REDUCING RISK

by Lenny Siegel
September, 1998

For all the effort that has gone into the development of tools for assessing risk from hazardous wastes, surprisingly little has been done to figure out how to use that information to determine remedial responses. I serve on a number of committees dealing primarily with toxic waste and the hazards of unexploded ordnance, and within those groups it's uniformly clear that we lack the tools to develop risk management strategies in a cooperative way.

I come from a background of environmental advocacy. The solution to pollution is to make the polluters pay to remove or destroy all detectable contamination. Even when that's not possible, the best way to protect public health and/or natural ecosystems is to work toward that goal. Not only do aggressive cleanup requirements reduce risk from old waste, but they encourage polluters to act vigorously to prevent repetition of their past mistakes. Though I recognize that there are often practical limits to this approach, I believe it's the most sustainable cleanup strategy. Unfortunately, the political winds are blowing in the other direction.

In recent years polluters and other responsible parties, as well as increasing numbers of environmental regulators and politicians, have concluded that try to reach pristine cleanup goals is infeasible, or at least not worth the cost. Their approach is to decide, up front, upon contaminant concentrations that represent an acceptable level of risk,. Then they figure out how to achieve that level by either by reducing the contamination or merely by interrupting the pathways that expose people (or ecological receptors) to that contamination.

In essence, the proponents of the "risk-based" approach to cleanup have captured the word "risk." Activists often don't want anything to do with risk assessment or any discussion of risk because they fear that those responsible for environmental contamination will use complex formulas and their political weight to "risk away" serious problems. Thus far, much of the debate has focused upon the adequacy of traditional risk assessment. Are all cultural groups, with their dietary and lifestyle differences considered? Are children, the aged, and pregnant women considered? Do assessments cover one source in a narrow timeframe, or do they evaluate multiple, cumulative, and synergistic impacts? Are risk assessors only measuring the risk of additional cancer deaths, or are they looking at other medical consequences such as reproductive disorders, retarded brain development, and liver disease? Are they considering the impact upon animals, plants, or the natural habitat as a whole?

It will take some time to solve those problems, but I believe risk assessment methodology is already undergoing significant improvement. Those who use it will always downplay the uncertainty in their estimates. But that's not the root of the problem. The real challenge is to come up with a better way to use risk estimates to decide upon cleanup goals and to select remedial responses: Risk management.

I see two fundamental shortcomings in prevalent approaches to risk management. First, risk managers often act as if they only have one shot at a response. In reality, a good response strategy emphasizes source reduction and/or pathway interruption early on, followed by long-term removal or treatment of contamination. As new information on a site is collected, or as new technologies emerge, it may make sense to revise the remedy more than once. For example, where there is groundwater contamination, it makes sense to supply alternate water supplies immediately, but in the long term efforts should be made to restore the aquifers. For unexploded ordnance, signs and fences may offer some protection in the short run, but range clearance is usually necessary in the long run.

Secondly, there is too much focus on absolute levels of acceptable risk, something which we don't know how to measure well. And the acceptability or risk depends upon who causes it, upon whether people undertake that risk voluntarily, and upon the emotional elements of the risk itself. There is no magic to the concept of 10^{-6} additional cancers. As far as I can tell, it's derived from the old figure of speech, "one in a million." Those responsible for paying for cleanup are usually willing to accept more risk—one in ten thousand, let's say. Knowing that polluters usually want to escape liability, I always argue in the other direction. But there is no magic level of acceptable risk.

I don't propose to eliminate the tension between polluters and the communities that they have encumbered, but I do think it's possible to create a framework that helps resolve that tension. The starting point is simple. I believe people on both sides of the debate share a simple precept: It doesn't make sense to spend a huge amount of time and money to reduce risk by a small amount, but it does make sense to expend relatively little time and money to get a bigger risk reduction bang. Of course, there is substantial disagreement over how much is a large or small resource commitment or a large or small reduction in risk.

Before I lay out my suggestions, I want to make clear that I am not suggesting the development of a different set of formulas that can be plugged into virtually every cleanup scenario in the hope of coming up with hard and fast answers. This is simply a tool that can be used to improve the somewhat political process inherent in multi-criteria frameworks such as the National Contingency Plan. I have spent the last several years promoting the significance of site-specific input from the people likely to be affected most by cleanup—or the lack thereof—so it is particularly important to me that a new risk management paradigm not be used to confuse or override the concerns of public stakeholders. Rather, it should offer them clearer opportunities to influence decisions before they are first made.

