Communicating Effectively about Risk Magnitudes: Bottom Line Conclusions and Recommendations for Practitioners

by Peter M. Sandman and Neil D. Weinstein[1]

1994

This is a brief summary of key conclusions and recommendations from three sequential pieces of research. For a longer summary of the Phase Three research, see the Executive Summary of Neil D. Weinstein, Peter M. Sandman, and William K. Hallman, *Communications to Reduce Risk Underestimation and Overestimation* (January 1994). For a longer summary of the Phase Two research, see the Executive Summary of Neil D. Weinstein, Peter M. Sandman, and Paul Miller, *Communicating Effectively about Risk Magnitudes, Phase Two* (September 1991). For a longer summary of the Phase One research, see the Executive Summary of Neil D. Wein-

---

[1]Peter M. Sandman was director of the Environmental Communication Research program at Rutgers University, and is now a risk communication consultant at 54 Gray Cliff Road, Newton Center, MA 02159 (617/630-0385). Neil D. Weinstein is Professor of Human Ecology, Cook Office Building, P.O. Box 231, Cook College, Rutgers University, New Brunswick, NJ 08903 (908/932-9169).

stein, Peter M. Sandman, and Nancy E. Roberts, *Communicating Effectively about Risk Magnitudes* (September 1989).

All three reports are available from the Risk Communication Project, Office of Policy, Planning and Evaluation, U.S. Environmental Protection Agency, or from the Center for Environmental Communication, Rutgers University.

Probably the most important thing to say about the conclusions and recommendations that will follow is that they are preliminary. The effects we have found are often small, taking a careful study with a big sample to find. They are based on people's reactions to hypothetical exposure data; we do not know if people respond similarly to real exposure data. They are based on studies of several different risks (radon, asbestos, etc.), but just a single health outcome, lung cancer. And they are based on just one or two studies; experienced social scientists know not to rely too heavily on a finding until it has turned up in several different studies using several different methodologies. In addition, participants in this research were much better educated than the general population. Finally, the research design we employed confronted people with personal choices about an individually remediable pollutant in their homes; a public risk controversy requiring large-scale mitigation by government agencies might have generated very different responses.

Of course, practitioners can rarely afford to wait for definitive research results. Since you have a job to do, a risk to describe, you are better off following the advice below than ignoring it. But see it as tentative.

It is worth emphasizing that this research effort focused on ways of explaining risk magnitudes more effectively — that is, ways to help people understand the size of their risk. A more controversial class of risk communication strategies attempt to influence risk responses by manipulating emotions or behavior rather than through improved understanding (examples include dramatic fear appeals, social pressure,

rewards for compliance, etc.). These non-cognitive approaches can be very effective — but many scientists object to them.

Seven Factors that Affect Risk Response

1. Data about the Actual Risk. With everything else held constant, subjects viewed the risk as more serious when the data they were given told them the actual probability of experiencing harmful effects was greater. (Henceforth in this summary we will use the phrases "actual risk" and "actual probability of experiencing harmful effects interchangeably.)

   In the Phase Two research, a ten-times-higher risk from geological radon affected risk perceptions but not mitigation intentions; a 24-times-higher risk from radon affected both. (The data provided included risk probability information plus comparisons to smoking.) In the Phase Three research, people facing a 40-in-1,000 radon risk expressed higher threat perceptions and action intentions than people facing a 1-in-100,000 risk. In a more demanding test, people facing a 40-in-1,000 risk still expressed lower threat perceptions and action intentions than those facing a 400-in-1,000 risk.

   This is an encouraging bottom-line conclusion: Telling people the size of the risk they face does help encourage an appropriate response. But large differences in actual risk yielded modest differences in perceived risk and action intentions. Moreover, the Phase One research showed that the effect of data about the risk can easily be swamped by other factors, such as an action standard or a risk ladder. And the Phase Three research showed that the risk probability effect also tends to disappear when the actual risk is low and "outrage" is high: Responses to nuclear power plant waste used in the construction of home basements were no weaker the risk was 1-in-1,000,000 than when it was 1-in-100,000.

2. An Action Standard. The Phase One research found that formats that included an action standard were superior to formats without an action standard in helping people respond in proportion to the actual risk from radon or asbestos. That is, the relationship between actual risk and perceived risk, and between actual risk and mitigation intentions, was stronger with an action standard than without. The effect of providing an action standard, in fact, was stronger than the effect of providing risk probability data.

