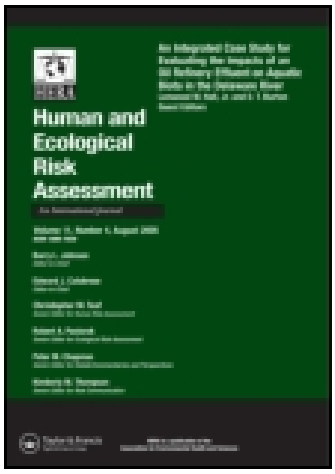

Human and Ecological Risk Assessment: An International Journal
Publication details, including instructions for authors and subscription information:
http://www.tandfonline.com/loi/bher20

Species Sensitivity Distributions Revisited: A Critical Appraisal

Valery E. Forbes [a] & Peter Calow [b]

[a] Department of Life Sciences and Chemistry, Roskilde University, Universitetsvej 1, 4000 Roskilde, Denmark.

[b] Department of Animal and Plant Sciences, The University of Sheffield, Sheffield S10 2UQ, UK

Published online: 03 Jun 2010.

To cite this article: Valery E. Forbes & Peter Calow (2002) Species Sensitivity Distributions Revisited: A Critical Appraisal, Human and Ecological Risk Assessment: An International Journal, 8:3, 473-492

To link to this article: http://dx.doi.org/10.1080/10807030290879781

Human and Ecological Risk Assessment: Vol. 8, No. 3, pp. 473-492 (2002)

Species Sensitivity Distributions Revisited: A Critical Appraisal

Valery E. Forbes[1] and Peter Calow[2]
[1]Department of Life Sciences and Chemistry, Roskilde University, Universitetsvej 1, 4000 Roskilde, Denmark. [2]Department of Animal and Plant Sciences, The University of Sheffield, Sheffield S10 2UQ, UK

ABSTRACT

We revisit the assumptions associated with the derivation and application of species sensitivity distributions (SSDs). Our questions are (1) Do SSDs clarify or obscure the setting of ecological effects thresholds for risk assessment? and (2) Do SSDs reduce or introduce uncertainty into risk assessment? Our conclusions are that if we could determine a community sensitivity distribution, this would provide a better estimate of an ecologically relevant effects threshold and therefore be an improvement for risk assessment. However, the distributions generated are typically based on haphazard collections of species and endpoints and by adjusting these to reflect more realistic trophic structures we show that effects thresholds can be shifted but in a direction and to an extent that is not predictable. Despite claims that the SSD approach uses all available data to assess effects, we demonstrate that in certain frequently used applications only a small fraction of the species going into the SSD determine the effects threshold. If the SSD approach is to lead to better risk assessments, improvements are needed in how the theory is put into practice. This requires careful definition of the risk assessment targets and of the species and endpoints selected for use in generating SSDs.

Key Words: application factors; extrapolation; probabilistic methods; risk assessment; uncertainty.

INTRODUCTION

Ecological risk assessment is a process that evaluates the probability or likelihood that adverse ecological effects will occur (or have occurred or are occurring) as a result of exposure to stressors from various human activities (USEPA 1992). One relevant definition of ecological harm is in terms of likelihood of a species population reducing to extinction. This can be effectively operationalized in terms of

* Corresponding author: Tel(voice): +45 46 74 27 23, Tel(fax): +45 46 74 30 11; vforbes@ruc.dk

Received March 2, 2002; accepted March 9, 2002

1080-7039/02/$.50

population growth rate (λ) such that if this is consistently below one it will ultimately lead to population density reducing to zero. Once an impact on a population has occurred, the likelihood and time required for population recovery will also depend on λ. Different species have inherently different population growth rates, and the sensitivity of population growth rate to changes in individual survival, reproduction and development time varies as a function of life-cycle type and the demographic state of the population (Calow *et al.* 1997; Akçakaya *et al.* 1999; Caswell 2001). Thus some species populations are inherently more susceptible to extinction than others, for example, because they start off with population growth rates close to one, because their survival, reproduction, and/or development are very sensitive to stressors, or because their life cycle is such that very small impairments in survival, reproduction and/or development lead to relatively large impacts on λ. Very often the aim of risk assessment is to determine risk to whole communities of species, and in such cases it becomes necessary to accurately quantify the variability in sensitivity of λs among species in the community of interest as a result of exposure to a stressor.

In the remainder of this article we restrict consideration to chemical stressors and define exposure in terms of chemical concentration in the environment. This is ideally represented as a distribution of the likelihood of exposure concentrations over space and/or time. If the exposure concentration below which there is no impairment in λ — effectively a λ NOEC or EC_{10} — could be determined for each species in the community, the resulting distribution of λ NOECs(EC_{10}s) would give an ecologically meaningful measure of the species sensitivity distribution for the community with respect to a particular stressor (*i.e.*, a community-SSD). The risk assessment would then involve a comparison between the community-SSD and the exposure distribution to give a likelihood of harm being done to species in the community.

One way that has been used to provide an assessment of risk in practice has been to divide a predicted or measured environmental concentration (PEC or MEC) by a predicted no effect concentration (PNEC), often referred to as a risk (or hazard) quotient (RQ). The PNEC is derived from some ecotoxicological endpoint (typically survival, reproduction, or growth) that is measured in individuals. When, as is often the case, the assessment endpoint is the population, the individual-level responses need to be extrapolated to population dynamics (ECOFRAM 1999). It is implicit that, since the effect endpoint chosen for input into the RQ is usually derived from the most sensitive species, it represents the left tail of the community-SSD. In recognition of the uncertainties in this assumption, the measured effect concentration (*e.g.*, L(E)C_{50} or NOEC(or EC_{10})) is divided by an application factor (AF). The Technical Guidance Document (EU TGD) used in association with existing and new substances legislation within the European Union (CEC 1996) provides specific recommendations on the size of application factors to be applied under different circumstances to derive PNECs for aquatic ecosystems. A similar set of factors is used in the U.S. by the U.S. Environmental Protection Agency's (USEPA) Office of Pollution Prevention and Toxics (OPPT) to set 'concern levels' (*i.e.*, the level of chemical exposure in the environment at or above which significant risks to aquatic organisms are likely) (Zeeman and Gilford 1993; Zeeman 1995). For pesticides, the mechanics of the calculations may differ somewhat (ECOFRAM 1999; CEC 1997), but the concept of reducing the measured effect concentration by a

fixed factor to account for the uncertainty in the risk assessment is essentially the same.

