Chapter 9 How to use early warning information

Ultimately, the goal of seismic early warning is to provide users with information that will help them in deciding the optimal course of action in the few seconds before the onset of some level of ground motion at the user site.

Consider the case of User A. User A would like to initiate a set of damage-mitigating actions if the ground motions at the user site exceed some level, a_{thresh}, above which damage occurs. Given the source estimates from an early warning system (estimated magnitude and location, as well as uncertainties on these estimates), User A can calculate the expected/predicted ground motion levels at her site via the appropriate attenuation relationships. The expected/predicted ground motion level is a_{pred}, with some uncertainty (standard deviation) σ_{pred}. (Also assume that User A will not initiate actions unless the uncertainty σ_{pred} is less than some reliability threshold, σ_{thresh}.) σ_{pred} is a function of the uncertainty of the early warning estimates as well as the uncertainty in the attenuation relationships. Bootstrapping methods can be used to estimate the effects of the uncertainties of the early warning estimates on the uncertainties on the predicted ground motions at site A. (Alternatively, the Bayesian predictive method described in Beck and Katafygiotis (1998) can be used, although these are assymptotic results that require a large number of data points.) As the uncertainty on the magnitude and location estimates decrease with additional observations, σ_{pred} approaches the uncertainty of the ground motion attenuation relationships. For the various envelope amplitudes discussed in Chapter 2, this limit is about a factor of 2, or about ± 0.3 log-units. (For convenience, and to be consistent with the typical attenuation relationships, the equations that follow are based on models that describe the *log* of ground motion amplitudes.)

Assuming a given source estimate, the probability density function of the actual peak motions at site A, a, is given by:

$$prob(a|\text{EW source est}) = \frac{1}{\sqrt{2\pi}\sigma_{pred}} \exp\left(-\frac{(a - a_{pred}(\text{EW source est.}))^2}{2\sigma_{pred}^2}\right) \tag{9.1}$$

The probability of exceedance, or the probability of the actual ground motions a exceeding the threshold a_{thresh}, is given by

$$Pr(a > a_{thresh}|a_{pred}) = \int_{a_{thresh}}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_{pred}} \exp\left(-\frac{(a - a_{pred}(\text{EW source est.}))^2}{2\sigma_{pred}^2}\right) da \tag{9.2}$$

Let us use the following definitions: 1) a false alarm corresponds to initiating action when it is ultimately not necessary, and 2) a missed alarm corresponds to not initiating action when it is ultimately necessary. Both refer to making less that optimal decisions under a set of circumstances. They arise because of the uncertainty in the relationship between the actual ground motions (which determine whether damage will occur or not) and the predicted ground motions (which determine whether to initiate action or not). The optimal decisions can always be made *if* the actual ground motions are known. However, this is often not the case. More typical is the situation where the user has a predicted level of ground motion on which a decision must be made. Only time will tell where the actual ground motions fall relative to the predictions, and whether the decision made was the correct one or not.

Even when the predicted ground motion level is lower than the threshold, $a_{pred} < a_{thresh}$, the probability of the actual ground motions exceeding the threshold is non-zero. This is shown in Figure 9.1. Assume that User A decides whether or not to initiate a set of damage-mitigating actions based on a_{pred}. User A initiates actions when $a_{pred} > a_{thresh}$. When $a_{pred} < a_{thresh}$, User A will not initiate actions, and the probability of exceedance is the probability of missed warning. The probability of missed alarm is dependent on the uncertainty on the predicted ground motion, σ_{pred}. Given a predicted ground motion level, the probability of a missed alarm is larger when the uncertainty on the predicted ground motion is larger. This is shown in Figure 9.2.

