LIFE EXPECTANCY CHARGE QUESTIONS

Charge Questions 1 & 2:[1]
What is the most appropriate methodology to use when valuing changes in mortality risk for persons with different remaining life expectancies? Is it appropriate to use a standard VSL to value reductions in mortality risk when information on remaining life expectancy is not available?

It is anticipated that EPA will need to issue rules affecting persons who differ in their remaining life expectancies in a relatively short time-frame. What does existing research imply about approaches to valuing mortality risk when people have life expectancies of varying lengths? How applicable and relevant is the existing literature and how does the existing theoretical and empirical literature inform these issues?

Economic theory dictates that the appropriate measure of the value in a reduction in mortality risk (e.g., in the probability of dying over a stated period) is what a person is willing to pay for it. This amount may be affected by a person's remaining life expectancy, but theory (e.g., the lifecycle consumption model with uncertain lifetime) has little to say about the relationship between willingness to pay (WTP) and remaining life expectancy. Only in very special cases can it be said that WTP should be an increasing function of remaining life expectancy.[2]

The relationship between WTP for mortality risk changes and remaining life expectancy is, therefore, an empirical matter. Unfortunately, this relationship is difficult to measure since remaining life expectancy is not observable while an individual is still alive. Individuals could be asked in a stated preference study what they would pay to reduce their probability of dying, *assuming different life expectancies*. However, this is a difficult question. In revealed and stated preference studies all that can be observed *ex ante* are correlates of life expectancy: viz., age and health status. So, one could try to measure how WTP varies with age and health status.

This suggests that EPA may want to allow WTP to vary with age and health status. It is, however, the committee's judgment that the empirical literature is not advanced enough at present to provide clear guidance on how age and health status affect WTP for changes in mortality risk. This suggests that EPA should, at present, use the standard VSL to value mortality risk reductions according to the conventional paradigm.

[1] These questions have been reworded to better reflect the Agency's intent, as discussed during the September 15 meeting.
[2] For example, in a lifecycle model with uncertain lifetime and perfect annuities markets, WTP for a change in the conditional probability of dying at any age will be an increasing function of discounted remaining life expectancy if an individual consumes at a constant rate throughout his lifetime and has a period utility function that is independent of age.

Charge Question 3:
Are there other areas of the literature that should be examined and how would they inform this issue in the short tem (i.e., less than 6 months)?

The committee agrees that the White Paper by Dockins, Maguire, and Simon covered the appropriate literature.

Charge Question 4:
What type of long-term research can inform these issues?

The committee agrees that willingness to pay for risk reduction is likely to be affected by remaining life expectancy, which is related to both age and baseline health status. The existing evidence on these relationships is weak and occasionally contradictory. The committee recommends that additional research be funded to improve these estimates.

Charge Question 5:
What paradigms should be considered in valuing changes in mortality risk for person with different life expectancies? How will these paradigms inform us in the short term?

One paradigm that is commonly used to allow remaining life expectancy to affect the value of a reduction in mortality risk is the *Value of a Statistical Life Year (VSLY).* The VSLY assumes that the value of mortality risk reductions is proportional to remaining life expectancy (or discounted remaining life expectancy) and uses this assumption to calculate a value per life year saved. The VSL is then computed by multiplying the average remaining life expectancy of the population affected by a regulation by the value per life year saved. As noted above, this procedure is difficult to justify on either theoretical or empirical grounds, if the appropriate valuation concept is what a person would pay to reduce his own risk of dying.

An alternative approach to valuation is a social approach, in which people are asked about what tradeoffs they would make if allocating resources from a social perspective. For instance, a senior's decision about the value of a risk reduction for him/herself may differ from the value that he/she would like government agencies to use when deciding how to allocate their resources between risk reduction for seniors and risk reduction for youth; a youth similarly may implicitly reveal a different value for risk reduction for him/herself than for the social decision. The White Paper by Dockins, Maguire, and Simon discusses the difference between these two approaches as the difference between "public preferences" and "private preferences" (p. 6), and it notes that there is often a divergence between the two.