In practice, the RQ approach has most often used single numbers for both P(M)EC and PNEC, ignoring the uncertainty associated with these estimates. However, there is always uncertainty in P(M)ECs and PNECs, and this can be represented by fitting the values to a statistical distribution. Thus, combining the distribution of PEC/MECs with the distribution of PNECs (*e.g.*, using resampling techniques such as Monte Carlo), the probability that the PEC/MEC exceeds the PNEC can be estimated, given the uncertainty in each (Calow and Forbes 1999; Campbell *et al.* 2000). This approach can be particularly helpful in borderline cases where risk quotients are close to one and there can be arguments about worst- and best-case assumptions (Calow and Forbes 1999).

In response to criticism of the application factors used in deriving PNECs, an alternative approach has been advocated that attempts to approximate a community-SSD by incorporating observations on a variety of species from ecotoxicological tests (Stephan *et al.* 1985; Aldenberg and Slob 1991; Wagner and Løkke 1991). The resulting distribution enables estimation of the concentration at which effects on some (large) fraction of the tested species will not occur. This value is then compared to the measured or predicted exposure concentration to get an estimate of risk. An often-cited criticism of the RQ approach is that it does not provide a quantitative probability of risk. In principle, an advantage of the SSD approach is that statistical techniques can be used to combine the uncertainty represented by the frequency distributions of effects and exposure concentrations to provide a statement of probability of harm to the selected group of species. However, since the species used for input into the sensitivity distributions generally are not derived from any known community, the ecological interpretation of the resulting risk probability is not obvious.

Thus, both the RQ and SSD approaches can lead to statistically based risk assessment that recognizes uncertainty in exposure and effect concentrations. The RQ approach has the advantage of being simpler to apply, but it has been criticized for being somewhat arbitrary. The SSD approach makes more use of all available information on effects and has been advocated as a more objective and scientific procedure. In what follows below we revisit the fundamental assumptions used in the theory behind the SSD approach and in putting the theory into practice. Our main questions are (1) Does the SSD approach clarify or obscure the setting of ecological effects thresholds for risk assessment? and (2) Does the SSD approach reduce or introduce uncertainty into the process of risk assessment?

Our conclusions are that if we could determine a community-SSD, this would provide a better estimate of an ecologically relevant effects threshold and therefore be an improvement for risk assessment. However, the SSDs that are often generated do not represent any known community, but are often interpreted as if they do. Good risk assessments should be transparent, by which we mean: (1) the quality of the input data can be readily assessed by an independent evaluator of the risk assessment (*i.e.*, the risk assessors should provide explicit quality criteria used in data selection); (2) all assumptions and decisions are clearly stated and justified (whether these are based on scientific evidence or are policy decisions); and (3) conclusions drawn represent a precise and accurate interpretation of the analysis. We shall

demonstrate in the following that risk assessments based on SSDs often have not been as transparent as they could (and should) be, and this tends to obscure the uncertainties going into the effects assessment. The SSD approach also introduces uncertainty into risk assessment in that many of its assumptions either have not been, or cannot be, tested in practice. These issues are particularly pertinent when it is remembered that risk assessments carried out for regulatory purposes are usually effected in a tiered way and that decisions about proceeding to higher tiers are based on the results of tests at each of the preceding tiers. Thus each step in the process needs to be robust enough to be applied routinely and conservative enough to ensure that the process is not stopped prematurely.

WHAT ARE THE ASSUMPTIONS?

In this section we list the major assumptions associated with the theory behind the species sensitivity distribution approach and the assumptions that need to be made in putting the approach into practice.

Assumptions Behind the Theory

T1. Interactions between species do not influence the sensitivity distribution — the sensitivity distribution represents each and every species within a community as an independent entity (Wagner and Løkke 1991). In reality, communities represent complex interactions between species such that the impairment of any one species may have knock-on effects on others, for example, through trophic and/or competitive interactions. However, the fewer the species in the community that are impaired, the lower the likelihood that interaction effects will have a major influence on the SSD.

T2. All species are weighted equally — the sensitivity distribution approach assumes that the loss of any species is of equal importance to the system, for example, in terms of stability properties and/or process properties, and therefore the effect values for the species going into the distribution receive equal weights. Keystone or other functionally important species (Paine 1966; Lawton and Brown 1993) are implicitly assumed to be distributed randomly in the sensitivity distribution and be as equally likely as other species to fall in the extreme left (most sensitive) tail of the distribution (Forbes and Forbes 1993).

T3. Structure is the target of concern — There is a distinction between community structure, represented by species composition, and ecosystems, as represented by the combination of community structure and underlying processes that involve fluxes of energy and matter (Cairns and Pratt 1995). The sensitivity distribution approach focuses on community structure and makes no direct connection to underlying ecosystem processes (Forbes and Forbes 1993). The available evidence suggests that whatever the form of the relationship between community structure and ecosystem process, that species composition is generally at least as sensitive to stress as changes in, for example, photosynthesis, respiration or decomposition (Pratt and Cairns 1996).

Assumptions in the Application

Assumptions and decisions made in putting theory into practice and that have to apply if the SSD approach is going to improve on the RQ approach:

P1. The sample of species used to construct the sensitivity distribution is an unbiased sample of the target group of species about which conclusions are to be drawn. For example, if the risk of a chemical for marine ecosystems is to be estimated, the species selected for input into the sensitivity distribution must be an unbiased (*i.e.*, statistically random) sample of the marine ecosystem(s) of interest. If the distribution is meant to represent 'the universe of species' (Klaine *et al.* 1996), then it should be an unbiased random sample of such. This assumption presents both theoretical challenges (including defining an ecosystem and deciding the basis (*i.e.*, number? biomass? functional role?) on which to draw a random sample from it) as well as practical challenges (*e.g.*, assuming that species could indeed be sampled randomly, it is unlikely that all, or even most, of them would be amenable to toxicity testing). Because the species used to construct sensitivity distributions are usually chosen from the available literature or databases, (1) they typically derive from a range of ecosystems (sometimes with freshwater and marine ecosystems pooled for analysis), and (2) their number and identity will vary among chemicals. There is therefore little reason to expect a consistent relationship between the SSDs that have been constructed to date (*e.g.*, those shown in Table 1) and impacts on actual communities or ecosystems. This is particularly problematic when the same SSD is combined with site-specific exposure data to compare relative risks of a chemical among sites (Cardwell *et al.* 1999). Such analyses do not provide a comparison of risk, but simply a comparison of relative exposure among sites.

P2. The endpoint is ecologically relevant. In selecting appropriate biological responses for input into a sensitivity distribution it must be assumed that these relate to harmful ecological effects. In practice responses have involved primarily acute or chronic survival, chronic reproduction, or growth. The relationship between changes in these individual-level responses and impacts on population dynamics varies considerably among species and as a result of the demographic starting point of the population (Forbes *et al.* 2001). This is generally ignored, and indeed often the data used to generate SSDs represent a mixture of endpoints.

