Earthquake Forecasts: The Life-Saving Potential of Last-Minute Warnings

by

Ármann Ingólfsson

B.S., Industrial Engineering, University at Buffalo, 1989

and

S.M., Operations Research, MIT, 1991

Submitted to the Department of Electrical Engineering and Computer Science
in partial fulfillment of the requirements for the degrees of

Doctor of Philosophy
in Operations Research

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 1994

Author ....
Department of Electrical Engineering and Computer Science
30 August 1994

Certified by ....
Arnold I. Barnett
Professor of Operations Research and Management
Thesis Supervisor

Accepted by ....
Robert M. Freund
Acting Co-Director, Operations Research Center

Earthquake Forecasts: The Life-Saving Potential of Last-Minute Warnings

by

Ármann Ingólfsson

Submitted to the Department of Electrical Engineering and Computer Science
on 30 August 1994, in partial fulfillment of the
requirements for the degrees of
Doctor of Philosophy
in Operations Research

Abstract

An attempt was made to estimate how many lives could be saved if unequivocal earthquake warnings were possible before major earthquakes. The estimates are for urban areas in California and are contingent on the best-case assumptions that (1) the best available communication strategies are used and (2) individuals react to the warnings by immediately taking the most beneficial action (e.g., exit the building), under the circumstances. The results of a survey of experts were combined with estimates of population fractions for each of a set of categories (e.g., "people at home, awake") and simple mathematical models to arrive at the estimates.

On taking the expert judgment at *face-value*, we estimate that 3 minutes are needed to halve the death toll and an additional 22 minutes reduce the remaining death toll by 50%. Under a *model-based* procedure, where the expert judgment was used to calibrate a set of first-principles based models, the corresponding halving times were estimated to be 1.5 and 28.5 minutes. With one minute of warning, we estimate that 45% of the deaths could be avoided, under both the face-value and model-based approaches. This scenario is of particular interest because with real-time earthquake monitoring such lead-times are thought to be possible for some regions.

The implications of three recent earthquakes were examined. Based on mortality patterns in the 1994 Northridge and 1989 Loma Prieta quakes we conclude that commonly quoted death toll estimates for "the Big One" expected to strike California may overestimate both the number of deaths (by a factor of 4 to 15) and the time-of-day variation (by a factor of two) in the death toll. The consequences of the 1988 Armenia earthquake lead us to believe that unequivocal earthquake warnings – were they possible – have greater life-saving potential in the third world in both absolute and percentage terms.

Thesis Supervisor: Arnold I. Barnett
Title: Professor of Operations Research and Management

Acknowledgements

On August 30th, the day of my thesis defense, I read in *the Tech* that a ribbon-cutting ceremony for the new Biology building (building 68) was scheduled for sometime in October. This news was important to me, because for the last three years I have been measuring my progress on this thesis against the pace at which this building has risen from the ground. For a while, it looked as if the building would be completed before my thesis. To console myself I first reasoned that the comparison was unfair, since many more people were involved in the construction of the biology building, whereas I was laboring alone on my thesis. Then, I came up with a much better rationalization: I should be comparing the date at which I defend my thesis with the date at which the Biology building is officially opened; not the date at which the first tenants move in.

Of course, my first rationalization was not very good: I was not alone. Many people have helped me complete this thesis and I would like to thank some of them. I would like to thank Arnie Barnett for supervising my thesis: thank you! I have learned many valuable lessons from you during the last three years; sometimes concerning matters that I thought I had nothing more to learn about. Working with you has been thoroughly enjoyable.

The other members of my thesis committee – Professors Al Drake, Amedeo Odoni, and Robert Whitman – all nudged the thesis in directions that cause me to be happier with the final product. I thank Professor Whitman for sharing his wealth of knowledge about earthquakes with me. I thank Al Drake and Amedeo Odoni not only for taking an interest in my thesis but also for providing me with financial support during its preparation. I would also like to thank Al for personal advice, which was not always solicited, but always appreciated.

This thesis could not have been completed without the effort of the many experts that completed my questionnaire. Most of these people I have never met; this makes me all the more grateful for their efforts. I have had valuable conversations with Professors Norm Rasmussen, Leonard Morse-Fortier, Daniele Veneziano, and Dick Larson; I thank them all.

I cannot think of a better environment to work in than the Operations Research Center. I thank Paulette, Laura, Cheryl, and other staff members before them. I know there are only two reasons why the place is not run perfectly: students and Professors. Even with those two groups around, you do a great job.

I thank the many students I have met for their friendship; I hope it will continue long after we have all graduated from MIT. Raghu: It has been great having someone to share the trials and tribulations of the final stages of thesis-writing with. Rob: Thanks for letting me run past you once in a while. It's been fun strategizing about how to deal with ... Oops, I better not mention that yet! Joe, Kerry, Mitchell, Sungsu, Pieter, Michael, Ediel, Alan, Rama, Rich, Rodrigo, and many others that I don't mention but should: It has been my good fortune to know you; I hope we will stay in touch.

Finally, I thank my parents, my brother, and my sister for their support.

Dedication

I dedicate this thesis to my grandfather, Ármann Dalmannsson, who was born 100 years before it was finished. I did not know him for long, but long enough to recognize that he was a great man.

Contents

Chapter 1

Introduction

Earthquakes seem to frighten and intrigue people at the same time. One's faith in the ground one walks on, as evidenced by such phrases as "solid ground," is among the items that are shaken. Or so the author imagines, having never personally experienced a significant tremor. Compounding this feeling of helplessness is the suddenness with which earthquakes occur. Compared to such natural disasters as floods and hurricanes, our current abilities to foresee when and where earthquakes will occur are limited indeed.

Worldwide, the risk of future disasters caused by earthquakes is probably growing. While the number of earthquakes capable causing widespread damage is likely to remain relatively stable, the size of the population at risk continues to increase. Most of the population increase occurs in the third world – often in seismically active regions – where the inhabitants cannot afford the seismically resistant construction practiced, for example, in California and Japan.

While understanding of the causes and nature of earthquakes has advanced tremendously since the days when they were thought to be "the instruments of displeasure of the Greek god Poseidon, the spiteful wriggling of the subterranean catfish *Namazu* in Japanese mythology, and punishment for sinners in Christian belief" [CS92], the state of the art has not reached the stage where seismologists are routinely able to pinpoint the location, time, and magnitude of future quakes. Some think that the processes that cause earthquakes are so complicated that this state will never be reached. And yet there are hopeful signs. Geological conditions are such in some places, notably Mexico City, that the epicenters of potentially destructive quakes are far enough away that news of their arrival, traveling at the speed of light, can reach people up to 50 seconds before the seismic waves [Lom94b] (this is referred to as *real-time earthquake monitoring*). A group of three Greek scientists claim to be able to identify signals of impending earthquakes by measuring the voltage difference between two electrodes implanted in the ground [VL91] (this claim is controversial [Bur85]). Significant research budgets continue to be allocated to earthquake prediction research in the United States and Japan.

While it is by no means certain that great progress in predicting earthquakes is around the corner, the mere fact that public funds are used to finance such research raises the question of how valuable foreknowledge of earthquakes would be. Reflecting a popular viewpoint that a little such knowledge can be a dangerous thing, the Economist [UNK91] notes:

> Better advance warning of an earthquake could cut the death toll – but there is always the chance that, in the hands of a panic-stricken populace, an accurate

forecast might actually do more harm than good.

Such an assessment, however, might be a bit glib. Sociological research indicates that most people do not react to earthquakes in a manner that is characterized by "sudden, unreasoning, hysterical fear," as *panic* is defined by Webster's dictionary. Experiences with earthquakes while working as a reporter in San Francisco prompted Mark Twain to "set it down ... as a maxim that the operations of the human intellect are much accelerated by an earthquake" [Twa91, page 45]. The suggestion that foreknowledge of an earthquake – communicated to the public by sirens, radio, and so on – could instill more widespread panic than the actual shaking of the ground is not obviously plausible.

With a few minutes of warning, people in or near vulnerable structures might reach adjacent open areas that are much safer. Freeway motorists about to enter a bridge or a series of overpasses might be able to avoid doing so. With state-of-the-art communications and training, even a short warning might be sufficient to reduce the earthquake's death toll substantially.

The premise of this thesis is that within the rich literature on earthquake prediction and preparedness, too little attention has been paid to what can optimally be accomplished with an accurate forecast. We have undertaken a series of endeavors that, while not purporting to be definitive, might nonetheless help clarify what can realistically be achieved taking account of existing technology and of the fallibility of human beings under stress. The focus will be on deaths caused by earthquakes, rather than on damage to buildings, injuries, or social disruption.

Large earthquakes occur inrequently at any given location; this makes studying them difficult. Even though more than twenty earthquakes of magnitude 7 or greater occur somewhere on Earth each year, on average, the time until earth scientists "catch" such a quake in a region that has been extensively equipped with measuring devices (as in Parkfield, California) may be very long. Studying the consequences of earthquakes on humans and societies is even more difficult because it is harder to generalize about the behavior of humans than it is to recognize patterns in the behavior of faults.

Because of these difficulties, a major element of the research was the canvassing of a variety of experts with a probing questionnaire. Most of the experts work in the United States and are probably most familiar with the likely effects of earthquakes in California. This is one reason why our efforts will be directed at earthquakes in California and their unique pattern of consequences. Other reasons for our focus on urban areas in California are: (1) the population of California has been exposed to considerable amounts of information about earthquakes, (2) research in the United States on the science of earthquake prediction – notably the Parkfield Prediction Experiment – is concentrated in California, and (3) the effects of unanticipated earthquakes striking San Francisco or Los Angeles have been studied extensively.

Nevertheless, we hope that our approach is amenable to estimating the life-saving potential of earthquake warnings elsewhere. In chapter 6 we discuss how the conclusions we reach for California would differ in the third world, based on the experience in Armenia, where a devastating earthquake struck in 1988. More generally, our approach may suggest how to proceed in evaluating the consequences of having a last-minute warning before an irreversible disaster, natural or man-made, occurs.

Except for chapter 7, which presents a model of mass evacuation, we will restrict our attention to *last-minute* warnings, i.e., we will not be concerned with lead-times beyond an hour or so. Two reasons for this emphasis are (1) the one currently feasible way of knowing when damaging seismic waves will arrive is real-time earthquake monitoring, which can provide lead-times on the order of a minute, and (2) scant attention has been paid to the unique considerations that arise when the length of the lead-time limits how many people will hear the warning and what they can do to increase their chances of survival. The research on human response to disaster warnings that we are aware of appears to assume implicitly that the lead-time will be measured in hours or days.

Our main conclusion is that lead-times of around three minutes may be sufficient to halve the number of deaths that can otherwise be expected if a major earthquake were to hit a Californian urban area. A twenty-five minute lead-time might reduce the death toll by three-quarters. This conclusion is contingent on a series of optimistic, or best-case, assumptions that we will discuss and whose effect on the main conclusion we will try to assess.

Next, we preview the spirit and discuss the motivation for the modeling approach we have adopted to tackle the problem. Then, we outline the specific steps we have taken and our findings. The chapter ends by describing how the rest of the thesis is organized.

1.1 Can Earthquake Prediction Save Lives? A Rudimentary Model

Imagine for a moment how earthquake predictions could be used to save lives in a perfect world. In this perfect world, seismologists are able to pinpoint exactly when and where a future major earthquake will occur. Call the lead-time of the prediction τ, i.e., at time t it is known with certainty that an earthquake will occur at a specified location at time $t + \tau$. As soon as a prediction becomes available, a warning is communicated to the people in the vicinity of the epicenter through various means, such as sirens, radio broadcasts, or pagers. On hearing the warning, every inhabitant acts immediately in a way that will maximize his or her chances of survival during the earthquake. An imperfection of this fantasy world is that earthquake hazards do exist: buildings, bridges, and freeway overpasses may collapse, landslides may occur, and so on. As a result, the lead time τ may not be sufficiently long for everyone to reach a perfectly safe location. Let $N_D(\tau)$ be the number of people that do not survive the earthquake. Our primary interest will be in how the death toll N_D depends on the prediction lead-time τ. Before thinking carefully about the relationship between N_D and τ (which will be the main subject of this thesis), consider the hypothesis embodied in the curve in figure 1-1. Like their counterparts in the real world, the inhabitants of the perfect world are not able to react instantly. Hence the conjecture that the death toll will not be significantly reduced until the lead-time reaches a threshold corresponding to a minimum human reaction time. With lead-times incrementally longer than the reaction time threshold, one would be able to improve one's chances of survival by getting under a desk or a table, by bracing oneself next to a load-bearing wall, by moving away from objects that might fall during an earthquake, by getting away from hazardous equipment, and so on. For longer lead-times (several minutes), the action of leaving a building for an open area would be the one taken by most people. For longer lead times still (several hours), the action of evacuating areas where there is a lack of open spaces might be contemplated.

Let us now step back to the real world. The reader may not consider the above an important contribution to the literature of fairy tales, and may ask what the purpose of constructing the non-existent perfect world is. The proposition that will be argued for in this thesis is that by describing the perfect world in unambiguous terms (as a mathemat-

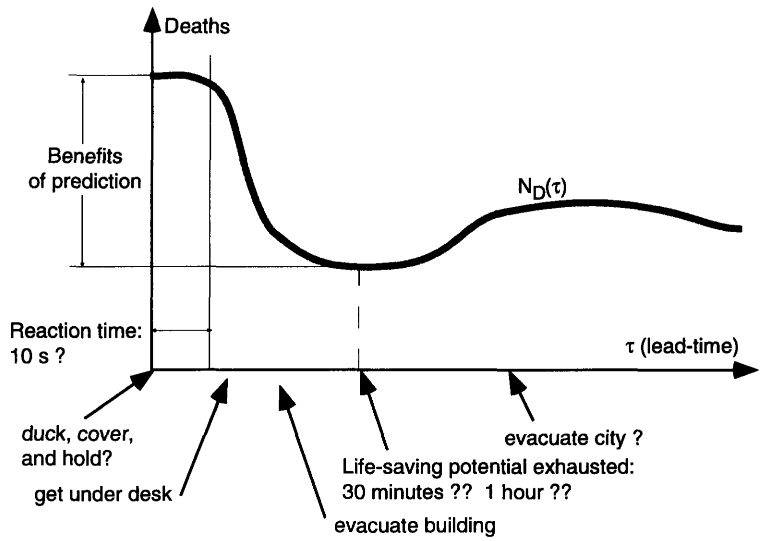

Figure 1-1: First hypothesis for shape of the death toll function.

ical model) and then adding imperfections, one by one, to bring the perfect world closer to reality, one can obtain a better understanding of the potential benefits of earthquake prediction. Such an understanding is a prerequisite for making informed decisions about the allocation of resources for research and development of earthquake prediction methods, for establishing the infrastructure necessary for dissemination of earthquake warnings, and for educating the public.

Even after all the imperfections that we plan to add to the perfect world have been added, what happens in the now less-than-perfect world will be an optimistic estimate of what would happen in reality. Thus, the final product of this study is best regarded as an upper bound on the number of lives that could be saved by predicting earthquakes.

If a lower bound on the cost of developing the science and technology necessary to predict earthquakes were available, one could compute a lower bound on the "cost per life saved" for earthquake prediction:

$$\text{cost per life saved} \geq C \equiv \frac{\text{lower bound on cost of developing technology}}{\text{upper bound on \# of lives saved}}$$

This cost estimate could then be compared to the costs per life saved for other ways to reduce earthquake risk (such as stricter enforcement of building codes or retrofitting of old structures) and more generally to other ways to improve public safety. If the lower bound C on the cost per life saved through earthquake prediction were larger than a reliable estimate of the cost per life saved through other reasonable means, then the analysis would have important and obvious policy implications. On the other hand, if C were about equal or smaller than the cost per life saved through other means, then a more detailed analysis of exactly how optimistic the upper bound on the reduction in casualties is would be called for.

In a sense, the burden of proof is put on earthquake prediction to show that it has the potential to improve public safety. The outcome of an analysis that estimates the lower bound C can be either (1) that more lives can be saved by spending money on other means of improving public safety, or (2) that earthquake prediction deserves more detailed study. Under no circumstances could this type of analysis imply, on its own, that earthquake prediction is a more cost-effective means of saving lives than, say, improvements in construction technology.

Estimating the reduction in fatalities that would be realized if earthquake prediction were possible is a challenging undertaking. The reasons for this are numerous: empirical data on earthquake fatalities under conditions comparable to those in the U.S. at present are scarce, the science and technology of earthquake prediction is in its infancy, human behavior is poorly understood, and the means that would be used to communicate earthquake warnings are uncertain. These difficulties provide one justification for taking the *best-case* approach. We will optimistically assume that people react to a warning in a way that increases their chances of survival, rather than trying to estimate the fraction of people that do so; we will assume that the most effective available methods of communication are used, instead of trying to predict what method would eventually be chosen; and so on. Our intent is less to predict how the process of social, political, and technical forces will shape the future evolution of earthquake prediction than it is to provide the players in the process with an assessment of what the potential benefits of earthquake prediction are, given that the players act "optimally:" people learn (or are trained) to respond appropriately to warnings and the most effective communication strategies are used. In short, the analysis is normative rather than descriptive.

We turn now to the specific steps we have taken to estimate the death toll reduction that might optimally be achieved, i.e., in a world that is perfect in the sense we have described.

1.2 Outline of the Thesis and Its Findings

Much of the work in this thesis is motivated by a simple model of how the death toll in an earthquake depends on the lead-time τ of an earthquake warning. Under this model, the population of the area of interest is divided into a set of groups I. If the fraction of residents that fall into group i is $F_i(t)$ when the time of day is t and the average fatality rate in that group is $D_i(\tau)$, then the overall fatality rate may be written as

$$D(\tau, t) = \sum_{i \in I} F_i(t) D_i(\tau)$$

Examples of groups in the set I are "people at home, awake" and "people in cars, on, under, or near a freeway overpass." The groups are chosen so that the death rate $D_i(\tau)$ in a group can reasonably be assumed not to depend on the time of day t; for example "people at home" are divided into two groups: those that are awake and those that are asleep. Under this assumption, time-of-day variations in the overall death rate $D(\tau, t)$ are accounted for entirely by variations in the population fractions $F_i(t)$.

If one expresses the death rate $D_i(\tau)$ as the product $D_i(0)(1 - R_i(\tau))$, where $D_i(0)$ is the "baseline" death rate that can be expected in a future earthquake that occurs without warning and $R_i(\tau)$ is the proportional reduction in the number of lives lost that can be achieved with a τ minute warning, then the information needed to "calibrate" our model can be divided into three categories: (1) information about how the population of an urban area shifts between groups over the course of a typical day (the $F_i(t)$'s), (2) information about the life loss that can be expected in future earthquakes that occur without warning (the $D_i(0)$'s), and (3) information about how much the life-loss could be reduced if warnings before earthquakes were possible (the $R_i(\tau)$'s).

Empirical data exist for the first two kinds of information, although the information is by no means plentiful. Various kinds of direct and indirect evidence can be used to estimate the population fractions $F_i(t)$. Death tolls in past earthquakes and earthquake loss studies that have been prepared for several U.S. urban areas can be used to obtain estimates of the baseline fatality rates $D_i(0)$ for earthquakes that occur without warning. This category includes all past earthquakes, since earth scientists are not at present able to reliably predict the time of occurrence of earthquakes. As a result, empirical data on which to base estimates of the death toll reduction factors $R_i(\tau)$ that might be possible for positive warning lead-times are nonexistent. For this reason a questionnaire was designed to elicit the judgment of experts regarding this issue.

Rather than ask the experts to estimate the risk reduction factors $R_i(\tau)$ directly, we expressed the death rates $D_i(\tau)$ in terms of more elementary quantities in the hope that this would make the expert's task easier. The expression for the death rates is contingent on a series of optimistic assumptions that were alluded to in our discussion of the "perfect world." In particular, we suppose that an unequivocal warning is available τ minutes before a major earthquake strikes an urban area. Furthermore, we suppose that the warning is broadcast through some means to the public as soon as it becomes available, and persons – once they hear (or see) the warning – take some action that is likely to increase their chances of survival in the imminent earthquake. Under these assumptions, the fatality rate $D_i(\tau)$, interpreted as the probability that a randomly chosen member M of group i will die in the earthquake, can be written as:

$$D_i(\tau) = A_i(\tau)C_i(\tau)D_i'' + (1 - A_i(\tau)C_i(\tau))D_i'$$

where $A_i(\tau)$ is the probability that M will hear the warning in time, $C_i(\tau)$ is the chance that M will complete an action a that improves his chances of survival, given that M has heard the warning, D_i'' is the fatality rate for persons who complete action a, and D_i' is the fatality rate for persons who do not complete action a. For example, action a might be "get under a desk" or "get outside, to an open area free of overhead wires." The questionnaire asks the experts to estimate the quantities $A_i(\tau)$, $C_i(\tau)$, D_i'', and D_i'.

The questionnaire has been completed by experts in seismology, earthquake engineering, architecture, emergency management, the sociology of human response to disaster, and other fields. In addition, a group of operations researchers has completed the questionnaire. The experts were encouraged to suggest actions that would be likely to increase chances of survival and strategies for effective communication of warnings, in addition to actions and strategies that were suggested in the questionnaire. A total of 67 questionnaires were mailed, of which 45 were completed and returned to us, for a response rate of 67%.

Two approaches were taken to analyzing the questionnaire data. The first approach took the expert estimates more or less at *face- value*, i.e., a minimal number of assumptions were injected into the processing of the data. The second approach, which we refer to as *model-based*, used the expert estimates to calibrate the parameters of "first-principles" based models. The models can be thought of as specifying functional forms for the distributions of two random variables: T_A, the time until a randomly chosen person M hears an earthquake warning, and T_C, the time M uses to complete the action most likely to enhance chances of survival, under the circumstances.

Overall risk reduction estimates based on the face-value and model-based approaches agree closely. These estimates indicate that the death rate in a major earthquake could potentially be halved with two to three minutes of unequivocal warning and that it might be reduced by three-quarters with 25 to 30 minutes of warning, under the optimistic assumptions mentioned.

When risk reduction estimates for individual population groups are examined, a dichotomy between groups involving people inside buildings on the one hand and people outside buildings (in cars or outdoors) on the other hand emerges, under both the face-value and model-based approaches. Generally, longer lead-times seem to be needed to reduce the death toll by a given proportion for people outside buildings.

Motivated by the possibility that the more optimistic experts might more reliably assess the best-case risk reduction potential, i.e., what could be achieved under our optimistic assumptions, we then searched for evidence of heterogeneity among the experts, in the sense of identifiable optimists and pessimists. The evidence was inconclusive, but we nevertheless categorized the experts as optimists or pessimists, based on their responses, and computed separate estimates for the two groups. The lead-time needed to halve the death toll was estimated at 1 minute using the optimists and at 5 minutes using the pessimists. The corresponding estimates for a three-quarters reduction in the number of deaths were 10 minutes (based on the optimists) and more than 30 minutes (based on the pessimists).

The experts were also asked to estimate baseline death rates, even though such estimates are available from other sources for at least some of the subpopulations we considered. Baseline death rate estimates based on the expert judgment range from 205 per 100,000 (for "people in cars on, under, or near overpasses or bridges") to 2.5 per 100,000 (for "people in cars on the open freeway"). For people at home, awake, the estimate was 21 per 100,000. These estimates were combined with estimates of the population fractions $F_i(t)$ to assess likely time-of-day variations in the death toll of future earthquakes. The resulting death rate estimates ranged from 25 per 100,000, in the early morning, to 50 per 100,000, in the late afternoon. The factor of two difference is smaller than commonly quoted estimates suggest, and much smaller than suggested by some in the press after the recent Northridge earthquake.

The 1994 Northridge quake is one of three recent quakes (the others are the 1989 Loma Prieta and the 1988 Armenia quakes) whose implications we consider. Comparisons between the Northridge quake, which occurred at 4:30 AM, and the Loma Prieta quake, which occurred at 5:04 PM, indicate that the factor of two variation between different times of day may not be far off. Extrapolations of the life loss suffered in these two earthquakes to scenarios on which commonly quoted death toll estimates [SAL$^+$80] are based indicate that these estimates may be too pessimistic regarding life loss in the "Big One" expected to strike California in the future. Death toll estimates based on the expert judgment generally fell between the extrapolated Northridge/Loma Prieta death tolls and the estimates from [SAL$^+$80].

We used our risk reduction estimates, based on expert judgment, to estimate how many of the 60 Northridge fatalities and 63 Loma Prieta fatalities might have been avoided had warnings been possible before those quakes. Between two and three minutes of warning lead-time were estimated as sufficient to save the first twenty lives in both cases; with an additional twenty to twenty five minutes we estimated that twenty more lives could have been saved. Based on the circumstances leading to the Northridge deaths, we argue that this projection is reasonable, under our best-case assumptions. For Loma Prieta, the projections may be too optimistic (even under our optimistic assumptions) but this is largely because of the special circumstances (collapse of an elevated freeway) leading to most of the Loma

Prieta deaths; circumstances whose future likelihood may have been reduced as a result of the Loma Prieta quake.

Finally, we argue that based on the pattern of fatalities in the 1988 Armenia quake, the life-saving potential of last-minute earthquake warnings may be greater, in both absolute and proportional terms, in third world countries than it is in California. There is not much doubt that the life-saving potential is greater in absolute terms, simply because the loss of life experienced when earthquakes occur in third world countries is orders of magnitude greater than for comparable quakes in industrialized countries. But earthquake warnings, if they were possible, may also reduce the death toll by a greater *percentage* in third world countries, because the fraction of deaths that could not possibly be averted by foreknowledge of the quake (such as deaths that occur after the quake) is likely to be smaller in third world countries.

The main contributions of the thesis are twofold. First, a modeling framework is presented, which improves understanding, provides a guide to the relative importance of different issues, and could be used to evaluate and incorporate new knowledge as it accumulates. Second, expert judgment is used to obtain rough estimates of important parameters in the models.

The next chapter provides background material on earthquake risk and reviews what has been written about earthquake prediction and its implications. It discusses the views of seismologists regarding the present status and future prospects for a science of earthquake prediction, it mentions social science research of how individuals react to disaster warnings, and it critically reviews two papers that attempt to put the costs and benefits associated with earthquake prediction into a comprehensive framework. The following three chapters (3, 4, and 5) form the core of the thesis; they describe attempts to estimate the mean death toll in a future earthquake as a function of the lead-time τ of an earthquake prediction. Chapter 3 discusses the scarcity of empirical data on earthquake casualties in the United States. A questionnaire that was used to elicit expert judgment of some critical parameters of the models presented in chapter 4 is described. Other sources of relevant data are also discussed. Chapter 4 fleshes out the models alluded to in this chapter. Chapter 5 presents the results of the survey of experts, under both the face-value and model-based approaches. The implications of three recent earthquakes are discussed in chapter 6. Chapter 7 is to a large extent independent of the rest of the thesis: it describes models of evacuation, which may be more broadly applicable to evacuation because of threats such as hurricanes, floods, or nuclear reactor accidents. Chapter 8 summarizes the findings of the thesis and offers concluding remarks.

Chapter 2

Background and Literature Review

This chapter describes earthquake risk in general terms and reviews what has been written about earthquake prediction and its implications.

The first section contains background material on the nature and extent of earthquake risk and how the consequences of earthquakes can be mitigated. The remaining sections review the literature. Efforts by seismologists to predict earthquakes will be described and assessments will be made of the present status and future outlook for earthquake prediction. Attempts by social scientists to characterize human response to disaster warnings in general and earthquake warnings in particular will be reviewed. Finally, attempts to quantify the costs of earthquakes and the benefits of various hazard mitigation measures (particularly earthquake prediction) will be described.

2.1 Earthquake Risk and How It Can Be Reduced

Following [KG81] we take *risk* to mean a spectrum of possible events along with their associated probabilities of occurrence and consequences. Accordingly, earthquake risk has three facets:

- There is a spectrum of earthquakes that may occur. The primary quantities used to describe the physical characteristics of an earthquake include magnitude (a measure of energy release as reflected by the amplitude of ground motion), epicentral location (latitude and longitude), and depth below the surface of the earth. Richter magnitude is historically the most common magnitude scale and while steadily being replaced is still widely used today. For example, the 1994 Northridge, California earthquake was reported [Ear94] to have a Richter magnitude of 6.4, a surface wave magnitude of 6.8, and a moment magnitude of 6.7. The surface wave magnitude and moment magnitude scales represent two efforts to construct a scale that is a more consistent measure of energy release in an earthquake than Richter magnitude.

- Corresponding to each possible earthquake is a likelihood of occurrence, relative to other possible earthquakes.

- Also associated with every possible earthquake is a set of consequences which includes the number of casualties and injuries caused by the earthquake, the extent of damage

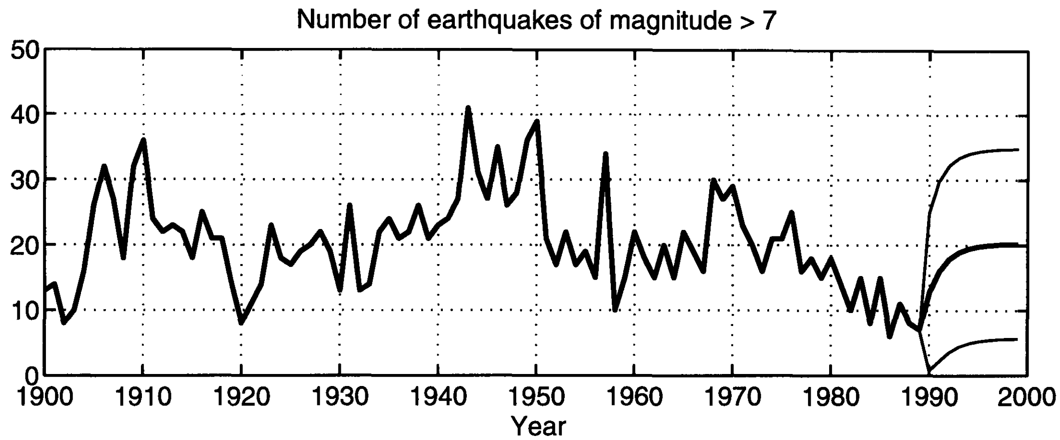

Figure 2-1: Number of earthquakes of Richter magnitude 7.0 or greater per year, throughout the world for the years 1900 - 1989, with 95% forecast intervals for the years 1990 - 1999.

to buildings and other structures, and loss of productivity. The Modified Mercalli intensity scale is designed to summarize the consequences of an earthquake at a specific location.

We now look at some measures of each of the three aspects of earthquake risk. The statistics to follow are meant to be suggestive rather than definitive. We will mention some of the caveats one must bear in mind when interpreting these statistics.

Distribution in time: The first two facets of risk concern the distribution of earthquakes in time and space, and across magnitudes. We first consider distribution in time. Figure 2-1 shows the number of earthquakes per year of Richter magnitude greater than 7.0 (anywhere in the world) for the years 1900 to 1989. It is reasonable to expect that all earthquakes this large in this century have been recorded [BHMS75, page 11]. The average number of such quakes is 20.2 per year, but we see that there is considerable variation around the mean, from a high of 41 to a low of 6. Furthermore, the series exhibits substantial autocorrelation (0.55 at lag 1). In fact, the series is well fit by the autoregressive model:

$$y_t - 20.2 = 0.56(y_{t-1} - 20.2) + e_t$$

The estimated standard deviation of the innovation term e_t is 6.2. Also shown in figure 2-1 are forecasts for the years 1990 to 1999 with (individual) 95% forecast intervals, based on the autoregressive model. The actual number of earthquakes of magnitude 7.0 or more in 1990 (12), 1991 (11), and 1992 (23) are in good agreement with the forecasts.

We will later in this thesis be concerned with whether death tolls are likely to be systematically higher in earthquakes that occur at particular times of day, say during the afternoon rush hour. It is difficult to think of any physical reason for earthquake frequency to differ based on the time of day, i.e., one would expect the time when an earthquake strikes to be uniformly distributed over the 24 hours of the day. The two histograms in figure 2-2 support this hypothesis. The histograms show the local time[1] of earthquakes for which both

[1]Local time was taken to equal Greenwich Mean Time plus the longitude divided by 15 for locations east of Greenwich and minus the longitude divided by 15 for locations west of Greenwich.

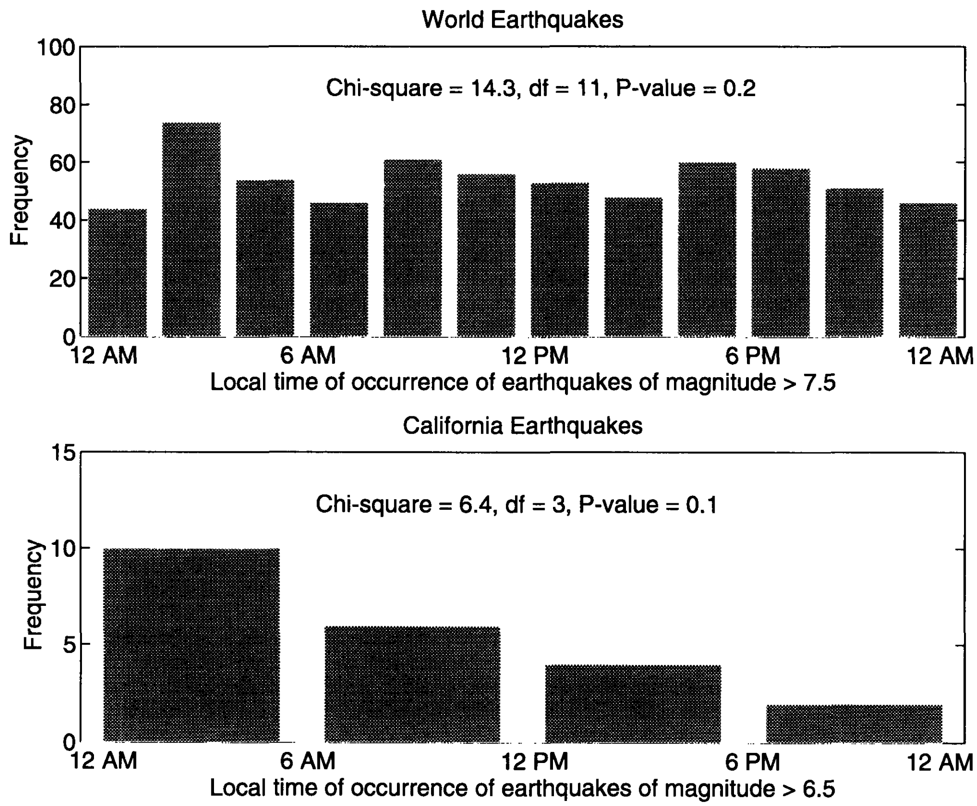

Figure 2-2: Distribution of the local time of occurrence of large earthquakes, for the World and for California. Source of data: NOAA (National Oceanic & Atmospheric Administration).

the magnitude and time-of-day are on record. The top histogram shows world earthquakes of magnitude greater than 7.5 and the bottom histogram shows California earthquakes of magnitude greater than 6.5. According to a χ^2 significance test of the hypothesis of uniformity there is a 20% chance of observing a more lopsided distribution than the one observed for world earthquakes conditional on the uniformity hypothesis being true. Similarly, for California earthquakes there is a 10% chance of observing a less uniform distribution than the one in figure 2-2 even if the underlying distribution is uniform. Thus we have little reason to doubt the uniformity hypothesis.

The histograms in figure 2-2 are based on a database of "significant earthquakes" maintained by the National Oceanic & Atmospheric Administration. An earthquake is included in the database if it meets at least one of the following criteria: (1) it caused damage in excess of about one million dollars, (2) it caused ten or more deaths, (3) it had a magnitude of 7.5 or greater, or (4) the maximum felt intensity was X or greater on the Modified Mercalli scale. These selection criteria lead to an overrepresentation of earthquakes that occur in densely populated areas, where quakes are more likely to cause damage and more likely to be in the historical record. By restricting attention to earthquakes of magnitude greater than 7.5 for which the time of day is on record we avoid possible biases caused by the selection criteria. For earthquakes in California the threshold was lowered to 6.5 in order to have a reasonable sample size for testing the uniformity hypothesis.

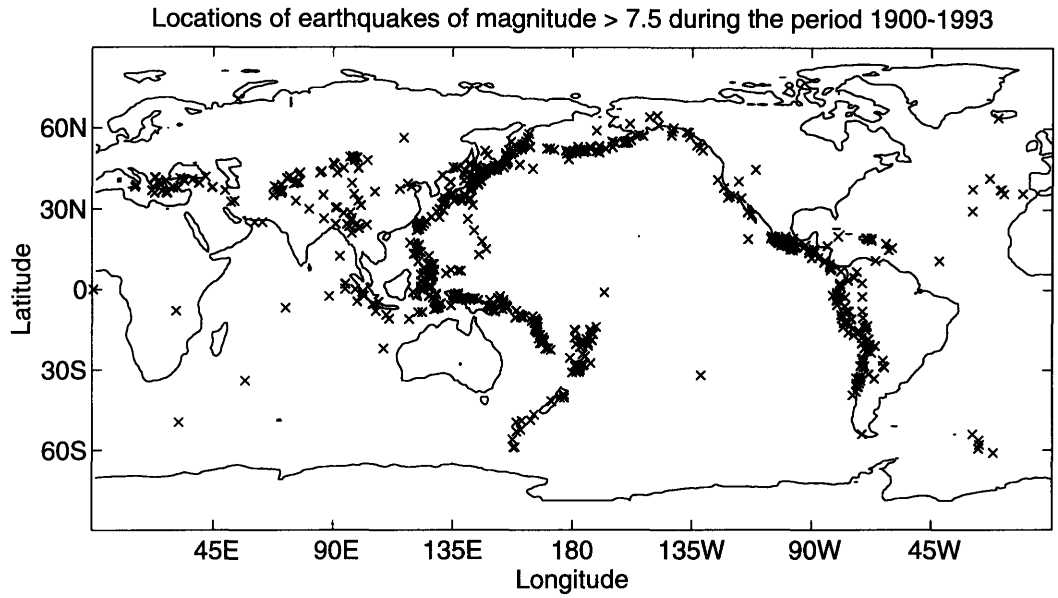

Figure 2-3: Locations of earthquakes of magnitude 7.5 or greater during the period 1900 – 1989. Source of data: NOAA (National Oceanic & Atmospheric Administration).

Distribution in space: The main pattern of the distribution of earthquakes in space is that earthquakes tend to occur along boundaries between tectonic plates. According to [GS84] almost 95% of all earthquakes occur at the edges of tectonic plates. This is illustrated in figure 2-3 which shows the locations of earthquakes in this century with magnitude greater than 7.5. The dominant feature on this map is the concentration of earthquakes along the perimeter of the Pacific Ocean. This earthquake band coincides with converging tectonic plates. In contrast, much fewer large earthquakes occur where tectonic plates diverge, as they do along the Mid-Atlantic Ridge. Besides the circum-Pacific belt, two regions experience large numbers of significant earthquakes: A region extending from the Eastern Mediterranean to Iran and a region which includes the Himalayas and Central Asia.

Distribution across magnitudes: The following relation is often postulated [Ric58] to describe the frequency of earthquakes as a function of magnitude:

$$\log_{10} N = A - bM \Rightarrow N = 10^A 10^{-bM}$$

Here N is the (expected) number of shocks of magnitude M or greater per unit time; A and b are constants. This formula is known as the Gutenberg-Richter law. In probabilistic terms, the Gutenberg-Richter law means that if X is the magnitude of a randomly chosen earthquake, then the probability density of X is proportional to the derivative of N (since N is proportional to the complementary CDF for X):

$$f_X(M) \propto \frac{d}{dM} 10^A 10^{-bM} \propto e^{-b'M} \Rightarrow f_X(M) = b'e^{-b'M} \text{ for } M \geq 0$$

where $b' = b \log_e 10$, i.e., the magnitude has an exponential distribution. Note that the constant A does not appear in the density function for X; this is because A is a measure of the the absolute frequency of earthquakes of any magnitude while the density $f_X(M)$

Magnitude	Western U.S. (excl. AK/HI)	Eastern U.S.	Alaska	Hawai
8 and higher	1	0	7	0
7.0 - 7.9	18	0	84	1
6.0 - 6.9	129	1	411	15
5.0 - 5.9	611	41	1886	36
4.0 - 4.9	3171	335	8362	315

Table 2.1: Number of earthquakes in the United States since 1900 located by the U.S. Geological Survey National Earthquake Information Center.

describes the relative frequencies of earthquakes of different magnitudes. [Ric58] gives an estimated value of 0.9 for the constant b, corresponding to an "average magnitude" of $1/b' \approx 0.48$.

However, [Ric58] also indicates that this simple model may have to be modified to account for the fact that there is an upper limit to earthquake magnitude, due to to the finite strength of the rocks that fracture during an earthquake. The suggested modification is to use one value for the constant b' for magnitudes below some threshold and a larger value for b' above the threshold.

Table 2.1 shows the distribution across magnitudes of earthquakes in the United States that have occurred in this century. Earthquakes with magnitude less than 4.0 are excluded since reporting of such events is uneven and incomplete. The table also gives an indication of the geographical distribution of earthquakes in the U.S. since it has separate columns for the western part of the contiguous United States, for the eastern part, for Alaska, and for Hawaii. A notable feature of this distribution is that more than half of all earthquakes in the U.S. in this century have occurred in Alaska. Nevertheless, efforts to understand earthquakes and to mitigate their effects focus predominantly on California, due to the higher population density in that state.

Having discussed the distribution of earthquakes in space, time, and across magnitudes, we now turn to the third facet of earthquake risk: the consequences of earthquakes.

The consequences of earthquakes: Table 2.2 lists the most deadly earthquakes on record. The only natural disasters in human history with a larger death toll than the 1556 Shansi, China earthquake are two floods which occurred in the Huang He river in China in 1887 and 1931 causing 900,000 and 3,700,000 deaths, respectively. Bruce Bolt [BHMS75, page 19] comments as follows on the 1556 Shansi earthquake:

> The earthquake which cost the greatest known loss of life anywhere in the world occurred on January 23, 1556 in Shensi near the city of Hsiăn. Dynastic records are available which give an estimate of 830,000 people who died from all causes in this earthquake, a death toll so great that one might question the validity of the figures. The explanation is that the earthquake struck a densely populated region, where the peasants lived mostly in caves in the loess hillsides (*Loess* is windblown dust that compacts into thick layers ....)
>
> The 1556 earthquake occurred at 5 o'clock in the morning, when the families were asleep indoors and dwellings collapsed on them. In addition, demoral-

Date	Location	Deaths	Magnitude
1556, 1/23	China, Shansi	830,000	
1737, 10/11	India, Calcutta	300,000	
1976, 7/27	China, Tangshan	255,000 - 655,000	8.0
1138, 8/9	Syria, Aleppo	230,000	
1927, 5/22	China, near Xining	200,000	8.3
856, 12/22	Iran, Damghan	200,000	
1920, 12/16	China, Gansu	200,000	8.6
893, 3/23	Iran, Ardabil	150,000	
1923, 9/1	Japan, Kanto	143,000	8.3
1730, 12/30	Japan, Hokkaido	137,000	
1908, 12/28	Italy, Messina	70,000 - 100,000	7.5
1290, 9/27	China, Chihli	100,000	
1667, 11/?	Caucasia, Shemakha	80,000	
1727, 11/18	Iran, Tabriz	77,000	
1755, 11/1	Portugal, Lisbon	70,000	8.7
1932, 12/25	China, Gansu	70,000	7.6
1970, 5/31	Peru	66,000	7.8
1268, ?/?	Asia Minor, Silicia	60,000	
1693, 1/11	Italy, Sicily	60,000	
1935, 5/30	Pakistan, Quetta	30,000 - 60,000	7.5
1988, 12/7	NW Armenia	55,000	6.8
1783, 2/4	Italy, Calabria	50,000	
1990, 6/20	Iran	50,000	7.7

Table 2.2: The deadliest earthquakes on record. Based on materials from the US Geological Survey and other sources.

> ization, famine, and disease which can follow such a great disaster no doubt accounted for a significant number of deaths.

The primary cause of the tragically large number of lives lost in earthquakes in some regions of the world is the collapse of poorly constructed dwellings. The problem of improving construction practices in the third world is discussed by Bolt [BHMS75, page 61]:

> Dwellings of people in seismic areas vary a great deal, from the adobe and torquezal materials of much of Latin America and the Middle East to timber frame dwellings in New Zealand and California and light wood houses of Japan. The great loss of life associated with historical earthquakes is usually traceable to both poor design and poor construction.
>
> ... One UNESCO study estimated that the additional cost to a one-family house in a medium intensity earthquake region is about 4 per cent, while in areas of major earthquakes the seismic resistant design may add 10 - 15 per cent. Such cost increases may be beyond the purse of poor people.

None of the deadly earthquakes listed in table 2.2 occurred in the United States. Indeed, the human loss suffered by the United States due to earthquakes is far from proportional to its share of the population of the world. Two reasons for this are that despite California's reputation as "Earthquake Country," the frequency of earthquakes is substantially higher in other parts of the world, as a cursory glance at the map of figure 2-3 reveals, and that construction in earthquake-prone areas of the United States is generally of higher quality than it is in most other seismically active regions.

Nevertheless, among natural disasters, earthquakes are among the major causes of death in the U.S. One way to summarize the consequences of different kinds of disasters is using so-called risk curves [KG81, page 13], which relate the number of lives lost per year x, due to a particular disaster, to the frequency with which x or more lives are lost per year. Thus, risk curves are simply complementary distribution functions, $\Pr\{X > x\}$, for the number of lives lost per year. Figure 2-4 shows empirical complementary distribution functions[2] for deaths per year in the U.S. due to earthquakes, tornados, hurricanes and blizzards, and floods. According to these curves, the chances that 100 or more people will die because of earthquakes in a given year are about one percent. For floods the figure is around ten percent, for hurricanes and blizzards it is approximately twenty-five percent, and for tornados the figure is thirty percent.

The curves in figure 2-4 are based on data from the last hundred years or so. The risk attributable to each of the four categories of disasters has evolved during that period and therefore the risk curves must be interpreted with caution. The risk has changed because the potential consequences of natural disasters (the third facet of risk) have changed. The size of the population has increased and its distribution has shifted, and scientific understanding, monitoring, and prediction of natural disasters has advanced.

There are reasons to believe that because of these changes, the risk curves in figure 2-4 may underestimate current earthquake risk relative to the other three categories shown. One reason is the overall westward shift in population during the last century, combined with the fact that earthquakes occur predominantly in the West. Hurricanes, on the other hand, are confined to the eastern seaboard, tornados and floods are most common in the midwest, and blizzards happen more frequently in the northern part of the U.S. Furthermore, great strides have been made in the monitoring and prediction of hurricanes, blizzards, and floods, while the prediction of earthquakes remains an elusive goal.

On the other hand, structures designed in accordance with current building codes are far better equipped to withstand earthquakes than the unreinforced masonry buildings dating from the early part of the century.

Generalizing from the last two paragraphs, the potential consequences of earthquakes are affected by at least four factors:

- The number and magnitude of potentially destructive earthquakes.
- The number of people that live in regions where potentially destructive earthquakes occur. This can be influenced by zoning, where areas with a high potential for seismic damage are identified.
- The quality of construction in areas of seismic hazard. Building codes can be used to ensure that buildings are designed to withstand earthquakes without endangering their occupants, assuming that the owners can afford this.

[2] The complementary distribution function for random variable X is $\bar{F}(x) \equiv \Pr\{X > x\}$. The *empirical* complementary distribution function for a sample $X_1, \ldots, X_n$ is $\hat{F}(x) \equiv \frac{1}{n}\sum_1^n 1(X_i > x)$, i.e., the fraction of the sample which is larger than x.

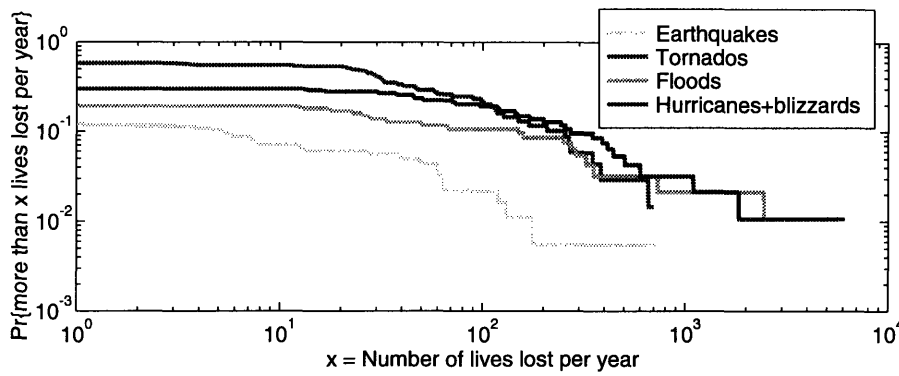

Figure 2-4: Empirical complementary distribution functions for number of deaths per year in the U.S. due to various categories of natural disasters. Note that both the horizontal and vertical scales are logarithmic.

- The population distribution at the instant an earthquake occurs, i.e., the proportion of people that are sleeping in their homes versus the proportion that is commuting to work or school, etc. Unequivocal earthquake warnings could be used to quickly shift this distribution.

Except for the first factor, which can be expected to remain relatively constant on a human timescale, all of the above factors change as people migrate, multiply, learn to build safer dwellings, and understand earthquake risk better. Current population trends seem to suggest that earthquake risk is increasing steadily. The population at risk is increasing in all parts of the world, with the exception of a few industrialized countries (such as France) that are in regions of low seismicity. In addition, there is reason to believe that the overall quality of construction is deteriorating, since population in the third world is increasing faster than it is in developed countries, and because the earthquake-resistance of third-world construction is generally lower than it is in developed countries.

This thesis will focus attention on earthquake prediction as a means of reducing the death toll caused by earthquakes, i.e., we will not be concerned with the other ways of reducing earthquake risk mentioned above.

The remaining sections of this chapter review the literature on earthquake prediction. To relate this literature to the perfect world was described in chapter 1, consider three of the imperfections of the real world that were ignored in our fantasizing about the perfect world:

- Seismologists are not able to routinely predict the time, location, and magnitude of future earthquakes. Even if major breakthroughs were to occur in the science of earthquake prediction in the future, seismologists seem to agree that predictions are likely to be fraught with uncertainty. The current status and future prospects for earthquake prediction will be discussed in section 2.2.

- Real humans may fail to take the action that maximizes their chances of survival for several reasons: They may not be aware of the warning, they may not believe it, they may not think that it applies to them, and they may not know what the most appropriate action is. Social scientists have studied how humans behave during disasters and in response to disaster warnings. Some of their conclusions will be described in section 2.3.

- Rational decisions about allocating resources to the development of earthquake prediction as a means of mitigating earthquake hazards must consider both the costs and benefits of prediction, and how prediction compares to other means of reducing risk. Engineers, economists, and others have attempted to develop comprehensive frameworks for balancing the costs and benefits of prediction. These efforts will be reviewed in section 2.4.

2.2 Is Earthquake Prediction a Realistic Future Possibility?

In 1958, the eminent seismologist Charles F. Richter, after whom the Richter scale is named, wrote:

> At present there is no possibility of earthquake prediction in the popular sense; that is, no one can justifiably say that an earthquake of consequence will affect a named locality on a specified future date. It is uncertain whether any such prediction will be possible in the foreseeable future; the conditions of the problem are highly complex [Ric58, pages 385 – 386].

Despite these reservations about the feasibility of earthquake prediction, research programs were established soon after these words were written, first in Japan and later in the U. S., the former Soviet Union, China, and elsewhere [Gel91]. In the 1970s, seismologists were optimistic about a possible breakthrough. Frank Press, former president of the National Academy of Sciences was quoted in 1975 as saying "prediction within a decade [is] a realistic goal" [Pre80]. Thirteen years later, Frank Press's optimism was more guarded:

> Some of the ideas we had in the 1970s didn't pan out ... I don't think anybody would claim that we have either a long-term or a short-term prediction capacity, except in probabilistic terms [Mos88, p. 354].

Today, opinions among seismologists about earthquake prediction are mixed. Some think that earthquakes may be impossible to predict:

> The available "positive" data on earthquake prediction are so sparse and often of such marginal character as to constitute one of the most telling and chilling aspects of our state of knowledge. One of the major propounded tenets of this paper is that the paucity of positive data ... is indicative either of failure to do the right things or of the impossibility of geophysicists obtaining the solution to the problem. For the moment, I remain an optimist and assume the former to be true [Eve82, p. S343].

Others think that given enough time and resources, success is inevitable. Muawia Barazangi, a geophysicist at Cornell University was quoted as saying "The answer will not be found in a year or two. It will require long-term commitment and investment in resources and funds. But it will happen eventually; there is no question in my mind." [Mos88, p. 355]. A third group claims to have already developed a working method to predict earthquakes [VL91], [Taz92], but many are skeptical of this claim [Bur85].

Part of the reason for the difference of opinion among seismologists may be semantic; the more restrictive a definition of what constitutes a prediction one uses, the less likely it is that prediction will ever be possible. For example, some seismologists take the word prediction to mean a more or less deterministic statement about the time, location, and magnitude of a future earthquake [Eve82]. In contrast, Agnew and Ellsworth define prediction as

> ... some statement about time of occurrence, ... the expected magnitude range and location. A sign of increasing maturity has been the realization that only rarely (perhaps never) are we justified in a bald assertion that an earthquake will occur (within some limits). Indeed, such a strong statement is unnecessary; more useful is a statement of the odds that we give of one occurring [AE91, p. 877]

One approach to understanding how faults rupture is the construction of idealized mechanical models that consist of sliders that are connected with springs. The behavior of such models has been shown to be similar to that of earthquake sequences on certain faults [Tur91]. Since the equations of motion of these systems exhibit chaotic behavior, some have concluded that earthquake prediction may be impossible. But as Julian states [Jul90, p. 482], "If earthquakes are truly chaotic, they may be easier to predict over short periods of time. On the other hand, because chaotic systems are sensitive to initial conditions in a way that increases exponentially with time, there would be an absolute limit to how far in the future earthquakes can be predicted ..." Other approaches to prediction include attempts to identify short-term precursors to major earthquakes [AE91] and attempts to model earthquake sequences as a renewal process, often with a lognormal inter-arrival time distribution [AE91].

At present, major research efforts are underway in Parkfield, California [Lin90] and in Japan [Mog86]. Earthquake prediction is commonly classified [AE91] as long-term (with a lead time of 30 years or more), intermediate-term (with a lead time of a few days to several years), or short-term (with a lead time of a few days or less). This distinction is based on physical considerations: short-term corresponds to times that are within four orders of magnitude of the rupture time of a fault (1 - 10 seconds) and it seems plausible that failure-related precursors might occur within this time. The next four orders of magnitude constitute intermediate-term prediction, and anything longer than that is long-term prediction, which is based on the theory of elastic rebound [AE91]. Agnew and Ellsworth comment as follows on long-term prediction in a 1991 review paper: "Socially, such lengths of time are also distinct from those much shorter, since in terms of reducing possible hazard anything much longer than a few decades is equivalent to forever [AE91, p. 877]."

As we said, whether earthquake prediction is possible or not is in part a question of semantics. For example, the U. S. Geological Survey in 1990 issued a statement [Wor90] that "the chance of one or more large earthquakes [magnitude 7 or greater] in the San Francisco Bay region in the coming 30 years is about 67 percent." This type of statement is perhaps not what most people instinctively think of as "earthquake prediction." But if the phrase "the coming 30 years" is replaced with "between 1983 and 1993" and the 67 percent probability is increased to 95 percent, then the statement begins to sound more like a prediction. Such a statement was issued [Lin90] by the U. S. Geological Survey in 1985, this time for a magnitude 6 earthquake in Parkfield, California. As of August 1994, the predicted Parkfield earthquake has yet to occur.

The area around Parkfield has been heavily instrumented to gather as much information as possible when the predicted earthquake occurs [Fin92]. This project, known as the

Parkfield Prediction Experiment has been criticized both for diverting funds ($19 million had been spent in 1992 [Fin92]) away from regions where quakes are likely to cause more damage and for having statistical flaws.

The statistical criticism centers on two issues. First is the issue of why the authors of the forecast used a two-sided 95% confidence interval that extended backwards in time beyond 1985, when the forecast was issued. Second, according to [Sav93], "the Parkfield prediction was based on an extrapolation of five of the six events in the 1857 to 1966 earthquake sequence; the 1934 event was omitted because it did not fit the regularity exhibited by the other data." According to [Fin92], "Although the authors of the 1985 study maintained that they had valid physical reasons for believing that the 1934 earthquake was uncharacteristic, [Mark Matthews, a statistician at the Massachusetts Institute of Technology] maintains that these justifications are 'simplistic and wishful' and 'not really scientifically defensible'."

Japanese earthquake prediction research is not free from controversy either. The Japanese program, which employs about 500 researchers and receives around US$56 million a year in government support in addition to salaries [Swi92], admits in a report issued in 1992 after the program had been assessed by outsiders for the first time, that "Society's expectations for earthquake prediction are high, but we cannot deny that there is a considerable gap between society's expectations and the present state of earthquake prediction" (quoted in [Swi92]).

However, in spite of these doubts about the current status of earthquake prediction, there is one way of knowing that damaging seismic waves are coming before they arrive: *real-time earthquake monitoring.*

2.2.1 Real-Time Earthquake Monitoring

According to a report on the "Technical and Economic Feasibility of an Earthquake Warning System in California," [HLR89]

> An [Earthquake Warning System (EWS)] is not an earthquake prediction system. Rather, it would provide users with a warning that an earthquake has begun. Depending on the distance of the user from the earthquake epicenter, the warning could be received some seconds or tens of seconds prior to the onset of strong shaking.

Among the findings of this report were

- For earthquakes of M7 and less, average warning times of 10 seconds or less could be provided in the significantly damaged areas (Modified Mercalli Intensity VIII or greater). For an earthquake of about M7.5 or greater, an average warning time of approximately 30 seconds could be provided in the significantly damaged areas.
- Based on the results from two surveys, potential users generally desire warning times of 30 seconds or greater. Thus, candidate earthquakes for an EWS should be M7.5 or greater. In southern California, earthquakes of M7.5 or greater would be limited, in all likelihood, to the southern San Andreas fault. The U.S. Geological Survey estimates that the annual probability of such an event is about 2 percent.
- An EWS is technically feasible and could be built for $3.3 million to $5.8 million in capital costs with annual operating costs of $1.6 million to $2.4 million, depending on the ultimate configuration of the system.

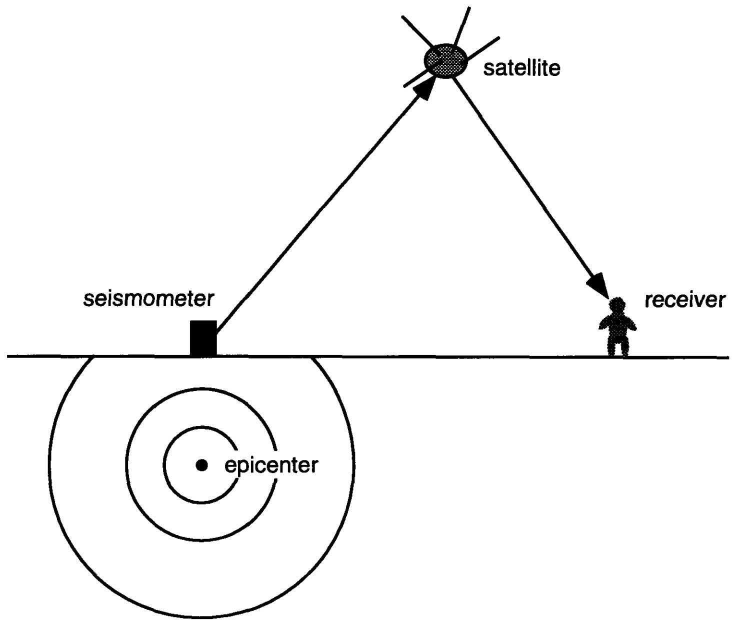

Figure 2-5: Real-time earthquake monitoring – a schematic view.

Figure 2-5 shows a schematic view of real-time earthquake monitoring where a satellite is used to relay the information that the occurrence of an earthquake has been detected from a network of seismometers to a central facility that may then send the information to people or organizations that might be affected by the earthquake.

There appears to be little doubt that this kind of earthquake monitoring is technically feasible. Indeed, such a system has been in operation [HLR89] in Japan since 1966 for the *shinkansen* or "bullet train." Initially, the system contained a number of seismic sensors distributed along the tracks; the sensors would cut off power to a particular section of the track when ground motion exceeded a specified value, thus providing warning (i.e., positive lead-time) to trains about to enter the section of track that was shut down. The system has now been upgraded to include sensors away from the track to provide longer lead-times. During the first 20 years of operation, the system would stop the train five times a year on average because of ground shaking and about once a year on average because of a false alarm. Since then, the triggering mechanism has been changed to reduce the number of train stoppages, since damage to the tracks has never been judged great enough to pose a danger to the train.

In Mexico City, a real-time earthquake warning system has been installed which broadcasts warnings to the public through commercial radio and television stations [DeP93]. This system can provide warning lead-times of around 50 seconds because of an unusual geological situation, where most potentially damaging earthquakes are centered about 200 miles away from the city (along Mexico's Pacific coast). Lead-times for a similar system in Los Angeles or San Francisco would be shorter for almost all residents. According to [DeP93], "The 50-second warning is thought to be time enough to turn off the gas and run out of the house."

However, the Mexican system is not currently in operation. According to Cinna Lomnitz, a geophysicist at the National Autonomous University of Mexico [Lom94a],

> The warning system was inaugurated on August 1 [1993] and suspended 3 months later because of malfunctioning. There were several alarms for minor (non-damaging) earthquakes and 2 major malfunctions: one of a large earthquake which did not set off the alarm and one of an alarm without any earthquake. The latter caused major inconvenience to the public as it occurred at 7:15 p.m. during the rush hour; a few people were hurt. The alarm system was officially interrupted hours later and has not been heard of since.
>
> ... I happen to believe that there is a potential for a good, pretty reliable warning system for Mexico City. Of course this fiasco makes is more difficult ...

The California feasibility study [HLR89] focused on use of a warning system by "public officials, schools, hospitals, police, fire stations, private industry, critical defense contractors and gas, oil and electrical industries" for such actions as shutdown of computer systems and production processes, switching of power or natural gas mains, removal of fire fighting equipment prior to strong shaking, and automatic disconnection of power to railroad lines. While most respondents to a survey of businesses and institutions indicated that times needed to complete such actions were greater than a real-time warning system could provide, the authors noted that "actions such as crawling under a desk or getting away from windows... require only a few seconds and could therefore be implemented with only a short warning." Even though a "public alerting system" was not one of the uses considered initially, the authors suggest it as one way of making use of the very short warning times that are possible. They comment that

> A public alerting system operates in the Tokyo/Kanto area. Loudspeakers are placed in schools and scattered throughout the cities and countryside. Public announcements notify listeners after an earthquake has occurred. The systems in Nakano, Chiba and downtown Tokyo are automatically activated to broadcast tape-recorded messages over loudspeakers. A pre-recorded message, informing the public that an earthquake is occurring or is about to occur and giving brief instructions on what actions to take, could be broadcast over an extensive public address system, after activation by a warning system. In addition, the Emergency Broadcast System (EBS), or other radio-frequency broadcasting systems, could be supplied with pre-recorded messages to be issued upon receipt of a warning signal.

The report estimates how much warning would be attainable with a warning system for several hypothetical future earthquakes. For example, if a magnitude 7 quake were to occur on the Newport-Inglewood fault, then approximately 80% of the area projected to experience Modified Mercalli intensities of VIII or more could receive warning from one to 15 seconds before the shaking begins. However, only 10 to 20% of that area would receive greater than 10 seconds of warning. Longer lead-times are possible for an earthquake that occurs on the San Andreas fault: San Fernando and Pasadena residents could know of the quake 25 and 35 seconds beforehand, respectively, and San Bernandino would receive 50 seconds of warning.

As for earthquake prediction in the sense of forecasting when and where a future earthquake will occur one might conclude, based on views expressed by seismologists, that such a capability is neither a future certainty nor a completely utopian concept. Given that many respected seismologists consider prediction to be a realistic future possibility, and given the magnitude of the social and economic impact such a capability might have, it seems that it would be appropriate to study how earthquake warnings could best be put to use, were they available.

2.3 A Sociological Perspective: Risk Communication Theory

We turn now to the question of how people might respond to earthquake warnings. Social scientists have studied human response to disasters in general and earthquakes in particular extensively. We will not attempt a comprehensive review of their findings (see [Dra86] for an inventory of findings); instead we will concentrate on *risk communication theory* and how it relates to our modeling approach.

Dennis Mileti (who has studied earthquake prediction from a sociological point of view) and Paul O'Brien describe the findings of risk communication theory as follows [MO93]:

> People who receive warnings of risk typically go through stages that shape their risk perception and behavior. This process is often modeled as a sequence; for example, hear-confirm-understand-believe-personalize-respond. The sequence may not be the same for every person, and each stage can be affected by characteristics of the people who hear warnings. These characteristics include age, gender, level of education, and other demographics, as well as the characteristics of the information, such as how frequently it is repeated, the source, and so on.
>
> The process begins when someone hears the risk information that is communicated. Second, people then typically attempt to confirm the warning, for example by checking with other people or seeking information from an alternative medium. Third, an understanding of risk is formed; individual meanings are attached to the information heard. The fourth stage is belief that the risk information received is accurate and that it is germane to the receiver. Usually, an individual must believe and personalize a warning in order to act. Fifth, people decide what to do and then perform that behavior. A person typically goes through the stages of the model each time that new warning or risk information is received. Response to communicated risk information thus follows from a series of perceptions.

Thus, risk communication theory postulates that most people go through the sequence of stages shown in figure 2-6.

Mileti and O'Brien continue by describing variables thought to impact how individuals process risk information, including (1) environmental cues, such as whether it is raining when a flood warning is received, (2) social attributes of the receiver, e.g., whether the person has a car to evacuate in and can afford to stay at a motel, and (3) psychological attributes, for example whether the person has experienced the hazard in question before. They conclude that

> The above research conclusions can be readily synthesized into a theory of public perception and response to risk information, and it can be summarized as follows. Public response to risk information is a direct consequence of perceived risk

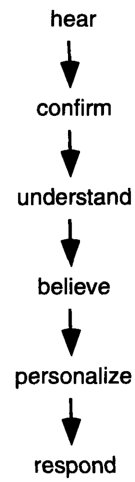

Figure 2-6: *Risk communication theory* postulates that people usually go through the sequence of stages shown in response to disaster warnings.