One justification for using the VSLY approach is that, by weighting the value of lives saved by remaining life expectancy, it may better reflect people's public preferences (citations). We therefore recommend that EPA present both VSL and VSLY estimates

regardless of whether information on the length of life expectancy changes is available. When information on the length of life expectancy changes is not available, EPA should develop VSLY estimates based on the average life expectancy of the affected population. The two estimates should be given equal prominence and emphasis in analysis and reports. In addition, we recommend that the EPA report the number of life-year and lives that would be saved under the regulations under consideration.

Charge Question 6:

More generally, based on the economics literature, under what conditions is it most important to provide information on life expectancy and baseline risks as part of an economic analysis of environmental policy? If the information cannot be incorporated directly into monetized benefits estimates, how might it best be provided as a supplemental analysis?

In general, the measure of benefits based on WTP should reflect the WTP of the population that is affected by the change. The central, default VSL of $6.1 million, which is based on 21 hedonic wage and 5 stated preference estimates, reflects the WTP of the population included in these studies. The wage studies obviously reflect a population of working individuals, which implies at least some minimal health status (i.e., healthy enough to work) and a specific age distribution (i.e., working age adults) with an associated average life expectancy (estimated to be 35 years). If the population most affected by an EPA regulation or policy change differs from the population represented in these studies, then the WTP estimates generated by these studies will be a biased estimated of the true WTP of the affected population. *Whenever* this is the case, it will be important to provide information on the life expectancy and baseline risk[3] of the affected population as part of an economic analysis of the policy.

Unfortunately, the current economics literature does not provide convincing evidence regarding the direction of the bias that would exist if the baseline risk and life expectancy of the affected population differ from those of the population included in the WTP studies, i.e., there is mixed evidence on whether increases in baseline risk or reductions in life expectancy increase or decrease WTP estimates (see discussion of other charge questions). Nonetheless, it is important to include these characteristics of the affected population for two reasons: (1) to highlight the potential for bias, even if it is not possible to predict its direction, and (2) to highlight the fact that the policy is likely to affect certain sub-populations disproportionately. While this latter information may not be formally incorporated into the benefit-cost analysis (e.g., by providing WTP estimates that are specific to affected sub-populations), it would provide the basis for an equity assessment, which in many cases is required by statute, executive order, or agency policy. Information about disproportionate impacts can be important input into policy decisions.

[3] Strictly speaking, the term baseline risk refers to the survival curves of the members of the population affected by the regulation.

The information about baseline risk and life expectancy is most useful if provided in the form of a distribution (rather than simply an average) across the affected population. This is particularly true when the distribution is bi-modal. A bi-modal distribution would exist, for example, in cases where the very young and the very old are susceptible to pollution effects.

META-ANALYSIS CHARGE QUESTIONS

Charge Question 1:
In light of the workgroup's findings, what approach or approaches are the most scientifically appropriate to derive summary estimates of mortality risk valuation for use in environmental policy analysis? Should meta-regression techniques be applied to selected estimates or are other methods (e.g., fitting distributions) more appropriate? Please specify which methods, aside from or in addition to meta-regression techniques the Agency should explore.

Meta analysis is a scientifically appropriate and useful method for deriving summary estimates of mortality risk valuation for use in environmental policy analysis. Meta analysis can be used to estimate a central tendency of a set of estimates, a range or probability distribution of estimates around the central tendency, and to estimate how the estimates depend on characteristics of the underlying studies, including both population characteristics and methodological choices. Fitting a distribution to a set of estimates can be viewed as a special case of meta analysis, in which study estimates are characterized as random draws from a common probability distribution (i.e., with mean and variance independent of study characteristics).