P3. The chosen level of protection is appropriate. In defining the appropriate protection threshold (*i.e.*, 1%, 5%, 10%, *etc.*) to put into the risk assessment, assumptions have to be made about the connection between the threshold and the sustainability of the system. Strictly speaking this could be considered a policy decision and not an assumption of the model. Nevertheless, if the SSD approach is to live up to its reputation as an objective and transparent approach to risk assessment, then the chosen protection criteria should in principle have a clearly articulated and scientifically based justification.

P4. Chosen confidence limits around the protection level are appropriate. The selected threshold can be defined with varying levels of confidence chosen by the investigator, and it is assumed that the level chosen for use will be sufficiently rigorous to deliver an appropriate protection level. The level of confidence is usually chosen to be either 50% or 95% (Aldenberg and Slob 1991; Wagner and Løkke 1991). It is often unclear how decisions about the size of confidence limits are made, but they may make a large (> factor 10) difference in the protection level (OECD 1992).

P5. The shape of the distribution is appropriate. It is assumed that the shape of the distribution that is fitted to the effects data approximates to the sensitivity

Table 1. Review of articles published in the Hazard/Risk Assessment section of *Environmental Toxicology and Chemistry* from January 1996 through March 2001 that employed a species sensitivity distribution (SSD) approach.
Articles listed chronologically. L(E)C50=concentration causing lethality/effect in 50% of test population; NOEC=no observed effect concentration; MAC=minimum algistatic concentration; MATC=maximum allowable toxicant concentration; CL=confidence limits

Reference	Source of species data	Endpoints used	Chosen Protection Level and Confidence Limits	Number of species in distribution	Definition of distribution
Klaine *et al.* 1996	Literature values for freshwater and marine species combined	Acute L(E)C50s	10%; no CL	14	Log-normal
Solomon *et al.* 1996	Literature values for freshwater and marine species combined	Acute L(E)C50s and chronic NOECs/MACs/MATCs analyzed separately	10%; no CL	Not stated; ≥ 10	Log-normal
Hall *et al.* 1998	Databases of various marine and freshwater species analyzed separately	Acutes and chronics separately; mixture of endpoints	10%; no CL	4-88	Log-normal
Cardwell *et al.* 1999	Computerized literature search; USEPA water quality criterion document for marine species	Acutes and chronics separately; most chronics predicted from acutes using an ACR=14.7; mixture of endpoints (most sensitive if >1 available)	None defined; centiles compared to give relative risk	Not stated; ≥ 8	Log-logistic
Suter *et al.* 1999	Specific tests carried out for the study plus literature values; only fish and daphnids included	Acutes and chronics separately; chronics included a mixture of endpoints	None defined; no CL	Not stated; appears to be 4-12	Empirical
Jones *et al.* 1999	Literature values for marine, estuarine and freshwater benthic species combined	Thresholds for lethality and community properties (*e.g.*, taxa richness and abundance) analyzed separately	< 20% defined as negligible risk; > 20% defined as marginal risk; > 50% defined as significant risk; no CL	Not stated	Empirical
Versteeg *et al.* 1999	Databases of a variety of freshwater species	Chronic data; mainly NOECs for various endpoints; some EC20, MATC and LOEC values used	5% with 95% CL	6-25	Log-logistic

Steen *et al.* 1999	Literature values for freshwater and marine species combined	Acute L(E)C50s and chronic NOECs analyzed separately	5% with both 50% and 95% CL	8-45	Log-logistic
Van de Plassche *et al.* 1999	Databases of a variety of freshwater and marine species combined	Acute L(E)C50s and chronic NOEC (for most sensitive endpoint per species) analyzed separately	5% with both 50% and 95% CL	≥ 5	Log-logistic
Newman *et al.* 2000	Literature, database and unpublished values for various freshwater and marine species analyzed separately	Acute L(E)C50 and chronic NOECs analyzed separately	5% with 95% CL	20-91	Bootstrap
Campbell *et al.* 2000	Literature, regulatory studies, studies from within federal and state regulatory departments; freshwater and marine species combined but fish and invertebrates analyzed separately	Acute L(E)C50 values	10%, no CL	Not stated; > 20	Log-normal
Morton *et al.* 2000	Data on estuarine and marine species given in a Ph.D. thesis	Acute L(E)C50s	5% with 95% CL	12-21	Log-logistic (gave most conservative values, used for risk assessment); log-triangular and log-normal also calculated
Crommentuijn *et al.* 2000	Databases of a variety of freshwater and marine species, analyzed together and separately; terrestrial species analyzed separately	Chronic NOECs based on mixed endpoints (lowest endpoint per species used)	5%, CL not stated	13-87	Log-logistic
Solomon *et al.* 2001	Data from literature, USEPA Pesticide Toxicity Database, and supplied by pesticide registrants; freshwater and marine species analyzed together and separately; arthropods and fish analysed separately, where possible, or pooled	Acute L(E)C50s; if > 1 value per species used geometric mean	10%, no CL	6-64	Log-normal

distribution of the target group of species. Choice of distribution is important since the differences among the most commonly employed distributions are largely in their tails, and it is here where the critical effect concentration is estimated. Several studies have addressed this issue of distributional shape choice and have used various statistical fitting methods to assess which distribution(s) is(are) the most appropriate for fitting single species NOEC data (Versteeg *et al.* 1999; Newman *et al.* 2000). These analyses clearly indicate that no single distribution consistently provides the best fit to these kinds of data. In practice, this means that every data set should be tested to determine which distribution provides the best fit. Problems here are that the choice will generally be constrained to a few standard distributions for which goodness-of-fit tests are available, and rarely are the data sets large enough that much confidence can be placed in the goodness-of-fit results. Some users of the SSD approach have assumed a distribution without apparently attempting to test whether it provides a good fit to the data (Steen *et al.* 1999). Newman *et al.* (2000) argued that bootstrap estimation provides a partial answer to the problems of ambiguity in selecting a specific distribution and estimation of the approximate number of species needed to precisely estimate the effects threshold. However, they also noted that bootstrap estimation requires a random sample of species sensitivities from the universe of possible species sensitivities about which one wishes to draw conclusions (and hence does not avoid the need to address P1). Although bootstrapping may offer an attractive alternative to assuming a specific distribution, it increases the technical requirements (computer and software availability as well as expertise) necessary to apply the sensitivity distribution method. Another, possibly simpler alternative is the so-called empirical distribution function (Jones *et al.* 1999; Suter *et al.* 1999) in which effects data are ranked in ascending order and the centiles estimated as (100 x rank)/(n + 1) where n is the number of data points in the distribution (Klaine *et al.* 1996). This approach has the advantage of not requiring a standard distribution to be fit to the data points. However it can easily be shown that for sample sizes smaller than 19, the 5th centile of affected species cannot be estimated, and a sample size of at least 9 is needed to estimate the 10th centile (Table 2). Another potential disadvantage of this approach is that confidence limits around the chosen centile are not readily derived. It should be noted that linear regression may also be used to estimate centiles from the ranked effects data (Giesy *et al.* 1999), but this then requires assumptions to be made about the underlying shape of the SSD.

