CASE STUDY

Evaluation of Remedial Alternatives for Contaminated Sediments

A Coherent Decision-Making Approach

John Connolly, *Quantitative Environmental Analysis, LLC*

As I talk about methods for evaluating contaminated sediments, a bias will come through. I want to acknowledge that this work is not mine alone but the combination of efforts by Dawn Foster, Warren Lyman, and me. The three of us have been involved in the trenches, evaluating sites and trying to come up with appropriate remedial alternatives to address contaminated sediments.

The goal that almost everyone has when looking at contaminated sediments is to try to find some permanent remedy, one that protects human health and the environment. There is a typical approach applied at most sites. Go into a site, look at data, and decide whether an unacceptable risk exists. That is a bit complicated and somewhat controversial because of how we define risk. I will not get into that here, but think about it, because an important issue in determining what we do at a site is how we define the risk. If there is an unacceptable risk, then in most cases, we immediately move to evaluating the feasibility of various remedial options—you have to do something now. We set out remedial action objectives, evaluate options relative to those objectives, choose an option, and then attempt to clean up the site.

At most sites, the preferred option is to remove the contaminated sediment. There is a presumption that removing sediment accelerates recovery. There is a presumption that, by taking the sediment out, we have eliminated a risk that some catastrophic event will occur that will reset the clock, as John Haggard said earlier, and bring to the surface sediments that may have been buried. I would like to challenge this approach by saying that it is not axiomatic that taking out sediments accelerates recovery, at least not in all cases. I will give two examples; I am sure there are others.

In 1994 and 1995, about half of the polychlorinated biphenyl (PCB) mass in New Bedford Harbor, in Massachusetts, was removed. There is a program in which caged mussels are sampled. They were sampled before, during, and after the dredging operation, through 1997. The caged mussels have shown no reduction in contaminant levels as a result of taking out half of the PCB mass. There were other reasons to go after the PCB mass in New Bedford besides accelerating recovery, because of the levels there. The other example is the Grasse River in New York, where 27 percent of the PCB base mass was removed by dredging in 1995. A resident fish sampling program has been going on since the early 1990s. That program has shown no effect associated with the removal of 27 percent of the PCB mass.

Why does mass removal not necessarily accelerate recovery? I will suggest a few reasons. It may be that the sediments taken out were not the dominant contaminant source for the ecosystem to begin with. That could happen if ongoing sources are part of the problem. We talked earlier about ongoing sources and how to address them. It also may be true that the source issue is a surface-area phenomenon as opposed to a "hot spot" phenomenon, and if we went in and removed the hot spots, then we may not have addressed the problem.

It is also possible that we have not substantially reduced surface sediment concentrations by taking the sediment out. That happens in places where dense, nonaqueous phase liquid (DNAPL) is present. When you remove sediment, DNAPL tends to move toward the bottom, because it is heavier than sediment. The removal efficiency for the oil would be less than the removal efficiency for the sediment. If the concentrations are much higher at depth than at the surface, then there is a good chance, or at least a chance, that the residual concentration left behind will be close to—or maybe even higher—than what was there at the start. Similarly, if the contamination extends down to hardpan, which means that the dredge cannot get an overbite with clean sediment, there is the potential of leaving contaminated sediment behind.

I will quickly discuss a few examples of these types of issues. First, an example of an ongoing source problem is Lavaca Bay in Texas, a mercury-contaminated site. Like a number of other sites with which I am familiar or have been involved, the initial focus was on the sediments. The sediments were the problem; the focus was on what we could do about the sediments. It was only after quantitative evaluations of what was going on in Lavaca Bay that it became clear that maybe contaminated sediments were not the real problem.

We made a vertical profile of mercury concentrations in the sediment core. Then, based on the history of mercury releases in the late 1960s, we developed a model predicting what the concentration profile would look like assuming that the only releases in the system were the original ones. That profile does not look anything like the measurements you get close to the surface of that sediment core. The reason is that the concentrations of mercury in the surface sediments of that core are due largely to ongoing sources as opposed to historical releases. At sites where there is not necessarily a point source that you can focus on right away, the issue is complicated and the source is sometimes not obvious.

With regard to the issue of hot spots versus surface area, it becomes important to look at problems in the right units. If we look at organic contaminants, for example, then the right units are normalized organic matter because that is what the organisms are seeing. The benthic organisms are eating so many grams of organic matter per day, so their dose of PCBs is related to the organic matter PCB content. In water, PCBs are controlled by what is on the particles of organic matter, so the fluxes from sediments depend upon what is on the organic matter.

If you look at PCB concentrations in the Hudson River, both in areas designated as hot spots (because they have dry weight concentrations significantly greater than other areas of the river) and in other areas, and you normalize the data to get micrograms of PCBs per gram of organic carbon, there is no difference. The hot spots and non-hot spots are comparable. In 1984, the numbers were essentially the same; in 1991, the number is slightly higher—statistically, it was not higher—in the non-hot-spot areas. In this case, we are looking at a surface-area problem. The hot spots in Thompson Island pool in the Hudson River represent 10 percent of the surface area. If you dredged out the hot spots, then you would have removed just 10 percent of the surface area. You would have left behind 90 percent of the surface area, which had the same concentration on an organic carbon basis as did the hot spots.