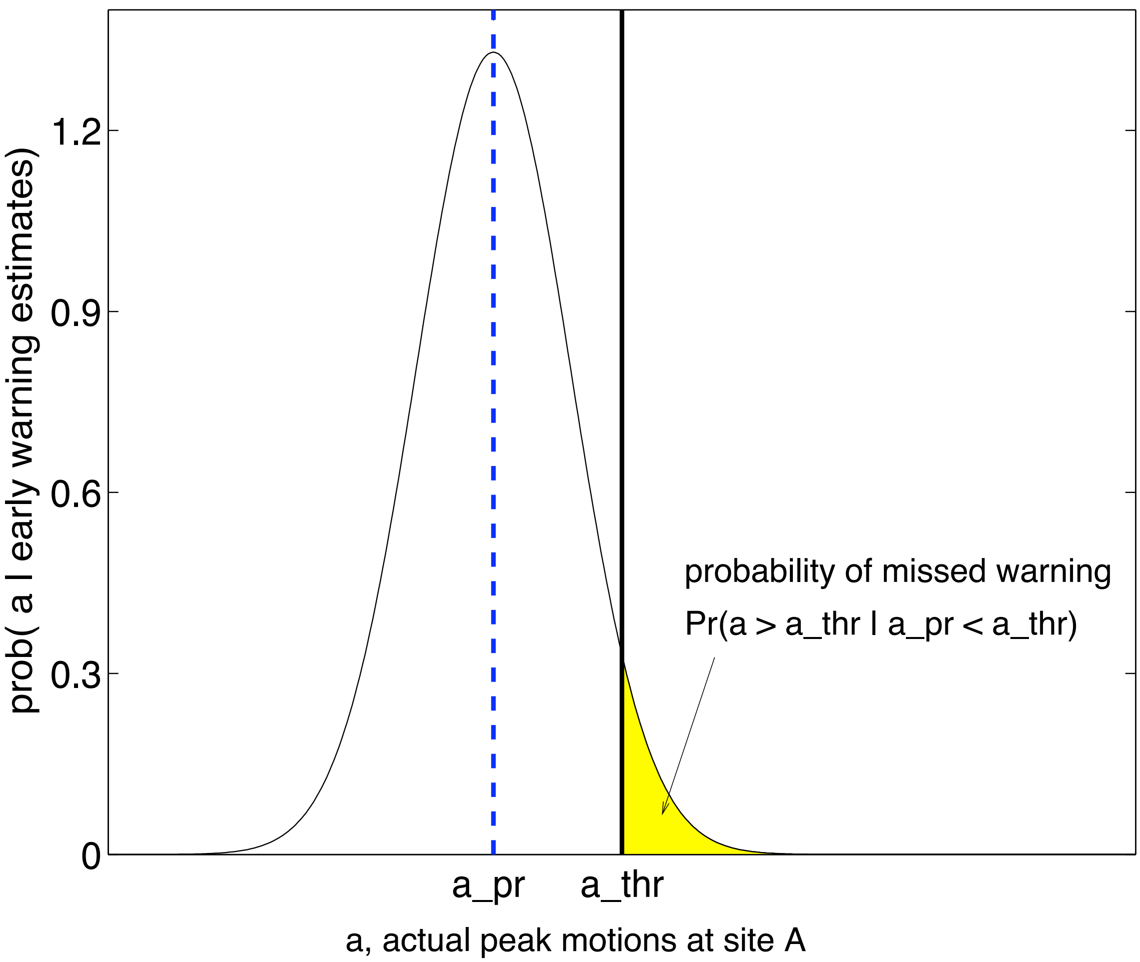

Figure 9.1: When the predicted ground motion based on some given early warning source estimates is less than the threshold, $a_{pred} < a_{thresh}$, the probability of exceedance is the probability of missed alarm.

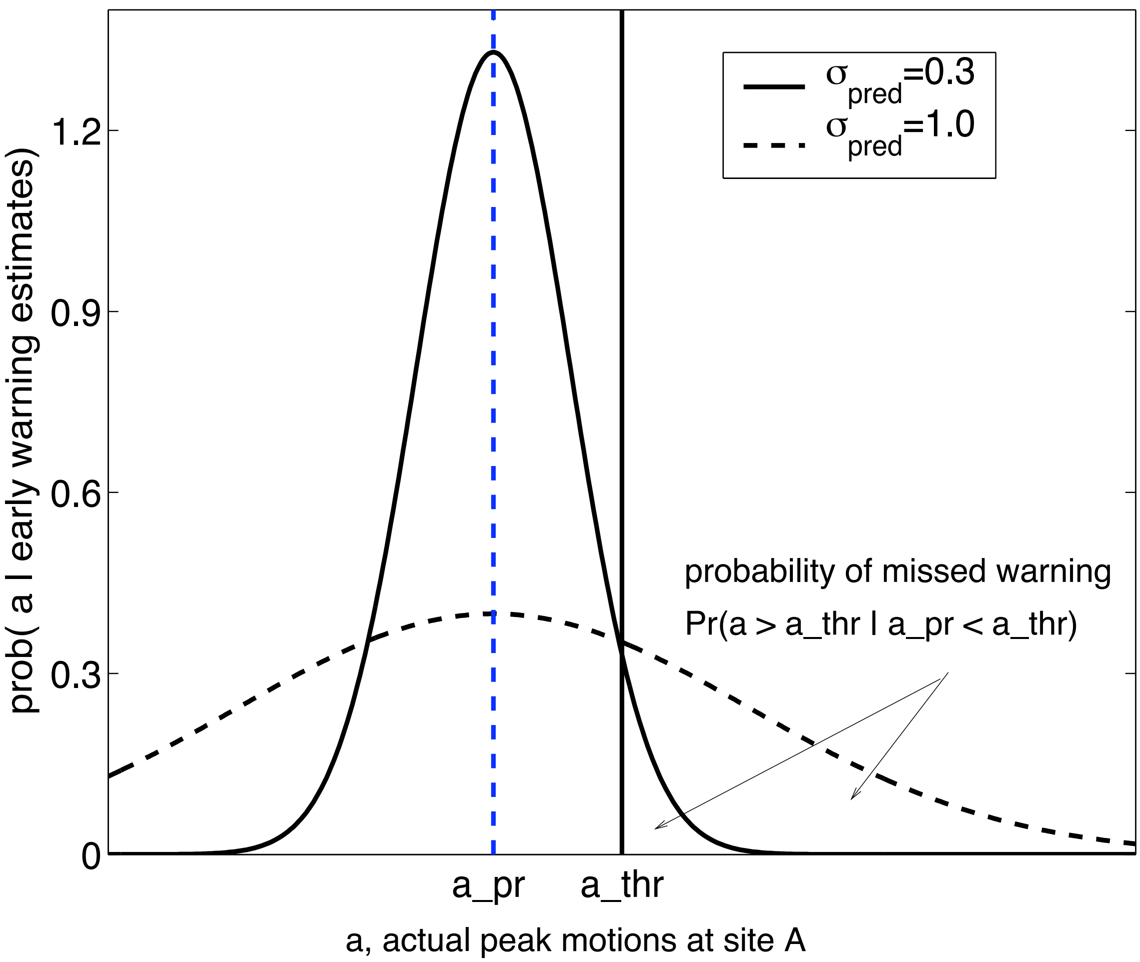

Figure 9.2: The probability of missed alarm is dependent on the uncertainty on the predicted ground motion, σ_{pred}, which is, in turn, dependent on the uncertainties of the early warning estimates and the attenuation relationships.