P6. The number of species used to fit the distribution is adequate. The assumption here is that the number of species used in the analysis will be sufficient to deliver Assumption P5.

HOW HAS THE METHOD BEEN USED IN PRACTICE?

To obtain an impression of the way the sensitivity distribution approach has been applied in practice, we reviewed the Hazard/Risk Assessment section of *Environmental Toxicology and Chemistry* from volume 15 (1996) through volume 20(3) (2001). This involved slightly more than 100 articles of which those shown in Table 1 employed a sensitivity distribution approach. The aim of this exercise has not been to carry out a comprehensive review, but to draw attention to problems by sampling

Table 2. Relationship between sample size and centile of the most sensitive species where centile is estimated as (100 x rank)/(n + 1) and n=sample size (Klaine *et al.* 1996).

When sample size is:	Most sensitive species (rank 1) will be at centile:
5	16.7
6	14.3
7	12.5
8	11.1
9	10
10	9.1
15	6.2
20	4.8
25	3.8
30	3.2

from a few influential studies. For each article we determined the source of the input data, endpoints used, the chosen protection level and its confidence limits, the sample size used as input, and the distribution that was assumed. It is possible to make the following generalizations:

1. Effects data were mainly taken from the literature and not from the target community/ecosystem of interest. In other words, none of the distributions was a community-SSD despite the fact that a number had site-specific communities as targets of the risk assessment. We also noted that the species were generally selected to represent major trophic- or taxonomic groups but that these were not represented in proportion to their abundance in actual systems. In some cases species from very different ecosystems (*e.g.*, freshwater and marine) were combined in a single distribution. This means that all violated Assumption P1. In one case the relative risks for very different regions (*i.e.*, Puget Sound, Narragansett Bay, and Lake Erie) and habitat types (*i.e.*, marinas, commercial harbors, shipyards, fish and shellfish habitats) were compared on the basis of the same SSD, despite the fact that the local species compositions would have differed greatly among sites (Cardwell *et al.* 1999). The SSD approach presumes that the sensitivity of a community depends on the sensitivity of the individual species of which it is composed. Because the species composition of a commercial harbor is likely to be very different than that of a fish and shellfish habitat, comparing the relative susceptibility of communities at such sites using the same literature-based SSD is questionable.

2. The effects represented in a single distribution were often based upon a variety of endpoints for different species going into the distribution. These are not all of the same ecological relevance given differences in the life cycles of the

species used (Forbes *et al.* 2001). This causes problems with regard to Assumption P2.

3. A variety of distributions was employed and not always with rigorous justification, so raising questions with regard to Assumption P5.

4. The percentage of species chosen for protection was often 5 to 10%, although in some studies no 'acceptable' percentage was defined, but rather centiles were used to compare relative risks. In one case effects on up to 20% of species were considered as 'negligible risk' (Jones *et al.* 1999). Solomon *et al.* (2001) stated that any centile could in principle be used, 'provided that this measure can be validated against a knowledge and understanding of ecosystem structure and function, or calibrated in tests conducted in microcosms or in the field'. Because the species going into the SSD mostly have been taken from the literature and not from an intact ecosystem, the relation between any distribution-derived protection level and ecosystem structure/function is arbitrary, and this raises questions under Assumption P3. We recognize that the risk quotients are based on the same data set of test species, but here the lack of precision in interpretation of the RQ is more obvious. Using an SSD to generate a precise probability statement that a certain fraction of species is likely to be affected by a chemical is only helpful to the risk assessment process to the extent that it is also accurate.

5. Confidence limits around the effects threshold either could not be defined, were not defined, or if defined were specified somewhat arbitrarily. So this raises doubts about whether Assumption P4 was satisfied and even if it was, given the importance of this assumption on the outcome, it was rarely justified clearly enough.

6. The number of species used was generally greater than five and sometimes more than 50, but data on them were usually taken from databases and this raises questions of relevance as already discussed above.

7. In addition, in at least one instance, application factors were used to convert acute effects endpoints to chronic values for input into the distribution, which introduces the sources of uncertainty associated with the RQ approach into the SSD.

RELATIVE MERITS OF THE TWO APPROACHES AS THEY HAVE BEEN USED

The ultimate arbiter of the two approaches would be in terms of which leads to better protection, that is, more accurate risk assessments, but it is not possible to make this comparison because neither of the approaches has been evaluated sufficiently for existing ecosystems. Both the RQ approach and the SSD approach may give an incorrect estimate of risk, for different reasons and to different and/or inconsistent degrees. Given the difficulty in determining which approach gives an

answer closer to reality, a precautionary approach would be to use whichever method leads to more conservative protection levels (*i.e.*, lower effects thresholds). Alternatively, it has been recommended that risk quotients remain in use as screening tools (in Tier 1) and SSDs be employed at higher tiers in the risk assessment process (ECOFRAM 1999). A common feature of all tiered regulatory processes is a progression beginning with conservative assumptions and moving toward more realistic estimates. In such a progression, the first tier is designed to be protective rather than predictive (ECOFRAM 1999). Thus, if RQs are to be used at tier 1 and SSDs at a higher tier it is clearly essential that RQs should always yield lower effects thresholds than corresponding SSDs. Otherwise chemicals could potentially pass the first tier but fail higher tiers in the risk assessment process, and this would defeat the purpose of the tiered approach. Previous reviews have indicated that if both RQ and SSD techniques are applied to the same data set, which of the two approaches gives a lower effects threshold depends on: (a) whether the RQ uses acute data divided by a large factor or chronic data divided by a small factor, (b) the range of interspecies variability for the chemical of interest, (c) the level of protection that is chosen for the SSD, and (d) the chosen confidence limits for the SSD. When such comparisons have been made, differences in protection levels between techniques can go either way and are often not very great (OECD 1992; Calabrese and Baldwin 1993; Forbes and Forbes 1993). This led Calabrese and Baldwin (1993) to recommend the use of SSDs, whereas Forbes and Forbes (1993) and Zeeman (1995) came to the opposite conclusion.

