Dose-Response: An International Journal

Volume 9 | Issue 2 Article 1

6-2011

PRACTICAL RISK ASSESSMENT AND MANAGEMENT ISSUES ARISING WERE WE TO ADOPT LOW-DOSE LINEARITY FOR ALL ENDPOINTS

Lorenz R Rhomberg
Gradient

Dose-Response, 9:144–157, 2011
Formerly Nonlinearity in Biology, Toxicology, and Medicine

ISSN: 1559-3258
DOI: 10.2203/dose-response.10-023.Rhomberg

International Dose-Response Society
www.Dose-Response.org

PRACTICAL RISK ASSESSMENT AND MANAGEMENT ISSUES ARISING WERE WE TO ADOPT LOW-DOSE LINEARITY FOR ALL ENDPOINTS

Lorenz R. Rhomberg, Ph.D., FATS □ Gradient

□ The 2009 National Research Council report, *Science and Decisions*, proposes harmonizing dose-response approaches for cancer and non-cancer endpoints, and for non-cancer quantitative risk assessment, this would usually take the form of a low-dose linear no-threshold dose-response curve. The soundness of this recommendation has been questioned, but I focus on its consequences if adopted, many of them apparently unintended. If most endpoints for most agents are assumed to have non-zero low-dose risks, then the critical–effect concept, choosing the one endpoint on which to calculate acceptable doses, loses its basis. All regulatory decisions, since they entail substituting some exposures (and their attendant risks) for others, become risk-risk trade-off decisions, and equity questions are raised since risk transfer is inevitably involved. A valid basis for estimating low-dose linear components is not evident, and upper-bound approaches fail to be reliably public health-protective owing to the risk trade-off decisions that need to be faced.

Key words: Science and Decisions; Non-cancer; Low-Dose Linearity; Critical Effect; Benefit-Cost Analysis; Regulatory Benefits; Additivity to Background

INTRODUCTION

A committee deliberating under the auspices of the National Research Council (NRC) published its recommendations in the NRC report, *Science and Decisions*: *Advancing Risk Assessment* (NRC 2009). Among the most discussed recommendations in this report is that the approaches to dose-response analysis and evaluation of safety of low-level exposures be "harmonized" across assessment of cancer and non-cancer toxicity endpoints. The committee recommended that, like non-cancer assessments, quantitative assessment of cancer risks includes an explicit allowance for interindividual variation in susceptibility to the impact of carcinogenic agents. In turn, for assessment of non-cancer toxicity, it was recommended that, as with the case for carcinogens, the evaluation of impacts of lower doses should be done with an extrapolated dose-response curve, assigning a varying level of estimated risk to each exposure level. This would stand in contrast to the current reference dose (RfD) approach used by the United States Environmental Protection Agency (US EPA) – which is similar to approaches used worldwide – in

Address correspondence to Lorenz R. Rhomberg, Ph.D., FATS, Gradient, 20 University Road, Cambridge, MA 02138; Phone: +1-617-395-5552; Fax: +1-617-395-5001; E-mail: lrhomberg@gradientcorp.com

which a dose is defined below which no individuals in an exposed population are expected to have adverse reactions.

The National Academy of Sciences (NAS) committee had several motivations for these recommendations. The first was interest in harmonization as such – bringing disparate and somewhat artificially separated approaches together so that the same set of considerations applies. But second, and more importantly, it aimed to make non-cancer assessments more useful to the analysis of benefits of regulation as well as to the evaluation of risk-risk trade-offs, comparative cost-effectiveness of alternative regulatory options, and benefit-cost analysis. A traditional RfD interpreted in the traditional way (or an acceptable daily intake or any similar measure) provides a "one-way" measure of exposure impact – lower exposures are deemed safe with a reasonable certainty, but higher exposures are not necessarily unsafe (owing to the use of uncertainty factors, the full extent of which may not be needed for the agent in question). The degree of certainty in population protection erodes in an unquantified way at ever higher exposure levels, followed eventually by the erosion of protection of more sensitive members of the whole variable population. But because there is no difference in calculated impact at different exposure levels, especially at levels below the RfD, there is no way to judge the marginal costs or benefits (in increased or decreased impact on an exposed population) of small adjustments in the permitted exposure level. If any exposure below the RfD is deemed equivalent in impact (*i.e.*, expecting no impact), then there is no calculable benefit gained from exposure reduction, nor any cost to allowing somewhat higher exposures than currently prevail.