When the predicted ground motion level is higher than the threshold, $a_{pred} > a_{thresh}$, the probability of a false alarm is given by $1 - Pr(a > a_{thresh}|a_{pred})$. This is illustrated in Figure 9.3. Like the probability of missed alarm, the probability of false alarm for given a predicted ground motion level is proportional to the uncertainty on the predicted ground motions.

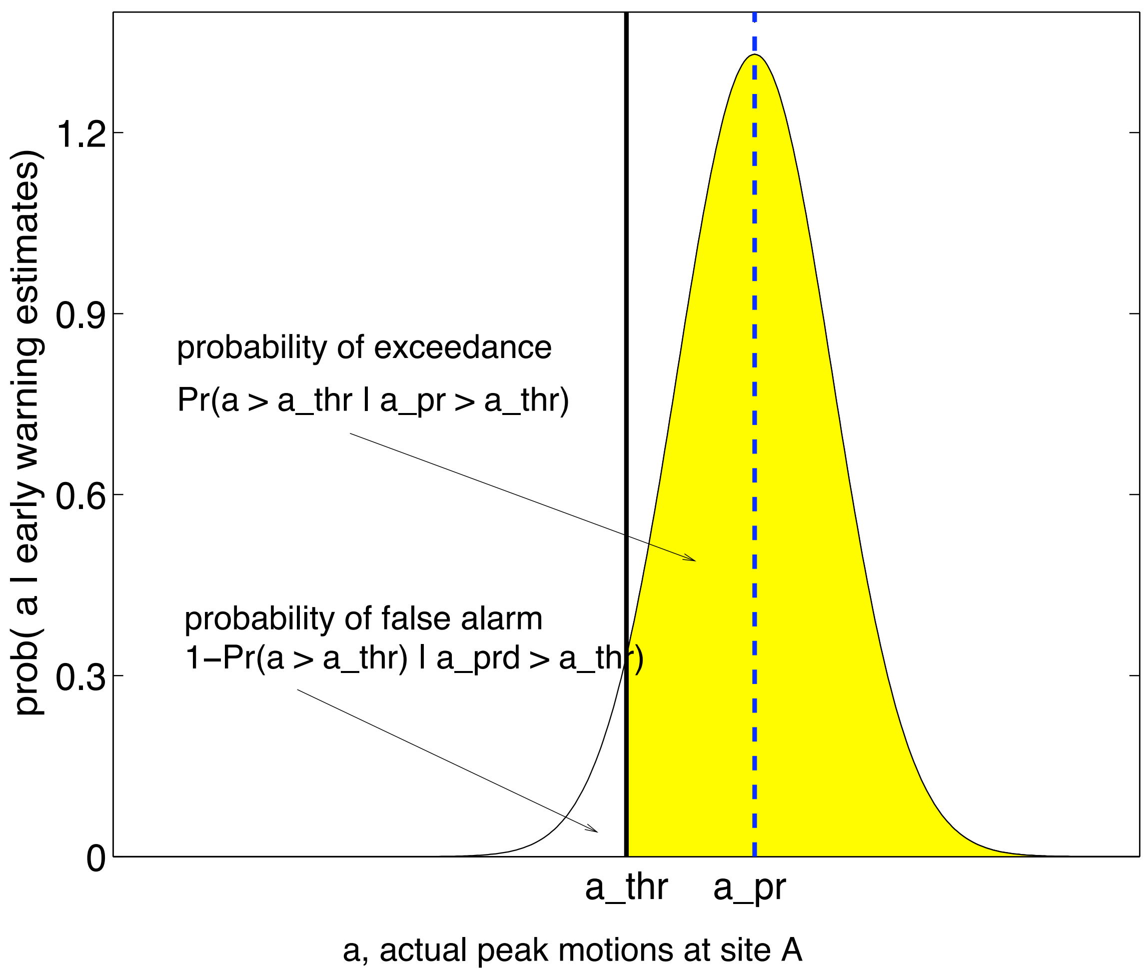

Figure 9.3: When the predicted ground motion based on some given early warning estimate is greater than the threshold level, $a_{pred} > a_{thresh}$, the probability of false alarm is $1 - Pr(a > a_{thresh}|a_{pred})$.

The decision of whether to take action or not should not only depend on the predicted ground motion level; intuitively, the relative costs the potential damage and that of initiating damage-mitigating actions should come into play through the application of decision theory.

9.1 A simple cost-benefit analysis

Assume that User A needs t seconds to initiate a set of damage-mitigating actions, and that User A must decide whether to initiate these actions or not based on the available source estimates (magnitude and location estimates, with uncertainties). A simple cost-benefit analysis for User A is presented using basic decision theory concepts (Grigoriu et al., 1979).

Let $H = h_i, i = 1, ..., n$ be the (exhaustive and mutually exclusive) set of possible states of nature. In this simple example, there are only 2 possible states ($n = 2$): 1) the actual ground motions are larger than the threshold, $a > a_{thresh}$, and 2), the actual ground motions are *not* larger than the threshold $a \leq a_{thresh}$. Let $B = b_j, j = 1, ..., m$ be the set of possible actions. In this simple example, there are only 2 possible actions ($m = 2$): 1) do nothing, and 2) initiate damage-mitigating actions. Let $C(b_j, h_i)$ be the cost of action b_j if the state of nature is h_i. Let p_i be the probability of the state of nature h_i. For brevity, let P_{ex} refer to the probability of exceedance, or the probability of the actual ground motions being larger than the threshold, $a > a_{thresh}$ (Eqn. 9.2). Let C_{damage} be the damage if the set of damage-mitigating actions are not initiated and $a > a_{thresh}$; this is the cost of a missed alarm. Let C_{act} be the cost of performing the sequence of damage-mitigating actions; this is the cost of a false alarm. The costs for the other states of nature and actions are listed in Table 9.1. This is a very simplified analysis in that it is assumed that C_{damage} and C_{act} are known. In practice, these are uncertain; probability models are required to describe these quantities.