> (understanding, belief, and personalization), the warning information received (specificity, consistency, certainty, accuracy, clarity, channel, frequency, source, and so on), and personal characteristics of the warning recipient (demographics, knowledge, experience, resources, social network, cognitions, and so on); and perceived risk is a direct function of both the warning information received and the personal characteristics of the warning recipient.

This theory, and all of the sociological research we are aware of, seems to make the tacit assumption that the warning lead-time is long enough not to be a limiting factor. If the lead-time is only a couple of minutes or less, then most people will not have time to, say, confirm a warning by asking a neighbor or checking a newspaper. So the theory that most people go through the stages shown in figure 2-6 would predict that for very short lead-times, most people would not attempt to act, because they wouldn't have enough time to do so.

In this thesis, we have attempted a different modeling approach that pays particular attention to the limits imposed by the length of the warning lead-time and by how quickly people can perform such actions as crawling under a desk or exiting a building. To relate our approach to risk communication theory, it is helpful to consider figure 2-7, which is meant to depict most of the possibilities for how an individual might respond to a disaster warning. It may be useful to think of this diagram as listing the states of a discrete state stochastic process, with the arrows indicating possible transitions. Every individual is assumed to start in the ignorant state, not knowing that a disaster has been predicted to occur. Especially for short lead-times, it is possible that the earthquake occurs before the individual becomes aware of the warning, hence the arrow that goes directly from the ignorant state to the "earthquake occurs" state. If the individual does hear the warning in time, three things might happen: He might not have time to do anything before the earthquake occurs, he might attempt to confirm the warning, or he might try to act immediately. The arrows emanating from the remaining states have similar interpretations. Ultimately, every individual will end up in one of the two absorbing states: either the earthquake oc-

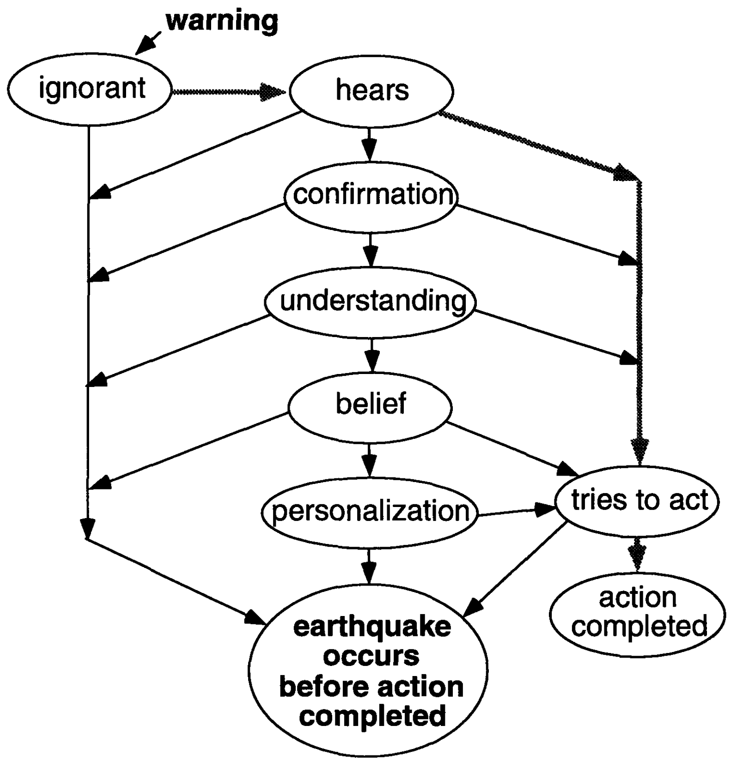

Figure 2-7: Possible responses of an individual to an earthquake warning.

curs before the individual is able to complete an action whose purpose is to increase the individual's chances of survival, or the individual is able to complete such an action.

One interpretation of risk communication theory is that it states that the most likely path through the diagram in figure 2-7 is the one down the middle, given enough time. For last-minute warnings, however, the path which is shaded gray is the most important. For very short lead-times, the people that are likely to benefit from the warning are the ones that act without delay, i.e., the ones that follow the gray path.

Our emphasis will be on estimating four probabilities associated with the diagram of figure 2-7: (1) the chances that a person hears the warning before the earthquake occurs (as a function of the warning lead-time), (2) the likelihood that a person who tries to act is able to complete the action before the shaking begins (also a function of the lead-time), (3) the probability that a person who completes an action dies in the earthquake, and (4) the probability that a person who is unable to complete an action before the earthquake occurs dies. The first two are probabilities for two of the three transitions along the gray path and the other two measure death risk for people in the two absorbing states. Note that we will not attempt to estimate the probability that an individual tries to act, given that he has heard the warning (the third probability along the gray path); we assume that this probability equals one. Because of this, our approach may be considered normative rather than descriptive. We discuss the dependence of our findings on this assumption in chapter 5.

Our review of the literature on earthquake prediction concludes by describing two frameworks that have been developed to assess the costs and benefits of earthquake prediction.

2.4 Cost-Benefit Frameworks for Earthquake Prediction

Two attempts at a comprehensive cost-benefit analysis of earthquake prediction are reported on in [PS79] and [SBHBD90]. It is notable that [PS79] was published eleven years before [SBHBD90], but nevertheless the modeling and analysis in [PS79] is in many respects more sophisticated. [PS79] is not cited in [SBHBD90].

The paper by Schulze, et al. [SBHBD90], is titled "Should We Try to Predict the Next Great U. S. Earthquake?". To answer this question, the authors define and estimate the following costs and benefits:

1. The value of lives saved by a successful prediction of a major earthquake, VLS.
2. The economic cost which can be attributed to a successful prediction, C_S. This excludes costs which can be attributed to the earthquake itself.
3. The economic cost of a false prediction, C_F.
4. The cost of a prediction program, $C(P_S, P_F)$. This value is obtained by discounting and summing the estimated costs of establishing and maintaining a prediction system throughout the time horizon considered. This cost is assumed to be a function of the two probabilities P_S (the probability that any given earthquake will be predicted) and P_F (the probability that any given prediction will be false). These probabilities are treated as the most important decision variables. However, a model of how $C(P_S, P_F)$ depends on P_S or P_F is not presented in the paper.

No costs are assumed to be incurred when an earthquake occurs but is not predicted, or when no earthquake occurs and none is predicted. The authors define the net benefits of earthquake prediction as

$$\int_0^T e^{-(r-n)t} \{P_1(VLS - C_S) - P_2 C_F\} \, dt - C(P_S, P_F)$$

where P_1 is the probability that an earthquake will be predicted and an earthquake occurs, P_2 is the probability of a false prediction, r is a discount rate, n is a rate of population growth, and T is the length of the time period to be considered. Presumably, the probabilities P_1 and P_2 are defined with respect to a time period of non-zero length (later in the paper the length of a time period is taken as one year), and thus a summation rather than an integral would perhaps be more appropriate, e.g.

$$\sum_{t=0}^{T} \left(\frac{1+n}{1+r}\right)^t \{P_1(VLS - C_S) - P_2 C_F\} - C(P_S, P_F)$$

Allegedly to allow the use of a time-varying discount rate, the authors describe in great detail a Monte Carlo simulation program used to evaluate the net benefits of earthquake prediction. It appears that the only stochastic element of this simulation is which of the four situations: (1) successful prediction, (2) false alarm, (3) an earthquake which is not predicted, or (4) no earthquake and no prediction, occurs in a given year of a 50 year period. But the probabilities of each situation have already been estimated by the authors. Therefore, all that is required to compute the net expected benefits is summing 50 terms of a series, no matter how the discount rate varies with time. So it is not clear what the purpose of developing the simulation was.

The probabilities of each of the four situations listed in the previous paragraph are completely determined by the values of P_S, P_F (which were defined previously), and $P_E(t)$, the probability that a major earthquake will occur in year t [3]. This probability is estimated by fitting a Weibull probability density to the time intervals between the last twelve major earthquakes on the southern San Andreas fault, via maximum likelihood.

In general terms, Schulze et al. have this to say about the potential benefits of the kind of earthquake warning they have in mind:

> The primary benefit of short term prediction is that public officials can take immediate action; in our case, it is assumed the population is required to remain in the relative safety of their homes for a 48-hour period. Although extensive emergency preparations probably cannot be made due to the short period between geologic indicators and an imminent event, many lives can be saved by keeping people away from dangerous structures ([SBHBD90, p. 249]).

Later, the authors apply their analytic framework to estimating the net benefits of earthquake prediction on the southern San Andreas fault. Their estimation of the life-saving potential VLS of such a prediction is as follows:

> The prediction used ... takes the following form: "A great earthquake is predicted to occur within the next 48 hours on the southern San Andreas fault near Los Angeles." This short-term prediction is aimed at saving lives, not at reducing property losses, and has the advantage of minimizing pre-event disruption. The appropriate response ... is not a panic evacuation; rather, residents are simply told to stay home, since most homes in the Los Angeles area are wood framed structures with a very low risk of death in a major earthquake. It is assumed that public officials (e.g., the National Guard) require the populace to remain inside residences for the 48-hour prediction period. ...
>
> ... The value of safety benefit is taken as the product of the estimated number of lives saved and the marginal value of safety, assumed to be one million dollars.
>
> Steinbrugge *et al.* [16] have estimated the expected number of deaths in Orange and Los Angeles Counties if an earthquake of magnitude 8.3 occurs on the southern San Andreas fault. Updated to reflect 1980 Census data, the following numbers of lives are lost for different times of occurrence: (1) 2:30 a.m., 3,080 lives; (2) 2:00 p. m., 11,906 lives; (3) 4:30 p.m., 13,007 lives. Given no prediction, it is assumed that on average the population spends 4 hours on freeways and entering/leaving work daily (the 4:30 p.m. risk), 8 hours at work (the 2:00 p.m. risk), and 12 hours at home (the 2:30 a. m. risk). Thus the time of occurrence weighted average number of deaths which would occur as the result of a large earthquake, given no warning, is approximately 7,800.
>
> If a prediction is made that a large event will occur within 48 hours and the population is required to return and remain home in response to the prediction, the 4:30 p.m. risk is assumed to apply for 1.5 hours during the initial response period, while the at-home (2:30 a. m.) risk applies for the remainder of the 48-hour period. The weighted average number of deaths that would occur, given a successful prediction warning is 3,400 in this case.

[3] Apparently, the authors do not consider the possibility that an earthquake *does* occur during the 50 year period under consideration. This event would change the probability $P_E(t)$.

> The difference between the two cases (7,800 - 3,400) implies that approximately 4,400 deaths would be saved ... by an earthquake warning. Assuming a marginal safety value of $1 million, the safety benefits ... are about $4.4 billion ([SBHBD90, pages 256 - 257]).

The economic costs of a successful prediction C_S and of a false alarm C_F are estimated to equal the loss of local output for one day (in the case of a successful prediction) and seven days (in the case of a false alarm). The authors cite [ENN86] to support their assertion that these cost figures are probably exaggerated. The scenario described in [ENN86] was that of a prediction by a U. S. scientist that a major earthquake was likely to occur in Peru at a specified time. However, this scenario does not appear comparable to the one envisioned in southern California by the authors, for three reasons: (1) The southern California prediction would come from the U. S. G. S. and would have been evaluated by the California Earthquake Prediction Evaluation Council [And85]. The Peru prediction was made by two independent scientists and was repudiated the National Earthquake Prediction Evaluation Council long before the first predicted event [All82]. (2) The Peru prediction was a long term one. The southern California prediction discussed in the paper has a lead time of about 24 hours. (3) The inhabitants of Lima, Peru, were *not* required or encouraged to stay home at the time when the earthquake was predicted to occur, in contrast to the situation described by the authors, where the population of southern California is required to stay at home, through the use of force if necessary.

The second paper, by Paté and Shah [PS79], is entitled "Public Policy Issues: Earthquake Prediction". It, and the companion paper [PS80] are based on Paté's doctoral thesis, which compares the costs and benefits of earthquake prediction on the one hand and earthquake engineering measures (e.g. stricter building codes) on the other.

The analytic framework used in [PS79] is similar to that used in [SBHBD90] in that the performance measure used is the net discounted benefits over a 50 year period. The main difference is that Paté keeps the economic benefits and the number of lives saved separate, and then computes a "cost per life saved", rather than explicitly assigning a dollar value to a life. The cost per life saved with earthquake prediction is then compared to the corresponding value for earthquake engineering, and various other policy options for improving public safety. Also, Paté's model includes the evolution of prediction capabilities, it accounts for earthquakes of different magnitudes, and it distinguishes between short-, medium-, and long-term predictions.

Paté's work is primarily a modeling effort. As she states in her thesis [PC78, page 267], "It is not in the scope of this study to emphasize the data research." In contrast, this thesis emphasizes the numerical estimation of the life-saving potential of earthquake prediction.

The human response to an earthquake prediction is discussed on page 1539 in [PS79]:

> Most critical to the results of the model is the set of measures (spontaneous or preplanned) that individuals, business, and governments are likely to take after an earthquake prediction of a given magnitude with a given lead time (in particular, high magnitude/long lead time, which appears to be the critical case in the long run). Such behaviors can be anticipated only through surveys and interviews of the respective groups ...

Later, on page 1546, the authors state:

> The third conclusion is that, because the benefit of a prediction to society is largely determined by the decisions of individuals ... it is most important that

> all information be available to the public and that the incentives from the State lead the people to take optimal measures for themselves.

A similar opinion is expressed in [SBHBD90], in discussing the use of expected utility theory in cost-benefit analysis, and the similarities between hurricane warnings and earthquake prediction:

> ... the question is not one of observed rationality so much as one of enforced rationality. That is, often police organizations and the national guard [in the case of hurricane warnings] enforce a warning through evacuation procedures or limits placed on travel.

These three quotes motivate the following observations:

- The utility of earthquake prediction is strongly dependent on how individuals in the area in question react to it.
- One has to be careful to distinguish between how individuals are likely to react to a prediction on the one hand, and what the "optimal" reaction would be on the other.
- An important issue is how, and to what extent, individuals can be encouraged to modify their behavior towards the "optimal" behavior.

The research that this thesis is based on differs from and complements the work reviewed in the last two sections in three important ways:

- We concentrate on last-minute warnings. Evaluating the potential benefits of such warnings requires a different modeling approach which pays particular attention to limitations imposed by the available time to act. At the same time, such warnings are of current interest because of the technical feasibility of real-time earthquake monitoring.
- We emphasize the development of defensible, numerical estimates of the life-saving potential of earthquake warnings. We hope that with the use of our estimates, cost-benefit assessments of earthquake warning systems can become more credible.
- Our approach is normative in that we attempt to evaluate what can be accomplished in the best case, i.e., assuming that individuals make the best possible use of the warning information. We believe that such an approach is a necessary prerequisite for deciding what actions people should be encouraged to take in response to earthquake warnings and how much effort should put into educating and training people to take advantage of such warnings.

  As pointed out in [SBHBD90], the concern might be less with whether people will react rationally if left alone, than it is with *enforced* rationality, e.g., telling people to stay off the freeways. But before anyone can enforce rationality, someone has to determine what the rational thing to do is.

Chapter 3

Gathering of Data

In this chapter, we discuss available sources of data on the death toll of earthquakes that occur without warning and on the fraction of citizens in different settings at the time of an earthquake (e.g., in cars). We then describe the design of a questionnaire whose purpose is to elicit expert judgment on the life-saving potential of knowing when the earthquake will occur, with lead times ranging from thirty seconds to thirty minutes.

A Death Rate Model

We propose to model the death rate $D(\tau)$ in a major earthquake, given that earthquake warnings are available and communicated to the public τ minutes before the earthquake as the sum of contributions from each of several subpopulations[1], i.e.,

$$D(\tau, t) = \sum_i F_i(t) D_i(\tau)$$

Here $F_i(t)$ is the fraction of the population that falls in the i-th setting when the time of day is t and $D_i(\tau)$ is the death rate in the i-th setting or subpopulation that can be expected with a τ minute lead-time. Note that the death rate in a particular setting is assumed not to depend on time of day; we will attempt to define the boundaries between the subpopulations so as to make this assumption tenable.

Suppose we let $R_i(\tau)$ be the *proportional* reduction in death risk that the τ minute warning achieves among individuals in setting i. Then $D(\tau)$ would be given by

$$D(\tau, t) = \sum_i F_i(t) D_i(0)(1 - R_i(\tau))$$

Three kinds of information are needed to replace the algebraic symbols in this relation with actual, defensible numbers:

Population distribution (the $F_i(t)$'s): We need information about how the population of a typical U.S. urban area is distributed across such groups as "people at home, awake," "people outdoors," or "children in school." Oddly, it appears that no statistical studies have been performed to address this question directly. Fortunately, various types of indirect evidence can be used instead to estimate the F_i's. The first section of this chapter defines and justifies the set of categories we have chosen and develops estimates

[1] This model will be discussed at greater length in chapter 4

of how people in a U.S. urban area are distributed across these categories at different times of the day. Besides being necessary for combining the death rate estimates $D_i(\tau)$, the population fractions $F_i(t)$ give an indication of which death rates $D_i(\tau)$ it is most important to estimate.

Baseline death rates (the $D_i(0)$'s): We need estimates of death rates that can be expected if a major earthquake were to strike a U.S. urban area *without warning.* We examine evidence from past earthquakes that have affected U.S. cities and review earthquake loss estimates that have been prepared for several U.S. urban areas. This is the subject of the second section.

Death rate reduction potential (the $R_i(\tau)$'s): We need estimates of how much the death rates could be reduced if earthquake warnings were available and were communicated to the public. In contrast to the first two kinds of information, no historical evidence is available to clarify this issue, since people have not had the benefit of warning before past earthquakes. In fact, one could argue that it is impossible to know exactly how people would respond to such warnings in our present state of knowledge. But rather than be paralyzed by this realization, it seems prudent make *informed guesses* of how effective a warning system might be, especially since the decision of whether a warning system is desirable has to be made before any experience with its use becomes available. To make the guesses as informed as possible we enlisted the help of people that have spent much of their professional lives thinking about and observing earthquakes and thus deserve to be called *experts* on the subject. The third section of this chapter describes the design and administration of a questionnaire that the experts were asked to respond to.

The three sections of this chapter correspond to these three categories of information.

3.1 Estimates of the Population Fractions $\{F_i(t)\}$

Below, we define the subpopulations that the experts were asked to consider. We concentrate on California for the reasons outlined in chapter 1: The experts are likely to be familiar with earthquake hazard in California, the level of awareness of earthquake hazard is comparatively high there, earthquake prediction research in the U.S. is concentrated there, and the effects of unanticipated earthquakes in California have been the object of considerable study (in particular, estimates for some of the baseline fatality rates $D_i(0)$ of interest to us are available).

1: People at home, awake: Consists of people that are awake, at home, in a *residential area in California.* For example, the person could be preparing a meal, eating, playing, reading, or watching television.

2: People at home, sleeping: Consists of people that are asleep, in a *typical wood-frame home.*

3: Office workers: Consists of people who are at work in office buildings.

People in cars: This category was divided into three subcategories.

4: Cars on or under a bridge or overpass: Consists of occupants of cars located *on a bridge or on or under a freeway overpass* in an *urban area in California.*

5: Cars on an open freeway, close to but not on or under an overpass: Consists of occupants of cars *on an open freeway,* possibly close to, but neither on nor under an overpass (remember that we are considering an urban area, so it is unlikely that the car would be far away from the closest overpass).

6: Cars on an urban street: Consists of occupants of cars *on urban streets* in California.

7: Children in school: Consists of people that are in school, specifically students in classrooms.

8: People outdoors, on urban streets: Consists of individuals walking outdoors on a street in an *urban area in California.*

Admittedly, the definitions above may not be sufficiently detailed to unambiguously classify any individual in any situation. It was necessary keep the number of categories small and the definitions relatively brief to avoid overburdening the experts, since the number of experts was limited and the amount of information they were asked to provide for each subpopulation was substantial.

Nevertheless, we hope that this set of categories covers a representative spectrum of situations. In particular, we hope that the average death risk in an earthquake is about equal for people that fall into one of the above groups and for those fall into none of them.

Recall that the model for the death rate $D(\tau, t)$ assumed that the death rate $D_i(\tau)$ for a particular setting i was independent of the time of day t when an earthquake strikes. Our choice of categories was guided by this assumption, i.e., we tried to define the subpopulations so as to make this assumption as reasonable as possible, while keeping the number of categories small. For example, a subpopulation defined as "People that are currently in their homes" (the union of the first two categories above) would not only vary in size throughout the day; one would also expect that the vulnerability of a randomly chosen member of this group would vary with time of day. By dividing the population of people in their homes into those that are sleeping and those that are awake, and more generally by categorizing a person based on *activity* rather than *location*, we hope to minimize dependence of the death rate in each subpopulation on the time of day when an earthquake strikes. As Bourque, Russell, and Goltz [BRG93] remark in their study of human behavior during and after the 1989 Loma Prieta earthquake,

> Where a person is at the time of an earthquake differs with the time of the year, day of the week, and time of day that the earthquake occurs. ... Although distributions of persons across locations will change, behavior within a given type of location may be constant across earthquakes.

We will return to this study in a moment, to examine where people in the San Francisco Bay area were at 5:04 PM on October 17, 1989, when the Loma Prieta earthquake disrupted a nationally televised World Series baseball game.

The rest of this section is devoted to reviewing various sources of data that are relevant to estimating the population fractions $F_i(t)$. Our final estimates, for eight equally spaced hours of the day, (presented in table 3.3) were adjusted by trial and error to be as closely in agreement with these sources of data as possible.

Location	Rest of Bay area	San Francisco/Oakland	Santa Cruz
Home (1)	49.9%	42.1%	62.4%
Work or school (3, 7)	25.5%	25.6%	12.0%
In transit (4, 5, 6)	15.7%	15.6%	11.1%
Public place or other location (8)	8.9%	16.7%	14.6%

Table 3.1: Distribution of a random sample of San Francisco Bay area residents (stratified by location) at the time of the Loma Prieta earthquake. Each of the four categories used in the study is followed by the numbers of the corresponding subpopulations under the scheme defined in this chapter. Adapted from [BRG93, table 4, page B10].

3.1.1 Where People Were During the 1989 Loma Prieta Earthquake

The quantities we wish to estimate are the fractions $F_i(t)$ of the population of a typical U.S. urban area that fall into each of the eight subpopulations listed at the beginning of this section as a function of time of day t. It would appear that statistics have not been collected specifically for this purpose in the U.S., with one exception that we are aware of and is particularly relevant. Borque, Russell, and Goltz [BRG93] surveyed a sample of 656 San Francisco Bay area residents an average of 224 days after the Loma Prieta earthquake. We will describe relevant portions of the results of that survey.

Table 3.1 (table 4 from [BRG93]) shows where the residents in the sample reported being when the shaking began. In parentheses, we show the subpopulations, according to our scheme, that correspond to the categories used in the Borque et al study. We see that around half the population appears to have been at home, around a quarter either at work or in school, 15% were "in transit," and the remaining 10% were in a public place or at other locations. The residents were also asked whether they were near a radio or television that was turned on at the time of the quake; the results are shown in table 3.2. By combining the results of the two tables, we can estimate the fraction of Bay area residents that would have heard an earthquake warning if one had been broadcast on all radio and television stations just before the quake:

$$\begin{aligned}
&\Pr\left\{\begin{matrix}\text{Randomly chosen person is near a ra-}\\ \text{dio or television that is turned on}\end{matrix}\right\}\\
&= \sum_i \left(\text{fraction in setting } i\right)\begin{pmatrix}\text{fraction of people in set-}\\ \text{ting } i \text{ with TV/radio on}\end{pmatrix}\\
&\approx 0.50 \times 0.55 + 0.25 \times 0.18 + 0.15 \times 0.47 + 0.10 \times 0.12\\
&= 40\%
\end{aligned}$$

This distribution of the sample of Bay area residents provides a useful snapshot of where people in a U.S. urban area are and what they are doing around 5 PM on a weekday. But to estimate how the distribution changes with time of day, we must consider additional evidence.

	At home	Work/School	In transit	Public place
Radio on	5.0%	11.5%	42.7%	6.0%
TV on	50.3%	7.0%	4.1%	5.8%
No media on	44.7%	81.6%	53.2%	88.3%

Table 3.2: Fraction of a random sample of San Francisco Bay area residents that had a radio or television on at the time of the Loma Prieta earthquake, stratified by where the respondent was when the quake occurred. Adapted from [BRG93, table 8, page B12].

3.1.2 Household Travel Survey Data

Glickman [Gli86] attempts to estimate time-of-day variations in the population distribution of an urban area. According to him, "The U.S. Census does not indicate how CED [census enumeration districts] vary throughout the day, and there is no other national database to turn to for that information." Instead, [Gli86] uses a 1968 survey of household travel patterns for Washington D.C. to estimate the population of various geographical areas in the Washington D.C. metropolitan area as a function of time-of-day. In figure 3-1 we present a simplified version of his findings in the hope that the overall patterns observed are roughly generalizable to other U.S. urban areas.

Figure 3-1 shows how the population of "a representative urban residential area" [Gli86] varies over the course of a weekday. The symbol a in figures 3-1 stands for the baseline population of the area, taken to be the 4 AM population. Figure 3-1 indicates that the population of a residential district decreases by about 40% between 6 AM and 11 AM (as people leave their homes in the morning), then stays relatively constant until 4 PM, and increases to its baseline value between 4 PM and 8 PM (as people return home at night). Note that the fact that the residential population is relatively constant between 11 AM and 4 PM does not mean that no traveling occurs in this time period; it only means that the flow of people to and from their homes is roughly balanced. The population of a residential area includes people in their homes, asleep or awake, and children in school. Thus, we would expect that $F_1(t) + F_2(t) + F_7(t)$ would follow a profile similar to the residential population profile in figure 3-1.

Also shown in figure 3-1 are time-of-day variations in the population of a commercial urban area. The population of this area grows by a factor of about eight between 6:30 AM and 10 AM (as people arrive for work, shopping, or other purposes), is stable between 10 AM and 3 PM, and then decreases to its baseline level between 3 PM and 7 PM. The population of a commercial area includes office workers and people outdoors, so we expect $F_3(t) + F_8(t)$ to follow the commercial population profile of figure 3-1. Since the baseline level a will include both people that live in the commercial area and people that come to work there, one would not necessarily expect $F_3(t) + F_8(t)$ to be exactly proportional to the commercial population profile. Rather, one would expect the fraction of people either in offices or outdoors to satisfy

$$F_3(t) + F_8(t) \approx \frac{N_c(t) - N_c'(t)}{N}$$

where $N_c(t)$ is the number of people that happen to be in commercial areas at time t, $N_c'(t)$ is the number of commercial area *residents* that are in commercial areas at time t, and N

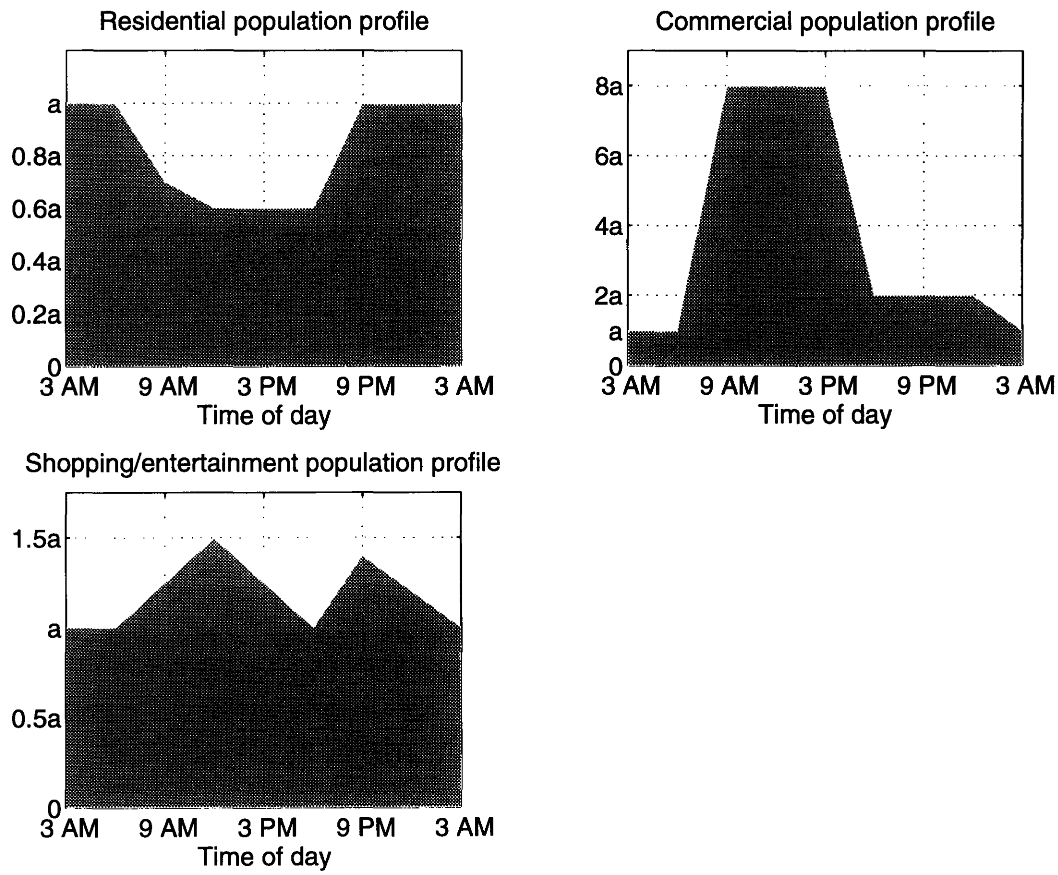

Figure 3-1: Population profiles for a representative residential, commercial, and shopping/entertainment area in a U.S. urban area. The symbol a represents the 4 AM population. Adapted from [Gli86].

is the total population of the urban area in question. At 4 AM, there are $N_c(4\text{ AM}) = a$ people in commercial areas, with a substantial fraction $N_c'(4\text{ AM})$ being residents. But around noon, the number of residents $N_c'(12\text{ PM})$ is probably much smaller than the total number $N_c(12\text{ PM})$, which includes office workers, shoppers, etc.

Finally, figure 3-1 shows time-of-day variations for a shopping/entertainment district. Here, we see two distinct peaks, one corresponding to people working and shopping during the day, and another corresponding to people who come to the area at night for entertainment. We expect that $F_8(t)$, the fraction of people outdoors, will follow a profile similar to the shopping/entertainment population profile in figure 3-1.

Glickman [Gli86] used microscopic data on travel patterns in a particular urban area. Aggregate data on travel patterns is available for the entire U.S. population and it can be used to estimate the time-average fraction of people $\bar{F}_4 + \bar{F}_5 + \bar{F}_6$ that are in their cars.

3.1.3 Aggregate Data on Commuting Patterns

In 1990, an estimated total of 2,316,455 million person miles of travel were undertaken by the U. S. population [Stu90]. Assuming that all trips were made at a constant speed v, the probability that a randomly chosen person is commuting (or traveling), averaged over time, can be estimated as

$$\Pr\left\{\begin{matrix}\text{A randomly chosen person is commuting}\\ \text{at a randomly chosen hour of the day}\end{matrix}\right\}$$
$$= \frac{2.3 \times 10^{12} \text{ person miles per year}}{v \text{ mph } \times 2.5 \times 10^{8} \text{ persons } \times 365 \times 24 \text{ hours per year}} = \frac{1.05}{v}$$

The median speed for "Urban Other Principal & Minor Arterial" roads was 54 mph for the year 1988 [DOT88], but people probably do a lot of their commuting at lower speeds, *especially* on the California freeways. Reasonable high and low estimates for the probability that a randomly chosen person is commuting at a randomly chosen hour might be 2.1% (for $v = 50$ mph) and 5.3% (for $v = 20$ mph).

To check the reasonableness of these estimates, suppose a person spends 30 minutes driving to work in the morning and 30 minutes driving back at night, every day. Then the person spends $1/24 = 4.2\%$ of his or her time driving.

Data on television viewing is readily available for the U.S. and we consider this data next.

3.1.4 Data on Television Viewing

The average number of hours per week spent watching television by people two years and older is 30 hours and 14 minutes, according to Nielsen ratings [Res93]. From this one can estimate the probability that a randomly chosen person is watching television to be

$$\Pr\left\{\begin{matrix}\text{A randomly chosen person is watching televi-}\\ \text{sion at a randomly chosen hour of the day}\end{matrix}\right\} = \frac{30 + 14/60}{7 \times 24} = 18\%$$

The Nielsen ratings also contain information about time-of-day variations in television viewing, as shown in figure 3-2. In estimating the fraction of the population at home, awake, during daytime hours, we used the profile shown in this figure as a lower bound. This maybe justified under two assumptions that seem reasonable. Let H and TV denote the events that a randomly chosen person is at home, or is watching television, respectively,

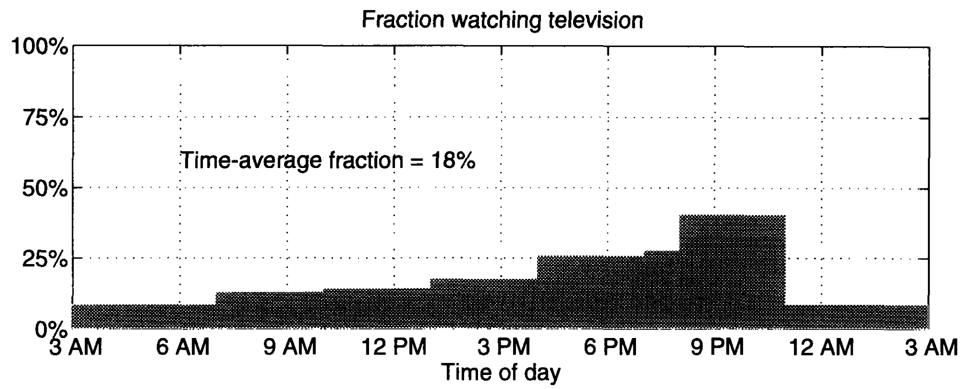

Figure 3-2: Fraction of the U.S. population that is watching television as a function of time of day. Adapted from [Res93].

during normal waking hours. Using the numbers in tables 3.1 and 3.2 the fraction of people watching television at 5 PM that are at home (i.e., $\Pr\{H|TV\}$) may be estimated (using Bayes theorem) to be about two thirds. If we assume that at least one out of three people that are at home is *not* watching television (i.e., if $\Pr\{\text{not TV}|H\} > 1/3$), then it follows[2] that the number of people at home, awake ($F_1(t)$) is at least as large as the number of people that are watching television.

Knowing what fraction of people are watching television at different hours of the day is also useful in its own right, since warning messages will presumably reach them very quickly (assuming that such warnings would be broadcast immediately on television).

3.1.5 Fraction of Cars on Freeways that are On, Under, or Near an Overpass

Estimates for the three subpopulations involving "people in cars" were fixed in a 1:2:2 ratio. Thus, if the total estimated fraction of the population in cars was 5%, then 1% would be assigned to "cars on or near overpasses," 2% would be assigned to "cars on the open freeway," and 2% would be assigned to "cars on urban street."

Maps of the San Fernando valley in California provide evidence that the 1:2:2 ratio may not be too far off. There are about 68.9 miles of freeways with a total of 110 overpasses in the San Fernando valley, for an average distance of 0.6 miles between overpasses. The overpasses range from bridges over minor streets that have no access to the freeway to intersections between freeways. Based on this average distance, and assuming that cars are uniformly distributed over the freeway network, one can estimate the fraction of cars on freeways in a typical U.S. urban area that are on or near an overpass as $d/(0.6 \text{ miles})$,

[2]Under the assumptions in the text,

$$\begin{aligned}\Pr\{H\} &= \Pr\{H, TV\} + \Pr\{H, \text{not } TV\} \\ &= \Pr\{H|TV\}\Pr\{TV\} + \Pr\{\text{not TV}|H\}\Pr\{H\} \\ &> (2/3)\Pr\{TV\} + (1/3)\Pr\{H\}\end{aligned}$$

which implies $\Pr\{H\} > \Pr\{TV\}$.

where d is an appropriately chosen distance. The distance d should include both the average length of an overpass and a threshold value below which a car would be considered "near" an overpass. The length of an overpass ranges from perhaps 15 m (approximately double the width of a two lane road) to maybe 150 m for some bridges at freeway junctions. Suppose the average length of an overpass is 50 m. For drivers approaching an overpass, we will take "near" an overpass to mean that the driver does not have time to stop before reaching the overpass. Suppose the average driver's speed is 20 m/s (about 45 mph) and that when braking, his deceleration is a constant 2 m/s^2 (20% of g). Then the driver would cover 100 m before coming to a stop, and the estimated fraction of cars on or near overpasses would be $(50 + 100)\text{m}/(0.6 \times 1609)\text{m} \approx 15\%$. Since overpasses are probably the most dangerous locations for a freeway traveler to be at during an earthquake, our assumption that a third of freeway travelers are near overpasses is likely to overestimate the number of deaths on freeways, if anything.

3.1.6 Various Educated Guesses

To complete the estimation of how the $F_i(t)$'s vary, the following educated guesses were used:

Sleeping: Surprisingly, we were not able to locate statistics on how long people in the U.S. sleep, on average. But it seems reasonable to assume an average of eight hours per night.

Working or at school: We assumed that on average, people spend eight hours a day, five days a week working or studying outside their homes. Thus, we estimate that $\bar{F}_3 + \bar{F}_7 \approx \frac{5\times 8}{7\times 24} = 23.8\%$.

Time spent on various activities at home: To fix the estimated average time spent at home, awake, we complemented the data reviewed already with the following guesses:

Meals, at home:	Assume one hour per day, or about 4%
Personal hygiene:	Assume 45 minutes per day, or about 3%
Work at home:	Assume 4 hours per day for half the population, 2 hours for the other half, or about 12% on average
Recreation:	(Reading, sports, hobbies, etc) Assume 2%

Table 3.3 and figure 3-3 show our estimates of how the population fractions $F_i(t)$ evolve over the course of a day. The estimates are based on the direct and indirect evidence and educated guesses we have presented. In chapter 5, these estimates will be combined with the judgment of experts to estimate the life-saving potential of earthquake warnings.

While data that bear directly on how much time members of the U.S. population spend on various activities is not available to our knowledge, such data is available for some other industrialized societies. We conclude this section by comparing such data from the Central Statistical Office of Finland with the estimates in table 3.3. The Finns are asleep 35.5% of the time, on average, compared to our estimate of 31.9% for the U.S. population. Eating, hygiene, and work in the home occupy 21.5% of the Finns' time, while we estimate that the U.S. population spends 33.9% of its time at home. But "free time" accounts for 26.5% of the time for the Finnish population and about half of that time is probably spent at home,

Subpopulation	6 AM	9 AM	12 PM	3 PM	6 PM	9 PM	12 AM	3 AM
People at home, awake	31.5%	25.0%	21.0%	20.0%	56.0%	77.5%	35.0%	5.0%
People at home, asleep	61.0%	21.0%	2.0%	2.0%	7.0%	10.0%	59.5%	92.0%
Office workers	4.0%	25.0%	30.0%	35.5%	15.0%	5.0%	2.0%	1.0%
Cars near bridge or overpass	0.5%	2.0%	1.0%	1.5%	2.0%	0.7%	0.5%	0.3%
Cars on open freeway	1.0%	4.0%	2.0%	3.0%	4.0%	1.4%	1.0%	0.6%
Cars on urban street	1.0%	4.0%	2.0%	3.0%	4.0%	1.4%	1.0%	0.6%
Children in school	0.0%	17.0%	35.0%	30.0%	5.0%	1.0%	0.0%	0.0%
People outdoors	1.0%	2.0%	7.0%	5.0%	7.0%	3.0%	1.0%	0.5%
Total:	100%	100%	100%	100%	100%	100%	100%	100%

Subpopulation	Time average ($\bar{F}_i$)
People at home, awake	33.9%
People at home, asleep	31.9%
Office workers	14.7%
Cars on bridge/overpass	1.1%
Cars on open freeway	2.1%
Cars on urban street	2.1%
Children in school	11.0%
People outdoors	3.3%

Table 3.3: Estimates of how the population of a U.S. urban area is distributed across the subpopulations defined at the beginning of this chapter. The estimated population fractions $F_i(t)$ are given at three hour intervals. The estimates are based on direct and indirect evidence and educated guesses presented in this chapter. Also shown are time average values over the 24 hours of the day. These estimates will be used to combine death rate estimates $D_i(\tau)$ for various warning lead-times into overall death rate estimates.

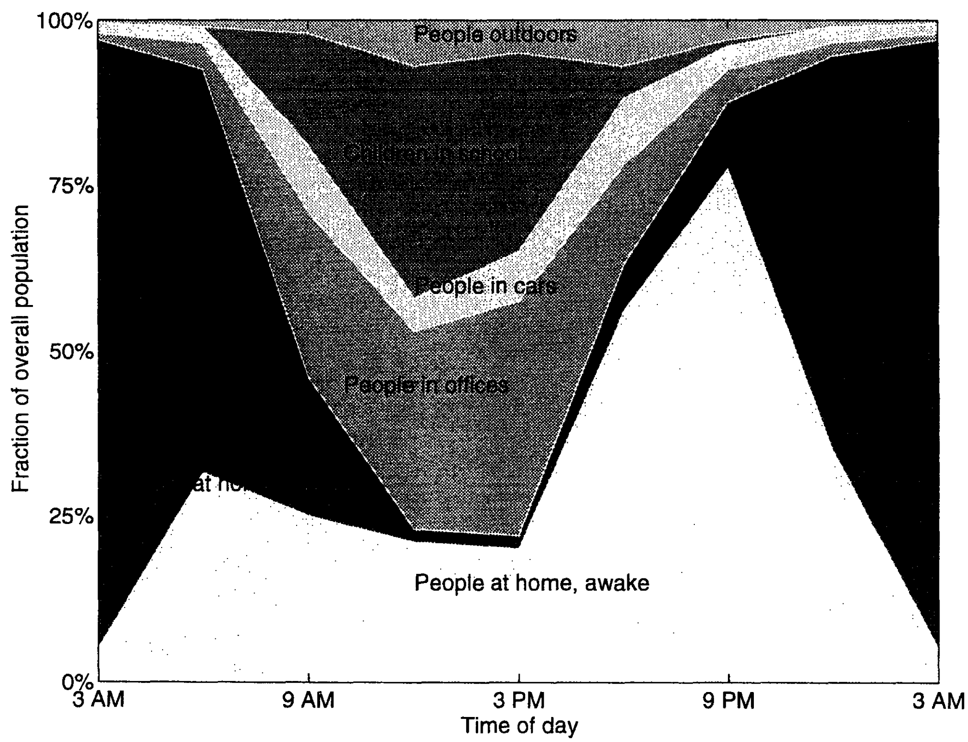

Figure 3-3: Estimates of how the population fractions $F_i(t)$ evolve with time of day.

so our estimate of the time spent "at home, awake" seems to be in approximate agreement with the Finnish data. Finally, the Finns are working or studying 16.5% of the time, while we estimate that U.S. residents spend 25.7% of their time in the office or at school. The difference may be partly justified, since the average work week is longer in the U.S. than it is in many European countries.

3.2 Earthquake Loss Studies: Estimates of the Baseline Death Rates $D_i(0)$

Numerous estimates of the likely impacts of future earthquakes on urban areas throughout the U.S. have been made during the last two decades. Such *earthquake loss studies* have been commissioned by several federal agencies: the National Oceanic and Atmospheric Administration (NOAA), the U.S. Geological Survey (USGS), and the Federal Emergency Management Agency (FEMA). According to a panel report on loss estimation methodology [Pan89],

> An earthquake loss estimate is a forecast of the effects of a hypothetical earthquake. Depending on its purpose, a loss study may include estimates of deaths and injuries; property losses; loss of function in industries, lifelines, and emergency facilities; homelessness; and economic impacts.

In this section, we will concentrate on estimates of death tolls. Various methodologies have been used to estimate losses from future earthquakes, but two common threads emerge in the estimation of life loss:

1. Extrapolations from available historical data, primarily from California earthquakes. In effect, death rates from a past earthquake were scaled up or down to reflect differences between the conditions in the area where the historical earthquake occurred and the area for which loss estimates are desired. The scaling factors were to a large extent based on the judgment of the team of experts assembled to prepare the study. This approach was used in loss studies for the metropolitan areas of San Francisco [ARD+72], Los Angeles [AHC+73], Puget Sound [HLS+75], and Salt Lake City [RAh+76], sometimes referred to collectively as the USGS/NOAA studies.

2. Predicting life loss using detailed predictions of building damage. Under this approach, damage to buildings of various types is classified into a small set of categories and the fraction of the building stock in the area of interest that will fall into each category is forecast. A set of casualty rates that are conditional on each of the damage categories is then used to predict the death toll. The casualty rates that have been most commonly used are based on expert judgment and were developed as part of a study by the Applied Technology Council (ATC), commonly known as the ATC-13 [RS85] project.

3.2.1 The USGS/NOAA Studies

We now describe the death rate estimates used in the USGS/NOAA studies. The first of these studies ([ARD+72]) reviews historical evidence from several U.S. earthquakes. Table 3.4, which is based on this review, shows estimated death rates. It is not completely clear from either the study itself or two reports on "Methodology" and "Methodology Improvements" ([Ste74]) that followed it exactly how these death rates were arrived at. In

Date	Location	Time	Death rate (per 100,000)	Magnitude (Richter)	Intensity (MMI)
1872, 3/26	Owens Valley, CA (Lone Pine)		8,000		
1886, 8/31	Charleston, SC	9:51 PM	45 (outright) 113 (total)		X
1906, 4/18	San Francisco, CA	5:12 AM		8.3	
	(San Francisco)		124		XI
	(Santa Rosa)		116		
	(San Jose)		80		VIII
1925, 6/29	Santa Barbara, CA	6:42 AM	45	6.3	VIII-IX
1933, 3/10	Long Beach, CA	5:54 PM	26	6.3	IX
1940, 5/18	Imperial Valley, CA	8:37 PM	18	7.1	X
1949, 4/13	Puget Sound, WA	11:56 PM	1	7.1	VIII
1952, 7/21	Kern County, CA (Tehachapi)	4:52 AM	500	7.7	XI
1952, 8/22	Bakersfield, CA	3:41 PM	3	5.8	VIII
1964, 3/27	Alaska (Anchorage)	5:36 PM	9	8.4	XI
1965, 4/29	Seattle-Tacoma, WA	7:29 AM	1.5	6.5	VIII
1971, 2/9	San Fernando, CA	6:01 AM		6.4	XI
	(excl. VA Hospital)		12		XI
	(incl. VA Hospital)		64		XI

Table 3.4: Death rates in selected earthquakes in the U.S. Based on [ARD+72]. As discussed in the text, it is not clear how the boundaries of the population to be considered in each earthquake were determined.

particular, it is not clear how the region whose population would be used as a denominator in the death rate calculation was chosen.

For example, 58 people died in the 1971 San Fernando earthquake, 47 of them in the collapse of the San Fernando Veterans Administration (VA) hospital. The death rates per 100,000 people given in table 3.4 for the San Fernando quake are 64, if the fatalities in the VA hospital collapse are included, and 12, if those fatalities are excluded. To obtain these death rates, one has to divide the actual death toll by approximately 90,000. The only information provided in [ARD+72] about where this number comes from is that it corresponds to "the heaviest shaken area of the 1971 San Fernando shock."

The USGS/NOAA studies estimate overall death rates by dividing the population according to the size of the dwelling they live in (see [ARD+72] and [Ste74]):

- Inhabitants of *single-unit dwellings*, which are assumed to be predominantly wood-frame houses, are assumed to face the 1971 San Fernando rate of *12 deaths per 100,000*, for those regions where the Modified Mercalli Intensity (MMI) is IX or greater. This rate is scaled down to 2, 4, and 6 per 100,000 for MMI VI, VII, and VIII, respectively.

Note that this is the death rate that excludes deaths caused by the VA hospital collapse.

- For inhabitants of *2 to 9 unit dwellings*, it is assumed that 50% are at 50% higher risk than the 12 per 100,000 San Fernando rate and the other 50% face a risk comparable to what was experienced in the 1933 Long Beach earthquake, i.e., a death rate of 26 per 100,000. These assumptions lead to an average rate of *22 deaths per 100,000* which is assumed to prevail where intensities are MMI IX or greater; rates of 4, 9, and 15 per 100,000 are assumed for MMI VI, VII, and VIII.

- Half the population of *dwellings with 10 or more units* is assumed to face a death rate of 64 per 100,000, which is the death rate for the San Fernando shock when VA hospital fatalities are included. The other half is assumed to live in old brick buildings and the death rate for such structures is assumed to be 500 per 100,000. This figure appears to have been obtained by adjusting the 1964 Alaska earthquake death rate for downtown Anchorage to what might have been if "the hour had been earlier and had the buildings under construction been completed." The average of these two rates is *282 deaths per 100,000*. The rate is reduced to 210 for areas 25 to 50 miles away from the epicenter and to 140 for areas more than 50 miles away.

These death rates form the basis for an estimate of the overall death toll at 2:30 AM. Estimates are also given for 2 PM and 4:30 PM. At 2 PM, 40% of the population is assumed to face the average 2:30 AM risk. Then, the population in "congested areas" is estimated; the death rate in such areas is assumed 500 per 100,000 for MMI VIII-IX, 400 per 100,000 for MMI VIII, 300 per 100,000 for MMI VII-VIII, and 200 per 100,000 for MMI VII and VI-VII. In addition, the number of buildings that pose a hazard from parapets, ornamentation, glass, and so on falling off the outside is estimated and a rate of 1 death per building is assumed. The balance of the population is assumed to face a death rate that is 25% higher than the 2:30 AM rate. The 4:30 PM death toll was assumed the same as at 2 PM, except for deaths due to collapses of hospital, schools, and freeways, which were estimated separately and added to the total. Also, the number of deaths due to falling debris was assumed higher at 4:30 PM.

Detailed data about hospitals and schools in the San Francisco Bay Area were used to estimate the number of fatalities in the event of a magnitude 8.3 earthquake on the San Andreas fault occurring at 2:30 AM, 2 PM, or 4:30 PM. It was then assumed [Ste74] that the ratio between deaths in hospitals, schools, and on freeways and deaths elsewhere would be the same for San Francisco and Los Angeles. It is not clear how estimates of the number of freeway fatalities were computed.

Table 3.5 shows estimated overall fatality rates for hypothetical magnitude 8.3 earthquakes on the San Andreas fault in the vicinity of San Francisco and Los Angeles, respectively, based on a 1980 update [SAL+80] of the loss studies for those areas originally published in 1972 and 1973. Also shown are estimates for a magnitude 7.5 quake on the Newport-Inglewood fault, which lies under downtown Los Angeles. The estimates range from 33 to 229 deaths per 100,000 with the day-time estimates generally three to four times larger than the night-time estimate. For comparison, estimated death rates for a magnitude 7.5 earthquake in the Puget Sound Area [HLS+75] ranged from 5 to 106 per 100,000 and magnitude 7.5 shock in the Salt Lake City Area [RAh+76] was estimated to cause fatality rates from 33 to 361 per 100,00, depending on the time of day.

It is a bit difficult to compare the estimation method used in the USGS/NOAA studies

Magnitude 8.3 quake on the San Andreas fault in the San Francisco Bay Area

	Deaths:			
Time of Day	Schools	Hospitals	Elsewhere	Death Rate
2:30 AM	0	600	2,500	61 per 100,000
2 PM	2,100	1,450	6,640	201 per 100,000
4:30 PM	200	1,450	9,720	224 per 100,000

Magnitude 8.3 quake on the San Andreas fault in the Los Angeles Area

Time of Day	Total Deaths	Death Rate
2:30 AM	2,965	33 per 100,000
2 PM	10,125	112 per 100,000
4:30 PM	11,045	122 per 100,000

Magnitude 7.5 quake on the Newport-Inglewood fault in the Los Angeles Area

Time of Day	Total Deaths	Death Rate
2:30 AM	4,380	48 per 100,000
2 PM	18,860	208 per 100,000
4:30 PM	20,755	229 per 100,000

Table 3.5: Estimated overall death rates for a magnitude 8.3 earthquake on the San Andreas fault in either the San Francisco Bay Area (top) and near Los Angeles (middle), or for a magnitude 7.5 shock on the Newport-Inglewood fault (bottom). Based on [SAL$^+$80].

Light steel and wood-frame construction	
Damage State	Death Rate
None	0
Slight	0.01 per 100,000
Light	0.1 per 100,000
Moderate	1 per 100,000
Heavy	10 per 100,000
Major	100 per 100,000
Destroyed	2,000 per 100,000

All other construction	
Damage State	Death Rate
None	0
Slight	0.1 per 100,000
Light	1 per 100,000
Moderate	10 per 100,000
Heavy	100 per 100,000
Major	1,000 per 100,000
Destroyed	20,000 per 100,000

Table 3.6: Death rate estimates for various construction and damage categories. Based on [RS85].

to the one adopted in this thesis because, as has been pointed out, some aspects of the USGS/NOAA method are not clearly defined in the USGS/NOAA studies or documents that accompany them ([Ste74]). However, the following remarks seem in order:

- The USGS/NOAA studies divide the "at home" population into three categories according to building size. The method described in this thesis divides this population into two categories: people that are awake and people that are asleep. This reflects the different purposes of the two methods: whether a person is awake probably has little effect on chances of survival in an earthquake that occurs without warning, but being awake is likely to improve one's chances of taking advantage of an earthquake warning.

- The USGS/NOAA studies assume that 40% of the population is at home at 2 PM and at 4:30 PM and 100% are at home at 2:30 AM. Our estimates of the fraction of people at home are 22% at 3 PM, 63% at 6 PM, and 97% at 3 AM.

- It would appear that the USGS/NOAA method double-counts people that are on freeway structures, in schools or hospitals, or at risk from debris falling off buildings. This is likely to exaggerate differences in death rates between day and night.

3.2.2 The ATC-13 Report

Twelve years after the first USGS/NOAA study was published (the San Francisco Bay Area [ARD+72] study) the Applied Technology Council published a report titled "Earthquake Damage Evaluation Data for California" [RS85], commonly known as the ATC-13 report. The life-loss estimation approach recommended in this report consists of estimating the number of people in buildings in each of several "damage states" and multiplying with a set of death rates given in table 3.6. These death rates are based on the same data as the NOAA/USGS studies (i.e., table 3.4), combined with considerable judgment elicited from a group of experts via the Delphi method. A 1989 National Research Council panel report on earthquake loss estimation [Pan89] states that "This does represent a rational approach to estimating casualties, and the panel recommends use of this method combined with careful judgment and comparison with historical data, where comparable cases pertain."

Among the loss studies that have used this approach is a FEMA sponsored study for St. Louis City and County, Missouri [Cen90]. This study estimates the death toll in the St. Louis Area in the event of a magnitude 8.6 earthquake to be 1,200 at 2 AM and 1,400 at 2 PM (the corresponding death rates are 78 and 91 per 100,000). The day-night difference in estimated death tolls is much smaller here than in the USGS/NOAA studies. In part, this may be because the double-counting alluded to earlier does not occur when the ATC-13 method is used[3].

Some of the advantages of the ATC-13 method seem to be that it is well-documented, it is compatible with the methods used to estimate damage to buildings (also documented in the ATC-13 report), and the apparent double-counting of the USGS/NOAA studies does not occur. The price one has to pay for using the ATC-13 method is that it does not explicitly account for deaths that occur on freeways or downtown streets.

We will refer back to the death rate estimates that have been quoted in this section in chapter 5, where the results of our survey of experts will be presented, and in chapter 6, which discusses the implications of three recent earthquakes.

3.3 The Design and Administration of the Questionnaire

A substantial effort was devoted to getting information from experts that, sensibly processed, would allow estimation of how much death rates could be reduced with earthquake warnings, i.e., estimation of the risk reduction factors $R_i(\tau)$. This section describes the development of a questionnaire that was used for this purpose. For each of the subpopulations listed at the beginning of this chapter and for each of four lead-times (30 seconds, 1 minute, 5 minutes, and 30 minutes) the questionnaire inquires about several crucial issues:

- Strategies for communicating an earthquake warning.
- The likely degree of success of such strategies.
- The most beneficial action(s) to attempt given τ minutes until the quake starts.
- The chances of successfully completing the action(s) in the available time.
- The probability of perishing given the action is successfully completed.
- The probability of perishing given no warning.

The questionnaire went through two iterations of testing, where both graduate students in operations research and persons who would have been considered qualified to respond to the final version of the questionnaire, i.e., experts, were asked to respond to the questions and comment on clarity, relevance, and whether it was reasonable to expect busy professionals voluntarily to make the effort of completing the questionnaire.

[3]Another contributing factor is that a larger fraction of St. Louis residents may live in unreinforced masonry buildings than is typically the case for the Western urban areas studied by the USGS and NOAA. The night-time risk for residents of unreinforced masonry buildings is probably not appreciably lower than the day-time risk, if at all.

3.3.1 The List of Experts

A list of potential respondents to the questionnaire was compiled from several sources. Participants in three workshops held in the U.S. in 1992 on earthquake hazard mitigation were put on the list. So were individuals whose names were suggested by Tom Tobin, Director of the California Seismic Commission, and by scholars working on earthquake-related research in the Boston area.

People on this list were contacted by telephone and asked whether they would be willing to receive the questionnaire. The overwhelming majority who indicated that they would (all except one) then received the questionnaire in the mail, along with a stamped return envelope and a postcard that could be used to indicate whether the respondent, after glancing through the questionnaire, thought that she or he would take the time to complete it. In some cases, people were asked to suggest names of additional persons that it would be appropriate to send the questionnaire to. The final list contained 62 names of experts in a variety of fields, including earthquake engineering, seismology, sociology, architecture, emergency management, urban planning, disaster medicine, and risk analysis.

As we shall see, the questions that the experts were asked require two kinds of knowledge to answer intelligently. First, there is the kind of knowledge that distinguishes experts on earthquakes and their effects from other intelligent people: knowledge of the mechanisms that cause earthquakes and knowledge about how structures and people respond to earthquakes. Second, even though we made every effort to not have the experts perform detailed calculations, a basic familiarity with conditional probability would have facilitated the thought processes required to answer the questions. For this reason, we added a group of 11 distinguished operations researchers to our list, most of whom have no special knowledge about earthquakes, but can be expected to be intimately familiar with the use of conditional probabilities. Our hope was that the operations researchers would be more likely to avoid subtle biases in the assessment of probabilities and thus their strengths would complement those of the earthquake experts.

Response rates: Of the 73 experts on our list (62 earthquake experts and 11 operations researchers), five could not be reached and one declined to receive the questionnaire. The remaining 67 experts received the questionnaire in the mail, with 45 of them eventually completing and returning it to us, for an overall response rate of 67%. In most cases, respondents that did not return the questionnaire within a month or two were contacted by us. Some of the experts were also contacted after we received their questionnaires to clarify certain responses, as we will explain later in this section. Table 3.7 shows response rate by field of expertise.

3.3.2 The Design of the Questionnaire

In designing the questionnaire, we strove to maintain a sensible balance between structured questioning to allow comparability of responses and open-endedness to allow the experts freely to raise possibilities that had not occurred to us. At the beginning of the questionnaire, the experts were told that

> The data from this questionnaire, along with data from other sources, will be used to obtain estimates of how many lives might be saved if short-term earthquake prediction were possible. In responding to the questions, we ask you to make the following **HIGHLY IDEALIZED** assumptions:

Field	Response rate
Earthquake engineering	7/8
Seismology	7/14
Sociology	2/3
Architecture	8/11
Emergency Management	3/4
Urban Planning	1/2
Earthquake engineering and architecture	2/2
Earthquake engineering and seismology	1/1
Operations research	10/11
Other fields	5/9
Unknown	0/2
Total:	45/67

Table 3.7: Response rates for the questionnaire, by field of expertise.

1. A capability to reliably predict earthquakes in the short term has been developed. The earthquake forecasts are completely reliable, meaning that they accurately specify the exact time τ which will elapse before the imminent earthquake occurs.
2. All people who learn of the earthquake forecast believe it and know what the lead time τ is.
3. In any given situation, all people know what to do in order to minimize their risk of death.

Furthermore, please assume for all sections of this questionnaire that

4. The type and quality of construction and the level of public awareness of earthquake hazards is comparable to what it is in California.

At the beginning of each section of the questionnaire, the experts were reminded to assume that the earthquake forecasts were completely reliable. They were also asked to assume that the local intensity of the imminent earthquake was VIII or more on the Modified Mercalli scale in the area in question. A definition of intensity level VIII from [Ric58] was provided in a footnote. Each section inquired about one subpopulation and generally had three subsections, titled "Communicating a Prediction," "Responding to a Prediction," and "Estimates of the Risk of Death." The section about "People in Cars" was organized slightly differently, and we will now reproduce parts of it for illustration:

People in Cars

This section of the questionnaire is about a typical occupant of a car, on a freeway or street in an *urban area in California,* when the earthquake warning is issued. ... The first subsection is about communicating an earthquake prediction to people in cars. The remaining subsections deal with three distinct cases. There are

two subsections for the case of a car which is located *on a bridge or on or under a freeway overpass,*

two subsections for the case of a car *on an open freeway,* and

two subsections for the case of a car *on an urban street.*

Communicating a Prediction

Consider different ways of communicating an earthquake prediction to a person in a car:

A: Install or use existing sirens which are sounded when an earthquake is imminent.

B: Install a system of red lights on freeways and blue lights on streets, to warn drivers of an imminent earthquake. The lights would be located a few hundred feet apart. They would start flashing as soon as an earthquake prediction had been made.

C: Equip every car with a device which sounds the horn of the car for some period, to warn the driver of an imminent earthquake. The devices would be triggered remotely, from an earthquake prediction center.

Assume all strategies are technologically feasible. Again, these are not mutually exclusive strategies. Also, assume that some indication of the length of the lead-time is given. For example, the red and blue lights could be complemented with digital signs which display the length of the lead time.

Can you think of of other effective ways of warning people in cars that an earthquake is about to occur?

D:

E:

In the table below, you are asked to estimate *the fraction of people in cars that would learn of the prediction,* given four different values for the lead-time before the earthquake will occur. You are asked to do this for each strategy separately, *assuming that no other means of communication are used.* In the last row of the table, you are asked to indicate which *pair* of strategies you think would be most effective, and to estimate the fraction of people in cars that would receive the warning, for different lead times, assuming that only those two strategies were used.

An example of how this table might be filled out, for each of the strategies, appears on the next page.

Fraction of people in cars that learn of prediction					
Strategy used \ Lead-time	30 sec.	1 min.	5 min.	30 min.	
A: Sirens	%	%	%	%	? ??
B: Lights on freeways	%	%	%	%	? ??
C: Sound horns	%	%	%	%	? ??
D:	%	%	%	%	? ??
E:	%	%	%	%	? ??
Most effective pair of strategies (please circle two): A B C D E	%	%	%	%	? ??

The single and double question marks next to the table provide participants who lack strong confidence in their numerical estimates with an easy way of signaling their queasiness. Most respondents were willing to hazard such estimates, but some of them circled either the single or double question marks, which signalled "limited" and "very little" confidence in the estimates, respectively, according to instructions given in the questionnaire. Every question had single and double question marks next to it. About 28% of the responses had the single question mark next to it circled, 21% had the double question mark circled, and 50% had neither circled, indicating that the respondent had no particular qualms about the estimate provided in answer to the question. The remaining 1% of the responses had triple question marks (invented by the expert) next to them, presumably to indicate an exceptionally low level of confidence. However, we did not weigh the expert estimates differently based on the level of doubt expressed by the experts. We expect that such signals may have as much to do with a person's cautiousness and awareness of how little is known about the issues we ask about as they have to do with the reliability of an expert's answers, in comparison with other experts.

In later chapters, we will use the following notation for the numbers provided by the experts in the above table:

$$A_{ij\tau} = \begin{array}{l} j\text{th expert's estimate of the fraction of people in subpopula-} \\ \text{tion } i \text{ that learn of a warning with a } \tau \text{ minute lead-time} \end{array}$$

For example, a number written in the top left cell of the table by expert number 4 would be denoted $A_{4,4,0.5}$ and would be interpreted as that expert's estimate of the fraction of people in cars[4] that would hear a siren blare within 30 seconds. In the numerical processing of the questionnaire data that will be reported on in chapter 5 we will present estimates that assume that sirens are the sole means of communication and thus only use estimates from the top row of the table above. The analysis will then be repeated using the largest fraction in each column in the above table, which usually corresponds to the "most effective pair of strategies."

Continuing with the "People in Cars" section, the experts were next asked what they thought were appropriate ways to respond to a prediction. They were given the following list of actions that a person in a car which is located on a bridge or on or under an overpass could take, to reduce his or her risk of death by taking advantage of the lead-time provided

[4] The three categories of people in cars (on, under, or near overpasses, on the open freeway, and on urban streets) were treated as one for purposes of estimating the effectiveness of different communication strategies.

by an earthquake prediction:

A: Stop the car, get out of the car, and run away from the bridge or overpass.

B: Keep driving at the same speed until a safe distance from the bridge or overpass has been reached, then stop the car but stay in the car.

As before, the experts were also asked to suggest other actions, and were then asked to rank the actions (the ones given by us, and the ones suggested by the expert) given each of the four lead-times below.

Next, the experts were asked for numerical estimates of the effectiveness of the actions they ranked first. The corresponding part of the questionnaire reads as follows:

Estimates of the Risk of Death – On Bridge or on or Under an Overpass

First, please estimate the average risk of death for people in cars on a bridge or on or under an overpass, given that an earthquake occurs without warning:

	Risk of death:	
no lead-time: earthquake occurs without warning	%	? ??

Second, consider the case when the earthquake is predicted. You are given the following scenario:

The occupant knows that an earthquake is predicted to occur in exactly τ minutes. The person believes the prediction and attempts to take the action you ranked as the most effective. Given the lead-time τ you are asked to estimate how likely the person is to successfully take the action (*Probability of successfully completing action* in the table). You are also asked to estimate the risk of death, conditional on having successfully completed the action. The same action choices are available as in the last question ...

