Engineering and Science
MAY-JUNE 1976

Defending Against Disaster

by ROGER G. NOLL

What we know—and what we don't know—about the adequacy of society's defenses against a major earthquake

PUBLIC POLICY-MAKING about earthquakes is very difficult for a number of reasons. A major problem is simply how to determine the magnitude of the threat. Devastating earthquakes are particularly difficult to plan for because they are extremely infrequent. Rational behavior is especially elusive to define when the threat is a tiny probability of a major catastrophe.

The first step in attacking the public-policy problems related to earthquakes is to state the magnitude of the damage in terms that make the threat of earthquakes comparable to other hazards that we face in everyday life. One such measure is the average annual destruction from earthquakes, which can be calculated by dividing the damage from major earthquakes by their frequency.

Unfortunately, because the earthquakes that we're most interested in — the ones that cause widespread damage — are so infrequent, we can make only the crudest estimates of the likelihood that one is going to happen this year. Probably on the order of once every 100-200 years an earthquake of major proportions, like the San Francisco earthquake of 1906, will strike a major population center in California. We can't be much more specific than that, because we don't have observations over a long enough period of time to say more.

A second part of the problem has to do with the nature of scientific inquiry into earthquakes. The people who study earthquakes, the seismologists and the geophysicists, do not focus their attention primarily on the earthquakes that government officials are most interested in for public policy-making purposes. Naturally, the scientists focus on the quakes that occur frequently enough to enable them to collect enough statistical data to test their hypotheses about the nature of the earth. And the earthquakes that happen most frequently are small ones that do little or no damage. As a result, estimates of the relationship between the frequency and the size of earthquakes are fairly good for quakes that public policy-makers don't care much about, but they're atrocious for the ones that matter most.

Another feature of the problem is that even if we knew how often earthquakes of magnitude 7.5 and up occur, we still would not know enough to estimate the damage they would cause. The relationship between the damage to a building standing in Los Angeles, and the magnitude number of an earthquake on the San Andreas fault is very loose. The exact amount of damage will depend upon the type of ground motion created by the earthquake, the time of day that it occurs, and numerous other uncertainties about the location and design of each building.

One way of getting at the likely damage is to examine the historical trends. The annual property loss from disasters of all kinds, including earthquakes, has risen fairly rapidly since 1900, for quite obvious reasons. We're a far richer society, and far more people live here than in the past. Still, despite the population growth that's taken place, the average annual number of people who are killed by a natural disaster has fallen quite sharply. Crude estimates of the expected annual death rate in California due to earthquakes come to about 30 people a year. That's a remarkably small number in comparison to accident rates from other things. For example, several thousand people are killed every year in auto accidents in the State of California.

The proper perspective to take in making decisions about the allocation of resources for safety is in terms of the relative magnitude of threats, and the threat of death from earthquakes is quite small compared to death by automobile accidents. We focus on earthquakes because a single event is so disruptive, possibly producing a larger number of deaths than numerous other threats that, over several years, claim a much higher toll.

To identify opportunities for improving public defenses against earthquakes it is useful to categorize the kinds of damage that will take place and the cause of each kind. Primary damage refers to the direct results of the quake — a building falls over on some people; a dam breaks, and people are drowned or property is washed away. Secondary damages occur after the earthquake, as a result of the disorganization of society that comes about because of it. For example, the earthquake might reduce the ability of the city fire department to fight fires; consequently, a substantial number of fires might go unchecked. Or it might disrupt the water supply, the sewage system, or medical care facilities, so that a few days after the quake an epidemic or other public health problem arises.