Figure 1 presents two graphical methods of viewing the cost effectiveness of risk reduction activities. The first simply plots risk, which might be an amalgam of potential additional illnesses and deaths, against the cost of remediation. Though graphs are rarely used, this is the way that remedial options are typically presented. At the y-axis (zero dollars), the line shows the estimated original risk to human and ecological receptors. As money is spent, the risk goes down, but at some point graph levels off as additional activity produces minimal or no risk reduction advantage.

Figure 1b presents the same reality, though I did not even attempt to have it match Figure 1a. The second figure plots risk reduction—that is, the difference in risk as a result of the remedial response—against cleanup expenditure. It should actually represent the absolute value of the slope of the first curve at each dollar (x-axis) value. When cleanup expenditures reach a point of diminishing returns, the y value (risk reduction) approaches zero. Note that time, often the x value in similar graphs, is not plotted. Depending upon the purpose of the exercise, that is both a strength and weakness of this method of representing risk.

The advantage of the Figure 1b method for standard-setting and remedy selection is simple. In the absence of other factors, greater y-values are “better”—that is, they deliver greater risk reduction for the dollar.

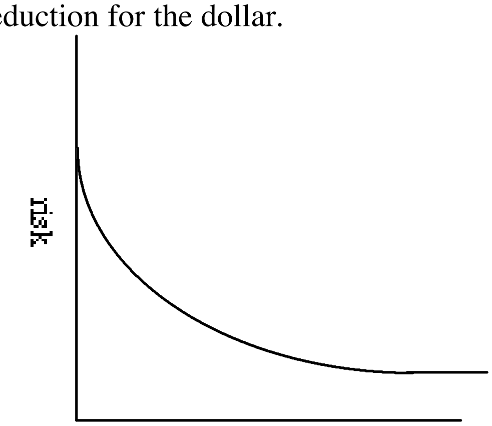

remediation cost ($)

Figure 1a

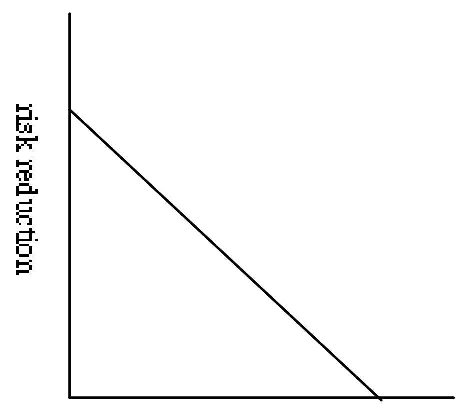

remediation cost ($)

Figure 1b

Of course, the graphs presented in Figure 1 are ideal or notional. Figure 2 represents what might be more realistic values. In the real world, actual risk reduction does not follow simple mathematical functions, and even in the best of circumstances we can only project ranges of effectiveness, not single values at each point.

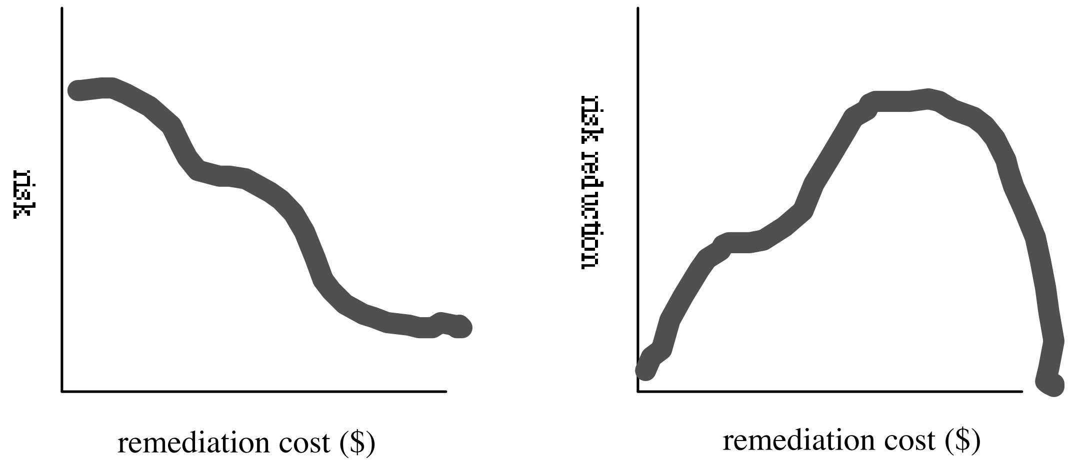

Figure 2a

Figure 2b

The risk vs. expenditure display found in Figures 1a and 2a provides a good way to visualize the predominant current paradigm for risk-based remediation. At the start of the process, decision-makers determine the exposure and contaminant concentrations likely to cause an acceptable level of risk. This is illustrated by a horizontal line in Figure 3. The remedial cost, therefore, is determined by the intersection of that line and the curve representing the remedial response, typically selected later.

