Astronomy and Geophysics*, *55*, 2.11–2.17.
- Stein, S., and A. Newman (2004), Characteristic and uncharacteristic earthquakes as possible artifacts: applications to the New Madrid and Wabash seismic zones, *Seis. Res. Lett.*, *75*(170–184), 2004.
- Stein, S., and E. A. Okal (2007), Ultralong period seismic study of the December 2004 Indian Ocean earthquake and implications for regional tectonics and the subduction process, *Bull. Seismol. Soc. Am.*, *87*, S279–S295.
- Stein, S., and E. A. Okal (2011), The size of the 2011 Tohoku earthquake needn’t have been a surprise, *Eos Trans. AGU*, *92*, 227–228.
- Stein, S., and J. L. Stein (2013a), Shallow versus deep uncertainties in natural hazard assessments, *EOS*, *94*(4), 133–140.
- Stein, S., and J. L. Stein (2013b), How good do natural hazard assessments need to be?, *GSA Today*, *23*, no. 4/5, doi: 10.1130/GSATG167GW.1.
- Stein, S., and J. L. Stein (2014), *Playing Against Nature: Integrating Science and Economics to Mitigate Natural Hazards in an Uncertain World*, Wiley/AGU.
- Stein, S., and M. Liu (2009), Long aftershock sequences within continents and implications for earthquake hazard assessment, *Nature*, *462*, 87–89.
- Stein, S., and M. Wyssession (2003), *Introduction to Seismology, Earthquakes, and Earth Structure*, Blackwell, Oxford.
- Stein, S., B. D. Spencer, and E. M. Brooks (2015a), Bayes and BOGSAT: issues in when and how to revise earthquake hazard maps, *Seismol. Res. Lett.*, *86*, 6–10.
- Stein, S., B. D. Spencer, and E. M. Brooks (2015b), Metrics for assessing earthquake hazard map performance, *Bull. Seismol. Soc. Am.*, *105*, 2160–2173.

- Stein, S., B. D. Spencer, and E. M. Brooks (2015c), Comparing the performance of Japan's earthquake hazard maps to uniform and randomized maps, *Seismol. Res. Lett.*, *87*, 90–102.
- Stein, S., M. Liu, E. Calais, and Q. Li (2009), Midcontinent earthquakes as a complex system, *Seismol. Res. Lett.*, *80*, 551–553.
- Stein, S., R. J. Geller, and M. Liu (2012), Why earthquake hazard maps often fail and what to do about it, *Tectonophysics*, *562–563*, 1–25.
- Stirling, M. W. (2012), Earthquake hazard maps and objective testing: the hazard mapper's point of view, *Seismol. Res. Lett.*, *83*, 231–232.
- Stirling, M. W., and M. Gerstenberger (2010), Ground motion-based testing of seismic hazard models in New Zealand, *Bull. Seismol. Soc. Am.*, *100*, 1407–1414.
- Stirling, M. W., and M. Petersen (2006), Comparison of the historical record of earthquake hazard with seismic-hazard models for New Zealand and the continental United States, *Bull. Seismol. Soc. Am.*, *96*, 1978–1994.
- Swafford, L., and S. Stein (2007), Limitations of the short earthquake record for seismicity and seismic hazard studies, in *Continental Intraplate Earthquakes*, Special Paper 425, edited by S. Stein and S. Mazzotti, pp. 49–58, GSA, Boulder, CO.
- Wang, Z. (2011), Seismic hazard assessment: issues and alternatives, *Pure Appl. Geophys.*, *168*, 11–25.
- Wyss, M., A. Nekraskova, and V. Kossobokov (2012), Errors in expected human losses due to incorrect seismic hazard estimates, *Natural Hazards*, *62*, 927–935.