Some regulatory programs are mandated to calculate the benefits of the regulations they promulgate in terms of actual expected improvement in public health – that is, in terms of reduction in the number of cases of adverse effects or in the severity of those effects. In some cases, a further analysis of monetized benefits *versus* costs is also needed, if not by statute or executive mandate, then at least for the social and political justification of the regulations. The estimates of benefits, as well as their balancing with costs, rely on continuous measures of impact, such that greater or smaller changes in exposure lead to greater or smaller changes in total population burden of induced disease. The inability of traditional non-cancer risk assessment approaches to provide such measures has long been seen as problematic for those concerned with benefits analysis.

It is fair to ask why the quantitative analysis of non-cancer toxicity should be changed from an approach used successfully around the world for decades only to accommodate the assumptions about dose-dependence that have been built into existing benefit-cost models. One answer is that, when we are faced with the question of how to adjust exposure in the range of or even above the RfD – a region where there is some rea-

sonable probability that impacts may exist – we do indeed benefit from having some notion of how probable it is that population risks of different degrees are being caused or avoided. The question is not only of the degree of population impact, but also of the confidence we have that the impact actually exists, in view of the uncertainties in using animal data, the possible existence of especially sensitive people in the population to be protected, and other factors addressed by the traditional uncertainty factors in the RfD process. In my view, this second issue – about gauging the uncertainties as to whether there are any impacts, rather than about the magnitude of the impacts themselves – is the more salient, at least until one considers exposures well above the RfD. Several authors have proposed probabilistic approaches to RfD determination, in place of the use of fixed uncertainty factors, as a way of understanding this (Evans *et al.* 2001; Baird *et al.* 1996; Price *et al.* 1997; Slob and Pieters 1998; Hattis *et al.* 2002). The compound-specific adjustment factor approach also helps in this regard, when the extrapolations covered by the uncertainty factors can be addressed with case-specific data to estimate them (WHO 2005). The NRC *Science and Decisions* report acknowledges these approaches, but it does not develop them much further. Such efforts show that progress on this front depends on clarifying our understanding and appropriately separating the specific extrapolations and sources of uncertainty, as is explored in a recent series of papers (Cooke 2010a,b; Goble and Hattis 2010; Louis 2010; Rhomberg 2010).

Once the issue of developing a continuous relation of change with exposure of the risk of inducing a non-cancer toxic effect has been introduced, the question arises, How is this to be done? A number of quantal dose-response models exist that can be fit to data on response levels at different doses, but the shape of these fitted curves when extrapolated well below the dose levels of the data points differs markedly among models and is considered an unreliable index of actual risk. In particular, most such statistical models assume as a necessary condition that, even though risk may fall off in a sublinear way with decreasing exposure, there is no exposure so low as to be without some added probability of response over the background rate prevailing in unexposed individuals. This property runs counter to the widely held view among toxicologists that, owing to the nature of non-cancer toxicity-generating processes, they should have exposure thresholds: levels of exposure below which the agent is incapable of sufficiently perturbing the biological system to engender an increased risk of the effect in question.

Belief in the existence of exposure thresholds for many non-cancer toxicity endpoints is a fundamental tenet for many toxicologists, and it forms the rationale for the nearly universally applied RfD or acceptable daily intake (ADI) approach, in which the data at higher doses are used to identify a dose level that – though not the threshold itself – can be con-

cluded to be near or below the actual population threshold with an adequate degree of confidence. The belief in such thresholds comes not merely from lack of observable effects at lower doses (which could be only an issue of statistical power to observe them reliably), but more importantly from an understanding of the nature of the underlying pathobiological changes, the degree of perturbation of normal biological processes necessary to have consequences on overall physiological function or dysfunction, and an appreciation of the robustness of the living system and its ability to counter or accommodate stressors without losing function.

