Berkeley Law

From the SelectedWorks of Robert Cooter

2016

Community Versus Market Values of Life

Robert D. Cooter
David E. DePianto

Available at: https://works.bepress.com/robert_cooter/175/

Truth-Revealing Mechanisms for Courts

by

ROBERT COOTER AND WINAND EMONS*

In trials witnesses often slant their testimony in order to advance their own interests. To obtain truthful testimony, the law relies on cross-examination under threat of prosecution for perjury. We show that perjury law is an imperfect truth-revealing mechanism. Moreover, we develop a truth-revealing mechanism for the same set of restrictions under which perjury rules operate. Under this mechanism the witness is sanctioned if a court eventually finds that the testimony was incorrect; the court need not determine that testimony was dishonest. We explain how truth-revealing mechanisms could combat distortions of observations by factual witnesses and exaggerations by experts, including "junk science." (JEL: D 82, K 41, K 42)

1 Introduction

Witnesses often have a material interest in the court's judgment. The plaintiff and defendant, for example, are interested in the stakes in the dispute, and an expert has an interest in future employment as a witness.[1] In deciding legal disputes, courts must rely on observers to report facts and experts to provide opinions. The interest of the witness in the case provides an incentive to distort testimony. To obtain undistorted testimony, witnesses must face legal sanctions for distortions that offset the gain.

The law relies on cross-examination under the threat of prosecution for perjury to deter distorted testimony. Cross-examination probes the quality of testimony by the witness, searching for internal inconsistencies or contradictions with testimony

* We thank Rafael Gely, Hans Peter Grüner, Martin Hellwig, Lewis Kornhauser, Benny Moldovanu, Eric Rasmusen, Urs Schweizer, Matt Spitzer, and two referees for helpful comments. Emons gratefully acknowledges the hospitality of the UC Berkeley Law School.

[1] For the rapid growth of economists acting as expert witnesses see POSNER [1999a], THORNTON AND WARD [1999], MANDEL [1999], and SLOTTJE [1999]. This form of consulting is now designated "forensic economics." Several associations such as, e.g., the National Association of Forensic Economics (NAFE) as well as a couple of journals like, e.g., the *Journal of Forensic Economics* have emerged due to this boom in the demand for economists as experts.

Journal of Institutional and Theoretical Economics
JITE 159 (2003), 259–279 © 2003 Mohr Siebeck – ISSN 0932-4569

by other witnesses. In a criminal trial for perjury, the plaintiff must prove that the defendant lied or recklessly disregarded the truth.[2]

Establishing guilt or liability often requires more information than anyone can prove in court, so perjury trials or civil trials of false witnesses are rare. In practice, a skillful witness can slant testimony without fear of prosecution or liability. Moreover, as we show formally, even when all this information is available, perjury rules are not truth-revealing.

We apply the theory of mechanism design to the courtroom and derive a system of sanctions that is truth-revealing. To allow for a fair comparison, our sanctions work under the same set of restrictions under which perjury rules operate.[3] Their implementation replaces the fault-based perjury rule with a rule of strict liability; the witness pays a sanction for testimony that proves inaccurate, regardless of whether inaccuracy was the witness's fault. Replacing a fault-based rule with strict liability reduces the information needed by the plaintiff to impose the sanction. In principle, a rule of strict liability can deter distortions by factual witnesses and exaggerations by experts, including "junk science."[4]

To motivate our analysis, consider the following examples of testimony by witnesses to which our model applies:

Example 1: A suit between two motorists over an automobile accident turns on who was at fault. A pedestrian, who is neutral between the parties, observed the accident. The pedestrian testifies on the question, "Was the stoplight yellow or red?"

Example 2: A woman maintains a sexual liaison with a handsome poor man and an ugly rich man. When a child is born, the mother needs to establish the father's identity to win a paternity suit by testifying on the question, "Who is the child's father?"

Example 3: The defendant in an antitrust suit considers whether to argue that he lacks monopoly power or, alternatively, to concede that he has market power and argue that he did not use it to raise prices. The second argument undermines the first argument, so the defense attorney does not want to make both of them. The defense retains an expert to answer the questions, "Is the defendant a monopolist?" "If he is a monopolist, did he use his market power to raise prices?" The defense attorney's strategy depends on how the expert will answer these questions.

Example 4: The side-effects of a drug injures a consumer, who sues the pharmaceutical company in a civil law country. The judge appoints an expert to answer the

---

2 For details on perjury law see COOTER AND EMONS [2000].

3 These restrictions imply that the existence of a truth-revealing mechanism does not follow trivially from the revelation principle.