People seemed to use the standard as an "anchor" to help them interpret their own levels. The standard was especially powerful in helping people distinguish levels above the standard from those below the standard — so powerful that it sometimes created an artificial discontinuity in risk perceptions at the standard. It also helped people distinguish levels just below the standard from those far below the standard. On the other hand, an action standard did not help people make the distinction between levels just above the standard and those far above the standard. Three ranges seemed to be psychologically meaningful when a standard was provided: "okay" (way below the standard), "possible trouble" (just below) and "deep trouble" (anywhere above the standard).

Practitioners should always provide a standard when one exists. Even though the standard may tend to distort people's understanding of the risk (way below versus just below versus above), this trichotomy is still closer to the actual risk than people's responses would be without the aid of the standard. It may help to qualify the standard, where appropriate, with warnings that risks just below the standard are nearly as risky as those just above, and that risks far above the standard are much more risky than those just above.

(An action standard without additional risk information is useful when an apathetic response is anticipated and the goal is to provoke more risk aversion. See Number Four below.)

3. Advice. People said they felt less uncertainty and had a better understanding of their risk when advice was provided. More importantly, although people in the Phase One study often said they would choose to mitigate at levels below the standard, those receiving action advice showed this tendency least. That is, adding advice to the standard made people less likely to "overreact" vis-à-vis the standard, more likely to accept the recommendation not to take action at low levels. Advice was not similarly useful at high levels; it did not increase the probability that those above the recommended action level would plan to act (most already said they planned to act).

Providing explicit advice is thus especially useful for panic prevention, to deter overreaction at low risk levels. Its value for increasing remedial action at high levels (beyond what would be expected with a standard alone) was not demonstrated.

4. A risk ladder. People felt more at risk when presented simply with a suggested "action level" at which mitigation is recommended than when presented with such a standard located midway up a risk "ladder." In the Phase One research, the ladder included mortality data and risk comparisons; in Phase Two it did not. In both studies, the context that the ladder provided — and the implication that levels higher than one's own are not rare — appeared to reassure subjects and reduce their perception of risk. In Phase Two, the presence or absence of a risk ladder, even without any additional information, had an effect on perceived risk equal to a several-fold difference in actual risk.

Like a standard, in other words, a risk ladder acts as a sort of "anchor." Assuming the ladder goes higher than the individual's risk level, the ladder communicates that things could be worse. If the communicator's goal is maximum risk aversion — that is, if the hazard is serious and the audience is inclined toward apathy — a standard without additional information is ideal; its very ambiguity generates the desired risk-averse response. If panic is a

problem and the goal is to provide reassuring context, on the other hand, a risk ladder is worth adding.

5. Location on the Risk Ladder. By modifying the risk ladder, the Phase Two research was able to locate the same hypothetical hazard exposure with the same risk information either one-quarter of the way up the ladder or three-quarters of the way up the ladder. The difference in location significantly affected perceived risk in two experiments, and mitigation intentions in one. This locational effect was roughly equivalent in size to the effect of a 10× difference in actual risk.

In Phase Three, a risk chart was created that combined the effects of the ladder (with location one-quarter or three-quarters of the way up, as appropriate) with a recommended action standard and risk comparisons. Compared to the no-chart conditions, the chart increased perceived threat when the actual risk was high and decreased perceived threat when the actual risk was low. The effect of the chart on risk perception was as great as or greater than the effect of a 10× difference in actual risk. The effect of the chart on action intentions was smaller; it was still significant when the actual risk was low, but not quite significant for the high-risk scenarios.

Risk information developed to guide laypeople is often arrayed on a risk ladder, and the structure of the ladder may be determined more or less arbitrarily. How low should the ladder begin? How high should it rise? Should the scale be linear or logarithmic? The answers to these questions are not obvious. What is clear from the data is that people's risk perceptions can be substantially altered — whether intentionally or inadvertently — by constructing the ladder so that their risk appears low or high on the page.