The Benchmark Dose approach (Crump 1984; US EPA 2000), increasingly used by regulatory agencies, introduces a means to reconcile the contradictions of no-threshold mathematical dose-response models with the biological understanding of exposure thresholds for most non-cancer toxicity. The models' results are recognized as applying reliably *only* in the region of exposures where they demonstrably fit the data used in their parameter estimation. The fitted curves are followed downward only to the limits of that reliability, at which point they define the dose associated with a "benchmark response" that is a response level on the curve that can be estimated with reliability and that is not pronouncedly different for different fitted models – that is, it is understood as part of the observable relationship at higher exposures and not an assumption-contingent low-dose extrapolation. The associated "benchmark dose" producing this level of response (or the lower statistical bound on such a dose) can then be treated in a way appropriate to our larger biological knowledge. Any inferences about the potential for impacts below the benchmark dose can be couched as science policy decisions about how to characterize that potential in view of available understanding, and not as a falsely claimed observation of actual risks at much lower doses. In short, the Benchmark Dose approach resolves the conflict between no-threshold statistical models and threshold toxicological processes by confining its attention to the span of exposure levels where they are not in conflict and eschewing the claim that such observations of frank effects at high doses can by themselves (*i.e.*, without consideration of additional knowledge regarding underlying modes of action) directly address what may happen at lower doses.

The *Science and Decisions* report (NRC 2009), however, takes a broader view of the issue of continuous risk variation with exposure level, aiming to make it (like cancer risk dose-response curves) applicable to all doses, no matter how low. To do this requires taking on the question of how to develop a basis for the plausible extrapolation of non-cancer effect dose-response curves well below the range of direct observation. *Science and Decisions* proposes three general options for extrapolating low-dose risks of both carcinogens and non-carcinogens (conceptual models

1, 2, and 3). Model 1 proposes thresholds for some individuals, but not on a population basis. Model 2 proposes both individual and population thresholds. Model 3 proposes linear and non-threshold responses for the individual and the population. In practice, one must essentially ascertain that the chemical's effects do not follow any of the mechanisms that are posited as causing linearity to accept it as a case for Model 2. In practice, it is likely that this will almost always lead to the adoption of either Model 1 or 3, both of which assume linearity at the population level.

The reasoning put forward for linearity at low doses, even for non-cancer toxicity endpoints usually assumed to have an exposure threshold, involves several lines of argument, but the primary one is additivity to background disease processes. This argument proposes that, if an agent acts to enhance or exacerbate any ongoing underlying process that is itself part of the causation of background cases of the endpoint in question (*i.e.*, cases appearing even in unexposed individuals), then even a small amount of the agent will, by moving this causative underlying process along ever so slightly, lead to a small marginal increase in the rate of appearance of the endpoint in the exposed population. In essence, even if there is a threshold, this argument proposes that, owing to variation among individuals in key causative underlying physiological states, some individuals are already past the threshold of tolerable variation in the states in question and so have "spontaneous" cases of the endpoint. There are others whose values for the key variables are just marginally adequate, and if an agent moves these values however slightly in the wrong direction, they too will pass the threshold and succumb to the endpoint.

My purpose here is not to argue the merits of this point of view, although I have done so elsewhere (*e.g.*, Rhomberg 2004, 2009), and further discussion is forthcoming. For the moment, it is sufficient to note that the logic for low-dose linearity of non-cancer-effect dose-response curves has not been universally accepted and has been a matter of much discussion since S*cience and Decisions* appeared. It is also important to note that the arguments for such linearity are arguments in principle for the existence of a low-dose linear component to the overall dose-response relationship – they do not themselves determine how large such a component should be or even lead to a way to estimate it. They only argue that some unmeasured effect ought to be supposed to exist.