Other methods that can be used to summarize estimates of mortality valuation include literature review and judgment and expert elicitation. A literature review can qualitatively describe the strengths and weaknesses of alternative studies then offer a judgmental summary of the best estimate, best range of estimates, or even a probability distribution characterizing the authors' judgment about the best estimates. Several such reviews have been published over the years (e.g., Fisher et al., 1986; Viscusi, 1993; Viscusi and Aldy, 2003) and the committee does not recommend using resources to conduct another at this time.

Alternatively, judgments of several valuation experts can be obtained through a formal expert judgment study. In such a study, experts should be identified and selected through a scientifically defensible process. Each expert is asked (in a structured interview) to provide a probability distribution characterizing his or her subjective beliefs about the best estimate of the valuation measure, or multiple distributions corresponding to different valuation contexts if the expert believes valuation differs among these contexts. EPA has recently conducted an expert judgment study of the relationship between fine particulate exposure and mortality risk. We suggest EPA consider whether an expert judgment study of valuation would be worthwhile.

Additional Material Re: Charge Question 1:

The need to aggregate VSL estimates from different studies arises because the studies yield different results. This can occur for two reasons: (1) the studies use different techniques to measure the VSL—for example, hedonic wage (HW) studies use different sets of covariates and/or different functional forms; (2) the populations in the

different studies differ in characteristics that may affect the VSL—for example, age, gender, and health status. Given these sources of variation in VSL estimates meta-analysis could be used in several ways:

1. A meta-regression could be used as an empirical literature review, to summarize what appears to matter in estimating the VSL for a given population: for example, to examine the effect of controlling for the worker's industry or for non-fatal accident risk on the size of the VSL in a HW study.
2. A meta-regression could be used to estimate a reduced-form model relating the VSL to population characteristics.
3. A meta-analysis (not necessarily a meta-regression) could be used to pool VSL estimates from different studies, treating the VSL as a parameter. (The VSL is treated the same as an average treatment effect in a randomized trial.)

Are we saying that all 3 of these are valid uses of meta-analysis? How should EPA use meta-analysis if we say it is OK? Should they use a meta-regression and then set the value of covariates to the most "desirable" values?

Charge Question 2:
Using the approach identified above, what measures/estimates should be combined? VSL estimates? The coefficient on fatal risk? Other? How should the Agency select the measures to be combined? Should a single, preferred estimate be selected from each study or should all estimates be included?

The answer to this question depends on what type of meta-analysis we urge the agency to use in question 1. For example, if we take the third approach above we would urge the agency to select a VSL estimate and associated standard error from each study. There seems to be a consensus that only one estimate should be selected from a study that reports several models all estimated from the same dataset. When a study reports results for independent sub-samples, it may be appropriate to take a VSL for each sub-sample (e.g., a VSL for men and a VSL for women).

Charge Question 3:
Should original studies be required to use a common empirical specification (functional form and choice of covariates) in order to be included in a meta-analysis? What data are required of the original studies to be included?

The committee recommends that original studies report the results of a common specification of the compensating wage regression, in addition to the author(s)' own specifications. The compensating wage study report or article should also report the estimate of the VSL calculated by the authors and its standard error—the latter being essential for creating the weights to be deployed in the meta-regression—and ample details on how exactly both were calculated.

We recommend that the compensating wage studies ultimately to be included in a meta-regression

- provide information on the source of data on risk, worker pay and worker characteristics;
- include codes for creating the sample used for the compensating wage regressions and for transforming variables;
- report average risk and average pay for the sample;
- explain whether the author(s) did or did not include non-fatal risks in the compensating wage regression, in addition to the fatal risks variable
- explain whether the sample contains only union workers, or if union membership was controlled for in the regression
- explain whether high-risk workers (e.g., police officers, firefighters, etc.) are included or excluded from the sample
- explain clearly whether the researcher(s) included a quadratic term in risk, and interactions between risk and other variables, in the regression.

Charge Question 4:
Given the various approaches used in the literature, what is the most scientifically appropriate measure to derive when combining estimates from multiple studies? A single central point estimate, a single distribution, or a range of estimates in economic analyses? How can such a measure best reflect the uncertainty and variability in mortality risk valuation estimates?