4 Martha Nussbaum's testimony in *Romer v. Evans* seems to be an example of expert perjury. She seems to have misleadingly cited the long superseded 1897 edition of a Greek–English lexicon listing no pejorative connotation of the Greek word *tolmêma* whereas in the later 1940 edition, which she normally cites in her academic work, "shameless act" is included as a possible translation of *tolmêma*. See Lingua Franca, Sept/Oct. 1996, http://www.linguafranca.com/9609/stand.html.

question, "Was the drug defective?" The expert knows that the judge wants to end the trial quickly.

In Example 1, the pedestrian who witnesses the accident is neutral in the sense that the decision of the court does not affect her material interests. In Example 2, the mother testifying about her child's paternity has a direct material interest in the case. In Example 3, a party to the case pays an expert witness, as is the usual practice in the "adversarial" systems of the common law countries. By advancing the interests of the party retaining her, an expert witness increases her prospects for employment in subsequent cases. In Example 4, the court selects an expert witness, as is the usual practice in the "inquisitorial" systems of the civil law countries. The judge wants to end the trial quickly. By advancing the interests of the judge in Example 4, the expert witness increases her prospects of being hired by courts in subsequent cases.

All of our witnesses observe a fact that is relatively good or relatively bad for a party in the case. The witness is either certain or uncertain about the observation's accuracy. In terms of Example 1, the pedestrian may have observed that the stoplight was red, but she may be uncertain because the sun was bright. In terms of Example 2, the mother may believe confidently that the poor man is the father. In technical language, a witness receives a signal that is better or worse with high or low precision. When testifying in court, a witness reports on the signal's content (better/worse) and precision (high/low).[5] An *honest* witness reports truthfully about content, and a *dishonest* witness reports falsely about content. A *candid* witness reports accurately about precision, and a *misleading* witness reports inaccurately about precision. We use the phrase "slanted testimony" to mean testimony that is dishonest or misleading.

The court uses the available evidence, including the testimony of witnesses, to decide the case. After a witness testifies, subsequent events may prove that the testimony's content was right or wrong. To illustrate by Example 1, after the pedestrian testifies that the stoplight was red, someone may discover a photograph proving conclusively that the stoplight was yellow. In Example 2, the mother may testify that the rich man was the father and, after the trial, subsequent developments in biology may produce a proof that she was wrong. Note that evidence about the content does not prove unambiguously the quality of the testimony. Proof of the poor man being the father makes it only more likely that the mother did not tell the truth when testifying that the rich man fathered the child.

[5] In our set-up the witness can lie, i.e., report false information. There is a related literature comparing the adversarial (partisan) procedure of the Anglo–Saxon law in which partisan advocates present their cases to an impartial jury with the inquisitorial procedures of Roman–Germanic countries in which judges take an active role in investigating a case (DEWATRIPONT AND TIROLE [1999] and SHIN [1998]). In these papers a party can conceal information but *cannot* report false information. SHIN [1998] justifies the assumption of no false evidence (all reported information is verifiable) with the effectiveness of perjury rules. Our results on perjury rules tend to qualify this assumption.

Our model stylizes these facts. We assume that, after the witness testifies, the court subsequently learns with positive probability whether the testimony's content was wrong or right. The truth-revealing mechanism sanctions the witness whose testimony's content was inaccurate.

A *truth-revealing mechanism* induces honest and candid testimony in all circumstances.[6] Our mechanism has a straightforward interpretation. A witness may gain from dishonest or misleading testimony. Against this gain, the witness must balance the probability and magnitude of a sanction. A truth-revealing mechanism imposes an expected sanction greater or equal to the gain from slanted testimony.

As an illustration, consider the pharmaceutical expert Example 4. Assume her tests indicate that the drug has no defects, but she is uncertain. To promote her future business, the expert can help the judge by testifying falsely that she is certain. By doing so, however, the expert runs the risk that someone will subsequently present irrefutable proof that the drug is defective. With our truth-revealing mechanism, the expected sanction increases when the expert asserts her conclusion with certainty, and the increase exactly equals the gain to the expert from more future business.

The sanctions in the truth-revealing mechanism can be interpreted as bond forfeited by the witness in the event that evidence disconfirms her testimony.[7] Assume in Example 4 the witness reports that she is certain the drug has no defects, but the plaintiff suspects that the witness is actually uncertain. The plaintiff, consequently, challenges the witness to bond her testimony. To retain credibility, the witness has to post bond. In principle, the court or the attorneys in the case can compute the minimal sanction from the witness's gain and the probability of disconfirming evidence.