Instead, my purpose here is to ask, If we were to accept this principle, and consequently establish as a science policy that most non-cancer toxicity endpoints have no exposure threshold, what effects would this have on risk assessment approaches and, more importantly, on risk management processes? After all, risk assessment is done for the purpose of aiding and enabling risk management actions based, to the extent possible, on scientific understanding of the nature and magnitude of risks at expo-

sure levels actually experienced. It is important to ask, before adopting a novel approach, how the new method will be usable in actual risk management decision processes, what issues it will solve, and what new issues it will raise. As will be shown below, several consequences of accepting widespread low-dose linearity for non-cancer toxicity do not yet seem to have been widely recognized and may be unintended by its proponents.

The current state of regulatory risk assessment avoids many of the issues discussed below by assuming that sufficiently low exposures will be without effect. If, as has been advocated by some, low-dose linearity for all endpoints is established as a general or default procedure – to be invoked unless sufficient proof of thresholds can be adduced – then the issues I discuss below will arise as a logical consequence of that assumption. This will be true whether or not actual regulatory practice is changed, but if processes are not changed, they will be logically at odds with the low-dose linearity approach being invoked.

(1) THE LOW-DOSE LINEAR COMPONENT CANNOT BE MEASURED

The additivity-to-background argument is about the potential existence in non-cancer dose-response curves of a low-dose linear component of unspecified magnitude, but even if one accepts the logic for its existence, the reasoning does not provide a means for measuring the size of the component – that is, the low-dose potency of the agent to cause increases in response rate over the background incidence of the endpoint in question among unexposed individuals. The magnitude of the increased risk with low exposures, if there are any, depends on interactions of small amounts of the agent with generally unknown underlying physiological processes that account for the existence of background cases, and it will be these interactions in the human population, involving the causes of human background cases, that will be at issue, rather than the high-dose interactions in experimental animals on which most inference of human hazard is based. The high-dose shape of the dose-response curve, and even its low-dose extrapolation to background response rates in experimental rodent systems, does not provide any substantial information about these matters.

At present, it is not clear how proponents of linear low-dose extrapolation plan to handle this obstacle. One can imagine various calculation strategies, but none avoids a high degree of arbitrariness, since the information needed is not available in the data one typically has at hand. In practice (in the interest of "harmonization"), it is likely that a process similar to that used by US EPA for carcinogens will be adopted, *i.e.*, a line will be drawn from the "point of departure" (benchmark dose and its response over background) down to the origin, and its slope used as the "upper bound" on a low-dose potency. This procedure is likely to vastly

overestimate the actual degree of risk, if any. As will be argued below, the usual defense of this process as "conservative" (*i.e.*, erring on the side of being overly protective of public health), is not tenable when it is applied to all agents.

The consequence is that a measure of risk is to be given great importance, yet it is not measurable other than in an arbitrary way.

(2) THE POTENCY IS NOT THE SAME FOR EVERYONE IN AN EXPOSED POPULATION

The existence of the low-dose linear component is alleged on the basis of the existence of marked variability among individuals in their internal physiological states, such that some people will be just on the verge of a "spontaneous" case of the disease in question and pushed over the line by the small exposure to the agent. Thus, even if this is so, only those very marginal people will be at risk, while the main body of people will not be. The factors that cause variation in vulnerability will vary among populations, and so they cannot reliably be expected to have the same impact from one population to another. Moreover, any actual distribution of vulnerabilities must include increasing proportions of the population that are farther away from the critical level – one is after all dealing with the extreme tail of a statistical distribution that, in order for the additivity-to-background argument to work, must represent continuous variation in degree of vulnerability. Put another way, the slope of the cumulative distribution function of vulnerabilities changes as one moves from the tail toward the middle, and it is this slope that dictates the marginal increase in risk with each increasing unit of exposure – that is, the potency. This means that the potency is different for smaller or larger exposures. The magnitude of the "linearity" arising from additivity-to-background and the interaction with population variability is very local to a particular exposure level, with the slope of the linear component (the marginal increase in risk for the next unit of added dose) changing for different dose levels. This effect would be completely missed by the point-of-departure linear extrapolation approach, which imposes a single constant potency on all doses below the benchmark dose. Any single linear potency estimate will misstate the actual potency to different degrees for different dose levels. It is hard to see how such a measure, combined with the inability to estimate even such a faulty measure, would be useful for the benefits-estimation purposes that constitute the main justification for the low-dose extrapolation approach articulated by the *Science and Decisions* report.