With regard to our ability to get stuff out, we have to be careful when there are high concentrations at depth. One example is a sediment core profile we did of PCBs in a river. The PCB concentrations were very low near the surface, although actually not that low from the standpoint of what most people would consider a risk-based evaluation. The surface concentrations were about 20 parts per million (ppm) in this core. About 107 m into the core, there was a peak PCB concentration of almost 1,300 ppm. The bottom of the core was hard material. We did not know if it was truly hardpan or not, but it certainly would be hard to dredge. Down at the bottom of this core, the concentration was almost 300 ppm. If we dredged here because of the high concentrations at the bottom, to the extent that this was hardpan, it would be difficult to reduce the concentration relative to what is already at the surface. Dredging might or might not have the intended effect.

When we evaluate sites, we need to consider all of these issues. It is not enough to say there is an unacceptable risk and therefore the presumptive remedy is dredging. Dredging may work. It works in some places, but it does not work everywhere. In cases where we are looking at significant risks and significant costs, we need to do what I call a prognostic risk assessment. We need to evaluate all of the alternatives in terms of how they reduce risk. We need to compare natural recovery to various other options, and we need to be frank with ourselves. Let us not presume that dredging will be effective; let us look at the things that might affect dredging to determine whether or not it would be effective, and then put it on the same plot as the other alternatives and look at risk reduction.

I will run through a proposed procedure for doing that type of a risk assessment. The first thing that we clearly need to do at all sites is to look at the distribution of contamination spatially and vertically, in three dimensions. We need to have the data appropriately normalized. To look at concentrations on a dry weight basis and conclude that it is high here and low there and, therefore, we have to address that, is missing the issue. If we are looking at organic contaminants, then we should carbon-normalize all the data to decide where the problem areas are. If we are looking at divalent metals, then maybe we want to normalize by acid volatile sulfides. We have to know what the contaminant levels are in the buried sediments, at what depth there are clean sediments, and whether we can get an overbite with a dredge.

In all cases, we have to determine the significance of ongoing sources. At many of these sites, the ongoing source is not obvious; there is no pipe sitting there with a permit that tells us it is putting out 20 pounds of contaminants per day or per year and that this is part of the problem. At many sites, the ongoing sources are non-point, groundwater sources that we may not even know about. To determine whether these sources exist, you can do some things with the data, to the extent you have data. The spatial and temporal trends in the data may reveal something about ongoing sources. We also can conduct mass balances. In the absence of knowing whether there is an ongoing source, can we balance all the sources and sinks, or is there a piece missing? Are we missing some particular source that we can use to balance all the sinks? When the sinks are a lot bigger than the sources, are we missing a source?

We need to establish the rate of natural recovery. If ongoing sources are not important, then we can establish this rate based on temporal trends. If we have data over time, and if contamination levels are going down, then we can use those data to establish the natural recovery rate. However, if there are ongoing sources, then the trend we see in time is not reflecting natural recovery; rather, it is reflecting the influence of the ongoing sources. Then we need to do more research. We need to look at things like burial rate—how fast are sediments accumulating, if they are accumulating? We need to look at degradation rates—does this compound degrade, and at what rate?

Because this is a prospective risk assessment, we will try to look at risk reductions in the future. We will use a model. I think we need to constrain ourselves to quantitative models, which by definition have to conform to physical laws. (Sometimes we create models in our heads that violate laws such as conservation of mass, and we never know it.) The nice thing about quantitative models is that they are testable—all the assumptions are defined explicitly; you can see them. (The models in our heads, however, make lots of assumptions but they are not necessarily explicitly defined.)

The other nice thing about quantitative models is that they take advantage of all the science. They use our full scientific understanding. We know a lot about PCBs, for example, and how they behave in the environment. All of that knowledge can be incorporated into a quantitative model. We can use the totality of the field data. We can integrate, for example, water column data, sediment data, and biota data in the context of a quantitative model and evaluate the consistency of all that data. It then becomes an objective tool—it does not know anything about politics—for projecting future concentrations; by using that objective tool, we have a basis on which to make remedial decisions.

This type of approach is not new; it is applied in many places, including rivers, bays, and large lakes. There are a lot of PCBs, but also other contaminants, such as Kepone in the James River and metals in the Patuxent River. The models allow us to test the efficacy of practical alternatives. We can get an estimate of risk reduction because we can predict the concentrations in water sediment in the future and use that as a basis for estimating risk in the future.

A model also allows us to look at the permanence of the remedy. Remember, we are looking for a permanent remedy, and there is always this nasty voice in the back of your head that says, "Well, if I leave the contaminant out there, then there is a risk that this will not be a permanent remedy." The model is an objective tool for evaluating that risk. The models have been used successfully to evaluate the impact of catastrophic events, such as floods and hurricanes, for example.

I will conclude by saying that, whatever we do, we should answer the following questions, and we should do so through a prognostic risk-assessment approach. First, we need to look at the appropriate remedial actions. How do we define the goal for the site? We have to ask ourselves, critically and quantitatively, whether removal will accelerate recovery. We have to address all the issues about ongoing sources, contamination at depth, and whether hot spots really are hot spots. Are other remedial options more effective in accelerating recovery? What is the risk associated with leaving contaminated sediments in place?

Lastly, we need to look at the collateral impacts of the remedial options. All options have collateral impacts—impacts on the ecosystem, on the community in which remedial option is occurring, and on human health. We need to keep all of these questions in our minds as we evaluate contaminated sediments. With my bias, I think that prognostic risk assessment, looking out into the future, is the approach that allows us to have all of these discussions.