But that's not really how it works. Often, the selected remedial options are incapable of reach the remedial goals. For example, the National Research Council (*Alternatives for Ground Water Cleanup,* 1994) found that only 8 of 77 operating pump-and-treat systems that it studied had achieved ground water cleanup goal, and more important, that 34 of the 77 locations were unlikely to ever reach their goals. Second, risk-based goals may be abandoned with a finding of technical impracticability, which includes not only problems that are technically infeasible to solve, but those where the remedy is prohibitively expensive or where cleanup would create other problems, such as ecological harm. Formal findings of technical impracticability are still infrequent, but often responsible parties informally argue that it's not worth pursuing the original cleanup goals. Third, responsible parties, often with regulator backing, often substitute institutional controls for active remediation. Rather than remove contamination, they assert that the controls will keep exposures down below acceptable levels for the life of the contamination. Institutional controls have their place, but rarely must the responsible party demonstrate, to win approval for them as remedies or partial remedies, that they will remain effective.

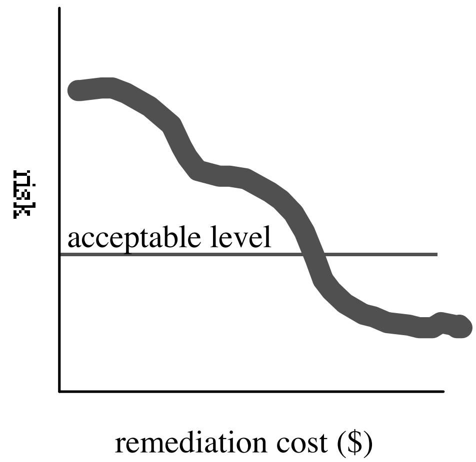

Figure 3

If one takes a more dynamic approach, in which the acceptability of expenditure on an activity is determined by the effectiveness of that activity in reducing risk, then the risk reduction vs. expenditure graphs (such as Figures 1b and 2b) are more helpful.

As Figure 4 shows, this method provides a graphic way to compare various remedial response strategies. The response that looks like it will produce more bang for the buck appears at the top. In this case, it's response C.

However, that doesn't mean that response C is necessarily the best strategy. It may be slower. It may cause ecological damage. It may create off-site risks—in transit, at land disposal sites, at incinerators, etc. (In theory, off-site risks could be calculated into the risk curve, but that's more of an improvement than I consider politically possible at this point.)

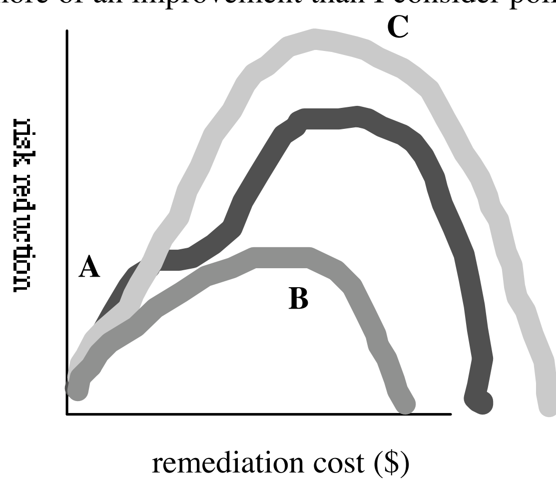

Figure 4

Once decision-makers decide upon a remedial strategy, using the risk reduction per dollar graph and all of the other criteria envisioned in the National Contingency Plan, there is still more work to be done. For example, if the best response appears to be the installation and operation of a pump-and-treat system, then it's still necessary to decide, or at least project, how much to spend on that remedy. That expenditure, represented by a point along the curve, corresponds to a level of remaining contaminant as well as a time period.

Points 1 and 5 should be rejected, since expenditures at those points would do little to reduce risk. Points 2, 3, and 4 appear to represent better choices. To maximize the risk-reduction per dollar, one would pick the top point (3) on the curve.

But there may be good reason to spend more, particular if the parties at risk are not those responsible for financing the cleanup. The area underneath the graph represents the cumulative risk reduction, so the area in Figure 5 covered by horizontal lines represents an advantage of going beyond the point of maximum projected efficiency.