In Example 3, the testimony is more a matter of opinion than knowledge. The concept of "monopoly" is probably too imprecise for a decisive test of the expert's testimony. The truth-revealing mechanism only applies to propositions that risk disconfirmation. The cross-examining attorney must, therefore, formulate a question whose answer risks disconfirmation. For example, the cross-examining attorney might challenge the witness by asking, "Would you bond the proposition that 3 out of 4 industrial economists who examined the same evidence as you would agree with your conclusion? If not 3 out of 4, how about 2 out of 3, or 1 out of 2?" In this case, bonding serves the purpose of forcing the witness to acknowledge the extent to which her testimony is eccentric. The expert gives the court perspective on her

---

6 Our aim throughout the paper is to reveal the truth. From a societal point of view there may be cases where lies are better than the truth, e.g., if somebody lies to protect a lady's honor, or if a government official lies about foreign policy event to protect ongoing dealings. In FLUET [2003] a court maximizes the *ex ante* surplus from a contractual relationship by not seeking the truth *ex post*. See POSNER [1999b] for a comprehensive economic analysis of the law of evidence with efficiency being the ultimate aim.

7 See COOTER AND EMONS [2000] for an elaborate discussion on how truth-bonding might work in practice.

testimony by acknowledging the extent to which other scientists disagree with her. An expert remains free to defend her own opinions, regardless of their eccentricity.

Let us now turn to a limitation of our truth-revealing mechanism. In all of our examples, for given precision, the witness benefits more from testifying that she observed "better" rather than "worse."[8] Given the content, she also benefits more from testifying that she is "certain" rather than "uncertain." Accordingly, payoffs increase whenever the report's content and certainty improve. This is analogous to a portfolio of stocks becoming more valuable whenever the mean increases and/or the variance decreases.

In Example 1 the defendant gains most if the witness testifies with certainty that the stoplight was red; the second-best testimony is with certainty that the stoplight was yellow. If the witness is uncertain, the court attaches a higher probability to a green stoplight which is bad for the defendant. In Example 2 the mother gains most from testifying with certainty that the rich man is the father. Next, she gains from testifying with certainty that the poor man is the father. Her prospects for winning in court are worse when she cannot testify with certainty about the father's identity.

Example 3 also fits this pattern of reasoning. The defendant's expert benefits the defense most by asserting that the defendant is not a monopolist. Next, the expert benefits the defense by asserting that the defendant is a monopolist who did not raise prices. Given the defendant's strategy, the expert benefits the defense least by asserting that she is uncertain whether the defendant is a monopolist. This pattern of benefits, which fits the conception of the defense in *Texaco v. Pennzoil*,[9] applies whenever the defendant prefers to take a stand on only one issue.

Our mechanism no longer works in scenarios where the witness benefits from being uncertain rather than certain. Change Example 3 such that the defense wants to argue that the defendant is not a monopolist and did not raise prices. An expert advances this defense more by testifying that the defendant probably has monopoly power but she is uncertain about it, rather than testifying that the defendant certainly has monopoly power. This pattern of payoffs violates a monotonicity requirement necessary for the existence of a truth-revealing mechanism that sanctions only wrong testimony. Without monotonicity, more possibilities to sanction are necessary to induce truth-telling. We sketch such a mechanism, which also sanctions right testimony.[10]

---

[8] For purposes of our analysis, the opposite is also acceptable, because "better" and "worse" are arbitrary.

[9] Texaco's lawyer refused to offer expert testimony on damages on the theory that doing so would undermine the confidence of the court that Texaco was not a fault. This proved a disastrous mistake because Pennzoil claimed damages equal to the sale value of the oil field, rather than the expected profits that the oil field would yield. See, e.g., PETZINGER [1999].

[10] In EMONS [2001] we explore this problem further. Here we have the additional problem that the defendant will not present a witness making unfavorable testimony in the first place. We have a trade-off quality versus quantity of testimony. We show that truth-revealing mechanisms produce little testimony of high quality; perjury produces

Finally, we use our model to analyze the fault-based perjury rule. Under the perjury rule if the testimony's content was wrong, the court must use this information to compute the probability that the witness was dishonest. If this probability exceeds the legal standard, the court imposes the sanction for perjury. As we show formally, these Bayesian inferences are difficult because they require much information. This fact provides one reason why perjury prosecutions are so rare.

Next we show that under perjury law a neutral witness will never report a high precision signal. Since a court is more likely to find perjury when testimony was given with certainty rather than uncertainty, a neutral witness minimizes the probability of being sanctioned for perjury by understating her certainty.

Finally, since a simple perjury rule does not adjust the sanction to the probability of detecting and prosecuting the perjury, it lacks the sophistication necessary to induce truthful testimony. Because of these limitations, a perjury rule is an imperfect truth-revealing mechanism.

The remainder of the paper is organized as follows. In the next section we describe our basic framework. In Section 3 we analyze the witness's incentives. In Section 4 we derive the truth-revealing mechanism. In Section 5 we discuss the perjury rule. The last section concludes. Proofs are relegated to the Appendix.