Meta analysis should be used to provide a description of the probability distribution of the estimates. The resulting probability distribution can be used for uncertainty analysis, and the expected value and other relevant point estimates (e.g., median, 5th and 95th percentiles) can be drawn from it.

A meta-analytic regression model can describe how the expected value and variance of an estimate depend on study features that are represented by the independent variables (e.g., study population and methodological choices). For policy evaluation, one must determine which estimates of VSL are most relevant. To the extent that a meta analysis shows that VSL depends on population characteristics (i.e., VSL is variable in the population), the model should be used to predict the expected VSL for the population that is relevant to the economic evaluation (i.e., the meta analysis can be used to transfer estimates of VSL from a study population to the policy-relevant population). To the extent that the meta analysis provides evidence that estimated VSL depends on methodological choices, it is necessary to determine which methodological choices are most appropriate (a question of uncertainty), and then to use the model to adjust estimates from studies that used other methodological choices. Note that it is not necessary to determine that one methodological choice is uniquely best. Instead, one can construct a weighted average of the estimates of VSL that the meta analysis suggests would be obtained from alternative methodological choices.

Charge Question 5:
How should stated preference studies and revealed preference studies be considered together in a scientifically appropriate method to derive summary estimates of mortality risk valuation?

The committee agrees with the expert panel that stated and revealed preference studies should not be combined in the same meta-regression. If the meta-regression is designed to examine the impact of different methodological factors on the VSL the nature of these factors will differ qualitatively in HW v. SP studies. For example, the nature of the models used to estimate the VSL are very different in the two types of studies, as are the covariates used.

Charge Question 6:
How should the Agency use studies based on specific sub-samples (e.g., elderly) in developing summary estimates of mortality valuation estimates for environmental policy analysis?

The answer to this question depends on what the agency wishes to measure. If the goal is to develop an estimate of the VSL for the entire population, then the question is how to combine estimates from different sub-populations. This seems to me a very difficult question. HW studies, for example, clearly exclude retirees and may often exclude women. How the results of these studies should be combined with (e.g.) stated preference studies that target older people is not clear. Even if one had a set of HW studies for difference subpopulations (e.g.) men and women, I'm not sure how they would be combined except, perhaps, in a meta-regression that allowed population characteristics on the RHS. If EPA believes that it should have a VSL estimate for a population not covered by HW studies, perhaps it will have to rely on SP methods.

Charge Question 7:
Most studies that combine estimates adjust the data from the original studies to some extent. For example, some studies adjust for after-tax wages, whereas others do not. Is there a set of such modifications that the SAB-EEAC believes to be critical when deriving summary estimates from the literature? Are there some data modifications that are generally incompatible with a sound approach to synthesizing existing estimates? What are the implications for interpreting results?

In synthesizing estimates from multiple studies, it is important to adjust as well as possible for differences among the studies that have predictable effects on their results. Such adjustments can be made to studies before including their estimates in a meta analysis, or the meta analysis can be used to estimate the appropriate adjustment.

One adjustment that can and should be made is to adjust for monetary inflation between studies by converting all nominal monetary values into real values. Because there is uncertainty about the best estimate of inflation over a period (reflected, for example, in different measures such as the various consumer price indices and the GDP deflator), the best adjustment is not clear. However, if one relies only on relatively recent studies (that are likely to be most relevant to the evaluation of current policy), differences between the alternative indices are likely to be modest and contribute little to uncertainty about the appropriate valuation compared with other factors. Similarly, if valuation estimates using other currencies are included, it is necessary to adjust for the exchange rate (again, uncertainty about the best exchange rate is not likely to be a major concern).