h_i	$p_i = Pr(h_i)$	b_1: do nothing	b_2: act
$h_1 : a > a_{thresh}$	P_{ex}	C_{damage}	C_{act}
$h_2 : a < a_{thresh}$	$1 - P_{ex}$	0	C_{act}

Table 9.1: Cost table

The expected cost of a given action is given by:

$$E[C_j] = \sum_{i=1}^{n} C(b_j, h_i) p_i \tag{9.3}$$

The probability of $a > a_{thresh}$ is the probability of exceedance; this is denoted as P_{ex}. The probability of $a < a_{thresh}$ is $1 - P_{ex}$. The optimal action is that which has the minimum cost. From Table 9.1, the optimal action will depend on the relationship between C_{damage} and C_{act}. Let us define $C_{ratio} = \frac{C_{damage}}{C_{act}}$. Given C_{ratio}, without loss of generality, we can let $C_{act} = 1$ and $C_{damage} = C_{ratio}$. The cost table with this substitution is shown in Table 9.2.

h_i	$p_i = Pr(h_i)$	b_1: do nothing	b_2: act
$h_1 : a > a_{thresh}$	P_{ex}	C_{ratio}	1
$h_2 : a < a_{thresh}$	$1 - P_{ex}$	0	1

Table 9.2: Cost table in terms of $C_{ratio} = \frac{C_{damage}}{C_{act}}$

The optimal action is that with the minimum cost. From Table 9.2, assuming that $C_{ratio} > 1$, it is clear what are the optimal actions if the true states of nature are known. If it is known that $a > a_{thresh}$, the optimal action is to initiate the damage-mitigating actions, or "act". If it is known that $a < a_{thresh}$, the optimal action is to "do nothing".

When there is uncertainty regarding the state of nature, the expected costs of actions are a function of the probability of a given state of nature. In this example, these probabilities are in terms of the probability of exceedance, P_{ex}. The expected cost of "do nothing" and "initiating actions", in terms of C_{ratio} and P_{ex}, are then given by:

$$E[\text{"do nothing"}] = P_{ex} C_{ratio} \tag{9.4}$$

$$E[\text{"act"}] = C_{act} \tag{9.5}$$

$$\tag{9.6}$$

We can solve for a critical probability of exceedance, P_{crit}, which is the value of P_{ex} where the costs of "act" and "do nothing" are equal. If $P_{ex} < P_{crit}$, to "do nothing" is the optimal action; if $P_{ex} > P_{crit}$, it is best to initiate the damage-mitigating actions. Setting Eqns 9.4 and 9.5 and solving for $P_{ex} = P_{crit}$, we find:

$$P_{crit} = 1/C_{ratio} \tag{9.7}$$

Thus, the probability of exceedance (which is a function of the predicted ground motion level and its uncertainty, a_{pred}, σ_{pred}, as described in Eqn. 9.2) beyond which it is better to "act" rather than "do nothing" depends on C_{ratio}. Since P_{crit} is a probability, it must be between 0 and 1, which means that C_{ratio} must be greater than or equal to 1. Recall that $C_{ratio} = \frac{C_{damage}}{C_{act}}$. $C_{ratio} \geq 1$ means that the cost of damage as a consequence of not acting must be equal to or greater than the cost of performing the actions. C_{ratio} will vary from user to user. There is no point to using seismic early warning information in applications where $C_{ratio} < 1$. (This can be seen from Table 9.2. If $C_{ratio} < 1$, "do nothing" will always have the minimum cost, and would always be the optimal action.)

This critical probability of exceedance, P_{crit}, can be related to a critical value for predicted ground motion $a_{pred,crit}$, beyond which it is optimal to act, by:

$$P_{crit} = 1/C_{ratio} = \int_{a_{thresh}}^{\infty} \frac{1}{\sqrt{2\pi}\sigma_{pred}} \exp\left(-\frac{(a - a_{pred,crit}(\text{EW source est.}))^2}{2\sigma_{pred}^2}\right) da \tag{9.8}$$

The predicted ground motion level, $a_{pred,crit}$, associated with the critical probability of exceedance $P_{crit} = 1/C_{ratio}$, can be solved for numerically, or can be expressed in terms of the error function erf, which is a commonly tabulated function (Sivia, 1996).

$$
\begin{aligned}
erf(z) &= \frac{2}{\sqrt{\pi}} \int_0^z \exp\left(-t^2\right) dt && (9.9) \\
\int_b^d \exp\left(-\frac{x^2}{2\sigma^2}\right) dx &= \sigma\sqrt{\frac{\pi}{2}} \left[erf\left(\frac{d}{\sigma\sqrt{2}}\right) - erf\left(\frac{b}{\sigma\sqrt{2}}\right) \right] && (9.10) \\
\text{let } x = a - a_{pred} \quad & \text{and} \quad b = a_{thresh} - a_{pred}, \qquad d = \infty \\
P_{crit} = 1/C_{ratio} &= \frac{1}{\sqrt{2\pi}\sigma_{pred}} \left[1 - erf\left(\frac{a_{thresh} - a_{pred,crit}}{\sigma_{pred}\sqrt{2}}\right) \right] && (9.11) \\
a_{pred,crit} &= a_{thresh} - \sigma_{pred}\sqrt{2} \left[erf^{-1}\left(1 - \frac{\sqrt{2\pi}\sigma_{pred}}{C_{ratio}}\right) \right] && (9.12)
\end{aligned}
$$

Given a user-specific $C_{ratio} = C_{damage}/C_{act}$, User A should initiate actions if the source estimates (magnitude and location, with associated uncertainties) result in a predicted ground motion level $a_{pred} \geq a_{pred,crit}$. Equivalently, the criteria can be expressed in terms of $P_{crit} = 1/C_{ratio}$: User A should initiate actions if the probability of exceedance given the early warning source estimates and their uncertainties, P_{ex}, exceeds $P_{crit} = 1/C_{ratio}$, or $P_{ex} \geq P_{crit} = 1/C_{ratio}$. It is clear that the decision of whether to act is related to $C_{ratio} = C_{damage}/C_{act}$, the relative cost of missed alarms to false alarms.