For helping people distinguish between high and low levels of a particular risk, X, the most effective ladder would be truncated at both ends, so that high levels of X appeared at the top of the ladder and low levels of X at the

bottom. For helping people see that risk X is actually less serious than risk Y, on the other hand, the ideal ladder would be extended upward, so that all levels of X clustered near the bottom of the ladder, with Y near the top. The best ladder to help people see that X is actually more serious than Z would be extended downward, clustering all the levels of X near the top, with Z near the bottom. A "universal ladder" incorporating all three risks would extend both upward and downward, and would cluster all the levels of X near the middle, with Y near the top and Z near the bottom. These three extended ladders would all be improvements on the original truncated ladder in encouraging an appropriate response to between-hazard risk differences (X versus Y versus Z) — but they would all be worse than the original in encouraging people to discriminate within-hazard risk differences (high versus low levels of X). Thus, it may be impossible to construct a risk ladder that makes optimal use of the locational effect for all risk levels included on the ladder.

6. Comparisons to Normal Background. The most surprising finding of the Phase Three research was the powerful impact of comparisons to normal background levels on threat perceptions and action intentions. Despite having no information at all about the likelihood of harmful consequences (that is, about risk), people responded strongly to the information that their radiation exposure was either 20× higher than normal background radiation or 200× lower than normal background. In fact, the comparison to normal background did a better job than risk information itself in helping people respond in proportion to the actual risk. The effect of the comparison to normal was equal to or greater than the effect of a 10-fold difference in actual risk.

Comparisons to normal were especially powerful when risk was small and the outrage substantial — precisely the situation when risk information itself was least powerful. Even more impressively, comparisons to normal seemed to affect the outrage itself. Subjects in the high-outrage "nuclear waste"

conditions were understandably angry. Their anger was reduced far more by the knowledge that the situation posed a risk 200× less than normal background than by the knowledge that the risk posed was a mere 1-in-100,000 or even a mere 1-in-1,000,000.

In effect, people may view the normal background exposure as the maximum safe exposure, regardless of the level of risk it presents. This exposure can then serve as an "anchor" for their risk judgments, with higher levels seen as "unsafe" or "unacceptable" and lower levels as "safe" or "acceptable." The potency of the comparison to normal background and its symmetry (that is, its effectiveness in both low-outrage, high-risk situations and high-outrage, low-risk situations) suggest that it may be a valuable piece of information to include in a risk communication.

Three qualifiers are needed on this advice, however. First, information on normal background levels is often not available. Second, the research focused on situations where background was a small fraction or a sizeable multiple of the risk under discussion; it isn't clear how people might respond to a non-natural risk about equal to natural background. Finally, it is important to note that comparisons to normal background can be misleading. For some hazards, normal background levels are sufficient to constitute a meaningful health risk, and even a small increment would be unwise if it were preventable. For other hazards, the risk due to normal background exposure is negligible, and an exposure many times background would still be negligible.

7. Outrage. In the Phase Three research, outrage substantially affected threat perceptions and action intentions. That is, subjects in the high-outrage, low-risk situation reported much higher perceived threat and higher action intentions than subjects in the low-outrage, low-risk situation, although the actual risk was identical.

When subjects received only risk numbers, the outrage effect was just as large as the 4,000-fold (!) difference in risk between the high-risk and low-risk conditions. In other words, people had about the same threat perceptions and action intentions whether they faced a low-outrage home radiation risk of 40-in-1,000 (from geological radon) or a high-outrage home radiation risk of 1-in-100,000 (from nuclear waste).

When communication was improved by comparisons to normal background levels or by the risk chart, however, the outrage effect, though still substantial, was smaller than the 4,000-fold difference in risk. We are encouraged by this apparent ability of some kinds of risk information to reduce threat perceptions and action intentions even in the presence of high outrage. Many practitioners have suggested that when people are outraged, explanations of the risk data are unlikely to prove fruitful. In the Phase Three research, outrage certainly increased threat perceptions and action intentions — but outrage did not diminish the ability of comparisons to background and risk charts to reduce threat perceptions and action intentions.

Five Factors that May Not Significantly Affect Risk Response

1. Risk Comparisons. In the Phase One research, the addition of risk comparisons to cigarette smoking had two effects: it made people feel the brochure was more helpful and they understood their risk better, and on some measures it made them less risk-averse (for example, the comparisons raised the highest level people would find acceptable). However, the comparisons had no effect on people's ability to distinguish high risks from low risks: no effect on the accuracy of illness probability estimates or on the relationship between actual risk and perceived risk or mitigation intentions. In other words, in most respects the impact of risk probability data was not improved by the inclusion of comparisons to smoking risks.