Lead time (τ)	Most effective action (please circle)	Probability of successfully completing action	Revised risk of death, *given that* action was completed:	
30 seconds	A B C D	%	%	? ??
1 minute	A B C D	%	%	? ??
5 minutes	A B C D	%	%	? ??
30 minutes	A B C D	%	%	? ??

Here, the experts are being asked to estimate two conditional probabilities, namely the probability that a person completes the best action, *given that* she attempts that action, and the probability that the person dies in the earthquake, *given that* she completes the best action. Subsequently, the jth expert's estimate of these two quantities for the ith subpopulation and a lead-time of τ minutes will be referred to as $C_{ij\tau}$ (for the completion probability) and $D''_{ij\tau}$ (for the "revised death rate"). The experts were also asked to estimate the probability that an individual dies in an earthquake that occurs without warning; this quantity will be referred to as the "baseline" death rate and the jth expert's estimate will be denoted D'_{ij}.

The phrasing of the passages we have quoted from the questionnaire represent a compromise between precision in stating the assumptions we wanted the experts to operate under and keeping the length of the questionnaire manageable. The experts are busy professionals who were being asked by us voluntarily to donate their time and attention to a task (probability assessment) that is probably unfamiliar to some of them and challenging to all of them. In designing and testing the questions we tried to be sensitive to this. As a result, some of the experts may have wanted additional details about the interpretation and intended use of the estimates they were asked to provide. We tried to accommodate such desires in three ways. First, the questionnaire included narrative examples of how the various tables might be filled out, with explanations of how we would interpret the numbers provided. Second, a "User's Manual" accompanied the questionnaire. The User's Manual (which is reproduced in an appendix to this chapter) attempts to clarify and expand upon the explanations in the questionnaire and to address questions and comments made by some of the persons that completed the questionnaire at the testing stage. Third, several ways to reach the authors of the questionnaire were included and the respondents were encouraged to contact us with any doubts they had.

Practical Difficulties with Interpreting the Responses: An Illustration

Built into the questionnaire were a number of ways to check an expert's responses for logical consistency. We now discuss one of those checks to illustrate some of the practical difficulties with assessing expert opinion with a questionnaire.

The logical criterion we will discuss concerns the estimation of revised fatality rates. The experts were asked to estimate the probability that a randomly chosen person dies in an earthquake, conditional on two different scenarios, that when a risk-reduction action has been completed and that when no warning is available. The first of these, the revised death rate, is defined as

$$D''(\tau) \equiv \Pr\{\text{death} \mid \text{action completed}\}$$

while the baseline death rate is defined as

$$D'(\tau) \equiv \Pr\{\text{death} \mid \text{no warning}\}$$

The unconditional death rate is then computed as

$$D(\tau) = A(\tau)C(\tau)D''(\tau) + (1 - A(\tau)C(\tau))D'(\tau)$$

where $A(\tau)$ is the probability that a randomly chosen person learns of the warning in time and $C(\tau)$ is the probability that the person completes a risk-reduction action, given that the warning was heard.

Our criterion was based on the presumption that if the same action is fully completed for different lead-times, say τ_1 and τ_2, then $D''(\tau_1)$ should equal $D''(\tau_2)$, i.e., once the action is completed the revised death rate $D''(\tau)$ is fixed from that time onward. So when a respondent gives different estimates for D'' for the same action but different lead-times, then either

- The respondent misunderstood the question, or
- The respondent disagrees with our presumption.

In the case of misunderstanding, the possibility that first comes to mind is that the respondent estimated $D(\tau)$, the unconditional death rate, instead of $D''(\tau)$, the death rate for people that complete a risk-reducing action.

In the second case, the respondent might reason as follows: In 5 minutes, 50% of the population of people "at home, awake" would be able to get "outside and to an open area, free of overhead wires" (call this action x). In 30 minutes, 95% of the population would be able to complete action x. But, the respondent might argue, with 30 minutes of lead time, not only would more people be able to complete action x, more people would also be able to get to exceptionally safe open areas, as opposed to just any open area. As a result, the respondent's estimates would have the property that $D''(5) > D''(30)$ even though action x was given as the best action for both lead-times.

We contacted all respondents that gave different revised death rate estimates for different lead-times for the same action and subpopulation. Most of them indicated that they had misunderstood the question, and their responses were adjusted accordingly, with their consent. However, two respondents indicated that they disagreed with the presumption that all people who "complete" the specified action were equally safe; the estimates of those people were not changed.

The experts are asked to provide a number of estimates for each subpopulation and this is likely to involve a substantial effort. In order to limit the length of the questionnaire to a reasonable level, each expert was only asked to consider two or three subpopulations. We now discuss how subpopulations were assigned to experts.

Experimental Design: Assigning Subpopulations to Experts

The questionnaire in its full length (i.e., including all eight subpopulations) is probably too long for one person to respond to in its entirety, for two reasons: (1) the longer the questionnaire, the less likely a person would be to complete and return it, and (2) as the questionnaire gets longer, the quality of responses may get poorer, as most people's concentration will start to falter after the length of the questionnaire reaches a certain point. This subsection describes a procedure that was used to generate several versions of the questionnaire, each of which contains a subset of a "master questionnaire." Different experts were then given different versions of the questionnaire.

Each version of the questionnaire will contain some subset of the following six sections:

1. People Sleeping,
2. People at Home,
3. People in Cars,
4. Children in Schools,
5. Office Workers, and
6. People Outdoors, on Urban Streets.

The task at hand is to decide how many versions of the questionnaire there should be, which sections each version should include, and which experts should answer each version.

Important considerations for deciding how many questionnaire versions there should be, which sections each version should include, and which experts should answer which version are:

- There should be fewer versions than there are experts (obviously!) and there should be enough versions to cover all sections of the population.
- Experts should answer questions about situations that are relevant to their expertise. The primary implication of this consideration is that questionnaire versions for highway transportation experts should include the "People in Cars" section (unfortunately, we were not able to locate any such experts).
- The questionnaire versions should be about equal in length. This was achieved by limiting the length of a version to three sections, with versions that include the "People in Cars" section having only two sections, since the "People in Cars" section is twice as long as any other section.
- There should be enough overlap between versions to allow comparisons between experts.

The questionnaire versions that we finally settled on are shown in table 3.8. For example, version 3.0 E (versions 1.0 and 2.0 were developed during testing, while 3.0 was the final version) includes the sections "People at Home," "Children in School," and "People Outdoors, on Urban Streets," i.e., the sections whose squares are shaded in the column below the cell labeled 3.0 E. The column totals show the number of sections in each questionnaire version. The row totals show the number of versions that include each of the sections. Note that all versions include "people at home, awake," providing a basis for comparing experts based on a set of questions that they were all asked.

Our plan was to assign the experts to questionnaire versions as follows: Highway experts should all answer version 3.0 B, the only version that includes the "People in Cars" section (As it turned out, no highway experts completed the questionnaire). All other experts were randomly assigned to versions in such a way that the total number of experts (i.e., including highway experts) assigned to each questionnaire version were about equal for versions 3.0 C, 3.0 D, 3.0 E, and 3.0 F, and twice as many were assigned to each of versions 3.0 A and 3.0 B. Since some questionnaires were not returned, this ideal state of affairs was not achieved, but the distribution among completed and returned questionnaires was not far off the one we hoped for, as shown in table 3.9.

The plan of table 3.8 is a solution to an optimization problem which we now describe. This mathematical program seeks to minimize the variance of the estimated overall death rate in a forecasted earthquake, denoted $\mathrm{var}[D(\tau)]$, by appropriately assigning questionnaire versions to experts. The problem we attempted to solve was

$$
\begin{array}{lll}
\text{Minimize} & \mathrm{var}[D(\tau)] & \\
\text{subject to} & 7\sum_{j=1}^{8} x_{ij} = n_i & \text{for all } i \\
 & x_{1j} + x_{2j} + 2x_{3j} + x_{4j} + x_{5j} + x_{6j} = 3 & \text{for all } j \\
 & x_{ij} \in \{0,1\}, \text{ for all } i \text{ and } j & \\
 & n_i \geq l_i, \text{ for all } i &
\end{array}
$$

Section \ Version	3.0 A	3.0 B	3.0 C	3.0 D	3.0 E	3.0 F	Total:
People Sleeping	■		■			■	3
People at Home	■	■	■	■	■	■	6
People in Cars		■					1
Children in School				■	■		2
Office Workers	■			■		■	3
People Outdoors, on Urban Streets			■		■		2
Total:	3	2	3	3	3	3	

Table 3.8: The allocation of subpopulations to questionnaire versions.

Version	3.0 A	3.0 B	3.0 C	3.0 D	3.0 E	3.0 F
Returned questionnaires	10	10	5	6	7	8

Subpopulation	Number of responses
People at home, awake	46
People at home, asleep	23
People in cars	10
Children in school	13
Office workers	24
People outdoors	12

Table 3.9: Number of returned questionnaires, by version, and the resulting numbers of experts that provided estimates for each of the subpopulations. There are minor differences between the numbers reported in this table and the sample sizes reported in chapter 5, due to questionnaire that were returned partially completed.

where

$$x_{ij} = \begin{cases} 1 & \text{if } i\text{th section is in } j\text{th version} \\ 0 & \text{otherwise} \end{cases}$$

The meaning of the decision variables and constraints is best understood with reference to table 3.8. The decision variable x_{ij} is an indicator variable for whether the cell in the ith column and jth row is shaded. The first constraint can be thought of as a book-keeping device; it ensures that n_i, the number of experts that answer the ith section, is equal to the number of experts assumed to answer each questionnaire version times the number of versions that include section i. We assumed that seven experts would answer each version, and that there would be eight versions, for a total of 56 experts. The second constraint enforces the restriction that all versions include exactly three sections, except that the section "People in Cars" counts as two sections. In terms of table 3.8, the second constraint ensures that the column sums are either two (if the "People in Cars" cell is shaded) or three. The third constraint imposes a lower bound on the number of experts that answer each section. For example, if one were satisfied with the minimal requirement that each section appear in at least one version, then one would set $l_i = 7$ for all i.

Of course, we did not know beforehand how to compute the variance of the overall death rate estimate $D(\tau)$, since it depends on the actual responses; furthermore it would be different for different lead-times τ. What we did was approximate the variance using our estimates of the population fractions F_i and our prior expectations of the cross-expert variation that might be expected for different subpopulations. Thus, we used the optimization problem to devise a sampling plan that deviated from one that devotes equal sampling efforts to all subpopulations, for two reasons:

1. Some subpopulations i deserve disproportionate attention because they contain a large fraction F_i of the overall population or because the baseline death rate $D_i(0)$ could be expected to be comparatively high. Thus, the attention devoted to a subpopulation was a function of our prior estimate of the number of potential earthquake deaths $F_i D_i(0)$ that might be expected in setting i in an earthquake that occurs without warning.

2. Other subpopulations deserved careful attention, in our view, because we expected the expert responses to be more widely scattered than for others. For example, the risk facing people in cars had in our assessment been the subject of less previous research than the risk facing people inside buildings, and this might be expected to lead to wider divergence of opinion among the experts.

The sampling plan that was eventually used (see table 3.8) was a solution to the optimization problem, obtained using the *Solver* feature that comes with Microsoft's Excel spreadsheet. This solution has the property that two pairs of versions are identical, which explains why table 3.8 shows six versions even though the mathematical program assumes that eight versions are to be produced. Furthermore, all versions happened to include the "people at home, awake" subpopulation; thus fortuitously providing an opportunity to compare the experts using responses to a common set of questions (this was not guaranteed by our formulation).

Appendix 3A: User's Manual for the Questionnaire

The following paragraphs clarify and expand upon the explanations in the questionnaire. Much of this is based on comments made by persons that have already filled out the questionnaire.

Some of the remarks below may not seem relevant until one has read and attempted to respond to the questionnaire. Thus, this document is meant to be used as a reference manual, to be consulted when in doubt.

Not all respondents get the same questionnaire

The questionnaire you have has either two or three sections. Each section is about a particular population. There is a total of six categories of interest to us. However, to limit the size of the questionnaire, we only ask you to consider a subset of those categories and ask other respondents to consider different subsets. In other words, the populations we ask you to consider are not the only ones of interest to us.

Averaging across a population

Each section of the questionnaire is about a certain population, or a particular situation. We recognize that there can be significant variability within each group or situation. In estimating parameters such as the risk of death, we ask that you attempt to average across the population being considered. If you think that such an average is a particularly poor representation for the group in question, for example because the group consists of two very different subpopulations of comparable size, then we would appreciate your comments.

Validity of assumptions

In responding to the questionnaire, we ask you to make a variety of assumptions; some may seem at best highly optimistic and at worst ludicrous. We recognize that the assumptions are idealized and some of them may not be close approximations of reality and should not be interpreted as such. Rather, they should be interpreted as the ambient conditions of a "thought experiment" to be carried out by you. There are three categories of assumptions that we ask you to make; we discuss each category in turn.

Assumption of completely reliable predictions

We recognize that last-minute earthquake predictions, if they become available, are likely to be highly uncertain, particularly about the lead-time of the prediction. We also recognize that this uncertainty is likely to influence how people react to an earthquake warning.

We believe it is important to take the effect of this uncertainty on human behavior into account, and we plan to do so. Nevertheless, we ask you to assume the lead time known with certainty, because we believe it to be easier to construct mental models under this assumption, thus making it easier to respond to the questionnaire. If you are interested in how we plan to extrapolate from estimates made under this assumption to the more realistic case of uncertain predictions, then please contact us.

Assumption that all communication strategies are feasible

We realize that there are important political, technological, and cost issues that need to be considered in deciding if it would be feasible or desirable to, say, make the installation of earthquake alarms mandatory in all buildings. In asking you to not consider these issues while responding to the questionnaire, we are not "assuming them away," rather, we are trying to separate the tangled web of issues.

In particular, our study will focus on estimating the potential benefits (deaths avoided) of last-minute earthquake warnings, given that various *technologically feasible* communication strategies are used. Whether those strategies are *politically feasible* is a separate issue. By thinking about what the benefits of using a particular strategy would be, one will presumably be better able to address the issue of whether politicians (or whoever else the critical agents are) should be encouraged to pursue the strategy.

If you can think of strategies that to you seem more effective than the ones on our list, then please write them down in the space provided in the questionnaire.

Assumptions about human behavior

We ask you to assume that all persons that hear a warning believe it, and that all persons know what the most appropriate action is, where "most appropriate" means the action that increases one's chances of survival in the earthquake the most. Obviously, many people will do neither. However, some people will believe earthquake warnings, and some of the believers will know what the appropriate way to act is. We hypothesize that the deaths avoided due to an earthquake warning can be attributed primarily to this group of "informed believers."

Thus, to a first approximation at least, assessing the life-saving potential of earthquake warnings involves estimating the fraction of persons that hear, believe, and attempt to act based on an earthquake warning and then estimating how much safer these persons would be. The questions in this questionnaire are designed to elicit estimates of some (but not all) of the factors needed. For example, you are not asked to estimate the fraction of people that believe a warning, given they heard it.

Choice of most appropriate action

We ask you to identify the most appropriate action for an individual to attempt, given different values for a warning lead-time. For the longer lead-times, a person may be able to perform more than one of the actions listed, or suggested by you. If you believe this to be true, please specify the action that would contribute most to improving the chances of survival of the person. Your estimates of how much safer a person would be, given that the action was successfully completed, should be based on the improvement in safety attributable to the action you specified only, and not to other actions the person may take. If you believe that those other actions would significantly change the risk of death, we would be interested in your comments.

"Becoming aware of" versus "understanding" a warning

We recognize that "hearing" or "becoming aware of" a signal (e.g., the sound of a siren) is not equivalent to recognizing the signal as an earthquake warning. For example, sirens might be used to signal both earthquake forecasts and accidents at a nearby nuclear reactor, and the most appropriate actions may be very different for the two scenarios. Some people might be momentarily confused as to what the sirens are signalling, even if the sirens signalling the two situations were designed to have distinctive sounds.

Thus, two communication strategies that are equally effective in making persons aware of a warning do not necessarily cause equal fractions of the population to take appropriate action. This may be because one warning contains more useful information than the other, because it is more believable, because it is less ambiguous, etc.

While we recognize this, we ask you to "assume it away" in your responses. Thus, your responses should be estimates of what fraction of the given population "becomes aware of the signal," rather than "understands and believes the warning." This does not mean that we plan to ignore this issue; it reflects the fact that we ask you to estimate some, but not all, of the parameters needed to estimate the number of deaths in a forecasted earthquake. In our analysis of the questionnaires we plan to take the "warning versus signal" issue into account by processing the data two ways. First, we will assume hearing a signal to be equivalent to understanding and believing an earthquake warning, and acknowledge that this leads to excessively optimistic estimates (a best-case scenario). Second, we will include conditional probabilities that a person understands, believes, and acts upon a warning, given the person has heard the warning, and use plausible values for those parameters (a more realistic scenario).

Chapter 4

Models of Reactions to Earthquake Warnings

This chapter describes models of the death toll in a major earthquake, as a function of warning lead-time. Chapters 3 and 5 develop estimates of the parameters in the models to be presented in this chapter. The models will be presented in decreasing order of generality and increasing order of detail. Thus, we start by presenting a model which is essentially a disaggreggation of the overall death toll for an urban area into the sum of the death tolls in each of several subpopulations. Then, the death rate in each subpopulation is expressed in terms of parameters that measure (1) the effectiveness with which the earthquake warning is propagated, (2) the time needed by people to complete an action whose purpose is to increase chances of survival, (3) the chances of survival given that such an action is completed, and (4) the chances of survival if such an action is not attempted. The resulting expression for the death rate is based on certain assumptions about the capabilities of humans and technology; assumptions which are optimistic in the sense that they are not meant to capture what humans *would* do under present circumstances but rather what humans *could* do with proper education, training, and motivation. At the highest level of detail, we present models that decompose some of the four components of the death rate mentioned above.

Finally, models are presented of how experts might estimate the parameters of the behavioral models in this chapter. The models are likelihood functions, i.e., they postulate what the probability distribution of expert estimates, conditional on the true values of the parameters, would be. These likelihood functions will be used in chapter 5 to estimate the parameters of the behavioral models using data from the survey of experts.

4.1 Description of a Model for the Death Toll

After an earthquake has occurred, various agencies have a responsibility or interest in estimating the human and economic costs of the earthquake. Such estimates are constructed by aggregating information obtained from various sources, e.g., hospitals, relief agencies such as the Red Cross, police departments, and public officials.

The procedures that have been used to estimate the damage caused by future earthquakes in studies performed for several U.S. urban areas [UNK93, Pan89] are similar to the procedures used for assessing damage after an earthquake has occurred, in that the total damage is estimated as the sum of damage estimates for various categories.

Motivated by these considerations, we propose to model the death toll $N_D(\tau)$ in a future earthquake which is forecast τ time units ahead of its occurrence by

$$N_D(\tau) = \sum_{i \in I} N_{D,i}(\tau) \tag{4.1}$$

where I is a set of conveniently chosen categories or subpopulations, which has the property that every individual that lives in the area of interest belongs to exactly one category, and $N_{D,i}(\tau)$ is the number of individuals in category i that would die as a result of the earthquake. Chapter 3 presented the categories that we chose.

If we let

$$\begin{aligned} D_i(\tau) &\equiv \text{Death rate for people in category } i, \text{ given } \tau \text{ time units of warning} \\ N_i &\equiv \text{Number of individuals that belong to category } i \end{aligned}$$

then the model (4.1) can be re-expressed as

$$N_D(\tau) = \sum_{i \in I} N_i D_i(\tau) \tag{4.2}$$

After an earthquake has occurred, equation (4.2) evidently holds as a counting relationship. Before an earthquake, it holds also, with the death rates $\{D_i(\tau)\}$ interpreted as random variables. Taking expected values on both sides of (4.2), one obtains (assuming $D_i(\tau)$ to be independent of N_i)

$$\mathrm{E}[N_D(\tau)] = \sum_{i \in I} \mathrm{E}[N_i]\mathrm{E}[D_i(\tau)] \tag{4.3}$$

Dividing both sides with the total population $N = \sum_{i \in I} N_i$ (assumed constant) of the area of interest, the following equation results:

$$\frac{\mathrm{E}[N_D(\tau)]}{N} = \sum_{i \in I} \frac{\mathrm{E}[N_i]}{N} \mathrm{E}[D_i(\tau)] \tag{4.4}$$

This form of the model can be interpreted as follows. Let X denote a randomly chosen individual from the population of the area of interest, i.e., X is equally likely to be any inhabitant of the area of interest. Then

$$\begin{aligned} &\Pr\left\{ \begin{array}{l} X \text{ dies in a future earthquake which} \\ \text{is forecasted with a lead-time } \tau \end{array} \right\} \\ &= \sum_{i \in I} \Pr\{X \text{ belongs to category } i\} \Pr\{X \text{ dies in quake} \mid X \text{ belongs to category } i\} \end{aligned}$$

Thus, the basic model (4.2) can be viewed both as a counting relationship and as a statement about conditional probabilities. However, its primary use is in pointing the way toward four important modeling and estimation activities, namely

1. Deciding on how to divide the population into categories.
2. Estimating the sizes of the subpopulations (the N_i's) or equivalently, the population fractions (the F_i's, where $F_i = N_i/N$).
3. Modeling the expected death rates in each category $\{\mathrm{E}[D_i(\tau)]\}$.
4. Estimating the expected death rates in each category.

In this chapter, the emphasis will be on the third activity. The first two activities were discussed in chapter 3, and the fourth activity is the subject of chapter 5.

4.1.1 Description of a Model for the Death Rates

Figure 4-1 illustrates a model for the death rate[1] $D_i(\tau)$ in category i. The model postulates that two important events will affect the chances of survival in an earthquake of any member of category i when the earthquake is forecasted with τ time units of lead time: (1) The person will either hear (or see) the warning before the earthquake occurs, or she will fail to do so, and (2) Given the person hears the warning, she will either succeed or fail in improving her chances of survival by performing some action, such as getting under a desk or getting outside. Formally, if

$$\begin{aligned}
A_i(\tau) &\equiv \Pr\{\text{ person hears warning} \mid \text{category } i\text{, lead time } \tau\} \\
C_i(\tau) &\equiv \Pr\left\{\begin{array}{l}\text{person completes action to} \\ \text{improve chances of survival}\end{array}\middle|\begin{array}{l}\text{person} \\ \text{tries to act}\end{array}\right\} \\
D_i'' &\equiv \Pr\left\{\begin{array}{l}\text{person dies in} \\ \text{the earthquake}\end{array}\middle|\begin{array}{l}\text{person com-} \\ \text{pletes action}\end{array}\right\} \\
&= \text{“revised” death rate} \\
D_i' &\equiv \Pr\left\{\begin{array}{l}\text{person dies in} \\ \text{the earthquake}\end{array}\middle|\begin{array}{l}\text{person does not} \\ \text{complete action}\end{array}\right\} \\
&= \text{“baseline” death rate}
\end{aligned}$$

then

$$D_i(\tau) = A_i(\tau)C_i(\tau)D_i'' + (1 - A_i(\tau)C_i(\tau))D_i'$$

This model can be thought of as a specialization and elaboration of the general model of risk communication theory, as embodied in figure 2-7, to the case of very short lead times, where the shaded path in figure 2-7 will be the most important one.

4.1.2 Assumptions

Implicit in the model for the death rates are several assumptions which we now discuss:

People respond without delay: Everyone who hears the warning understands it immediately and decides to attempt an action designed to improve chances of survival in the earthquake. That is, we assume that

$$\Pr\{\text{ warning understood} \mid \text{warning heard}\} = 1$$

and

$$\Pr\{\text{ action attempted} \mid \text{warning heard and understood }\} = 1$$

[1] In the rest of this chapter and in other chapters of the thesis, we will take $D_i(\tau)$ to mean the *expected death rate.*

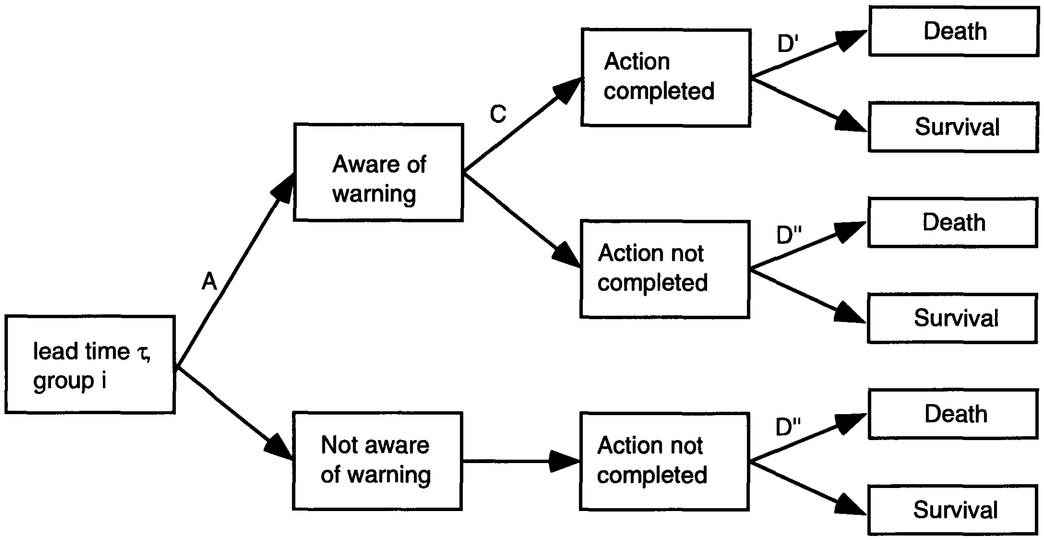

Figure 4-1: Sequential sample space for a randomly chosen individual in category i, given a lead-time τ.

Note that the model could easily be generalized to relax these assumptions; one would simply add two more levels to the tree of figure 4-1. There are two reasons for not doing so. First, it would complicate the model and increase the number of parameters that would need to be estimated. Second, and more importantly, the two probabilities assumed to equal one above depend on how individuals respond to earthquake warnings, whereas the two probabilities that appear on the tree in figure 4-1 (probability of hearing a warning and probability of completing an action) are determined by the constraints of available technology and human reaction times. Human response to warnings is not constant; it can be influenced by education and training (such as earthquake drills in schools and workplaces). Since a primary motivation of this study is to determine how people *should* respond to earthquake warnings rather than how they would react under the status quo we make this assumption.

Having said that, it must be admitted that even if the cost of educating and training people were no object, it is optimistic to expect that everyone will be capable of immediately understanding and acting on an earthquake warning. Sensitivity analysis will be used in chapter 5 to assess the consequences of such optimism.

Chances of survival depend only on whether action is completed: One's chances of survival depend only on whether one completes an action, such as getting under a desk, intended to reduce risk. In particular, the chances of survival are assumed the same for persons who do not hear the warning and for persons who hear the warning but are not able to complete an action.

This is probably an optimistic assumption. To see why, consider the action of leaving a building, and consider two groups of people: (1) people that hear the warning and attempt to leave the building they are in, but are unable to do so before the earthquake, and (2) people that do not hear the warning and hence do not attempt to leave the building before the earthquake occurs. Under our assumption, both groups will be assumed to face the baseline death rate D'. However, it might be argued that people in the first group are likely to be at higher risk. The argument would proceed as follows: A substantial fraction of people in the first group will be in the process of walking out of the building when the earthquake occurs. Some people in the second group may attempt to leave the building after the earthquake has started and will likely be on the way out when the strongest motion occurs. While it might be safer to be outside than inside when the earthquake occurs, being on the way out is likely to be more dangerous than either, because of objects falling off the outside of the building. Thus, there are really three levels of risk: inside, on the way out, and outside. The two groups of people we are considering consist of people that are either inside or on the way out and we expect the first group (people who heard the warning but did not manage to get outside in time) to have a higher proportion of people on the way out. As a result, one would expect the death rate to be higher in the first group.

People only have time to attempt one action: The lead time τ is short enough that people do not have time to execute complicated strategies. We assume that a person either attempts to perform a simple action or does nothing. As an example of behavior which we exclude, consider a person who hears the warning while driving, begins by calling other members of the family to make sure they are aware of the warning, then drives home, stopping on the way to purchase emergency supplies. The person might then turn off gas and electricity mains, secure objects that might be damaged in the earthquake, and so on.

The lead time τ is known with certainty: This is an assumption both about the reliability of the forecasting system and the effectiveness with which a warning is communicated, i.e., we assume both that the body which issues a prediction knows what τ is and that anyone who hears the earthquake warning will know what τ is.

As discussed in chapter 2, seismologists seem to agree that even if breakthroughs were to occur in the science of earthquake prediction there would continue to be substantial uncertainty associated with any statement about the instant in time when a future earthquake will occur. So once again, we have an optimistic assumption; perhaps the one with the most serious consequences. But an analysis performed under this assumption lays the groundwork, so to speak, for an analysis which takes the uncertainty about the lead time into account; this is discussed further in chapter 5. Furthermore, a statement about the potential utility of a perfect prediction is of independent value, particularly if it turns out that improvements in public safety equal to this optimistic scenario can be achieved by other means at lower cost, say, by improving pre-natal care or driver education.

All of the above assumptions can be characterized as optimistic, and thus an estimate of the life-saving potential of earthquake warnings, derived under these assumptions, should be interpreted as an upper bound, or best-case estimate. In chapter 5, sensitivity analysis will be used to assess how important some of these assumptions are.

The next three sections of this chapter will present models for the branch probabilities at the three levels of the tree shown in figure 4-1, i.e., the fraction of people aware of the warning $A_i(\tau)$, the fraction of those that hear the warning that are able to complete an action $C_i(\tau)$, the baseline death rate D'_i, and the revised death rate D''_i. Frequently, the category subscript i, or the lead-time argument τ, or both, will be omitted for clarity.

4.2 Models for Warning Propagation

This section will describe a simple model for how a warning is communicated to a population through a particular channel, e.g., using sirens. Let t index time, with $t = 0$ being the instant at which the warning is issued and $t = \tau$ the instant at which the earthquake occurs. Let $\tilde{A}(t)$ be the fraction of the population in category i which is aware of the warning at time t.

There are two mechanisms by which an individual can become aware of the warning:

1. Hearing (or seeing) the warning as it is broadcast through the channel being used.
2. Being told by another individual.

One would expect the first mechanism to dominate initially and the second mechanism to eventually take over. The function $A(\tau)$ (cf. figure 4-1) that is ultimately of interest is the fraction of the population that becomes aware of the warning before the earthquake occurs as a function of the prediction lead-time, rather than as a function of clock time t. An assumption that the functions $\tilde{A}$ and A are identical is not plausible in general. To take an extreme example, compare two warnings, one with a lead-time of 7 days and the other with a 5 minute lead-time. Assuming that the same methods are used to communicate the two warnings (at least initially) and that the value of the lead-time is part of the warning message, one would presume that the 5-minute warning would be perceived as much more urgent. As a result, the fraction of people aware of the 5-minute warning one minute after it is issued is likely to be greater than the fraction of people aware of the 7-day warning one minute after it is issued.

However, for lead times of about the same order of magnitude, say less than 30 minutes and more than 30 seconds, assuming that $\tilde{A} = A$ is perhaps harmless, and this assumption will be in effect for the rest of this section.

The model to be described will be in terms of an ordinary differential equation for the population fraction A. Suppose that the rate of change of A is the sum of two components (α and β), one due to each of the two mechanisms. Suppose that the component due to the first mechanism is proportional to the fraction of people that are unaware of the warning, i.e., $\alpha \propto 1 - A$. For the second mechanism, suppose that the component is proportional to the number of possible contacts between people that are aware and people that are unaware of the warning, i.e., $\beta \propto NA(1 - A)$.

The proportionality with $NA(1-A)$ is justified under the following homogeneous mixing assumption: the probability that a randomly chosen pair of people, of which one is aware and the other not aware of the warning, is constant and independent of all other such pairs. Later in this section, the implications of assuming instead a certain kind of spatial dependency will be explored.

Furthermore, suppose that each component of the rate of change is proportional to an "urgency factor" $U_i(t, \tau), i = 1, 2$ which could, in general, be a function of both the clock time t and the lead-time τ.

In its most general form, the model described above can be stated as

$$\frac{dA(t)}{dt} = \alpha + \beta = U_1(t,\tau)(1 - A(t)) + U_2(t,\tau)NA(t)(1 - A(t)) \quad (4.5)$$

$$A(0) = 0 \quad (4.6)$$

To interpret the parameters of the model, it is helpful to re-express it in terms of the number of people, $Y(t)$, that are aware of the warning at time t. If the total number of individuals is N, then $A = Y/N$, and

$$\frac{dY}{dt} = U_1(N - Y) + U_2Y(N - Y)$$
$$\Rightarrow Y(t + \Delta t) \simeq Y(t) + U_1\Delta t(N - Y) + U_2\Delta tY(N - Y)$$

One interpretation of this equation is the following. Let Q and R be two randomly chosen members of the population. Then

$$\Pr\left\{\begin{matrix}Q \text{ hears warning from pri-} \\ \text{mary source in } (t, t + \Delta t]\end{matrix}\middle|\begin{matrix}Q \text{ does not know} \\ \text{at time } t\end{matrix}\right\} \approx U_1\Delta t$$

and

$$\Pr\{R \text{ tells } Q \text{ in } (t, t + \Delta t]\,|R \text{ knows, } Q \text{ doesn't, at time } t\} \approx U_2\Delta t$$

Note that both U_1 and U_2 have units of minute^{-1} (if time is measured in minutes) and can be interpreted as rates of transmission of information. If there is no person-to-person contact ($U_2 = 0$) and the rate U_1 at which the warning is transmitted from the primary source is constant, then $1/U_1$ can be interpreted as the mean time until a randomly chosen member of the population hears the warning.

Based on this interpretation of the urgency functions U_1 and U_2 as transmission rates, it is reasonable to restrict them to be non-negative and bounded, i.e., the transmission rates cannot be negative and they cannot be infinite. To avoid trivialities, we also assume that the primary source is turned on at time zero, i.e.,

$$U_1(t) > 0 \text{ for } t \in [0, \epsilon] \text{ for some } \epsilon > 0 \tag{4.7}$$

Isolated individuals

The model (4.5)-(4.6) can easily be extended to model a belief that some individuals in the population are unlikely to hear the warning within a reasonable amount of time. This might for example be the case for some elderly people that live on their own, homeless people, or recent immigrants that face a language barrier[2]. Suppose that the fraction of people that are *not* isolated is p. Then the differential equation (4.5) would be modified to

$$\frac{dA}{dt} = U_1(p - A) + U_2NA(p - A); A(0) = 0$$

The fraction of people $1 - A$ that has not yet heard the warning has been replaced with the fraction $p - A$ of the population that has not yet heard the warning, but will eventually do so.

We will now describe a set of assumptions under which the differential equation (4.5) is an appropriate description of how the warning propagates through a population. These assumptions should be taken to apply to the non-isolated members of the population; the isolated members are assumed never to hear the warning.

[2]Of course, measures could and probably should be taken to avoid such isolation. The point is that the model can be used to understand the trade-offs involved in taking such measures.

4.2.1 Behavioral Assumptions

People don't have time to forget the warning: If the lead-time is short (less than an hour, say) and the warning is perceived as urgent, this is a reasonable assumption. For longer lead-times and for situations where the warning is not perceived as urgent, for example because of frequent false alarms in the past, this assumption is harder to justify.

Homogeneous mixing: At any time t, there is homogeneous mixing of the population, i.e., all pairs of individuals in the population are equally likely to meet and communicate the warning (if one knows and the other doesn't). This assumption is made primarily for tractability; it is difficult to envision a situation where it would hold literally. In appendix 4B, a model which relaxes this assumption will be analyzed. The result of this analysis is that the model (4.5) remains appropriate under certain departures from this assumption, i.e., the homogeneous mixing assumption is sufficient but not necessary for the differential equation (4.5) to be appropriate.

The system has no memory, in the sense that the probability that any member of the population hears the warning in the time interval $(t, t + \Delta t]$ (assuming he has not heard it yet) depends only on how many members of the population have heard the warning by time t. In particular, this probability does not depend on the order in which people that have heard the warning by time t became aware of the warning or when these people heard the warning.

4.2.2 Solution of the Differential Equation

The differential equation (4.5) can be solved by a change of variable that turns it into a linear differential equation. Let $w = A/(1 - A)$. Then $A = w/(1 + w)$ and $dA/dt = (dw/dt)/(1 + w)^2$. Substituting for A in (4.5) and multiplying all terms with $(1 + w)^2$ results in

$$\frac{dw}{dt} = (U_1 + U_2 N)w + U_1;\ w(0) = \frac{A(0)}{1 - A(0)} = 0 \tag{4.8}$$

The general solution to equation (4.8) is

$$\begin{aligned} w(t) &= w(0)e^{\int_0^t (U_1(s)+U_2(s)N)ds} + \int_0^t e^{\int_s^t (U_1(r)+U_2(r)N)dr} U_1(s)ds \\ &= \int_0^t e^{\int_s^t (U_1(r)+U_2(r)N)dr} U_1(s)ds \end{aligned} \tag{4.9}$$

(dependence of the U_i on τ has been suppressed.) Using $A = w/(1+w) = 1 - 1/(1+w)$, a functional form for $A(t)$ is obtained:

$$A(t) = 1 - \frac{1}{1 + \int_0^t e^{\int_s^t (U_1(r)+U_2(r)N)dr} U_1(s)ds} \tag{4.10}$$

Next, we will explore solutions for two special cases. The two special cases correspond to the two models that will be calibrated in chapter 5 using expert judgment.

Solution for Constant Urgency

If the urgency functions U_1 and U_2 do not depend on time, one obtains:

$$\begin{aligned} A(t) &= 1 - \frac{1}{1 + \int_0^t e^{\int_s^t (U_1 + U_2 N) dr} U_1 ds} \\ &= 1 - \frac{1}{\frac{U_1}{U_1 + U_2 N} e^{(U_1 + U_2 N)t} + \frac{U_2 N}{U_1 + U_2 N}} \end{aligned} \tag{4.11}$$

As one might expect, the solution A has the properties that $A(0) = 0$ and $\lim_{t \to \infty} A(t) = 1$, if one assumes that $U_1 > 0$ (as required by (4.7)). In appendix 4A, we will show that $A(t)$ has these properties if and only if the urgency functions do not decay too rapidly. More precisely, the condition is $\int_0^\infty (U_1(s) + U_2(s)) ds = \infty$.

One of the methods we use to estimate the parameters U_1 and U_2 in this model requires the inverse function A^{-1}; it is

$$A^{-1}(a) = \frac{1}{U_1 + U_2 N} \log \frac{U_1 + a U_2 N}{U_1 (1 - a)} \tag{4.12}$$

We note that the fraction $A(t)$ of people that have heard the warning by time t can be thought of as the cumulative distribution function for the time T_A until a randomly chosen individual hears the warning. We were not able to get a simple expression for the mean time until a randomly chosen person hears the warning, but to first order in $U_2 N$, we have

$$\mathrm{E}[T_A] \approx \frac{1}{U_1} \left(1 - \frac{1}{2} U_2 N \right)$$

Figure 4-2 shows three possible $A(t)$ curves. Note that depending on the values of U_1 and U_2, the $A(t)$ curve can range from an S-shaped curve similar to a logistic curve to a curve similar to the cdf for an exponential random variable. These are in fact the limiting functional forms; when U_1 approaches zero, the differential equation approaches

$$\frac{dA}{dt} = U_2 N A (1 - A)$$

whose solution is the logistic function, and when U_2 approaches zero the differential equation approaches

$$\frac{dA}{dt} = U_1 (1 - A)$$

whose solution is $A(t) = 1 - e^{-U_1 t}$.

Calibration of the constant urgency model using expert judgment suggested that the parameter U_2 was close to zero and also that the constant urgency model did not provide a particularly good fit to the expert estimates. This motivated us to consider another special case of the general model (4.5), which we discuss next.

Solution for Decreasing Urgency and No Person-to-Person Contact

Suppose that the rate of transmission from the primary source decreases with time according to

$$U_1(t) = \frac{\lambda}{t + \gamma}$$

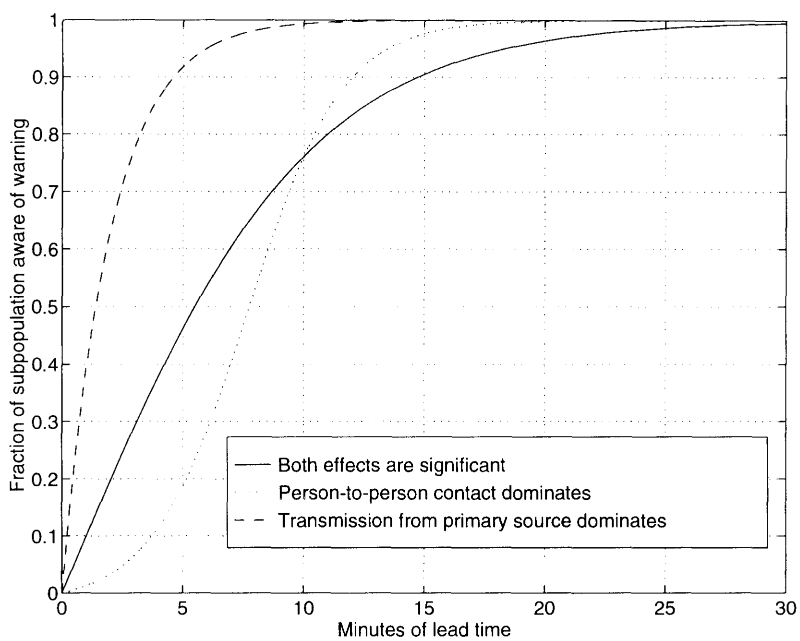

Figure 4-2: Three solutions to the differential equation $dA/dt = U_1(1 - A) + U_2NA(1 - A)$ corresponding to $(U_1, U_2N) = (0.1, 0.1)$ (both transmission from primary source and person-to-person contact are significant), $(0.01, 0.5)$ (person-to-person contact dominates), and $(0.5, 0.001)$ (transmission from primary source dominates).

where $\lambda > 0$ and $\gamma > 0$, and there is no person-to-person transmission of the warning ($U_2 = 0$). Recall that the urgency factor $U_1(t)$ may be thought as being proportional to the conditional probability that an individual will hear the warning in the next infinitesimal time interval $(t, t+\Delta t]$, given that the person has not yet heard the warning. Here, we allow this conditional probability to decrease with time. For last-minute warnings we believe this to be plausible. For example, knowing that a randomly chosen person has not heard the blare of sirens one minute after the sirens were turned on would heighten the chances that the person is fast asleep or unlikely to hear the sirens in the near future for other reasons.

The phrase "no person-to-person contact" that we associated with the condition $U_2 = 0$ should not be taken too literally. Under the assumptions we made to justify the differential equation (4.5) we associated the transmission rate U_2 with communication that occurs when one person hears the warning from the primary source and then, after a random amount of time has passed, happens to run into another person who has not yet heard the warning. The parameter U_2 was meant to be proportional to the rate at which such random encounters result in a person becoming aware of the warning. It appears plausible that if the warnings are only broadcast for a few minutes prior to the earthquake, then this process of random encounters will not be a significant contributor.

However, one can envision a different scenario, for example where a group of children in a classroom hear an earthquake warning through a PA system. It is then possible that some of the children do not hear the warning, but are informed by their peers immediately, rather than after a random period of homogeneous mixing. The special case we consider here does allow for this kind of person-to-person contact.

The solution to (4.5) for the decreasing urgency special case is

$$A(t) = 1 - \left(\frac{\gamma}{t+\gamma}\right)^{\lambda}$$

The inverse function is

$$A^{-1}(y) = \frac{\gamma}{(1-y)^{1/\lambda}} - \gamma$$

and the mean time until a randomly chosen person hears the warning is

$$\mathrm{E}[T_A] = \frac{\gamma}{\lambda - 1} \text{ for } \lambda > 1, \infty \text{ otherwise}$$

Figure 4-3 shows three possible solutions corresponding to different values of λ and γ.

Additional material on the spatial propagation model appears in appendices 4A and 4B; this material is not used in other chapters of the thesis and is not needed to understand them. In appendix 4A, we establish general properties of solutions to the differential equation 4.5 and in appendix 4B we present and analyze a model that relaxes the homogeneous mixing assumption and show that it reduces to the homogeneous mixing model under certain conditions.

4.3 Models for the Time Needed to Complete the Best Action

We now describe a model for the probability that a person who hears a warning, understands it, and decides to act on it, is able to complete the action chosen. To put this model in context, let M be a randomly chosen member of a particular category and define the

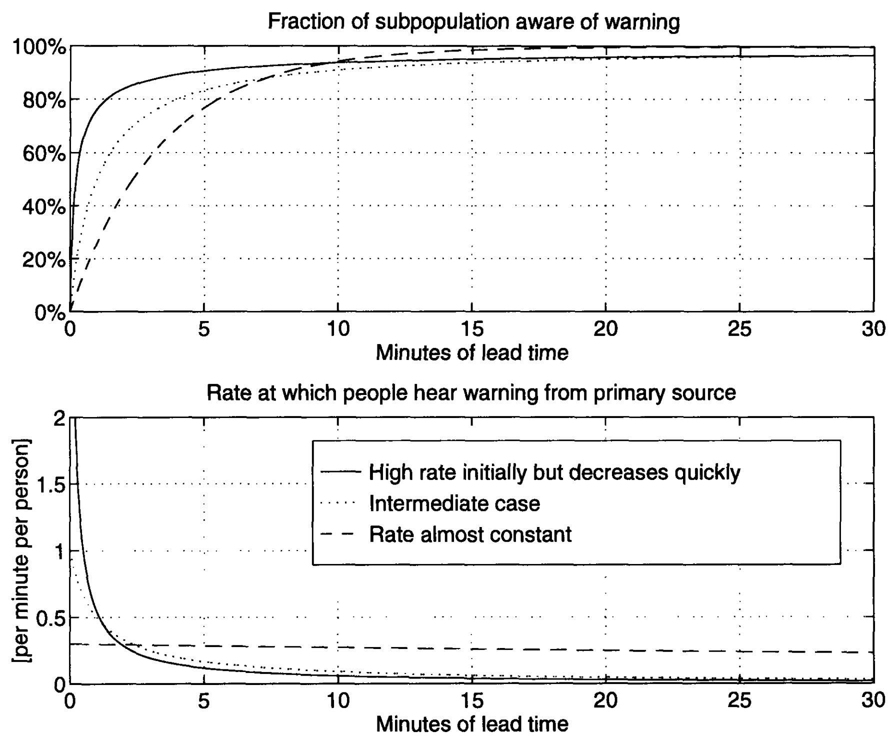

Figure 4-3: Three solutions to the differential equation $dA/dt = U_1(t)(1-A)$ where $U_1(t) = \lambda/(t+\gamma)$. The three solutions correspond to the transmission rate from the primary source being high initially but decreasing rapidly ($\lambda = 0.6, \gamma = 0.1$), an intermediate case ($\lambda = 1, \gamma = 1$), and the transmission rate being essentially constant ($\lambda = 30, \gamma = 100$).

following two random variables that describe the sequence of events that may occur between the time $t = 0$ when an earthquake warning is first communicated to the public and the time $t = \tau$ when the earthquake occurs:

$$\begin{aligned} T_A &\equiv \text{time until } M \text{ hears the warning.} \\ T_C &\equiv \text{time needed by } M \text{ to complete the action he chooses.} \end{aligned}$$

Our assumptions are that individual M will face the revised death rate D'' if and only if $T_A + T_C < \tau$, i.e., if M has time to both hear the warning and complete the best action before the earthquake occurs (recall that we assume that M understands the warning immediately and decides to act on it). Otherwise, i.e., if $T_A + T_C > \tau$, M will face the baseline death rate D'.

The preceding section presented a model stated in terms of the fraction $A(t)$ that have heard the warning by time t, and we noted that $A(t)$ can be interpreted as the cdf for random variable T_A. In this section we will present one model for the probability distribution of T_C, the time to complete an action.

Our model is based on the following two assumptions:

"Slow" and "quick" individuals: We assume that the people that fall into any category or subpopulation can be divided into people that can be expected to be relatively slow in completing the best action and those that can be expected to complete it comparatively quickly.

Exponentially distributed completion times: We assume that the time needed by a slow person to complete the best action is an exponential random variable with mean $1/\lambda_1$ and the time needed for completion by a quick person is an exponential random variable with a mean $1/\lambda_2 < 1/\lambda_1$.

These two assumptions imply that the distribution for T_C, the time needed by a randomly chosen individual to complete the best action, is a mixture of exponentials. Thus, the pdf for T_C is

$$f_{T_C}(t) = p\lambda_1 e^{-\lambda_1 t} + (1-p)\lambda_2 e^{-\lambda_2 t} \tag{4.13}$$

where p is the fraction of people that are "slow," and the cdf is

$$F_{T_C}(t) = p(1 - e^{-\lambda_t}) + (1-p)(1 - e^{-\lambda_2 t})$$

Under the model-based approach to processing the expert estimates (see chapter 5), we will estimate the parameters of the models presented in this section and the last section. The resulting parameter estimates will provide us with estimates for the distributions of T_A and T_C; these distributions will then be convolved to obtain estimates of $\Pr\{T_A + T_C < \tau\}$,the probability that a randomly chosen individual hears the warning and completes the best action before the earthquake occurs, for different warning lead-times τ.

The model 4.13 has the advantage of simplicity; it has only three parameters (p, λ_1, and λ_2) that we will estimate in chapter 5 using expert judgment. We hope that it also captures enough of reality to provide an adequate first-order approximation.

For other purposes, even if the slow/quick dichotomy turns out to be too crude, we hope that this model at least points the way toward a more realistic description. For example, the slow/quick classification could be expanded to "quick," "average," "slow," and "not likely to complete action in the near future." Also, the exponential distributions could be replaced with other distributions if this was thought to be more appropriate.

4.4 Estimation of the Baseline and Revised Death Rates

In this section, we discuss how to combine the baseline and revised death rates that the experts are asked to estimate for a variety of circumstances.

Suppose that n experts estimates a death rate D (either a baseline or revised death rate) for people in a particular setting. Let Y_i be the i-th expert's estimate of D, and suppose that

$$Y_i = D\epsilon_i$$

where ϵ_i is a random variable. Thus, we assume that the estimate given by the i-th expert is the true death rate D, multiplied by an error term ϵ_i. Since both the true death rate D and the estimates Y_i are restricted to lie between zero and one and will generally be quite close to zero, a multiplicative error term seems more appropriate than an additive one.

We will present two estimators $\hat{D}_1$ and $\hat{D}_2$ that combine the n expert estimates $\{Y_i\}_1^n$ into one. Some properties of these estimators will be established under the assumption that the error multipliers ϵ_i are independent and lognormally distributed with parameters $\mu = 0$ and $\sigma > 0$. Some reasons for making the lognormal assumption are

- It provides a tractable model.
- It has the intuitively pleasing consequence that any expert, in estimating D, is equally likely to do so by dividing the true rate D by some number as she is to multiply D by the same number. That is, for any number $x > 0$, the expert's estimate Y_i is equally likely to be D/x or Dx.
- A counterpart to the Central Limit Theorem guarantees, roughly speaking, that if the estimate Y_i of D is the result of multiplying together many independent factors, none of which dominates the overall product, then Y_i will be approximately lognormally distributed. It is not difficult to imagine that some of the experts may have gone through thought processes that involved successively multiplying some initial number. For example, for the subpopulation "people at home, awake" one might start with an estimate of the total number of people that fall into that group in some urban area one is familiar with. Then one might multiply that number with one's estimate of the fraction of homes that would collapse in an earthquake. The resulting product might then be multiplied with an estimate of the conditional probability of perishing, given that one is in a home that collapses.
- Histograms of the logarithms of the baseline death rate estimates provided by the experts (see figure 5-1) are not hostile to the lognormal assumption.

Under this model, the experts are biased estimators of the true death rate D, since

$$\mathrm{E}[Y_i] = D\mathrm{E}[\epsilon_i] = De^{\sigma^2/2} > 1 \text{ for } \sigma > 0$$

However, the *median* of the distribution for Y_i equals D. This has implications for combining the expert estimates into one estimate of D: one should attempt to estimate the median of the distribution of Y_i rather than its expected value.

We now consider the problem of estimating D under the above assumptions. If ϵ_i is lognormal, then Y_i will be lognormal also, with parameters $\log D$ and σ, i.e., $\log Y_i \sim$

$N(\log D, \sigma^2)$. Therefore, a "natural" method of estimating D would be to estimate $\log D$ as the sample mean of the $\log Y_i$'s and then estimate D by raising e to the appropriate power, i.e.,

$$\hat{D}_1 = \exp\left\{\frac{1}{n}\sum_{i=1}^{n}\log Y_i\right\}$$

Straightforward algebra shows that this is in fact the geometric mean of the original sample:

$$\hat{D}_1 = \left(\prod_{i=1}^{n} Y_i\right)^{1/n}$$

The mean and variance of $\hat{D}_1$ are computed in appendix 4C, they are

$$\mathrm{E}[\hat{D}_1] = De^{\sigma^2/(2n)} \text{ and } \mathrm{var}[\hat{D}_1] = D^2\left(e^{2\sigma^2/n} - e^{\sigma^2/(2n)}\right)$$

Thus, the geometric mean is biased upward by a factor of $\exp(\sigma^2/(2n))$, but it is asymptotically unbiased and consistent. It is also the maximum likelihood estimator of D.

An unbiased alternative to the geometric mean, which we refer to as a *modified geometric mean*, is derived in appendix 4C. This estimator has the form

$$\hat{D}_2 = \hat{D}_1 f_n(S^2)$$

The complicated expression $f_n(S^2)$ (given in appendix 4C) multiplying the geometric mean $\hat{D}_1$ adjusts it downward by an amount that depends on the sample size n and the sample variance $S^2 = n^{-1}\sum_1^n(\log Y_i - \overline{\log Y})^2$ of the logarithms of Y_i. This downward adjustment results in $\hat{D}_2$ being an unbiased estimator of D.

The function $f_n(S^2)$ is plotted in figure 4-4 for $n = 5$, 10, and 20. It is disturbing that it is possible for the function to take on negative values, since this would result in negative estimates of D, which must of course be positive. However, negative values are unlikely, unless the true variance σ^2 (of $\log Y$) is large or the sample size is small.

We show in appendix 4C that the modified geometric mean $\hat{D}_2$ is the minimum variance unbiased estimator of D. Of course, this does not mean that the modified geometric mean $\hat{D}_2$ has lower mean squared deviation than the geometric mean $\hat{D}_1$. All we know is that $\hat{D}_2$ has less mean squared deviation than any other unbiased estimator, and $\hat{D}_1$ does not belong to that class. Unfortunately, computing the variance of $\hat{D}_2$ presents algebraic difficulties. While it might be possible to compute the first few terms of an asymptotic expansion of the variance, this would not be particularly useful, since ignoring higher order terms in $(1/n)$ is justified only for large sample sizes. For large sample sizes, the estimators are essentially the same, since they are both functions of a complete sufficient statistic (see appendix 4C) and since their expectations are asymptotically the same.

To get around these difficulties, simulation was used to compare the performance of the two estimators. One thousand samples were generated under six different circumstances: The sample size n was set at either 5, 10, or 20, and the true median D was set at either 10^{-3} or 10^{-6}. The variance parameter σ^2 was fixed at one. These values were chosen to be representative of the sample sizes and likely ranges of parameter values for the death rates estimated by the experts.

Table 4.1 presents the results of the simulation. Four performance measures were computed. The first is the reduction in bias achieved by using the modified geometric mean.

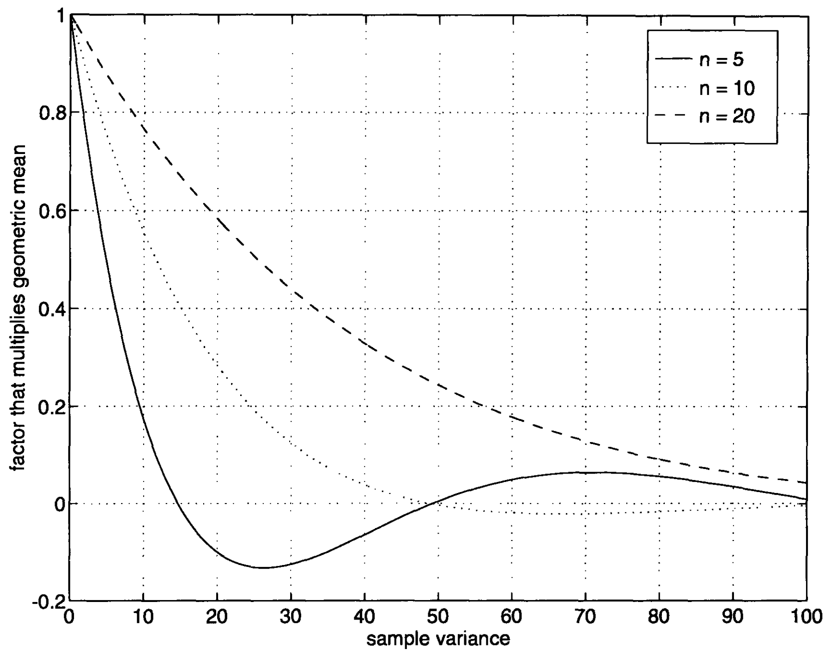

Figure 4-4: The function $f_n(S^2)$ for $n = 5, 10, 20$. We note that for large sample variances and small sample sizes, the function can take on negative values.

The results here are not surprising since the expectations of both estimators have been computed. The overall pattern is one of an increase in bias with reduced sample sizes and lower values of D. More generally, one would expect bias to increase as $\sigma / \log D$ increases, but in the simulation σ was fixed. The second performance measure is the percentage reduction in root mean squared deviation of the estimates from the true value D that is achieved by using the modified geometric mean. We see that here the pattern is the same as for bias: The modified geometric mean always does better, and the proportional improvement is larger when the sample size is smaller or when the true median D is smaller. This suggests that the two estimators may have about the same variance, so that any reduction in bias would directly reduce the mean squared deviation. The third performance criteria is the fraction of samples for which the modified geometric mean is negative. This never happened in the simulation, which involved the generation of a total of six thousand samples, suggesting that negative estimates of D are not a matter of practical concern. The fourth and final performance measure was the fraction of samples for which the modified geometric mean was closer to the true median D than the geometric mean. On this dimension, the geometric mean fares better in all cases: the geometric mean is closer to D than the modified geometric mean is for more than 50% of the samples in all six cases. Appendix 4C shows analytically that this last observation is not an aberration: the probability that the geometric mean is closer than the modified geometric mean is to the true value D is always greater than 50%.

This means that for most samples, the geometric mean does a little bit better than the modified version, but occasionally, the geometric mean will be far worse than the modified geometric mean. Thus, the modified geometric mean provides insurance against occasional bad behavior; we use it in chapter 5 to estimate baseline and revised death rates.

Sample size	$D = 10^{-6}$	$D = 10^{-3}$
$n = 20$	2.4% reduction in bias 2.9% reduction in rmsd no negative estimates 46.9% of estimates closer to true value	2.6% reduction in bias 3.1% reduction in rmsd no negative estimates 47.5% of estimates closer to true value
$n = 10$	3.6% reduction in bias 5.6% reduction in rmsd no negative estimates 46.8% of estimates closer to true value	1.3% reduction in bias 5.0% reduction in rmsd no negative estimates 45.6% of estimates closer to true value
$n = 5$	10.6% reduction in bias 11.3% reduction in rmsd no negative estimates 45.9% of estimates closer to true value	6.3% reduction in bias 10.5% reduction in rmsd no negative estimates 45.1% of estimates closer to true value

Table 4.1: Simulation results comparing the performance of the geometric mean $\hat{D}_1$ and the modified geometric mean $\hat{D}_2$ as an estimator of the median D of a lognormal distribution. Four performance measures are shown: (1) percentage reduction in bias achieved by using the modified geometric mean, (2) percentage reduction in root mean squared deviation from D achieved by using the modified geometric mean, (3) fraction of samples for which the modified geometric mean was negative, and (4) fraction of samples for which the modified geometric mean is closer to the true median than the geometric mean is.

4.5 Estimation of the Parameters in the Propagation Model

In this section we provide a detailed description of how the parameters in the differential equation 4.5 were estimated using expert estimates. Implicit in the methods to be described is a model of how the experts generate their estimates and we will attempt to justify this model. First, let us establish a notation for the expert estimates. Let

$$A_{ij\tau} \equiv \begin{array}{l}\text{Expert } j\text{'s estimate of the fraction of people in subpopulation}\\ i \text{ that would hear a warning that is first communicated to the}\\ \text{public } \tau \text{ minutes before the earthquake}\end{array}$$

We will use N_i to denote the set of experts that provide estimates for the i-th subpopulation and n_i to denote the number of experts in that set. Since each expert does not provide estimates for every subpopulation, we need this extra bit of notational complexity. We will use $L = \{0.5, 1, 5, 30\}$ to denote the set of lead-times τ that the experts were asked to provide estimates for.

Everything we do in this section will be subject to the assumption that the differential equation

$$\frac{dA}{dt} = U_1(t)(1 - A) + U_2(t)NA(1 - A)$$

and its solution $A(t)$ (which depends on the particular urgency functions U_1 and U_2 chosen) has been accepted as a valid model of how earthquake warnings propagate through a population. Basically, we will assume that the experts agree that the differential equation really does describe the propagation of warning information through a population. Of course the experts may not agree on this point, but it is important to recognize that the model-based approach implicitly makes this assumption.

Let Θ denote the vector of unknown parameters to be estimated. If the urgency functions are assumed constant, then Θ would equal (U_1, U_2); under the decreasing urgency case that was also discussed in this chapter Θ would equal (λ, γ). We will describe three approaches to estimating the unknown parameters: (1) least-squares, (2) maximum likelihood, and (3) Bayesian estimation.

4.5.1 Least-Squares Estimation

The least-squares approach is simple to describe. Let $A(\tau; \Theta)$ denote the solution to the differential equation when the values in Θ are substituted for the unknown parameters. Under the least-squares approach, the parameter estimates for subpopulation i were chosen as the values that minimize the sum of squares

$$\sum_{i \in N_i} \sum_{\tau \in L} (A_{ij\tau} - A(\tau; \Theta))^2$$

Before describing the maximum likelihood and Bayesian approaches we need to develop a likelihood function. Suppose we use $A_i(t)$ to denote the solution to the differential equation that is appropriate for the i-th subpopulation. Then the expert estimate $A_{ij\tau}$ is an estimate of $A_i(\tau)$. We need to specify, in probabilistic terms, exactly how we think the estimate $A_{ij\tau}$ and the "true value" $A_i(\tau)$ relate to each other. In essence, we need to postulate a model of how clairvoyant the experts are.

Our model is the following:

$$A_{ij\tau} = A_i(\tau\epsilon_{ij\tau}) \tag{4.14}$$

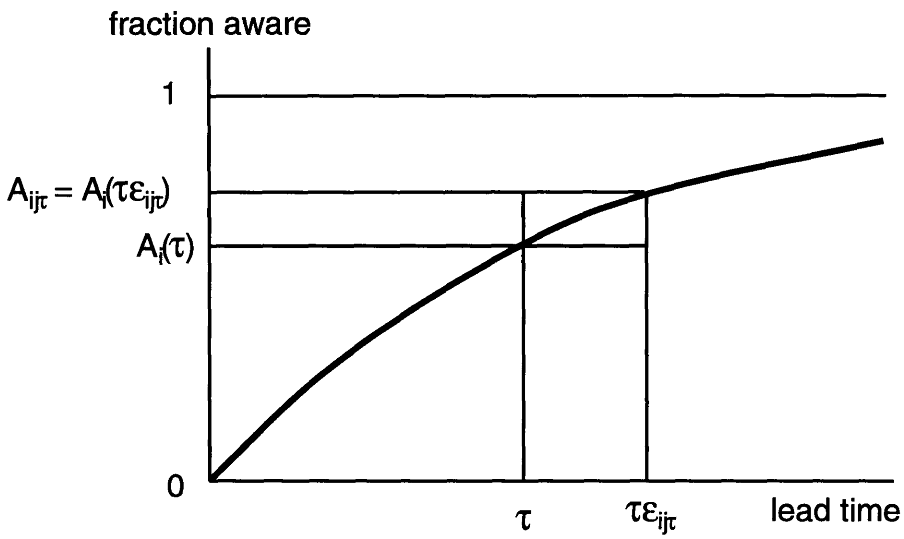

Figure 4-5: Illustration of the model relating the true fraction $A_i(\tau)$ of people aware of an earthquake warning by time τ to an expert's estimate $A_{ij\tau}$ of that fraction.

Figure 4-5 illustrates what this relation implies. Interpreted literally, this model says that the experts randomly perturb the lead-time τ, and then return the "true" value of A_i for the perturbed lead-time $\tau\epsilon_{ij\tau}$.

Of course, the experts do not do this. More reasonably, the experts might be expected to return the true value $A_i(\tau)$ corrupted by noise, i.e., $A_i(\tau)+\nu_{ij\tau}$, where $\nu_{ij\tau}$ is an error term. The error term arises because the experts have limited empirical evidence and imperfect theoretical models to base their judgments on. It would need to have the property that

$$0 \leq A_i(\tau) + \nu_{ij\tau} \leq 1$$

for all i, j, and τ. This requirement makes it difficult to find a parsimonious distribution for $\nu_{ij\tau}$; for one thing the distribution would have to depend on τ. But note that $A_i(\tau) + \nu_{ij\tau}$ equals the true value of A_i for *some* time value, which we might as well denote by $\tau\epsilon_{ij\tau}$. In this sense, it is immaterial whether one assumes that time is perturbed, or the value of $A_i(\tau)$ is perturbed. By assuming that time is perturbed, as we do, one can use a parsimonious distribution for the error term, while satisfying the following requirements:

$$\mathrm{E}[A_{ij\tau}] \approx A_i(\tau) \text{ and } \mathrm{var}[A_{ij\tau}] \to 0 \text{ as } \tau \to 0 \text{ or } \infty$$

The first requirement is that the model assume that the experts make correct judgments "on average," i.e., their judgments are at least approximately unbiased. Note that any explicit assumption about the expert judgments being systematically too high or too low would require the modeler to assume that he has more knowledge than the experts do. The second requirement enforces the initial condition $A_i(0) = 0$ associated with the differential equation and the asymptotic value $\lim_{t\to\infty} A_i(t) = 1$ of any solution to the differential equation that we are willing to consider.