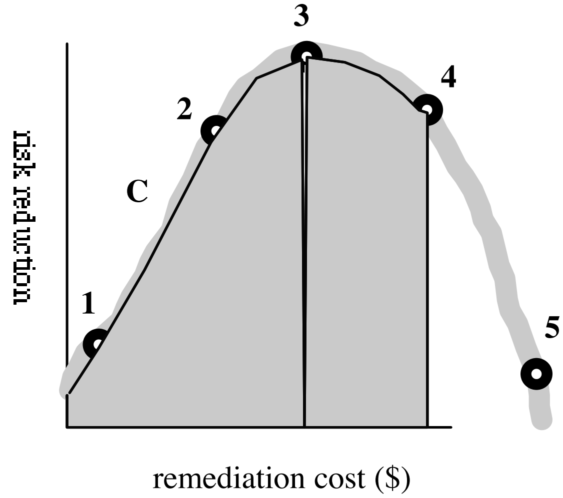

Figure 5

Figure 4, in which I compared three remedial strategies, may be too simple. In the real world, multiple strategies may not compare so easily. That is, as I have illustrated in Figure 6, the curves representing the cost effectiveness of remedial options may cross one or more times. In such cases, it may be necessary to move back and forth between the selection of the best curve and the choice of a point along the selected curves (Figures 4 and 5).

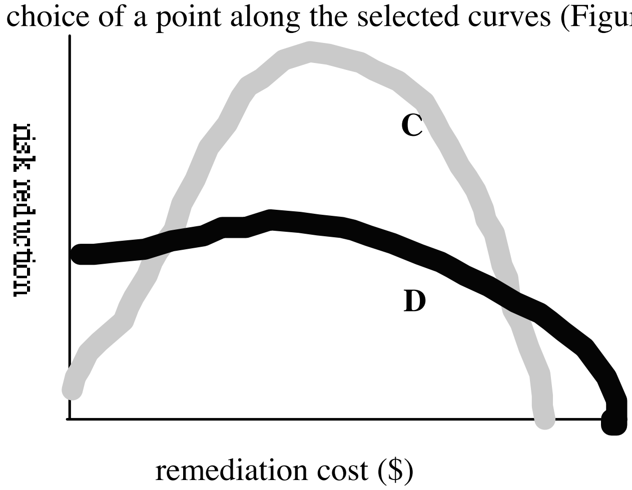

Figure 6

While this appears at first to present a complication in the decision-making process, it actually provides a mechanism for integrating the temporal dimension into the remedy selection process. (As I explained at the start of this article, that's one of the principal problems with remedy selection, as currently practiced.) I don't think it provides all the answers, however, since time is not directly illustrated on the graphs.

Figure 7 illustrates how this methodology might help evaluate interim responses or removal actions into the decision-making process. These are actions which are often designed to reduce risk quickly by, for example, breaking contamination pathways. They offer significant risk reduction for a small expenditure, but as time goes on they lose their risk reduction value. In fact, some interim actions—such as fencing or signs designed to keep people out of contaminated areas—may break down. Recalling that the graph doesn't not show the time variable, the curve may simply end—at point 6, for example—when there are no more funds spent at the site, or if funds are spent the risk reduction may actually go negative. To fully evaluate these consequences, it may be necessary to supplement these graphs with risk reduction vs. time curves.

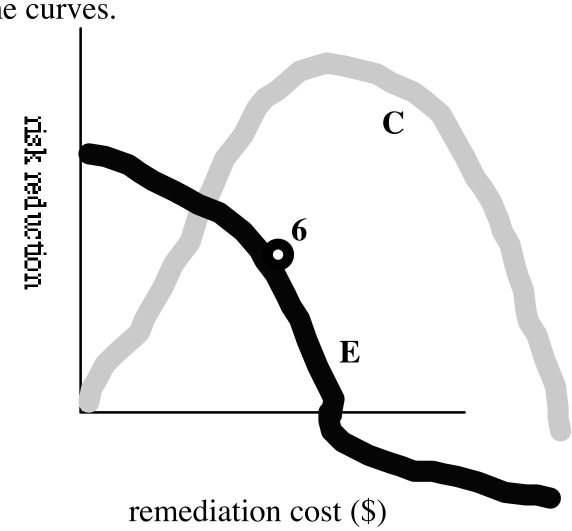

Figure 7

The graphs offer another time-dependent advantage. It is possible, in fact likely in many situations, that new, improved technologies will become available after the initial remedy is in place. Figure 8 is designed to illustrate this contingency. If the new technology is so advantageous that investing in its use—on top of existing expenditures—brings a greater risk reduction per dollar, then it should be considered, no matter how late in the process.