The RQ approach has been criticized because it does not use all of the available toxicity data, and since only the most sensitive endpoint is used in its derivation, the only consequence of including more data from more tests is that the effects threshold can be reduced. The SSD approach, on the other hand, uses all available data, not just the lowest value. This means that as data accumulate the protection level should become more precise and not necessarily lower. However, we noted above that when species sensitivities are empirically fit to a cumulative distribution and rank centiles calculated, a sample size of at least 19 is necessary to estimate the 5th centile and a sample size of at least 9 is necessary to estimate the 10th centile (Table 2). In both cases these centiles are given by the most sensitive species in the distribution. In North America it is common practice to compare the 10^{th} centile of effects data with the 90th centile of exposure data to quantify risk (SETAC 1994; Solomon *et al.* 1996). Unless sample sizes of effects data are substantially greater than 19 the outcome of the risk assessment will be largely determined by the one or few most sensitive species, and despite claims to the contrary, the remaining data in the sensitivity distribution have very little influence on the outcome of the risk assessment. For example, for sample sizes of 30, the third most sensitive species will fall at the 10th centile and for sample sizes of 50 the 5th most sensitive species will fall at the 10th centile. This means that the range in sensitivity of the remaining 27 and 45 species, respectively, does not necessarily influence the outcome of the risk assessment.

Another argument in favor of the SSD approach is that it defines levels of protection in terms of likelihood of a certain percentage of species being protected. This is more quantitative than the RQ approach, in which the risk (hazard) quotient makes no explicit connection with species protected. Apart from the problem of

deciding which percentage of species to protect (see above) the quantitative estimate of risk offered by the SSD approach is only an improvement if it provides an accurate representation of the target of the risk assessment. There is little reason to expect haphazard collections of literature data to accurately reflect the percentage of species at risk in actual communities in nature. The concern here is that we may be led to generate very precise but inaccurate risk probabilities that impede rather than improve the risk assessment process.

It has been argued that the SSD approach offers a clear advantage over the RQ method in providing 'an explicit, objective and quantitative basis for dealing with uncertainty' (Hart 1999). Although in principle, the approach has the potential to do all of these, unfortunately our survey of the literature indicates that it often has not. Many users of the SSD approach have not been explicit, but have obscured several sources of uncertainty during estimation of the distribution and its subsequent interpretation. The SSD approach could offer a more (if not entirely) objective approach to risk assessment if decisions about which species and endpoints to include, the shape of the distribution to use, the level of protection and its associated confidence limits are supported by observations on actual risk assessment targets. Because this remains to be done, all of these decisions have been at least partly subjective. The SSD approach does provide a quantitative estimate of uncertainty, but such an estimate is only helpful for risk assessment to the extent that it is accurate and can be interpreted (*i.e.*, what does it mean for an ecosystem if 5% of the species NOECs are exceeded?). At present there are question marks over the extent to which the distributions generated by the SSD approach capture the uncertainty accurately and are relevant for the communities/ecosystems they are intended to represent.

IMPROVING THE SSD APPROACH IN PRACTICE

Here we offer some practical suggestions for improving implementation of the SSD approach to enhance transparency and reduce uncertainty in risk assessment.

P1. The simplest and most transparent way to address the issue of appropriate sampling is to ensure that the risk assessment targets are defined such that they more accurately reflect the data on which they are based. For example, if an SSD is generated from a literature survey of available aquatic toxicity data then the appropriate target of the risk assessment is 'all of the aquatic species for which data could be collected' and not 'the universe of species' or 'the community of aquatic species in lake X'.

Since an important use of SSDs is in comparative risk assessment (*e.g.*, comparing risks of the same chemical among different sites, or comparing the relative risks among different chemicals for the same kinds of communities) an important priority for research should be an in-depth analysis of the consequences of species input selection on the resulting SSD. How different are SSDs that accurately reflect the taxonomic and trophic composition of specific community types compared to those based on the distribution that happens to be available in a collection of literature data? When comparing risks among chemicals is it important that the SSDs to be compared are based on identical sets of taxa? Moreover, it is currently

an open question as to which taxonomic level is the most appropriate for the analysis.

As a first step toward exploring these issues we examined the trophic composition of 11 SSDs published by Versteeg *et al.* (1999). We chose this article because (1) the authors provided the raw data used as input in tabular form, (2) they selected their input values from existing data sets, (3) they used genus geometric mean values that reduced bias in favor of standard test species (*e.g.*, *Daphnia*), and (4) they included a range of chemical classes in their analyses. We assigned each genus in the data sets a trophic code. In principle this should be 1 for primary producers, 2 for herbivores, and 3 for carnivores. However, as a first approximation we allocated all invertebrates to level 2 and all fish to level 3. There was a large amount of variability in trophic composition of the data sets for the different chemicals, but average values across chemicals were 27.5% primary producers, 34.7% invertebrates, and 37.8% fish (Table 3). In contrast, the ecological literature suggests that most food chains are between three and four levels and that the number of species in a wide variety of aquatic community types increases by a factor of between two and three at each decline in trophic level from primary producers (Rosenzweig 1995; Sand-Jensen 2000). Assuming a simple three-level food chain and taking a value of 2.5 for the factor change in species number between trophic levels, then starting with 100 species in level one, there should be 40 in level 2 and 16 in level 3. Thus 64% of the species will occupy the first trophic level, 26% will occupy the second level, and 10% will occupy the third level. This will be independent of total species number and might be used as a general ecological rule for species distributions amongst trophic levels (cf. the situation may be somewhat more complicated in terrestrial systems [Rosenzweig 1995]). Using the simple ecological rule for aquatic systems we com-

Table 3. Trophic composition of taxa within data sets from Versteeg *et al.* (1999) as compared with the composition expected for aquatic communities. First trophic level=primary producers, 2nd trophic level = invertebrates, 3rd trophic level = fish. See text for details.

Chemical	Number of species	% 1st trophic level	% 2nd trophic level	%3rd trophic level
Cd	22	36.4	13.6	50.0
Cu	25	44.0	28.0	28.0
Zn	7	0	42.9	57.1
LAS	16*	37.5	31.2	31.3
Ammonia	11	18.2	18.2	63.6
Atrazine	17	70.6	5.9	23.5
Chlorine	6	16.7	66.6	16.7
3,4 DCA	12	16.6	41.7	41.7
C_{12} TMAC	8	37.5	50.0	12.5
Lindane	10	0	50.0	50.0
Phenol	12	25.0	33.3	41.7
Average	13.3	27.5	34.7	37.8
Range	6-25	0-70.6	5.9-66.6	12.5-63.6
Ecologically realistic		64	26	10

*a single bacterial species was omitted from the analysis

pare these percentages to those in the SSDs shown in Table 3. It is clear that most of the data sets underrepresented primary producers and overrepresented fish.