Contrary to the imaginings of the motion picture industry, the principal primary threat is not the collapse of buildings

Contrary to the imaginings of the motion picture industry, the principal primary threat is not the collapse of buildings. Since the 1933 Long Beach quake, building codes have required that structures be able to withstand a major shock. Furthermore, the common architectural style in southern California — one-story frame buildings — is ideal for rolling with the punch of even a major quake. Widespread structural collapse of residences does occur frequently in other parts of the world, and notably in Latin America, but the reason is that houses are typically constructed of adobe bricks. Masonry structures are extremely vulnerable to earthquakes, but are relatively rare in California. Actions to demolish these remaining old structures would relegate structural collapse to the status of a relatively minor problem.

In the Los Angeles area the major primary threat is that dams might break

In the Los Angeles area the major *primary* threat is that dams might break. If a substantial number of dams in Los Angeles or Orange County break and break quickly, either during the earthquake or immediately thereafter, about 30,000 people could, conceivably, be killed. But if no dams break, then, even in the year it occurs, the number of people killed by an earthquake will be fewer than the number killed by automobiles.

Historically, the next most important threat has been the secondary threat of fire. A not-so-commonly-known fact about the San Francisco earthquake is that about 95 percent of the damage was due to the fire that broke out after the quake, which the fire department was unable to fight because of inadequate water supplies. Fires are less of a threat in spread-out Los Angeles, but breaks in natural gas and petroleum pipes could still make fires a serious problem.

Another potentially serious problem is maintaining public health with effective relief programs. In addition to medical help for the injured, immediate actions must be taken in response to broken water mains, sewage pipes, and utility lines. A distribution system for water, food, and portable toilet facilities must be set up within a few hours of the quake, in addition to marshalling forces to get normal systems working again.

One public fear of earthquakes — commonly presumed to be true but actually false — is what might be called the disaster-movie syndrome. Panic and riots are expected to be the common reaction in the wake of a major natural disaster, and some public policy in the past has been predicated on such a belief.

But in fact this does not happen. In general, during a natural disaster and for the first minute or two afterwards, people see the situation in very personal terms. They see themselves as the focus of the disaster, and they see the major threat as personal, affecting themselves and their family and friends.

Their first action is to try to get home, to make certain that their family and friends are safe, and to be comforted by familiar surroundings and people.

The second response is to engage in constructive activities to cope with the damages in a direct sense — to keep busy, to work. Over a fairly long period of time — a few weeks — people try to come to grips with repairing the damage in a quite rational and straightforward way. Violent reactions or depressions that immobilize people come later, if at all.

It's interesting to compare the findings of social scientists with the actual response to the Alaska earthquake. The police department there held the incorrect view that the chief threat was looting, rioting, and chaos. So they strengthened themselves in the downtown area where, because store windows were broken, the potential for looting was high. But this kind of antisocial behavior never happened, and it was several hours before it occurred to someone that the police had nothing to do and could effectively be used to rescue people who were trapped in collapsed buildings.

I think the payoff for investing our resources in disaster relief is substantially higher than for making stricter building codes or taking other preventive actions

Meanwhile, several hours were wasted that could have been used to find and treat injured people.

Aside from assessing the magnitude of the damage correctly, a major policy problem is how to choose the extent to which one wants to be protected — that is, how much, in terms of property and human injury and death, are we going to try to save? The common view among public officials and engineers is, first, that the object is to save human beings from death or injury, and second, that the life-saving benefits cannot be valued in monetary terms. Hence, it is said that economic analysis is irrational or inappropriate for aiding in determining public policy.

This conclusion does not follow from the first two observations. It is certainly true that economic analysis cannot be used to determine an optimum policy; that is, it will not identify the right amount of earthquake defense to undertake, since the benefits of the program (lives) cannot be measured in the same units as are the costs (dollars). So I'm not about to tell you exactly how strict the building codes should be or exactly how much we should spend on disaster relief. But comparisons can be made of the extent to which additional expenditures in various safety programs would differ in the number of deaths and injuries they would avert.