2 *The Model*

A court's decision in a case depends on the outcome of a random event. This random variable $\tilde{X}$ can take the two realizations $X = A$ and $X = B$. To illustrate, a drug may have two side-effects, one (B) somewhat worse than the other (A). Or B might mean, "the defendant is a monopolist but did not raise prices" and A might mean, "the defendant is not a monopolist." The court has some information about the likelihood of the two events which we denote by $\text{Prob}(X = B) := \text{Prob}(B)$ and $\text{Prob}(X = A) := \text{Prob}(A) = 1 - \text{Prob}(B)$.[11]

A party to the dispute (we will take the defendant in what follows) can base his defense on either A or B. His case is somewhat stronger for A than for B. A good defense for A is, however, a pretty bad one for B and vice versa. The defendant is, therefore, interested in obtaining as much information as possible about which state of the world will materialize. If the two states are equally likely, he prefers, of course, A.

A witness observes a fact with an attached probability, which we call a signal, that is relevant to the court's decision.[12] To illustrate, the witness observes a medical fact and infers a definite probability about the occurrence of the side-effects. Or the

---

a different amount of testimony with lower informational content. From an *ex ante* perspective all parties prefer truth-revealing mechanisms.

[11] If the court has no further information and thus no reason to discriminate between the events, by LaPlace's Principle of Insufficient Reason $\text{Prob}(B) = \text{Prob}(A) = 1/2$.

[12] As in SHIN [1998] we treat the information collection process as exogenous in order to focus on the incentives to disclose the evidence. In DEWATRIPONT AND TIROLE [1999] information gathering is costly; their focus is on the incentive to gather information.

witness runs some econometric tests and infers a definite probability of a monopoly. In notation, the witness has observed a signal (Y, P), $Y \in \{A, B\}$, $P \in \{L, H\}$, meaning that the state of the world Y will occur with probability P. The signal (B, H), for example, means that the state B will occur with probability H, which of course implies that the state A occurs with probability $(1 - H)$. Similarly, the signal (A, L) means that the state A occurs with probability L, etc.

First of all we assume that all four signals reduce uncertainty; furthermore, it is convenient to define H and L so that one is more precise than the other in the sense of conveying more information, or formally, $1 > H > L > \max[\text{Prob}(B); \text{Prob}(A)] \geq 1/2$. By this condition, a signal (B, L) or (A, L) is *low precision*, and a signal (B, H) or (A, H) is *high precision*.

We do not further model how the court reaches its decision. All we assume is that the court is more likely decide in favor of the defendant for high rather than low precision, and for A rather than B signals. Therefore, the defendant prefers (A, H) to (B, H) to (A, L) to (B, L) to no signal at all.[13]

We can formalize the defendant's desire for precision by defining the indicator variable $U(A) = 1$ and $U(B) = 0$. The defendant's utility function is then given as

$$\begin{aligned} V &= E(U(\tilde{X}) \mid \textit{testimony}) - \alpha \text{Var}(U(\tilde{X}) \mid \textit{testimony}) \\ &= \text{Prob}(A \mid \textit{testimony}) - \alpha \text{Prob}(A \mid \textit{testimony}) \text{Prob}(B \mid \textit{testimony}), \end{aligned}$$

where E stands for the expected value, Var for the variance of the indicator variable, $\text{Prob}(X \mid \textit{testimony})$ for the probability the court attaches to state X, $X = A, B$, given the testimony, and α measures the (dis-)taste of the associated risk. We assume α to be "sufficiently large," or formally, $\alpha > (L - (1 - H))/(L(1 - L) - H(1 - H))$. If α satisfies this condition, the defendant prefers (A, H) to (B, H) to (A, L) to (B, L) to no signal at all.

As an example consider tossing two dice. The state A corresponds to the sum of the two dice exceeding 7, while B occurs when the sum is 7 or less. Accordingly, $\text{Prob}(A) = 5/12$ and $\text{Prob}(B) = 7/12$. After the two dice are rolled, the witness observes something about the number on one of them. The witness observes either one of the two *high precision signals* {5,6} and {1,2,3,4}. The information that the first die is either 5 or 6 translates into the signal that the good state occurs with probability 3/4, i.e., $(A, 3/4)$ and the information that the first die is less than 5 corresponds to the signal $(B, 3/4)$. Alternatively, the witness observes one of the *two low precision signals* {4,5,6} corresponding to $(A, 2/3)$ and {1,2,3,4,5} corresponding to $(B, 2/3)$.

The witness testifies in court on her private information. She announces a state of the world and the probability with which this state will occur.[14] To avoid confusion,

[13] Our analysis carries over to the case where the defendant's preferences are $(A, H) \succ (A, L) \succ (B, H) \succ (B, L)$. In Proposition 4 we show that truth-revelation is impossible if we sanction only wrong testimony when monotonicity in the precision is violated.