To explore the implications of such ecologically unrealistic sampling on the effects assessment, we estimated the concentrations affecting 5% of the species in the distribution (HC5) for both log-logistic (Aldenberg and Slob 1991) and log-normal models (Wagner and Løkke 1991) (both with a 95% confidence level). We also estimated the centiles of the ranked data (Table 4). We then chose two of the larger data sets (LAS and Cd) and randomly resampled from the original data (with replacement) to generate new distributions that conformed with the ecological rule specified above. The results of these analyses are shown in Table 5. For Cd, the HC5 from the ecologically realistic distribution of trophic levels was a factor 3.4 lower than the HC5 for the haphazard collection of species whereas for LAS it was a factor of 3.2 higher. We believe that this kind of result with changes in sensitivity in both directions is to be expected because the original data being used in the SSDs are taken from available databases. They may therefore either over- or underestimate the sensitivity of more ecologically realistic distributions of species to a degree that is not possible to specify *a priori*.

It should also be noted that in this analysis we have adjusted distributions for trophic composition and not given any consideration to the relevance of the taxa to the ecological situations under consideration. This could very well exacerbate the problems illustrated by the readjustment of the HC5s on the basis of trophic composition.

Table 4. Probabilistic effects thresholds for 11 chemicals derived from the original data in Versteeg *et al.* (1999). HC5 values are the lower 95% confidence limits, assuming either log-logistic or log-normal distributions. 5^{th} and 10^{th} centiles were estimated as 100 x rank/(n + 1) where n= number of taxa in the distribution. The 5^{th} and 10^{th} centiles were estimated by linear interpolation between the ranks that were just above and just below these centiles. The dashes in the final column indicate that the most sensitive species in the sample had a centile greater than the 5^{th} or 10^{th} centile. See also Table 2.

Chemical	Log-logistic HC5	Log-normal HC5	5^{th} and 10^{th} centiles
3,4 DCA	0.017	0.024	-, -
Ammonia	0.194	0.264	-, 2.66
Atrazine	0.186	0.223	-, 1.22
C_{12} TMAC	6.180	7.44	-, -
Cd	0.050	0.061	0.13, 1.42
Chlorine	0.075	0.099	-, -
Cu	0.254	0.308	2.99, 4.69
LAS	80.9	92.4	-, 198.61
Lindane	0.010	0.014	-, 0.22
Phenol	6.66	8.70	-, 242.66
Zn	0.48	0.63	-, -

Table 5. Comparison of HC5 values (with 95% confidence limits) calculated for the original data in Versteeg *et al.* (1999), assuming a log-normal distribution and for a distribution adjusted to reflect a more realistic trophic composition as defined in the text. The final column gives the extent to which the adjusted HC5 differs from the original HC5.

Chemical	HC5 for original data	HC5 adjusted	Factor difference
Cd	0.061 μg/L	0.018 μg/L	3.4 times lower
LAS	81 μg/L	256 μg/L	3.2 times higher

P2. Endpoints for input into SSDs should be more rigorously chosen. When the assessment endpoint is the persistence of species populations, individual-level endpoints measured in ecotoxicological tests should be translated to likely population-level impacts and the resulting values used as input into the SSD. A variety of population models, from very simple to very complex, may be used for this purpose (Calow *et al.* 1997; Akçakaya *et al.* 1999; Caswell 2001).

P3. Decisions about the fraction of species to protect should be based on considerations of the total number of species per taxonomic and/or functional group in the target system. Systems with fewer species are likely to have less functional redundancy, and this may argue for protecting a higher fraction of them. The identity of the most sensitive species in the SSD should also be considered, and if there are indications that the left tail is biased toward certain taxa or functional groups, the 'acceptable' effects level could be adjusted downwards if necessary. In any case, justification for the chosen protection level should be stated clearly and the uncertainty associated with the decision articulated.

P4. Although confidence limits can provide a useful measure of the statistical uncertainty in a sample parameter, there remain substantial biological uncertainties that may influence the distribution of species sensitivities, but not be reflected in the calculated parameters of an SSD. Therefore interpreting confidence limits for SSDs requires great care. Comparison of the HC5 values (that assume a 95% confidence level) in Table 4 with estimated 5th centiles show that the latter (for those data sets that they could be calculated) are about an order of magnitude greater than the former.

P5. As a general rule, data should not be assumed to follow a particular distribution without explicit justification, for example, provided by a goodness-of-fit test. However, recognizing that the number of data points used to generate SSDs may be too small to place much confidence in lack-of-fit statistics, empirical distribution curves may be a more transparent alternative, provided that the original data points (and not just the fitted curve) are presented.

P6. Debate as to the minimum number of data points necessary to generate reasonable SSDs is ongoing (Newman *et al.* 2000). Unfortunately, it seems that the issue of quantity has taken precedence over the issue of quality (*i.e.*, appropriate endpoints measured in an appropriate sample of species), which is equally if not more important to reducing the uncertainty in effects assessment. The effect of adding more species to an SSD may depend on whether the added species represent

a new taxonomic or functional group. It is generally assumed that members of the same taxonomic group are more similar to each other in sensitivity than to members of other taxonomic groups. This assumption can be assessed by examining the relationship between the number of taxa used to generate an SSD and its associated spread. Taking data from Crommentujn *et al.* (2000) on chronic NOECs for four metals (Cd, Cu, Pb, Zn) measured in freshwater, marine, and terrestrial species, we have plotted the spread in each SSD (4 metals x 3 community types = 12 points) as a function of the number of taxa going into each distribution (Figure 1). There is a significant positive correlation ($P = 0.007$, $n=12$, $r^2=0.53$) indicating, not surprisingly, that the more taxa going into an SSD the wider the distribution of species sensitivities.

If the species selected for input into an SSD are truly an unbiased random sample of the distribution of target species, then as the sample size increases the spread of the resulting SSD should decrease. However, if the additional species are intentionally chosen to represent new taxonomic or functional groups, then it is likely that the spread of the SSD will increase as more species are added. This is an important consideration given that the application of SSDs is meant to reduce uncertainty in the risk assessment process.

QUO VADIS?