Society has scarce resources to devote to safety, and ought to devote them to activities that are most effective. No one is about to propose that all of the gross national product, or every one of our working hours, ought to be devoted to nothing but protecting ourselves against hazards of all kinds. Because resources are limited, if there is a vast difference between the cost of improving the building code and the cost of, say, making automobiles and highways safer in order to save one more life, a serious public-policy question has at least been raised.

Some estimates have been made about how much money it would take to make buildings sufficiently safe so that the *annual* equivalent number of lives lost due to earthquakes would fall from 30 to fewer than 5 — in other words, to save perhaps 25 lives per year. That is really a big number. Multiply it by 100 or 200 (which provides an estimate of the toll from a big earthquake), and you're talking about several thousand lives.

It turns out that the value of a human life implicit in upgrading building codes, even though it amounts to only a 3-5 percent increase in building costs, is approximately $1 million per person. An intriguing cost comparison is with mandatory airbags for automobiles. It is estimated that airbags would save on the order of 10,000 lives per year at a cost of approximately $300,000 per life saved. Both figures, of course, are subject to considerable debate (just like the figures with regard to earthquake costs and benefits), but the orders of magnitude are probably right.

These figures suggest a strange public-policy dichotomy. Building codes implicitly place a far higher value on human life ($1 million) than do mandatory airbags ($300,000), and the latter is a policy which the nation has thus far been unwilling to adopt. Numerous other examples could be cited that would illustrate the same disparity.

The point of this is not to say that we definitely ought to have mandatory airbags, or that it's problematical whether we ought to upgrade building codes. The point is that we're not doing a very good job of rationalizing safety. We're not allocating resources among these alternatives so as to achieve the maximal saving of human pain and suffering from the resources we're devoting to it.

The third major policy area to be faced in making decisions about earthquakes deals with how to organize disaster relief effectively. Unfortunately there's not enough information to permit estimates of the number of lives that would be saved from more disaster relief than is currently available. As far as I know no one has even begun to ask that question in a sensible, coherent way, but I think the payoff for investing our resources in disaster relief is substantially higher than for making stricter building codes or taking other preventative actions.

Recently, in California, there has been a substantial increase in local and state government planning for disaster relief immediately following an earthquake. This is, it seems to me, quite laudable. Planning doesn't require much in the way of resources, and — by causing people to think carefully about how to use their resources if an earthquake comes — it can avoid grievous mistakes.

In the Los Angeles area the focus of these plans is primarily on the chain of command that will operate if a disaster occurs and the allocation of responsibilities for various kinds of activity. There's also some focus on maintaining communications so that the people who are at the top of the responsibility chain can have adequate information on which to base decisions.

While this is well and good, there are still some problems. First, there is a tendency to rely for disaster relief on highly structured organizations. And there's also a tendency to rely upon existing institutions to take on different and additional responsibilities in case of a disaster. One difficulty of this approach is that the more hierarchically structured an organization is, and the more complete are the rules and regulations governing its behavior, the less flexible it is likely to be in responding to a new, unusual, and unexpected circumstance. Unfortunately, we don't know what's going to happen if a major earthquake hits Los Angeles. We literally don't know exactly what kinds of damages will be suffered or what demands will be placed on public institutions. As a result, a substantial amount of decision-making flexibility will be needed.

In plans for the Los Angeles metropolitan area, much attention has been given to improving communications. Communications are important in a disaster, and a major problem is to maintain them within relief organizations. This is particularly true after an earthquake, when the normal channels are likely to be disrupted and when the scope of the damage cannot be known in advance. But the dilemma resides in the fact that too much information can be as devastating as too little. And if too much of the decision-making authority is at the top of the organization chart, then decision-makers spend all their time receiving information and very little of it making decisions.

This is not to say that organizations should not be hierarchical. It is simply to say that there is value in decentralization of some decision-making. For example, a simple issue is how to cache emergency supplies — medical equipment, water, food. One possibility is to have a relatively few large caches located where they're easily accessible to the people at the top of the organization, who will then order them to be distributed where needed in response to factual statements by people in the field about the extent of damage in various parts of the community.