Other adjustments that are in principle desirable are more difficult to make, and so the committee does not view these as critical. One is to adjust for differences in real income and wealth between study populations. Since the value of reducing mortality risk increases with income and wealth, differences in these factors are expected to yield differences in estimated valuation. However, the appropriate magnitude of adjustment is not clear, because of uncertainty about the value(s) of the income elasticity and very little empirical evidence concerning the relationship between wealth and mortality valuation.

A second potential adjustment is to convert all estimates into marginal changes in consumer income (net of taxes and benefits). In hedonic-wage studies, workers choices are in principle driven by comparing the total incremental compensation with the total incremental risk between jobs, where total compensation includes wages, health insurance, retirement income, compensation conditional on injury, and other benefits, all evaluated post tax. In stated-preference studies, respondents are likely to view payments as coming from post-tax income (in principle, respondents may be asked about payments that would be made using either pre- or post-tax income; this detail is usually not specified but may be inferred from question wording). Adjustment for these factors is difficult because of variation in marginal tax rates and benefit schedules across populations, and so the committee does not view it as critical, but suggests that research attention be directed toward determining whether such adjustments can be made.

Charge Question 8:
What reporting and other protocols should the researchers conducting the combination study follow? How should the analysis handle zero or negative mortality risk valuation estimates from studies that otherwise meet its selection criteria for inclusion?

The purpose of the study should be clearly defined. Is it to average over different methodologies (eg RP versus SP) ? Or over technical differences in how a given methodology is implemented? Or over different populations?

There should be an explicit description of the rules for the inclusion and exclusion of items in the combination study, and also of the search rule by which the candidate studies were identified in the first place.

There should be systematic coding of important features of the items included in the combination study, for example the metric in which wages are expressed, the metric for risk itself, the other variables included in models involving VSL, and the populations for which VSL was evaluated.

There are two reasons why it is legitimate, in general, to exclude studies that otherwise meet the selection criteria for inclusion but yield zero or negative estimates of VSL. One reason applies to both HW and CV; the other applies specifically to HW.

The main reason is that, in the case of VSL, the researcher is likely to have a very strong prior expectation of the correct sign for what is being estimated, namely that the VSL is positive – i.e. that in general people are neither indifferent to whether they live or die, nor do they prefer death to life. Imposing a prior belief as a filter for the acceptance of observational data is reasonable only to the extent that the strength of the prior belief seems unassailable, but that would seem to hold in the particular case of VSL. Hence, it seems reasonable to conclude that, if a particular empirical study reveals a zero or negative estimate for VSL, this is not an accurate indication of the preferences held by the particular individuals who were the subject of that study but, rather, that it reflects some flaw or anomaly in the study itself.

There is an additional consideration in the case of an HW study that reveals a zero or negative derivative of wage with respect to risk. An HW study does not directly measure preference (the demand to avoid risk, the willingness to accept to endure it). Instead, it measures supply – the opportunity cost of avoiding risk that confronts a worker choosing an occupation. The conventional assumption is that there is an interior solution to the worker's utility maximization decision, as a result of which the worker's marginal rate of substitution between income and risk is equated to the marginal cost of income in terms of increased risk. If an HW study reveals that the marginal cost of avoiding risk is zero or negative and it is believed that this a real phenomenon, and not an artifact caused by a flaw or anomaly in the study, then it is unlikely that the worker will be at an interior solution to the utility maximization problem. If the worker really can raise his wage with no concomitant increase in risk – or with an actual reduction in risk – he will surely choose a corner solution involving the maximum possible wage. Hence, the slope of the HW function conveys no useful information on the worker's willingness to pay to avoid risk.

Charge Question 9:
What future research or additional data would offer the most improvement in the Agency's ability to derive summary estimates of mortality risk valuation for environmental policy analyses over the short run? What longer-term research is most needed for improved summary mortality risk valuation estimates?