The distributional assumption we will make is that $\epsilon_{ij\tau}$ is a lognormal random variable, i.e., $\log \epsilon_{ij\tau} \sim \mathrm{N}(0, \sigma^2)$. Thus, the median of random variable $\epsilon_{ij\tau}$ will equal one. This, combined with equation 4.14, implies that the median of the distribution we have assumed for the expert estimate $A_{ij\tau}$ will coincide with the true value $A_i(\tau)$. Thus, the first requirement above is satisfied, and it is not difficult to see that the second requirement is met also.

The parameter σ is a measure of how widely the expert estimates $A_{ij\tau}$ are scattered around the true values $A_i(\tau)$; it will need to be added to the vector of unknown parameters Θ.

There are various ways in which this model could be extended. One could attempt to model dependencies between experts, cognitive biases that the experts might have in estimating probabilities that are close to zero or one, or tendencies to choose round numbers, e.g., 0.01% rather than 0.000974%. However, all of these extensions require the introduction of additional parameters that need to be estimated using a limited amount of data. Therefore, we will limit ourselves to the model as stated, with the additional assumption of independence. Specifically, we assume that all estimates that $A_{ij\tau}$ that are provided for a particular subpopulation i, i.e., the random variables in the set $\{A_{ij\tau} : i \in N_i, \tau \in L\}$ are mutually independent.

The Likelihood Function

The likelihood function $L(\Theta)$ is the probability density function for the expert estimates $\{A_{ij\tau} : j \in N_i, \tau \in L\}$, evaluated at the observed values of the expert estimates. By our independence assumption, this probability density function is the product of the marginal density functions for the individual estimates $A_{ij\tau}$. We now compute the marginal density function for $A_{ij\tau}$, conditional on a specific value θ of the parameter vector Θ. First, we compute the cumulative distribution function for $A_{ij\tau}$. All probabilities are conditional on $\Theta = \theta$, although this is not always indicated explicitly.

$$\begin{aligned}
F_{A_{ij\tau}}(a|\Theta=\theta) &= \Pr\{A_{ij\tau} \leq a|\Theta=\theta\} = \Pr\{A_i(\tau\epsilon_{ij\tau}) \leq a|\} \\
&= \Pr\{\tau\epsilon_{ij\tau} \leq A_i^{-1}(a)\} \text{ (since } A \text{ is strictly increasing)} \\
&= \Pr\left\{\epsilon_{ij\tau} \leq \frac{A_i^{-1}(a)}{\tau}\right\} = \Pr\left\{\log \epsilon_{ij\tau} \leq \log \frac{A_i^{-1}(a)}{\tau}\right\} \\
&= \Phi\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\}
\end{aligned}$$

(σ_0 is the value taken on by σ in the vector θ) The likelihood function is obtained by differentiating with respect to a:

$$\begin{aligned}
l_{ij\tau}(\theta) &= f_{A_{ij\tau}}(a|\Theta=\theta) = \frac{d}{da}F_{A_{ij\tau}}(a) = \frac{d}{da}\Phi\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\} \\
&= \frac{d}{da}\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\}\phi\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\} \\
&= \frac{\tau}{\sigma_0 A_i^{-1}(a)}\frac{d}{da}A_i^{-1}(a)\phi\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\} \\
&= \frac{\tau}{\sigma_0 A_i^{-1}(a)\frac{dA}{dt}(A_i^{-1}(a))}\phi\left\{\frac{1}{\sigma_0}\log\frac{A_i^{-1}(a)}{\tau}\right\}
\end{aligned}$$

Here, Φ and ϕ are the cdf and pdf of a unit normal random variable, respectively. Thus, to compute the contribution of one expert estimate $A_{ij\tau}$ to the likelihood function for a particular value θ of the parameter vector Θ, one needs to be able to compute the solution $A_i(t;\theta)$ to the differential equation, it's inverse $A_i^{-1}(a;\theta)$, and the derivative of the solution $\frac{dA_i}{dt}(t;\theta)$.

The overall likelihood function for expert estimates given for subpopulation i will then be the product

$$L(\theta) = \prod_{j \in N_i} \prod_{\tau \in L} l_{ij\tau}(\theta)$$

The following algorithm can be used to compute the individual likelihood functions $l_{ij\tau}(\theta)$:

Step 1: Read the expert estimate $A_{ij\tau}$, the lead-time τ, and the parameter vector θ.

Step 2: Compute $t^* = A_i^{-1}(A_{ij\tau};\theta)$.

Step 3: Compute $f_1 = \frac{\tau}{\sigma_0 t^* \frac{dA}{dt}(t^*)}$ and $f_2 = \phi\left(\frac{1}{\sigma_0}\log\frac{t^*}{\tau}\right)$.

Step 4: Return $l_{ij\tau}(\theta) = f_1 f_2$.

Alternatively, one can compute the log-likelihood, to avoid overflow:

Step 1: Read the expert estimate $A_{ij\tau}$, the lead-time τ, and the parameter vector θ.

Step 2: Compute $t^* = A_i^{-1}(A_{ij\tau};\theta)$.

Step 3: Compute $\log f_1 = \log\tau - \log\sigma_0 - \log t^* - \log\frac{dA}{dt}(t^*)$ and $\log f_2 = \frac{1}{2}\left(\frac{1}{\sigma_0}\log\frac{t^*}{\tau}\right)^2$.

Step 4: Return $\log l_{ij\tau}(\theta) = \log f_1 + \log f_2$.

Now that we have specified the likelihood function, we can describe the two estimation methods that use it: maximum likelihood and Bayesian estimation.

4.5.2 Maximum Likelihood Estimation

The maximum likelihood estimation method chooses as parameter estimates the value of the unknown parameter vector Θ that maximizes the likelihood function $L(\Theta)$. We performed the maximization using a gradient search procedure; the resulting estimates are presented in an appendix to chapter 5.

4.5.3 Bayesian Estimation

If we are willing to treat the unknown parameters as random variables, with a prior distribution $p(\theta)$, then we can use Bayes law to compute the posterior distribution for the parameter vector, conditional on the expert estimates:

$$f_{\Theta|\{A_{ij\tau}: j\in N_i, \tau\in L\}}(\theta) = \frac{L(\theta)p(\theta)}{\int_\theta L(\theta)p(\theta)d\theta}$$

Point estimates of the unknown parameters can then be obtained by taking expectations of the posterior distribution.

Bayesian estimates were computed by discretizing the prior distribution over a rectangular grid. For example, for the constant urgency case, the parameter vector is (U_1, U_2, σ).

A uniform distribution was assumed over a three-dimensional grid (one dimension for each parameter). The value of the likelihood function was computed at each point in the grid, and the values were then normalized to produce the posterior distribution. The extent of the grid was adjusted by trial and error until it encompassed the region where the likelihood function took on a value substantially greater than zero, compared to the accuracy of the computer. The resulting parameter estimates are reported in an appendix to chapter 5.

Appendix 4A: Properties of the Homogeneous Mixing Model

In order for a solution $A(t)$ to the homogeneous mixing model of equation 4.5 to be interpreted as the fraction of the population aware of the warning, it should always be between zero and one. We show below that the assumptions we have made guarantee this. Since we have assumed that people do not forget the warning, we would expect $A(t)$ never to decrease, and this is also shown below. Furthermore, we show that the $A(t)$ approaches one as t increases if and only if the urgency functions U_1 and U_2 "do not decay too quickly" (this is stated more precisely below).

Proposition 1 *Suppose the urgency functions $U_1(t)$ and $U_2(t)$ are non-negative and bounded for $t \in [0, \infty)$ and suppose that $U_1(t)$ is strictly positive over some interval $[0, \epsilon]$ where $\epsilon > 0$. Then the solution $A(t)$ to the boundary value problem (4.5)-(4.6) has the following properties:*

1. $0 \leq A(t) \leq 1$.

2. *$A(t)$ is increasing.*

Furthermore, $\lim_{t\to\infty} A(t) = 1$ *if and only if*

$$\int_0^\infty (U_1(s) + U_2(s))ds = \infty \tag{4.15}$$

To interpret what the proposition says, one should interpret $A(t)$ as the fraction of *non-isolated* individuals that have heard the warning. What the proposition shows is that it is possible for $\lim_{t\to\infty} A(t)$ to be strictly less than one even if U_1 and U_2 are always positive. For example, if $U_1(t) = U_1(0)e^{-\gamma_1 t}$ and $U_2(t) = U_2(0)e^{-\gamma_2 t}$, where $\gamma_1, \gamma_2 > 0$, then $A(t)$ will not approach one.

Both the constant and decreasing urgency models that we consider have the desired property that $A(t)$ approaches one.

Proof of proposition 1: Consider the expression (4.9) for $w(t)$. All terms in the integral are non-negative by assumption, so we conclude that $w(t) \geq 0$ for all t. Also, we see that (4.7) implies that $w(t)$ is strictly positive for $t > 0$. Since $A(t) = w/(1+w)$, it follows that $0 \leq A(t) \leq 1$ for all t. Furthermore, $A(t)$ is strictly positive unless $w(t) = 0$, which occurs only at $t = 0$, and $A(t)$ is strictly less than one unless $w(t) = \infty$, which is not possible for any finite t. But then $1 - A(t)$ and $A(t)(1 - A(t))$ will also be bounded strictly between zero and one for $0 < t < \infty$, which implies that $dA/dt = U_1(1-A) + U_2NA(1-A)$ will be strictly positive for all $0 < t < \infty$. For $t = 0$, we have $dA/dt = U_1(0)(1 - A(0)) + U_2(0)NA(0)(1 - A(0)) = U_1(0) > 0$. Thus, $A(t)$ is increasing for all $t \geq 0$.

Suppose that 4.15 holds. Then it must be true that either $\int_0^\infty U_1 ds = \infty$ or $\int_0^\infty U_2 ds = \infty$; we will show that either condition implies that $A(t)$ has the desired property. If $\int_0^\infty U_1(s)ds = \infty$ is true, then we have

$$\begin{aligned}
\frac{dA(t)}{dt} &\geq U_1(t)(1 - A(t)) \\
\Rightarrow \frac{dA(t)/dt}{1 - A(t)} &\geq U_1(t) \qquad \text{(because } 1 - A(t) > 0 \text{ for } t \geq 0\text{)} \\
\Rightarrow \int_0^t \frac{dA(s)/ds}{1 - A(s)} ds &\geq \int_0^t U_1(s)ds
\end{aligned}$$

$$
\begin{aligned}
\Rightarrow -\ln(1-A(t)) &\geq \int_0^t U_1(s)ds \\
\Rightarrow 1-A(t) &\leq \exp\left(-\int_0^t U_1(s)ds\right) \\
\Rightarrow A(t) &\geq 1-\exp\left(-\int_0^t U_1(s)ds\right) \\
\Rightarrow \lim_{t\to\infty} A(t) &\geq 1
\end{aligned}
$$

Similarly, if $\int_0^\infty U_2(s)ds = \infty$ is true, then

$$
\begin{aligned}
\frac{dA(t)}{dt} &\geq U_2(t)A(t)(1-A(t)) \\
\Rightarrow \frac{dA(t)/dt}{A(t)(1-A(t))} &\geq U_2(t) \qquad \text{(for } t>0) \\
\Rightarrow \int_{0+}^t \frac{dA(s)/ds}{A(s)(1-A(s))}ds &\geq \int_{0+}^t U_2(s)ds \\
\Rightarrow \ln\frac{A(t)}{1-A(t)} &\geq \ln\frac{A(0^+)}{1-A(0^+)} + \int_{0+}^t U_2(s)ds \\
\Rightarrow \frac{A(t)}{1-A(t)} &\geq c'\exp\left(\int_{0+}^t U_1(s)ds\right) \qquad \text{(where } c' = \tfrac{A(0^+)}{1-A(0^+)} > 0 \text{ by (4.7))} \\
\Rightarrow A(t) &\geq \frac{c'\exp\left(\int_{0+}^t U_1(s)ds\right)}{1+c'\exp\left(\int_{0+}^t U_1(s)ds\right)} \\
\Rightarrow \lim_{t\to\infty} A(t) &\geq 1
\end{aligned}
$$

So in either case the limiting value of A is greater than or equal to one. But since $A(t) \leq 1$ for all t, it must be that $A(t)$ converges to one.

Conversely, suppose that 4.15 does not hold, i.e., both urgency functions decay quickly enough that their integrals are finite. Let $M = \int_0^\infty (U_1(r) + U_2(r)N)dt$. Taking limits in (4.9), we obtain

$$
\begin{aligned}
\lim_{t\to\infty} w(t) &= \int_0^\infty e^{\int_s^\infty (U_1(r)+U_2(r)N)dr} U_1(s)ds \\
&\leq \int_0^\infty e^{\int_0^\infty (U_1(r)+U_2(r)N)dr} U_1(s)ds \\
&= M\int_0^\infty U_1(s)ds \leq M^2 < \infty
\end{aligned}
$$

As a result,

$$
\lim_{t\to\infty} A(t) = \frac{\lim_{t\to\infty} w(t)}{1+\lim_{t\to\infty} w(t)} \leq \frac{M^2}{1+M^2} < 1
$$

■

Appendix 4B: A Model that Incorporates Spatial Dependence

As we mentioned at the beginning of this section, it is difficult to believe that the homogeneous mixing assumption will hold literally. We now present and analyze a model where the chances that two people meet and the warning message is transferred between them depend on the "distance" between the two individuals. The main implication of the analysis is that the differential equation 4.5 remains an appropriate description under certain departures from the homogeneous mixing assumption.

Consider a set of N individuals that at time $t = 0$ are not aware of the warning. Let $x_i(t)$ be an indicator variable for the i-th individual, i.e., $x_i(t)$ equals one if individual i is aware of the warning at time t and zero otherwise. The initial condition is that $x_i(0) = 0$ for all i.

We will model the vector $(x_1, x_2, \ldots, x_N)$ as a continuous time Markov process, whose state space is the set of all binary vectors with N elements. The only transitions that are possible are when a zero entry in the i-th position changes to a one, at the moment when individual i hears the warning. The chances that individual i hears the warning in the time interval $(t, t+h]$, given that he has not heard it yet (i.e., given $x_i(t) = 0$) will be assumed to equal

$$\tilde{U}_1 h + \tilde{U}_2 h \sum_{j=1}^{N} d_{ji} x_j(t) + o(h)$$

The first term ($\tilde{U}_1 h$) represents the chances that individual i hears the warning from the primary source. The second term is a sum of contributions from every individual j that has heard the warning at time t, so that $x_j(t) = 1$. For every such individual j, there is a probability $\tilde{U}_2 d_{ji} h$ that j communicates the warning to i in the time interval $(t, t+h]$. The "distances" d_{ji} are measures of how likely individuals i and j are to come into contact with each other that results in the warning message being communicated from j to i. We restrict d_{ii} to equal zero. The distance from j to i may not equal the distance from i to j. For example, i may know j's phone number but not vice versa.

One way to interpret this model is in terms of a collection of Poisson processes that are "turned on" and "turned off" at random times. Specifically, whenever $x_i(t) = 0$ (individual i has not yet heard the warning), there will be up to N Poisson processes that are turned on and directed at individual i: one for the primary communication source with rate $\tilde{U}_1$, and one for each individual j in the population that *does* know of the warning at time t, with rate $\tilde{U}_2 d_{ji}$. As soon as an arrival occurs for one of the Poisson processes directed at individual i, individual i's indicator variable switches from zero to one and the processes directed at him are turned off.

To relate this model to the differential equation 4.5 that we have been discussing, we will interpret the fraction $A(t)$ of people that have heard the warning by time t as $\mathrm{E}[\sum_1^N x_i(t)]/N$. We will derive a set of coupled differential equations for the expected values $\mathrm{E}[x_i(t)]$ and then we will show that under certain conditions this set of differential equations "collapses" into one differential equation of the form 4.5.

Let $e_i(t) \equiv \mathrm{E}[x_i(t)]$ and note that $e_i(t)$ is simply the probability that individual i has heard the warning by time t. By conditioning on whether individual i has heard the warning by time t or not we obtain

$$\begin{aligned} e_i(t+h) &= \sum_{k=0}^{1} \mathrm{E}[x_i(t+h)|x_i(t)=k]\Pr\{x_i(t)=k\} \\ &= e_i(t) + \Pr\{x_i(t+h)=1|x_i(t)=0\}(1-e_i(t)) \end{aligned} \tag{4.16}$$

We have already specified the conditional probability

$$\Pr\{x_i(t+h)=1|x_i(t)=0, x_j(t), j\neq i\} = \tilde{U}_1 h + \tilde{U}_2 h \sum_{j=1}^{N} d_{ji}x_j(t) + o(h)$$

By taking expectations over the states of every individual but individual i, we obtain

$$\Pr\{x_i(t+h)=1|x_i(t)=0\} = \tilde{U}_1 h + \tilde{U}_2 h \sum_{j=1}^{N} d_{ji}e_j(t) + o(h)$$

Upon substituting this expression in equation (4.16) and rearranging, we arrive at

$$\frac{e_i(t+h)-e_i(t)}{h} = \left(\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{ji}e_j(t) + \frac{o(h)}{h}\right)(1-e_i(t))$$

In the limit as h approaches zero, this becomes a differential equation:

$$\frac{de_i}{dt} = \left(\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{ji}e_j(t)\right)(1-e_i(t)) \text{ for } i=1,2,\ldots,N \tag{4.17}$$

When does this model reduce to the earlier one?

In the main text of this chapter, a model of information propagation wherein the only state variable was $A(t)=\sum_1^N e_i(t)/N$ (cf. equation (4.5)) was presented. A set of assumptions under which that model is reasonable was listed. One of the assumptions was that of homogeneous mixing of the population, which in the notation developed in this subsection may be expressed as $d_{ji}=1$ for all $i\neq j$. We will now state and prove a proposition that states that a substantially weaker condition is sufficient to justify the single differential equation model (4.5).

Proposition 2 *The following conditions (4.19) and (4.18) are equivalent.*

- $$d = \sum_{j=1}^{N} d_{ji} \textit{ for all } i=1,2,\ldots,N \tag{4.18}$$

  where d is a constant, independent of i.

- *The system (4.17) has the unique solution*

  $$e_i(t) = A(t) \textit{ for all } i=1,2,\ldots,N \textit{ and } t\geq 0 \tag{4.19}$$

  where $A(t)$ solves

  $$\dot{A} = (\tilde{U}_1 + \tilde{U}_2 dA)(1-A);\ A(0)=0 \tag{4.20}$$

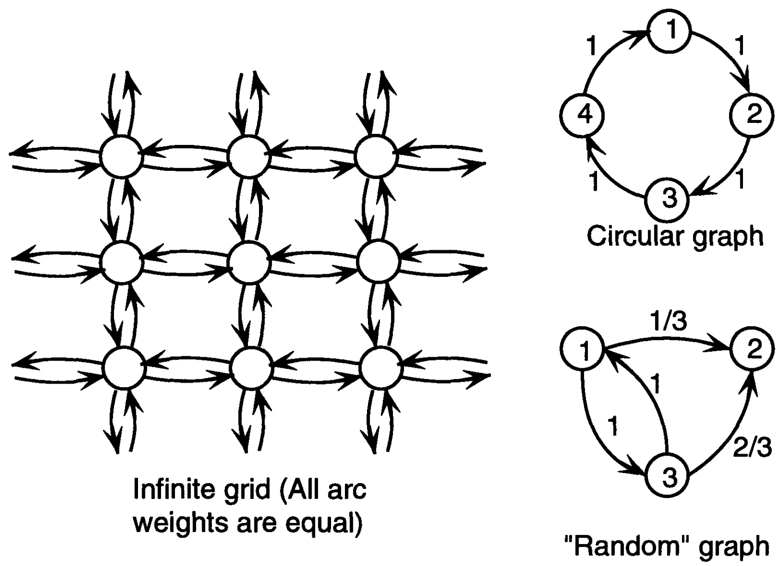

Figure 4-6: Three situations in which the spatial propagation model reduces to the homogeneous mixing model. Each node represents an individual, an arc from i to j means that i may tell j about the warning, and the arc weight is proportional to the probability that this occurs. In the infinite two-dimensional grid, four arcs enter each node (one from each of the four nearest neighbors). Since all arc weights on the infinite grid are equal, the sum of the arc weights entering any node is the same for all nodes. The same is true for the less regular circular and "random" graphs shown and this implies that the expected fraction of individuals that have heard the warning obeys the differential equation 4.5 originally proposed.

for some constant d.

Proposition 2 is proven at the end of this appendix. The necessary and sufficient condition (4.18) is best interpreted in graph theory terms. Imagine a graph with one node for each individual and arcs with weights d_{ji}. Then condition (4.18) simply requires that the sum of the weights of all arcs entering node i be the same for all nodes i. If this is the case, then the spatial propagation model formulated in this subsection reduces to the differential equation (4.5), in the sense that the expected fraction of individuals $\sum_1^N \mathrm{E}[x_i(t)]/N$ obeys the differential equation.

Figure 4-6 shows three examples for which condition (4.18) holds. It is perhaps not surprising that when there are infinitely many individuals, located on an infinite, spatially homogeneous grid, then the fraction of individuals that has heard the warning evolves as if the homogeneous mixing assumption held, i.e., it obeys the differential equation (4.5). It is less obvious that this will also be the case for the other two graphs shown in figure 4-6, but proposition 2 guarantees this.

The rest of this appendix is devoted to proving proposition 2. Before proving the proposition, we state and prove a uniqueness theorem for solutions of a system of first-order ordinary differential equations which is needed to prove proposition 2. Much more general results on the existence and uniqueness of solutions to boundary value problems are available (see, for instance, [CL55]) but the result stated below will suffice for our purposes.

For $x \in R^n$, let $|x|$ denote $\max_{i=1,2,\ldots,n} |x_i|$.

Lemma 1 *Suppose that $f : [a,b] \to R^n$ is differentiable on $[a,b]$, that $f(a) = 0$, and that for some constant A*

$$|\dot{f}(t)| \leq A|f(t)| \text{ for all } t \in [a,b]$$

Then $f(t) = 0$ for all $t \in [a,b]$.

Proof of the lemma: Fix $t_0 = \min(a + 1/(2A), b)$ and let

$$M_0 = \sup_{a \leq t \leq t_0} |f(t)| \text{ and } M_1 = \sup_{a \leq t \leq t_0} |\dot{f}(t)|$$

Then for any $t \in [a, t_0]$

$$\begin{aligned} |f(t)| &= \left| f(a) + \int_a^t \dot{f}(s)ds \right| = \left| \int_a^t \dot{f}(s)ds \right| \\ &\leq \int_a^t |\dot{f}(s)|ds \leq \int_a^{t_0} |\dot{f}(s)|ds \\ &\leq M_1(t_0 - a) \end{aligned}$$

Also, by the hypothesis of the lemma we have $M_1 \leq AM_0$, so $|f(t)| \leq AM_0(t_0 - a)$ for all $t \in [a, t_0]$, which implies

$$M_0 \leq AM_0(t_0 - a)$$

But

$$A(t_0 - a) = A\left(\min(a + 1/(2A), b) - a\right) \leq A(a + 1/(2A) - a) = 1/2$$

which implies $M_0 = 0$, so $f(t)$ must equal zero on $[a, t_0]$.

Next, fix $t_1 = t_0 + 1/(2A)$. The same argument as before, with M_0 and M_1 now redefined to be supremums over $[t_0, t_1]$, shows that M_0 equals zero and $f(t) = 0$ on $[t_0, t_1]$. Continue this, with $t_i = \min(a + (i+1)/(2A), b)$, until $t_i = b$. At each step, the preceding argument is used to show that $f(t) = 0$ on $[t_{i-1}, t_i]$, so the conclusion is that $f(t) = 0$ on $[a, t_0] \cup [t_0, t_1] \cup \cdots \cup [t_{i-1}, b] = [a, b]$. ■

The lemma will be used to prove the following uniqueness theorem:

Theorem 1 *Consider the initial value problem*

$$\dot{x} = g(x), \ x(0) = c$$

where $g : E \to R^n$, with $E = \{x : a_i \leq x_i \leq b_i, \ i = 1, 2, \ldots, n\}$. Suppose there is a constant M such that

$$|g(x_1) - g(x_2)| \leq M|x_1 - x_2|$$

whenever $x_1 \in E$ and $x_2 \in E$. Then the initial value problem can have at most one solution on $t \in [0, b]$, for any $b > 0$.

Proof of the theorem: Suppose that $\alpha(t)$ and $\beta(t)$ are two solutions of the initial value problem. Let $\Delta(t) = \alpha(t) - \beta(t)$. Then Δ is differentiable (since α and β must be) and

$$\dot{\Delta}(t) = \dot{\alpha}(t) - \dot{\beta}(t) = g(\alpha(t)) - g(\beta(t))$$

By the hypothesis of the theorem, we have

$$|\dot{\Delta}(t)| = |g(\alpha(t)) - g(\beta(t))| \leq M|\alpha(t) - \beta(t)| = M|\Delta(t)|$$

Since $\Delta(0) = \alpha(0) - \beta(0) = c - c = 0$, we see that the function Δ satisfies the conditions stated in the lemma, so according to the lemma, $\Delta(t) = 0$ for all $t \in [0, b]$ which implies that $\alpha(t) = \beta(t)$ for $t \in [0, b]$, for any positive, finite b. ■

Proof of proposition 2 : Suppose that (4.19) holds. Substituting A for e_i and e_k in the system (4.17) results in

$$\dot{A} = (\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{ji} A)(1 - A) = (\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{jk} A)(1 - A) = (\tilde{U}_1 + \tilde{U}_2 d A)(1 - A)$$

$$\Rightarrow \sum_{j=1}^{N} d_{ji} A(1 - A) = \sum_{j=1}^{N} d_{jk} A(1 - A) = dA(1 - A)$$

This equation holds for all $t \in [0, \infty)$. Since $0 < A(t)(1 - A(t)) < 1$ for $t > 0$ (by proposition 1), we can pick a $t > 0$, and divide through the above string of equalities with $A(t)(1 - A(t))$, to get

$$d = \sum_{j=1}^{N} d_{ji} = \sum_{j=1}^{N} d_{jk}$$

for all i and k.

Next, suppose that (4.18) holds. Then it is easily verified that one solution to (4.17) is obtained by setting $e_i(t) = A(t)$, where $A(t)$ is a solution to (4.20). Furthermore, the system (4.17) satisfies a Lipschitz condition for $(e_1, e_2, \ldots, e_N) \in E = [0, 1]^n$. More precisely, define a function $f : E \to R^N$ by

$$f_i(e_1, e_2, \ldots, e_N) = de_i/dt = (\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{ji} e_j)(1 - e_i)$$

so that

$$\frac{\partial f_i}{\partial e_j} = \begin{cases} \tilde{U}_2 d_{ji}(1 - e_i) & \text{for } j \neq i \\ -(\tilde{U}_1 + \tilde{U}_2 \sum_{j=1}^{N} d_{ji} e_j) & \text{for } j = i \end{cases}$$

which implies that

$$\left| \frac{\partial f_i}{\partial e_j} \right| \leq \tilde{U}_1 + \tilde{U}_2 n \max_{i,j} d_{ij} < \infty$$

for all $(e_1, e_2, \ldots, e_N) \in E$. Therefore, theorem 1 applies and the solution to (4.17) is unique. ■

Appendix 4C: Estimating the Median of a Lognormal distribution

Suppose that $Y_1, Y_2, \ldots, Y_n$ are i.i.d. lognormal random variables with parameters μ and σ. Let $X_i \equiv \log Y_i$. Since the Y_i's are i.i.d. lognormal with parameters μ and σ, the X_i's are i.i.d. $\mathrm{N}(\mu, \sigma^2)$. The median of the lognormal distribution for Y_i is $D = e^{\mu}$. We wish to estimate D.

The Geometric Mean

Let $\bar{X}$ be the sample mean of the $\{X_i\}$ and let $S^2 = n^{-1}\sum_1^n (X_i - \bar{X})^2$ (the maximum likelihood estimator of σ^2). Since $\bar{X}$ is the maximum likelihood estimator of μ, the maximum likelihood estimator of $e^{\mu} = D$ is

$$\hat{D}_1 = e^{\bar{X}} = e^{n^{-1}\sum_1^n \log Y_i} = \left(\prod_1^n Y_i\right)^{1/n}$$

i.e., the geometric mean of the $\{Y_i\}$. The mean and variance of $\hat{D}_1$ can be computed using the moment generating function of $\bar{X}$. Recall that the moment generating function of a $\mathrm{N}(\mu, \sigma^2)$ random variable X is

$$M_X(s) = \mathrm{E}\left[e^{sX}\right] = e^{s\mu + s^2\sigma^2/2}$$

Therefore,

$$\mathrm{E}[\hat{D}_1] = \mathrm{E}\left[e^{\bar{X}}\right] = M_{\bar{X}}(1) = e^{\mu + \sigma^2/(2n)} = De^{\frac{\sigma^2}{2n}}$$

Thus, $\hat{D}_1$ is always positively biased by a factor of $e^{\sigma^2/(2n)}$. As n increases, this factor converges to one, so $\hat{D}_1$ is asymptotically unbiased. The second moment of $\hat{D}_1$ equals the moment generating function of $\bar{X}$ evaluated at $s = 2$:

$$\mathrm{E}[\hat{D}_1^2] = \mathrm{E}\left[e^{2\bar{X}}\right] = M_{\bar{X}}(2) = e^{2\mu + 2^2\sigma^2/(2n)} = D^2 e^{2\sigma^2/n}$$

Thus, the variance of $\hat{D}_1$ is

$$\mathrm{var}[\hat{D}_1] = D^2 e^{2\sigma^2/n} - \left(De^{\sigma^2/(2n)}\right)^2 = D^2 e^{\sigma^2/n}\left(e^{\sigma^2/n} - 1\right) \sim \frac{D^2\sigma^2}{n}$$

A Modification of the Geometric Mean

Suppose we modify the maximum likelihood estimator of D by multiplying it with some function f_n of S^2 that may depend on the sample size n, i.e.,

$$\hat{D}_2 = \hat{D}_1 f_n(S^2) = e^{\bar{X}} f_n(S^2)$$

The motivation for doing this is that it is possible to construct an unbiased estimator of the mean of a lognormal random variable which has this form. We now try to do the same for the median of a lognormal random variable . Since $\bar{X}$ and S^2 are independent, the expected value of $\hat{D}_2$ can be written as

$$\mathrm{E}[\hat{D}_2] = \mathrm{E}\left[e^{\bar{X}}\right]\mathrm{E}[f_n(S^2)] = De^{\sigma^2/(2n)}\mathrm{E}[f_n(S^2)]$$

For unbiasedness, we require that $\mathrm{E}[f_n(S^2)] = \exp\{-\sigma^2/(2n)\}$.

Suppose that f_n can be represented as a power series around $S^2 = 0$, i.e., $f_n(S^2) = \sum_0^\infty a_i(S^2)^i$. Then

$$\mathrm{E}[f_n(S^2)] = \sum_0^\infty a_i\mathrm{E}[(S^2)^i]$$

To compute the moments of S^2 we use the fact that $nS^2/\sigma^2 \sim \chi^2_{n-1}$. Let u be a χ^2_ν random variable. Its moment generating function is $(1-2s)^{-\nu/2}$; the i-th derivative of this function is

$$\frac{d^i}{ds^i}M_u(s) = \prod_{j=0}^{i-1}(\nu+2j)(1-2s)^{-\nu/2-i}$$

The moments of u follow:

$$\mathrm{E}[u^i] = \left.\frac{d^i}{ds^i}M_u(s)\right|_{s=0} = \prod_{j=0}^{i-1}(\nu+2j) = \frac{2^i\Gamma(\nu/2+i)}{\Gamma(\nu/2)}$$

The i-th moment of S^2 is then

$$\mathrm{E}[(S^2)^i] = (\sigma^2/n)^i\mathrm{E}[(nS^2/\sigma^2)^i] = (\sigma^2/n)^i\frac{2^i\Gamma(\frac{n-1}{2}+i)}{\Gamma(\frac{n-1}{2})}$$

We require $\mathrm{E}[f_n(S^2)] = \exp(-\sigma^2/(2n))$. Both sides can be represented as a power series in σ^2:

$$\sum_0^\infty a_i\left(\frac{2\sigma^2}{n}\right)^i\frac{\Gamma(\frac{n-1}{2}+i)}{\Gamma(\frac{n-1}{2})} = \sum_0^\infty\frac{1}{i!}\left(\frac{-\sigma^2}{2n}\right)^i$$

Equating coefficients we obtain an expression for a_i:

$$a_i = (-1)^i\left(\frac{\sigma^2}{2n}\right)^i\left(\frac{n}{2\sigma^2}\right)^i\frac{\Gamma(\frac{n-1}{2})}{i!\Gamma(\frac{n-1}{2}+i)} = \left(-\frac{1}{4}\right)^i\frac{\Gamma(\frac{n-1}{2})}{i!\Gamma(\frac{n-1}{2}+i)}$$

Finally,

$$f_n(S^2) = \Gamma\left(\frac{n-1}{2}\right)\sum_{i=0}^\infty\frac{(-S^2/4)^i}{i!\Gamma(\frac{n-1}{2}+i)} = \frac{\Gamma(\frac{n-1}{2})}{(S/2)^{(n-3)/2}}J_{(n-3)/2}(S)$$

where $J_m(z)$ is a Bessel function. We will refer to the estimator $\hat{D}_2$ as the modified geometric mean.

Properties of the Modified Geometric Mean

The modified geometric mean is unbiased by construction. We now show that it is the minimum variance unbiased estimator of the median D.

First, we show that $\mathbf{T}(\mathbf{Y}) = (\bar{X}, S^2)$ is a sufficient statistic for the sample $\mathbf{Y} =$

$(Y_1, Y_2, \ldots, Y_n)$. The joint pdf of the sample is

$$\begin{aligned}
f(\mathbf{y}) &= \prod_1^n \frac{1}{y_i \sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\frac{\log y_i - \mu}{\sigma}\right)^2\right\} \\
&= \frac{1}{(\sigma^2 2\pi)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}\sum_1^n (x_i - \mu)^2\right\} \prod_1^n \frac{1}{y_i} \\
&= \frac{1}{(\sigma^2 2\pi)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}\left(ns^2 + n(\bar{x} - \mu)^2\right)\right\} \prod_1^n \frac{1}{y_i} \\
&= h(\mathbf{T}(\mathbf{y}); \mu, \sigma^2) k(\mathbf{y})
\end{aligned}$$

where $k(\mathbf{y}) = 1/\prod_1^n y_i$, and h has the obvious definition. The factorization criterion [KJR82, page 73, volume 8] now implies that $\mathbf{T}$ is a sufficient statistic. Since the statistic $\mathbf{T}'(\mathbf{Y}) = (\sum_1^n X_i^2, \sum_1^n X_i)$ is an invertible function of the sufficient statistic $\mathbf{T}(\mathbf{Y}) = (\bar{X}, S^2)$, $\mathbf{T}'$ is a sufficient statistic also. This will be useful in the following.

Second, we show that $\mathbf{T}'$ is a complete sufficient statistic. To do this, we again manipulate the joint pdf of the sample:

$$\begin{aligned}
f(\mathbf{y}) &= \prod_1^n \frac{1}{y_i \sigma\sqrt{2\pi}} \exp\left\{-\frac{1}{2}\left(\frac{\log y_i - \mu}{\sigma}\right)^2\right\} \\
&= \frac{1}{(\sigma^2 2\pi)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}\sum_1^n (x_i - \mu)^2\right\} \prod_1^n \frac{1}{y_i} \\
&= \frac{1}{(\sigma^2 2\pi)^{n/2}} \exp\left\{-\frac{1}{2\sigma^2}\left(\sum_1^n x_i^2 - 2\mu\sum_1^n x_i + n\mu^2\right)\right\} \prod_1^n \frac{1}{y_i} \\
&= k(\mathbf{y}) p(\mu, \sigma^2) \exp\left\{\mathbf{c}(\mu, \sigma^2)\mathbf{T}'(\mathbf{y})^T\right\}
\end{aligned}$$

where k is as before,

$$p(\mu, \sigma^2) = (\sigma^2 2\pi)^{-n/2} \exp\left(-\frac{n\mu^2}{2\sigma^2}\right)$$

and

$$\mathbf{c}(\mu, \sigma^2) = (-1/(2\sigma^2), \mu/\sigma^2)$$

The set spanned by $\mathbf{c}$ as the parameters μ and σ^2 range over their allowable values is

$$\{\mathbf{c}(\mu, \sigma^2) : -\infty < \mu < \infty, \sigma^2 > 0\} = (-\infty, 0) \times (-\infty, \infty)$$

i.e., the left half of the plane. This set contains an open set and this implies, using "Exponential criterion II" from [KJR82, page 77, volume 8], that $\mathbf{T}'$ is complete sufficient. Since $\mathbf{T}$ is an invertible function of $\mathbf{T}'$, $\mathbf{T}$ is also a complete sufficient statistic.

Third, we note that the estimator $\hat{D}_2$ is a function of the complete sufficient statistic $\mathbf{T}$, since it depends on the sample only through $\bar{X}$ and S^2, and it is an unbiased estimator of D. These two facts guarantee, via a theorem due to Lehmann and Scheffé [KJR82, page 76, volume 8], that $\hat{D}_2$ is the minimum variance unbiased estimator of D, and furthermore, for any convex loss function, $\hat{D}_2$ minimizes the risk among all unbiased estimators.

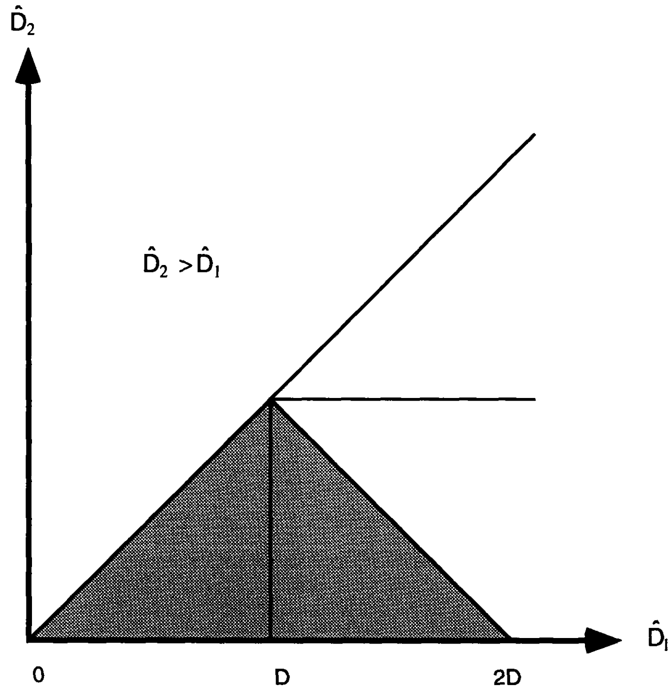

Figure 4-7: The shaded region represents the event that the modified geometric mean $\hat{D}_2$ is closer than the geometric mean $\hat{D}_1$ is to the true median D. Note that we need not consider the region where $\hat{D}_2 > \hat{D}_1$, since this never happens. The region extends below the horizontal axis, since $\hat{D}_2$ can be negative, i.e., the event can be written as $\{\hat{D}_1 + \hat{D}_2 \leq 2D\} \cap \{\hat{D}_1 \geq 0\} \cap \{\hat{D}_2 \leq \hat{D}_1\}$.

We end by showing analytically that

$$\Pr\{|\hat{D}_2 - D| > |\hat{D}_1 - D|\} > 1/2$$

i.e., the chances that $\hat{D}_1$ is closer to D than $\hat{D}_2$ are always better than even. From figure 4-7 we see that the event $|\hat{D}_2 - D| > |\hat{D}_1 - D|$ is equivalent to $\hat{D}_1 + \hat{D}_2 \leq 2D$. Therefore,

$$\begin{aligned}
\Pr\{|\hat{D}_2 - D| > |\hat{D}_1 - D|\} &= \Pr\{\hat{D}_1 + \hat{D}_2 \leq 2D\} \\
&= \Pr\{\exp(\bar{X})(1 + f_n(S^2)) \leq 2D\} \\
&= \Pr\left\{\exp(\bar{X}) \leq \exp(\mu)\frac{2}{1 + f_n(S^2)}\right\} \\
&= \Pr\{\bar{X} \leq \mu + \log 2 - \log(1 + f_n(S^2))\}
\end{aligned}$$

We know that $f_n(S^2) < 1$, and this implies that $\log 2 - \log(1 + f_n(S^2)) > 0$. Thus, the probability that $\hat{D}_1$ is closer than $\hat{D}_2$ is to D equals the probability that $\bar{X}$ is less than a quantity which is greater than μ, i.e., it is the probability that a normal random variable is less than a quantity larger than its mean, so the probability must be greater than one half.

Chapter 5

Results of the Survey of Experts

We now turn to analyzing the expert responses to the questionnaire described in chapter 3. The aim of the analysis is to estimate the death rate resulting from an unanticipated major earthquake for a typical U.S. urban area, and the potential reduction in death rates if earthquake warnings were available.

5.1 Introduction

We start by reviewing the quantities the experts were asked to estimate. The experts were instructed to estimate the following quantities for people that live in urban areas in California, since they are likely to be most familiar with earthquake hazard in that setting. Also, the likely effects of earthquakes that strike Los Angeles or San Francisco *without warning* have been the subject of considerable study.

5.1.1 The Questionnaire Data

As discussed in chapter 3, the experts provided estimates of the following quantities, for each of four lead-times τ (30 seconds, 1 minute, 5 minutes, and 30 minutes) and each[1] of eight subpopulations i:

A: The fraction of a subpopulation that becomes aware of the warning before the earthquake occurs. The experts were asked to estimate this fraction for several communication strategies suggested in the questionnaire and were also encouraged to suggest additional strategies. The analysis in this chapter will use estimates that treat sirens as the sole means of communicating the warning. We will also show how the results change if the most effective communication strategy (as judged by each expert) is used.

C: The chances that a person who hears the warning and attempts to complete the best action, under the circumstances, is able to complete that action before the earthquake occurs. As discussed in chapter 3, the experts chose what they considered the best action, e.g., getting under a sturdy desk (if inside) or getting out of the building.

D'': The **revised** death rate, i.e., the probability that a person who completes the best action dies nonetheless.

[1]Individual experts did not provide estimates for all eight subpopulations, see chapter 3.

D': The **baseline** death rate, i.e., the probability that a randomly chosen person (whose revised death rate would be D'') would die in an earthquake that occurs without warning.

The overall death rate D (again, for a particular subpopulation and lead-time) can be computed as

$$D = ACD'' + (1 - AC)D'$$

This computation implicitly makes the optimistic assumptions that were discussed in chapter 4, namely that as soon as a person becomes aware of the warning, she will strive to improve her chances of survival by attempting the best action, that all persons who complete the best action before the quake strikes face the revised death rate D'', and that all persons who either do not attempt the best action (because they did not hear the warning soon enough) or are not able to complete the best action in time face the baseline death rate D'. Note that we assume that all people who attempt an action take the *best* action, under the circumstances. Under these assumptions AC is the fraction of the subpopulation that completes the best action and faces the revised death rate D''; the remaining fraction $1 - AC$ faces the baseline death rate D', and the formula above follows.

More precisely the data may be thought of as a three-dimensional array:

$$\begin{aligned} D_{ij\tau} &\equiv j\text{-th expert's estimate of death rate in } i\text{-th group, given lead time } \tau \\ &= A_{ij\tau}C_{ij\tau}D''_{ij\tau} + (1 - A_{ij\tau}C_{ij\tau})D'_{ij} \end{aligned}$$

The experts were instructed to worry more about how much the death rate might be reduced with positive lead-times than about how accurate their baseline death rate estimates were. Therefore, we decompose the rates for positive lead-times into the product of the *relative* death rate

$$G_{ij\tau} \equiv \frac{D_{ij\tau}}{D_{i,j,0}}$$

and the baseline death rate, i.e., $D_{ij\tau} = G_{ij\tau}D_{i,j,0}$. The relative death rates $G_{ij\tau}$ measure the fraction of risk that remains even if τ minutes of warning are available. An equivalent but perhaps more optimistic measure is the proportional risk reduction factor

$$R_{ij\tau} \equiv 1 - G_{ij\tau}$$

which measures what percentage of the baseline risk level can be avoided with τ minutes of warning.

In the next section, we discuss the two approaches to a statistical analysis of the expert responses that we will pursue (the "face-value" and "model-based" approaches) and reasons for considering both approaches. The third and fourth sections present the results of the face-value and model-based approaches, respectively, and the fifth section compares them. The sixth section assesses the consequences of some of the optimistic assumptions that we have adopted and section seven summarizes the findings of this chapter.

5.2 Two Uses of Expert Opinion

We will process the expert assessments two ways. First, we will essentially take the responses at **face-value** and combine them into point estimates with associated confidence intervals. Second, we will analyze the same data under additional assumptions about how humans react to earthquake warnings and about how the experts form their judgment; this is our **model-based** approach.

The two approaches treat the data in entirely different ways. For example, under the face-value approach, we will only compute estimates for the four lead-times (30 seconds, 1 minute, 5 minutes, and 30 minutes) that the experts were asked to provide estimates for; except for straight lines connecting the estimates we will not attempt to extrapolate to other lead-times. In contrast, the assumptions we make under the model-based approach will allow us to estimate the expected death rate $D(\tau)$ for any lead-time. Depending on one's point of view, this could be viewed as a strength or a weakness. Our hope is that the strengths of the two approaches will complement each other so that their results, when viewed together, will provide a more complete picture of the consensus view of the experts, or else indicate a lack of consensus.

5.3 Face-Value Analysis

The analysis in this section will proceed as follows. First, the baseline death rates for each subpopulation will be estimated using the experts' judgment. In chapter 6, the resulting estimates will be compared to historical data and estimates from other sources. The estimates will also be combined into a baseline death rate estimate for the overall population, using estimates of how the population is distributed across the subpopulations at various times of day. Next, we turn our attention to the potential reduction in the baseline rates that can be achieved. After examining frequency distributions for the relative death rates for each subpopulation and lead time, we construct confidence intervals for the overall (across all subpopulations) death rate reduction – based on the expert assessment – as a function of lead-time.

Then, we assess how sensitive our estimates are to (1) perturbations to the estimated population distribution and (2) assumptions about how the warning is communicated.

The section concludes by examining whether the expert estimates provide evidence in support of the notion that the group of experts is heterogeneous, with some being consistently optimistic while others are consistently pessimistic.

5.3.1 Baseline Death Rates

Figure 5-1 shows histograms of the baseline death rates for each subpopulation (for example "People at home, awake"), as estimated by the experts. The horizontal axes in the figure are logarithmic, with a death rate of one in a hundred equidistant from one in ten and from one in a thousand. Thus, each histogram shows the frequency distribution of $\log_{10} D_{i,j,0}$ over all experts j, with the subpopulation i fixed. Since the baseline estimates are generally quite close to zero, a logarithmic scale seems more illuminating than a linear one.

The expert estimates cover the range from one in ten to one in a million for all subpopulations except "People in offices" and "People in cars on or near a bridge or overpass," where the lowest estimates are around one in ten thousand.

The expert assessments of the baseline death rate, for each subpopulation, were combined into one estimate using the modified geometric mean that was developed in chapter 4; the resulting estimates are shown in figure 5-2. According to these estimates, the safest place to be in an earthquake is in a car on the open freeway, where the death rate is only 3 per 100,000. In contrast, being in a car on or near a freeway overpass is judged to be about 70 times more hazardous, for a death rate of 203 per 100,000. Estimates for the remaining

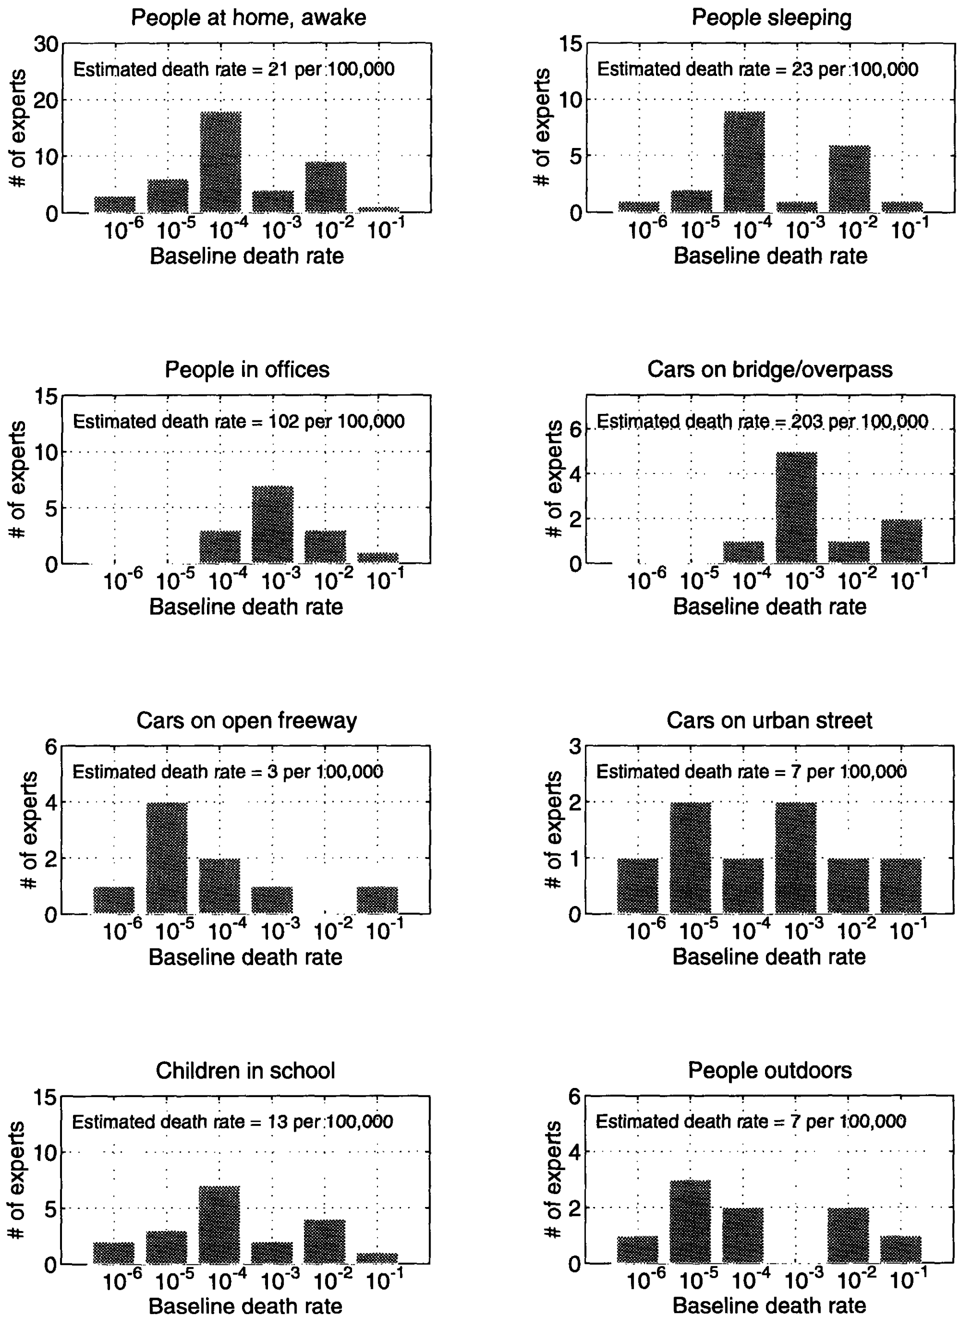

Figure 5-1: Histograms of expert estimates of the baseline death rate for the eight subpopulations. Note that the horizontal scale is logarithmic. The estimated death rates are computed using the modified geometric mean described in chapter 4.

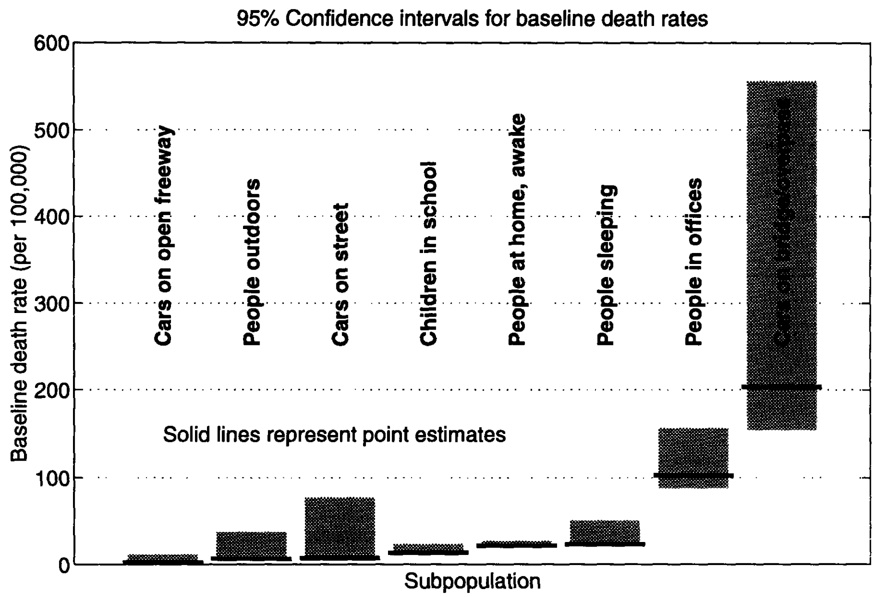

Figure 5-2: Baseline death rate estimates (solid lines) with two-sided 95% confidence limits (shaded boxes) for each subpopulation.

six subpopulations lie between these two extremes. For people in their homes, the death rate for those that are asleep is estimated to be only 10% higher (23 per 100,000) than for those that are awake (21 per 100,000). This does not negate the often-made assertion that lower death tolls can be expected for earthquakes that occur at night[2], when most people are asleep; for during the day many people will not be in the relative safety of their homes. Instead, many of them will be working in offices, where the estimated death rate is 102 per 100,000. School children are judged to be relatively safe (estimated death rate 13 per 100,000), probably because strict construction codes for schools have been in effect in California since the 1930's (the 1933 Field Act, see for example [Bol93, page 258] and [BHMS75, page 42]). People in cars on urban streets (death rate 7 per 100,000) are judged to be less safe than people in cars on the open freeway but safer than those on or near bridges or overpasses. Being outdoors is judged to be the second-safest place to be (tied with cars on urban street), with a death rate of 7 per 100,000.

Obviously, these estimates of baseline death rates should not be taken as definitive. One of the caveats one should keep in mind is the small number of experts that provide estimates for some of the subpopulations. To see whether this particular concern is of practical significance, one is motivated to estimate the sampling variation around the estimated baseline death rates. This was done by computing standard confidence intervals (of the form $\bar{X} \pm t_{\alpha/2,n-1}S/\sqrt{n}$) for the logarithmically transformed death rates (cf. figure 5-1) and then transforming back to the original scale. The resulting upper and lower 95% confidence

[2]For example, a preliminary reconnaissance report of the 1994 Northridge earthquake, which occurred at 4:31 AM, states that "If the earthquake had occurred during the day there would have been terrible casualties in stores, shopping centers and parking structures" [Ear94, page 32] and "The time of day the earthquake occurred – 4:31 in the morning on a holiday – clearly contributed to the relatively low death toll" [Ear94, page 87].

limits are shown in figure 5-2. The intervals differ in width because the number of experts that provided estimates differs between subpopulations (note the narrow confidence interval for people at home, awake, a subpopulation that every expert provided estimates for) and because the sample variances differ between subpopulations.

A National Research Council panel [Pan89] on earthquake loss estimation, in reviewing casualty estimates for future earthquakes for several U.S. urban areas (see chapters 3 and 6) suggests that

> Estimates of casualties ... should be regarded as having an order of magnitude (factor of 10) uncertainty ... These obviously are ... matters for which far more data from actual earthquakes are required to advance the state of the art.

In light of this quote, the confidence limits in figure 5-2 are surprisingly narrow. For the "People at home, awake" subpopulation, the upper 95% confidence limit is only 30% greater than the lower limit. None of the upper confidence limits are more than a factor of eleven greater than the corresponding lower confidence limit. The factor of eleven is for the "People in cars on urban street" subpopulation; the remaining subpopulations have upper limits that are no more than seven times greater than the lower limits.

Not only do the experts agree to a surprising degree on the absolute magnitude of the baseline death rates; the ranking of the subpopulations in terms of baseline death rates also appears relatively stable. For example, the confidence interval for the subpopulation with the highest estimated death rate (people in cars on or near overpasses) overlaps only slightly with the interval for the runner up (people in offices) and not at all with the other six intervals. Of the nine experts[3] that provided estimates for "cars on bridge/overpass," seven considered that to be the riskiest (or tied for riskiest) place to be among the settings considered by those experts (none of those experts were asked to consider "people in offices"). This is added evidence that the experts generally believe freeway overpasses to be the most dangerous places to be during an earthquake. Similarly, the confidence interval for people in cars on the open freeway, which has the lowest estimated death rate, overlaps with only two subpopulations (people outdoors and people in cars on urban streets). Of the nine experts asked about "cars on the open freeway," seven considered it the safest (or tied for safest) of the categories considered by those experts.

Of course, the experts could be wrong even though they agree. Skeptics might hypothesize that the experts are simply expressing the conventional wisdom of the community of scientists and professionals that concern themselves with earthquake hazards. Under this hypothesis, the narrow confidence intervals simply reflect effective communication and substantial agreement among people in this group. In particular, it does not mean that the estimates are right, in the sense of being close to the death rates that can be expected in a future major earthquake, i.e., the conventional wisdom might be wrong.

While this hypothesis is plausible in some respects, it may overstate the case. The experts that responded to the questionnaire are intelligent people, many of whom are acutely aware of the dangers of putting faith in a number simply because it has been published. Some of them have expressed concerns about an issue closely related to the validity of the hypothesis above: That risk estimates based on limited information become widely known and start having a life of their own.

[3]These nine experts all received the same version of the questionnaire, which asked about "cars on bridge/overpass," "cars on open freeway," "cars on street," and "people at home, awake."

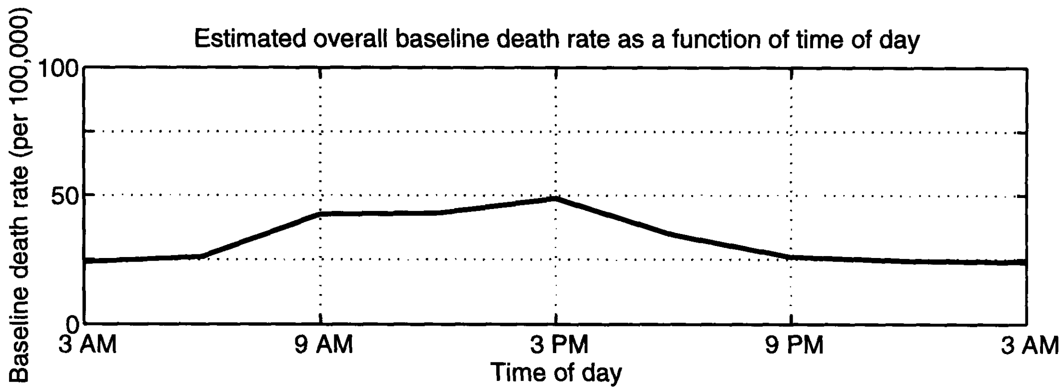

Figure 5-3: The estimated overall baseline death rate as a function of time of day. This estimate is computed using estimates of the baseline death rates (based on expert judgment) and estimates of the fraction of the population that falls in each subpopulation that the experts were asked about at different times of day (derived from various sources).

Since limited empirical evidence exists and since many of the experts are scientists who value original thought, they are quite likely to disagree with their colleagues. On the other hand, empirical evidence is not completely non-existent, and the experts are well aware of this evidence, so the suggestion that the estimates are wildly off the mark loses credibility.

The point is that a priori it was quite plausible that the expert estimates would vary widely. In fact, one of the experts reminded us to be sure to report "the wide variance in estimates that must certainly exist in the responses." In this light, the degree to which the experts agree is significant.

Overall Baseline Death Rate: The baseline death rate estimates for the individual subpopulations can be combined into an overall baseline death rate estimate using the population fraction estimates developed in chapter 3, as follows:

$$\hat{D}_0(t) = \sum_{i=1}^{m} F_i(t)\hat{D}_{i,0}$$

Here $\hat{D}_{i,0}$ is the estimated baseline death rate for subpopulation i, $F_i(t)$ is the fraction of the overall population in subpopulation i when the time of day is t, $\hat{D}_0(t)$ is the estimated overall baseline death rate at time t, and $m = 8$ is the number of subpopulations.

Figure 5-3 shows how this estimate varies with time of day. The overall death rate is lowest during the period from 9 PM to 6 AM, with a value around 25 per 100,000. During this period, most people are at home according to the population fraction estimates we are using (87.5%, 95.5%, 97%, and 92.5%, respectively, at 9 PM, midnight, 3 AM, and 6 AM). The overall death rate is generally higher during the day, with values in the neighborhood of 50 per 100,000. The peak at 3 PM occurs when the fraction of people in offices (the second most dangerous place to be) is highest, at 35.5%. The fraction of people in their cars on or near freeway overpasses (the most dangerous place to be) peaks at 2% during the morning and afternoon rush-hours, but this is not enough to pull the death rate above the 3 PM peak.

The relatively small range covered by the estimated overall death rate in figure 5-3 makes one question statements such as the following, which appeared in the *Los Angeles Times* shortly after the 1994 Northridge quake: "At another hour, on a normal business day, last week's casualties might easily have been multiplied a hundredfold." In a similar vein, U.S. Geological Survey seismologist Ross Stein is quoted in *Earth* [Dav94] magazine as saying "Northridge is an unremarkable quake for Southern California. It is remarkable for only one reason: It happened at 4:30 in the morning. If it had happened later, we would be counting the lives lost not in the tens but in the thousands. Our luck is going to run out at some point." In chapter 6 we pay closer attention to the Northridge quake and the 1989 Loma Prieta quake, which was of a similar magnitude but occurred around 5 PM, and thus provides a natural way to compare death rates at different times of day. We estimate that for regions experiencing intensity VII or less on the Modified Mercalli scale, the Loma Prieta death rate was six times higher[4] than the corresponding Northridge death rate. For the intensity VIII region, our estimate for the Loma Prieta death rate exceeds its Northridge counterpart by only 10%. Even if we use the factor of six to multiply the Northridge death toll of sixty, we do not get close to a death toll that is measured in the thousands.

In chapter 6, we also compare the actual death toll of the Northridge quake with an estimate (see table 6.2) based on the expert responses and the population distribution we have assumed for 3 AM. We conclude that if most of the experts, in making their assessments, were thinking of a scenario where intensity IX or higher on the Modified Mercalli scale prevails throughout the area of interest, and if the effects of the Northridge quake close to its epicenter were representative of what can be expected under such a scenario, then the experts were right on target. While the two premises on which this observation rests are not obviously true, it does reinforce our belief that the baseline death rate estimates are realistic.

We now turn to the expert estimates of how much the death toll could be reduced if earthquake warnings were available.

5.3.2 Relative Death Rates

Figure 5-4 shows histograms of the estimated relative death rates $G_{ij\tau}$ for two of the eight subpopulations that the experts were asked about. The histograms for the "people at home, awake" subpopulation show the general pattern most clearly, since the largest number of experts provided estimates for this subpopulation. Therefore we will direct our attention to the four histograms corresponding to this subpopulation. As figure 5-4 shows, the relative death rate estimates are concentrated close to 100% of the risk without warning when the lead-time is 30 seconds, indicating that most of the experts think that a lead-time this short has limited potential to reduce the death rate. As the lead-time increases to one minute and then to five minutes, the expert estimates begin to spread out over the range from 0% to 100%. With a thirty minute lead-time, the distribution of expert estimates has become concentrated at the lower end of the range, implying that the experts think a half-hour is sufficient to save most lives. A similar but less pronounced pattern is apparent in the histograms for "People outdoors," and for the remaining six subpopulations. Histograms

[4]This estimate is heavily influenced by a single event: the collapse of an elevated freeway. We argue in chapter 6 that this kind of collapse may be less likely to occur in future quakes. Without the deaths caused by this event the factor of six drops to a factor of two.

for the remaining subpopulations are shown in figures 5-15 to 5-17 (in appendix 5A).

Looking more closely at the histogram for people at home, awake, faced with a 30 second lead-time, we see that the arithmetic average of the relative death rate estimates is 76% and the sample standard deviation is 22%. These statistics can be used to construct a 95% confidence interval:

$$\hat{G}_{i=1,\tau=0.5} \pm t_{0.025,38-1}\text{s.e.}(\hat{G}_{1,0.5}) = 0.76 \pm 2.027\frac{.22}{\sqrt{38}} = (69\%, 83\%)$$

Here $\hat{G}_{i=1,\tau=0.5}$ is the point estimate 0.76 of the relative death rate for subpopulation one ("People at home, awake") and a half-minute lead-time ($\tau = 0.5$) and s.e.$(\hat{G}_{1,0.5})$ is the standard error of that estimate, i.e., the sample standard deviation divided by the square root of the sample size of 38. With 95% confidence we can state that the relative death rate is between 69% and 83%.

Such confidence intervals were computed for each histogram in figure 5-4 and figures 5-15 to 5-17, to facilitate comparisons between subpopulations, and to better discern general patterns. These intervals are shown in figure 5-5.

Comparing the trends in figure 5-5 one is tempted to divide the subpopulations into two groups: Those that consist of **people inside buildings** (people at home (awake or asleep), people in offices, and children in schools) on the one hand and **people outside buildings** (people in cars and people outdoors) on the other. According to the point estimates of relative death rates,

- three minutes suffice to halve the death rate and 25 minutes reduce the death rate by at least three-quarters for each of the *inside building* subpopulations. With one minute of warning, the estimated death rate is reduced by between 30 and 50%.

- In contrast, between 10 and 25 minutes are needed to halve the death rate and more than 30 minutes would be needed to reduce the death rate by three-quarters for the *outside building* subpopulations; one minute reduces the estimated death rate by between 15 and 25%.