Few people, in or out of government, know who will do what in case of a major disaster. That is a mistake that should be rectified

Another extreme is to have numerous small caches that are the responsibility of, say, local fire stations, or local police precincts, or even local community organizations (since people tend to congregate around their own neighborhoods). A counterpart to the old civil defense system could be set up, with civilians having emergency disaster relief equipment on a neighborhood basis.

Decentralization of some resources and responsibilities permits communication of less information to the people at the top of the emergency decision-making system. In any given earthquake, most of the Los Angeles area won't have much of a problem. The difficulty of learning who's in trouble and who isn't is substantially reduced if most of these areas can take care of themselves.

To accomplish decentralization of responsibility requires that plans for emergency relief, including those covering the distribution of emergency resources, be made known to government employees in the field and to the general public. As yet, planning information has not filtered down to any appreciable extent. Few

How should public policy-makers deal with this wonderful new technology being invented by the geophysicists and seismologists—predicting earthquakes?

people, in or out of government, know who will do what in case of a major disaster. That is a mistake that should be rectified.

Public policy-makers now face still another mind-bending problem. How should they deal with this wonderful new technology being invented by the geophysicists and seismologists — predicting earthquakes? This is a strange and wonderful witchcraft. Apparently certain kinds of earthquakes are very close to, if not now, "predictable." In this case predictable means that geophysicists can tell us with a reasonable probability what is going to happen. Instead of saying the chance is 1 in 200 of an earthquake of magnitude 8 occurring in L.A. this year, the scientists might be able to detect when the real probability is one in a million, and when it is 1 in 10, or 5, or 2. That's different, of course, from being able to tell you exactly *when* there's going to be an earthquake, but still prediction information is a significant gain over simple historical frequencies.

Of what possible use is this kind of information? How should we disseminate it? How should it affect decision-making. Some have criticized my courageous Caltech colleague, James Whitcomb (see page 8), for being so open about his research on prediction. They argue that this kind of information should not be made public because people will behave irrationally in response to it — they will evacuate the area in masses, or they will do all kinds of nasty things to each other.

History does not support this view. Studies of warning systems for bombings in World War II, and for tornadoes and hurricanes in the U.S., show that people behave calmly and rationally in response to predictions — particularly if they are used to them and know what they should do to protect themselves. All that need be avoided is sudden changes in the kind of information people are given, because they will not know instantaneously how to respond to it. I suggest that an official and regular process — like the weather report — be established. Perhaps once a week or once a month geophysicists would issue a press release on the current state of prediction. Usually the information released would be that nothing new is known, or that an earthquake of very small magnitude is predicted in some uninhabited area. This information is of no value directly, but it *is* useful.

If a geophysicist says that a magnitude 2.3 earthquake will hit eastern Riverside County (which would do no damage, even if anyone lived there), no one is going to change his way of life. But it helps people learn to make independent judgments on the quality of prediction technology, which is really the key to making rational decisions about such information. Right now no historical index is available to enable us to assess how fast scientists are learning to predict earthquakes. Society will start to pay attention as soon as it acquires enough historical information to judge the validity of prediction information.

At the present time there are some dangerous incentives *not* to make public predictions. One reason — the sort of perverse thing an economist would think of — is that the value of the information is far greater if it's private than if it's public. For instance, suppose you were the only person who knew there was going to be a recurrence of an earthquake of the kind that shook the San Fernando Valley in 1971. If you owned property there, it would behoove you to sell that property immediately and buy some somewhere else. That's a simple, straightforward use of the information. If everyone knows the information at the same time, however, it can't be used to take advantage of other people.

When the technology gets sufficiently good, all kinds of people and firms can be expected to want to have the information for their exclusive use. From the point of view of equity as well as efficiency, we should make the information public as soon as possible so that the private uses of it don't have any unfair economic consequences.