Figure 9.4 shows $a_{pred,crit}$ as a function of σ_{pred} for various values of C_{ratio}. Figures 9.5 and 9.6 show the optimal actions as a function of P_{ex} for different values of C_{ratio}. The probability of exceedance, P_{ex}, is a function of the ground motion threshold for the particular application, a_{thresh}, the predicted level of ground motion, a_{pred}, and the uncertainty on this prediction, σ_{pred}. These predicted quantities in turn depend on the early warning source estimates and their associated uncertainties. When $P_{ex} > 0.5$, the predicted ground motion exceeds the threshold, $a_{pred} > a_{thresh}$. When $P_{ex} < 0.5$, the predicted ground motion is lower than the threshold, $a_{pred} < a_{thresh}$. Depending on C_{ratio}, it is sometimes optimal to "do nothing" even when the predicted ground motions exceed the threshold, and sometimes optimal to "act" even when the predicted ground motions are below the threshold level. This can be seen from Eqn. 9.12: $a_{pred,crit}$, the predicted ground motion level above which User A should initiate action, can be greater or less than a_{thresh}, the ground motion level at which damage occurs, depending on the value of $C_{ratio} = C_{damage}/C_{act}$.

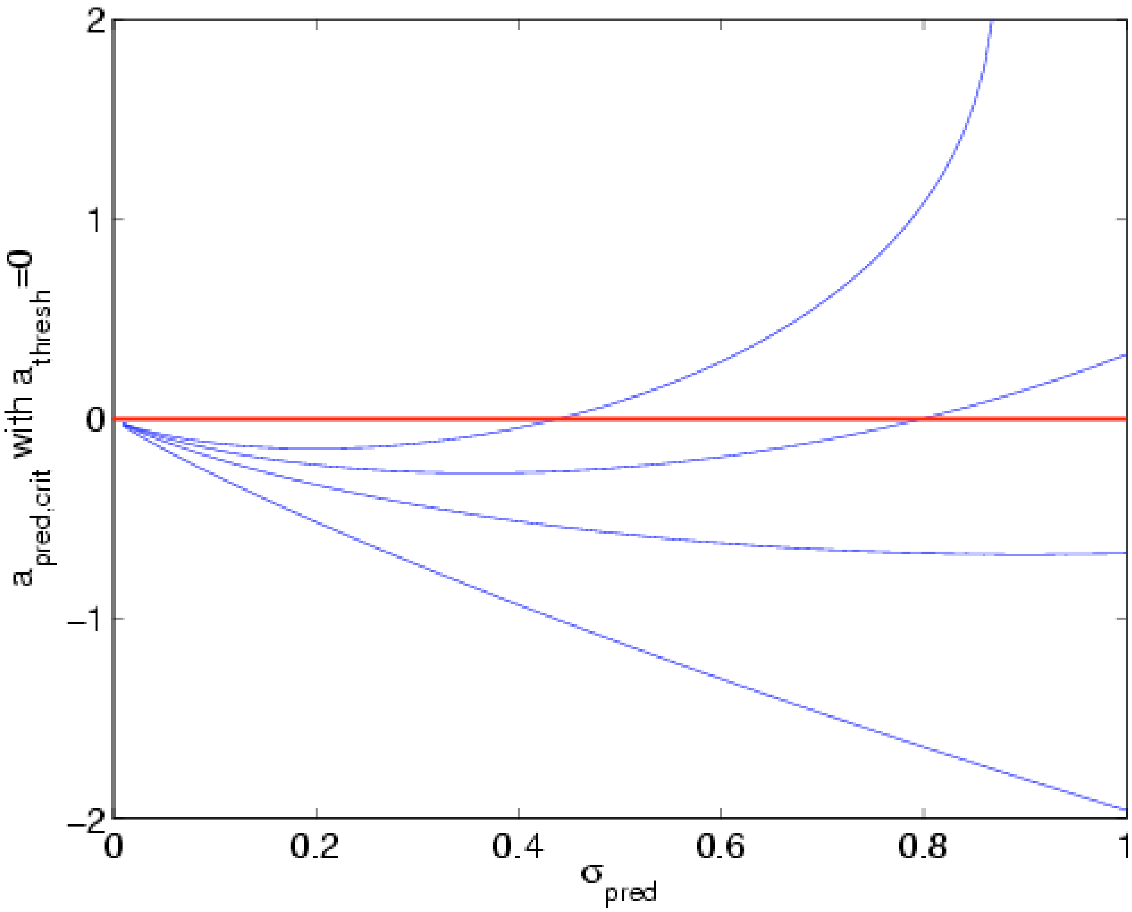

Figure 9.4: $a_{pred,crit}$, the predicted ground motion level above which the user should initiate action, as a function of the uncertainty on the predicted ground motions, σ_{pred}, and the relative costs of false and missed alarms, C_{ratio}.