  It should be noted that with the exception of the "Cars on bridge/overpass" subpopulation, the *baseline* death rate was lower for the "outside building" subpopulations than for the "inside building" ones (see figure 5-2).

The risk reduction estimates for a one minute lead-time are of particular interest because such lead-times may be achievable in some regions with real-time earthquake monitoring (See chapter 2).

A diminishing returns effect is apparent for all subpopulations: The marginal reduction in the death rate decreases as the lead-time increases. In chapter 6, the relative death rate curves in figure 5-5 will be used to predict how the number of deaths in three recent earthquakes might have been reduced had earthquake warnings been available.

How plausible are these findings? The dichotomy between people that are inside and outside seems reasonable, since the range of actions likely to increase chances of survival differs substantially depending on whether one is inside or outside a building. The actions differ both in terms of how much time is needed to complete them and by how much they are likely to increase chances of survival. Let us now consider the set of actions that it would be appropriate to take in response to earthquake warnings for people *inside buildings* and *outside buildings*, respectively.

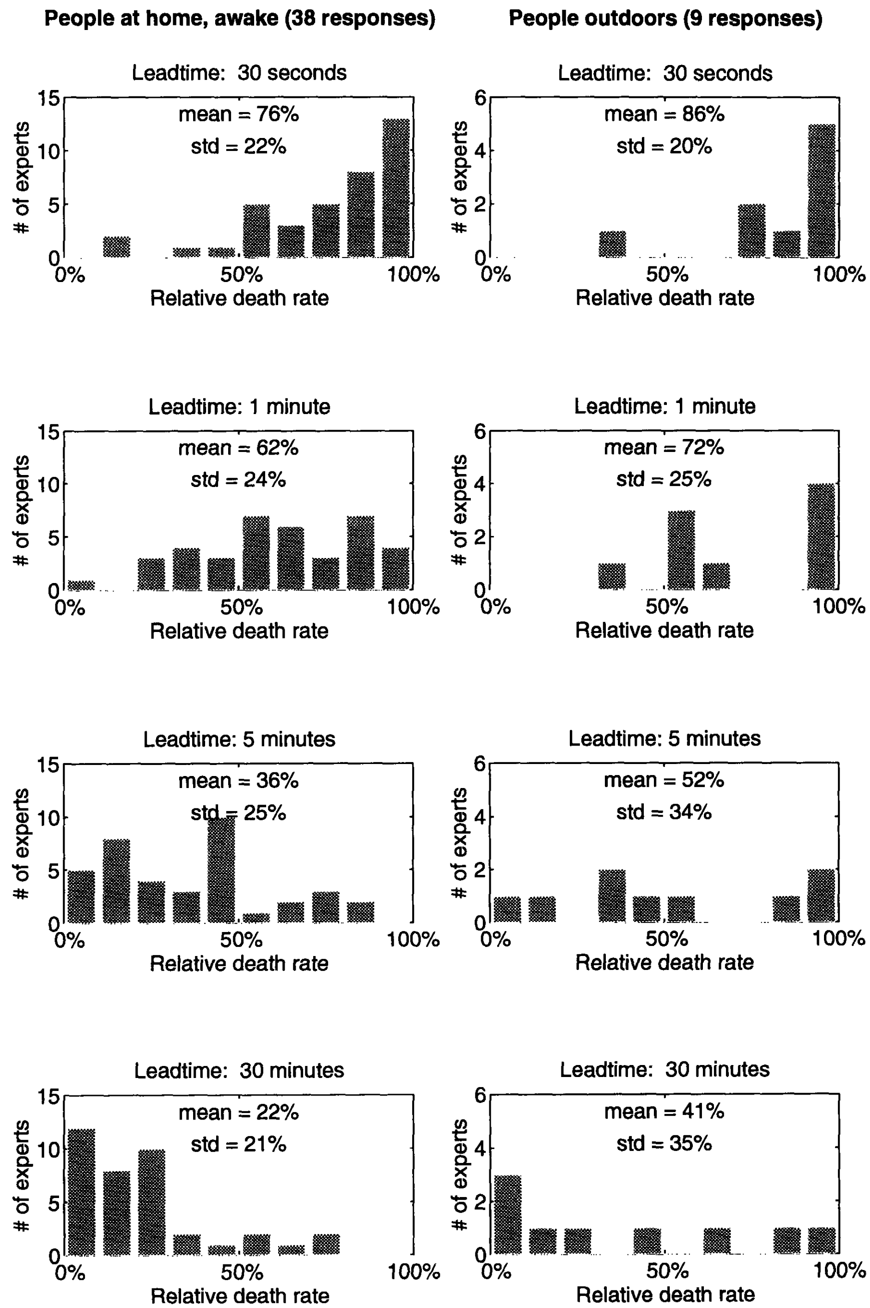

Figure 5-4: Histograms of relative death rates for "People at home, awake" and "People outdoors." Every expert did not provide estimates for every subpopulation. However, all experts did provide estimates for the "People at home, awake" subpopulation to allow for comparability of responses; hence the factor of four difference in the number of respondents.

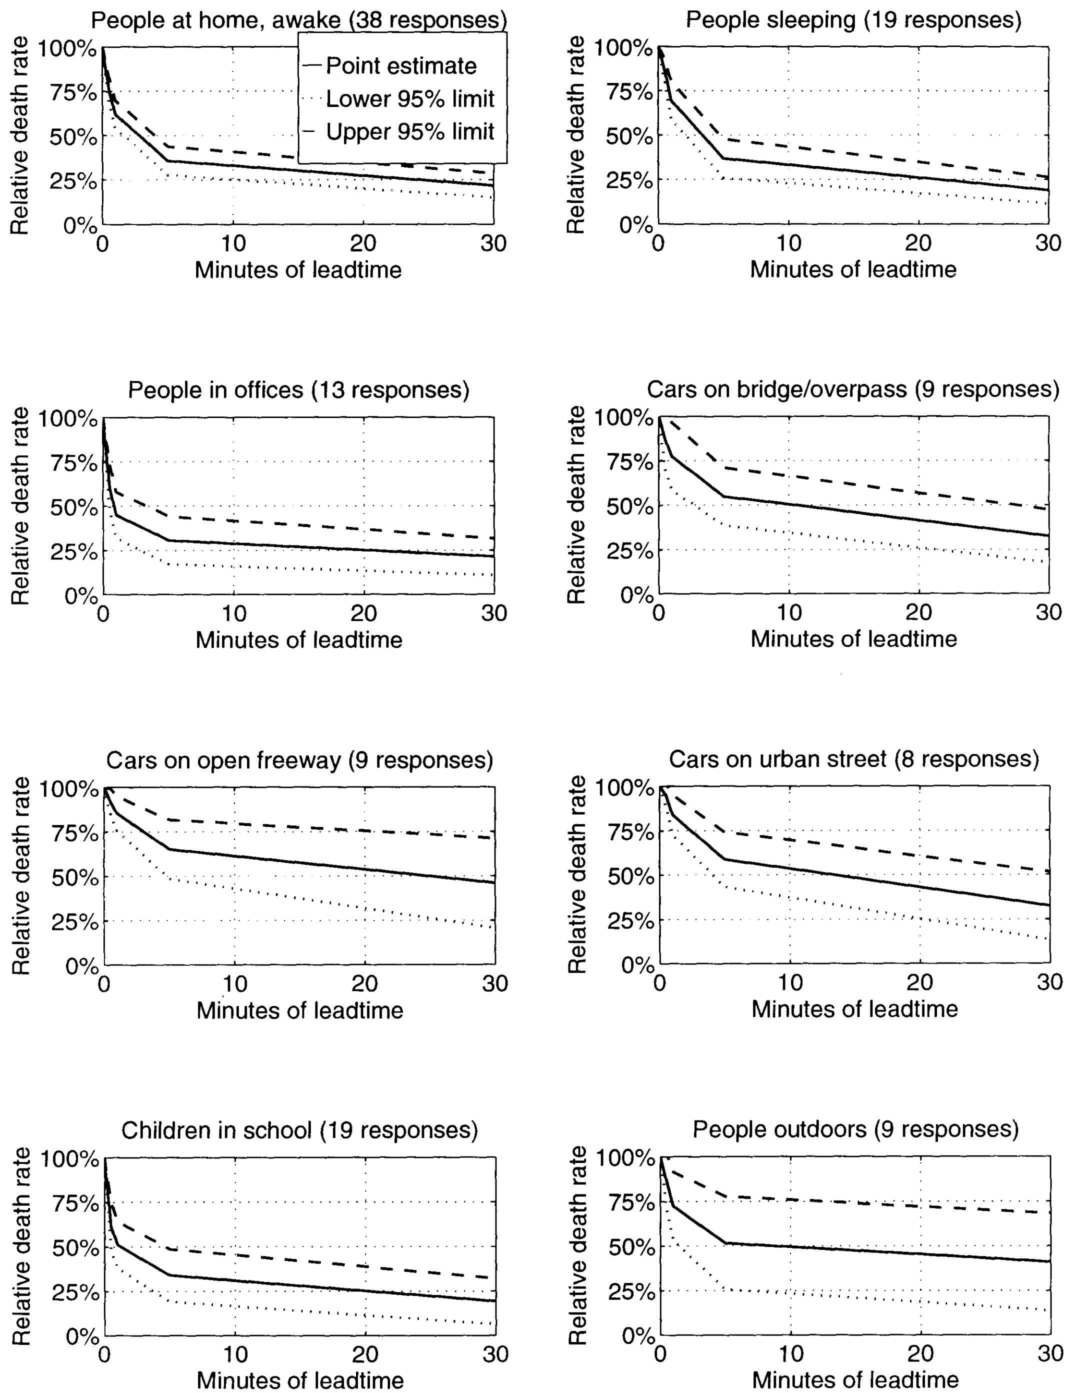

Figure 5-5: Confidence intervals for relative death rates in each subpopulation. Point estimates and 95% confidence limits for each of four lead-times (0.5, 1, 5, and 30 minutes) are connected with straight lines.

Appropriate Actions for People Indoors and Outdoors

There are two main classes of actions available to people inside buildings: seek shelter inside, or get outside the building. The main cause of deaths in earthquakes is being hit by parts of a collapsing building. The two obvious ways of avoiding this fate are to either not be in the building when it collapses or to be in a place inside the building likely to form a cavity in the ruins of the collapsed building. Cavities might for example exist under sturdy desks, inside door frames (if the door frame happens to be stronger than the walls adjacent to it), or in corners where two sturdy walls meet. In this light, the pattern of death rate reduction for the "inside building" subpopulations seems plausible. A three minute lead-time would probably suffice for most people to seek shelter inside a building, in a location likely to form a cavity in the event of collapse, and being in such a location may well halve the probability that the person will die in the earthquake. With a 25 minute lead-time, most people could exit the building they are in, and the death rate outside may well be less than a quarter of what it is inside.

It is instructive in this regard to consider a recent gas pipeline explosion that occurred in New Jersey. The explosion occurred near an apartment complex of two- and three-story brick-and-wood buildings where about 1500 residents lived (*New York Times*, March 25th, 1994). Most of the residents were asleep when, just before midnight on March 24th, 1994, they "heard a loud blast ... followed by a 7-10 minute gap before the buildings caught fire, giving them a few minutes to flee and averting a major disaster that could have cost many lives." Only one person died (from a heart attack) as a result of the explosion (although about 100 were injured). For our purposes, the important lesson is that people who lived in the apartments closest to the explosion were able to rouse themselves from sleep and evacuate the building in less than 10 minutes. Of course, there are many differences between this situation and an earthquake warning scenario (people are more likely to notice and take seriously a gas explosion than the blare of a siren) but it nevertheless demonstrates what people are capable of.

But, one might argue, the situation is different for people in offices. Indeed, after a terrorist bomb exploded in the World Trade Center in New York on February 26th, 1993, there were reports of people taking several hours to walk down from the upper floors (New York Times, February 27th, 1993). A counter-argument is that office buildings, such as the World Trade Center, are filled and emptied of people twice every working day, when people arrive for work in the morning and leave in the afternoon. Remember that the evacuation we are envisioning would take place *before* the earthquake; hence there is no reason to believe that the stairwells would be smoke-filled and without lights, as they were after the World Trade Center bombing (the elevators might not be operational, however, unless there is strong faith that the warning lead-time does not underestimate the time until the earthquake). Furthermore, even if evacuating a skyscraper takes more than thirty minutes, this is not necessarily inconsistent with the relative death rate profile shown in figure 5-5 for people in offices, because tall buildings are usually designed to higher standards of earthquake resistance than other buildings. Most of the potential earthquake deaths among people in offices are probably in old unreinforced masonry buildings. These buildings are typically not very tall, and could probably be evacuated in less than 25 minutes.

For people outside buildings, the range of actions depends more heavily on the situation, but again it is instructive to consider the most likely causes of death. For people in cars, two scenarios come to mind: Drivers losing control of their vehicles because of the earthquake and collapsing highway bridges. One can envision a variety of actions that could be taken to avoid either situation. With a sufficiently long lead-time, a driver could simply drive to a comparatively safe location. With lead-times between 10 and 25 minutes, most drivers should be able to make sure that they are not on, under, or close to a highway overpass at the instant the earthquake occurs, either by driving away, or by leaving the car if a traffic jam occurs. If the lead-time is shorter than 10 minutes, a driver has more limited options. He may attempt to stop the car and pull to the side of the road, where he may either stay in the car or venture outside. These actions can reduce the danger from the causes mentioned already. However, depending on the circumstances, taking these actions could also expose the occupants of the car to new dangers. Conceivably, other drivers may decide not to slow down (either because they are not aware of the warning or because they do not consider slowing down to be a sensible thing to do). Attempting to bring a car to a halt on a crowded freeway where other drivers have no intention of slowing down is of course dangerous. And even if one has safely stopped the car and pulled over to the side of the road, one is still at risk from passing cars. Sometimes the occupants of the car might be able to leave the car for a safe, open area, but the physical layout of the highway and lack of time can prevent this.

We certainly do not wish to suggest that drivers can do nothing to improve their chances of survival when the lead-time is short. If most drivers heed the warning and decide to stop, the others will probably do so to avoid collisions. But the options are less clear-cut and more situation-dependent than for people inside buildings. Sometimes the best action might be to continue driving at a steady, perhaps reduced, speed. The expert judgment is consistent with these observations: it takes at least twice as long to halve the death rate for the three subpopulations of people in cars, compared to people inside buildings.

The results for "people outdoors" are a bit puzzling. In an urban setting, this subpopulation would contain both people in open areas where nothing but the sky could possibly fall on a person and people on crowded downtown streets where they are in danger of being struck by objects falling off the outside of buildings during an earthquake; groups whose *baseline* death rates are perhaps at opposite ends of the spectrum. Presumably, most fatalities incurred among people outdoors during an earthquake that occurs without warning would come from the latter group. One would also presume that with lead-times of ten minutes or more most people that are outdoors should be able to reach a park or other comparatively open space, thus reducing the potential death toll substantially. Based on these considerations, one would expect that the death rate for people outdoors could be reduced at least as much as for people indoors with lead-times of a few minutes. Yet the experts seem to suggest that even with a half hour of warning, the death rate for people outdoors would only be reduced by about 60%.

There are a few possible explanations for this. First, the experts generally thought that it would be more difficult to communicate the warning to "people outdoors" using sirens (as has been assumed in the analysis so far) than in other settings. Second, some experts may not have included the high-risk group of people on crowded urban streets, perhaps next to old unreinforced masonry buildings, in this category. The large standard deviations of the relative death rate estimates for people outdoors for lead-times of five and thirty minutes (see figure 5-4) suggest that perhaps the nine experts that provided the estimates were making different assumptions about whom to include among "people outdoors" and about what being "outdoors" means. Third, one should note that the confidence limits around the relative death curve in figure 5-5 are wider for "people outdoors" than they are for any other subpopulation (because the sample standard deviations are large and the sample size is small, compared to the other subpopulations). If instead of looking at the point estimates, one focuses on the lower 95% confidence limit, then a halving of the death rate with about three minutes of warning and a three-quarters reduction with less than 10 minutes of warning are seen.

Overall, the confidence intervals in figure 5-5 are rather narrow, especially for those subpopulations where estimates from ten or more experts were available. This indicates that there is substantial agreement among the experts about the how much the risk can be reduced; a conclusion that was far from obvious before the experts were surveyed.

However, this observation raises the following question: Is it reasonable for the width of the confidence intervals to continue to decrease towards zero as estimates from more and more experts are obtained? Or to put it differently, suppose that 100 experts provide estimates that, when processed according to the procedures described above, result in a 95% confidence interval of $(9\%, 11\%)$ for the relative death rate among office workers when given a lead time of 30 minutes. Would we then believe that the chances were no more than one in 20 that the relative death rate in a future earthquake differed by more than one percentage point from 10% for officer workers, given a 30 minute lead-time? This question pertains to all of the confidence interval calculations in this chapter, not just the ones for the relative death rate.

To answer this question we have to examine the assumptions underlying the confidence interval calculation. These assumptions are that the expert estimates of relative death rates are independent draws from a normal distribution (a lognormal distribution was assumed for the baseline death rates) whose mean coincides with the "true" relative death rate. Consider first the assumption that the distribution is centered on the true death rate. Any specific alternative assumption, such as assuming that the experts systematically underestimate the relative death rate, would be stronger since it would require one to believe that one knows more than the experts do. The same can be said for the independence assumption: While it is likely that the expert estimates are correlated (for example because of similar educational backgrounds, familiarity with the same research literature, etc.) to some degree, it is difficult to feel confident about any specific number assigned to the correlation. Therefore, while we may suspect that some experts are biased (in the sense of systematically either over- or underestimating the true relative death rate) and that different "measurement errors" may be correlated, we have no solid basis for assigning values to either the bias or the correlation and hence we use the neutral value of zero.

The implications of these two decisions for the validity of the confidence intervals are a bit different. If a specific strictly positive number had been chosen for the correlation between experts, then wider confidence intervals would have resulted. The choice of one number for the bias with which the experts estimate the relative death rates, on the other hand, would not affect the width of the confidence intervals because one would simply subtract the assumed bias from every estimate to make it unbiased. However, if instead of one number, a distribution of values were used for the bias, then wider confidence intervals would result.

A distribution of values is probably more appropriate, given the variety of backgrounds of the experts. As discussed in chapter 3, we canvassed experts from a variety of fields, because we believe that many different kinds of knowledge are relevant to estimating the quantities of interest to us. For example, consider the estimation of the fraction of people at home, awake, that hear a siren blare within 30 seconds. Among the kinds of knowledge that might be helpful in estimating this quantity are

- Knowledge of the physiology of hearing.

- Knowledge of the physics of sound propagation.
- Knowledge of how a system of sirens is typically designed, e.g., where sirens are usually located, how far apart they are, etc.
- Knowledge of how to properly average over groups of people with different characteristics, i.e., avoiding common mistakes and being aware of potential cognitive biases.

No one expert is likely to possess all of these skills. While we hope that the strengths of the different experts complement each other and their weaknesses do not overlap, we must recognize that the confidence intervals in figure 5-5 (and elsewhere in this chapter) only account for some of the uncertainty around the point estimates of relative death rates.

At the end of this section we examine whether the data provide evidence for one specific alternative to the assumption of independent, unbiased experts: that some experts are consistently optimistic while others are consistently pessimistic. While the evidence is not clear-cut, there appears to be little reason to worry about such an effect.

We now discuss how to combine the relative death rates for the eight subpopulations into one overall relative death rate.

5.3.3 Overall Relative Death Rates

The overall relative death rate at lead-time τ is the fraction of the overall baseline death rate that remains if τ minutes of warning are available:

$$G_\tau = \frac{D_\tau}{D_0} = \frac{\sum_{i=1}^{m} F_i D_{i,\tau}}{D_0} = \frac{\sum_{i=1}^{m} F_i G_{i,\tau} D_{i,0}}{D_0} = \sum_{i=1}^{m} \left(F_i \frac{D_{i,0}}{D_0} \right) G_{i,\tau} \qquad (5.1)$$

We see that the overall relative death rate is a weighted sum of the relative death rates $G_{i,\tau}$ for each subpopulation. However, the weight for each subpopulation is not simply the population fraction F_i. The population fraction F_i has to be scaled up or down by the ratio of the baseline death rate $D_{i,0}$ for subpopulation i to the overall baseline death rate D_0. For example, at 9 AM, the weight given to the relative death rate for people in cars on or near overpasses would be the fraction of the population that belongs to this group at nine in the morning (2%) multiplied with 203 per 100,000 (baseline death rate for people in cars on or near freeway overpasses) and divided by 50 per 100,000 (the overall baseline death rate at 9 AM). This scales the 2% population fraction up to around 8%. Essentially, even though only 2% of the population is in cars near freeway overpasses, these 2% include 8% of the potential earthquake deaths. Thus, the success in reducing this risk via a warning has a disproportionate effect on the overall death rate.

Figure 5-6 shows the estimated overall relative death rate, based on expert estimates of baseline and relative death rates and independent estimates of the population fractions, as a function of lead-time. The population fractions used to compute this estimate are time-averages over the twenty-four hours of the day. According to this curve,

- it takes around three minutes of lead-time to halve the death rate and about 25 minutes of lead-time to reduce the death rate by three quarters. A one minute lead-time reduces the estimated death rate by about 45%.

Once again, there is a noticeable diminishing returns effect: it takes only three minutes to halve the death rate from its baseline level, but an additional twenty-two minutes are needed

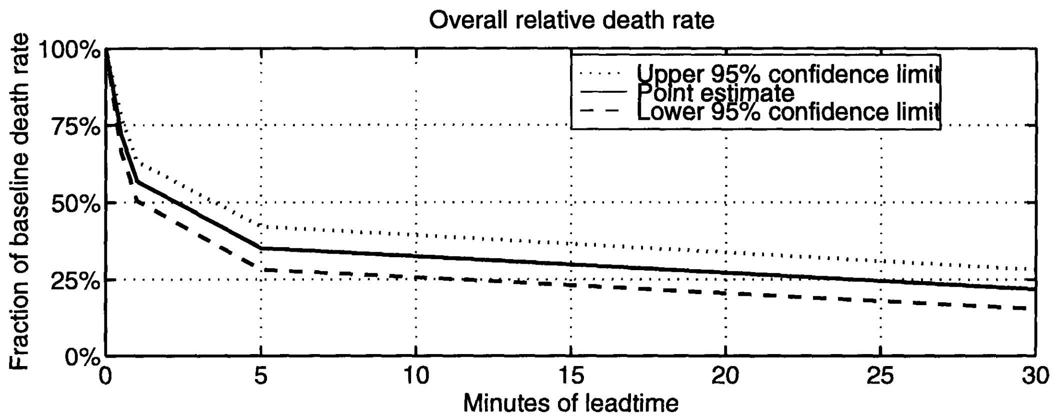

Figure 5-6: Estimated overall relative death rate with 95% confidence intervals. The estimate is computed by combining expert estimates of death rates in different subpopulations and estimates (based on data from various sources, see chapter 3) of how the population is distributed across the subpopulations.

to halve it again. Furthermore, 45% of the 50% reduction attained with three minutes of warning are achieved within the first minute.

The estimated overall relative death rate curve looks very similar to the corresponding curves for each of the "inside building" subpopulations. This simply reflects the fact that a majority of the population is estimated to belong to one of those subpopulations at all hours of the day. Table 5.1 shows the estimated time-average population fractions and how they are adjusted to account for differences between the subpopulations in the baseline death rate. We see from this table that when averaged over the 24 hours of the day, the "inside building" subpopulations are estimated to include 91.5% of the population and 92.1% of the potential earthquake deaths.

Figure 5-6 shows 95% confidence limits around the estimated overall relative death rate. The lead-time required to halve the death rate is seen to range from one minute (based on the lower confidence limit) to four minutes (based on the upper confidence limit). Similarly, the lead-time necessary to reduce the death rate by three-quarters ranges from 12 minutes to around 30 minutes.

We note that at any given time of day, the estimated overall relative death rate curve will look slightly different from the time-average curve of figure 5-6. Or computations indicate that during the day (9 AM to 6 PM), a greater than average one-minute risk reduction (around 50%) is achieved but it takes longer than average (between 25 and 30 minutes) to reduce the estimated death rate by three quarters. At night (12 AM to 6 AM) the situation is reversed: the one-minute risk reduction is smaller than average (about 35%) but the time needed to reduce the estimated death toll by three quarters is below average (between 20 and 25 minutes).

We now turn our attention to analyzing how sensitive the relative death rate curve in figure 5-6 and the confidence intervals around it are to (1) perturbations to the estimated population fractions F_i from chapter 3 and (2) the means by which the earthquake warning is communicated.

Subpopulation	F_i	$D_{i,0}/D_0$	$F_iD_{i,0}/D_0$
People at home, awake	0.339	0.635	0.215
People at home, asleep	0.319	0.682	0.218
People in offices	0.147	3.029	0.445
Cars on bridge/overpass	0.011	6.004	0.066
Cars on open freeway	0.021	0.076	0.002
Cars on urban street	0.021	0.217	0.005
Children in school	0.110	0.393	0.043
People outdoors	0.033	0.196	0.007
Total for people inside:	0.915		0.921

Table 5.1: Time-average (over the 24 hours of the day) estimated population fractions and their adjustment to account for different estimated baseline death rates. Numbers greater than one in the second column indicate that the corresponding subpopulation has "more than its share" of the potential earthquake deaths.

5.3.4 Sensitivity Analysis

Sensitivity to Estimated Population Fractions: The algebraic form of the confidence limits (developed in Appendix 5B) allows one to assess the effect of varying the estimated population fractions. The population fraction estimates were after all based in part on educated guesses, so it would be comforting to know that the curves in figure 5-6 are not extremely sensitive to the population fraction estimates, and this turns out to be the case.

The confidence intervals for the relative death rates have the following form:

$$\left(\hat{G}_\tau - t_{\alpha/2,\nu}\mathrm{SE}_\tau, \hat{G}_\tau + t_{\alpha/2,\nu}\mathrm{SE}_\tau\right)$$

The point estimate of the overall relative death rate $\hat{G}_\tau$ is computed using equation 5.1, replacing true values with estimates where appropriate. The standard error SE_τ has the form

$$\mathrm{SE}_\tau = \sqrt{\sum_{i=1}^{m} \frac{F_i^2}{n_i}} \sqrt{\sum_{i=1}^{m} \frac{n_i - 1}{\nu} S_i^2}$$

with $m = 8$ being the number of subpopulations, n_i the number of experts that provide estimates for the i-th subpopulation, $\nu = \sum_{i=1}^{m}(n_i - 1)$ the number of degrees of freedom, and S_i^2 the sample variance of the relative death rates for subpopulation i (and lead time τ).

Changes in the estimated population fractions $\{F_i\}$ will affect both the point estimates $\hat{G}_\tau$ (as dictated by equation 5.1) and the width of the confidence intervals (through the quantity $\sqrt{\sum_1^m F_i^2/n_i}$). The point estimates are most sensitive to changes in the population fraction F_i for the subpopulation for which the quantity $D_{i,0}G_{i,\tau}/D_0 = D_{i,\tau}/\sum_1^m F_iD_{i,0}$ is largest (cf. equation 5.1), i.e., the subpopulation with the largest absolute death rate. The width of the confidence intervals is most sensitive to changes in the subpopulation fraction F_i for which F_i/n_i is large, i.e., subpopulations for which the smallest number of experts provided estimates.

To explore how sensitive the relative death rate curve in figure 5-6 is to the estimated

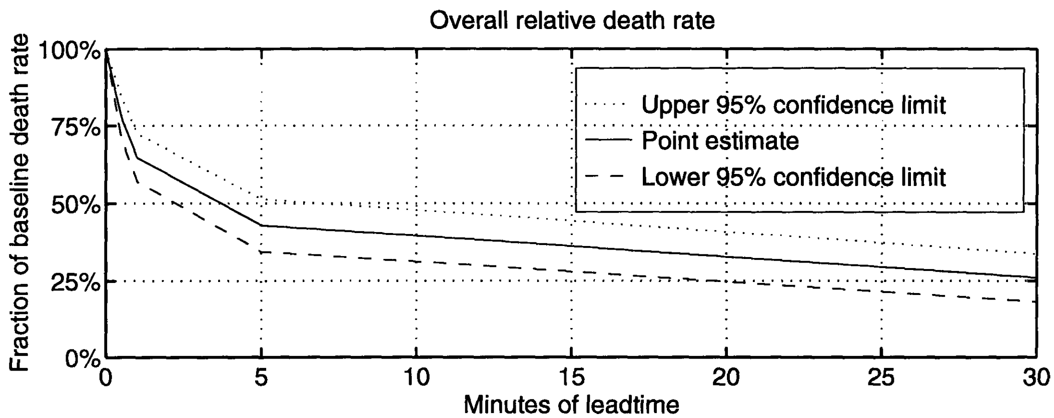

Figure 5-7: Estimated overall relative death rate with 95% confidence intervals. This figure is the same as figure 5-6, except that 10% of the population has been moved from their homes to cars on or near highway overpasses. Notice that while the confidence intervals have widened somewhat as a result of this move, the point estimates have changed very little.

population fractions F_i, we moved some of the population from the "people at home, awake" subpopulation, which has one of the lower estimated absolute death rates ($\hat{D}_{i,\tau}$) and the largest sample size n_i, to the "people in cars, on or near a highway overpass" subpopulation, which has the largest estimated absolute death rate and the second smallest sample size (sample size here refers to the number of experts that provided estimates). The estimated fraction in the "people in cars, on or near highway overpasses" subpopulation was increased from 1.1% to 11.1% (surely an overestimate) and the estimated fraction for the "people at home, awake" subpopulation was reduced from 33.9% to 23.9%. Figure 5-7 shows the effect of this change. The width of the confidence intervals has indeed increased but the point estimates have not moved much. The lead-time needed to halve the death rate has increased from three minutes to four minutes, and the time needed to reduce the number of deaths by three quarters has increased from 25 to 30 minutes; hardly a drastic change. We conclude that the fact that there is some uncertainty associated with the population fractions estimates arrived at in chapter 3 is not a matter of great concern.

Sensitivity to the Means of Communication Used: So far, the face-value analysis has assumed that sirens were the sole means of communication used. Systems of sirens exist in many communities for the purpose of alerting the population to an imminent hazard. Therefore, it is to be expected that if last-minute earthquake warnings were to become a reality, sirens would at least be one of the means used to communicate such warnings to the public. All experts were asked to estimate the fraction of people that would hear the warning before the earthquake occurs, assuming sirens were the sole means of communication, for all subpopulations and lead-times. However, the experts were also asked to estimate the same fraction under the assumption that some other technology (for example lights on freeways or "earthquake alarms" to be installed in all dwellings) were the exclusive means of communication. Furthermore, the experts were asked to suggest other means of communication, and to estimate the fraction that would hear the warning in time under the assumption that the most effective *pair* of communication strategies (e.g. sirens *and* earthquake alarms) was used. In the spirit of estimating what could be achieved in the best case,

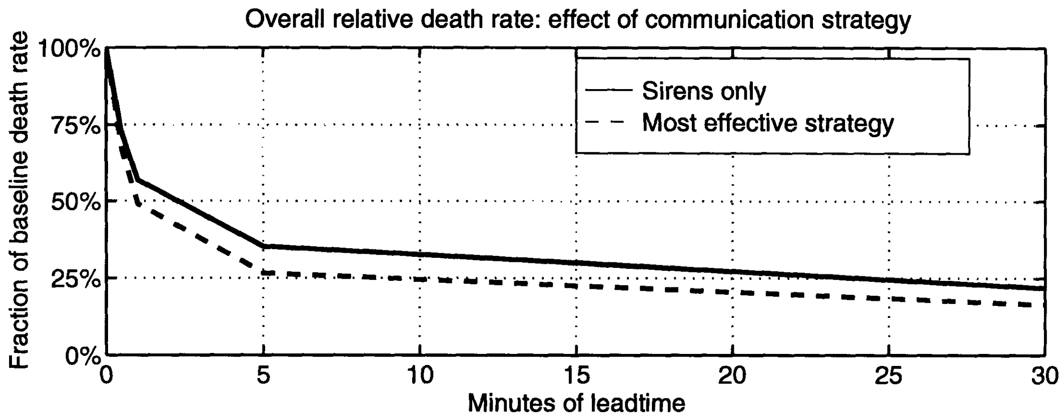

Figure 5-8: Comparison of estimated overall relative death rate curves under the assumption that sirens are the sole means of communication (as in figure 5-6) and under the assumption that the most effective communication strategy is used.

i.e., if the cost of installation were no object, the estimation of the overall relative death rate was repeated with the fraction $A_{ij\tau}$ of people estimated to hear the warning in time if only sirens were used replaced with the *largest* such fraction given by expert j. Usually, the largest fraction corresponds to the most effective pair of communication strategies, unless the expert did not provide estimates for this case. In that event, the fraction corresponding to the most effective single communication strategy was used.

The results are shown in figure 5-8. We see that the change in the estimated relative death rate is sizable. The lead time needed to halve the estimated death rate decreases from three minutes to around one minute, and 10 minutes of lead-time now suffice to reduce the estimated death toll by three-quarters, compared to 25 minutes under the "sirens only" assumption. The diminishing returns aspect of the curve is even stronger than it was in figure 5-6, suggesting that if effective communication strategies are used and if people can be educated or trained to react quickly enough to earthquake warnings, then most of the life-saving potential of earthquake prediction could be realized with as little as one minute of warning.

We conclude this section on "face-value" analysis by examining the data for signs of significant heterogeneity among the experts.

5.3.5 Heterogeneity: Are Some Experts Consistently Optimistic?

There are two main reasons for our interest in whether some of the experts are consistently more optimistic than others. First, our approach throughout has been a best-case one, i.e., our intent is to estimate the life-savings that can *optimally* be achieved if warnings before earthquakes were possible. Continuing in this spirit, it could be argued that experts in the "optimistic camp" are likely to assess the best-case life-saving potential more reliably than the other experts, perhaps because the judgment of the pessimists was affected by their discomfort with the best-case assumptions they were asked to base their assessments on. If one believes this argument, then it provides an incentive for identifying the optimistic experts and computing separate risk reduction estimates based solely on their judgment. Even if one does not believe the argument it might be of interest to see whether there are signs of substantial heterogeneity among the experts, and if so, the extent of the disagreement between the optimists and the pessimists.

Second, in the computation of confidence intervals around the relative death rate estimates, we have implicitly made a series of assumptions including independence and unbiasedness of the expert estimates. By unbiasedness we mean that "on average", the expert estimates are correct, i.e.,

$$\mathrm{E}[D_{ij\tau}] = \begin{array}{l}\text{true death rate for people in subpopulation } i \\ \text{when } \tau \text{ minutes of lead-time are available}\end{array}$$

These assumptions would be negated by an optimism/pessimism effect, if it exists.

Of course, the assumptions underlying the confidence interval calculations are idealizations that one would never expect to hold exactly in reality. The important question is not whether these assumptions hold exactly, but whether the results of computations based on them are sensitive to plausible departures from the assumptions.

In appendix 5C, we examine a specific optimism/pessimism hypothesis. Under this hypothesis, an optimistic expert would not only judge the baseline death rate $D_i(0)$ for any given category to be low, compared to other experts, he would also give a comparatively high estimate for the risk reduction $(D_i(0) - D_i(\tau = 30\text{ s}))/D_i(0)$ possible with 30 seconds of warning, and for the *additional* risk reduction possible with a one minute warning, i.e., $(D_i(1\text{ min}) - D_i(30\text{ s}))/D_i(30\text{ s})$, and so on. In particular, we do not simply test whether experts that estimate a baseline death rate as comparatively low also tend to give low death rate estimates for positive warning lead-times.

The conclusion of the calculations in appendix 5C is that the pattern of expert estimates for the "people at home, awake" subpopulation provides little evidence that some experts are consistently optimistic or pessimistic. However, data for some of the remaining seven subpopulations do support the notion of an optimism effect.

Despite not having reached a clear-cut conclusion that some experts are consistently optimistic and others pessimistic, we assumed this to be the case, and explored what effects such an assumption would have on the overall relative death rate estimates shown in figure 5-6. The experts were divided into optimists and pessimists using a procedure described in appendix 5C. Then, we computed estimates of the overall relative death rate using only the optimists and only the pessimists, respectively. Figure 5-9 shows the results (and the original overall relative death rate estimates, for comparison). According to the optimists, the death rate is halved with one minute of lead-time and is reduced by three-quarters with five minutes of lead-time. Even according to the pessimists, the reduction in death rates is substantial: a 50% reduction with five minutes of lead-time and almost three-quarters with 30 minutes of warning.

This concludes the "face-value" analysis of the expert responses. Next, we take a different approach to analyzing the same data. This approach will be referred to as "model-based;" it entails postulating certain parametric models of how humans react to earthquake warnings and of how the experts form their judgment. The expert responses are then used to estimate the parameters of the postulated models.

5.4 Model-Based Analysis

In the face-value analysis, we calculated point estimates and confidence intervals for the relative death rate in each subpopulation and for each lead-time (30 seconds, 1 minute, 5 minutes, and 30 minutes) that the experts were asked to provide estimates for. The

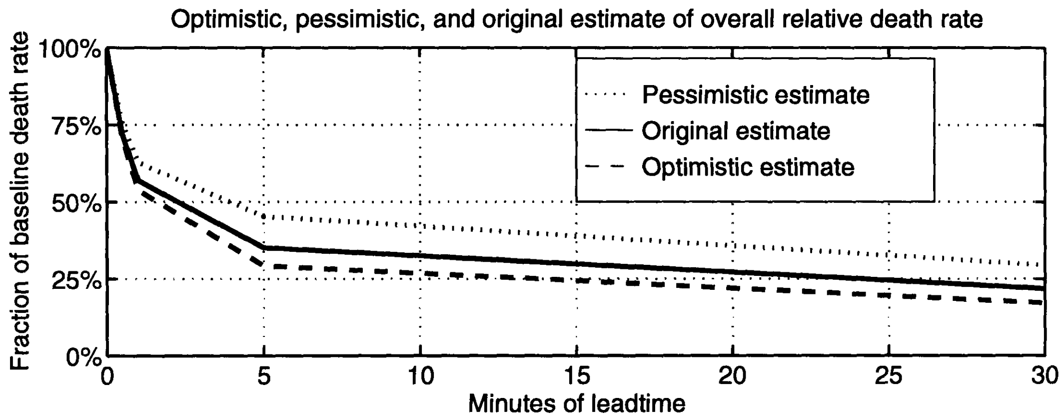

Figure 5-9: Optimistic, pessimistic, and original (i.e., based on the judgment of all the experts) overall relative death rate estimates.

computations were based on a minimum of additional assumptions. For a somewhat trivial example, there was nothing in the analysis to prevent the relative death rate estimate with 30 minutes of lead-time from being higher than the relative death rate estimate with 30 seconds of warning, even though such an outcome would be at odds with our expectations. Fortunately, this did not happen.

Not imposing unverifiable assumptions certainly has great merit, unless it has the effect of stifling progress. We feel that there is much to be learned from making specific assumptions about how humans would react to earthquake warnings if such warnings were available, expressing these assumptions as mathematical models, and then trying to calibrate the models with the responses of the experts. For one thing, the development of such models forces one to think carefully about the human reaction process. One must focus on the most important aspects of this process to obtain tractable models; identifying and then forcing oneself to ignore the less important details can lead to valuable insights. And even though the assumptions are unverifiable in the sense that it is impossible to know whether the models provide a satisfactory approximation to what would actually occur if reliable public earthquake warnings were to become a reality, one can at least observe how well the models fit the expert judgment.

Each expert probably used something akin to a model to arrive at estimates of the quantities that he was asked about in the survey. That is, an expert would probably try to identify critical quantities, settle on values for the critical quantities, and then somehow combine these values into one number. A thought process of identifying critical quantities and choosing a rule for obtaining a prediction based on the critical quantities is of course nothing but a process of building a model.

Different experts will choose different critical quantities and different rules for combining them. Therefore, the models chosen by us will not coincide exactly with the thought processes of all the experts that responded to the survey. In a sense, the purpose of the model-based analysis is to infer logical risk estimates from the expert judgments. Our models may not be the only logical way to proceed, i.e., there may be other equally plausible models, but we believe the approach we have taken has the advantages that

- the models are tractable,
- the models are flexible in that they can be extended easily and naturally to capture

greater complexity (some possible extensions were discussed in chapter 4), and

- most of the model parameters have a clear physical interpretation, such as "the mean time it takes an able-bodied individual to get under a sturdy desk," or "the average rate per minute the endangered population hears a siren blare."

The models of human behavior that we have developed were discussed at length in chapter 4. Here, we will review the models briefly, relate the models to the quantities that the experts were asked to estimate, and then use the expert estimates to calibrate the models. The resulting estimates of overall relative death rates and relative death rates for individual subpopulations will then, in the following section, be compared to the estimates obtained using the "face-value" approach.

The behavioral models describe the propagation of the warning message through a population, and the time it takes members of the population to complete an action whose purpose is to increase chances of survival. Let us define two random variables to facilitate the interpretation of the models:

$$T_A(i,\tau) \equiv \begin{cases}\text{The time until a randomly chosen individual in subpopulation } i \text{ hears a} \\ \text{warning that is first broadcast } \tau \text{ minutes before an earthquake occurs.}\end{cases}$$

$$T_C(i,\tau) \equiv \begin{cases}\text{The time needed by a randomly chosen member of subpopulation } i \text{ to} \\ \text{complete an appropriate action whose purpose is to enhance chances of} \\ \text{survival in an earthquake, if warnings are available } \tau \text{ minutes before the} \\ \text{earthquake occurs.}\end{cases}$$

(The arguments i and τ will often be deleted for clarity). The behavioral models can be thought of as specifying functional forms for the probability distributions for random variables T_A and T_C. The parameters of these distributions could in principle depend on both the subpopulation i and the lead-time τ, but we will assume that they depend only on the subpopulation. Thus, for example, the time a person uses to get under a desk is assumed independent of whether the person thinks he has 30 seconds or 30 minutes at his disposal before the earthquake occurs.

An individual will hear the warning and complete an action before the earthquake occurs if and only if the total time until the action is completed is less than the lead-time τ, i.e., if $T_A + T_C < \tau$. The notation used in the previous section, where $A(\tau)$ denoted the probability that an individual hears the warning before the earthquake occurs and $C(\tau)$ denoted the conditional probability of being able to complete an appropriate action before the quake occurs, given that one has heard the warning in time, relates to the random variables T_A and T_C as follows:

$$AC = \Pr\{T_A + T_C < \tau\}$$

People that hear the warning and complete an appropriate action face the revised death rate D'', while those that do not hear the warning or are unable to complete an appropriate action before the earthquake occurs face the baseline death rate D', so the overall death rate in a particular subpopulation will be

$$D(\tau) = ACD'' + (1 - AC)D' = \Pr\{T_A + T_C < \tau\}D'' + \Pr\{T_A + T_C \geq \tau\}D'$$

We will use the expert responses to calibrate probability distribution functions for the random variables T_A and T_C, for each subpopulation, and then these distribution functions will be convolved[5] to obtain the probability distribution function for the total time $T_A + T_C$ from the instant a warning is first issued until a randomly chosen person has heard the warning and completed an appropriate action in response to the warning. This allows one to estimate $D(\tau)$ for lead-times other than the four lead-times that the experts were asked about, in contrast to the face-value analysis which did not explicitly allow for extrapolation.

5.4.1 Distribution of the Time to Hear Warning, T_A

Chapter 4 presented a differential equation that was postulated as a model of how warning information propagates through a population. The differential equation described the evolution of the fraction $A(t)$ of a population that is aware of the warning by time t, where t measures time from the instant that the warning is first broadcast. The fraction $A(t)$ can be interpreted as the probability that a randomly chosen member of the population has heard the warning by time t, which equals the probability that the time T_A until the randomly selected individual hears the warning is less than t. Thus, there is a straightforward relationship between the fraction A and the distribution function for the random variable T_A:

$$A(t) = \Pr\{T_A < t\} = F_{T_A}(t)$$

We assume that $A(t)$ follows a differential equation of the following general form (cf. equation (4.5)):

$$\frac{dA}{dt} = U_1(t)(1 - A(t)) + NU_2(t)A(t)(1 - A(t))$$

The function $U_1(t)$ is the instantaneous rate at which the warning information is being transmitted from a primary source, such as a system of sirens, while $U_2(t)$ measures the rate at which "person-to-person" transmission of the warning information occurs among the N members of the population of interest. As discussed in chapter 4, the identification of the second term in the differential equation with "person-to-person" transmission should not be taken too literally. For example, suppose that one student in a class-room does not initially hear an announcement on the PA system, but his attention is turned to it immediately by other students. This kind of person-to-person interaction may be considered to be modeled by the first term ($U_1(1 - A)$) in the differential equation.

Collectively, we will refer to the functions $U_1(t)$ and $U_2(t)$ as urgency functions, since the rates of transmission can be expected to depend heavily on how urgent the warning is perceived to be by the public.

Two special cases of the differential equation were calibrated using the expert responses: (1) the constant urgency case, where U_1 and U_2 are constants, and (2) the decreasing urgency and no person-to-person contact case, where $U_1(t) = \lambda/(t+\gamma)$ (for some constants λ and γ) and $U_2 = 0$. The constant urgency model has the advantage of simplicity, but as we will see it does not provide a good fit to the expert responses. The decreasing urgency model gives a better fit to the expert responses, possibly because it models the following phenomenon. Consider the probability that a randomly chosen person that has not heard the warning by time t hears it during the infinitesimal interval $(t, t + \Delta t]$. The fact that the person has not heard the warning yet increases the chances that the person for some reason is not likely to hear the warning in the near future, for example because the person is listening to loud music or is in a deep sleep. Therefore, we expect the conditional probability of hearing the

[5] We assume that T_A and T_C are independent.

warning in the next time interval of length Δt, given that the individual has not heard the warning yet, to decrease with time. But this is exactly what happens in the decreasing urgency model, since $U_1(t)\Delta t$ *is* the conditional probability referred to above.

The relation between the quantity $A_{ij\tau}$ that the experts were asked to estimate and the distribution of $T_A(i,\tau)$ is straight-forward:

$$
\begin{aligned}
A_{ij\tau} &= \text{j-th expert's estimate of the fraction of people in subpopulation i that hear an earthquake warning before the earthquake occurs if the warning lead-time is τ.} \\
&= \text{j-th expert's estimate of } \Pr\{T_A(i,\tau) < \tau\}
\end{aligned}
$$

Occasionally, experts estimated that 100% of a particular subpopulation would hear an earthquake warning in time. If taken literally, such an estimate is difficult to believe. Presumably, these experts are of the opinion that the fraction $A(t)$ is 100% "for all practical purposes," i.e., their estimate has passed some mental threshold of closeness to 100%. A similar problem will be discussed later, that of experts estimating death rates as being equal to zero. The zero estimates were replaced by a number that was one-tenth of the smallest non-zero estimate provided by the expert in question and a similar strategy was used to replace the 100% fractions; they were replaced with numbers strictly less than 100%, but larger than any other estimate of the fraction $A(t)$ provided by the expert.

One reason for replacing the 100%'s is the same as that for replacing the zeros: To avoid difficulties with the estimation schemes adopted. In particular, a 100% estimate for $A(t)$ is inconsistent with the differential equation model we are postulating. Under that model, the fraction $A(t)$ will approach 100% as $t \to \infty$, but it will be strictly less than 100% for any finite t. It could be argued that this is simply an imperfection of the model, but we would argue that it is a good description of reality: It is hard to imagine that *every inhabitant* of an urban area, including the homeless and recent immigrants, could be reached in half an hour, say.

Three different approaches were taken to calibrating the constants in the differential equation. Suppose we let Θ denote the vector of unknown parameters. In the constant urgency model, Θ would equal (U_1, U_2) and in the decreasing urgency model it would equal (λ, γ). The three approaches are:

Least squares fit: The constants were adjusted so as to minimize the sum of squared deviations of the expert estimates $A_{ij\tau}$ from the solution $A_i(\tau;\Theta)$ to the differential equation. Thus, for each subpopulation i, the constants in Θ were chosen to minimize

$$\sum_{j \in N_i} \sum_{\tau \in L} (A_{ij\tau} - A_i(\tau;\Theta))^2$$

where N_i is the set of experts that provided estimates for the i-th subpopulation and $L = \{0.5, 1, 5, 30\}$ is the set of lead-times the experts were asked to consider. The least-squares parameter estimates for the constant urgency model are shown in table 5.8 and least-squares estimates for the decreasing urgency model are shown in table 5.11.

Maximum likelihood fit: This approach requires an explicit model of the probability distribution of the expert estimates $A_{ij\tau}$, conditional on the true values of the constants in the differential equation, i.e., the likelihood function. Then, the parameters in Θ are chosen to maximize the likelihood function

$$L(\Theta) = f\left(\{A_{ij\tau}\}_{j\in N_i,\tau\in L};\Theta\right)$$

The resulting parameter estimates are shown in tables 5.9 (for the constant urgency model) and 5.12 (for the decreasing urgency model). The likelihood function has a parameter σ, which is a measure of how widely scattered the expert estimates are, and needs to be added to the parameter vector Θ.

Bayesian estimation: If one is willing to assign a prior probability distribution $p(\Theta)$ on the constants in the differential equation, then this prior distribution can be updated via Bayes theorem to a posterior distribution that is conditional on the expert responses. A likelihood function is required in this case also. Tables 5.10 and 5.13 list posterior expected values for the constant and decreasing urgency models, respectively, under a uniform prior distribution[6].

A detailed description of the likelihood function adopted and the assumptions on which it is based is in chapter 4. Note that in tables 5.8 to 5.13 (which appear in Appendix 5E), the three subpopulations that involve "people in cars" have been combined into one, since the experts were asked to estimate the effectiveness of communication strategies for "people in cars," regardless of whether the car was on an open freeway, near an overpass, or on an urban street.

A few general patterns emerge from inspection of tables 5.8 to 5.13. If we consider the constant urgency model first, we see that the rate of person-to-person transmission of the warning information is estimated at zero for all subpopulations under the least-squares and maximum likelihood estimation schemes, and the estimates are close to zero under the Bayesian scheme. The implication is that the experts seem to think that for lead-times as short as the ones they were asked to consider, most people that hear the warning will hear it from a primary source[7], rather than from another person that has already heard the warning some time ago; a proposition that appears plausible. This observation, along with the fact that the constant urgency model does not provide a particularly good fit to the expert estimates, motivated the development of the decreasing urgency model.

The maximum likelihood and Bayesian parameter estimates are generally very close to each other. This is to be expected, since the maximum likelihood estimates coincide with the mode, or maximum, of the likelihood surface, while the Bayesian estimates (posterior means) coincide with the "center of mass" of the likelihood surface, since a uniform prior distribution was assumed. A parameter σ, which is a measure of the dispersion among the expert estimates, was estimated under the maximum likelihood and Bayesian schemes. The estimates of σ are fairly stable across subpopulations, models, (constant or decreasing urgency) and estimation schemes (maximum likelihood or Bayesian), ranging from 1.15 to 3.46. This suggests that the scatter of the expert estimates is not radically different for

[6]Technically, an improper prior was used, i.e., $p(\Theta)$ was set equal to a uniform distribution over the space $[0,\infty)^3$ (since both the constant and decreasing urgency models have three parameters, all of which are restricted to non-negative values). In the numerical implementation, the likelihood function was computed over a grid whose dimensions were adjusted until the region where the likelihood function was greater than an arbitrary, small value had been encompassed, and then the values of the likelihood function were normalized to turn it into (a discrete approximation to) the posterior distribution

[7]Possibly after someone else has directed their attention to it, cf. our earlier discussion how we interpret the phrase "person-to-person" communication

Subpopulation	Median time to hear warning (minutes)					
	Constant urgency model			Decreasing urgency model		
	LS	ML	Bayes	LS	ML	Bayes
People at home, awake	0:30	1:24	1:21	0:19	0:17	0:19
People at home, asleep	0:49	2:06	2:01	0:50	0:45	0:50
People in offices	2:17	0:55	0:52	0:04	0:05	0:07
People in cars	0:16	3:09	2:56	2:05	2:10	2:29
Children in school	2:09	0:58	0:55	0:07	0:02	0:05
People outdoors	1:38	1:23	1:13	0:04	0:07	0:26

Table 5.2: Estimated median (min:sec) time until a randomly chosen individual hears an earthquake warning that is broadcast using sirens, under two different models, three different estimation procedures, and the expert data.

different subpopulations.

Table 5.2 shows estimated median times for a randomly selected individual to hear a warning, for each subpopulation and under each of the two models and three estimation schemes. Admittedly, the estimates cover a wide range. The most extreme example is the "children in school" subpopulation, where the estimated median time to hear a warning ranges from 0.03 minutes to 2.15 minutes. But as mentioned earlier, the constant urgency model does not fit the expert estimates very well. The marked difference between the least-squares estimates on the one hand and the maximum likelihood and Bayesian estimates on the other for the constant urgency model is one reflection of this lack of fit. In contrast, all three estimation methods result in similar estimates for the decreasing urgency model. We will focus on the numerical estimates from the decreasing urgency model from now on.

The rankings of subpopulations according to how long it takes to hear the warning are rather stable. The "people in offices" and "children in school" subpopulations are usually ranked first and second, i.e., the median time to hear the warning is shortest for these two subpopulations under most of the estimation schemes. This is perhaps a reflection of the fact that schools and offices are likely to have well-maintained alarm systems and that people are likely to be relatively alert when they are at school or at the office. The estimated median times to hear warning (under the decreasing urgency model) range between two and eight seconds for these two subpopulations. Next come the "people at home, awake" and "people outdoors" subpopulations, typically ranked third or fourth with estimated median times to hear warning between three and 27 seconds. Then come "people at home, asleep," with estimated median times to hear warning that exceed the corresponding estimates for "people at home, awake" by about half a minute; it seems plausible that most people would be roused from sleep by sirens in that timespan. "People in cars" have the longest estimated median times to hear warning in all cases but one, with estimates between two and two and a half minutes. Remember that these estimates are made under the assumption that sirens are the sole means of communication. The sirens may not be audible inside a car whose windows are closed and whose driver is listening to the radio. On the other hand, the warning would probably be broadcast on the radio also, so the median times to hear warning for "people in cars" are likely to be underestimated.

In the calculation of relative death rates later in this section, we will use the decreasing urgency model with the Bayesian parameter estimates (see table 5.13 and the last column of table 5.2). The estimated relative death rates will of course be different if the constant urgency model or another estimation scheme is used. However, the overall relative death rate estimates based on different models and estimation schemes differ by less than ten percent of the baseline death rate for all lead-times less than thirty minutes.

Assuming that a person hears an earthquake warning, believes it, and decides to attempt an appropriate action in response to it, there is the question of whether the person will complete the action before the earthquake occurs. We now describe a model for how long it will take people to complete such actions.

5.4.2 Distribution for the Time to Complete Action, T_C

The survival-enhancing actions suggested to the experts and by the experts cover a wide spectrum. However, most of the actions involve physical movement, for example getting under a sturdy desk or getting outside a building. Members of certain segments of the population will take significantly longer than average to carry out such actions, because of disability, advanced age, or other reasons. This observation is the basis for a model of the probability distribution for the time $T_C(i, \tau)$ that a randomly chosen individual in subpopulation i will need to complete an action whose purpose is to improve chances of survival in an earthquake that is forecasted with τ minutes of lead-time. Assume that

- There are two kinds of people: "slow" and "quick". The fraction of "slow" (mobility-impaired) people is p.
- The time it takes a randomly chosen "slow person" to complete the specified action is an exponential random variable with mean $1/\lambda_1$. The same is true for a "quick person," but the mean time $1/\lambda_2$ for "quick people" is shorter than for "slow people."

Under these assumptions, the probability density function for T_C has the form

$$f_{T_C}(t) = p\lambda_1 e^{-\lambda_1 t} + (1-p)\lambda_2 e^{-\lambda_2 t};\ t \geq 0$$

The basic idea underlying this model has substantial intuitive appeal: Categorize people based on their general mobility, then postulate that people have different expected times to complete action depending on their mobility category but allow for variation around the expected value due to varying circumstances and unpredictable events. Many of the experts will surely have applied simple variants of this idea in estimating the fraction of people that complete a specified action under specified circumstances. Nevertheless, let us mention two caveats that should be kept in mind when interpreting the model and the parameter estimates to be presented shortly. First, using only two categories ("slow" and "quick") is perhaps an oversimplification. Second, the exponential distribution is a bit suspect, since it implies that the most likely time to complete an action is zero minutes, while the time required to complete an action such as getting out of a building will surely always be above some positive threshold. However, if the threshold is small, the exponential distribution may provide an adequate approximation. While these two caveats should be kept in mind, it should be recognized that the model could easily be extended to include more than two categories of mobility and distributions other than exponential.

Subpopulation	p^{LS}	$1/\lambda_1^{\mathrm{LS}}$	$1/\lambda_2^{\mathrm{LS}}$
People at home, awake	16.9%	19:20 min.	0:14 min.
People sleeping	17.3%	16:30 min.	0:31 min.
People in offices	15.6%	44:30 min.	0:00.1 min.
Cars on bridge/overpass	17.1%	38:10 min.	0:46 min.
Cars on open freeway	7.5%	21:50 min.	0:19 min.
Cars on street	17.9%	28:50 min.	0:05 min.
Children in school	12.4%	26:00 min.	0:10 min.
People outdoors	17.5%	35:10 min.	0:33 min.

Table 5.3: Estimated parameter values for the probability distribution of the time to complete the best action, under the circumstances. The first column lists estimates of the fraction p of mobility-impaired or "slow" people, the second column lists the estimated mean time to complete action $1/\lambda_1$ for "slow" people, and the third column lists the estimated mean time to complete action $1/\lambda_2$ for "quick" people. Least-squares estimation was used.

The data consist of estimates $C_{ij\tau}$ of the probability of completing the best action, under the circumstances, for a given subpopulation i and lead-time τ. We will assume that

$$C_{ij\tau} = \text{j-th expert's estimate of the fraction of people in subpopulation i that complete best action in less than τ minutes.}$$

$$= \text{j-th expert's estimate of } \Pr\{T_C(i,\tau) < \tau\}$$

Put differently, we assume that experts "set the clock at zero" when the person becomes aware of the warning and thus estimate the probability $\Pr\{T_C < \tau\}$ that the action can be completed in τ minutes, rather than the more difficult to grasp conditional probability $\Pr\{T_C + T_A < \tau \mid T_A < \tau\}$ that a person who hears the warning before the earthquake will be able to complete the action before the earthquake occurs.

Different experts chose different actions as the most appropriate one under the circumstances, and experts usually chose different actions for different lead-times, but this will not be accounted for in the analysis. The main reason for this limitation is insufficient sample sizes for any one action. Note, however, that this procedure is in accord with the strategy of obtaining optimistic estimates of relative death rates. Tables 5.16 and 5.17 (in appendix 5G) show the frequency with which different actions were chosen as most appropriate, for different subpopulations.

The parameters p, λ_1, and λ_2 were estimated using a least-squares approach, i.e., by minimizing the following sum of squared deviations with respect to p, λ_1, and λ_2, for each subpopulation i:

$$\sum_{j\in N_i}\sum_{\tau\in L}(C_{ij\tau} - \Pr\{T_C < \tau\})^2$$

The resulting parameter values are shown in table 5.3.

We see from table 5.3 that the estimated proportion of "slow" people ranges between 7.5% and 17.9%, which is perhaps about right for the fraction of the population that is disabled or mobility-impaired for other reasons. According to [RA93],

> ... between 12 and 14 percent of the [U.S.] working population are disabled and ... about another 10 percent have some activity limitation (U.S. Bureau of Census, 1980). For persons with disabilities related to motor functioning, the estimate is around 23.4 million (9.6 percent of U.S. population) (Elkind, 1990). ... In relation to earthquake situations, Aroni and Durkin (1985) reported that 50 out of the 133 injured persons they contacted after the 1983 Coalinga (California) earthquake had some type of disability. At nearly 38 percent, this percentage is larger than the overall rate of disability for that community.

The statistics reported for the Coalinga earthquake do suggest that disabled people take longer to move out of harm's way. However, the orders of magnitude differences between the estimated mean times in table 5.3 to complete the best action for slow and quick people are rather extreme, and may suggest that a model with more categories or different distributional assumptions would be more appropriate. In particular, the estimated mean time of a tenth of a second for "quick" people in offices to complete the best action is difficult to take seriously, indicating that the model may be too simplistic to summarize the views of the experts.

However, the relative stability of the estimated fraction of slow people across subpopulations is interesting. For five out of eight subpopulations the estimated fraction of slow people is close to 17%. There are plausible explanations for the three deviations, which we now outline. In all three cases, we postulate that "slow" can have different meanings under different circumstances.

The lowest estimated fraction of "slow" people, 7.5%, is for people in cars on the open freeway. There, the action most often chosen as best was "stop the car and stay in the car," and it may be that people that would be "slow" in completing an action requiring movement of the whole body would be able to complete this action as quickly as anyone. After all, if a disabled person is driving a car, one would expect the car to be specially equipped to allow the driver to react as quickly as other drivers to dangers on the road.

The subpopulations of people in offices and children in school have estimated fractions of 15.6% and 12.4%, respectively. In both cases, those that take unusually long to complete the best action would typically take long because of either congestion or physical distance, say from the tenth story of an office building to an open area outside the building, as opposed to mobility-impairedness.

The methods we have described, for estimating the model parameters using data from the questionnaires, were developed concurrently with the design of the questionnaire and were not completed until after the first questionnaires had been mailed. With the benefit of hindsight, some aspects of the questionnaire that could be improved upon were identified, and we will now discuss those aspects briefly.

5.4.3 Calibrating the Models: Lessons Learned

Two changes would have facilitated estimation of the distribution of T_A and T_C:

- Instead of asking "given a lead-time τ, what fraction of the population do you think would learn of the prediction (or: complete the most beneficial action) in time" one could have asked "how much time do you think will elapse before half the population will have heard the warning (or: completed the most beneficial action)." The experts could then be instructed to concentrate their attention on the "first half" of the population and estimate how much time would elapse before half of that half would hear the warning or complete the action, and so on. In essence, the experts would be asked to estimate the quantiles of the distributions for T_A and T_C rather than values taken on by the cumulative distribution function at various times. This change would have facilitated the processing of the data under the model-based approach. Perhaps, it would also have made the experts' task easier.

- The experts were instructed to estimate the probability of completing the most beneficial action, given the lead-time τ and assuming that the person has heard the warning, believed it, and decided to attempt the action. Taken literally, these instructions might suggest that the quantity to be estimated is $\Pr\{T_A + T_C < \tau \mid T_A < \tau\}$. We consider it unlikely that any of the experts were in fact thinking of such a complicated object; more likely they were simply estimating $\Pr\{T_C < \tau\}$, the probability that the time needed to complete the action is less than τ minutes, and the analysis assumes this to be the case. But in retrospect, it would have been better to unambiguously instruct the experts to estimate the latter quantity.

The model-based relative death rate estimates are weighted averages of the baseline and revised death rates, for each subpopulation. The distributions for the two random variables T_A and T_C, whose estimation we have discussed, provide us with the information needed to compute the weight $\Pr\{T_A + T_C < \tau\}$ to be assigned to the revised death rate. But before presenting the results of this computation we need to resolve a difficulty that arose in estimating the revised death rates.

5.4.4 Revised Death Rate Estimates

In a few instances, an expert would give zero as an estimate of the revised death rate $D''_{ij\tau}$ (zero was never given as an estimate of the baseline death rate). These estimates would have caused problems in using the modified geometric mean (see chapter 4) to combine the expert estimates for a particular subpopulation: a single zero estimate would cause the modified geometric mean to equal zero.

If the revised death rate is interpreted in an expected value sense, such that $D''N$ is the expected number of deaths in a population with N members, each of whom faces the death rate D'', then it is difficult to convince oneself that an expert truly believes the revised death rate to be zero. It seems more plausible that an expert would write zero because he considers it overwhelmingly likely that no one would die under the circumstances in question. Such a belief is not inconsistent with a small but positive probability that one or more people will die. In other words, the expert may consider the death rate to be zero for all practical purposes, without completely ruling out the possibility of deaths. One may think of the expert as having a mental threshold, with values below the threshold replaced by zero.

Experts that do give zero death rate estimates provide partial information about what their thresholds are, since every non-zero estimate they give must be above the threshold. Let L_j be the smallest non-zero death rate estimate given by expert j. The zero estimates given by the j-th expert were replaced with $L_j/10$. The factor of 10 is arbitrary, of course, but results obtained using a factor of two or a factor of 100 instead turn out to be almost indistinguishable anyway.

As an example, table 5.4 shows the death rate estimates provided by one expert. In this case, the lowest non-zero death rate estimate given was 0.0005% and the zero entries were replaced with $0.0005\%/10 = 0.00005\%$.

Lead-time	People at home, awake	Children in school	People outdoors
No warning	0.001%	0.001%	0.001%
30 seconds	0.0005%	0.0005%	0.0005%
1 minute	0	0.0005%	0
5 minutes	0	0	0
30 minutes	0	0	0

Table 5.4: An example of an expert's estimates of baseline and revised death rates. The smallest non-zero death rate given by this expert was 0.0005%. The zero entries were replaced with $0.0005\%/10 = 0.00005\% = 0.5 \times 10^{-6}$.

5.4.5 Relative Death Rates

Figure 5-10 shows estimated relative death rates as a function of warning lead-time, for each of the eight subpopulations. The solid curves in this figure are estimates computed under the model-based approach described in this section, while the 'x's are point estimates based on the face-value approach described earlier in this chapter.

The estimated relative death rates for individual subpopulations were combined with estimated population fractions (averaged over the day) in the same way that the face-value estimates were to estimate the overall relative death rate as a function of warning lead-time. The resulting curve is shown in figure 5-11. For comparison, the face-value estimates are also shown (as 'x's) in this figure. The agreement between the two approaches is even closer than it was for the individual subpopulations, indicating that the minor discrepancies observed in figure 5-10 tended to cancel each other out.

In the next section, we will interpret the model-based results in figures 5-10 and 5-11 and compare them to the face-value results.