[14] We thus confine our attention to direct revelation mechanisms. The revelation principle (see, e.g., MYERSON [1985]) tells us that in Bayesian decision problems without

we use small letters (rather than formally correct capital letters) for her reported values. Formally, the witness announces (y, p), $y \in \{a, b\}$ and $p \in \{l, h\}$. We will use the following semantics: If $y = Y$, testimony is *honest;* otherwise, testimony is *dishonest*. If $p = P$, testimony is *candid*; otherwise, testimony is *misleading*. Our aim is, quite naturally, to get an honest and candid testimony.

Depending on her reported values, the witness receives a remuneration (wage) $w(y, p) \geq 0$ from a third party. Taking future consequences into account, remuneration is higher when the testimony is more favorable to the party for whom the testimony is given. We view testimony from the viewpoint of the defendant, who prefers high (h) to low (l) precision signals and better (a) to worse (b) news. Consequently, $w(a, h) \geq w(b, h) \geq w(a, l) \geq w(b, l)$.

The wage depends upon the legal and contractual status of the witness. An *interested witness* receives a wage for testimony that increases with the strength of her testimony. Formally, for an interested witness all three of the above wage inequalities are strict. Typically, an interested witness is a party to the suit or an expert paid by a party to the suit. Under U.S. rules, expert witnesses are interested. A *neutral witness* receives a constant wage for testifying, meaning that the equality holds in all of the above weak inequalities. If this constant wage is zero, we will call the witness *disinterested*. Under European rules, expert witnesses are more often neutral. Under European and American rules, witnesses to the facts are typically unpaid for testimony, so they are neutral unless connected to the plaintiff or defendant.[15]

After the witness has testified, further developments in the trial may reflect upon the accuracy of the witness's testimony. We stylize this fact by assuming that the court observes the true state of the world after the trial's end.[16] We will say the testimony is *right* if $X = y$; otherwise, the testimony is *wrong*. Conditional on the relationship between the testimony and the court's observation, the witness can be rewarded or sanctioned. Formally, we denote a sanction/reward by $S(X, y, p)$, where $S > 0$ is a sanction and $S < 0$ a reward. We want to derive mechanisms working under the same set of restrictions as the perjury rule.

---

loss of generality one can restrict attention to direct mechanisms under which agents report truthfully. In such a Bayesian decision problem players' utilities depend on the state of the world and any decision made by other players. In contrast, in our set-up wages depend on direct reports and we impose further restrictions the mechanism has to satisfy. Accordingly, truth-revelation does not automatically follow from the revelation principle.

15 Note that we do not further analyze the relationship between the defendant and the witness. We have specified the defendant's preferences only to motivate the witness's wage schedule.

16 MILLER [2001] uses the same timing of events as we do; he shows that perjury rules should give greater weight to information that surfaced after the witness testified. Since our witness is risk neutral, it is straightforward to extend the analysis to the case in which the court observes the true state only with a probability $\gamma < 1$. Then all the sanctions in the truth-revealing mechanism have to be multiplied by $1/\gamma$; this assumes of course that the witness has sufficient wealth. Note that we could work with any imperfect signal of the witness's observation (Y, P). To save on notation we have chosen the true state of the world.

Therefore, we set the sanction equal to zero whenever the testimony is right, i.e., $S(B, b, l) = S(A, a, l) = S(B, b, h) = S(A, a, h) = 0$.

The witness's expected payoff equals her wage minus the expected sanction. Formally, the payoff is given as $w(y, p) - E(S(\tilde{X}, y, p)|Y, P)$, where $E(S(\tilde{X}, y, p)|Y, P)$ stands for the expected sanction given her reported testimony (y, p) and the true information (Y, P). She chooses her reported testimony (y, p) so as to maximize her expected payoff.

3 The Incentive Constraints

We want to derive a system of sanctions that induces the witness to be honest and candid. We call such a mechanism *truth-revealing*. This means that reporting the true signal must generate at least as much payoff as announcing any other signal. Formally, this requirement means

$$w(Y, P) - E(S(\tilde{X}, Y, P)|Y, P) \geq w(y, p) - E(S(\tilde{X}, y, p)|Y, P)$$
$$\forall (y, p) \in \{b, a\} \times \{l, h\}, \forall (Y, P) \in \{B, A\} \times \{L, H\}.$$

Consider, for example, the case in which the true signal is $(Y, P) = (B, L)$. Here one of our tasks is to guarantee that announcing $(y, p) = (b, l)$ is at least as good as reporting (a, l). Formally, this means $w(b, l) - (1 - L)S(A, b, l) \geq w(a, l) - LS(B, a, l)$.[17] If the witness tells the truth, she receives the wage $w(b, l)$. With probability $(1 - L)$ the state A materializes and the witness has to pay the sanction $S(A, b, l)$. If, in contrast, she reports (a, l), she receives the (higher) wage $w(a, l)$. Now the sanction is $S(B, a, l)$, triggered by the state B which occurs with the (high) probability L if (B, L) is the true signal. Similarly, if the true state is $(Y, P) = (B, L)$, we must guarantee that the message $(y, p) = (b, l)$ is at least as good as the reports (b, h) and (a, h).

