

Limitations of a Young Science

The centennial of the 1906 San Francisco earthquake is a natural time to reflect on the past, present, and future of seismology. We've come far in little more than the hundred years since the first seismometers were developed in the late 1800s. We've learned a lot about Earth structure and are making progress in using that knowledge to understand the Earth's composition and dynamics. Similarly, we've learned much about the phenomenon of earthquakes and a reasonable amount about the tectonic processes that cause them, and are starting to learn about the physics of the earthquake process.

Still, it's useful to recall how short one hundred years is compared to the time scales of many Earth processes. The planet is billions of years old, plate motions that cause earthquakes haven't changed much in the past few million years, and earthquake histories on many faults seem complicated on scales of many hundreds or thousands of years. As a result, inferences we draw from the short earthquake history available—even adding historical and paleoseismic data—have serious limitations and leave many questions unanswered.

We all know this and can think of examples. While discussing the 1906 earthquake, we recognize that modern seismicity maps don't show the southern segment of the San Andreas on which its 1857 cousin occurred. We realize that we know of this earthquake from historical accounts, and know from paleoseismic data that this fault segment has had a complicated and irregular earthquake history over the past 2,000 years.

This example from the most studied fault system on Earth illustrates why it's worth bearing in mind the short records available, recognizing the limitations they pose for our ability to understand earthquakes, and accepting that the Earth will continue to surprise us. Much of what we expect will prove right—but some will prove wrong.

The great December 2004 Sumatra earthquake is a good example. The segment of the trench between Sumatra and the Andaman Islands wasn't particularly active seismically, wasn't considered particularly dangerous, and wasn't high-risk on seismic gap maps. Little thought had been given to the possibility of a giant earthquake and devastating megatsunami.

In hindsight, the short record biased ideas about where such earthquakes occur. Earthquakes like this probably have

happened before, but with long enough recurrence times that they left no conventional cultural record. Now that scientists know to look, they're starting to find possible paleotsunami records.

Although the earthquake surprised us, it shouldn't have. We know that the earthquakes that occur at trenches and other major faults can be highly variable. For example, the long earthquake history at the Nankai trough shows that sometimes the entire region has slipped in very large earthquakes, whereas at other times only parts have slipped in smaller events years apart. Another example is the trench segment that produced the giant 1960 Chilean earthquake. If this event were to recur as often as the 400-year record suggests, the seismic slip rate would exceed the plate convergence rate. Hence either the 1960 event was bigger than some of the other large historical earthquakes, or the mean recurrence of such events is longer.

Thus ideas about plate-boundary dynamics may also be biased by the short earthquake history. Many of the apparent differences between subduction zones, such as some trench segments but not others being prone to great earthquakes, may reflect the short history rather than differences in underlying physics. The seismic coupling hypothesis characterizes these differences in terms of either the largest earthquakes or the highest fraction of the plate motion released as earthquakes, and seeks to relate them to parameters such as convergence rate and plate age. No clear correlation has emerged, however. Moreover, the Sumatra earthquake would not have been expected from the idea that strong coupling and thus giant ($M_w > 8.5$) earthquakes occur only when young lithosphere subducts rapidly. Hence we don't know whether trenches where we haven't seen large thrust earthquakes are weakly coupled and thus unlikely candidates for great earthquakes and megatsunamis—or simply have a long recurrence time. Where some of the expected seismic moment release is missing either there's a seismic gap—or there isn't. Detailed studies will be needed to resolve this issue. For example, GPS data show that sites in the southern Lesser Antilles move as though they were part of the Caribbean plate, implying that the effects of the earthquake cycle and coupling are weak.

More generally, the fact that the short earthquake history makes it hard to distinguish aseismic motion from long recurrence times makes it hard to assess whether the apparent regional differences between the fractions of plate motion or

[I]nferences we draw from the short earthquake history available ... have serious limitations and leave many questions unanswered.

intraplate deformation that are released seismically are real. It looks like essentially all of the expected motion occurs seismically on the San Andreas and in continental interiors, whereas trenches, oceanic transforms, and continental plate boundary zones often appear to have significant aseismic motion. This may tell us a lot about differences in rheology and deformation—or it might be a sampling artifact.