5.5 Comparison of Results From the Two Approaches

The face-value and model-based estimates agree closely for most subpopulations(see figure 5-10). The greatest disagreement between the two approaches is for the "cars on open freeway" and "cars on urban street" subpopulations and this can be attributed to the small number of experts that were asked to provide estimates for these subpopulations. There was no prior reason to expect such close harmony between the two approaches to summarizing the expert judgment. On the one hand, this reinforces the conclusions reached based on the face-value estimates, and on the other hand, one's confidence in the models espoused in this section is improved.

The dichotomy between the *inside* subpopulations ("people at home" awake or asleep, "people in offices," and "children in school") and the *outside* subpopulations ("people in cars" and "people outdoors") is borne out by the model-based approach. As was the case for the face-value approach, the relative death rate profiles in figure 5-10 indicate that between one and five minutes suffice to halve the death rate and between 10 and 30 minutes are needed to reduce the death rate by three quarters for the "inside" subpopulations. For the "outside" subpopulations, the model-based estimates are a bit less optimistic than the face-value estimates are: At least 10 minutes are needed to halve the death rate and more than 30 minutes are required to reduce it by three quarters. Still, the "outside" relative

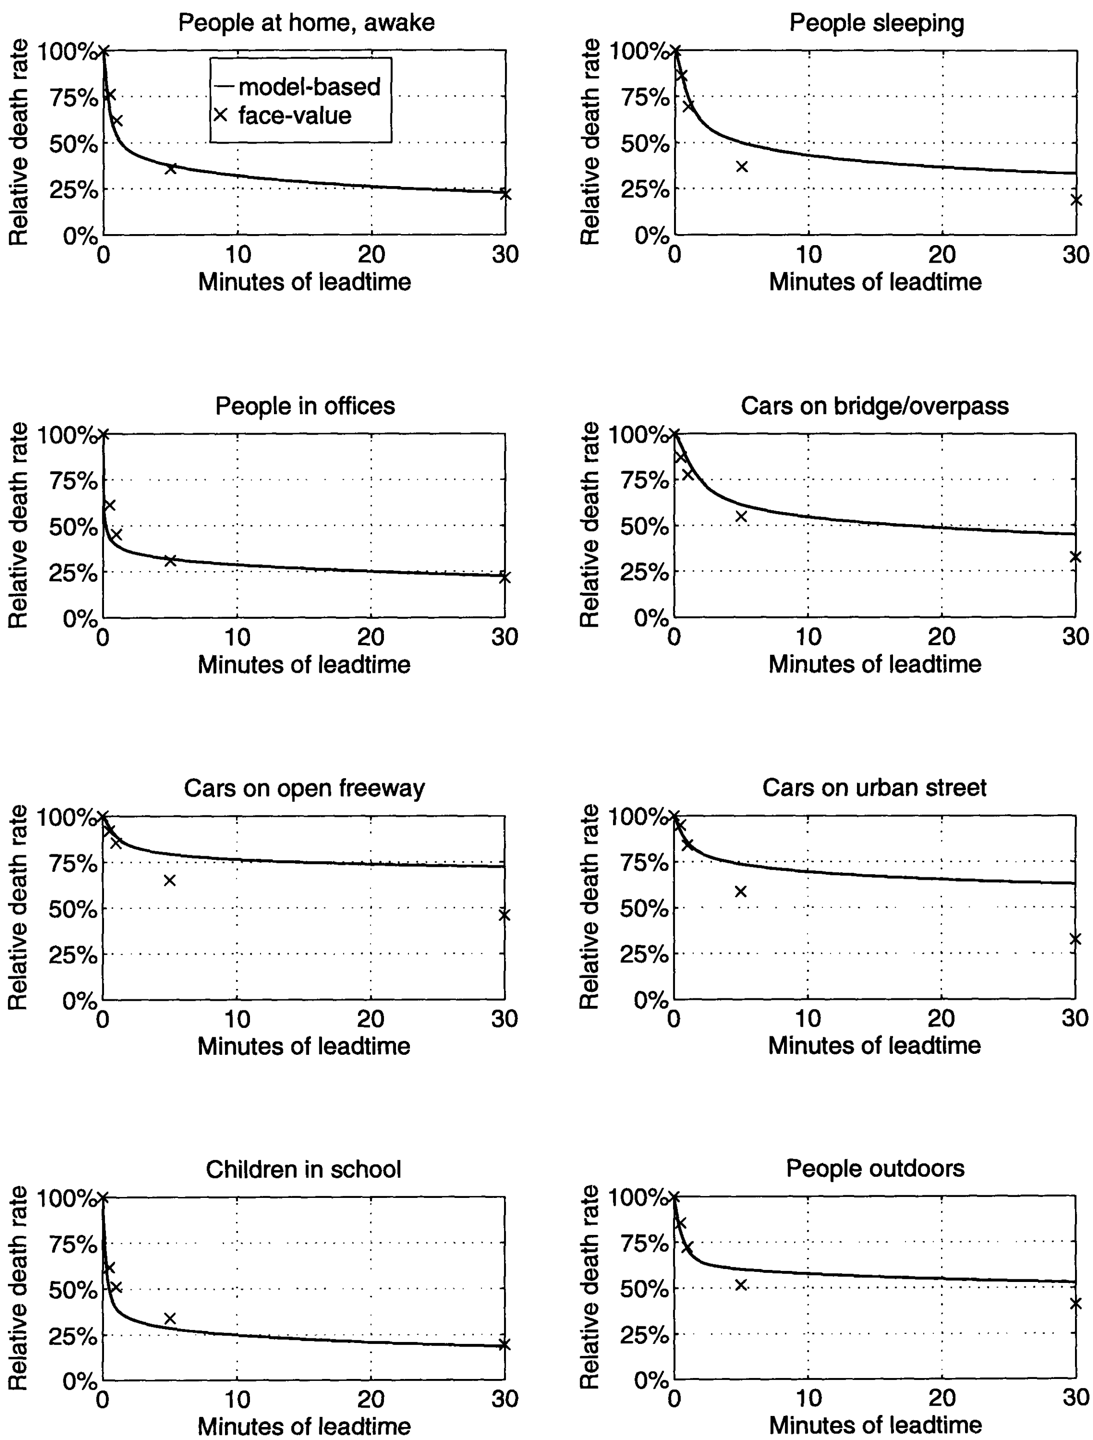

Figure 5-10: Model-based estimates of the relative death rate as a function of warning lead-time (solid curves), for each subpopulation. Face-value estimates are shown also, as 'x's.

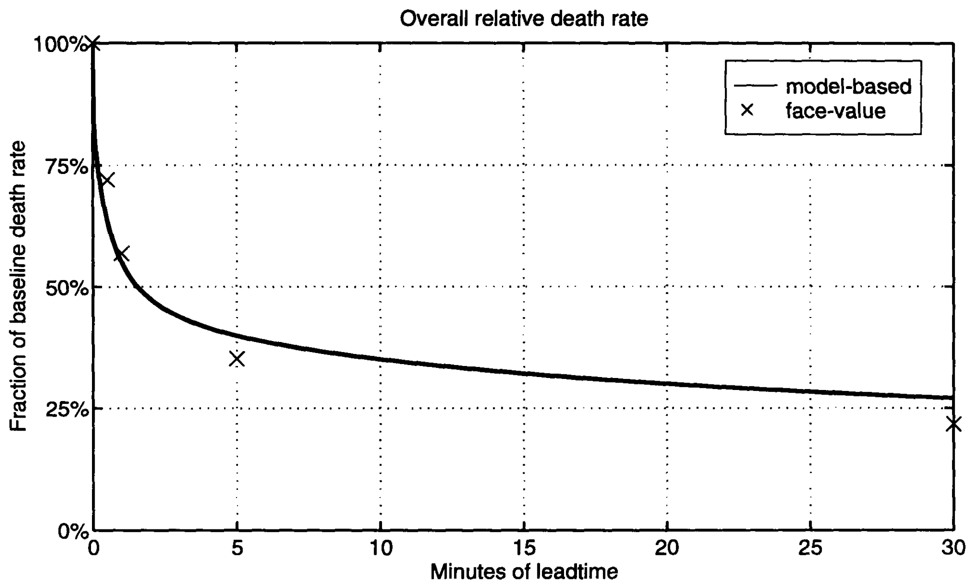

Figure 5-11: Model-based estimate of the overall relative death rate as a function of warning lead-time (solid curve). The 'x's represent face-value estimates of the relative death rate which are shown for comparison.

death rate profiles (those for people in cars or outdoors) are as similar to each other as they were under the face-value approach.

The lead-time needed to halve the estimated overall death rate, according to the model-based curve in figure 5-11 is around one and a half minute (compare to three minutes for the face-value approach) and 30 minutes suffice to reduce the estimated overall death rate by three quarters (compare to 25 minutes for the face-value approach). Some of the discrepancy between the face-value and model-based approaches can be attributed to the straight lines that were used to interpolate between the four face-value point estimates, while the model-based approach delivered a continuous convex curve. So to be fair, the difference between the estimated "halving times" under the two approaches is even smaller.

This concludes the analysis of the survey responses. Before summarizing the main findings of this analysis, we will assess the consequences of two of the optimistic assumptions we have made.

5.6 Optimistic Assumptions and Their Consequences

Two of the optimistic assumptions on which both the face-value and model-based approaches are based are (1) all persons who hear the warning will immediately attempt an action and (2) the time until the earthquake occurs is known with certainty. In this section we show how our estimates would change under a particular departure from the first assumption and indicate how our results might be used as input to an analysis that relaxes the second assumption.

5.6.1 Sensitivity to the Assumption that People React Immediately to a Warning

As discussed in the introduction to this chapter, when the absolute death rate estimates $D_{ij\tau}$ are computed an implicit assumption is made that a person that belongs to subpopulation i and hears a warning which is first broadcast τ minutes before the earthquake occurs will attempt, without delay, to act in a way to increase her chances of survival. Suppose we assume instead that any person who hears the warning before the earthquake occurs will act immediately with probability P, and will not act at all with probability $1 - P$. That is, if X is a randomly chosen person from a particular subpopulation, then we let

$$P = \Pr\left\{ \begin{array}{l} X \text{ understands and believes the warning,} \\ \text{and attempts to act without delay} \end{array} \middle| \begin{array}{l} X \text{ hears the} \\ \text{warning in time} \end{array} \right\}$$

and the computation of the absolute death rate (under the face-value approach) estimate is modified to

$$D_{ij\tau} = A_{ij\tau} P C_{ij\tau} D''_{ij\tau} + (1 - A_{ij\tau} P C_{ij\tau}) D'_{ij} \tag{5.2}$$

Of course, assuming that P is a constant is a bit simplistic. In particular, the fraction of people that act on a warning can be expected to depend strongly on the available lead-time, and probably also on the situation the person is in, i.e., P would be different for different subpopulations. But assuming P to be constant allows one to observe in a simple manner the effect on the relative death rate curve of varying the fraction of people that act without delay, as we shall demonstrate.

Changes in P affect the absolute death rate estimate $D_{ij\tau}$ (for a particular subpopulation, lead-time, and expert) linearly, as equation 5.2 shows. This linear effect propagates all the way to the overall relative death rate estimate $\hat{G}_\tau$:

$$\hat{G}_\tau(P) = (1 - P) + P\hat{G}_\tau(P = 1)$$

The same linear dependence results for the model-based approach. The effect is perhaps better expressed in terms of the estimated *reduction* in death rate, $\hat{R}_\tau$:

$$\hat{R}_\tau(P) = 1 - \hat{G}_\tau(P) = P\hat{R}_\tau(P = 1)$$

The quantity $\hat{R}_\tau(P = 1)$ is the estimated increase in chances of survival possible if everyone acts without delay after hearing the warning, i.e., under our best-case assumption. We see, as one might have expected, that the life-saving potential of a last-minute earthquake warning depends strongly on how quickly people decide to act on the warning, since by definition there are only a few minutes available to ponder the warning.

An interesting and challenging problem is to estimate P and how it depends on the individual and the circumstances the individual is faced with. However, a more important task, from the perspective of achieving the potential benefits of earthquake warnings, is to examine how the fraction P of people that act without delay upon hearing a warning can be increased through education and training. After all, P is not a never-changing physical constant to be determined; rather it is a measure of how human beings with free wills make use of potentially life-saving information.

The second half of this section describes how the results of the survey of experts could be helpful in an analysis that takes into account the considerable uncertainty that will likely be associated with the lead-time τ of earthquake warnings, if they were to become available.

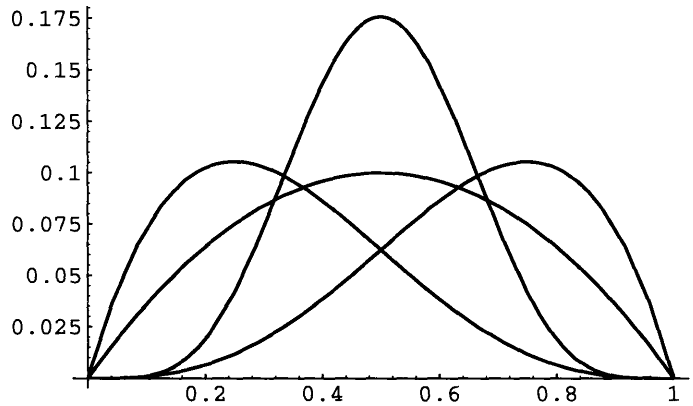

Figure 5-12: Contrived numerical example: The four utility functions shown are beta functions. From the graph one can see that the action corresponding to one of the curves is never the optimal one if τ is known with certainty. But if τ is uniformly distributed between zero and one, then the expected utility of that action is 0.067, while the expected utility of the two skewed functions is 0.05 and that of the symmetrical function with the highest mode is 0.065.

5.6.2 Accounting for Uncertain Warning Lead-Times

In this subsection we will indicate how uncertainty about the value of τ can be taken into account, i.e., how the face-value and model-based analyses in the first two sections of this chapter could serve as a basis for such an analysis. While we will not carry out such an analysis, we will mention new possibilities that emerge and give an example of the kinds of calculations that would be needed.

Suppose that a complete analysis of how a person in some setting should react to the information that "a major earthquake is *certain* to strike in τ minutes" has been performed. In particular, suppose that the choice facing the person is which of k courses of actions to take, and that utility functions $V_i(\tau)$ have been determined for each of those k courses of action. Then the optimal reaction to a deterministic forecast is simply to follow that course of action i for which the utility V_i is highest, given the (deterministic) value of τ.

Borrowing terminology from Decision Analysis, we will refer to the lead time τ as the unknown "state of nature;" it's probability distribution is assumed given. The optimal course of action is then the one that maximizes the expected utility $E[V_i(\tau)]$, with the expectation taken with respect to the probability distribution of τ.

When the lead-time is random, a course of action which would never be chosen if τ were known with certainty could turn out to be the optimal action. Numerical and graphical examples are easy to construct, see for example figure 5-12. However, the fact that such examples are easy to construct does not mean that such circumstances are likely to arise in reality. We now consider an example where the emphasis is on plausibility rather than on demonstrating possible pathologies.

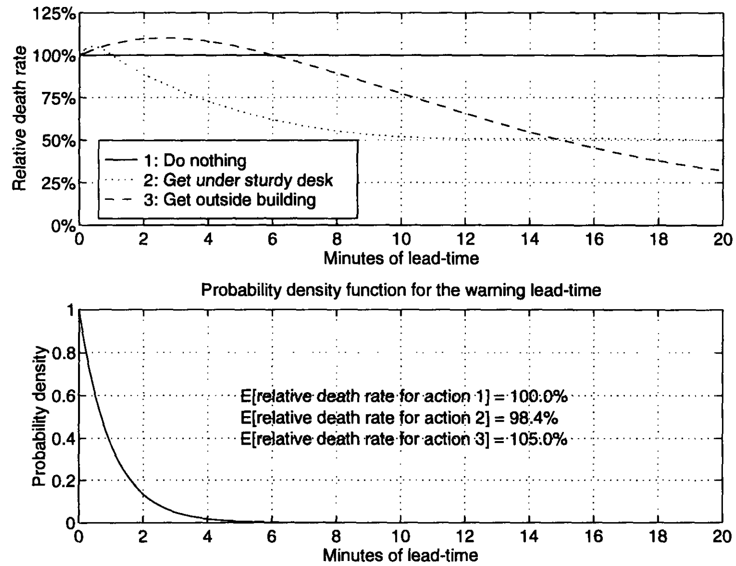

Figure 5-13: Hypothetical relative death rates for "people at home, awake" under three different actions. Also shown is a hypothesized probability density function (exponential distribution with a mean of one minute) and the expected relative death rate for each action under this distribution.

Example: Action Choices for "People at Home"

Consider the choices available to an individual who is at home, awake, and has just learned that a major earthquake will strike in τ minutes. Assume that the individual has information about the probability law for the lead-time τ, that he wishes to maximize his expected chances of survival, and that the following three options are open to him:

1. Do nothing, i.e., stay put and wait for the quake to occur.
2. Attempt to get under a nearby sturdy desk before the earthquake.
3. Attempt to get outside the house, to an open area free of overhead wires before the shaking starts.

Maximizing the chances of survival is equivalent to minimizing the relative death rate. Figures 5-13 and 5-14 show two hypothetical scenarios. In both cases, the relative death rate for the "do nothing" action equals 100% of the baseline rate no matter what the lead-time is. The relative death rate for the "get under sturdy desk" action may be above 100% for lead-times less than one minute, because of the possibility of bumping one's head into the desk if the earthquake happens to occur as one is bending down. As the lead-time increases, the relative death rate for those that attempt to get under a sturdy desk is assumed to decrease towards 50%. The relative death rate curve for the "get outside" action has a similar shape: it first ascends above 100%, due to the danger of being on the way out when the quake occurs, and then descends down to a quarter of the relative death rate for those that do nothing.

Figure 5-13 shows what happens if the lead-time has an exponential distribution with mean equal to one minute. The expected relative death rate is lowest for the "get under

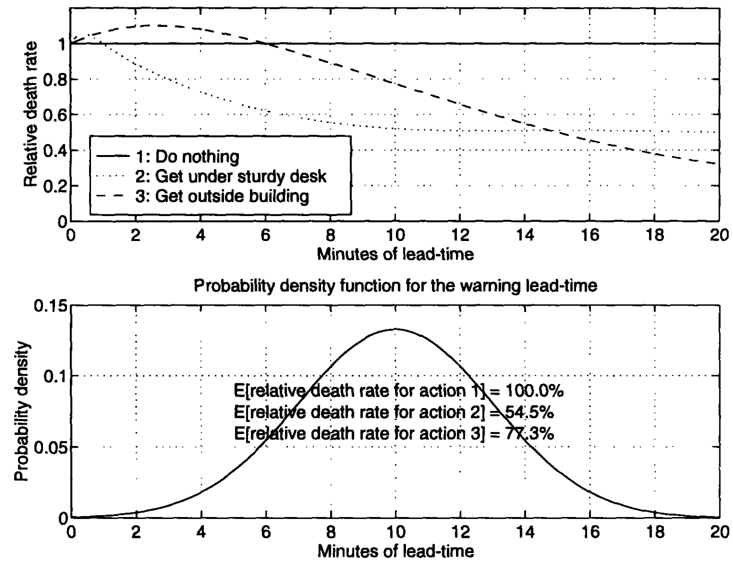

Figure 5-14: Hypothetical relative death rates for "people at home, awake" under three different actions. Also shown is a hypothesized probability density function (normal distribution with a mean of 10 minutes and standard deviation of three minutes) and the expected relative death rate for each action under this distribution.

sturdy desk" action; it is 98.4% of the death rate facing those that do nothing in response to the warning. For comparison, if the lead-time has the deterministic value of one minute, then the "do nothing" and "get under sturdy desk" options are tied at 100%. But for lead-times longer than one minute, the relative death rate curve for the "get under sturdy desk" option descends sufficiently rapidly and there is sufficient probability left in the lead-time distribution to break the tie in favor of attempting to get under a sturdy desk.

In figure 5-14, the lead-time distribution has been changed to a normal distribution, with a mean of 10 minutes and standard deviation of three minutes. Again, the "get under sturdy desk" action is the winner, with an expected relative death rate of 54.5%, which is quite close to what the relative death rate under that option would be if the lead-time was ten minutes, with certainty.

Of course, this is an artificial example, but the relative death rates were chosen to be at least plausible, given the results of the survey of experts. If nothing else, the example reminds one that while it is mathematically possible for an analysis that assumes the lead-time to be deterministic to be misleading, this will not necessarily happen.

5.7 Summary of Main Findings

In this chapter, we have combined expert judgment with population distribution estimates to assess how many lives could potentially be saved if warnings before major earthquakes were possible. The main findings of the analyses in this chapter are:

- Two different procedures for processing the expert assessments, the face-value and model-based approaches, led to estimates of the overall reduction in risk possible with earthquake warnings that agree closely. Model-based and face-value estimates of the warning lead-time needed to halve the overall death rate were 1.5 and 3 minutes, respectively, while 30 and 25 minutes of lead-time were estimated as sufficient to reduce the death toll by three quarters under the two approaches.

In fact, the agreement between the two approaches is greater than these numbers suggest; the discrepancy is in part an artifact of a linear interpolation scheme used in the face-value approach.

- Results of the face-value and model-based approaches were also similar for individual subpopulations. Under both approaches, a dichotomy between people inside buildings and outside buildings (in their cars or outdoors) was apparent, with longer lead-times needed to reduce the death toll by a given proportion for people outside buildings.

- The overall death toll reduction estimates, under the face-value approach, were relatively insensitive to perturbations to the population distribution estimates used. However, greater sensitivity to the means of communications assumed was observed: When estimates based on the most effective means of communication (as judged by each expert) were used, instead of estimates that assumed sirens to be the sole means of communication, the lead-time needed to halve the estimated death rate dropped from 3 to 1 minute and 10 rather than 25 minutes were estimated as sufficient to reduce the number of deaths by three quarters.

- Evidence for heterogeneity among the experts, in the sense of identifiable optimists and pessimists, was inconclusive. Nevertheless, the experts were categorized as optimists and pessimists, based on their responses, and separate estimates were computed for the two groups. The lead-time needed to halve the death toll was estimated at 1 minute using the optimists and at 5 minutes using the pessimists. The corresponding estimates for a three-quarters reduction in the number of deaths were 10 minutes (based on the optimists) and more than 30 minutes (based on the pessimists).

- Estimates of baseline death rates that might be expected if a major earthquake were to strike an urban area in California, based on the expert judgment, ranged from 203 per 100,000 (for people in cars on, under, or near overpasses or bridges) to 3 per 100,000 (for people in cars on the open freeway). For people at home, awake, the estimate was 21 per 100,000. Using estimates of population shifts over the course of a day, we estimated that the overall baseline death rate might range from 25 per 100,000 (in the early morning hours) to 50 per 100,000 (during the late afternoon). The factor of two difference is smaller than some other commonly quoted estimates (see chapters 3 and 6) suggest.

Appendix 5A: Relative Death Rate Histograms

Figures 5-15 to 5-17 show histograms of the relative death rates for six of the eight subpopulations. Histograms for the remaining two subpopulations ("people at home, awake" and "people outside") are shown in figure 5-4.

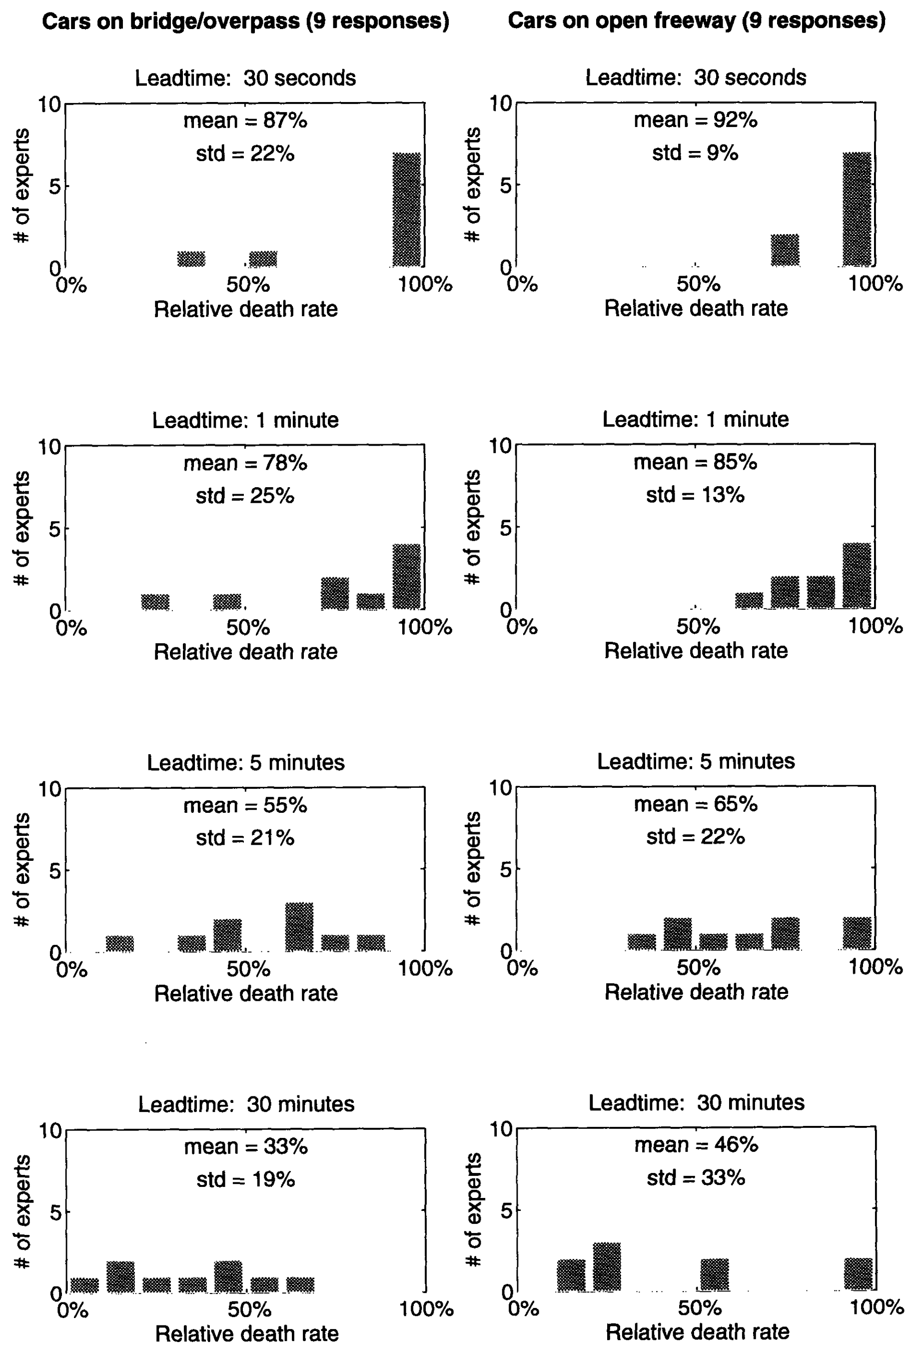

Figure 5-15: Histograms of relative death rates for "People in cars on a bridge or overpass" and "People in cars on open freeway."

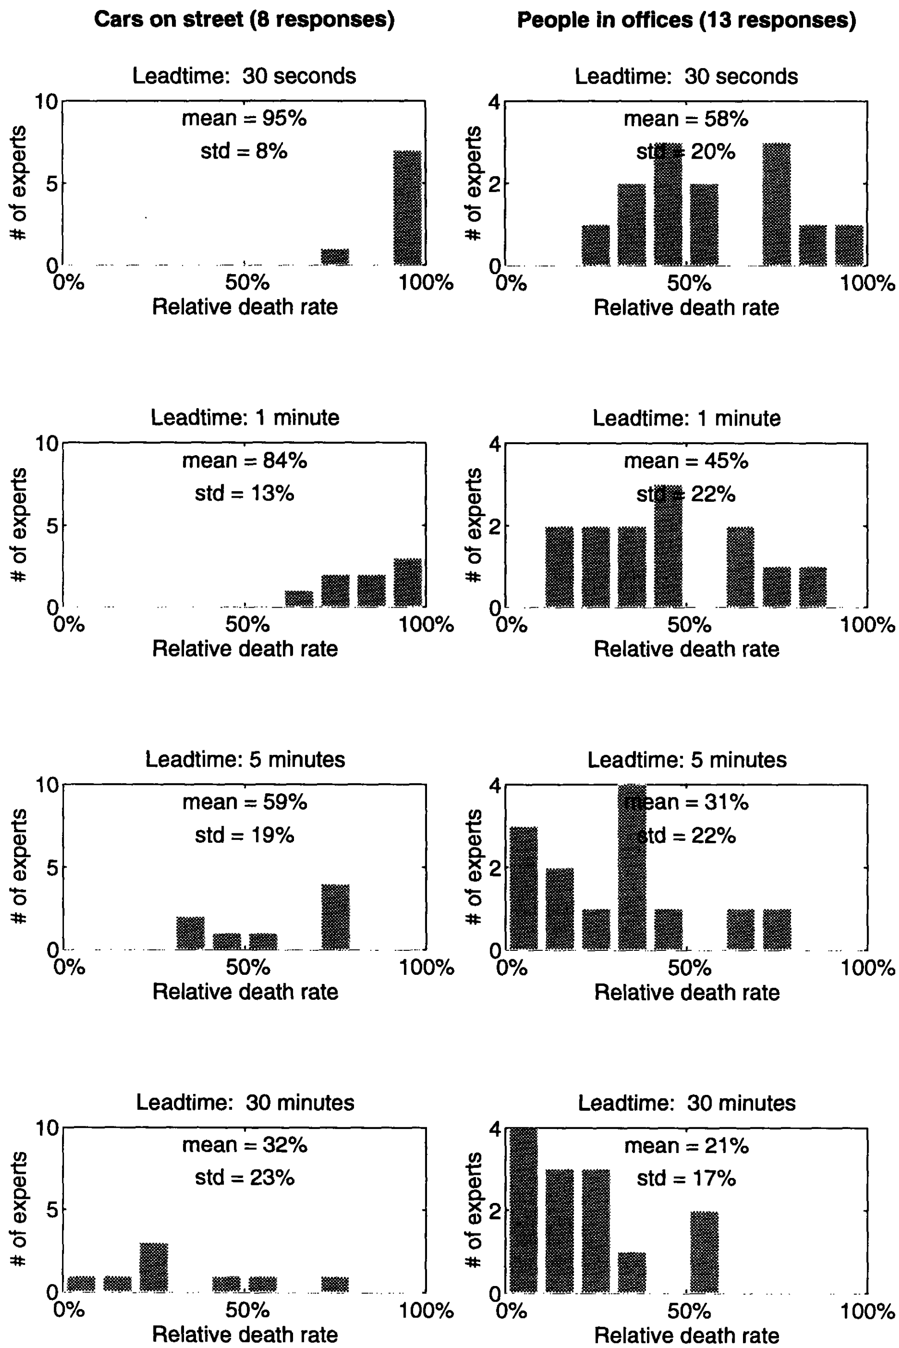

Figure 5-16: Histograms of relative death rates for "People in cars on urban street" and "People in offices."

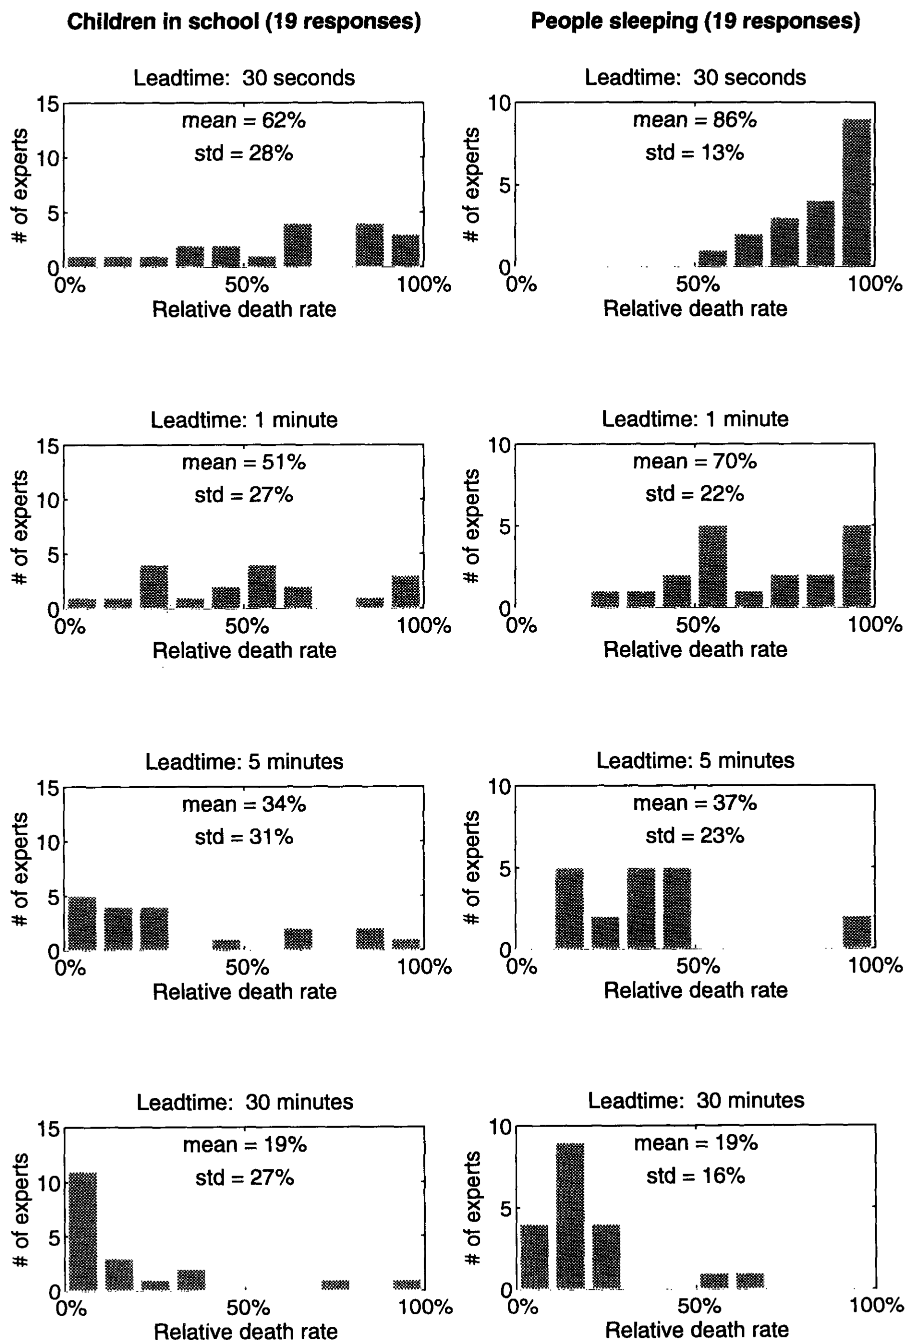

Figure 5-17: Histograms of relative death rates for "Children in school" and "People at home, asleep."

Appendix 5B: Confidence Intervals for Relative Death Rates

Fix a lead-time τ, and let G_{ij} be the j-th expert's estimate of the relative death rate in the i-th group, and suppose that n_i experts provide estimates for the i-th group. We wish to construct a confidence interval for the overall relative death rate $\mu = \sum_i F_i \mu_i$, where F_i is the fraction of the population in the i-th group and μ_i is the true relative death rate in the i-th group. The confidence interval will be developed under the following assumptions:

- The G_{ij}'s are independent for all i and j.
- The G_{ij}'s are normally distributed with mean μ_i and common variance σ^2.
- The F_i's are known constants.

Under these assumptions, the sample mean for group i, $\bar{G}_i = \frac{1}{n_i}\sum_{j=1}^{n_i} G_{ij}$ will be normally distributed with mean μ_i and variance σ^2/n_i. Therefore, $\bar{G} \equiv \sum_{i=1}^{m} F_i \bar{G}_i$ will be normally distributed with mean μ and variance $\sigma^2 \sum_i F_i^2/n_i$. Define $\kappa \equiv \sqrt{\sum_i F_i^2/n_i}$. Then the standardized statistic $\frac{\bar{G}-\mu}{\sigma\kappa}$ will be N$(0,1)$.

The following statistic will have a chi-square distribution with $n_i - 1$ degrees of freedom:

$$\frac{n_i - 1}{\sigma^2} S_i^2$$

where S_i^2 is the sample variance of the relative death rate estimate for the i-th group:

$$S_i^2 = \frac{1}{n_i - 1}\sum_{j=1}^{n_i}(G_{ij} - \bar{G}_i)^2$$

By the reproductive property of chi-square random variables, the statistic

$$\sum_{i=1}^{m} \frac{n_i - 1}{\sigma^2} S_i^2$$

is chi-square with $\nu \equiv \sum_1^m (n_i - 1)$ degrees of freedom. Finally, the statistic

$$T = \frac{(\bar{G} - \mu)/(\sigma\kappa)}{\left(\sqrt{\sum_{i=1}^{m} \frac{n_i-1}{\sigma^2} S_i^2}\right)/\sqrt{\nu}}$$

will have a Student's t distribution, with ν degrees of freedom. The expression for T can be simplified to

$$T = \frac{\bar{G} - \mu}{\kappa\sqrt{\sum_{i=1}^{m} \frac{n_i-1}{\nu} S_i^2}}$$

Thus, a two-sided $1 - \alpha$ confidence interval for μ is

$$\left(\bar{G} - t_{\alpha/2,\nu}\text{s.e.}(\mu), \bar{G} + t_{\alpha/2,\nu}\text{s.e.}(\mu)\right)$$

where

$$\text{s.e.}(\mu) \equiv \kappa\sqrt{\sum_{i=1}^{m} \frac{n_i - 1}{\nu} S_i^2}$$

Appendix 5C: Tests for Heterogeneity Among the Experts

Here we examine a specific hypothesis that departs from the independence and unbiasedness assumptions and can be described as postulating that some experts are optimists and others are pessimists. Namely, perhaps some experts consistently make comparatively optimistic estimates of the life-saving potential of earthquake warnings in all sets of circumstances (subpopulation/lead-time combinations) that they are asked about, while other experts are consistently pessimistic.

We will first operationalize the optimism/pessimism hypothesis by introducing quantities that we shall call ***risk ratios***. Then we shall investigate whether the data support this hypothesis. As it turns out, the evidence is inconclusive, but in the spirit of sensitivity analysis we nevertheless compute separate optimistic and pessimistic estimates of how the overall relative death rate depends on lead-time. Even though the existence of a significant optimism effect is in doubt, it is of interest to see how big this effect might be, if it were real.

One way to decompose the absolute death rate with 30 minutes of lead-time is as a product of the baseline death rate and four risk ratios, as follows:

$$\begin{aligned} D_{i,j,30} &= \underbrace{\frac{D_{i,j,30}}{D_{i,j,5}}}_{\theta_{i,j,30}} \underbrace{\frac{D_{i,j,5}}{D_{i,j,1}}}_{\theta_{i,j,5}} \underbrace{\frac{D_{i,j,1}}{D_{i,j,0.5}}}_{\theta_{i,j,1}} \underbrace{\frac{D_{i,j,0.5}}{D_{i,j,0}}}_{\theta_{i,j,0.5}} D_{i,j,0} \\ &= \theta_{i,j,30}\theta_{i,j,5}\theta_{i,j,1}\theta_{i,j,0.5}D_{i,j,0} \end{aligned}$$

Thus, we define the risk ratios as:

$$\theta_{i,j,\tau_k} = \frac{D_{i,j,\tau_k}}{D_{i,j,\tau_{k-1}}}$$

In words, the risk ratio is the fraction of the death rate $D_{i,j,\tau_{k-1}}$ at lead-time τ_{k-1} which remains at lead-time τ_k, where $\{\tau_i\} = \{0, 0.5, 1, 5, 30\}$. A consistently optimistic expert would be one that judges the risk ratios $\theta_{i,j,\tau}$'s to be consistently small, i.e., such an expert would think that a lead time of 30 seconds could substantially reduce the death rate from its baseline level *and* an increase from 30 seconds to 1 minute of lead-time could reduce it still further *and* he would hold a similar opinion on opportunities for risk reduction for other subpopulations.

Consider a particular subpopulation, say "people at home, awake" ($i = 1$), and two lead-times, say one minute and five minutes ($\tau = 1$ and $\tau = 5$). If some experts are consistently optimistic and others are consistently pessimistic in their assessment of how the chances of survival for "people at home, awake", can be enhanced as the lead-time increases from 30 seconds to one minute and from one minute to five minutes, then a positive association is to be expected between the sets of risk ratios

$$\{\theta_{1,1,1}, \theta_{1,2,1}, \theta_{1,3,1}, \ldots \theta_{1,j,1} \ldots\} \text{ (one minute lead-time)}$$

and

$$\{\theta_{1,1,5}, \theta_{1,2,5}, \theta_{1,3,5}, \ldots \theta_{1,j,5}, \ldots\} \text{ (five minute lead-time)}$$

That is, if the optimism/pessimism effect exists, we expect experts whose risk ratios are high for a one minute lead-time to have high risk ratios for a five minute lead-time also (those would be the pessimists) and similarly for experts whose risk ratios are low (the optimists).

For each subpopulation i we have four risk ratios for each expert that was asked about the subpopulation (one for each lead time). The statistical question is whether there is significant positive association between ratios for different lead times, i.e., whether experts whose ratios for one lead time are high (relative to other experts) tend to have high risk ratios also for other lead times. Standard measures of association (for example, see [KS61]) between two paired samples are the Pearson product-moment correlation coefficient and its non-parametric counterpart: Spearman's rank correlation coefficient. Pearson's product-moment correlation coefficient is primarily a measure of linear association, and its distributional theory assumes that the two populations are bivariate normal. There are no prior grounds to believe that either linearity or normality are reasonable assumptions about the risk ratios, and the data do not support those assumptions either. Thus, we concentrate our attention on the non-parametric alternative.

The Spearman rank correlation coefficient is simply the regular (Pearson product-moment) correlation coefficient, computed using the *ranks* of the two samples rather than the actual sample values. Its value may be thought of as measuring the strength of monotonicity in the relation between the two samples, since (in the absence of noise) an arbitrary sample of argument-value pairs from a strictly monotonic function will have a perfect linear relationship between the ranks, causing the rank correlation coefficient to be $+1$ (if the function is increasing) or -1 (if the function is decreasing). The null hypothesis is that the two rankings are independent random permutations of the integers 1 to n (the sample size); other than independence, no assumptions are made about the distribution of the two populations.

A common measure of association between three or more rankings is the average Spearman rank correlation coefficient, with the average taken over all distinct pairs of rankings. Appendix 5D contains mathematical details regarding this statistic and its use in testing the null hypothesis of no association.

Table 5.5 shows the pairwise rank correlations between the four risk ratios and the baseline death rate for the subpopulation "people at home, awake." Every expert was asked to provide estimates for this subpopulation. The pairwise rank correlations range from -0.15 to 0.49. Only the largest pairwise rank correlation of 0.49, between the risk ratios for five minute and 30 minute lead-times, is significant at the 0.05 level (note that the 30 minute lead-time risk ratio measures the proportional reduction in risk *beyond* the five minute lead-time level). It is interesting to note that three of the four pairwise rank correlations with the baseline death rate are negative, suggesting that experts who thought the baseline death rate was relatively high were more optimistic than average about how much the death toll could be reduced if earthquake warnings were available. However, none of these negative correlations are significantly different from zero at the 0.05 level.

The average rank correlation for the risk ratios is 0.12, as shown in table 5.6. Two approximation procedures, one based on the χ^2 distribution and the other on the normal distribution, were used to test the hypothesis of no association (see appendix 5D). The significance levels obtained were 0.12 and 0.18, respectively, suggesting that the hypothesis of no association cannot be rejected. Table 5.6 also shows the average rank correlation for both the risk ratios and the baseline death rate; it is only 0.04 and again the approximate significance levels indicate a non-significant deviation from zero.

The hypothesis of no association between the risk ratios and the baseline death rates was also tested using a different statistic, which we now describe. The results of applying

Pairwise rank correlations for "people at home, awake"	30 seconds ($\theta_{1,j,0.5}$)	1 minute ($\theta_{1,j,1}$)	5 minutes ($\theta_{1,j,5}$)	30 minutes ($\theta_{1,j,30}$)
baseline ($D_{1,j,0}$)	0.13 (0.44)	-0.08 (0.63)	-0.03 (0.87)	-0.12 (0.47)
30 seconds ($\theta_{1,j,0.5}$)		0.22 (0.18)	0.09 (0.57)	-0.15 (0.37)
1 minute ($\theta_{1,j,1}$)			0.01 (0.95)	-0.13 (0.43)
5 minutes ($\theta_{1,j,5}$)				0.49 (0.00)

Table 5.5: Pairwise rank correlations between risk ratios and baseline death rates for "people at home, awake." Significance levels for the null hypothesis of no association are given in parentheses.

Average rank correlation among risk ratios only	0.09 (0.12, 0.18)
Average rank correlation between risk ratios and baseline death rate	0.04 (0.21, 0.39)

Table 5.6: Average rank correlations for "people at home, awake." Two averages were computed: one for risk ratios only and another for both risk ratios and baseline death rates. Approximate significance levels for the null hypothesis of no association are given in parentheses; the two values are based on a χ^2 approximation and a normal approximation. These approximations are discussed in appendix 5D.

this method are consistent with the rank correlation analysis.

Suppose that the ranking of experts for a given lead time is replaced with a sequence of zeros and ones, with experts whose rank is greater than $\lceil n/2 \rceil$ (i.e., greater than the median rank) being assigned a one and experts whose rank is smaller being assigned a zero. Under the hypothesis of no association, a randomly chosen expert's assignment is a one with probability

$$p = \lceil n/2 \rceil / n = \begin{cases} \frac{1}{2} & \text{if } n \text{ is even} \\ \frac{1}{2} + \frac{1}{2n} & \text{if } n \text{ is odd} \end{cases}$$

If this is is done for every lead time (and for the baseline death rate), and each expert's assignments of zeros and ones are added up (call the sum S_j for the j-th expert), then under the null hypothesis of no association S_j would have a binomial distribution with parameters $n = 5$ (the number of rankings) and p as defined above. If some experts were consistently pessimistic and others were consistently optimistic, then one would expect S_j to have a distribution with more weight on extreme values, such as zero (expert's ranking always lower than the median) or five (expert's ranking always above the median). The null hypothesis can then be tested using the χ^2 goodness-of-fit statistic. That test is usually applied under the assumption of independent observations. In this case, there is dependence between the S_j's, caused by the constraint that

$$\sum_{j=1}^{n} S_j = \begin{cases} 2n & \text{if } n \text{ is even} \\ 2n+2 & \text{if } n \text{ is odd} \end{cases}$$

Because of this, we subtract one degree of freedom from the χ^2 distribution used to compute significance levels. The merit of this approach is that it displays the evidence in a very intuitive and graphic manner. Of course, it has the drawback that some of the information in the sample is ignored.

The results of applying this method are shown in figure 5-18 and they are consistent with the rank correlation analysis. The histograms in the figure compare the distribution of the number of times an expert's rank is above the median rank (the outlined bars) with the probability distribution for the number of heads in five tosses of a fair coin (the shaded bars). The picture give little support to the notion of consistently optimistic or pessimistic experts, which one would expect to produce a distribution less peaked than the binomial. A significance level of 0.17 confirms this visual assessment.

We conclude that the pattern of expert estimates for the "people at home, awake" subpopulation provides little evidence that some experts are consistently optimistic or pessimistic. However, data for some of the remaining seven subpopulations do support the notion of an optimism effect, as table 5.7 shows. This table shows the average rank correlations for each subpopulation, along with three values for the significance level of the hypothesis of no association. The significance levels are computed using three different methods that have been mentioned already. We see that for three of the eight subpopulations, the no association hypothesis is rejected at the 0.05 level according to at least one of the three significance level calculations, and for one subpopulation ("people in cars on or near a bridge or overpass") the significance level is 0.02 or lower under all three calculations.

The overall conclusion is that there is some evidence that some experts are consistently optimistic or pessimistic, at least under certain circumstances (when estimating the chances of survival for "people in cars" and "people outdoors") but the evidence is by no means conclusive.

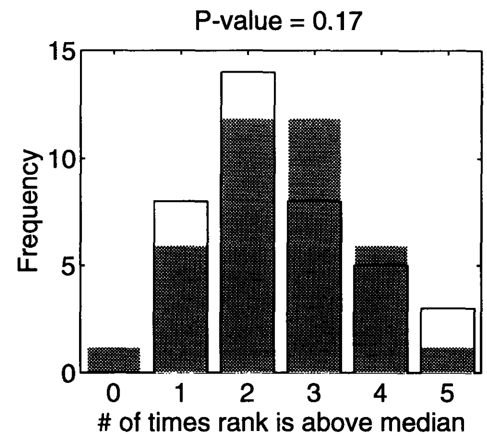

Figure 5-18: The distribution of the number of times an expert's rank is above the median rank (outlined boxes) compared to a binomial reference distribution (shaded boxes). The experts were ranked based on their risk ratios and baseline death rate estimates for "people at home, awake." The observed distribution does not differ significantly from the "no association" reference distribution ($p = 0.17$).

Subpopulation	Avg. rank correlation	Significance level		
		(1)	(2)	(3)
People at home, awake	0.04	0.21	0.39	0.17
People at home, asleep	0.08	0.17	0.31	0.59
People in offices	0.12	0.12	0.18	0.16
Cars on bridge/overpass	0.31	0.02	0.00	0.01
Cars on open freeway	0.26	0.04	0.02	0.50
Cars on urban street	0.02	0.37	0.86	0.59
Children in school	0.13	0.08	0.09	0.15
People outdoors	0.35	0.01	0.09	0.01

Table 5.7: Average rank correlations between risk ratios and baseline death rate estimates for each of the eight subpopulations. Also shown are three different significance levels for the hypothesis of no association. The first and second are based on approximations to the sampling distribution of the average rank correlation (one based on the χ^2 distribution, the other on the normal distribution). The third significance level is computed by comparing the distribution of the number of times an expert's rank falls above the median rank to a binomial reference distribution (see appendices 5C and 5D for details).

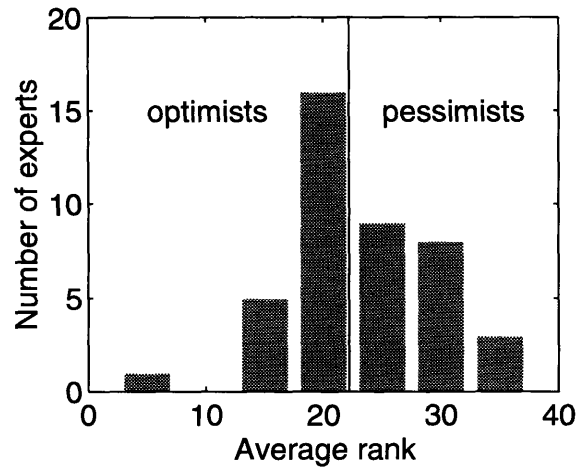

Figure 5-19: Histogram of average ranks. A vertical line at the median average rank divides the optimists (those that gave low estimates of the baseline death rate and risk ratios) from the pessimists (those that gave high baseline death rate estimates and risk ratios).

Despite not having reached a clear-cut conclusion that some experts are consistently optimistic and others pessimistic, let us assume this to be the case, and explore what effects such an assumption would have on the overall relative death rate estimates shown in figure 5-6. First, we have to somehow divide the experts into optimists and pessimists. This was done as follows. The experts were ranked based on their baseline death rate estimates $D_{i,j,0}$ and their risk ratios θ_{i,j,τ_k}. For subpopulations other than "People at Home," where only a subset of the experts provided estimates, the rankings were scaled to make them comparable across subpopulations. Then, an average rank was computed for each expert.

Figure 5-19 shows a histogram of the average ranks. A vertical line shows the median average rank, which was used to separate the optimists from the pessimists[8]. We computed estimates of the overall relative death rate using only the optimists and only the pessimists, respectively. The results are shown in figure 5-9, which is discussed in the main text.

[8]It is worth noting that this histogram gives additional information about the plausibility of the optimism/pessimism hypothesis. The histogram represents the distribution of an average of 15 or so (this was the most common number of situations, i.e., subpopulation/lead-time combinations, that an expert was asked to provide estimates for) ranks for a randomly chosen expert. Under the null hypothesis of no association, each of the 15 ranks should be uniformly distributed over the possible values. Hence the distribution of the average rank should be approximated by a normal distribution, by the central limit theorem. Alternatively, if the experts can be divided into two clearly defined groups of optimists and pessimists, then the distribution of the average rank of a randomly chosen expert should look something like a mixture of two approximately normal distributions: the optimistic distribution and the pessimistic distribution. A visual inspection of the histogram in figure 5-19 suggests that if there are separate distributions for optimists and pessimists, then they are very similar.

Appendix 5D: Rank Correlation Measures

In this section, we state some definitions and results from [FV92] that are used to measure and test for the presence of association between the risk ratios for different experts, for a given subpopulation. Let n be the number of expert judgments available for the subpopulation of interest and let m be the number of lead-times under consideration (m is always equal to 4 for our purposes). Then the data consist of four (in general, m) permutations of the integers $1, 2, \ldots, n$ which we will denote by a_1, a_2, a_3, and a_4, i.e., the rank of expert j for the first lead-time is $a_1(j)$, etc. Suppose μ and ν are two arbitrary permutations of $1, 2, \ldots, n$. Then the following function can be thought of as measuring the distance between μ and ν:

$$d(\mu, \nu) = \frac{1}{2} \sum_{j=1}^{n} (\mu(j) - \nu(j))^2$$

The Spearman rank correlation between μ and ν may be written in terms of this distance d as

$$\alpha(\mu, \nu) = 1 - \frac{2d(\mu, \nu)}{M}$$

where $M = n(n^2 - 1)/6$ is the maximum distance between two permutations.

The null hypothesis of interest is that the four permutations are independent and random. Under this hypothesis, the expected value of the Spearman rank correlation between any two permutations is

$$\mathrm{E}[\alpha] = 1 - \frac{2\frac{n(n^2-1)}{12}}{\frac{n(n^2-1)}{6}} = 1 - 1 = 0$$

This is obtained by noting that the average distance is one half the maximum distance. The test statistic most commonly used is the average pairwise correlation:

$$\begin{aligned} \bar{\alpha} &= \sum_{l>j} \alpha(a_l, a_j) \Big/ \binom{m}{2} \\ &= \frac{1}{6} \left(\alpha(a_1, a_2) + \alpha(a_1, a_3) + \alpha(a_1, a_4) + \alpha(a_2, a_3) + \alpha(a_2, a_4) + \alpha(a_3, a_4) \right) \end{aligned}$$

We now state four results regarding the distribution of $\bar{\alpha}$ under the null hypothesis:

1. The expected value of $\bar{\alpha}$ is of course the same as the expected value of any one of the pairwise correlations, i.e., $\mathrm{E}[\bar{\alpha}] = 0$.

2. The variance of $\bar{\alpha}$ is

$$\mathrm{var}[\bar{\alpha}] = \frac{1}{(n-1)\binom{m}{2}}$$

   For $m = 4$, we get $\mathrm{var}[\bar{\alpha}] = 1/(6(n-1))$.

3. As $m \to \infty$,

$$(n-1)((m-1)\bar{\alpha} + 1)$$

   converges in distribution to χ^2_{n-1}. In our case, m is fixed at 4, so this result is of limited interest.

4. As $n \to \infty$,

$$\frac{\bar{\alpha}}{\sqrt{\mathrm{var}[\bar{\alpha}]}}$$

converges in distribution to $\mathrm{N}(0, 1)$.

Appendix 5E: Parameter Estimates for the Warning Propagation Model

This appendix shows parameters estimates for the distribution of the "time to hear warning" under the constant urgency model (tables 5.8 to 5.10) and the decreasing urgency model (tables 5.11 to 5.13).

Subpopulation	U_1^{LS}	NU_2^{LS}	median$(T_A \mid \Theta^{\mathrm{LS}})$
People at home, awake	1.40	0.00	0.50
People at home, asleep	0.81	0.00	0.86
People in offices	2.29	0.00	0.30
People in cars	0.27	0.00	2.57
Children in school	2.15	0.00	0.32
People outdoors	1.63	0.00	0.43

Table 5.8: Least squares estimates of the parameters of the constant urgency model. The first column shows estimates of the estimated rates (per minute) at which individuals hear an earthquake warning that is communicated exclusively with sirens. The second column shows estimates of the rates at which person-to-person communication occurs (N is the number of individuals). The third column shows the median time (in minutes) until a randomly chosen individual hears the warning, assuming that the constant urgency model is valid and the least-squares estimates are correct.

Subpopulation	U_1^{ML}	NU_2^{ML}	median$(T_A \mid \Theta^{\mathrm{ML}})$	σ^{ML}
People at home, awake	0.50	0.00	1.39	1.37
People at home, asleep	0.33	0.00	2.10	1.27
People in offices	0.75	0.00	0.92	1.40
People in cars	0.22	0.00	3.15	1.15
Children in school	0.72	0.00	0.96	1.46
People outdoors	0.50	0.00	1.39	1.62

Table 5.9: Maximum likelihood estimates of the parameters of the constant urgency model. The first two columns show estimates of the rate at which the warning is propagated from the primary source (sirens) and via person-to-person contact, respectively. The third column shows the median number of minutes until a randomly chosen individual will hear the warning obtained by assuming that the maximum likelihood estimates (and the constant urgency model) are correct. The fourth column shows maximum likelihood estimates of the parameter σ, which is a measure of the dispersion of the expert estimates.

Subpopulation	$E[U_1]$	$NE[U_2]$	median$(T_A \mid E[\Theta])$	$E[\sigma]$
People at home, awake	0.51	0.01	1.35	1.38
People at home, asleep	0.34	0.01	2.02	1.29
People in offices	0.78	0.05	0.87	1.44
People in cars	0.23	0.02	2.94	1.20
Children in school	0.74	0.04	0.92	1.49
People outdoors	0.56	0.05	1.21	1.69

Table 5.10: Bayesian estimates of the parameters of the constant urgency model. The first two columns show posterior expected values for the rate of propagation from the primary source (sirens) and via person-to-person contact, under a uniform prior distribution. The third column shows the median number of minutes until a randomly chosen individual hears the warning, using the Bayesian parameter estimates and assuming the constant urgency model to be correct. The last column shows posterior expected values for the parameter σ, which measure the dispersion between the expert estimates.

Subpopulation	λ^{LS}	γ^{LS}	median$(T_A \mid \Theta^{LS})$
People at home, awake	0.38	0.06	0.31
People at home, asleep	0.42	0.20	0.84
People in offices	0.34	0.01	0.07
People in cars	0.32	0.27	2.09
Children in school	0.43	0.03	0.12
People outdoors	0.17	0.001	0.06

Table 5.11: Least squares estimates of the parameters in the decreasing urgency model. The first two columns show estimates of the parameters λ and γ which appear in the functional form $\lambda/(t+\gamma)$ for the decreasing rate (per minute) of warning propagation. The third column shows the median time until a randomly chosen individual hears the warning, assuming that the least-squares estimates and the decreasing urgency model are correct.

Subpopulation	λ^{ML}	γ^{ML}	median($T_A \| \Theta^{\text{ML}}$)	σ^{ML}
People at home, awake	0.58	0.12	0.28	2.08
People at home, asleep	0.50	0.25	0.75	1.69
People in offices	0.55	0.03	0.08	2.32
People in cars	0.66	1.16	2.16	1.40
Children in school	0.47	0.01	0.03	2.83
People outdoors	0.45	0.03	0.11	3.46

Table 5.12: Maximum likelihood estimates of the parameters in the decreasing urgency model. The first two columns show estimates of the parameters λ and γ that determine how the urgency function $U_1(t) = \lambda/(t+\gamma)$ decreases with time. The third column shows the median number of minutes until a randomly chosen person hears the warning, assuming that the decreasing urgency model and the maximum likelihood estimates are correct. The last column shows maximum likelihood estimates of the parameter σ which measures dispersion among the expert estimates.

Subpopulation	E[λ]	E[γ]	median($T_A | $E[$\Theta$])	E[σ]
People at home, awake	0.61	0.15	0.31	2.07
People at home, asleep	0.53	0.31	0.84	1.69
People in offices	0.63	0.06	0.12	2.22
People in cars	0.76	1.67	2.49	1.41
Children in school	0.54	0.03	0.08	2.66
People outdoors	0.63	0.22	0.44	3.03

Table 5.13: Bayesian estimates of the parameters in the decreasing urgency model. The first two columns show the posterior expected values of the constants λ and γ which determine how rapidly the urgency function $U_1(t) = \lambda/(t+\gamma)$ decreases. The third column shows the median number of minutes until a randomly chosen individual hears the warning, assuming the decreasing urgency model and the Bayesian estimates are valid. The last column shows the posterior mean for σ, which measures dispersion among the expert estimates.

Appendix 5F: Bayesian Parameter Estimates for the "People at Home" Subpopulation

The expert estimates for the fraction of people at home that hear a warning when only sirens are used were used to compute a likelihood function on a three-dimensional grid, corresponding to a discretization of the ranges for the parameters U_1, U_2, and σ in the constant urgency warning propagation model. The expert estimates were assumed independent of each other, i.e., the overall likelihood function was computed as the product of the individual likelihood functions.

If a uniform prior over the three-dimensional grid is assumed, then the likelihood function is also the posterior distribution for (U_1, U_2, σ). The marginal posteriors and two-dimensional joint posteriors under this assumption are shown in figure 5-20. Table 5.14 shows the posterior moments for the parameters. Figure 5-21 shows a region of 95% posterior probability in (U_1, U_2) space. Figure 5-22 shows two curves that encompass all $(t, A(t))$ curves that correspond to values of U_1 and U_2 from inside the 95% region.

Figure 5-23 shows the marginal posteriors under the assumption that the priors for σ, U_1, and U_2 are exponential with mean 1/50. The corresponding posterior moments are shown in table 5.15. We see that the posterior distributions shown in figure 5-23 are much closer to the corresponding posteriors in figure 5-20 than they are to a an exponential distribution with mean 1/50, indicating relative insensitivity of the posterior distribution to the prior distribution.

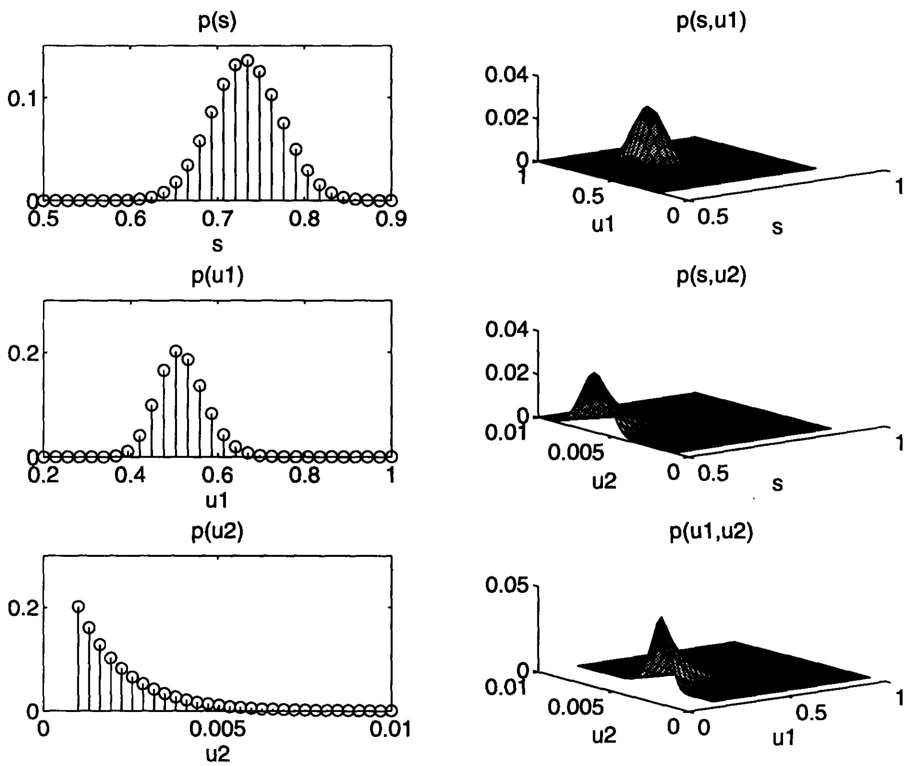

Figure 5-20: Marginal and two-dimensional joint posterior distributions for U_1, U_2, and σ.

Parameter	Mean	Standard Deviation
σ	0.73	0.04
U_1	0.52	0.06
U_2	0.00	0.00

	U_1	U_2
σ	-0.04	-0.00
U_1		0.09

Table 5.14: Posterior means, standard deviations, and correlations for U_1, U_2, and σ.

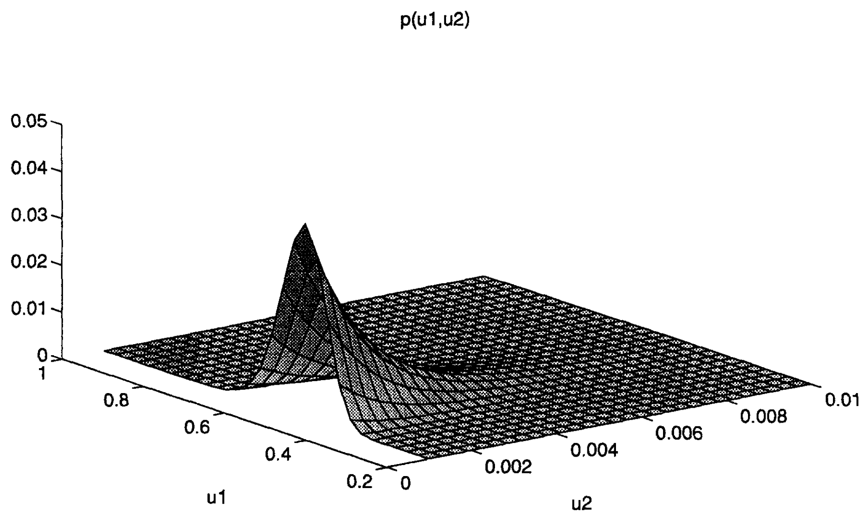

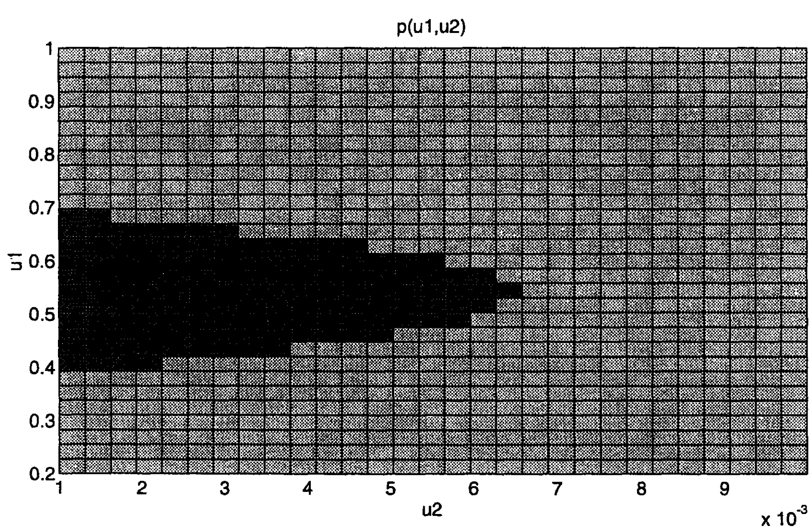

Figure 5-21: Region of 95% posterior probability in (U_1, U_2) space.