Analogous incentive constraints hold for the other 3 signals so that overall we end up with 12 incentive constraints. After some algebraic manipulation and rearranging we get the following 6 chains of weak inequalities.

(1) $$(1 - L)S(B, a, h) - (1 - L)S(B, a, l) \geq$$
$$w(a, h) - w(a, l) \geq (1 - H)S(B, a, h) - (1 - H)S(B, a, l),$$

(2) $$(1 - L)S(A, b, h) - (1 - L)S(A, b, l) \geq$$
$$w(b, h) - w(b, l) \geq (1 - H)S(A, b, h) - (1 - H)S(A, b, l),$$

(3) $$LS(B, a, l) - (1 - L)S(A, b, l) \geq$$
$$w(a, l) - w(b, l) \geq (1 - L)S(B, a, l) - LS(A, b, l),$$

(4) $$HS(B, a, h) - (1 - H)S(A, b, h) \geq$$
$$w(a, h) - w(b, h) \geq (1 - H)S(B, a, h) - HS(A, b, h),$$

[17] Recall that we set the sanction to zero whenever the testimony is right, so $S(B, b l) = S(A, a, l) = 0$.

(5) $$LS(B,a,h) - (1-L)S(A,b,l) \geq w(a,h) - w(b,l) \geq (1-H)S(B,a,h) - HS(A,b,l),$$

(6) $$LS(A,b,h) - (1-L)S(B,a,l) \geq w(b,h) - w(a,l) \geq (1-H)S(A,b,h) - HS(B,a,l),$$

call the first inequality in such a chain (a) and the second one (b).

Before deriving the truth-revealing mechanism in detail, we can already state a preliminary result, namely that truth-revealing mechanisms differ for interested and neutral witnesses. For a neutral witness, sanctions must be constant, whereas for an interested witness the sanctions increase with the strength of the testimony.

Proposition 1: Truth-revealing sanctions for interested witnesses satisfy $S(A,b,l) < S(A,b,h)$ and $S(B,a,l) < S(B,a,h)$. If the witness is neutral, $S(A,b,l) = S(A,b,h) = S(B,a,l) = S(B,a,h)$.

The intuition for this result is straightforward. An interested witness's wage increases with the strength of the testimony, being maximal for the reported value (a,h). If the sanctions were, say, constant, an interested witness would always report (a,h). To compensate for the increasing wage schedule, sanctions must increase with the strength of the testimony. Conversely, if the witness is neutral, the wage schedule provides no incentives not to tell the truth. In order not to distort the wage schedule's proper incentives, the sanctions must be neutral too.

4 The Truth-Revealing Mechanism

Let us now determine our truth-revealing mechanism. We focus on mechanisms employing minimal sanctions and no rewards. This means first that, given no rewards, we set as many sanctions as possible to zero; second, we set those sanctions, which need to be positive, to the minimal values still providing proper incentives. Formally, we look for the mechanism $S^*(\cdot)$ satisfying $S^*(X,y,p) \geq 0 \; \forall (X,y,p)$ and for any other truth-revealing sanctions using no rewards $S'(\cdot)$ it is true that $S^*(X,y,p) \leq S'(X,y,p) \forall (X,y,p)$. We make the sanctions as low as possible in order to minimize the monetary strain on the witness. See also the following discussion on individual rationality in Proposition 3. The reasons why we do not work with rewards are as follows. First, we want to keep the cost of the judicial system low and rewards are costly. A second problem arises if rewards become so high that before having observed the signal the witness knows she will receive an expected reward. Then frivolous witnesses without any knowledge of the case may try to be called upon simply to cash in on the expected reward. Third, we want to compare our truth-revealing mechanism with the perjury rule which does not use rewards either.[18]

[18] See EMONS AND SOBEL [1991] for a more elaborate discussion of the problems generated by expected rewards.

Proposition 2: The truth-revealing mechanism using minimal sanctions and no rewards is given by

$$S^*(X, y, p) = \begin{cases} (w(a,h) - w(a,l))/(1-L) + \\ \qquad (w(a,l) - w(b,l))/L, & \text{if } X = B, y = a, p = h; \\ (w(b,h) - w(b,l))/(1-L), & \text{if } X = A, y = b, p = h; \\ (w(a,l) - w(b,l))/L, & \text{if } X = B, y = a, p = l; \\ 0, & \text{otherwise.} \end{cases}$$

The truth-revealing mechanism obviously reflects Proposition 1. If the witness is neutral, all sanctions are zero. If the witness is interested, sanctions increase with the strength of the testimony.