(3) THE CONCEPT OF A "CRITICAL EFFECT" WILL NO LONGER BE TENABLE

Non-cancer risk assessment has depended on the concept of the "critical effect": the endpoint that produces an adverse impact at the lowest dose among all the endpoints that may have been observed. This critical effect can serve as the basis for the reference dose, on the logic that any control of exposure that protects against it will protect against all other effects as well, since they appear only at yet higher doses. But if all doses are presumed to engender some risk – a risk that may diminish at lower doses but never entirely vanishes – then there is no longer a basis for focusing only on one endpoint, since all endpoints will have risks that are extrapolated to low doses without thresholds. Endpoints appearing at measurable levels at higher doses than the "most sensitive" one nonetheless have impacts at lower doses, albeit perhaps with lower frequencies of occurrence at any given low dose. The probability of being unaffected by an exposure is the complement of developing at least one of the endpoints that appear at higher doses, all of which each person would be presumed to be at some dose-dependent risk of developing.

If additivity-to-background and consequent low-dose linearity is presumed to apply to most critical effects (as they have been formerly defined), then there is no obvious basis to deny the principle's applicability to other endpoints, and the total risk of an adverse effect at low doses will always be bigger than that calculated from only the "most sensitive" endpoint by an amount that depends on how many other endpoints there are and what the low-dose risk projection from each individual endpoint may be.

Once this door is opened, a host of other difficult questions arises. Are the several endpoints statistically independent of one another in their occurrence? If so, a ready (though complex) means of calculating their joint probability is available (one minus the product of n terms, each of which is one minus the risk for each of the n endpoints, and then corrected by a series of two-way, three-way, *etc.* cross-product terms). Statistical independence would be expected if each of the endpoints arises at low doses because of susceptible people being vulnerable in different underlying physiological states, but if to any extent the state of the same contributing variables is influential for more than one of the endpoints, they ought to be partially correlated, and the total risk calculation would need to account for this, even though the basis for doing so is not clear. Negative as well as positive correlations are biologically plausible, so there is no clear basis to define even a dependably conservative answer to this.

Should all endpoints observed in any animal system be included in this calculation? For instance, should risks based on rat liver enlargement be combined with risks based on mouse behavioral changes? And should these be combined with risks from rat heart arrhythmias seen in a study

that failed to detect liver enlargement? When one can identify a critical effect (in a most-sensitive sex, strain, and species), as long as that effect is deemed sufficiently reliable to serve as the basis of the assessment, one does not need to adjudicate the further relevance and bearing of each and every endpoint that has been observed in any test system. As a consequence, elaborate weight-of-evidence and hazard-identification schemes to evaluate particular endpoints and their consistency across test systems have not been seen as necessary in non-cancer risk assessment. But if all endpoints are potentially sources of low-dose risk, it is necessary to make a weight-of-evidence decision about the human relevance of each endpoint separately, only eliminating endpoints from the overall risk calculation if they are deemed insufficiently supported to constitute an indication of a potential human risk. Moreover, for each endpoint that is deemed relevant, it is necessary to identify the corresponding human adverse effects and their spontaneous background rates of occurrence, because it is these human effects that human doses are being presumed to exacerbate based on the analogy of effects in animals. In short, if low-dose linearity for all (valid) endpoints is to be entertained, the need for an elaborate method to judge the applicability of each potential endpoint is created.

The critical-endpoint concept affects the design and conduct of animal testing, and with its loss, difficult test-design questions are created. Current testing is largely aimed at identifying critical endpoints for each test system, and so it is not necessary to run many at much higher doses to see if additional endpoints crop up. Again, the reasoning is that, if endpoints have thresholds, then effects seen at higher doses will only appear at dose levels already higher than would be allowed in view of the critical endpoint. But if all endpoints have a low-dose linear component, then this justification is removed. The lack of observation of increases at lower doses for a high-dose endpoint would be seen only as a matter of chance, and it would seem appropriate to require testing of higher and higher doses to determine whether additional endpoints become statistically detectable. Indeed, failure to do so could be argued to constitute failure to identify sources of low-dose risk, since these high-dose endpoints would presumably contribute to low-dose risk estimation.