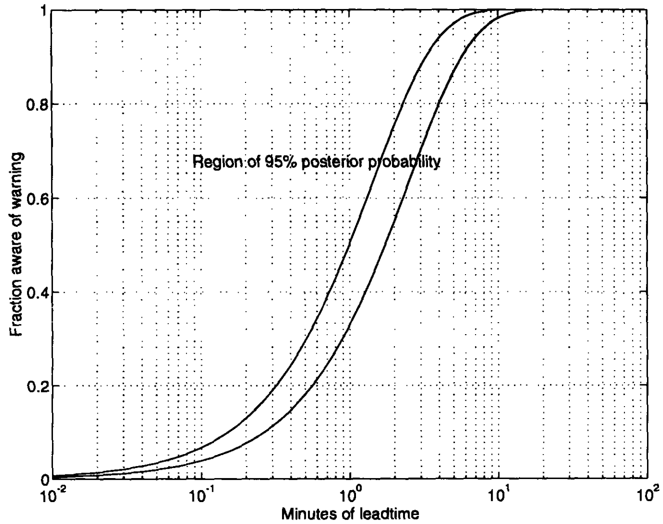

Figure 5-22: Envelope of 95% posterior probability for the "Fraction aware of warning" curve.

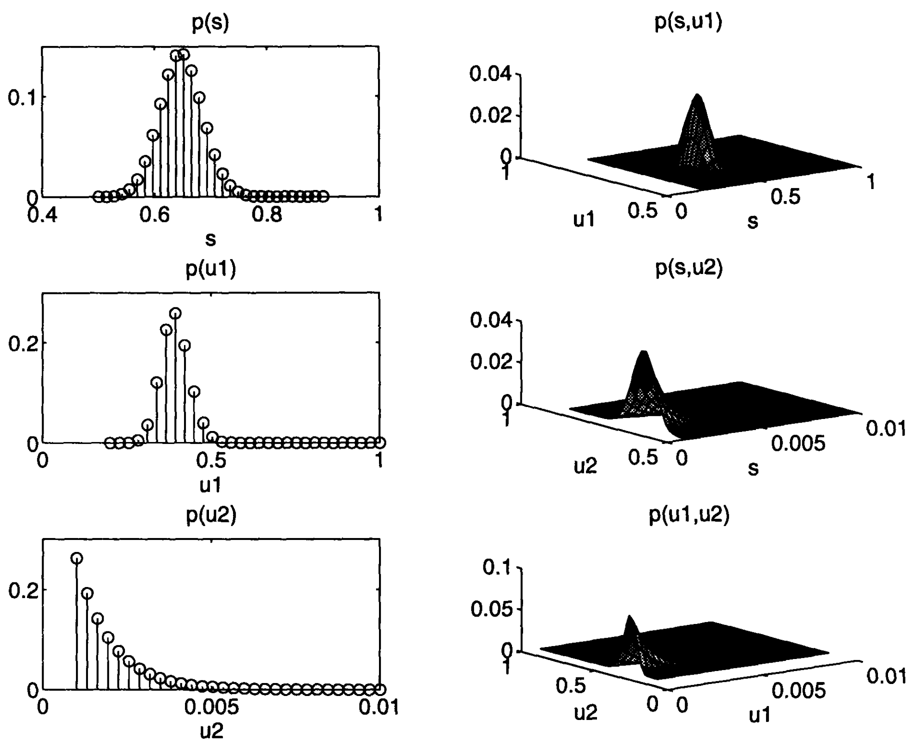

Figure 5-23: Marginal and two-dimensional joint posterior distributions for U_1, U_2, and σ, with exponential prior distributions for all parameters, with mean 1/50.

Parameter	Mean	Standard Deviation
σ	0.65	0.04
U_1	0.39	0.06
U_2	0.00	0.00

	U_1	U_2
σ	0.20	0.02
U_1		0.09

Table 5.15: Posterior means, standard deviations, and correlations for U_1, U_2, and σ, when all parameters have exponential prior distributions with means $1/50$.

Appendix 5G: Tables of Actions Chosen as "Best"

This appendix contains tables 5.16 and 5.17, which show the frequency with which different actions were chosen as the most appropriate one, for different subpopulations.

People at home, awake:

Action	Number of times chosen
Get outside the building, to an open area, free of overhead wires	94
Get under a sturdy desk	54
Get inside a door frame	12
Other	4

People at home, sleeping:

Action	Number of times chosen
Get outside the building, to an open area, free of overhead wires	51
Get under a sturdy desk	22
Get inside door frame	7
Get inside an "earthquake-proof room"	1

People in offices:

Action	Number of times chosen
Get under a sturdy desk	26
Get outside of the building, to an open area, free of overhead wires	22
Get inside an "earthquake-proof room"	11
Get inside a door frame	1

Children in school:

Action	Number of times chosen
Get under a sturdy desk	37
Get outside the building, to an open area, free of overhead wires	32
Get inside an "earthquake-proof room"	9
Other	2

Table 5.16: Frequency with which various actions were chosen as the most appropriate one.

People outdoors:

Action	Number of times chosen
Get to an open area, free of overhead wires	18
Get to the middle of the street	14
Lie down on the sidewalk	2
Enter the nearest building	1
Drive to safety	1

Cars on or near freeway overpasses:

Action	Number of times chosen
Keep driving	24
Keep driving, then stop	8
Stop the car, or don't enter the area of potential danger	4
Other	2

Cars on open freeway:

Action	Number of times chosen
Stop the car and get out of the car	18
Stop the car and stay in the car	16
Drive to safety	2

Cars on urban street:

Action	Number of times chosen
Stop the car and get out of the car	16
Stop the car and stay in the car	14
Drive to safety	6

Table 5.17: Frequency with which various actions were chosen as the most appropriate for people in cars.

Chapter 6

Implications of Recent Earthquakes

In this chapter we discuss three recent earthquakes: the 1994 Northridge earthquake, the 1989 Loma Prieta earthquake, and the 1988 Armenia earthquake. Two of these earthquakes (Loma Prieta and Northridge) occurred in California, under conditions similar to those that the experts were asked to envision when responding to the questionnaires. The Armenia quake graphically illustrates how different the dimensions of earthquake risk are in the third world: The official death toll in that quake was about 25,000, while the official death tolls for the Northridge and Loma Prieta quakes (whose magnitudes were similar to the Armenia quake) were 60 and 63, respectively. We will argue that last-minute earthquake warnings have a larger life-saving potential under conditions as in Armenia, compared to California, both in absolute terms (because death rates in third world earthquakes are orders of magnitude higher) and also in relative terms, at least for some lead-times (because the typical cause of death in third world quakes is building collapse; a cause of death that can be avoided with lead-times long enough for people to evacuate the building they are in).

We will attempt to learn from these earthquakes in two ways. On the one hand we try to see whether the death tolls in these earthquakes are consistent with estimates of baseline death rates based on the expert judgment (presented in chapter 5) and with estimates developed in existing earthquake loss studies (reviewed in chapter 3). On the other hand, we use the analysis performed in the last chapter to estimate, as a function of warning lead-time, how many of the deaths caused by these earthquakes might have been avoided. Information about the circumstances leading to some of the fatalities is used to judge whether the life-saving projections appear plausible.

6.1 The 1994 Northridge, CA Earthquake

The Northridge earthquake occurred at 4:30 AM on January 17, 1994. It had a moment magnitude[1] of 6.7 and a Richter magnitude of 6.4 [Ear94] and was described as follows in the *Los Angeles Times* the following day:

> A deadly magnitude 6.6 earthquake – the strongest in modern Los Angeles history – ripped through the pre-dawn darkness Monday, awakening Southern California with a violent convulsion that flattened freeways, sandwiched buildings,

[1] A magnitude scale that is more directly related to energy release than the Richter scale

ruptured pipelines, and left emergency crews searching desperately for bodies trapped under the rubble.

The official death toll of the Northridge quake, i.e., the number of deaths attributed to the quake by the coroners of Los Angeles, San Bernandino, and Kern counties, was sixty.

Despite the substantial human and economic losses suffered in the Northridge earthquake, it was not comparable in its effects to the likely impact of "the Big One," an earthquake of magnitude 8 or greater on the Southern California portion of the San Andreas fault. We will present one way of extrapolating the pattern of fatalities in the Northridge quake to hypothetical scenarios that are deserving of the label "the Big One" and form the basis for earthquake casualty estimates for the Los Angeles area [AHC+73].

6.1.1 Extrapolation of the Northridge Death Toll

The rationale for the extrapolation method we propose is that even though the overall effects of the Northridge earthquake were not comparable to the Big One, its **local** effects in the region closest to the epicenter were. In order to properly "scale up" the effects of the Northridge quake, we will attempt to relate the death rate with *intensity*, on the Modified Mercalli Intensity (MMI) scale. Figure 6-1 shows an intensity map for the Northridge quake. The maximum intensity was IX in three small regions close to the epicenter. A substantial portion of the San Fernando valley and fragmented outlying areas experienced intensity VIII.

The experts, in responding to the questionnaire, were instructed to assume that the intensity in the region of interest was VIII or greater. It is likely that most of them did not assume the intensity was constant throughout the affected region. Instead they may have assumed that some subset of the region experienced intensity IX or greater, a subset of that subset experienced intensity X or greater, and so on up to the maximum intensity of XII, where the MMI scale ends. Rather than imagining the likely effects in a region experiencing intensity VIII, the experts probably were "taking a weighted average" of regions experiencing various intensity levels, all of them at or above VIII.

Suppose we consider a densely populated region of area A_{tot} that is hit by a shock whose maximum intensity is IX (as in the Northridge earthquake). Let $A_{(\text{VIII})}$ and $A_{(\text{IX})}$ be the areas of the subsets of A_{tot} that experience MMI VIII and IX, respectively. If the population density is roughly even throughout the affected region, then an appropriate set of weights in the weighted average of death rates would be $A_{(\text{VIII})}/A_{\text{tot}}$ and $A_{(\text{IX})}/A_{\text{tot}}$. Generally speaking, one would expect that as the earthquake magnitude increases, the ratio $A_{(\text{IX})}/A_{\text{tot}}$, which can be interpreted as the fraction of the affected area that experiences intensity IX or more, will increase also, and as a result, the appropriate weights would shift towards the higher intensities.

Using 1990 census information, an intensity map for the Northridge quake prepared by the USGS (see figure 6-1), and newspaper accounts of where the fatalities occurred, we obtained the following table:

MMI	population	fatalities	death rate
VII or less	9,916,503	17	0.2 per 100,000
VIII	808,014	17	2 per 100,000
IX	100,896	24	24 per 100,000
Total for MMI VIII and higher:	908,910	40	4 per 100,000

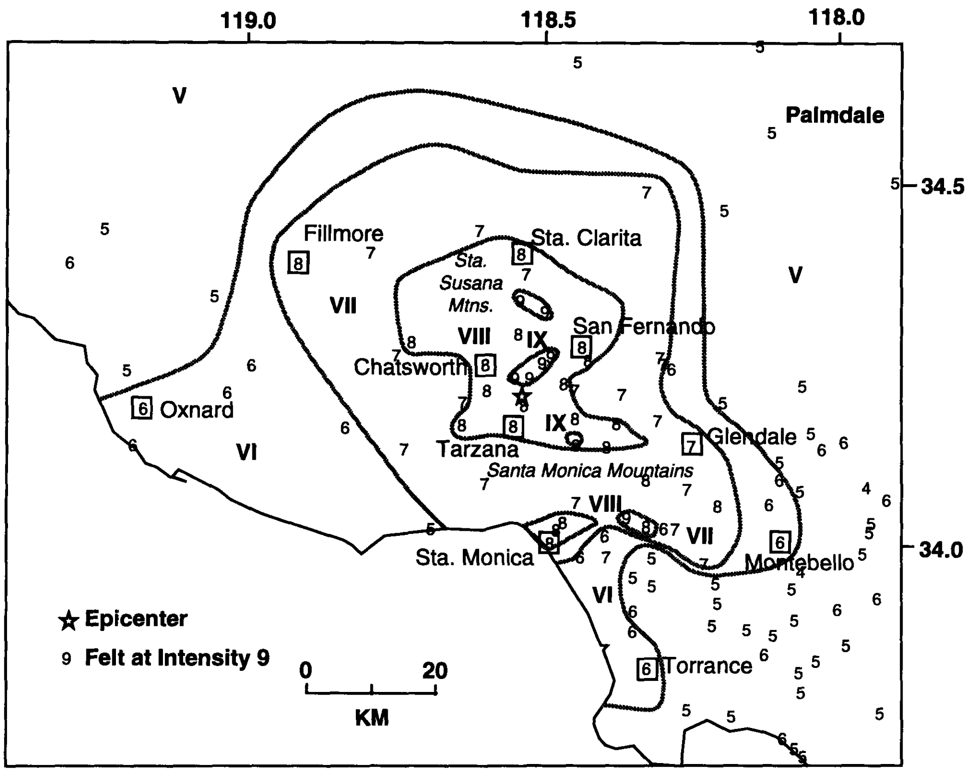

Figure 6-1: Preliminary intensity map for the Northridge earthquake. The roman numerals give average MMI intensities within isoseismals. Most arabic numerals represent intensities in individual communities, individual zip-code areas, or sections of Los Angeles that span many blocks. Some plotted intensity 9 observations correspond to heavily damaged individual structures. Squares denote towns labeled in the figure. Source: USGS [Dew94].

The population for the MMI VII or less region is the 1993 population of Los Angeles, San Bernandino, and Kern Counties, less the population exposed to MMI VIII or IX. Those three counties reported deaths that were officially attributed to the quake. We were unable to determine where two of the deaths occurred. This explains why only 58 of the 60 deaths officially attributed to the quake appear in the table.

As we shall soon see, for purposes of estimating the death rate in a large earthquake centered in an urban area, it makes little difference whether the death rate in the MMI VII or less region is 0.2 per 100,000 or zero, since the death toll will most likely be dominated by deaths that occur in regions experiencing MMI IX or greater.

To extrapolate these death rates to "the Big One," we use two scenarios described in a study of future earthquake losses in the Los Angeles area [AHC$^+$73], which was published in 1973. In this study, which was discussed in chapter 3, estimates of life loss were computed for a hypothetical magnitude 8.3 quake on the San Andreas fault and a magnitude 7.5 quake on the Newport-Inglewood fault. Intensity maps for these hypothetical earthquake scenarios were prepared for the Los Angeles basin [AHC$^+$73, figures 3 and 5]. We measured the areas expected to experience MMI IX, VIII, and VII or less under each scenario, and used the resulting fractions to weigh the death rates computed for the Northridge quake.

Intensity (MMI)	Northridge death rates	San Andreas scenario	Newport-Inglewood scenario
VII or less	0.2 per 100,000	62.7%	30.0%
VIII	2 per 100,000	31.0%	24.5%
IX	24 per 100,000	6.3%	45.5%
Total:		2.3 per 100,000	11.5 per 100,000
USGS/NOAA estimate:		33 per 100,000	48 per 100,000

Table 6.1: Death rates in the Northridge earthquake (which occurred at 4:30 AM) extrapolated to hypothetical earthquakes on the San Andreas and Newport-Inglewood faults. Estimates (for 2:30 AM) from the USGS/NOAA loss study for Los Angeles [SAL+80] are shown for comparison.

The results are shown in table 6.1.

Comparison with USGS/NOAA Loss Studies

The extrapolated death rate for the San Andreas scenario is 2.3 per 100,000. For comparison, a 1980 update [SAL+80] of the USGS/NOAA earthquake loss study for Los Angeles [AHC+73] estimated a death rate of 33 per 100,000 if the hypothetical magnitude 8.3 shock on the San Andreas were to occur at 2:30 AM (recall that the Northridge quake struck at 4:30 AM); a factor of fifteen difference. The difference is not as pronounced for the Newport-Inglewood scenario: our extrapolation results in an estimated 11.5 deaths per 100,000 while the USGS/NOAA estimate is 48 per 100,000; a factor of four difference. For both scenarios, the majority of the estimated number of deaths comes from the MMI IX region; setting the estimated death rate in the MMI VII or less region to zero would have an insignificant effect on the overall death rate extrapolations.

The difference between our extrapolations and previous estimates is striking. Presumably, when "police and city authorities said the number of fatalities could have run into the hundreds – perhaps thousands – if the earthquake had struck during a normal weekday" (*New York Times*, January 18, 1994) they were basing their assessment on the USGS/NOAA loss estimates.

Possible reasons for the pronounced difference between our extrapolations and previous estimates are that (1) the population of Los Angeles was unusually lucky on January 17, 1994, and a larger number of deaths could be expected if an exact clone of the Northridge earthquake were to occur again at the same time of day, (2) previous estimates were too pessimistic, or (3) improved construction practices have decreased earthquake risk in the Los Angeles area since 1973, when the USGS/NOAA loss study was published. We will consider each possibility in turn.

Were Angelenos Lucky? The population of the Los Angeles Area and the San Fernando valley was indeed lucky that the Northridge earthquake occurred in the middle of the night rather than during a busy weekday. But the comparison we are making is not between the risk at different times of day, but rather between the risk at 4:30 AM on January 31st, 1994 and the risk at 2:30 AM on a randomly chosen date if an earthquake were to strike; as is assumed for the USGS/NOAA estimate quoted in table 6.1. There does not appear to be any reason to believe that the population was at lower risk than usual on the night when the Northridge quake struck. One might think that the quality of construction in the San Fernando valley would cause the death risk to be lower than the average risk for the Los Angeles area, perhaps because most of the high-risk unreinforced masonry buildings are outside the San Fernando valley. But the authors of the USGS/NOAA study do not appear to have believed this: their estimated death rates for the San Fernando valley are slightly higher than the average death rate estimate for the Los Angeles area.

Still, the residents of the San Fernando valley might have been lucky in the sense that a person who wins in the lottery is lucky, even if the prior odds of survival were not unusually favorable. Suppose we assume that the 24 deaths in the intensity IX region, which dominate our extrapolation, is a sample from a Poisson distribution. Then an approximate 95% confidence interval for the expected number of deaths in this region in an exact clone of the Northridge quake[2] is $24 \pm 2 \times \sqrt{24} = (14.2, 33.8)$. The factor of two difference between the upper and lower confidence limits is smaller than the discrepancy between our extrapolations and the USGS/NOAA estimates, suggesting that fluctuations from the mean can not be blamed for the difference. However, the Poisson assumption probably underestimates the variance of the death toll; after all sixteen people died in the collapse of one apartment complex and their prior chances of dying can hardly be assumed independent as required for a Poisson process of fatalities. In an appendix to this chapter we develop an approximate confidence interval of $(2, 46)$ for the number of deaths in the intensity IX region, under an assumption that the number of *structures* that fail follows a Poisson process and the number of *fatalities per structure* is a geometric random variable[3] with mean equal to 3. This substantially wider confidence interval makes it more difficult to rule out the possibility that the USGS/NOAA estimates are realistic and the lower-than-expected death toll in the Northridge quake was simply a random fluctuation.

Are the USGS/NOAA Estimates too Pessimistic? In chapter 3 we mentioned what appeared to be double-counting in the estimation of fatalities that occur on freeways and on urban streets, outside buildings. But this criticism did not apply to the night-time estimates in [AHC+73], since the study appears to assume that everyone is at home at 2:30 AM.

Still, even if the methods used were internally consistent, that does not mean that they were not overly pessimistic and our extrapolation of the Northridge death toll suggests that this may be the case.

Has Earthquake Risk Decreased Since 1973? There are at least two reasons for believing that the risk has decreased and continues to decrease. First, engineers have steadily advanced their knowledge of how to design structures to withstand earthquakes and this knowledge has gradually been incorporated in building codes. As urban areas sprawl and older buildings are replaced, the fraction of structures built according to the latest building codes increases while the number of buildings built

[2]Meaning that the physical effects and timing of the clone are the same as for the original quake. Where people happen to be and what they happen to be doing would be different.

[3]The assumption that the number of fatalities per structural failure has a geometric distribution was made for tractability. We simply want to indicate roughly by how much the uncertainty associated with the observed death rate is underestimated under the assumption that the total number of fatalities follows a Poisson distribution.

Subpopulation	Estimated fraction of population at 3 AM	Estimated death rate
People at home, awake	5%	21 per 100,000
People at home, asleep	92%	23 per 100,000
People in offices	1%	102 per 100,000
Cars near bridge/overpass	0.3%	203 per 100,000
Cars on open freeway	0.6%	3 per 100,000
Cars on urban street	0.6%	7 per 100,000
Children in school	0%	13 per 100,000
People outdoors	0.5%	7 per 100,000
Estimated overall death rate:		24 per 100,000

Table 6.2: Estimation of the overall death rate that can be expected in a major earthquake that occurs at 3 AM, using the expert estimates of baseline death rates from chapter 5 and estimates of how the population is distributed at 3 AM from chapter 3.

before building codes were instituted continues to decrease. The second reason might be described as "survival of the fittest buildings." The fatality rate estimates in the USGS/NOAA loss studies for the San Francisco and Los Angeles areas [AHC$^+$73, ARD$^+$72] were based on extrapolations of death rates in past California earthquakes, notably the 1971 San Fernando earthquake which affected an area that largely coincided with the area that shook in the Northridge temblor. But the structures that collapsed in the 1971 San Fernando quake obviously did not have the opportunity to collapse again in the Northridge quake; and if the structures were replaced, then it is likely that the replacements were better able to withstand groundshaking.

Comparison with Expert Judgment We can also compare extrapolations based on the Northridge death rates with an estimate obtained by combining the expert responses to our survey that were presented in chapter 5 and our estimates of how the population would be distributed at 3 AM (see chapter 3). This computation is shown in table 6.2. The overall estimate of 24 deaths per 100,000 lies between the Northridge extrapolations and the NOAA/USGS estimates. Many of the experts were probably familiar with the NOAA/USGS loss studies and perhaps they were of the opinion that these studies overestimated the death risk. Note that the overall estimate of 24 per 100,000 coincides with the death rate in the MMI IX region in the Northridge quake. If most of the experts were envisioning a scenario where intensity IX or higher prevails throughout an urban area, and if one believes that the damage near the epicenter of the Northridge quake was representative of what can be expected under such a scenario, then the experts were right on target[4].

We can also use the expert estimates, combined with our population distribution estimates, to project how many people might have died in a "clone" of the Northridge quake, had it occurred at a different time of day (January 17, the date of the Northridge quake, was a holiday, but our projections are for a weekday). Figure 6-2 shows the result of this projection. Our death rate estimate from chapter 5 was scaled up to coincide, at 3 AM, with the official Northridge death toll of sixty. Death rate estimates for other times of day

[4]Almost all questionnaires were completed and returned before the Northridge quake occurred

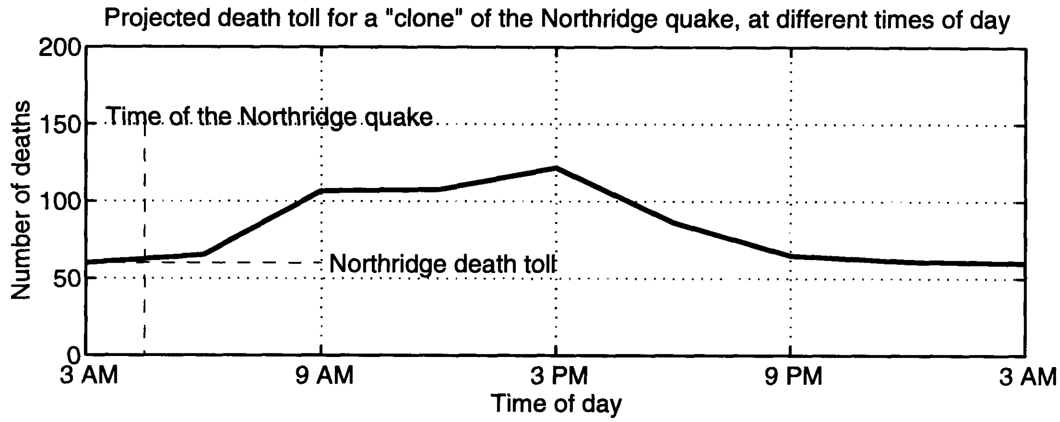

Figure 6-2: A 24 hour projection of the number of fatalities that could be expected if the Northridge quake had occurred at a different time of day. The projection is based on expert judgment (see chapter 5) and estimates of the fraction of the population that falls in each of several subpopulations (see chapter 3) at different times of day.

were then scaled up by the same amount, and the 24 hour profile of figure 6-2 is the result.

The projected death toll reaches its maximum at 3 PM, at 120 deaths. This is in stark contrast with speculations of a "casualties ... multiplied a hundredfold" and "lives lost ... in the thousands" [Dav94] that we mentioned in chapter 5. It is possible that these speculations were based not only on available estimates, for example from the USGS/NOAA study for Los Angeles, but also on the collapse of a department store and the partial collapse of parking garages that were caused by the quake. Had the tremor occurred on a busy weekday, it is indeed likely that these collapses would have increased the death toll, but it is difficult to believe that the increase would have been measured in the thousands. For example, suppose that a parking garage has space for 500 cars, and that it is filled to capacity. At any given time, only a fraction of the *occupants* of these cars will be inside the garage. Let's assume that on average, a person leaves his car in the garage for a time period five times longer than the time he spends in the garage, so that about 150 people (assuming 1.5 occupants per car) might be in the garage when it is filled to capacity. Even in a total collapse, it is likely that some of these 150 people would survive. Thus, to bring the death toll into the thousands, one would have to believe that about six structural failures comparable to our hypothetical garage collapse would occur.

6.1.2 Projected Death Toll Reduction as a Function of Warning Lead-Time

We now use the estimated relative death rates from chapter 5 to project how many lives might have been saved if earthquake warnings had been available before the Northridge earthquake occurred. The projections are based on our estimates of how the population would be distributed at 3 AM (see table 3.3), expert responses, and the optimistic assumptions we have been making throughout: That the warnings are unequivocal regarding the time and location of the earthquake, that a system of sirens is used to communicate the warning to the public, and that people who hear the warning immediately attempt an

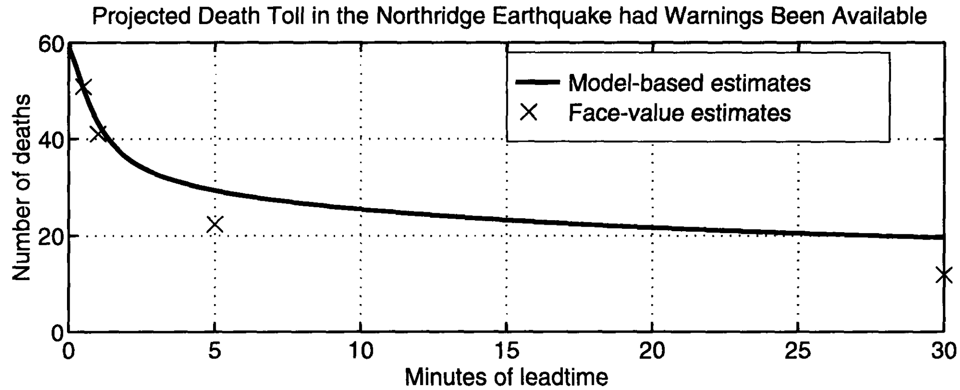

Figure 6-3: Projections of how many of the sixty deaths officially attributed to the Northridge earthquake could have been avoided if earthquake warnings had been available before the shaking started. The projections are based on the model-based and face-value analyses reported in chapter 5 and the "best-case" assumptions described there and in chapter 4.

appropriate action to increase their chances of survival. The actual time of occurrence of the Northridge quake, 4:30 AM, lies between two times (3 AM and 6 AM) for which we estimated the population distribution; we used the 3 AM estimate because in our judgment it is more similar to what can be expected at 4:30 AM than the 6 AM estimate.

The resulting projection is shown in figure 6-3, where the projected number of fatalities is plotted versus the length of the warning lead-time in minutes. The projections are simply the relative death rate estimates developed in chapter 5 multiplied with the official Northridge death toll of 60. The figure is not identical to figure 5-11 because that figure is based on a population distribution that is an average over the hours of the day, whereas figure 6-3 assumes that the time of day is 3 AM. We see that the model-based projections are not as optimistic as the face-value projections are: A five minute warning suffices to save 30 of the 60 lives that were lost according to the model-based projection, while the face-value approach predicts that about 35 lives could be saved. With 30 minutes of warning, 40 lives could be saved according to the model-based projection and 48 according to the face-value estimate.

To judge the reasonableness of these projections, let us consider the circumstances that led to some of the fatalities. According to a reconnaissance report published by the Earthquake Engineering Research Institute (EERI) [Ear94],

> Structural failure was the underlying cause of most of the fatalities directly attributable to the earthquake. Sixteen persons died in the collapse of a three-story, wood frame Northridge apartment building located close to the epicenter. The building was essentially a soft-story configuration with the ground floor housing garden apartments and carports. All those killed lived in first-floor apartments in those building areas that completely collapsed. About 180 residents evacuated this building despite the separation and collapse of exit stairs. Four other deaths occurred in three separate single-family homes that collapsed and slid down hillsides.

The authors of this report thought most of the heart attack deaths included in the Coroner's official death toll had a "tenuous" relationship with the earthquake and could not be directly attributed to it; we will discuss this matter later in this section.

Each of the twenty deaths caused by structural failure could probably have been avoided if the victims had been warned far enough ahead of the quake that they had time to leave the building. Since the victims lived either on the first floor of an apartment building or in single-family homes, a couple of minutes would suffice for most people for the physical action of getting outside the building. Adding three minutes for hearing the warning and rousing oneself from sleep at 4:30 AM, one obtains an estimate of about five minutes of warning as the lead-time required to save the twenty victims of structural collapse.

Shorter lead-times would have been of little use to the four victims that were in homes that slid down hillsides. But according to the *Los Angeles Times* (January 18, 1994), 10 people were rescued alive from the ruins of the Northridge Meadows apartment complex and 160 evacuated the building successfully on their own. The fact that 10 of the 26 people trapped in the building were rescued alive suggests that some locations inside the building were safer than others, probably because cavities were formed in the collapsed first story. Cavities might form next to load-bearing walls or sturdy furniture and it appears reasonable to expect that a couple of minutes would have sufficed for most of the residents to move to a location (perhaps under a sturdy desk) where their chances of survival would have been substantially greater. How much greater is difficult to know, but a reasonable guess might be that 8 of the 16 deaths could have been avoided with two minutes of warning.

The reconnaissance report continues:

> The 13 additional deaths not attributable to structural failure included a couple who died when buried under hundreds of pounds of books, model trains, and other collectables in their home; a woman who apparently tripped and hit her head on a crib while checking on her baby; a downtown Los Angeles hotel resident who fell from a sixth-story window; a Los Angeles police officer who died when his motorcycle plunged off a collapsed freeway overpass; an elderly woman who died in a mobile home fire; a 25-year-old man who was electrocuted; and two men who perished in a helicopter crash three days after the earthquake while checking for earthquake damage to pipelines.

The couple that was buried under their possessions could certainly have avoided that fate by leaving their home; an action that could probably have been accomplished in less than five minutes. The same can be said for the woman who died in her mobile home. Had the woman that was checking on her baby started running two minutes *before* the quake she might have been less likely to trip. According to the *Los Angeles Times* (January 23, 1994), the man who was electrocuted was attempting to save a child trapped in a car with a live powerline dangling over it. It is conceivable that had the mother of the child (who was standing nearby when the victim died) known that the quake was coming, she would have left the car in a less vulnerable spot. However, it must also be admitted that foreknowledge of the quake would not have saved the police officer who plunged off the freeway *after* the shaking had stopped, or the two men in the helicopter. There were a total of 21 deaths that were attributed neither to structural failure nor to heart attacks; it might be reasonable to suppose that half of those deaths could have been avoided with a five minute lead-time and perhaps a quarter with a two minute lead-time. Thus, if we consider deaths from causes other than heart attacks, it does not seem unreasonable to expect that about 13 lives could have been saved with two minutes of warning and that this number would have risen to

30 with five minutes of warning. This accounts for most of the projected lives saved for lead-times of five minutes or less (see figure 6-3).

As the lead-time increases from five to thirty minutes, an additional 10 to 13 lives could be saved according to our projections. Perhaps the number of fatal heart attacks induced by the earthquake could have been reduced by this number if individuals had enough time to calmly prepare for the event and could avoid sudden exertion during or immediately after the quake. Of course, others might argue that the knowledge that an earthquake is about to occur could cause extreme anxiety.

On balance, we believe that the life-saving projections of figure 6-3 are not in clear conflict with the conditions under which the deaths in the Northridge earthquake occurred, as long as one accepts the best-case assumptions that we have been making throughout (unequivocal warnings and immediate reaction to warnings).

We conclude our discussion of the Northridge earthquake with a few remarks about the criteria that are used to decide whether a fatality should be attributed to the quake.

6.1.3 Was the Official Death Toll Realistic?

The EERI's reconnaissance report [Ear94] offers the following assessment of the official death toll:

> Although the earthquake's official death toll stands at 60, the number of fatalities actually caused by this event is probably closer to 33. Early estimates included 19 deaths from heart attacks; however, with the exception of a 62-year-old male who apparently died of a heart attack after extricating himself from his collapsed apartment, the relationship of these fatalities to the earthquake is tenuous. The circumstances surrounding eight of the reported fatalities are still under investigation.

The USGS/NOAA loss study for San Francisco [ARD$^+$72] remarks that available records on past earthquakes are less reliable "with respect to deaths and injury information as ... for building and other property damage. Heart attack deaths may or may not be included, and the text leaves the matter unclear in most cases. Injuries leading to deaths may be included under injuries or under deaths." This problem is not limited to distant countries or the distant past: According to [WJS94], "even the scientific literature could not agree on a total count [of deaths in the 1989 Loma Prieta earthquake], offering a range of 60 to 67 deaths."

In the judgment of the authors of the reconnaissance report, most of the fatal heart attacks attributed to the quake by the coroner were not clearly related to the earthquake. Since we do not have detailed information about the circumstances leading to each of the fatal heart attacks, we do not pretend to bring this matter to a definite resolution. Nevertheless, we believe that the data we are about to present are illuminating. Figure 6-4 shows the number of fatal heart attacks in Los Angeles County for one week before the January 17 quake and two weeks after the quake.

These data indicate that the occurrence of the earthquake was accompanied by a marked jump in the number of fatal heart attacks. The mean number of cases per day over the whole period is 11.3. If one assumes that the number of cases per day are independent Poisson random variables with mean 11, then the chances of seeing one or more days with 37 (the number of fatal heart attacks reported on the day of the quake) or more cases are about 4×10^{-9}. But we also note that the number of cases was at or below the postulated mean of

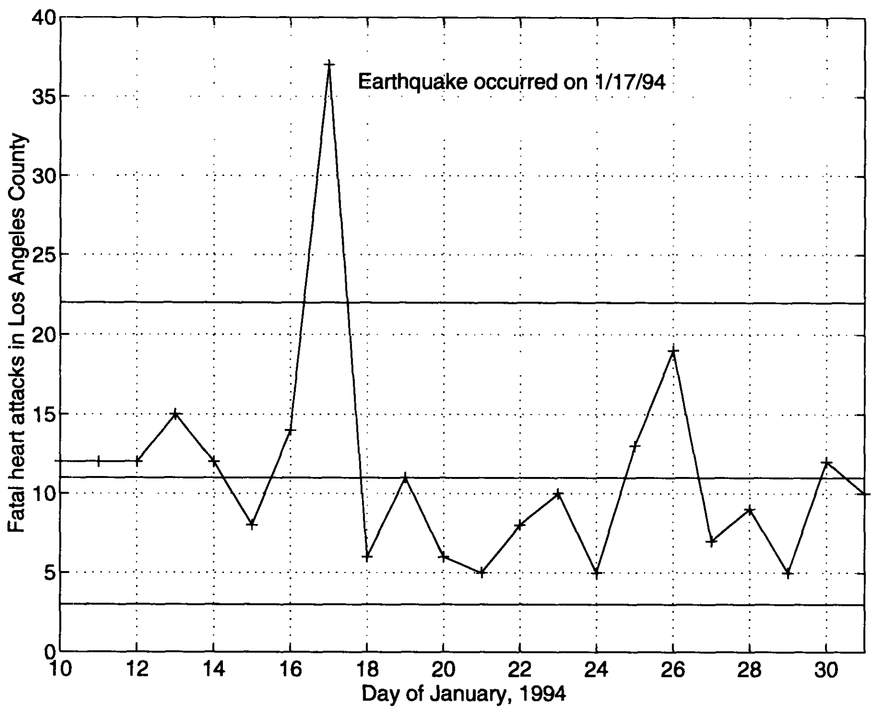

Figure 6-4: Fatal heart attacks in Los Angeles county for the period from 10 to 31 January, 1994. The mean number number of fatal heart attacks per day over this period is 11.3. If the true mean were 11 and if the number of cases per day were a Poisson random variable, then the actual number of cases on any given day would fall between the lines at 22 and 3 with 99.8% probability on any given day.

11 for seven consecutive days following the earthquake. Furthermore, the number of fatal heart attacks during the week immediately preceding the quake was 85, the corresponding number for the week starting with the day of the earthquake was 83, and for the third week the number of cases was 70. Applying the Poisson assumption to the weekly counts, the expected number of fatal attacks per week would be 77, and the observed counts of 85, 83, and 70 fall at the 83-rd, 77-th, and 23-rd percentiles of a Poisson distribution with that mean. We conclude that the data are consistent with a hypothesis that the quake did not cause a significant number of heart attacks that would not otherwise have occurred within about a week. Of course, there may be other more plausible explanations of the effects of earthquakes on the risk of heart attack, but any such explanation would have to explain the observed pattern in the data.

The main conclusion is that there was a substantial increase in fatal heart attacks on the day of the earthquake. While the data are not hostile to a hypothesis that the people who died from heart attacks were "due," i.e., they would have died anyway within a week or so it is not clear what relevance this hypothesis – even if true – has for the problem of counting the number of deaths attributable to the earthquake. Everyone is "due" to die eventually, and a conjecture that an earthquake shortened a life by a only a week as opposed to, say, ten years does not seem to be an acceptable basis for deciding that the death should not be attributed to the earthquake.

Further questions about the validity of the official count were raised six months after the quake. The *Los Angeles Times* (July 4, 1994) reported that the state of California, through a program funded in part by FEMA, had reimbursed the families of 117 people, whose deaths were "directly attributable to the quake or hastened by the disaster" for funeral expenses. An additional 138 requests were still pending. The 117 deaths included about half the fatalities counted by the Coroner (grants were not given if funeral expenses were covered by insurance) and more than half were caused by heart attacks. According to the article,

> One explanation for the difference between the official toll and the number of funeral grants given is that the coroner's office does not investigate every death. The coroner is required to investigate deaths of people who had not seen a doctor for 60 days or died as a result of trauma, such as being trapped in a collapsed building.
>
> But deaths from natural causes, such as heart attacks, are not routinely subject to coroner's review, and the determination of whether the death was quake-related is left to the attending physician.

The specific cases mentioned in the article (e.g., a man "who hanged himself over the loss of his job" and a person "who died of exposure after refusing out of fear to go inside after the quake") underline the difficulty of deciding, in particular cases, whether a death can be reasonably attributed to the earthquake. These difficulties suggest that the use of statistical information of the kind presented in figure 6-4 could help in predicting the death toll of future earthquakes. Of course, this does not solve the dilemma of deciding who deserves reimbursement for funeral expenses!

The last major earthquake to strike a densely populated region in California prior to the Northridge earthquake was the 1989 Loma Prieta earthquake, which is the concern of the next section.

6.2 The 1989 Loma Prieta, CA Earthquake

According to *Seismicity of the United States* [SC93], the Loma Prieta earthquake, which occurred on October 18, 1989, at 5:04 PM, "caused 63 deaths, 3,757 injuries, and an estimated $6 billion in property damage [and] was the largest earthquake to occur on the San Andreas fault since the great San Francisco earthquake in April 1906." As we alluded to in the last section, sources do not agree on the number of deaths caused by the earthquake; according to a preliminary EERI reconnaissance report [Ear89] the total number of fatalities was 67 and injuries numbered 2,435. In the final version of the EERI report [Ear90], the death toll is revised to 62 and the number of injuries to 3,757. The same report indicates that the moment magnitude of the shock was 6.9 while the Richter magnitude was 7.0; this implies that about twice as much energy was released as in the Northridge earthquake[5].

The organization of this section will largely parallel the section on the Northridge quake. We start by estimating death tolls in different intensity regions and then use the estimates to extrapolate to scenarios postulated in the USGS/NOAA loss study [ARD$^+$72] for the San Francisco Bay Area. As it turns out, this extrapolation is more problematic than was the case for the Northridge quake, for reasons that we will discuss. The extrapolation results will be compared to the USGS/NOAA estimates and to an estimate based on expert judgment. Since the Loma Prieta quake occurred in the late afternoon, while the Northridge quake occurred in the early morning, the two events provide a useful comparison of the death tolls that can be expected at different times of day. The section ends with a projection of how many of the Loma Prieta fatalities could have been avoided had earthquake warnings before the tremor been possible.

6.2.1 The Pattern of Fatalities and Its Extrapolation

Figure 6-5 shows the intensity distribution of the Loma Prieta quake. The Santa Cruz region was assigned intensity VIII on the MMI scale while most of the Bay Area was assigned intensities VI and VII. The map of figure 6-5 is from [SC93], which states that "intensity IX was assigned to San Francisco's Marina District, where several houses collapsed, and to four areas in Oakland and San Francisco, where reinforced-concrete viaducts collapsed: Nimitz Freeway (Interstate 880) in Oakland, and Embarcadero Freeway, Highway 101, and Interstate 280 in San Francisco." However, these isolated incidents of intensity IX were not deemed extensive enough to warrant an intensity IX region on the isoseismal map. Similarly, an early intensity map for the Northridge earthquake, prepared by the EQE International company, showed a region of intensity X near the epicenter, while the maximum intensity shown on the USGS isoseismal map of figure 6-1 is MMI IX.

The *Journal of the American Medical Association* [CP$^+$89] reported on the deaths attributed to the Loma Prieta earthquake by the county coroners of the various counties in the Bay Area. Table 6.3 shows how the fatalities were distributed across counties.

The intensity VIII region is entirely contained within Santa Cruz county and except for the sparsely populated northern third of the county, the intensity VIII region covers the whole county. Therefore, we use the Santa Cruz death rate as a proxy for the death rate in the intensity VIII region. On computing the average death rate for the remaining counties,

[5]The Northridge earthquake had a moment magnitude of 6.7. A one unit increase on the logarithmic moment magnitude scale corresponds roughly [Bol93] to a 30-fold increase in energy release. Thus, an increase from 6.7 to 6.9 in magnitude corresponds to an increase of $30^{6.9-6.7} \approx 2$ in energy release.

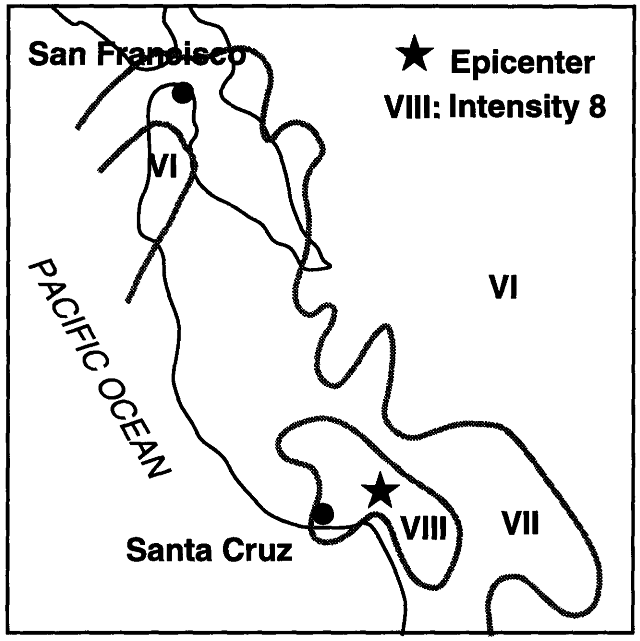

Figure 6-5: Isoseismal map for the Loma Prieta, California earthquake of October 18, 1989. Source: [SC93].

County of death	No. deaths	Population	Death rate	Intensity (MMI)
Alameda	42	1,252,400	3.4 per 100,000	VI, VII
Santa Cruz	5	229,900	2.2 per 100,000	VIII, VII
San Francisco	13	731,700	1.8 per 100,000	VII
Monterey	2	349,300	0.6 per 100,000	VI, V, VII
Santa Clara	1	1,440,900	0.1 per 100,000	VII, VI
San Benito	0	35,500	0	VI, V
San Mateo	0	632,800	0	VII, VI
Total:	63	4,672,300	1.3 per 100,000	

Table 6.3: Number of deaths, population, and death rate, by county for the Loma Prieta earthquake. Population estimates are for January 1, 1989. Intensity values are ordered in decreasing order of area covered. For example, intensity VI prevailed throughout more than half the area of Alameda County, with the remainder experiencing intensity VII. Adapted from [CP+89].

Intensity (MMI)	Loma Prieta death rates	San Andreas scenario: fractions of built-up areas
VII or less	1.3 per 100,000	19.1%
VIII	2.2 per 100,000	46.6%
IX	2×24^a per 100,000	28.2%
X	x	12.2%
	Total:	$14.08 + 0.12x$ per 100,000
USGS/NOAA estimate:		224 per 100,000

[a] from Northridge quake

Table 6.4: An attempt to extrapolate death rates in the Loma Prieta earthquake to a hypothetical magnitude 8.3 earthquake on the San Andreas fault near San Francisco. Since intensities IX and X were not experienced (except for isolated locations) in the Loma Prieta quake, we have used a death rate from the Northridge earthquake for intensity IX (multiplied by two to account for the different time of occurrence of the Northridge quake) and the symbol x for the unknown intensity X death rate. An estimate (for 4:30 PM) from the USGS/NOAA loss study for the San Francisco Bay Area [SAL+80] is shown for comparison.

we obtain the following counterpart to the death rate table presented for the Northridge quake:

MMI	population	fatalities	death rate
VII or less	4,442,400	58	1.3 per 100,000
VIII	229,900	5	2.2 per 100,000

The USGS/NOAA loss study for the San Francisco Bay Area [ARD+72] presents death toll estimates for a hypothetical magnitude 8.3 earthquake on the San Andreas fault (and for other scenarios). We measured the fractions of built-up areas in the Bay Area projected to experience various intensity levels as a result of this hypothetical quake. Table 6.4 shows an attempt to extrapolate from the Loma Prieta death rates to this scenario. The problem we encounter is that while large regions on the San Francisco peninsula are expected to suffer intensity IX and X under the magnitude 8.3 scenario, Bay Area residents were fortunate enough to experience no more than intensity VIII during the Loma Prieta earthquake (with the isolated exceptions already noted). Thus, we have no intensity IX or X death rates to extrapolate. However, we have an intensity IX death rate for the Northridge earthquake, and for a first approximation we have used double that rate in table 6.1. The factor of two is meant to account for the different time of occurrence (4:30 AM) of the Northridge earthquake; it lies between the observed differences of 10% (for MMI VIII) and 600% (for MMI VII or less) between the Northridge and Loma Prieta death rates, and coincides with our estimate from chapter 5 (see figure 5-3). For intensity X, neither the Northridge nor the Loma Prieta earthquake provide a basis for estimating a death rate, so we have used the symbol x for that unknown death rate.

Comparison with the USGS/NOAA Estimate

Even though we don't know what x is, we can draw limited inferences about whether the Loma Prieta death toll is consistent with the USGS/NOAA estimate that an overall death rate of 224 per 100,000 can be expected if a magnitude 8.3 earthquake were to occur at 4:30 PM on the San Andreas fault close to San Francisco. For our extrapolated death rate to equal 224 per 100,000, the unknown intensity X death rate would have to equal 1,740 per 100,000. For comparison, two estimates of the death rate in the great 1906 San Francisco quake are 124 per 100,000 (quoted in table 3.4) and 667 per 100,000 (quoted in [CP$^+$89]). San Franciscans are probably better prepared for earthquakes now than they were before the 1906 disaster and the houses they live in are likely to be better able to withstand earthquakes than pre-1906 dwellings were[6]. So our extrapolation suggests that the USGS/NOAA loss study [ARD$^+$72] overestimates the number of deaths that can be expected if a great earthquake were to occur in the Bay Area in the future, but this conclusion is tentative because conditions near the epicenter of the Loma Prieta earthquake were not comparable to the conditions that are expected to prevail in "the Big One."

Our extrapolation is suspect for several reasons, especially when compared to our extrapolation of the Northridge death toll. We now discuss reasons why the extrapolation should be interpreted with caution.

Based on the time of occurrence of the Loma Prieta earthquake (5:04 PM) one would expect the death rates to be higher than for the 4:30 AM Northridge quake. Indeed, the intensity VIII death rate for Loma Prieta is 10% higher than the Northridge rate (2.2 versus 2 per 100,000) and the intensity VII or less death rate for Loma Prieta is six times larger than the Northridge rate (1.3 versus 0.2 per 100,000).

The intensity VII or less death rate for the Loma Prieta quake is dominated by 41 deaths that occurred in the collapse of double-decked freeway structure in Oakland (the Cypress Street viaduct), constructed in the 1950's. The only structures of similar design in California are in San Francisco; these structures were damaged by the quake but did not collapse [Bol93]. Some of these structures have been demolished (the Cypress Street and Embarcadero viaducts) while others have been strengthened (Interstate 280), so it appears that the chances of a similar occurrence in the future have been lowered. If the Cypress Street viaduct deaths are excluded, the intensity VII or less death rate for the Loma Prieta quake becomes 0.4 per 100,000, or double the corresponding rate for the Northridge quake.

But there are other reasons to expect that *larger* time-of-day differences than those observed between the Loma Prieta and Northridge death rates can be expected in future quakes:

Heart attack deaths were not included: None of the deaths reported by the county coroners as related to the Loma Prieta quake were caused by heart attacks, whereas one third of the Northridge fatalities were from heart attacks. Presumably this is because coroners in the Bay Area used more stringent criteria for classifying a death as earthquake-caused. As noted in [CP$^+$89], "There is no universally accepted definition of an "earthquake-related death"; for this report, the determination was made by each county [Medical Examiner or Coroner]."

Short duration: According to [Ear90], the duration of the shaking for the Loma Prieta quake was unusually brief for an earthquake of this magnitude.

In addition, the authors of the EERI reconnaissance report [Ear90] are of the opinion that "the casualties would have been much greater had many people not gone home early to

[6]Older buildings in San Francisco and Oakland are perhaps no better than pre-1906 buildings, but the overall level of earthquake resistance is probably higher now, since at least some of the buildings were built after earthquake building codes were instituted.

Subpopulation	Estimated fraction of population at 6 PM	Estimated death rate
People at home, awake	56%	21 per 100,000
People at home, asleep	7%	23 per 100,000
People in offices	15%	102 per 100,000
Cars near bridge/overpass	2%	203 per 100,000
Cars on open freeway	4%	3 per 100,000
Cars on urban street	4%	7 per 100,000
Children in school	5%	13 per 100,000
People outdoors	7%	7 per 100,000
Estimated overall death rate:		35 per 100,000

Table 6.5: Estimation of the overall death rate that can be expected in a major earthquake that occurs at 6 PM, using the expert estimates of baseline death rates from chapter 5 and estimates of how the population is distributed at 6 PM from chapter 3.

watch the World Series baseball game. ... Under "normal" circumstances, we could have expected a much higher death toll." However, we are not aware of data that support this assertion. Indeed, the survey results we reported in chapter 3 regarding where people where at the time of the earthquake did not appear to be substantially different from what one would expect based on data from the other sources discussed in chapter 3.

On balance, we conclude that our extrapolation in table 6.4 is likely to underestimate the death rate that can be expected if a magnitude 8.3 earthquake were to occur close to San Francisco, during the afternoon rush hour.

Comparison with Expert Estimates

Table 6.5 illustrates the computation of a death rate estimate based on expert judgment and our estimate of how the population is distributed at 6 PM. The resulting estimate of 35 deaths per 100,000 is less than one-sixth of the USGS/NOAA estimate of 224 deaths per 100,000.

6.2.2 Projected Death Toll Reduction as a Function of Warning Lead-Time

Figure 6-3 shows our projections of how many lives could have been saved if it had been possible to warn the public in advance of the Loma Prieta earthquake. According to the projection, less than a minute of lead-time would have sufficed to save the first 23 lives and with an additional fourteen minutes of warning, the next twenty lives could have been saved. After that, the projected number of lives saved does not appear to increase much; with 30 minutes of warning an estimated 45 lives could be saved.

To judge how plausible these projections are, we have reproduced in table 6.6 a summary from [CP$^+$89] of the circumstances under which the deaths caused by the Loma Prieta quake occurred. More than half (41) of the deaths were caused by the collapse of the Cypress Street Viaduct. An additional twelve fatalities were due to structural failure of brick walls or dwellings, two people fell from towers, and one person perished in a landslide onto a

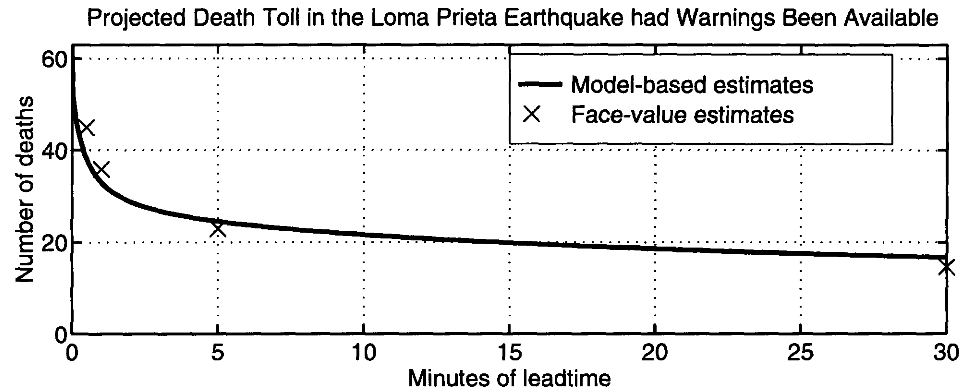

Figure 6-6: Projections of how many of the sixty three deaths officially attributed to the Loma Prieta earthquake could have been avoided if earthquake warnings had been available before the shaking started. The projections are based on the model-based and face-value analyses reported in 5 the "best-case" assumptions described there and in chapter 4.

coastal highway. The remaining 7 deaths occurred after the main shock was over; it is unlikely that foreknowledge of the earthquake could have prevented those deaths.

For the deaths caused by collapse of dwellings, one would expect that with as little as two minutes of warning, people could at least get to a comparatively safe place within the house and five minutes should suffice for most people to leave the house. If the people that fell from towers had known the shock was about to occur, they could at least have tried to hold onto something and five minutes might have been sufficient to climb down from the tower. Thus, we expect that of the 15 deaths that occurred *during* the quake and were not caused by the freeway collapse, about half could have been saved with two minutes of warning and all of them with five minutes of warning.

Consider now the question of whether foreknowledge of the quake could have saved the people that were crushed in their cars on the Cypress Street Viaduct. The section that collapsed was about a mile long, with the lower deck around 20 feet above ground. Thus, jumping off the freeway does not seem like a realistic option. If one supposes that cars on the lower deck would have had to drive up to two miles to exit the freeway and one assumes that the cars were traveling at 20 mph or faster, then 6 minutes would have sufficed to clear the lower deck of the elevated section that collapsed. One can imagine signs along the freeway that in the event of an earthquake warning would automatically instruct drivers to leave at the next exit[7]. With the technology envisioned by proponents of Intelligent Vehicle Highway Systems (IVHS) such instructions could even be communicated to each driver individually, for example with a computer display or an audio signal. The point is that warning drivers of an imminent potentially dangerous event and providing instructions on how best to avoid the danger is within the limits of currently available technology.

Thus, based on the circumstances that led to the fatalities in the Loma Prieta quake,

[7]We are assuming that leaving the freeway at the next exit would be the most appropriate action because of the upper deck that might collapse. However, given the demolition of the Cypress Street viaduct and structures of the same design, such conditions might not be very common in future California quakes. Hence, other actions might be more appropriate.

Circumstance	County	Number of deaths
Collapse of elevated freeway section[a]	Alameda	41
Brick wall collapse onto automobiles[a]	San Francisco	5
Brick wall collapse[a]	Santa Cruz	3
Brick wall collapse[a]	Monterey[b]	1
Dwelling collapse[a]	San Francisco	3
Fall on stairway[c]	San Francisco	2
Fall from tower[a]	Santa Clara	1
Fall from tower[a]	Monterey	1
Landslide on coastal highway[a]	Santa Cruz	1
Smoke inhalation from gas fire[d]	San Francisco	1
Gunshot wound while directing traffic[e]	San Francisco	1
Motor vehicle collision with damaged bridge section[a]	San Francisco	1
Motor vehicle collision with horse on highway[e]	Santa Cruz	1
Carbon monoxide inhalation from emergency generator[e]	Alameda	1

[a] Occurred within 2 minutes of the earthquake.
[b] One person injured in Santa Cruz County died in Monterey County.
[c] Occurred within 8 hours after the earthquake
[d] Preliminary determination
[e] Presumed indirectly earthquake related.

Table 6.6: Earthquake-related deaths – as determined by county coroners – for the Loma Prieta earthquake, by circumstance and county. From [CP+89].

City/region	Population	# Extricated	# Dead	Death rate
Leninakan	232,000	16,959	9,974	4,300 per 100,000
Kirovakan	171,000	4,317	420	246 per 100,000
Spitak	18,500	13,990	9,733	52,610 per 100,000
Stepanavan	21,000	108	63	300 per 100,000
Rural areas	146,500	4,421	4,352	2,971 per 100,000
Total:	589,000	39,795	24,542	4,167 per 100,000

Table 6.7: Number of deaths in the Armenian cities and towns that suffered in the 1989 earthquake. This data is quoted in [Abr89, page P-449, table 5] and is obtained from the Armenian Civil Defence.

we suggest that with two minutes of warning, perhaps about 6 deaths could have been avoided and five or six minutes of warning could have saved 53 lives. That is, the life-saving potential of warnings might have been even greater than is indicated by the projection of figure 6-6 (where an estimated 45 lives could be saved with 30 minutes of warning) for lead-times longer than five minutes. For lead-times shorter than five minutes, the projected death toll reduction may be too optimistic.

We now turn our attention to an earthquake that occurred in Armenia in 1988 and provides a stark contrast to the two California earthquakes we have described.

6.3 The 1988 Armenia Earthquake

A magnitude 6.8 earthquake struck Northern Armenia on December 7, 1988, at 11:41 AM [YG92]. While the magnitude of this earthquake was similar to that of the Northridge and Loma Prieta quakes, the number of fatalities caused by it was far higher. The author of an epidemiological study of the Armenia quake [Noj91], whose results we will examine, observes:

> The Armenian earthquake was of less magnitude than the 1989 Loma Prieta earthquake in California. However, its consequences were incomparably greater, primarily because of the design and quality of construction of buildings in the area. The primary cause of death, injury, and destruction was the total collapse of buildings that were not adequately designed for earthquake resistance.

According to official Civil Defence figures [Abr89], 41,666 people were extricated from collapsed buildings, of which 15,457 were alive and 26,209 were dead. Table 6.7 shows the death toll[8] in the various cities and towns affected. The statistics for the town of Spitak, which is located very close to the epicenter, are particularly sobering. Of the 18,500 inhabitants of the town, 13,990 were trapped in collapsed buildings and 9,733 were dead upon extrication. So according to these figures, more than half the population of Spitak died in the earthquake; the proportion is even higher if other estimates are to be believed.

[8]Once again, sources do not agree on the death toll. In fact, the figures of 26,209 (quoted in the text) and 24,542 (from table 6.7) total deaths are both from the Armenian Civil Defence. Other published estimates are as high as 55,000 [Rob93].

Location	Number of people	Number of deaths	Death rate	Rel. risk
Outside	651	57	8,760 per 100,000	1.0
Inside	7,120	3,923	55,100 per 100,000	6.3
Unknown	729	222	30,450 per 100,000	
Total	8,500	4,202	49,440 per 100,000	

Table 6.8: Effect of location on survival for three rural towns affected by the Armenia earthquake. From [Noj91].

The overall death rate for the region affected by the quake was 4,167 per 100,000. Even if we accept the USGS/NOAA estimates of life loss if a major earthquake were to strike either San Francisco or Los Angeles (we have argued in this chapter that the USGS/NOAA loss studies may have overestimated the death toll), this death rate is more than an order of magnitude larger. This is a common observation when comparing the life loss of two earthquakes of similar magnitudes where one occurs in a developing country and the other in an industrialized country. Thus, even if earthquake warnings could only reduce the death toll by 10% in a developing country, a larger fraction of the total population would be saved than is possible in, say, California.

We will not attempt to compute death rates for different intensity regions for the Armenia earthquake. Instead, we will examine the results of an epidemiological study of the pattern of fatalities, and we will use this information to speculate about how many lives might have been saved had the public been warned of the earthquake ahead of its occurrence.

6.3.1 Mortality Patterns

The results of two studies of earthquake-related mortality and morbidity in the Armenia quake are described by Eric Noji [Noj91]. The first study was a survey of three rural towns conducted immediately after the earthquake. In the second study, residents of the city of Leninakan that were hospitalized because of injury were interviewed, along with a sample of non-hospitalized residents. The researchers attempted to match the two samples in terms of age, sex, and city block of residence. Since the second study was concerned with morbidity patterns and we are primarily concerned with fatalities, we will only show the results of the first study.

The total population of the three rural towns was approximately 8,500 people prior to the earthquake. Table 6.8 shows the effect of location (inside or outside) on survival in the earthquake. More than half the people that were inside during the earthquake died, a death rate that is more than six times higher than for people that were outside.

A subset of the people that were inside were *trapped*, i.e., they were not able to leave the building unassisted after the earthquake. As table 6.9 shows, about 80% of people that were trapped died; a death rate that is 67 times higher than for people that were not trapped. By comparing tables 6.8 and 6.9 we see that the number of deaths among people that were not trapped is lower than the number of people that were outside and did not die. Thus, some people that were outside when the quake occurred must have been trapped, perhaps under debris that fell of buildings as they collapsed.

Most of the buildings that collapsed in the Armenia quake could be classified into one

	Number of people	Number of deaths	Death rate	Rel. risk
Not trapped	3,390	34	1,000 per 100,000	1.0
Trapped	5,110	4,160	81,410 per 100,000	67.3
Total	8,500	4,202	49,440 per 100,000	

Table 6.9: Effect of entrapment status on survival for three rural towns affected by the Armenia quake. From [Noj91].

of three types of construction [Noj91]:

Stone masonry: Almost all residential dwellings were one-story unreinforced masonry structures.

Precast concrete frame: Most industrial facilities were built of precast concrete elements that were welded or tied together to form frames.

Precast concrete panel: A few buildings were made of precast concrete panels, i.e., they did not have a frame.

Table 6.10 shows statistics from the Armenian Ministry of Internal Affairs and the State Committee for Construction regarding the effect of building type on survival for the Armenian town of Nalbad. The eight precast concrete frame buildings included a sewing factory, where 205 of the 212 workers were killed, and a school, where 285 of the 305 children perished. The death rate in precast concrete frame buildings was close to 90%, while the death rate in stone masonry dwellings was almost seven times lower. According to [Noj91],

> Infill masonry, panels and bricks often fell off, killing persons both inside and outside and the frequent collapse of stairways made it particularly difficult for people to escape, since many of these buildings had only one stairway. In all three building types, the collapse of non-structural elements such as parapets caused many serious injuries. The total collapse ("disintegration") of the precast-concrete-frame buildings was associated with particularly high mortality rates (greater than 90 %) because the characteristic failure pattern of this type of construction greatly complicated the search and rescue effort and reduced significantly the opportunity for occupant survival. We observed that the fragmentation of the floor system resulted in very tight packing of the rubble with no cavities or "void spaces" for possible survival of victims.

6.3.2 How Many Victims Could Have Been Saved had They Known the Earthquake Was Coming?

The statistics and comments we have quoted from [Noj91] suggest that residents that lived in one story stone masonry dwellings could have reduced their risk considerably by knowing of the earthquake in advance, even with very short lead-times. Getting out of a one story structure is unlikely to take more than a couple of minutes. The situation is different for people that were in multi-story precast concrete frame buildings. For these structures, it appears that unless the lead-time had been long enough for most people to exit the building, Effect of building type on survival in Nalbad

Building type	# of buildings	# of occupants	# of deaths	Death rate	Rel. risk
Stone masonry	38	415	53	12,770 per 100,000	1.0
Precast concrete panel	2	40	19	47,500 per 100,000	3.7
Precast concrete frame	8	577	502	87,000 per 100,000	6.8
Total	48	1,032	574	55,600 per 100,000	

Table 6.10: Effect of building type on survival in one rural town (Nalbad) affected by the Armenia earthquake. Source: [Noj91].

foreknowledge would have been of little use. This is because of the way these buildings collapsed: there were very few cavities in the ruins that might have allowed victims to survive.