In the cost-benefit analysis presented, there are just 2 possible actions: 1) "do nothing", and 2) "act". This describes the situation where the time to the estimated onset of the peak predicted motions is equal to the time necessary to perform the damage-mitigating actions. Waiting for an updated estimate with presumably lower uncertainties is not an option.

When the time to the estimated onset of the peak ground motions is greater than the time necessary to perform the damage-mitigating actions, waiting is potentially

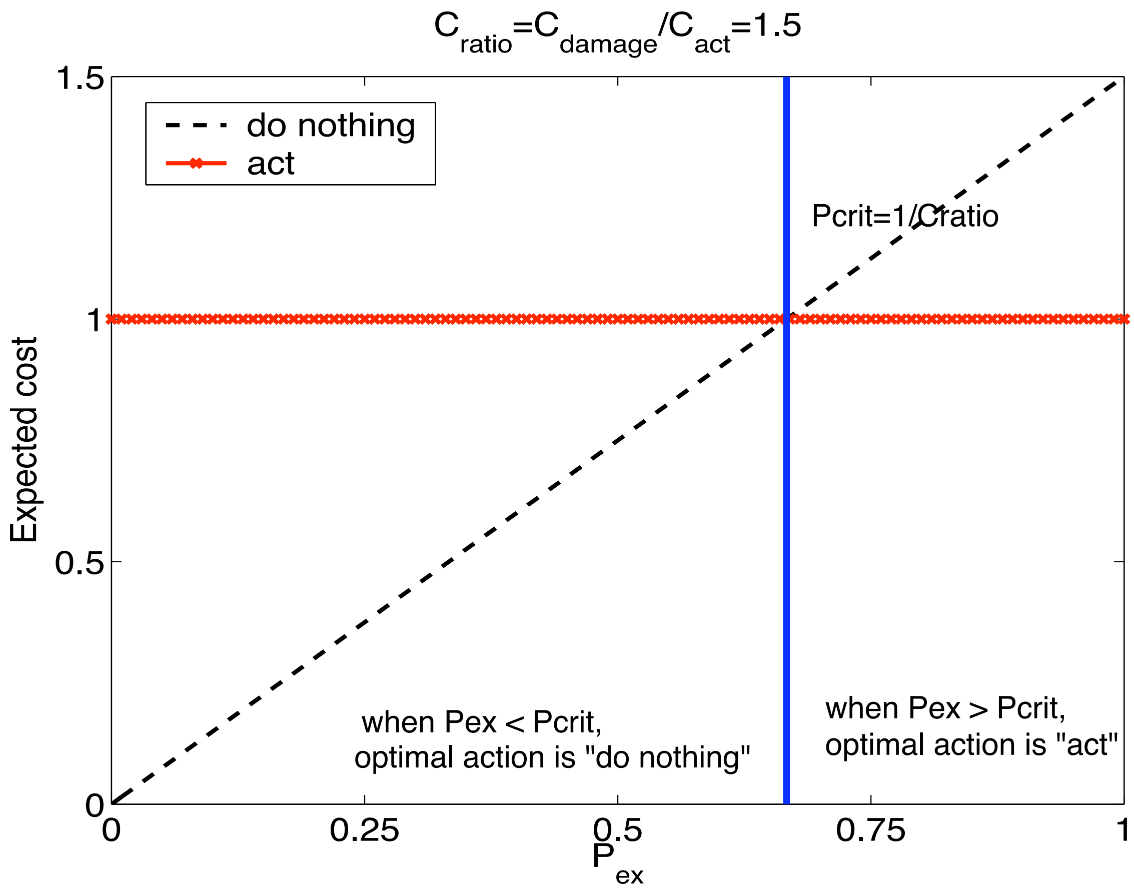

Figure 9.5: Optimal actions as a function of the probability of exceedance, P_{ex}, for $C_{ratio} = 1.5$. Recall that $P_{crit} = 1/C_{ratio}$. When $P_{ex} < P_{crit}$, the optimal action is to "do nothing". When $P_{ex} > P_{crit}$, the optimal action is to perform the sequence of damage-mitigating actions, or "act". Note that $P_{ex} = 0.5$ corresponds to the situation when the predicted value is equal to the threshold ground motions for the application. For $C_{ratio} = 1.5$, there are situations where even if the predicted ground motion exceeds the threshold, it is still optimal to "do nothing". The costs of false alarms is relatively high.