If this tack were taken, however, there is no clear way to keep it from devolving into meaninglessness. At ever higher dose levels, increasing physiological failure will be engendered, increasingly massive toxicity will be observed, and, at some point, one needs to concede that the effects being caused are not ones that appear with some probability at all dose levels but are high-dose-only overwhelming assaults on the living system – *i.e.*, they are endpoints with (very high) exposure thresholds. There is no clear basis, however, to draw a line between endpoints in the "regular" higher-dose range (which would be presumed to have a linear low-dose

component, according to the presumed generality of additivity-to-background) and endpoints in the "too high" dose range (which would be presumed to have exposure thresholds).

In short, under linear low-dose extrapolation, all high-dose effects have some impact on the total risk at low doses, since none has a threshold, and so added pressure is put on the question of determining all the effects that exist and on deciding their relevance to humans. The practice of figuring that protecting from the most sensitive effect also protects against all others loses its justification, since the most sensitive effect alone contributes only part of the total low-dose risk. How much practical difference this makes depends on how many additional toxicities are brought into consideration and how much less may be their individual contributions to estimates of total low-dose risk, with the possibility that, although each is small, they total to a substantial amount of impact.

(4) CONTROL DECISIONS WOULD BECOME RISK-RISK TRADE-OFF DECISIONS

Whenever one considers limiting the exposure to a compound on the grounds of its potential to cause public health impact, this is either implicitly or explicitly a simultaneous consideration to cause greater exposure to some other agents in the form of substitutes or alternative product formulations or even just results of shifts in activity patterns. We are always exposing ourselves to some set of compounds in the environment, day in and day out without cease, and this is in some sense a zero-sum game; any diminution of one source of exposure leads to increases in others. Sometimes, at least with some thought, the exposure shifts may be obvious and even measurable; for instance, banning a crop-protection chemical leads to greater exposure to other chemicals used for the same purpose, or, if not that, to greater levels of secondary metabolites produced by crop plants under the stress of attack by pests, or, if not even that, the exposures to chemicals associated with alternative foodstuffs to which a consumer switches. In current practice, only substitutes satisfying the "reasonable certainty of no harm" criterion would be permitted, but with low-dose linearity for all endpoints, there would be no substitutes that did not raise some additional risk question to be factored in to the decision. In other cases, the shifts in exposure may be very difficult to identify with any specificity, much less to measure.

Every compound shows some kind of toxicity at sufficiently high doses. If these high-dose effects are presumed to occur with some lower probability at any lower dose, then *every* exposure to *every* compound is associated with some degree of risk of the effects seen for those compounds at higher doses. Since every decision about limiting exposure to a particular agent entails greater exposure to other agents, such a deci-

sion is also a decision to prefer to lower the risks from the agent being controlled at the expense of raising the exposures to, and hence the risks from, other chemicals. That is, every decision to alter an exposure necessarily entails a risk-risk trade-off. Whether the trade-off provides a net public health benefit or not depends on the particulars – what is being avoided and what is being increased. The difficulty is that, as we have already noted, the actual levels of risk cannot really be estimated even for the agents specifically under scrutiny, and this problem is compounded when the trade-offs entail larger exposure to unnamed, perhaps unstudied, compounds with an unknown degree of increased risk that would be attributed to them by the no-threshold presumptions being broadly applied.

We avoid this quandary today by believing that small shifts in exposures that are well below recognized thresholds for adverse effects do not entail any alteration of risk. Thus, we can entertain a regulation to limit an exposure, and hence reduce the public health impact, for a specific chemical with some confidence that we may be making things better and are unlikely to be making them worse. Most of the time (with identifiable exceptions), we believe that the small shifts in exposure to other agents that arise as a consequence of the action to limit one particular agent leave all of the other exposures well in their sub-threshold range and hence provide no countervailing problem to be balanced with the benefit of the chemical-specific limit being contemplated. But if everything has linear risks at low doses, then this assurance cannot be had, and there is no alternative (if one seeks net public health benefit) to identifying the exposures and risks being raised so that the one target exposure and risk can be lowered, and further to measure the plusses and minuses to arrive at a judgment about whether a net benefit is being created.