Based on these considerations, we can speculate that with two minutes of warning, most of the inhabitants of one story stone masonry structures could have been saved, while as long as ten minutes might have been required to evacuate some of the multi-story precast concrete buildings. Even though we do not have the benefit of expert judgment that applies directly to conditions in Armenia, let us attempt to estimate how many of the 25,000 lives that were lost might have been saved if earthquake warnings had been possible. We will assume that the fatalities were equally divided between one story stone masonry dwellings (note from table 6.10 that in the town of Nalbad, about 40% of the deaths occurred in stone masonry houses). The death toll in stone masonry dwellings will be assumed to decay as $e^{-a\tau}$ and as $e^{-b\tau}$ for precast concrete dwellings. The time constants a and b were chosen so that the stone masonry death toll would be halved with 2 minutes of warning and the precast concrete death toll would be halved with 10 minutes of warning. Thus, our projection is given by

$$25,000 \times \left(\frac{1}{2} \exp^{-\tau \log 2/2} + \frac{1}{2} \exp^{-\tau \log 2/10} \right)$$

Figure 6-7 shows a projected death toll based on these assumptions.

The estimated life-savings in figure 6-7 are far greater, in absolute terms, than our projections for Los Angeles and San Francisco, simply because many more people died in the Armenia quake than can be expected to die in a California earthquake. But we also estimate the *proportional* reduction in the death toll to be greater, at least for lead-times beyond ten minutes or so, because the fraction of people whose deaths could not have been prevented with earthquake warnings (for example deaths caused by the Loma Prieta quake that occurred *after* the quake) was much lower in Armenia.

6.4 Summary of Estimates and Findings

- We attempted to extrapolate the official death tolls of the Northridge and Loma Prieta earthquakes to scenarios often referred to as the "Big One." Our extrapolations,

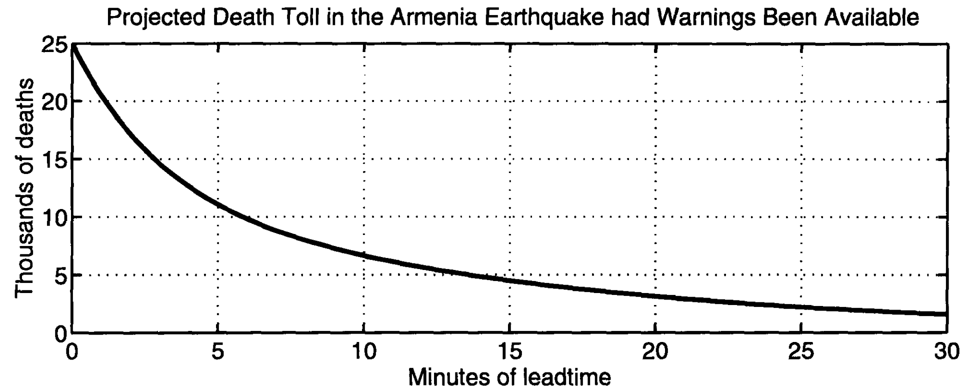

Figure 6-7: Projected death toll in the Armenia quake if earthquake warnings had been available before the shaking started. The projection assumes that the fatalities that occurred were equally divided among stone masonry dwellings and precast concrete structures, that the death toll in both types of buildings decays exponentially, and that the lead-times needed to halve the death toll are two minutes and ten minutes for stone masonry and precast concrete, respectively.

which are based on the presumption that the *local* effects near the epicenters of the Northridge and Loma Prieta quakes were comparable to what might be expected in the scenarios we considered, suggest that previous life-loss estimates (the USGS/NOAA loss studies) for these scenarios may be too pessimistic.

For example, one scenario postulates a magnitude 7.5 earthquake on the Newport-Inglewood fault under downtown Los Angeles. Our extrapolation suggests a death rate of 11.5 per 100,000 if the quake were to occur at 4:30 AM; the USGS/NOAA estimate is 48 per 100,000. An estimate based on expert judgment and our estimates of how the population would be distributed at 3 AM lies between the extrapolation and the USGS/NOAA estimate, at 24 per 100,000. The ordering of the estimates (extrapolation < expert estimate < USGS/NOAA estimate) was the same for a hypothetical magnitude 8.3 quake on the San Andreas fault near Los Angeles. The third scenario assumes a magnitude 8.3 quake on the San Andreas fault near San Francisco. While our extrapolation for this scenario was more tentative than for the others, because conditions near the epicenter of the Loma Prieta quake were less severe than assumed by this scenario, the ordering of the estimates seems to be the same in this case also.

- Comparisons between the patterns of fatalities in the Northridge quake (which occurred at 4:30 AM) and the Loma Prieta quake (which occurred at 5:04 PM) reinforce a conclusion reached in chapter 5 based on expert judgment and estimates of population shifts over the course of a day: that speculations about about how many deaths would have been caused by the Northridge quake had it occurred during a busy weekday may be exaggerated. These comparisons, and results from chapter 5 (figure 5-3), suggest that the difference between early-morning and late-afternoon death tolls in comparable future earthquakes might be expected to lie between 10% and 600%.

- We used the findings of chapter 5 to project how many of the fatalities in the

Northridge and Loma Prieta quakes might have been avoided if warnings had been available ahead of the shaking. For Northridge, we estimate that around 20 of the 60 deaths could have been avoided with about one minute of warning and an additional 20 lives might have been spared with 20 additional minutes of lead-time. For Loma Prieta, we project that the first 23 (out 63) deaths might have been avoided with less than one minute of warning and the next 20 lives could perhaps have been saved with 14 additional minutes of warning. Based on the actual circumstances leading to the deaths in the Northridge quake, we argue that the projection is plausible (in the sense of being within the bounds of what people can do). For Loma Prieta, the projection may be too optimistic, primarily because of the collapse of an elevated freeway that caused most of the deaths, but may be less likely to recur in the future (since most elevated freeways of similar design in California have now been torn down).

- Finally, we reviewed studies of the pattern of fatalities in the 1988 Armenia earthquake and used the results of those studies to prepare a tentative projection of the reduction in death toll that might have been possible with earthquake warnings. It seems clear that the life-saving potential of earthquake warnings in the third world is far greater than in California, simply because earthquakes in the third world cause many more deaths (the magnitude of the Armenia quake was similar to the Northridge and Loma Prieta quakes, but the death toll was 400 times greater). But we argue that the *proportional* death toll reduction may be greater also, because the fraction of third world earthquake deaths that cannot be prevented with earthquake warnings, for example because the death occurs after the earthquake, is lower than in California.

Appendix 6A: Compound Poisson Process Fatality Model

A simple model for the pattern of fatalities in an earthquake is that deaths occur according to a spatial Poisson process, whose intensity may vary as a function of population density and construction quality. But experience from recent earthquakes in the United States seems to suggest that this model is inappropriate: Each of the 1971 San Fernando, 1994 Northridge, and 1989 Loma Prieta earthquakes had one event (the Veterans Administration hospital, the Northridge Meadows apartment complex, and the Oakland freeway collapse) that caused a substantial fraction of the death toll. This seems to invalidate the assumption that the chances of dying for any two individuals are independent.

A better description would be to assume that the number of fatal *events*, i.e., structural failures, follows a Poisson process and the number of fatalities per event is a random variable. The pattern of fatalities would thus be assumed to follow a compound Poisson process.

To be explicit, let $N_{(\mathrm{IX})}$ be the number of people that experience MMI IX, and define the population sizes for other intensities in the obvious way. Suppose that the number of people among the $N_{(\mathrm{IX})}$ who die, $D_{(\mathrm{IX})}$, can be written as

$$D_{(\mathrm{IX})} = M_1 + M_2 + \cdots + M_k$$

where k is a Poisson random variable with mean $\lambda_{(\mathrm{IX})}N_{(\mathrm{IX})}$ and the M_i's are i.i.d. random variables. Think of k as the number of structures that fail and M_i as the number of people that die because of the i-th structural failure. Then the expected number of deaths in the MMI IX region is

$$\mathrm{E}[D_{(\mathrm{IX})}] = \lambda_{(\mathrm{IX})}N_{(\mathrm{IX})}\mathrm{E}[M]$$

and the variance is

$$\mathrm{var}[D_{(\mathrm{IX})}] = \mathrm{E}[k]\sigma_M^2 + \mathrm{E}[M]^2\sigma_k^2 = \lambda_{(\mathrm{IX})}N_{(\mathrm{IX})}\mathrm{E}[M^2]$$

so

$$\mathrm{var}[D_{(\mathrm{IX})}]/\mathrm{E}[D_{(\mathrm{IX})}] = \frac{\mathrm{E}[M^2]}{\mathrm{E}[M]}$$

Suppose that M has a geometric distribution with parameter p; then it has mean $1/p$ and variance $(1-p)/p^2$. The expected death toll would then be $\lambda_{(\mathrm{IX})}N_{(\mathrm{IX})}/p$, with variance $\lambda_{(\mathrm{IX})}N_{(\mathrm{IX})}(2-p)/p^2$. If the average number of deaths per structural failures is three (i.e., $p = 1/3$), then we have

$$\mathrm{var}[D_{(\mathrm{IX})}]/\mathrm{E}[D_{(\mathrm{IX})}] = \frac{2-p}{p} = 5$$

Thus, if we observe 24 deaths and accept all of the above assumptions, then we would estimate the mean of $D_{(\mathrm{IX})}$ to be 24 and the standard deviation to be $\sqrt{5 \times 24}$. An approximate 95% confidence interval would then be $24 \pm 2\sqrt{5 \times 24} = (2, 46)$.

Chapter 7

An Evacuation Model

Up to now, we have concentrated on warning lead-times shorter than an hour. With longer lead-times, a greater variety of risk-reducing actions become possible. This chapter offers a preliminary analysis of one such action: mass evacuation of the region that is expected to suffer in the earthquake. Besides earthquakes, our discussion may be relevant to planning for floods, hurricanes, nuclear reactor accidents, and other potentially catastrophic events where advance notice might be possible.

We present a simple model of the flow of people during an evacuation from their homes and places of work to a safe location, via a road network. According to the baseline death rate estimates developed using expert judgment in chapter 5, being in a car when an earthquake strikes is less safe than being at home: The average[1] death rate for "people in cars" is 45 per 100,000 while the death rate estimate for "people at home" is 21 per 100,000. Thus, if most of the population is at home when an earthquake warning accompanied by an evacuation advisory is issued, then we would expect that the population would be at higher risk just after the evacuation begins than just before the earthquake warning was issued. Hence an evacuation would only be recommended if the lead-time were long enough to allow most of the population to flow through the road network to safer locations away from the danger zone. If the exodus of people away from the city results in a paralyzing gridlock, then the lead-time needed for evacuation to be a sensible option could be quite long. The question we will explore with the model is: Under what conditions would evacuation be likely to lower the earthquake death toll rather than raise it?

7.1 The Model

This section presents a model of the flow of people during an evacuation. The model is in the form of a Markov chain (see figure 7-1), whose states are labeled "home[2]," "on the road," "broken down," and "safe area." Everyone is assumed to be either at home (with probability p_0) or on the road (with probability $1 - p_0$)) when an earthquake warning and evacuation advisory are issued. Then, people (or perhaps more appropriately, cars) flow from the "home" state to the "on the road" state, and from there to either the "broken down" or the "safe area" state, according to a set of transition rates.

[1]Assuming, as in chapter 3, that 20% of people in cars are on, under, or near bridges or overpasses, 40% are on the open freeway, and 40% are on urban streets.

[2]This state is assumed to contain everyone who is not initially in a car. For brevity, we refer to this group of people as the "at-home population."

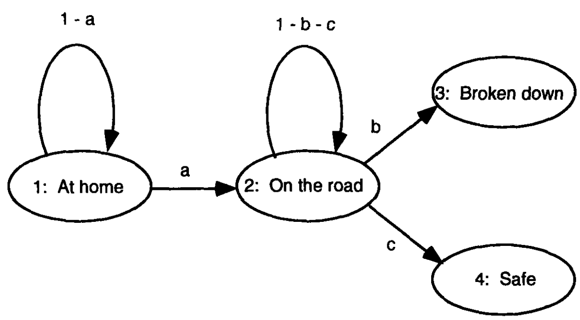

Figure 7-1: State transition diagram for a Markov chain model of evacuation flow. The "at home" state includes people at home, at work, or in school. The "broken down" state is inhabited by people whose cars run out of gas or stop moving for other reasons.

The three parameters of the model are a, b, and c, where a is the fraction of the at-home population at time t that is on the road by time $t+1$, and b and c are the fractions of the on-the-road population that run out of gas (or break down on the road for other reasons) and reach safety, respectively, during each time unit. We assume that $0 < a, b, c < 1$ and $0 < b + c < 1$, i.e., none of the indicated transition rates are zero. Note that we interpret the state probabilities $\pi(t)$ as fractions of the population that are in different states, rather than as probabilities of a particle being in different states.

The distribution across the states of the Markov chain evolves according to

$$\pi(t+1) = \begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \\ \pi_4 \end{bmatrix}(t+1) = \begin{bmatrix} 1-a & 0 & 0 & 0 \\ a & 1-b-c & 0 & 0 \\ 0 & b & 1 & 0 \\ 0 & c & 0 & 1 \end{bmatrix} \begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \\ \pi_4 \end{bmatrix}(t) = P^T \pi(t) \tag{7.1}$$

Taking time zero to be the time at which an earthquake warning is issued, the initial condition for the Markov chain becomes

$$\begin{bmatrix} \pi_1 \\ \pi_2 \\ \pi_3 \\ \pi_4 \end{bmatrix}(0) = \begin{bmatrix} p_0 \\ 1-p_0 \\ 0 \\ 0 \end{bmatrix}$$

Our purpose in analyzing the model is to assess the vulnerability of the population at the time of the earthquake, under different scenarios. The measure of vulnerability we will use is the expected fraction of the population that perishes in the earthquake, i.e., the death rate $D(\tau)$, which we will express as

$$D(\tau) = \sum_{i=1}^{4} D_i \pi_i(\tau) \tag{7.2}$$

where the D_i's are the death rates in each of the four states of the Markov chain in figure 7-1 and the $\pi_i(\tau)$'s are population fractions in each of the four states at time $t = \tau$ when the earthquake is assumed to occur. We will assume that the evacuation rates a, b, and c are

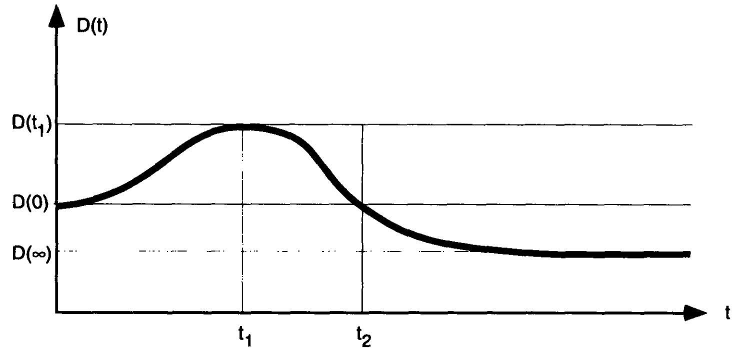

Figure 7-2: Generic shape of the overall death rate $D(t)$ as a function of the time t available for evacuation.

not affected by the warning lead-time τ.

If it is true that the population is less safe just after the evacuation begins than just before the warning is issued, then there will be a threshold time t_2 such that if the available lead-time is longer than t_2 $(\tau > t_2)$, then evacuation is likely to save lives but if the available lead time is shorter than t_2, the population will be less safe if evacuation were to be attempted. This is illustrated in figure 7-2. There are two interesting values of t: t_1, the time at which the expected death toll takes on its highest value, and the threshold time t_2.

Appendix 7A derives the following formula for $D(t)$:

$$D(t) = \alpha(1-a)^t + \beta(1-(b+c))^t + \gamma$$

where

$$\begin{aligned}
\alpha &= \frac{p_0}{a-(b+c)}\{D_1(a-(b+c)) - D_2 a + D_3 b + D_4 c\} \\
\beta &= \frac{1-p_0-a/(b+c)}{a-(b+c)}\{-D_2(b+c) + D_3 b + D_4 c\} \\
\gamma &= \frac{1}{b+c}\{D_3 b + D_4 c\}
\end{aligned}$$

From this we see that in the limit, as the evacuation proceeds for a very long time, the death rate approaches $\gamma = \frac{1}{b+c}(D_3 b + D_4 c)$, a weighted average of the death rates D_3 (in the broken down state) and D_4 (in the safe state). The appendix also shows that the function $D(t)$ must have either a shape similar to that in figure 7-2 or a shape that is obtained by reflection around a horizontal axis (see figure 7-6). The latter shape, where the death rate decreases initially and then increases towards it asymptotic value, might arise under the unlikely circumstances where being on the road is safer than being at home or in the safe state. This set of circumstances is not ruled out by the assumptions we have made so far, but even if it were possible for them to arise in reality, it is unlikely that mass evacuation would be seriously considered.

To focus attention on the more interesting possibilities, let us assume that the death rate in the safe state is zero ($D_4 = 0$), that the death rates on the road and in the broken down state are both equal to D_r ($D_3 = D_4 = D_r$) and let us refer to the at-home death rate

Time of day (t)	Fraction not in car ($p_0(t)$)
6 AM	97.5%
9 AM	90.0%
12 PM	95.0%
3 PM	92.5%
6 PM	90.0%
9 PM	96.5%
12 AM	97.5%
3 AM	98.5%

Table 7.1: Fraction of the population of a U.S. urban area that is *not* in their cars, at different times of day. Based on estimates developed in chapter 3.

as D_h ($D_1 = D_h$). In the next section, we will present a numerical example, using estimates of the transition rates that we hope are at least plausible, and then we will explore for which lead-time values evacuation would be recommended.

7.2 A Numerical Example

We start be developing estimates of the transition rates that we hope can at least be defended as plausible.

The fraction of the population initially at home: p_0

We have already made similar estimates, in chapter 3. Using those estimates, and interpreting "at home" to mean "not in a car," we obtain table 7.1.

The rate of leaving home a:

Under our Markov chain model, the time until a randomly chosen person leaves home will be a geometric random variable with mean $1/a$ time units. Suppose we fix the time unit at one minute, and let us assume that half the at-home population will have started the evacuation after half an hour, i.e., $(1-a)^{30} = 1/2$, which implies that $a \approx 0.02$. Thus, we will assume that about 2% of the population that has not yet started to evacuate will have done so one minute later.

The breakdown rate b:

The break-down rate b is the fraction of drivers on the road whose cars will cease to function for some reason during the next time unit. The reasons for break-down might be running out of gas, mechanical failure, or an accident. Let us consider the rate at which cars might run out of gas, and assume that the overall breakdown rate is double the running-out-of-gas rate.

Assume that at any time the distribution of gas tank levels in the population of cars is uniform between zero and one. Also assume that the size of gas tanks is constant, say 15 gallons. Finally, assume that all cars achieve the same gas mileage, say 20 mpg. Then the distance that a randomly chosen car can drive on the amount of gas it currently has will be

$$U \times 15 \text{ gallons } \times 20 \text{ mpg } = 300U \text{ miles}$$

where U is a random variable uniformly distributed between zero and one. Assuming that people do not stop for gas, the running-out-of-gas rate $b/2$ will depend on both the average speed v of vehicles and on the length of the time unit Δt, which we have fixed at one minute. The distance traveled per time unit will be $v\Delta t$. The break-down rate will be the fraction of cars that have less than $v\Delta t$ "miles of gas" in their tanks, or

$$b/2 = \Pr\{300U \leq v\Delta t\} = \Pr\{U \leq v\Delta t/300\} = v\Delta t/300$$

Suppose $v = 30$mph. Then

$$b/2 = \frac{30\text{mph} \times (1/60)\text{hours}}{300} = \frac{1}{600}$$

Thus, we will assume that the overall break-down rate equals $b = 1/300$ per minute.

Of course, this is a very rough estimate. For example, we have assumed that the running-out-of-gas rate does not change with time. If people do not stop for gas once they have started to evacuate, and the gas tank level is uniformly distributed at the beginning of the trip, then one might expect the rate of running out of gas to increase with time. On the other hand, some people probably will stop for gas and not all people will start the evacuation at the same time, so perhaps our assumption results in an adequate approximation.

The rate of reaching safety c:

Let us assume that cars are uniformly distributed over a "distance to safety" d. The distance traveled per time unit is $v\Delta t$, so

$$c = \frac{v\Delta t}{d}$$

is the fraction of cars that reach safety per time unit. As before, assume that $v = 30$ mph and $\Delta t = 1$ minute. Then, if $d = 50$ miles, we obtain

$$c = \frac{30 \times (1/60)}{50} = \frac{1}{100}$$

Thus, we will assume that the rate of reaching safety is three times larger than the break-down rate.

Example: Evacuation advisory issued at 3 PM

If we use the transition rate estimates developed above and the estimate $p_0(3\text{ PM}) = 92.5\%$ and assume that the "at home" death rate is 21 per 100,000 while the "on the road" death rate is 45 per 100,000, then under our assumptions the death rate will depend on the time available for evacuation in the manner illustrated in figure 7-3.

We see that for these numerical values (and assuming that all of our assumptions are reasonable), evacuation would be recommended if it was considered likely that at least 1 hour and 54 minutes would be available for evacuation before the earthquake struck. If the evacuation were to run for a very long time, then, according to our model, a quarter of the

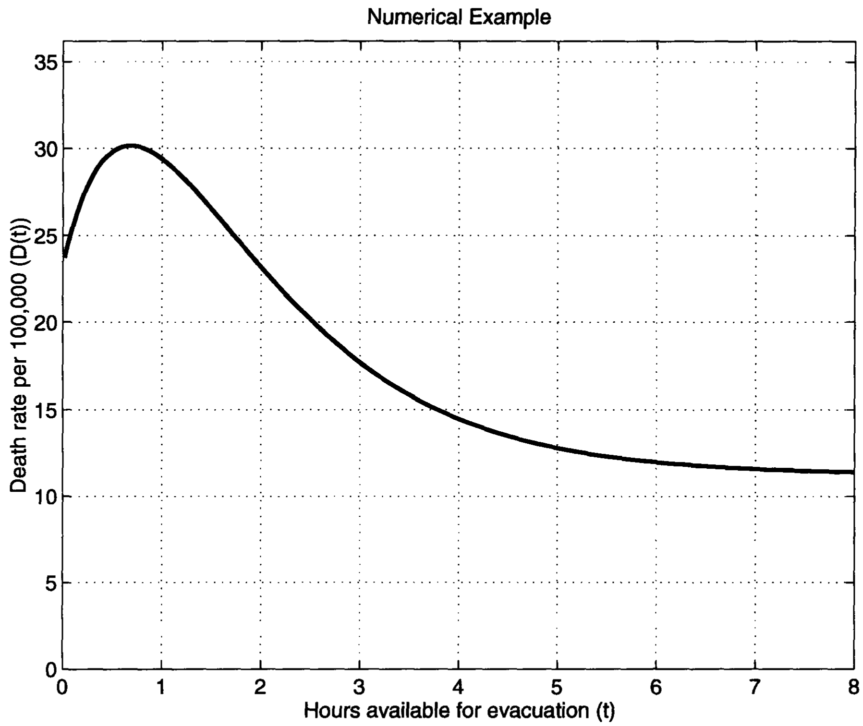

Figure 7-3: The death rate $D(t)$ that can be expected if t hours are available for evacuation, for the case when the transition rates are $a = 0.02$, $b = 1/300$, and $c = 1/100$, $p_0 = 92.5\%$, and $D_h = 21$ per 100,000, $D_r = 45$ per 100,000. The death rate is largest when 41 minutes are available for evacuation. If more than 1 hour and 54 minutes are available for evacuation, the death rate will be lower at the time of the earthquake than if evacuation had not been attempted (under the conditions of this numerical example).

population would eventually end up in the broken down state ($b/(b + c) = 1/4$) facing a death rate $D_r = 45$ per 100,000, while the other three quarters would reach the safe state. Thus the overall death rate would approach $45/4 = 11.25$ per 100,000, or about half the initial death rate $D(0)$.

We conclude this chapter by indicating how the model we have presented could be changed to allow for greater realism.

7.3 Generalizations of the Model

There are a variety of objections that could be raised to the simple model presented in this chapter, for example

1. The capacity of the roads depends on how many broken down cars there are. In terms of the diagram of figure 7-1, this means that the transition rate c, at time t is a function of $\pi_3(t)$, the fraction of cars that are broken down at time t.

2. The rate of people attempting to evacuate, a, may depend on time. One hypothesis is that the probability that a family attempts to evacuate, given that it has not already

done so, increases initially, as the threat of earthquake becomes more real, but then starts to decrease as the chances of being able to evacuate in time start to decrease.

3. People that are on the road at time zero may go home before trying to evacuate. This can be taken into account by allowing transitions from the on the road state to the home state, with a rate that decreases with time.

4. Cars don't necessarily stay in the broken down state forever. One might want to allow transitions from the broken down state back to the on the road state.

5. Assuming that the fraction of people that reach safety per time unit is proportional to the fraction of people that are currently on the road is too crude an approximation. *Where* on the road network people are needs to be taken into account also.

All but the last of these objections can be overcome, in principle, by allowing all possible transitions between the four states and by allowing the transition rates to depend both on time t and on the current distribution $\pi(t)$. The defining relation for such a model could for example be a differential equation of the general form

$$\dot{\pi}(t) = g(\pi(t), t)$$

As before, the death rate would be expressed as

$$D(t) = \sum_i D_i(t)\pi_i(t)$$

where we now allow the death rates D_i in each state to depend on time. The problem, of course, would be to estimate the transition rates and how they depend on time and on the fractions of the population currently in various states. In practice, one would have to make simplifying assumptions, but they would perhaps not need to be as rigid as in our initial model.

The last objection, that the rate at which people reach safety depends on exactly where on the road network people are, can only be addressed by adding additional states to the diagram of figure 7-1. A simple example is shown in figure 7-4, where the "on the road" state has been replaced by a series of states. The "broken down" state can be reached from any one of the "on the road" states, but the "safe" state can only be reached from the last "on the road" state. At the highest level of detail, the "on the road" state could be replaced with a faithful replication of the actual road network, i.e., there would be one state for any road segment in the area.

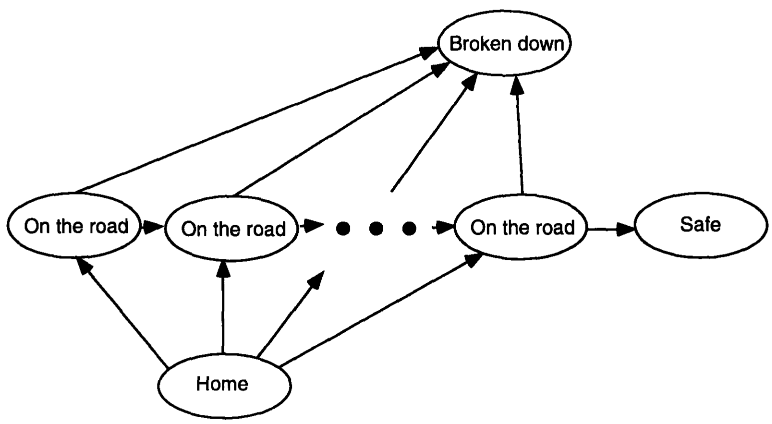

Figure 7-4: A possible generalization of the evacuation model of figure 7-1 to account for the topology of the road network in the area of interest.

Appendix 7A: Mathematical Details

Explicit Formula for $D(t)$

We first derive an explicit formula for the death rate function $D(t)$. To do this, we first diagonalize the transition matrix P:

$$
\begin{aligned}
P^T &= \begin{bmatrix} 1-a & 0 & 0 & 0 \\ a & 1-b-c & 0 & 0 \\ 0 & b & 1 & 0 \\ 0 & c & 0 & 1 \end{bmatrix} \\
&= \begin{bmatrix} \frac{a-b-c}{c} & 0 & 0 & 0 \\ \frac{-a}{c} & \frac{-(b+c)}{c} & 0 & 0 \\ \frac{b}{c} & \frac{b}{c} & 1 & 0 \\ 1 & 1 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1-a & 0 & 0 & 0 \\ 0 & 1-b-c & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \frac{c}{a-b-c} & 0 & 0 & 0 \\ \frac{ac}{(b+c)(-a+b+c)} & \frac{-c}{b+c} & 0 & 0 \\ \frac{b}{b+c} & \frac{b}{b+c} & 1 & 0 \\ \frac{c}{b+c} & \frac{c}{b+c} & 0 & 1 \end{bmatrix} \\
&= M\Lambda M^{-1}
\end{aligned}
$$

Then we express the initial condition $\pi(0)$ in terms of the columns m_i of M (which are independent eigenvectors of P^T), i.e.,

$$\pi(0) = My = \sum_1^4 y_i m_i \Leftrightarrow y = M^{-1}\pi(0)$$

$$\Leftrightarrow \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \end{bmatrix} = \begin{bmatrix} \frac{c}{a-b-c} & 0 & 0 & 0 \\ \frac{ac}{(b+c)(-a+b+c)} & \frac{-c}{b+c} & 0 & 0 \\ \frac{b}{b+c} & \frac{b}{b+c} & 1 & 0 \\ \frac{c}{b+c} & \frac{c}{b+c} & 0 & 1 \end{bmatrix} \begin{bmatrix} p_0 \\ 1-p_0 \\ 0 \\ 0 \end{bmatrix} = \begin{bmatrix} \frac{p_0 c}{a-(b+c)} \\ \frac{c(1-p_0)-ca/(b+c)}{a-(b+c)} \\ \frac{b}{b+c} \\ \frac{c}{b+c} \end{bmatrix}$$

The solution of $\pi(t+1) = P^T\pi(t)$ (equation (7.1)) can now be written as

$$
\begin{aligned}
\pi(t) &= \left(P^T\right)^t \pi(0) = \left(M\Lambda M^{-1}\right)^t My \\
&= M\Lambda^t M^{-1}My = M\Lambda^t y = \sum_1^4 \lambda_i^t y_i m_i \\
&= \frac{(1-a)^t p_0}{a-(b+c)} \begin{bmatrix} a-(b+c) \\ -a \\ b \\ c \end{bmatrix} \\
&\quad + \frac{(1-(b+c))^t(1-p_0-a/(b+c))}{a-(b+c)} \begin{bmatrix} 0 \\ -(b+c) \\ b \\ c \end{bmatrix} + \frac{1}{b+c} \begin{bmatrix} 0 \\ 0 \\ b \\ c \end{bmatrix}
\end{aligned}
$$

Now the death rate $D(t)$ can be written as

$$D(t) = D_1\pi_1(t) + D_2\pi_2(t) + D_3\pi_3(t) + D_4\pi_4(t)$$

$$
\begin{aligned}
&= \frac{(1-a)^t p_0}{a-(b+c)} \{D_1(a-(b+c)) - D_2 a + D_3 b + D_4 c\} \\
&\quad + \frac{(1-(b+c))^t(1-p_0-a/(b+c))}{a-(b+c)} \{-D_2(b+c) + D_3 b + D_4 c\} \\
&\quad + \frac{1}{b+c} \{D_3 b + D_4 c\}
\end{aligned} \tag{7.3}
$$

If $D_4 = 0$ (the risk of death in the "safe area" is zero), $D_3 = D_2 = D_r$ (the risk of death for people in cars is independent of whether the car is moving or not), and $D_1 = D_h$, then the expression for $D(t)$ simplifies to

$$
\begin{aligned}
D(t) &= \frac{(1-a)^t p_0}{a-(b+c)} \{D_h(a-(b+c)) + D_r(b-a)\} \\
&\quad + \frac{(1-(b+c))^t(1-p_0-a/(b+c))}{a-(b+c)}(-D_r c) + \frac{D_r b}{b+c}
\end{aligned}
$$

The baseline death rate is $D(0) = D_h p_0 + D_r(1-p_0)$.

The Shape of $D(t)$

From the preceding discussion, we see that for appropriate constants α, β, and γ we can write

$$D(t) = \alpha(1-a)^t + \beta(1-(b+c))^t + \gamma$$

If we let $\lambda_1 \equiv 1-a$ and $\lambda_2 \equiv 1-(b+c)$, we can write

$$D(t) = \alpha\lambda_1^t + \beta\lambda_2^t + \gamma \tag{7.4}$$

We have assumed that $0 < a, b, c < 1$ and $0 < b+c < 1$, which implies that $0 < \lambda_1, \lambda_2 < 1$. As originally defined, $D(t)$ takes values only for $t = 0, 1, \ldots$. But it will be convenient to take (7.4) as the definition of a function whose domain is the real line. Properties of this extended function will be deduced and then used to infer properties of the original function defined on the non-negative integers.

The n-th derivative of D is

$$D^{(n)}(t) = \alpha(\ln\lambda_1)^n\lambda_1^t + \beta(\ln\lambda_2)^n\lambda_2^t \text{ for } n = 1, 2, \ldots \tag{7.5}$$

Suppose we set $D^{(n)}(t) = 0$. This equation has a unique solution if α and β have opposite signs:

$$t(n) = \frac{\ln(-\beta/\alpha) + n\ln(\lambda_2/\lambda_1)}{\ln(\lambda_1/\lambda_2)} \tag{7.6}$$

If α and β have the same sign, then $D^{(n)}(t) = 0$ has no solutions.

From the last three equations we may deduce the following:

Proposition 3 *The function $D(t)$ defined by (7.2) has the following properties:*

1. $\lim_{t\to\infty} D(t)$ *exists.*

2. *Every derivative $D^{(n)}$ of D changes sign at most once on $(-\infty, \infty)$. In particular, if α and β have opposite signs, then $D^{(n)}$ changes sign once for all n, otherwise it never changes sign.*

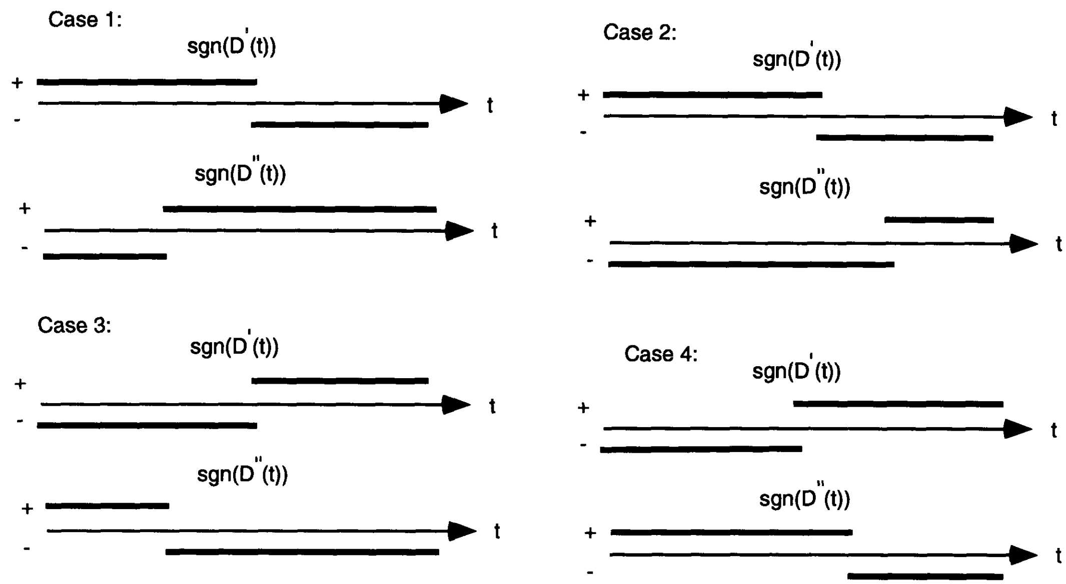

Figure 7-5: Potentially possible cases for the signs of the first two derivatives of D, and the times at which the signs change.

3. *The value at which $D^{(n)}$ changes sign is monotonic in n.*
4. *The derivatives of D alternate in sign in the following sense: If $D^{(n)}(t)$ changes from positive to negative as t increases, then $D^{(n+1)}(t)$ changes from negative to positive as t increases.*

Proof: The first property is obvious from (7.3); as t approaches infinity, D approaches $\frac{D_3 b + D_4 c}{b+c}$. The second and third properties follow from (7.6), i.e., from the uniqueness of the solution to $D^{(n)}(t) = 0$ and from the fact that $t(n)$ is linear in n with a nonzero slope.

To prove the fourth property, first note that as $t \to -\infty$, λ_1^t and λ_2^t both approach ∞. Also,

$$\lim_{t \to -\infty} \left(\frac{\lambda_1}{\lambda_2}\right)^t = \begin{cases} \infty & \text{if} \quad \lambda_1 < \lambda_2 \\ 0 & \text{if} \quad \lambda_1 > \lambda_2 \end{cases}$$

In words, the smaller eigenvalue will start to dominate as $t \to -\infty$. Therefore,

$$\begin{aligned} \lim_{t \to -\infty} D^{(n)}(t) &= \lim_{t \to -\infty} \alpha(\ln \lambda_1)^n \lambda_1^t + \beta(\ln \lambda_2)^n \lambda_2^t \\ &= \begin{cases} \operatorname{sgn}\left(\alpha(\ln \lambda_1)^n\right) \infty & \text{if} \quad \lambda_1 < \lambda_2 \\ \operatorname{sgn}\left(\beta(\ln \lambda_2)^n\right) \infty & \text{if} \quad \lambda_1 > \lambda_2 \\ \operatorname{sgn}\left((\alpha + \beta)(\ln \lambda_1)^n\right) \infty & \text{if} \quad \lambda_1 = \lambda_2 \end{cases} \end{aligned}$$

Since λ_1 and λ_2 are assumed to lie between 0 and 1, their logarithms will be negative. Therefore, the terms $(\ln \lambda_1)^n$ and $(\ln \lambda_2)^n$ will alternate in sign as n increases. ■

If we now restrict attention to the first two derivatives of D, we see that there are at most four possible cases, as shown in figure 7-5. However, it is not hard to see that case 1 is not possible. Let t_1 be the time at which D'' changes from negative to positive. We see

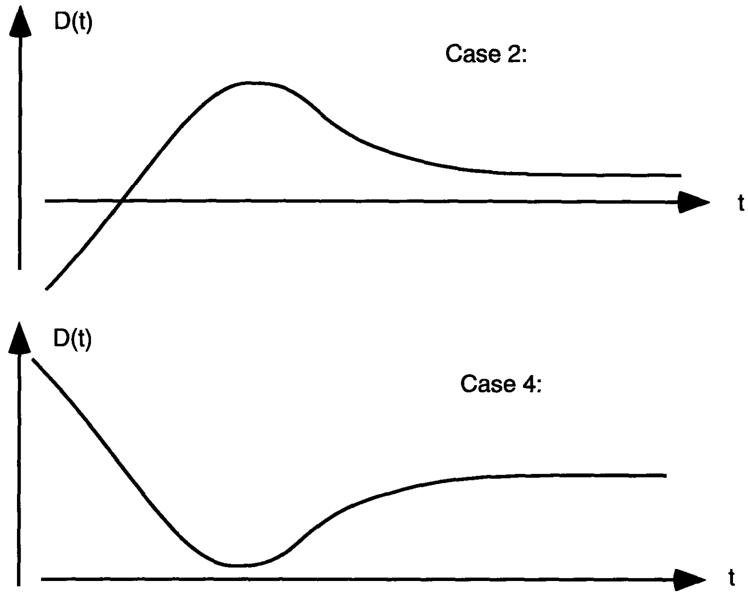

Figure 7-6: Possible cases for the shape of the function D, on $(-\infty, \infty)$.

that $f'(t_1) \geq 0$. Therefore, for all $t \geq t_1$ we have

$$D'(t) = D(t_1) + \int_{t_1}^{t} D''(s)ds \geq 0$$

since $D'(t_1) \geq 0$ and $D''(s) \geq 0$ for $s > t_1$. But this contradicts the indicated sign change for D'. A similar argument shows that case 3 is not possible either. The shape of D under the remaining two cases (cases 2 and 4) is shown in figure 7-6. The shape of D restricted to the non-negative integers will then be some "subset" of one of the two shapes shown in figure 7-6. Which subset depends on where the origin $t = 0$ falls on the shape.

For example, if D takes on the shape labeled Case 2, and $t = 0$ lies to the right of the maximum, then the restriction of D to the non-negative integers will be monotonically decreasing. This would happen if it were true that $p_0 = 0$ (everyone is on the road at time zero) and $D_2 = D_3 = D_r > D_4 = 0$. Then the decrease per unit time would be c times D_r times the fraction of people currently on the road, but not in the broken down state.

Chapter 8

Summary and Final Remarks

In this thesis, we have tried to estimate how many lives might be saved by a few minutes of unequivocal advance warning that a great earthquake is about to strike a California urban area. Our estimates were optimistic in that they were contingent on (1) the use of the most effective technologically feasible means of communicating the warning and (2) people having been trained to quickly recognize earthquake warnings and to react to them by taking an action that is likely to increase chances of survival, given the circumstances. By "unequivocal" we mean that our estimates are made under the assumption that there is no uncertainty associated with when and where the forecasted earthquake will occur.

We focus on last-minute forecasts for two reasons. First, such forecasts may be more realistic to hope for in the near future: The one form of earthquake "prediction" that seems technologically feasible at present is *real-time earthquake monitoring*, where warning of an earthquake that has already occurred is relayed ahead of the seismic waves from areas close to the epicenter to outlying areas that might experience damage. In certain locations, notably Mexico City, such technology could provide citizens with about 50 seconds of warning. Second, most research on human reactions to disaster warnings appears to assume warning lead-times measured in hours or days, either implicitly or explicitly. With only a few minutes to act, the constraints imposed on human reactions by time, space, and available communication technologies become so much tighter that an approach that is different from the one usually adopted by social scientists in their study of earthquake prediction seems called for.

There are clear difficulties involved in making such estimates. The most obvious is that seismologists are not currently able to issue warnings of the sort we envision; indeed they may never be able to pinpoint the time and place of future earthquakes as exactly as our best-case estimates assume. But this difficulty is also a prime motivation for our study: Public funds are used to further the science of earthquake prediction and this research is justified, at least in part, on the grounds that it would provide a social good. Recognizing that a finite amount of resources are available for improving public health and safety, it seems prudent to ask: What is the life-saving potential of earthquake prediction and how does this potential compare to other means of improving public safety? While we do not carry out such a comparison, our estimates could be helpful in such an analysis.

Our estimation procedures were guided throughout by a model of how the expected death rate D experienced in a large earthquake depends on the lead-time τ citizens of the area of interest have for preparing themselves for the shaking:

$$D(\tau, t) = \sum_{i \in I} F_i(t) D_i(0)(1 - R_i(\tau))$$

Here I is a set of categories such as "people at home, awake" and "people in cars on urban streets," $F_i(t)$ is the fraction of the population that can be expected to fall into category i when the time of day is t, $D_i(0)$ is the death rate that could be expected in category i if the earthquake were to occur without warning, and $R_i(\tau)$ is the proportional reduction in the expected death rate that could be achieved with τ minutes of warning.

We estimated key quantities in the above model in various ways. Our population fraction estimates $F_i(t)$ were based on data from several sources that were complemented with informed guesswork where necessary. Estimates of the baseline death rates $D_i(0)$ arose from historical data and estimates of losses in future earthquakes that have been prepared for the Los Angeles and San Francisco areas. However, baseline death rate estimates were not available for all the categories that we used, and estimates of the proportional risk reduction factors $R_i(\tau)$ were not available elsewhere, to our knowledge, since experience with unequivocal earthquake warnings is non-existent. For this reason, we surveyed various experts regarding these quantities. Of 67 questionnaires mailed to such experts as seismologists, earthquake engineers, architects, sociologists, and disaster managers, 45 (i.e., 67%) were returned.

We used two parallel but entirely different approaches to combine the expert estimates with our population fraction estimates. To begin, we accepted the expert estimates at *face-value* and computed point estimates and confidence intervals based on the expert judgment and a minimum of additional assumptions. Then, we went beyond the face-value approach to a *model-based* procedure, which involved certain assumptions about human behavior in response to earthquake warnings and contained parameters that were estimated using expert judgment. The behavioral assumptions were translated into functional forms for two important random variables: T_A, the time until a randomly chosen person M in a particular category hears the earthquake warning, and T_C, the time M uses to complete an action designed to improve chances of survival in the earthquake (for example "get under a sturdy desk" might be appropriate for people inside buildings if the lead-time is too short to leave the building).

Among our primary findings were:

- The experts' estimates of baseline death rates (i.e., without warning), when combined into one estimate for each category, ranged from 203 per 100,000, for "people in cars on, under, or near freeway overpasses or bridges" to 3 per 100,000 for "people in cars on the open freeway." The estimate for "people at home, awake" was 21 per 100,000.

- The baseline death rate estimates for individual categories, based on expert judgment, were combined into an estimate for the overall population using the estimated population fractions $F_i(t)$. The resulting estimate ranged from 25 per 100,000 (at 3 AM) to 50 per 100,000 (at 3 PM). The factor of two difference is smaller than commonly quoted estimates [SAL$^+$80] indicate and far smaller than speculations that appeared in newspapers after the recent Northridge earthquake. Not only is the time-of-day variation of our estimates smaller than for death rate estimates [SAL$^+$80] based on hypothetical earthquake scenarios for Los Angeles and San Francisco deserving of the label "the Big One;" the absolute magnitude of our estimate is also smaller: the estimates reported in [SAL+80] range from 33 to 61 per 100,000, for 2:30 AM, and from 122 to 229 per 100,000, for 4:30 PM.

- The face-value and model-based results about the life-saving value of warnings were similar at the level of individual categories. Under both approaches, a dichotomy between "people inside buildings" and "people outside buildings" (outdoors or in their cars) was apparent, with slightly longer lead-times generally required to reduce the estimated death rates by a given proportion for people outside buildings.

- With *one minute* of warning lead-time – as thought feasible for some regions with real-time earthquake monitoring – we estimate, based on expert judgment and our assessment of the population fractions $F_i(t)$, that the overall death toll caused by a major earthquake striking an urban area in California could be reduced by *45%*, if citizens used that time optimally (in the sense we have described). The model-based and face-value approaches resulted in almost identical estimates of the one-minute life saving potential.

- Furthermore, we estimate that 1.5 to 3 minutes of unequivocal warning could halve the death rate; an additional 22 to 28.5 minutes would be required to reduce the remaining death rate by half (i.e., the total risk by three-quarters), according to our estimates. Thus, we expect a substantial "diminishing returns" effect: most of the one and a half to three minute potential is achieved during the first minute (45 out of 50 percent), and most of the 25 to 30 minute potential is achieved during the first one and a half to three minutes (50 out of 75 percent). The overall (i.e., for the whole population) results of the face-value and model-based approaches agreed closely: the 1.5 and 28.5 minute figures above are based on the model-based approach and the 3 and 22 minute figures on the face-value approach.

After preparing estimates of the number of lives that might be saved if earthquake warnings were possible before a major earthquake strikes a California urban area, we studied the implications of three recent earthquakes both for our estimates of the life-saving potential of warnings and for available estimates of life loss in future quakes that occur without warning.

The three earthquakes were the Northridge earthquake, which occurred at 4:30 AM on January 17, 1994, the Loma Prieta earthquake that struck the San Francisco Bay Area at 5:04 PM on October 18, 1989, and an earthquake that leveled towns in northern Armenia, at 11:41 AM on December 7, 1988.

The two California quakes, while not comparable in their overall effects to "the big one" expected in the future, caused damage levels close to their epicenters similar to what might be expected to prevail over a larger area when the big one arrives. This provides an opportunity to extrapolate the death tolls (60 deaths for Northridge, 63 for Loma Prieta) experienced in those two quakes to scenarios on which available life loss estimates are based. Furthermore, the different times of day at which the two earthquakes occurred provide a natural opportunity to assess how much the number of deaths in future quakes can be expected to depend on time of day. Finally, the circumstances leading to the deaths caused by those two earthquakes allow one to judge the plausibility of our estimates of how many lives could have been saved, in the best case, had unequivocal earthquake warnings been available before those tremors.

The Armenia quake, which was of a similar magnitude to the two California quakes, caused around 25,000 deaths according to official estimates and thus illustrates how much more devastating earthquakes generally are in the third world.

Among the findings that emerged from our study of these three earthquake are:

- Extrapolations of the Northridge and Loma Prieta death tolls to earthquake scenarios thought to be representative of "the big one" resulted in death rate estimates *far lower* than commonly quoted estimates [SAL+80]. The Northridge death toll, when extrapolated to a hypothetical magnitude 8.3 quake on the San Andreas fault near Los Angeles, results in a projected 2 deaths per 100,000; a magnitude 7.5 tremor on the Newport-Inglewood fault under Los Angeles results in a projected 12 deaths per 100,000. These may be compared with death rate estimates from [SAL+80] of 33 and 48 per 100,000, respectively, if the hypothetical quakes were to occur at 2:30 AM (the Northridge quake occurred at 4:30 PM). An extrapolation for the Loma Prieta quake, while more tentative, led to similar conclusions.

  Conceivably, Angelenos were lucky that the death toll was not higher in the Northridge quake and higher death tolls can be expected in similar earthquakes, striking at the same time-of-day in the future. However, it is also possible that the original estimates were overly pessimistic. Furthermore, it is likely that buildings and structures in California are better able to withstand earthquakes now, on average, than in the 1970's, when the original estimates were prepared.

- Loma Prieta death rates were 10% higher (for regions experiencing Modified Mercalli intensity (MMI) VIII) and 600% higher (for regions experiencing MMI VII and lower) than in the Northridge quake. The factor of 600% was heavily influenced by a single event: the collapse of a stretch of elevated freeway. The demolition of some and strengthening of other freeways of similar design may have reduced the likelihood of such tragedies in future California quakes. Without the freeway collapse, the factor of 600% decreases to 200%. For comparison, death rate estimates [SAL+80] for the scenario earthquakes we have been discussing are generally about four times higher at 4:30 PM than at 2:30 AM. The time-of-day variation in estimates based on expert judgment and our population fraction assessments is about 200%. Thus, there is reason to suspect that a factor of four difference between death tolls that can be expected at different times of day is an overestimate.

- Our findings regarding how many lives could have been saved had warnings been possible before the Armenia quake were more provisional than for California since we did not have the benefit of expert judgment directly applicable to Armenian conditions. However, there is little doubt that the absolute life-saving potential of earthquake warnings is far greater in the third world than in industrialized countries, simply because the earthquake death tolls experienced in the third world are much greater. Using data from an epidemiological study [Noj91] of the mortality pattern in the Armenia quake, we argued that the *proportional* life-saving benefit may also be greater, at least for warning lead-times greater than ten minutes. This is because the death rate for people inside buildings was far higher for the Armenia quake than for the two California quakes, while the death rate for people in open areas was presumably very low in both cases. Thus, people in Armenia could have reduced their death risk more than Californians by exiting buildings. In particular, we projected that a ten minute warning might have prevented 18 thousand of the 25 thousand deaths officially attributed to the Armenia quake and thirty minutes would have sufficed to save about 24 thousand lives.

To put our findings in perspective, imagine that a magnitude 7.5 earthquake on the Newport-Inglewood fault (or any of the other faults under Los Angeles) strikes in the next few years. Based on our extrapolation for the Northridge quake, expert judgment, and our estimates of how the population is distributed across categories at various times of day, the overall death toll suffered by the 15 million inhabitants of the Los Angeles metropolitan area could range from about 1,800 (if the quake were to strike in the early morning hours) to 3,600 (if the big one were to hit during the afternoon rush hour). Corresponding estimates from [SAL$^+$80], scaled up to account for population increases, are 7,200 at 2:30 AM and 34,400 at 4:30 PM. If our estimates are closer to the truth and the quake strikes in the middle of the night, the death toll of 1,800 would make this the second deadliest[1] natural disaster in U.S. history – unless the big one strikes San Francisco first.

If the population of Los Angeles were warned of the quake one minute ahead of its occurrence and if most people believe the warning instantly and attempt to increase their chances of survival (for example by exiting buildings, if there is enough time to do so, or seeking shelter inside a building otherwise), then our estimates indicate that the death toll could be reduced by about 600 out of 1,800 (if the quake hits at 3 AM) or 1,600 out of 3,600 (if the quake strikes at 3 PM)[2] . A real-time earthquake monitoring system for the Los Angeles area is considered technically feasible; however such a system would provide lead-times shorter than one minute for most residents if the earthquake were centered within the area. If the earthquake were to occur on the San Andreas fault, 30 miles north-east of downtown Los Angeles, then some residents might receive warning close to one minute before the onset of strong shaking.

A scientific break-through that enables earth scientists to warn Angelenos of the big one five minutes before the shaking starts could – by our estimates – prevent 1,000 out of 1,800 or 2,300 out of 3,600 deaths, depending on when the big one occurs. Such a break-through could presumably also save lives if a major earthquake were to hit the San Francisco Bay region and death tolls in less catastrophic earthquakes could be reduced also.

Once again, these estimates are contingent on many simplifying and often optimistic assumptions. They may be thought of as measuring the *possible* life-saving benefits individuals could derive from unequivocal earthquake warnings, rather than the *probable* benefit under present levels of public awareness. The estimates are uncertain and provisional, as all estimates of losses in large earthquakes must be, because fortunately such catastrophes are rare. And yet sensible policy decisions about how to mitigate earthquake hazard require such estimates. We hope that this thesis contributes to the debate about protection from earthquakes in a constructive manner.

[1]A hurricane accompanied by a flood is estimated to have killed 6,000 people in Galveston, Texas, in the year 1900.

[2]The overall proportional reduction estimates are slightly different for different times of day.

Bibliography

[Abr89] Joel I. Abrams. Detection and Extrication in the Armenian Earthquake. In *International Workshop on Earthquake Epidemiology for Mitigation and Response*, pages P–435 – P–449, July 1989.

[AE91] Duncan Carr Agnew and William L. Ellsworth. Earthquake Prediction and Long-Term Hazard Assessment. *Reviews of Geophysics*, 29, April 1991.

[AHC+73] S. T. Algermissen, Margaret Hopper, Kenneth Campbell, W. A. Rinehart, David Perkins, Karl V. Steinbrugge, Henry J. Lagorio, Donald F. Moran, Lloyd S. Cluff, Henry J. Degenkolb, C. Martin Duke, George O. Gates, Dan W. Jacobson, Robert A. Olson, and Clarence R. Allen. *A Study of Earthquake Losses in the Los Angeles, California Area.* National Oceanic & Atmospheric Administration, 1973.

[All82] Clarence R. Allen. Earthquake Prediction – 1982 Overview. *Bulletin of the Seismological Society of America*, 72(6), December 1982.

[And85] Richard Andrews. Earthquake Prediction and Public Policy: Recent Experiences in California. *Earthquake Prediction Research*, 3, 1985.

[ARD+72] S. T. Algermissen, W. A. Rinehart, James Dewey, Karl V. Steinbrugge, Henry J. Lagorio, Henry J. Degenkolb, Lloyd S. Cluff, Frank E. McClure, Stanley Scott, and Richard F. Gordon. *A Study of Earthquake Losses in the San Francisco Bay Area.* National Oceanic & Atmospheric Administration, 1972.

[BHMS75] B. A. Bolt, W. L. Horn, G. A. MacDonald, and R. F. Scott. *Geological Hazards.* Springer Verlag, New York, NY, 1975.

[Bol93] Bruce A. Bolt. *Earthquakes.* W. H. Freeman and Company, New York, NY, 1993.

[BRG93] Linda B. Borque, Lisa A. Russell, and James D. Goltz. Human Behavior During and Immediately After the Earthquake. In Patricia Bolton, editor, *The Loma Prieta, California, Earthquake of October 17, 1989 – Public Response*, Professional Paper 1553-B, pages B3 – B22. U. S. Geological Survey, 1993.

[Bur85] Paul W. Burton. Geophysics – Electrical Earthquake Prediction. *Nature*, 315, 30 May 1985.

[Cen90] Central United States Earthquake Preparedness Project (CUSEPP). *Estimated Future Earthquake Losses for St. Louis City and County, Missouri.* Earthquake Hazards Reduction Series 53. Federal Emergency Management Agency (FEMA), 1990.

[CL55] Earl A. Coddington and Norman Levinson. *Theory of Ordinary Differential Equations.* McGraw-Hill, New York, NY, 1955.

[CP+89] D. P. Cain, C. C. Plummer, et al. Earthquake-Associated Deaths – California. *Journal of the American Medical Association*, 262(22):3112 – 3115, December 8 1989.

[CS92] Andrew Coburn and Robin Spence. *Earthquake Protection.* John Wiley & Sons, New York, NY, 1992.

[Dav94] Keay Davidson. Learning from Los Angeles. *Earth*, 3(5):41 – 47, September 1994.

[DeP93] Anthony DePalma. Mexico Turns Off Its Quake-Warning System. *New York Times*, page 12, December 19 1993.

[Dew94] Jim Dewey, April 1994. Personal communication.

[DOT88] U. S. DOT/FHWA. Highway Statistics, 1988.

[Dra86] Thomas E. Drabek. *Human System Responses to Disaster: An Inventory of Sociological Findings.* Springer-Verlag, New York, NY, 1986.

[Ear89] Earthquake Engineering Research Institute. Loma Prieta Earthquake October 17, 1989: Preliminary Reconnaissance Report, 1989.

[Ear90] Earthquake Engineering Research Institute. Loma Prieta Earthquake Reconnaissance Report, 1990. Supplement to volume 6 of *Earthquake Spectra.*

[Ear94] Earthquake Engineering Research Institute. Northridge Earthquake January 17, 1994: Preliminary Reconnaissance Report, 1994.

[ENN86] Julio A. Echevarría, Kathryn A. Norton, and Roger D. Norton. The Socio-Economic Consequences of Earthquake Prediction: A Case Study in Peru. *Earthquake Prediction Research*, 4, 1986.

[Eve82] Jack F. Evernden. Earthquake Prediction: What We Have Learned and What We Should Do Now. *Bulletin of the Seismological Society of America*, 72(6), 1982.

[Fin92] Robert Finn. Rumblings grow about Parkfiled in wake of first earthquake prediction. *Nature*, 359, October 1992.

[FV92] A. Fligner and Joseph S. Verducci, editors. *Probability Models and Statistical Analyses for Ranking Data*, chapter 8: Rank Correlations and the Analysis of Rank-Based Experimental Designs by M. Alvo and P. Cabilio. Springer-Verlag, 1992.

[Gel91] Robert J. Geller. Shake-up for Earthquake Prediction. *Nature*, 352, 25 July 1991.

[Gli86] Theodore S. Glickman. A Methodology for Estimating Time-of-Day Variations in the Size of a Population Exposed to Risk. *Risk Analysis*, 6:317 –324, 1986.

[GS84] James M. Gere and Haresh C. Shah. *Terra Non Firma.* Stanford Alumni Association, Stanford, California, 1984. Reprint.

[HLR89] Richard Holden, Richard Lee, and Michael Reichle. *Technical and Economic Feasibility of an Earthquake Warning System in California.* Special Publication 101. California Department of Conservation, Division of Mines and Geology, March 1989.

[HLS+75] Margaret Hopper, Charley J. Langer, William J. Spence, A. M. Rogers, S. T. Algermissen, Bruce C. Olson, Henry J. Lagorio, and Karl V. Steinbrugge. *A Study of Earthquake Losses in the Puget Sound, Washington, Area.* Open-File Report 75-375. U. S. Geological Survey, 1975.

[Jul90] Bruce R. Julian. Are Earthquakes Chaotic? *Nature*, 345, 7 June 1990.

[KG81] Kaplan and Garrick. On the Quantitative Definition of Risk. *Risk Analysis*, 1, 1981.

[KJR82] Samule Kotz, Norman L. Johnson, and Cambell Read. *Encyclopedia of Statistical Sciences.* Wiley, New York, NY, 1982.

[KS61] Maurice G. Kendall and Alan Stuart. *The Advanced Theory of Statistics*, volume 2: Inference and Relationship. Hafner Publishing Company, New York, NY, 1961.

[Lin90] Allan G. Lindh. Earthquake Prediction Comes of Age. *Technology Review*, February/March 1990.

[Lom94a] Cinna Lomnitz, April 1994. Personal communication.

[Lom94b] Cinna Lomnitz. *Earthquake Prediction.* John Wiley & Sons, New York, NY, 1994.

[MO93] Dennis S. Mileti and Paul W. O'Brien. Public Response to Aftershock Warnings. In Patricia Bolton, editor, *The Loma Prieta, California, Earthquake of October 17, 1989 – Public Response*, Professional Paper 1553-B, pages B31 – B41. U. S. Geological Survey, 1993.

[Mog86] Kiyoo Mogi. Recent Earthquake Prediction Research in Japan. *Science*, 233, 18 July 1986.

[Mos88] Charles Moseley. Slow Progress in Earthquake Prediction. *Editorial Research Reports*, 2(2), July 15 1988.

[Noj91] Eric K. Noji. Epidemiologic Studies from the 1988 Armenia Earthquake: Implications for Casualty Modelling. In *Workshop on Modelling Earthquake Casualties for Planning and Response*, pages 37 – 64, December 1991.

[Pan89] Panel on Earthquake Loss Estimation Methodology, Committee on Earthquake Engineering, Commission on Engineering and Technical Systems, National Research Council. *Estimating Losses From Future Earthquakes.* National Academy Press, Washington, D. C., 1989.

[PC78] Marie-Elisabeth Paté-Cornell. Public Policy in Earthquake Effects Mitigation: Earthquake Engineering and Earthquake Prediction. Technical Report 30, The John A. Blume Earthquake Engineering Center, Department of Civil Engineering, Stanford University, May 1978.

[Pre80] Frank Press. Earthquake Prediction. *Earthquakes and Volcanoes: Readings from the Scientific American*, 1980.

[PS79] Marie-Elisabeth Paté and Haresh C. Shah. Public Policy Issues: Earthquake Prediction. *Bulletin of the Seismological Society of America*, 69(5):1533 – 1547, October 1979.

[PS80] Marie-Elisabeth Paté and Haresh C. Shah. Public Policy Issues: Earthquake Engineering. *Bulletin of the Seismological Society of America*, 70(5):1955 – 1968, October 1980.

[RA93] Mansour Rahimi and Glenn Azevedo. Building Content Hazards and Behavior of Mobility-Restricted Residents. In Patricia Bolton, editor, *The Loma Prieta, California, Earthquake of October 17, 1989 – Public Response*, Professional Paper 1553-B, pages B51 – B62. U. S. Geological Survey, 1993.

[RAh+76] A. M. Rogers, S. T. Algermissen, W. W. hays, D. M. Perkins, David O. Van Strien, Hoffman C. Hughes, Richard C. Hughes, Henry J. Lagorio, and Karl V. Steinbrugge. *A Study of Earthquake Losses in the Salt Lake City, Utah Area.* Open-File Report 76-89. U. S. Geological Survey, 1976.

[Res93] Nielsen Media Research. 1992 – 1993 Report on Television, 1993.

[Ric58] Charles F. Richter. *Elementary Seismology.* W. H. Freeman and Company, San Francisco, California, 1958.

[Rob93] Robert Farnighetti (editor). *The World Almanac and Book of Facts, 1994.* Press Pub. Co., Mahwah, NJ, 1993.

[RS85] Cristopher Rojahn and Ronald L. Sharpe. *Earthquake Damage Evaluation Data for California (ATC-13).* Applied Technology Council (ATC), 1985.

[SAL+80] Karl V. Steinbrugge, S. T. Algermissen, Henry J. Lagorio, Lloyd S. Cluff, and Henry J. Degenkolb. *Metropolitan San Francisco and Los Angeles Earthquake Loss Studies: 1980 Assessment.* Open-File Report 81-113. U. S. Geological Survey, 1980.

[Sav93] J. C. Savage. The Parkfield Prediction Fallacy. *Bulletin of the Seismological Society of America*, 83(1):1 – 6, February 1993.

[SBHBD90] William D. Schulze, David S. Brookshire, Renatte K. Hageman, and Shaul Ben-David. Should We Try to Predict the Next Great U.S. Earthquake? *Journal of Environmental Economics and Management*, 18, 1990.

[SC93] Carl W. Stover and Jerry L. Coffman. *Seismicity of the United States, 1568 - 1989 (Revised).* Professional Paper 1527. U. S. Geological Survey, 1993.

[Ste74] K. V. Steinbrugge. Reports on "Methodology" and "Methodology Improvements". An unpublished report, 1974.

[Stu90] Nationwide Personal Transportation Study, 1990.

[Swi92] David Swinbanks. Japanese report shows cracks in predicting earthquakes. *Nature*, 358, July 1992.

[Taz92] Haroun Tazieff. *Earthquake Prediction.* McGraw - Hill, New York, New York, 1992.

[Tur91] Donald L. Turcotte. Earthquake Prediction. *Annual Review of Earth and Planetary Sciences*, 19:263 – 281, 1991.

[Twa91] Mark Twain. *Gold Miners & Guttersnipes.* Chronicle Books, San Francisco, CA, 1991. Tales of California by Mark Twain, selected by Ken Chowder.

[UNK91] UNKNOWN. Waiting for the Big One. *The Economist*, 321(7736):101 – 103, December 7th 1991.

[UNK93] UNKNOWN. *Assessment of the State of the Art Earthquake Loss Estimation Methodologies.* Risk Management Software, Inc. (RMS) and California Universities for Research in Earthquake Engineering (CUREe), 1993. Task 1 Report.

[VL91] P. Varotsos and M. Lazaridou. Latest Aspects of Earthquake Prediction in Greece Based on Seismic Electric Signals. *Tectonophysics*, 188(3 - 4):321 — 347, 1991.

[WJS94] Robin M. Wagner, Nicholas P. Jones, and Gordon S. Smith. Risk Factors for Casualty in Earthquakes: The Application of Epidemiologic Principles to Structural Engineering. *Structural Safety*, 13:177 – 200, 1994.

[Wor90] Working Group on California Earthquake Probabilities. *Probabilities of Large Earthquakes Occurring in the San Francisco Bay Region, California.* Circular 1053. U. S. Geological Survey, 1990.

[YG92] Mishac K. Yegian and Vahé G. Ghahraman. The Armenia Earthquake of December 1988. Technical Report CE-92-11, Northeastern University, Boston, MA, October 1992.