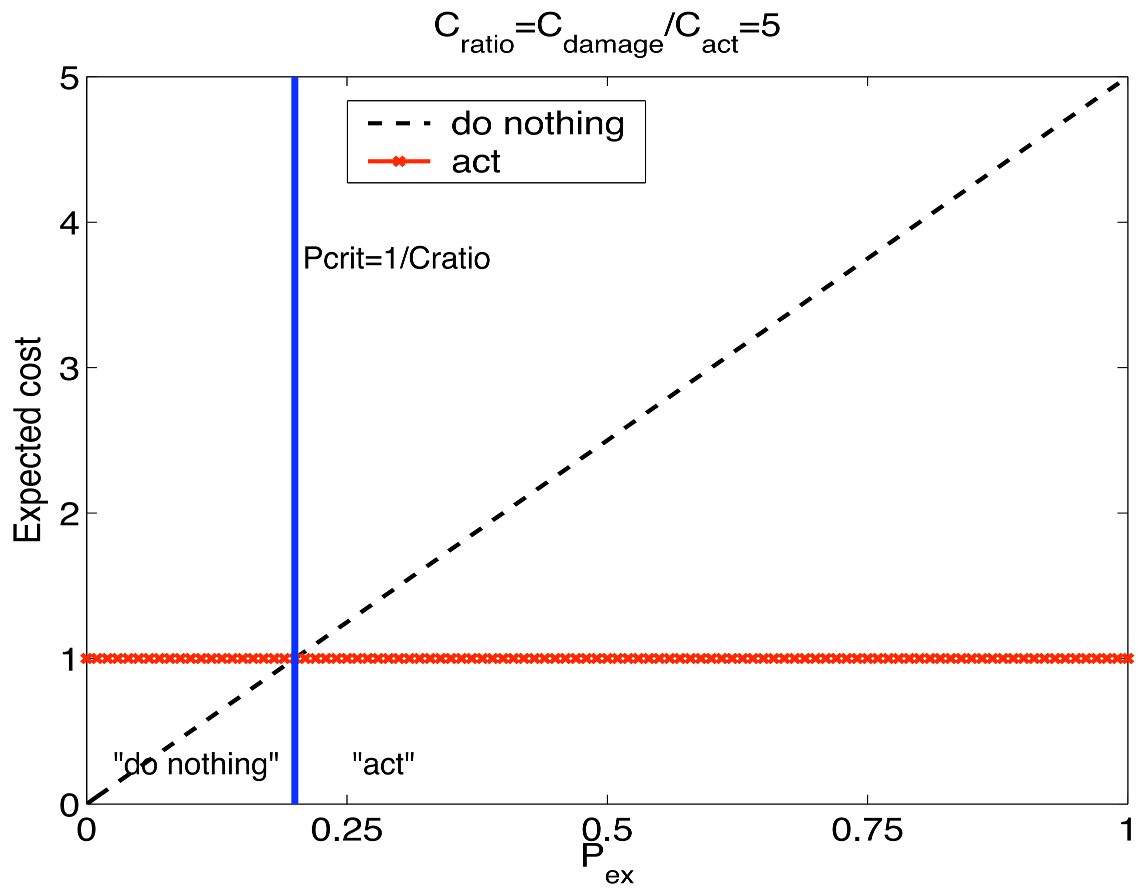

Figure 9.6: For $C_{ratio} = 5$, $P_{crit} = 0.2$. When $P_{ex} < P_{crit}$, it is optimal to "do nothing". When $P_{ex} > P_{crit}$, it is optimal to initiate the damage-mitigating actions, or "act". $P_{ex} = 0.5$ corresponds to the situation where the predicted ground motion level is equal to the user-defined threshold. For $C_{ratio} = 5$, there are situations where even if the predicted ground motion is less than the threshold, it is optimal to "act". The consequences of not acting when the actual ground motions exceed the threshold are getting larger relative to the cost of false alarms.

an option. However, there is no benefit (or reduction in cost) due to initiating actions sooner rather than later. Earlier estimates typically have larger uncertainties than the subsequent updates. Larger uncertainties translate to larger P_{ex}. For a given predicted level of ground motion, the earlier estimates are more likely to optimize "act", since they will typically have larger uncertainties. (Recall that to initiate the damage-mitigating actions is optimal when $P_{ex} > P_{crit}$.) On the other hand, if it is possible to wait, in the situation where the source is estimated to be far enough such that the time to the arrival of the peak ground motions is larger than the time necessary to "act", then it is always optimal to wait for an updated estimate. Later estimates typically have smaller uncertainties than earlier ones, which means that there is potential for P_{ex} to be lower than P_{crit}. Therefore, it is best to wait for updated estimates until the estimated onset time of the peak motions is equal to the time necessary to perform the set of actions. At the time of this critical decision, the only possibilities are those we have considered: "do nothing" or "act".

9.2 Using VS estimates with and without the Gutenberg-Richter relationship

Consider the situation of User A, who has some specified value of C_{ratio}. Recall that C_{ratio} is the relative cost of missed warnings to false alarms. At the time of the critical decision, the VS estimates *with* the Gutenberg-Richter built in predicts a ground motion level has a probability of exceedance P_{ex}^{G-R}. The VS estimates *without* the Gutenberg-Richter gives a predicted ground motion level with probability of exceedance $P_{ex}^{\text{no G-R}}$. Assume that User A requires that the reliability of the early warning estimates satisfies the criteria: $\sigma_{pred} < \sigma_{thresh}$. Assume also that the uncertainty is primarily on the magnitude estimate. Three scenarios are possible:

- $P_{ex}^{G-R} > P_{crit}$ and $P_{ex}^{\text{no G-R}} > P_{crit}$. The VS magnitude estimates with G-R are always lower than the VS magnitude estimates without G-R and almost always lower than the actual magnitude. The optimal action is to "act". This is the only unambiguous scenario.

- $P_{ex}^{G-R} < P_{crit}$ and $P_{ex}^{\text{no G-R}} > P_{crit}$. Should User A base its actions on the VS estimate with G-R, which optimizes "do nothing", or initiate the damage-mitigating actions based on the VS estimate without G-R? To address this properly, one needs to take into account how the VS estimates typically evolve with time. VS magnitude estimates with the G-R are always lower than the actual magnitude and hence always increase with time. On the other hand, in 1 out of 4 times (Parkfield), the VS magnitude estimate without the G-R was larger than the actual magnitude. Is how much the VS magnitude estimate with G-R typically increases enough to make $P_{ex}^{G-R} > P_{crit}$?

- $P_{ex}^{G-R} < P_{crit}$ and $P_{ex}^{\text{no G-R}} < P_{crit}$. Both VS estimates (with and without the Gutenberg-Richter) optimize "do nothing". However, in 4 sample events considered, the VS magnitude estimate with G-R is smaller than the actual magnitude. Again, the optimal action will depend on how likely it is that the VS magnitude estimates will increase such that P_{crit} will be exceeded.