A. Short Run Research and Data Collection

1. The meta analyses need expanded in a number of ways, including developing meta summaries with a variety of samples that omit CV results and address the questions posed by the review panel. It would also be interesting to consider alternative transformations of the dependent variable –based on the models estimated. Smith et al. (2003) demonstrate a relationship between VSL and the elasticity of labor supply for a specific indirect utility function. Similar potential relationships could be derived for specific assumed functional forms for the indirect utility function underlying the ex ante marginal rate of substitution that underlies the VSL. This strategy may provide alternative sources for VSL measures and a basis for testing how VSL changes with demographic characteristics. There is a large set of labor supply elasticity estimates for different demographic groups.

2. If CV results are to be combined with labor market estimates, then a consistent structural description of the marginal rate of substitution underlying the VSL needs to be estimated jointly. Under these conditions the willingness to pay function for risk reductions would be estimated in one equation and the VSL estimates associated with a second equation. Both would be derived from the same assumed preference structure and would require a comparable treatment of the risk measures. While this strategy imposes significant structure as maintained hypotheses, it does assure consistent treatment of the probabilities and the monetary tradeoffs measured with wage hedonic and stated preference studies. At present only the Smith, Pattanayak and Van Houtneu study (Ecological Economics 2006) illustrates the method with limited calibration findings. The logic could be expanded and the parameters estimated within a generalized method of moments framework with limited incremental effort using the data already assembled for the existing meta analyses..

3. There appear to be significant opportunities to use large existing data sets to consider the effects of age, health conditions, and risk measures. For example, the NHANES sample and other large scale health oriented surveys appear to have significant data to consider hedonic wage models. For NHANES this is a conjecture based on a briefing to EPA staff and associated researchers approximately a year ago. Further information should be collected. It appears there was limited follow-up. The most direct suggestion that can be made without a complete record of the variables collected in these data bases would be to inventory health oriented surveys that could offer the prospects for estimating wage models with different groups. At present, confidentiality limitations prevent researchers from linking measures of risk and other geo-coded information to individual records that contain detailed health conditions.

4. There does not appear to have been an extensive effort to exploit the structure of panel data such as the PSID (together with the newly available BLS age specific job risk measures) to investigate the implications of the age/VSL literature. These data offer the potential to observe survey respondents over their full life cycle. Once again this is a conjecture that may be worth assembling labor economists who familiar with these surveys to judge its feasibility. The proposal is not necessarily to estimate structural

models based on the inter-temporal structures providing different conclusions for the VSL/age relationship, but rather to test specific hypotheses about how the consumption/saving patterns evolve over the life cycle and their relationship to labor supply and risk tradeoffs as would be implied by these models.

B. Longer Term Research

1. At present the air pollution/health models are treated completely independently of the economic models used to estimate the tradeoffs people would make to reduce risk. Moreover, the methods used to apply the survival functions in policy adjust baseline cases of deaths by specific groups (age groups for example) by the implied proportionate reduction in deaths caused by the reduced air pollution. To my knowledge there has been no attempt to use a survival function that allows for age related differences in the effects of air pollution. There is very limited information confirming that the effects would be disproportionately experienced by older groups. The estimates of differentials in the cases avoided for older populations is a consequence of applying the same proportionate adjustment to a higher number of baseline cases of death for older groups. There may be significant differences in these groups behavior that would cause them to have lower exposure profiles. Moreover, we do not know whether these groups undertake more or less mitigating behaviors that would reduce their exposure.

2. The Cameron-DeShazo analysis is the first serious effort to deal with morbidity conditions, reduced risk of death, time periods and recovery as the "object of choice" most relevant to the policy decisions. However, their formulation presents very complex cases whose elements have not been individually evaluated. Equally important these analyses could be undertaken jointly with actual choices of the respondents involved.

3. Real choices in response to the assumed behavior underlying both the wage hedonic and the stated preferences need to be considered. The special focus needs to be how age and other demographic factors relevant to the effects of air pollution affect risk taking behaviors.

4. Confirming the actual links between health conditions and the epidemiological conditions needs to be expanded. Evans and Smith (JEEM 2005) was a limited, special purpose investigation which could be expanded with access to the geo-coding in NHANES.