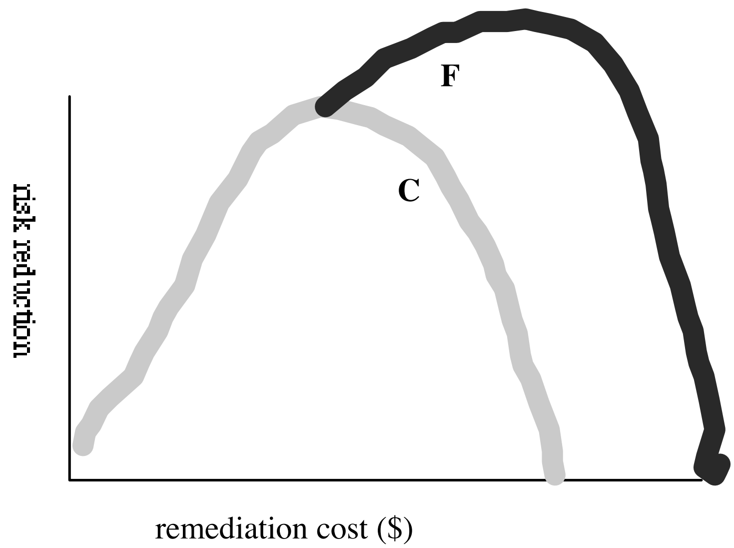

Figure 8

If, as I have often regarded, that a good remedial strategy often means applying a sequence of remedial responses, then the curves represented by the last three illustrations can be combined to make the curve Figure 9. This represents an initial pathway-breaking response, followed by use of a traditional cleanup technology, complemented by an innovative technology that further improves risk reduction per dollar. If the curves accurately represent the use of the technologies, then it makes sense to continue cleanup at least until point 7. Cumulative risk reduction, illustrated by the shaded area beneath the curve, is improved by combining the curves.

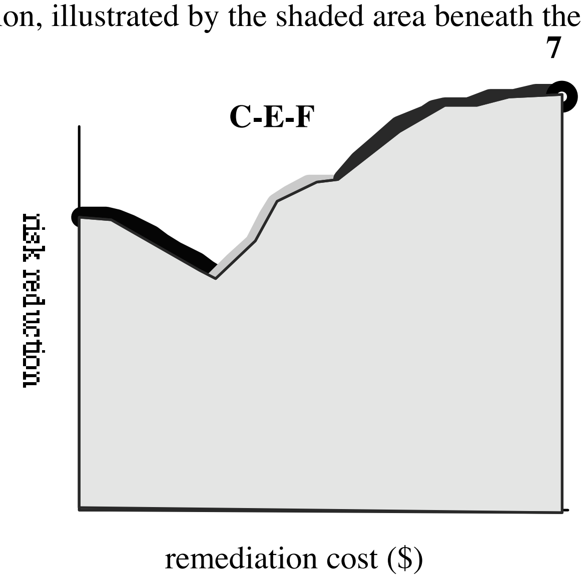

Figure 9

I believe that this framework can substantially improve the way we compare and choose remedial options. It should make it easier to visualize the costs and benefits of different cleanup strategies, particularly over time. It helps relate short-term actions, such as pathway interruption, to long-term cleanup, and it provides a mechanism for determining when continued action is no longer effective. In fact, it provides at the remedy-section stage a tool for considering the financial costs and risk reduction value of long-term operation.

In its simplest form, this framework provides a tool for qualitative comparison. However, formulas can be developed to project quantitatively the risk reduction per dollar offered by a range of remedial responses. This seems like an enormous task, since most risk assessment looks only at absolute risk. However, there are likely to be many situations in which it is much easier to quantify the risk reduction resulting from a particular action than it is to put a number on absolute risk.

In conclusion, I want to stress that I am proposing a framework to help evaluate remedial options against some of the criteria established by the National Contingency Plan (or similar lists). I am not proposing to replace it. And I strongly believe that there are many other factors, not represented on my graphs—symbolic threat, cultural values, future use, etc.—that must also be considered. And I am not proposing to reopen decisions that have already been made.

I expect that some of my friends will be disappointed because I acknowledge that we can't clean up everything. In fact, I'm trying to extend cleanup dollars as far as they will go.

Finally, this construct doesn't pretend to eliminate conflict—that is, to supply a magic answer based upon site-specific input. Rather, it offers tools for defining the debate. I invite others to challenge or augment what I propose. To me, the requirement is clear: We need a better way to integrate risk reduction goals into the recurring process of selecting remedial responses.