As noted earlier, it is very problematic to measure the low-dose risks being traded off in such a setting. "Upper-bound" estimates of low-dose risk potential for each agent make poor bases for judging the trade-off, because the degree of over-estimation of risk is likely to differ, perhaps sharply, for the different agents being considered, and even the relative size of the overestimates will rarely be apparent. For this reason, use of upper-bound low-dose risk estimates is not dependably protective of public health, and the use of such upper bounds makes the outcomes of the risk-risk trade-off analysis essentially arbitrary.

(5) UNANSWERABLE EQUITY AND ENVIRONMENTAL JUSTICE CONSIDERATIONS WILL BE RAISED

The assumption of low-dose linearity for non-cancer effects rests on the notion that there are people in an exposed population who are just on the margins of sustainable values of some key physiological variables,

and that small pushes by environmental agents will push them over into having an adverse response. Such people must be classified as vulnerable sub-populations if they are somehow identifiable, either as individuals or as a class; it is their protection that is sought by limits on low-level exposure that would be tolerable by most other people. But the risk-risk trade-offs inherent in a world where all endpoints are assumed to be low-dose linear (as outlined in the preceding paragraphs) means that the protection of the people vulnerable to a given chemical, obtained by limitation of its use, results not only in other risks, but in other risks *to other people* – those with different vulnerabilities that leave them at risk to the increased exposure to substitute chemicals. That is, not only is there an issue of risk-risk trade-off, there is a further issue of the equity and justification of reducing one person's risk by increasing another's. Some questions are raised by this realization of the inevitability of risk-transfer. Does it matter whether one risk is of a minor adversity to many people while another is a major impact on a few? How is net benefit to be calculated? What is the responsibility of a regulator to try to identify the kinds of people who will have risk reduced and especially those who will have risk increased? In particular, is it acceptable simply to avoid investigating the question of whose risk is being increased and by how much? Does it ameliorate the problem if the vulnerable cannot be identified but are just an unidentified fraction of the population as a whole?

The concept of environmental justice is predicated on the idea that certain ethnic or economic groups might have special vulnerabilities, or different frequencies of individuals with vulnerabilities, and so the transfer of risk from those with presumed vulnerabilities to the agent to those with vulnerabilities to the substitutes is relevant. In any case, there is an equity consideration even if the vulnerabilities cannot be assigned to particular recognizable sub-populations.

DISCUSSION

Lest any reader be mistaken, let me reiterate that I am not convinced by the arguments put forth by those who would assert low-dose linearity for all endpoints. This is not a matter of fleeing the conclusions, but rather one of questioning the applicability of the arguments to non-cancer toxicity processes as we currently understand them. I have not devoted the present paper to the critique of the basis for the universal low-dose linearity argument, which is discussed elsewhere, but rather to a thinking-through of the consequences of such a stance, should it be accepted as a regulatory policy despite the qualms I and others have expressed about its soundness.

Most of the problems I discuss above arise as a result of the proposed broad application of the no-threshold/low-dose linearity principle. If it is

taken as a default or if it is applied to most environmental agents and most toxicity endpoints, it runs into the problem that every action affects everything else, and every change results in a bad outcome for someone. In my view, the proponents have not thought through the consequences of applying this idea to more than one compound and endpoint at a time, even though the logic of their position dictates that they need to do so.

By looking only at one dimension at a time – one agent to which more or less exposure can occur and one endpoint affected by it – it is possible to think that one is being protective by assuming a linear/no-threshold dose-response. But if this idea is widely applied, then the world is cast as sitting on a multidimensional knife's edge in which any change in one dimension affects everything else with no level of tolerance. The logical consequence of this view is that any action of any kind will result in some good to some and some ill to others, without any evident way of calculating how much of each is engendered and to whom fall the benefits and dangers.