The second reason for not providing information is that if one's business is the scientific prediction of earthquakes, one will, for professional reasons, strive for scientific certainty before going public. Conservatism in evaluating experimental data is a necessary component of successful science, and because the data that are used in earthquake predictions are still sparse, the community of seismologists and geophysicists is reluctant even to talk publicly about prediction.

Of course, the requirements of scientific proof are often quite different from what we might want societally. Airplanes were built long before scientists understood flight, and wireless was extensively used before it was understood. More to the point, the earthquake in Oroville in August 1975 was predicted, but outside of a small group of scientists, no one knew it. Several geologists (not at Caltech) observed that some small earthquakes were occurring on a long-dormant fault. The situation was similar to one in which the Chinese claimed that they had been able to predict an earthquake, but the geologists who were aware of this situation had no real explanation for the mechanics of this kind of earthquake prediction. And they didn't feel that they should make a statement either to the scientific community or to the press that they were expecting a quake near Oroville. Several of them did station themselves in the area so they could see if anything happened, and an earthquake did indeed occur.

In this particular case there was no substantial damage that could have been avoided had the general public known the earthquake was going to happen. But if it had broken the Oroville Dam, the ethics involved in withholding this information from the public would, indeed, have been dubious. In fact, there is an interesting principle of liability law to the effect that if you possess information that another person could use to avoid damage to himself, and if you withhold that information from him, you may well be liable for that damage.

On those grounds there's not likely to be a lawsuit against people who predict earthquakes, because the technology is simply not good enough for any reasonably certain prediction. But the point remains that some individuals might well regard themselves as being able to make use of that information — or at least would regard themselves as being better off if they could see for themselves the extent to which prediction technology is improving, and thereby take their own action in response to it.

What can government do with prediction information? First, it can expect a couple of things to happen that will change its responsibilities. For example, the existence of prediction technology will have quite a devastating effect on earthquake insurance. If eventually earthquakes can be predicted accurately, the absence of a prediction will be valuable information, too, because it will mean that an earthquake is extremely unlikely, and so there is no reason to have insurance. If a prediction comes along, then insurance is attractive, but no sensible insurance company would sell it.

If that's the case, the demands are greater on government to act as an implicit insurer in the form of disaster relief programs. One of the consequences of the development of prediction technology is likely to be transferral of part of the responsibility for compensation for earthquake damage from the private insurance companies to disaster relief programs.

In addition, numerous little things can be done if a prediction takes place. Dams can be drained, and engines can be removed from the fire station so the firehouse won't collapse on them and make them inoperative. But the number of things you could or would want to do in response to this information is quite limited compared to the damage that would take place.

Perhaps the most important consequence of accurate predictions will be the stimulus they will provide to take simple precautions, such as bolting bookcases to the wall. Because major quakes are so infrequent, there is not much risk in postponing defensive actions for a little while. But with predictions, the risk will rise. Plans will be taken more seriously, and people will be more responsive to instructions regarding damage prevention and relief. Furthermore, predictions may well reduce the psychological stress caused by earthquakes by eliminating some of the surprise and the sense of helplessness one feels during and immediately after a quake.

No stupendous, all-encompassing response at either the public or private level will come about from the existence of predictability. When people have access to this information, and when most of it turns out to be about minor earthquakes that do not threaten them or require them to make immediately cataclysmic decisions, it seems reasonable to expect them to respond rationally. Thus, there is no particular reason not to make predictions public. Furthermore, there are all kinds of good reasons to make it public *now,* even if the people who could provide this information do not have the proper incentive to do so.

As nice as prediction is scientifically, as nice as it looks as a research topic today, the likelihood that anybody is going to get a substantial amount of benefit from it in the short run is very small. Even in the long run it is not too great. But the fears of giving this information to the public are completely without foundation. The best way to deal with predictions is to make them open and public as quickly as possible.□