Risk comparisons may of course prove more helpful in ways not examined in this research — different comparisons, different situations — but the research so far provides little guidance on how to deploy risk comparisons usefully.

2. Graphical Presentation. A bar graph showing risk probabilities at different exposure levels functioned in the Phase One research exactly like risk comparisons. It improved people's ratings of the helpfulness of the brochure and their certainty about their risk, and made them somewhat less risk-averse. However, there were no significant differences between graphical and strictly quantitative presentations of risk data in the extent to which people distinguished high levels from low levels or radon from asbestos. Graphical displays, in other words, did not strengthen the relationship between the actual risk and people's responses to that risk.

In Phase Three, a matrix of dots (to represent the denominator of a risk probability fraction) and a matrix of dots and X's (to represent the denominator and numerator, respectively) were tested for two hypothetical health risk decision problems. Neither form of graphical presentation had any effect on people's responses.

It is of course possible that different graphical devices would show a greater impact on risk response.

3. Magnitude of Test Numbers. The Phase Two research tested the hypothesis that people respond to risk data in terms of the magnitude of the test numbers themselves, quite apart from the risk represented by those numbers. By expressing asbestos risk alternatively in fibers per liter and in fibers per cubic foot, a 30-fold difference in numerical magnitude was achieved without any difference in risk (as presented in terms of probabilities plus smoking comparisons). No significant effects of the magnitude manipulation were found.

This somewhat surprising finding is reassuring. Concentration levels for radon in water, for example, are typically much greater than for radon in air, although the waterborne risk is usually lower. It is encouraging that homeowners are apparently able to disregard the misleading test magnitude cue, at least when mortality information and smoking comparisons are also provided.

4. Simultaneous Presentation. An additional factor tested in the Phase Two research was the possibility that the simultaneous presentation of asbestos and radon risks on the same ladder might help subjects understand that the asbestos risk was less serious than the radon risk. This hypothesis was rejected. There were no significant differences between the joint and separate presentations for either radon or asbestos.

Of course it is possible that a different use of simultaneous presentations might help owners take note of risk differences — for example, presentations that included the different action levels for the two hazards, or presentations that directed readers' attention to the differences more forcefully or interactively.

5. Information Overload. One of the formats tested in Phase One presented more information than any other (risk probability data, risk comparisons, an action standard, advice, verbal labels, a risk ladder). Yet it scored as well as or better than the other formats on almost all measures of communication success, including people's certainty about the risk and their evaluations of the amount of information provided and the helpfulness of the brochure.

"Information overload" may be an issue for still more complex presentations of risk information, or for audiences that are less interested or less educated. But no evidence of overload has been found for the formats tested so far. For most uses, in fact, this "maximum information" condition is probably optimal. The likely exceptions would be cases where no action standard

exists or where apathy is a major problem and the communicator wishes to encourage maximum risk aversion.

The Bottom Line

In general: Don't worry about information overload in written materials, so long as each piece of information is relatively simple and clearly presented. Always include an action standard if one exists. Except where maximum risk aversion is your goal, always include risk probability data (if they are available), an appropriately constructed risk ladder, and advice for different levels. Unless it is misleading, always include a comparison to normal background if the information is available.

If you are worried about apathy and want to encourage maximum risk aversion: Give people a standard and no other risk information.

If you are worried about panic and want to encourage minimum risk aversion: Give people advice for different levels, specifying at what levels you recommend action and at what levels you recommend doing nothing. Include a risk ladder that extends to levels higher than those your audience will experience. If the situation seems likely to generate outrage out of proportion to the risk, include a comparison to background if at all possible.

If you are trying to help people distinguish high from low levels of a single hazard: Construct a risk ladder with the high levels at the top and the low levels at the bottom.

If you are trying to help people distinguish between hazards or understand that all levels of a particular hazard are relatively high or low in risk: Construct a ladder that is extended upward or downward, so that the hazard you want to depict as low has all its levels near the bottom of the ladder, and the hazard you want to depict as high has all its levels near the top.