We have deliberately contrasted the RQ and SSD approaches as alternatives because this is often the way they have been portrayed in the literature. However,

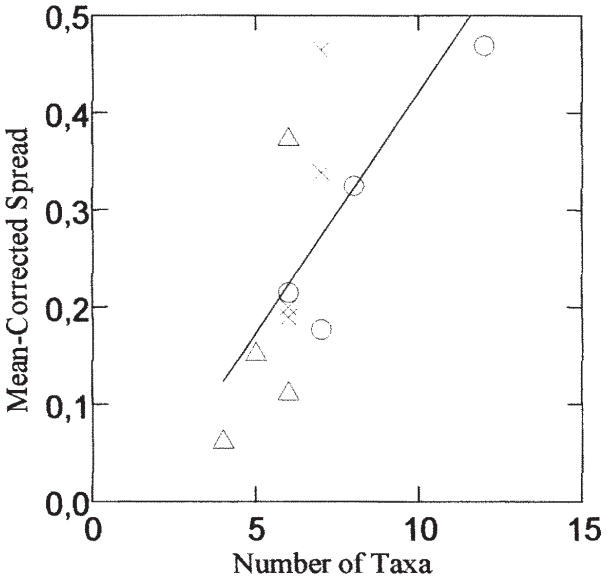

Figure 1. Number of taxa used to generate SSDs for freshwater (circles), marine (X) and terrestrial (triangles) species by Crommentujn *et al.* (2000; Table 1). The Mean-Corrected Spread was obtained by dividing the distribution width parameter (β) by the mean of the SSD (α) to account for differences in toxicity among metals.

this need not be the case. In a tiered risk assessment it seems entirely sensible that a less data-hungry approach be used as an initial screen involving worst-case assumptions and risk quotients, followed by a more refined approach based upon SSD (as has been recommended in ECOFRAM [1999]) for pesticide risk assessment).

Indeed, as indicated in the introduction, were we able to define species sensitivity distributions for ecological targets of interest, then by definition this would clearly be a better way of carrying out an ecological risk assessment (subject to the concerns raised in the Assumptions T1 to T3). Nevertheless, we have wanted to draw attention to three issues in the foregoing in the context of the rising popularity of SSDs in the regulatory and scientific communities. First, if used uncritically (without careful attention to Assumptions P1 to P6) the SSD approach can potentially lead to conclusions based on foundations no more substantial than the RQ approach. Second, and related, the Assumptions P1 to P6 (and their violations) are often obscured by the apparent sophistication of the models and their outputs. This can lead to less, not more, transparency than provided by the RQ approach. Third, given the greater data requirements of constructing an SSD, it is likely that risk quotients will remain an important screening tool at the first tier of the risk assessment with SSDs being used at higher tiers. Comparisons of the two methods show that neither method provides consistently lower risk estimates (OECD 1992; Forbes and Forbes 1993). This suggests that either the application factors used to generate risk quotients are not large enough for screening purposes or that the SSDs are overestimating the variability in sensitivity among species. Moreover, Versteeg *et al.* (1999) showed that between 10 and 52% of the genera in literature-derived SSDs were predicted to be affected at concentrations corresponding to mesocosm NOECs indicating that SSDs generated in this way may often overestimate risk relative to intact systems.

If the SSD approach is to lead to improved (*i.e.*, more accurate) risk assessments, it is essential that the distributions generated accurately reflect the variability among species that are the target of the risk assessment, whether this be a generalized type of ecosystem or a community at a specific site. For existing and new chemicals legislation it is often necessary to carry out risk assessments in terms of generalized ecosystems. Is it possible to approach these by constructing a sensitivity distribution for a generalized community? In principle this would need to be done in terms of the distribution of taxonomic, functional, and life-cycle types in nature across the region(s) concerned. It would seem unlikely that the arbitrary distributions constructed from available databases would be satisfactory in these terms, but this should be explored. For site-specific risk assessments, sensitivity distributions should be generated for communities in the specific receiving ecosystems. In other words, the sensitivity distribution would need to reflect the species from the natural systems rather than be based on a haphazard collection of literature data. At the very least this means that the taxonomic and trophic composition of the target community should be adequately represented. One way that this might be approached would be through assigning weights to the input values to adjust the distribution to reflect the relative abundance of different taxa.

The overriding conclusion is that if we want to move forward from risk/hazard quotients in a substantiated rather than just pragmatic way we need to know more about ecological organization and the sensitivity distributions that go with it. Both

the limitations and the challenges of the SSD approach need careful consideration. The sensitivity distribution approach should be encouraging us to think more carefully about how to sample species from ecological systems of interest. Faced with the need to determine the sensitivity distribution for an ecological system consisting of a large number of species requires that some decision is taken on the relative importance of species and hence how to characterize their sensitivity. Relative importance might be considered in terms of numbers, biomass, or functional roles within the system, and this therefore raises fundamental questions that are rarely, if ever, addressed in routine ecological risk assessments. Although probabilistic methods such as SSDs have the potential to lead to better risk assessments than simple risk quotients, they will not do so unless they are rigorously and transparently applied and interpreted. Using the sensitivity distribution approach in a measured way, that carefully makes explicit and takes account of the assumptions associated with it (T1-T3 and P1-P6), therefore ought to take us forward to more relevant assessment of risks arising from human activities for ecological systems.

REFERENCES

Akçakaya HR, Burgman MA, and Ginzburg LR. 1999. Applied Population Ecology, 2nd ed. Sinauer Associates, Inc., Sunderland, MA, USA

Aldenberg T and Slob W. 1991. Confidence Limits for Hazardous Concentrations Based on Logistically Distributed NOEC Toxicity Data. Report No. 719102002. National Institute of Public Health and Environmental Protection (RIVM), Bilthoven, The Netherlands

Cairns J and Pratt JR. 1995. The relationship between ecosystem health and the delivery of ecosystem services. In: Rapport DJ, Gaudet CL, and Calow P (eds), Evaluating and Monitoring the Health of Large-Scale Ecosystems, pp 63-76. Springer-Verlag, Berlin, Germany

Calabrese EJ and Baldwin LA.1993. Performing Ecological Risk Assessments. CRC Press, Lewis Publishers, Boca Raton, FL, USA

Calow P and Forbes VE. 1999. Manual on Risk Assessment of Subregional Sea Areas: A Practical Guide for Tropical Ecosystems. MPP-EAS Technical Report 21, GEF/UNDP/IMO. Regional Programme for the Prevention and Management of Marine Pollution in the East Asian Seas, Quezon City, Philippines

Calow P, Sibly RM, and Forbes VE. 1997. Risk assessment on the basis of simplified life-cycle scenarios. Environ Toxicol Chem 16:1983-9

Campbell KR, Bartell SM, and Shaw JL. 2000. Characterizing aquatic ecological risks from pesticides using a diquat dibromide case study. II. Approaches using quotients and distributions. Environ Toxicol Chem 19:760-774

Cardwell RD, Brancato MS, Toll J, *et al.* 1999. Aquatic ecological risks posed by tributyltin in United States surface waters: pre-1989 to 1996 data. Environ Toxicol Chem 18:567-77