At a fundamental level, the short earthquake history means we still don't know the most basic question about earthquake recurrence: whether it's time-dependent or time-independent. Most of us think the former—in the earthquake cycle or elastic rebound model the probability of a major earthquake increases with time since the last one on a fault segment—and so we think in terms of seismic gaps. Experiments based on the idea, however, either global or at specific sites such as Parkfield, haven't entirely succeeded.

Moreover, the recurrence intervals shown by paleoseismic records can be interpreted in various ways. A time-dependent model predicts quasiperiodic earthquakes, whose recurrence times have a standard deviation smaller than the mean. In contrast, a time-independent model predicts clusters of earthquakes—because a large earthquake can occur shortly after another—and so recurrence times with a standard deviation close to the mean. Two of the longest records we have lead to opposite conclusions: Large earthquakes on the Nankai trough appear quasiperiodic, whereas large earthquakes at Pallett Creek on the San Andreas look clustered. The latter may show either time independence, elastic rebound perturbed by stress transfer from neighboring faults or fault segments, or something else. This is going to be a hard question to settle, especially because the answer may differ between faults. Perhaps trench earthquakes occur in a simpler geometry, are less affected by stress interactions, and are thus more regular than continental transform ones. Much longer earthquake records, some of which can come from paleoseismology, are needed to address these questions. Simulations of fault histories can also help.

Even so, we face the fundamental challenge that because earthquake cycles for large earthquakes are longer than 100 years, our seismological and geodetic observations do not span even one earthquake cycle. As a result, our models are derived from combinations of shorter data spans on different faults. This process seems to give a reasonable average picture, but observations of multiple cycles on individual faults are needed before we know if we've missed some crucial features.

The short history makes it hard to understand earthquake recurrence at plate boundaries, where earthquakes episodically release strain that geodetic data show accumulates smoothly over time. The challenge is even more complicated for plate interiors. We've learned from GPS that continental interiors deform very slowly, less than 1–2 mm/yr, which limits the long-term steady-state seismic moment release rate. We've also learned from paleoseismology that faults in continental interiors can turn “on” and “off”, such that where the small strains

are released varies. This makes sense, since the slow deformation isn't enough to generate large earthquakes on many fault systems. Hence the past hundred years of seismicity may not tell us much about long-term patterns.

For example, although the New Madrid seismic zone has had the most visible seismicity in the North American continental interior, we've seen only a two-hundred-year snapshot. Geologic data indicate that this zone has been active for less—probably much less—than a million years. We don't know why the large earthquakes started, or how long they'll continue. Rock-friction studies imply that aftershock sequences should be longer in intraplate areas than at plate boundaries due to slow loading, so today's small earthquakes are likely to be aftershocks of the large events of 1811–1812. Hence although this is a likely place for continued small earthquakes, we don't know if future large earthquakes are more

likely here than in other regions within the continental interior that may be equally or more susceptible to strain concentrations. Paleoliquefaction studies show prior sequences around A.D. 900 and A.D. 1450, but we do not know if such sequences will continue. After all, nothing appears to be very special about the Reelfoot Rift. Recent data show little or no heat flow anomaly across it, so it's unlikely to be much hotter or weaker than its surroundings. Lots of similar fossil structures are out there, and it is possible that seismicity migrates among them. The Meers Fault in Oklahoma—with no present seismicity but with evidence of motion in the past few thousand years—may be just one such structure.

Similarly, although New Madrid is today's best example of large earthquakes in a continental plate interior, I doubt it's very different from other such zones. Intraplate seismic zones in Australia, northwest Europe, and the Pannonian basin deform at similar rates, and so have similar maximum earthquake magnitudes and recurrence rates.