To resolve the ambiguous situations (2 out of 3 possible scenarios) requires some statistics on how the VS estimates evolve with time. In particular, robust statistics (based on more than 4 events) are necessary to quantify 1) how much the VS magnitude estimates with the G-R are likely to increase with time and 2) how likely the VS magnitude estimates without the G-R are to be larger than the actual magnitude. These can be obtained by running the VS method for early warning on many earthquakes.

It should be mentioned that these ambiguities would only arise if σ_{thresh} is large enough such that there are differences in the VS estimates with and without the G-R. From the 4 sample events, as additional data becomes available, these two estimates eventually converge. However, if the reliability threshold is too high, (or equivalently, the required σ_{thresh} is too low), little or no warning time would be available.

Typically, the initial VS magnitude estimates (with and without the Gutenberg-Richter magnitude-frequency relationship) are lower than the actual magnitude. The envelope attenuation and ground motion ratio relationships, which ultimately define the mapping from observed amplitudes to early warning estimates, are based on peak P and S wave amplitudes. For large events, the peak P-wave amplitudes are usually not available at the time of the initial 3 second estimates. The M=6.0 Parkfield event is an exception; the initial estimates are dominated by station PKD, which is in the forward directivity direction of the rupture. The VS magnitude estimate without G-R is greater than 6.0; the VS magnitude estimate with G-R is less than 6.0. For smaller events, for instance, the M=4.75 Yorba Linda event, the peak amplitudes (at least on the vertical components) are available at the time of the initial estimates, and the VS estimates are very close to the actual magnitude.

At the time of this writing, most research in seismic early warning has focused on the source estimation problem. It is only recently that attention has been brought to the questions of how users might response to early warning information. Aside from the work described in this thesis, Grasso et al. (2005) have begun to address how early warning information might be used for structural control applications.

9.3 Some user requirements

For a potential user to even consider subscribing to an early warning system, a few requirements must be met. First of all, the user must determine its C_{ratio}, which can be seen as the relative cost of missed alarms (C_{damage}) to false alarms (C_{act}). C_{ratio} must be greater than 1 for a user to consider seismic early warning. C_{ratio} less than 1 means that it will always be optimal to "do nothing", since the relative cost of false alarms is prohibitively high. Next, the user must also determine how much warning time is necessary to initiate the desired damage-mitigating actions. Obviously, the warning time must be on the order of seconds. The necessary warning time dictates the level of uncertainty on the early warning estimates (and hence the predicted ground motions) that the user must accept. Actions which require longer warning times must tolerate higher uncertainties on the early warning estimates. For actions requiring relatively short warning times, users can afford to wait for updated estimates that will usually have better reliabilities. In situations when waiting is an option, it is always better to wait for an updated estimate. With regards to the user reliability threshold σ_{thresh}, it must be kept in mind that even if the magnitude and location estimates are very precise, the lowest possible uncertainty on the predicted ground motions is determined by the uncertainty of the attenuation relationships. For the envelope attenuation relationships, this uncertainty is about a factor of 2 (± 0.3 in log-units) without station corrections, and about a factor of 1.7 (± 0.24 in log-units) with station corrections. If this uncertainty is too high for the user, one possibility is to install a station at the user site and develop attenuation relationships tailored to the ground motions recorded at the user site. These user-specific attenuation relationships could take into account other predictors, such as faulting type, azimuth, and source-to-station path; these additional predictors may decrease the uncertainty on the attenuation relationships.

9.4 Some comments

If the user wants to ensure that the appropriate actions are initiated for "the Big One", the user must be prepared to accept a certain number of false alarms. Such users would have $C_{ratio} = C_{damage}/C_{act} \gg 1$. Stopping elevators at the closest floor would be an example of an application with $C_{ratio} \gg 1$. On the other hand, perhaps shutting down a nuclear power plant would be an application with C_{ratio} closer to 1.

In Chapter 1, it had been mentioned that there was some question about whether the *most probable* source estimates given the available ground motions were the most appropriate information for an early warning system to broadcast to its users. After much deliberation (between myself and my advisor), we concluded that the Gutenberg-Richter has to be included at some point - either in the source estimation or the user decision-making process. If the VS method for seismic early warning were to be implemented, it would perhaps be useful to provide two source estimates - with and without the G-R. Depending on their particular risks, users would probably weigh one type of estimate more than the other. For instance, a user with a high

C_{ratio} who needs relatively long warning times will need to act on earlier estimates. Such a user should use the VS estimates without the G-R, as these are more often closer to the actual magnitude than those with the G-R built in. Acting on these estimates (without the G-R) exposes the user to more false alarms. However, the number of false alarms to expect or to be tolerated by the user is ultimately set by C_{ratio}, the relative cost of missed alarms to false alarms. If the goal of the user is to act appropriately for the large events, they have to accept the false alarms. On the other hand, a user with a relatively low C_{ratio} who needs short warning times should include the Gutenberg-Richter relationship in it's decision process. A low C_{ratio} means that false alarms are relatively expensive. Using the Gutenberg-Richter removes the risk of false alarms, at the expense of introducing the possibility of missed warnings. The probability of missed alarms based on estimates at a given time in the estimation process can be quantified by taking into account how far off the estimate is from the true magnitude, and what decrease in the uncertainties can be expected with time. This requires robust statistics (based on more than 4 events) regarding how the VS estimates evolve with time.

It could be argued that it is confusing to have two potentially different source estimates at a given time. This confusion can be reduced if the user understands how the source estimation process works and how to optimally use these two different source estimates. This again points to the need to move away from the perception of the source estimation and user response problem as two distinct problems, and towards a more integrated paradigm in which the source estimation process and the user are part of a single (albeit very complex) system.