Caswell H. 2001. Matrix Population Models: Construction, Analysis, and Interpretation, 2nd ed, Sinauer Associates, Inc., Sunderland, MA, USA

CEC (Commission of the European Communities). 1996. Technical Guidance Documents in Support of the Commission Directive 93/67 EEC on Risk Assessment for New Substances and the Commission Regulation (EC) No 1488/94 on Risk Assessment for Existing Substances. Parts I through IV. Commission of the European Communities, Brussels, Belgium

CEC (Commission of the European Communities). 1997. Guidance Document on Aquatic Ecotoxicology in the Frame of the Directive 91/414/EEC; 8075/VI/97. Commission of the European Communities, Brussels, Belgium

Crommentuijn T, Polder M, Sijm D, *et al.* 2000. Evaluation of the Dutch environmental risk limits for metals by application of the added risk approach. Environ Toxicol Chem 19:1692-1701

ECOFRAM. 1999. ECOFRAM Aquatic and Terrestrial Final Draft Reports. USEPA, June 1, 1999. Available at http://www.epa.gov/oppefed1/ecorisk/index.htm

Forbes VE, Calow P, and Sibly RM. 2001. Are current species extrapolation models a good basis for ecological risk assessment? Environ Toxicol Chem 20:442-7

Forbes TL and Forbes VE. 1993. A critique of the use of distribution-based extrapolation models in ecotoxicology. Funct Ecol 7:249-54

Giesy JP, Solomon KR, Coats JR, *et al.* 1999. Chlorpyrifos: ecological risk assessment in North American aquatic environments. Rev Environ Contam Toxicol 160:1-129

Hall LWJr, Scott MC, and Killen WD. 1998. Ecological risk assessment of copper and cadmium in surface waters of Chesapeake Bay watershed. Environ Toxicol Chem 17:1172-89

Hart ADM. 1999. Options for testing and risk assessment. In: Assessing Risks to Non-target Terrestrial Plants, Commission N. PN0923. Ministry of Agriculture Fishery and Food of England and Wales (MAFF), Pesticides Safety Directorate, London, UK

Jones DS, Barnthouse LW, Suter GWII, *et al.* 1999. Ecological risk assessment in a large river reservoir: 3. Benthic invertebrates. Environ Toxicol Chem 18:599-609

Klaine SJ, Cobb GP, Dickerson RL, *et al.* 1996. An ecological risk assessment for the use of the biocide, dibromonitrilopropionamide (DBNPA) in industrial cooling systems. Environ Toxicol Chem 15:21-30

Lawton JH and Brown VK. 1993. Redundancy in ecosystems. In: Schultze E-D and Money HA (eds), Biodiversity and Ecosystem Function, pp 255-70. Springer-Verlag, Berlin, Germany

Morton MG, Dickson KL, Waller WT, *et al.* 2000. Methodology for the evaluation of cumulative episodic exposure to chemical stressors in aquatic risk assessment. Environ Toxicol Chem 19:1213-21

Newman MC, Ownby DR, MÈzin CA, *et al.* 2000. Applying species-sensitivity distributions in ecological risk assessment: Assumptions of distribution type and sufficient numbers of species. Environ Toxicol Chem 19:508-15

OECD (Organisation for Economic Cooperation and Development). 1992. Report of the OECD Workshop on the Extrapolation of Laboratory Aquatic Toxicity Data to the Real Environment. Environment Monograph No. 59. Paris, France

Paine RT. 1966. Food web complexity and species diversity. Am Nat 100:65-75

Pratt JR and Cairns JJr. 1996. Ecotoxicology and the redundancy problem: understanding effects on community structure and function. In: Newman MC and Jagoe CH (eds), Ecotoxicology: A Hierarchical Treatment, pp 347-70. CRC Press, Boca Raton, FL, USA

Rosenzweig ML. 1995. Species Diversity in Space and Time. Cambridge University Press, Cambridge, UK

Sand-Jensen K. 2000. Økologi og Biodiversitet. G.E.C. Gads, Copenhagen, Denmark

SETAC (Society of Environmental Toxicology and Chemistry). 1994. Final Report: Aquatic Risk Assessment and Mitigation Dialogue Group. Pensacola, FL, USA

Solomon KR, Baker DB, Richards RP, *et al.* 1996. Ecological risk assessment of atrazine in North American surface waters. Environ Toxicol Chem 15:31-76

Solomon KR, Giddings JM, and Maund SJ. 2001. Probabilistic risk assessment of cotton pyrethroids: I. Distributional analyses of laboratory aquatic toxicity data. Environ Toxicol Chem 20:652-659

Steen RJCA, Leonards PEG, Brinkman UATh, *et al.* 1999. Ecological risk assessment of agrochemicals in European estuaries. Environ Toxicol Chem 18:1574-81

Stephan CE, Mount DI, Hansen DJ, *et al.* 1985. Guidelines for deriving numerical national water quality criteria for the protection of aquatic organisms and their uses. PB85-227049. U.S. Environmental Protection Agency, Washington, DC, USA

Suter GW, II, Barnthouse LW, Efroymson RA, *et al.* 1999. Ecological risk assessment in a large river-reservoir: 2. Fish community. Environ Toxicol Chem 18:589-98

USEPA (U.S. Environmental Protection Agency). 1992. Framework for Ecological Risk Assessment. EPA/630-R-92/001. Risk Assessment Forum, Washington, DC, USA

Van de Plassche EJ, de Bruijn JHM, Stephenson RR, *et al.* 1999. Predicted no-effect concentrations and risk characterization of four surfactants: linear alkyl benzene sulfonate, alcohol ethoxylates, alcohol ethoxylated sulfates, and soap. Environ Toxicol Chem 18:2653-63

Versteeg DJ, Belanger SE, and Carr GJ. 1999. Understanding single-species and model ecosystem sensitivity: Data-based comparison. Environ Toxicol Chem 18:1329-46

Wagner C and Løkke H. 1991. Estimation of ecotoxicological protection levels from NOEC toxicity data. Wat Res 25:1237-42

Zeeman MG. 1995. Ecotoxicity testing and estimation methods developed under Section 5 of the Toxic Substances Control Act (TSCA). In: Rand GM (ed), Fundamentals of Aquatic Toxicology, 2nd ed, pp 703-16. Taylor & Francis, Philadelphia, PA, USA

Zeeman MG and Gilford J. 1993. Ecological hazard evaluation and risk assessment under EPA's Toxic Substances Control Act: An introduction. In: Landis WG, Hughes JS, and Lewis MA (eds), Environmental Toxicology and Risk Assessment, pp 7-21. ASTM (American Society for Testing and Materials), Philadelphia, PA, US