The sanctions are constructed as follows. When the witness works out, for instance, whether to report the true weak or a false strong signal, she compares the increase in the wage with the increase in the expected sanction. Accordingly, all we have to do is to ensure that the increase in the expected sanction is at least as great as the increase in the wage. This task is somewhat tedious due to the stochastic nature of our problem; sanctions appear in several incentive constraints at the same time. This generates several lower bounds for certain sanctions, and of these we have to take the maximum. With this type of construction, for a certain deviation the witness is just indifferent while for other deviations the incentives are strict. Finally, we have to check that we did not overdo it, i.e., set the sanctions so high that they distort the witness' incentives elsewhere.

After all this technical parlance it seems a good idea to illustrate the truth-revealing mechanism using the dice example. Recall that $L = 2/3$ and $H = 3/4$. Let $w(b,l) = 0$, $w(a,l) = 6$, $w(b,h) = 8$, and $w(a,h) = 10$. Then $S(A,b,l) = 0$, $S(A,b,h) = 24$, $S(B,a,l) = 9$, and $S(B,a,h) = 21$.

If the witness has observed, for example, (A, L), expected sanctions for the possible reports are given by Table 1. Given the true report (a, l), the "marginal" expected sanctions are greater or equal the "marginal" wages, making any deviation from reporting the true signal unattractive.

Here the surprising feature is that the highest sanction is imposed when the witness has reported (b, h) and A materializes. This result follows immediately from (2a). Reporting (b, h) rather than (b, l) increases the wage by a steep 8. This increase has

Table 1
Truth-Revealing Mechanism for the Signal (A, L)

Report	(a,h)	(b,h)	(a,l)	(b,l)
Wage to witness	10	8	6	0
Expected sanction given the signal (A, L)	7	16	3	0
Net payoff	3	−8	3	0

to be compensated by $S(A, b, h)$ which is imposed only with the (low) probability $(1 - L) = 1/3$. We have chosen the example deliberately such that the monotonicity of the wage scheme is not entirely reflected in the sanctions; note that we are talking now about actual rather than expected sanctions. Wages are monotone increasing in the strength of the testimony, but sanctions are not. The highest sanction is imposed for the second highest testimony. Note that such non-monotone incentive schemes are the rule rather than the exception. In the principal–agent problem, for example, the agent's remuneration is typically not monotone in the outcome. See, e.g., GROSSMAN AND HART [1983].[19]

Another unpleasant feature of truth-revealing mechanisms is that agents often do worse if they participate in the mechanism than if they do not participate. In mechanism design jargon, participation in an incentive scheme with this feature is not *individually rational*. Agents do better if they stay out of the incentive scheme than if they take part: they must be forced to participate.[20] Translated into our problem, (interim) individual rationality requires that the witness's expected payoff is non-negative whatever signal she received.[21] To put it differently, expected sanctions may not exceed the wage, or formally

$$w(Y, P) - E(S(\tilde{X}, Y, P)|Y, P) \geq 0 \quad \forall (Y, P) \in \{B, A\} \times \{L, H\}.$$

Fortunately, our mechanism is well behaved.

Proposition 3: The mechanism $S^*(X, y, p)$ defined in Proposition 2 is individually rational.

This result, which follows more or less immediately from our construction has the implication that the witness, whatever signal she receives, need not be forced to testify in court.[22] She will do so voluntarily because her expected payoff from doing so is non-negative. *Ex post*, however, when the testimony has actually turned out to be wrong and the witness is sanctioned, she may end up with a negative payoff as can be seen by our example. We may, therefore, conclude that from an interim point of view it is individually rational to testify in court although *ex post* the witness may regret to have done so.

---

[19] $S(A, b, h) > S(B, a, l) > S(A, b, l)$ holds for interested witnesses under our assumptions. $S(B, a, h) > S(A, b, h)$ holds, however, iff $w(a, h) - w(a, l)(2L - 1)/L > w(b, h) - w(b, l)(2L - 1)/L$. This condition is satisfied if either the wage increase for announcing high precision is higher for a than for b, or if L is sufficiently close to $1/2$.

[20] See, e.g., EMONS [1994] for a detailed discussion of individual rationality.

[21] Recall that under our rule the witness reports truthfully all signals. See HOLMSTRÖM AND MYERSON [1983] for a definition of the *ex ante*, *interim*, and *ex post* concepts.

[22] For a neutral, disinterested witness all truth-revealing sanctions are zero so that individual rationality is trivially satisfied.

Let us conclude this section by showing that truth-revelation is impossible when the wage is decreasing in the precision so that, e.g., $w(b, l) > w(b, h)$, meaning the defendant prefers bad news with low rather than with high precision.