As noted at the outset, a large part of the stated motivation for adopting low-dose linear extrapolation, as expressed in the *Science and Decisions* report, is to enable the measurement of the benefit of regulations and to aid matching costs of controls to the degree of benefits received. It is ironic that the means being proffered cannot produce measurements of the risks invoked, cannot identify the impacts that are implied, and in general cannot achieve the aims that at least nominally motivated them in the first place. What is really needed is characterization of the *non*linearities in the local region of exposure levels in play in the risk management application, since it is the interaction of curve shapes that yield insights into optimal actions. Moreover, since trade-offs among influences are inevitably involved, it is important to have central estimates of these effects, or, even better, to have distributional estimates of tenable descriptions of them along with characterization of our uncertainty about the reliability of the estimates.

Do we really need universal low-dose linearity to address these risk management analyses? All that is really needed is continuous (not necessarily linear) measures of risk magnitude as a function of dose, and these are only needed in the local region of exposure levels that one is considering for the immediate question at hand. Linearizing everything over a wide low-dose range is actually destructive of the applicability to the motivating risk-management questions. As I have attempted to show, there are also a host of unintended consequences of adopting this stance, and it is important to keep these larger questions in view as science policy options are evaluated in the near future.

ACKNOWLEDGEMENTS

This paper was produced by the author without financial support.

REFERENCES

Baird SJS, Cohen JT, Graham JD, Shlyakhter AI, and Evans JS. 1996. Noncancer risk assessment: A probabilistic alternative to current practice. Hum. Ecol. Risk Assess. 2:79-102

Cooke R. 2010a. Conundrums with uncertainty factors. Risk Anal. 30(3):330-339

Cooke RM. 2010b. Response to reviewers. Risk Anal. 30(3):353

Crump KS. 1984. A new method for determining allowable daily intakes. Fundam. Appl. Toxicol. 4:854(18):89-1579

Evans JS, Rhomberg LR, Williams PL, Wilson AM, and Baird SJS. 2001. Reproductive and developmental risks from ethylene oxide: A probabilistic characterization of possible regulatory thresholds. Risk Anal. 21(4): 697-717

Goble R and Hattis D. 2010. Are conundrums with uncertainty factors an obstacle to developing probabilistic interpretations of noncancer risk from chemicals? Risk Anal. 30(3):340-345

Hattis D, Baird S, and Goble R. 2002. A straw man proposal for a quantitative definition of the RfD. Drug Chem. Toxicol. 25(4):403-436

Louis TA. 2010. Discussion of "Conundrums with uncertainty factors." Risk Anal. 30(3):346-348

NRC (National Research Council). 2009. Science and Decisions: Advancing Risk Assessment. National Academies Press, Washington, DC

Price PS, Keenan RE, Swartout JC, Gillis CA, Carlson-Lynch H, and Dourson ML. 1997. An approach for modeling noncancer dose-responses with an emphasis on uncertainty. Risk Anal. 17:427-437

Rhomberg LR. 2004. Mechanistic considerations in the harmonization of dose-response methodology: The role of redundancy at different levels of biological organization. In: McDaniels T, Small MJ, (eds), Risk Analysis and Society: An Interdisciplinary Characterization of the Field. Cambridge University Press, Cambridge, UK

Rhomberg LR. 2009. Linear low-dose extrapolation for noncancer responses is not generally appropriate. Environ. Health Perspect. 117(4):A141-A142

Rhomberg LR. 2010. Uncertainty factor conundrums: What lessons should we draw? Risk Anal. 30(3):349-352

Slob W and Pieters MN. 1998. A probabilistic approach for deriving acceptable human intake limits and human health risks from toxicological studies: General framework. Risk Anal. 18:787-798

US EPA (US Environmental Protection Agency). 2000. Benchmark Dose Technical Guidance Document (External Review Draft - EPA/630/R-00/001, October 2000). Risk Assessment Forum. US Environmental Protection Agency, Washington, DC

WHO (World Health Organization). 2005. Chemical-specific adjustment factors for interspecies differences and human variability: Guidance document for use of data in dose/concentration-response assessment. (IPCS Harmonization Project Document No. 2). World Health Organization, Geneva