As a result, the short record poses a major challenge for efforts to estimate future seismic hazards, especially where the plate motion or intraplate deformation rate is slow enough that the recurrence times of large earthquakes are long compared to the record. Almost every aspect of hazard estimation faces this challenge, because hazard estimates seek to quantify the shaking expected during periods of time (once in 500 years in California and other most countries, once in 2,500 years in the central and eastern U.S.) that are much longer than the seismological records.

The first issue is deciding where large earthquakes are likely. Seismic hazard maps for places such as the North African coast, North America's eastern continental margin, and the St. Lawrence Valley sometimes show “bull's-eyes” of high predicted hazard where we know from instrumental or historic records that moderate to large earthquakes have occurred. There's no reason to believe, however, that these sites are more likely to have future large earthquakes than other sites on the same structures that should deform at similar rates. In fact, stress transfer

[W]e still don't know the most basic question about earthquake recurrence: whether it's time-dependent or time-independent.

arguments imply that large earthquakes at the other sites may be more likely. Hence although U.S. earthquake hazard maps are still based primarily on the short earthquake record, other countries are starting to also use the geology in maps that predict more uniform hazard along similar structures.

The second issue is inferring the maximum sizes and recurrence intervals of future earthquakes in a given area. Where long earthquake records from plate-boundary segments exist, they show variability in the sizes and recurrence times of large earthquakes. Hence a short earthquake record from an area with long recurrence times is likely either to miss the largest earthquake entirely, or preferentially to detect large earthquakes with recurrence times shorter than the average. As a result, frequency-magnitude (*b* value) studies are likely either to underpredict the sizes of the largest earthquakes, or conclude that they are “characteristic”—more common than expected from the rate of smaller earthquakes. Adding historical and paleoseismic data is very valuable, but combining these data with seismological data is tricky. Historical studies add events with known dates, but with considerable uncertainty in magnitudes. For example, magnitude estimates for the 1906 San Francisco earthquake—based on early seismological data—have been as high as 8.3, compared to the typical current value of 7.9. The challenge is even greater for pre-instrumental data; recent results suggest low M7 magnitudes for the largest 1811–1812 New Madrid earthquakes, but estimates still range from low M7 to over M8. Paleoseismic studies have uncertainties both in the estimated dates and in recurrence times due to possibly missed events—and even larger uncertainties in estimated magnitudes. For example, coauthors and I have concluded that paleoliquefaction analysis for New Madrid has overestimated the sizes of paleoevents, producing apparent characteristic earthquakes.

We can assume earthquakes are most likely in parts of a seismic zone where they've happened recently, more likely where they haven't happened recently, or equally likely throughout the zone.

A third issue is that we lack instrumental records of strong ground motion in most areas of low seismicity. Hence hazard maps for areas such as northwestern Europe and the eastern U.S. depend crucially on the assumed ground motion model.

A fourth issue is that because the earthquake record is too short to resolve whether earthquake recurrence is time-dependent or time-independent, it's not clear what to assume in hazard maps. We can assume earthquakes are most likely in parts of a seismic zone where they've happened recently, more likely where they haven't happened recently, or equally likely throughout the zone. The predicted hazards vary: Time-independent models predict the same probability of a large earthquake

regardless of the time since the last one, whereas time-dependent models predict lower probabilities for the first two thirds of the mean recurrence interval, and then higher probabilities as the earthquake is “due.” There's no standard choice: Some California maps have been based on time-dependent probabilities, whereas the central U.S. maps are based on time-independence. In each region

these opposite assumptions tend to predict higher probabilities than the alternative, due to the longer recurrence time in the central U.S.

Problems resulting from the short earthquake record are going to be tough, if not impossible, to fully solve on time scales shorter than thousands of years unless we learn a great deal more about the underlying earthquake physics. Naturally, we won't stop trying. All we can do is keep doing our best, hoping to do better as we learn more, while maintaining healthy humility about what we don't know in the face of the complexities of nature. No matter what we do, our young science will continue to encounter major surprises such as the Sumatra earthquake. ■

Seth Stein
Northwestern University
seth@earth.northwestern.edu