Proposition 4: If $w(b, l) > w(b, h)$ and/or $w(a, l) > w(a, h)$, no truth-revealing mechanism with $S(X, y, p) \geq 0$ if $X \neq y$ and $S(X, y, p) = 0$ if $X = y$ exists.

If the wage increases when the reported precision decreases, the probability that triggers the sanction if the witness is misleading decreases with the wage. But then it is impossible that the expected sanction increases with the wage as is necessary for truth-telling.

The non-existence problem arises because we sanction the witness if and only if testimony is wrong. The probability of a wrong testimony is lower for the high than for the low precision signal. To provide proper incentives, however, the expected sanction for wrong testimony must be higher for the high rather than the low precision signal.

The non-existence phenomenon disappears if we also allow the witness to be sanctioned when the testimony is right, i.e., if we let $S(X, y, p) > 0$ for $X = y$. This can be seen by the following extension of our mechanism. Consider the case where $w(a, h) \geq w(a, l) \geq w(b, l) > w(b, h)$. We introduce sanctions $\sigma(y, p)$ which are levied in addition to $S(X, y, p)$. We charge the witness $\sigma(y, p)$ simply for reporting (y, p), independently of whether this report turns out to be right or wrong. Let $\sigma(a, h) = \sigma(a, l) = \sigma(b, h) = 0$ and $\sigma(b, l) = w(b, l) - w(b, h)$. Then the witness's wage net of σ, $W := w - \sigma$, satisfies $W(a, h) \geq W(a, l) \geq W(b, h) = W(b, l)$. Construct $S^*(X, y, p)$ as in Proposition 2, using W rather than w. Obviously, the extended mechanism $(\sigma(\cdot); S^*(\cdot))$ is truth-revealing.[23]

If we sanction the witness only for wrong testimony, we do not have enough leverage to induce the report (b, h) rather than (b, l). Letting the witness pay $\sigma(b, l) > 0$ gives us this leverage (alternatively, we can also reward the witness with $\sigma(b, h) < 0$). We plan to explore this case further in a subsequent paper.

5 *The Perjury Rule*

Let us now compare our truth-revealing mechanism S^* with the perjury rule. We start with a straightforward formalization of the perjury rule.[24] As in our mechanism under a perjury rule the sanction is zero whenever the testimony is right. If, however, the testimony is wrong, the court uses this information to compute the probability ϕ

[23] Even more to the point is the mechanism $\sigma(y, p) = w(y, p)\ \forall (y, p)$ and no further sanctions for wrong testimony. Under this mechanism the witness's payoff equals zero for all reports.

[24] We model a Bayesian court's decision process. There are also indications that a trial court process of fact finding and aggregation is not purely Bayesian but is constrained by rules of evidence and procedure; see, e.g., POSNER [1999b]. Therefore, DAUGHETY AND REINGANUM [2000a], [2000b] use axiomatic methods to model information and decisions in court.

that the witness did not tell the truth. If this probability exceeds a legal standard $\overline{\phi}$, the court imposes a sanction $s > 0$; if the probability ϕ is below the legal standard, the sanction is zero.[25] Formally,

$$S_p(X, y, p) = \begin{cases} s, & \text{if } \phi(X, y, p) \geq \overline{\phi}; \\ 0, & \text{otherwise.} \end{cases}$$

Computing the probability ϕ of not having reported, the truth turns out to be tricky. First, the court has to know the probabilities Prob(B) and Prob(A) with which the two states of nature occur. Note that we did not use this piece of information for our mechanism S^*. Moreover, the court needs to know the probability distribution over the signals (Y, P) which we denote by Prob$((Y, P))$.[26] To have some structure, let us make the reasonable assumption that low precision signals are at least as likely as the high precision signals, i.e., $\text{Prob}((B, H)) \leq \text{Prob}((B, L))$ and $\text{Prob}((A, H)) \leq \text{Prob}((A, L))$. Note once again that for our mechanism S^* we did not use the probability distribution Prob$((Y, P))$.

If the witness has reported, say, (a, h) and nature has chosen B, the probability of not having told the truth is[27]

$$\begin{aligned} \phi(B, a, h) &= 1 - \text{Prob}((A, H)|B) = 1 - \frac{\text{Prob}(B \cap (A, H))}{\text{Prob}(B)} \\ &= 1 - \frac{\text{Prob}(B|(A, H))\,\text{Prob}((A, H))}{\text{Prob}(B)} = 1 - \frac{(1 - H)\,\text{Prob}((A, H))}{\text{Prob}(B)}. \end{aligned}$$

The probability that (a, h) was not the true signal given B equals the sum of the probabilities that the witness has observed (a, l), (b, l), and (b, h) given B which in turn equals 1 minus the probability that (a, h) was the true signal given B. Analogously, we compute

$$\phi(B, a, l) = 1 - \frac{(1 - L)\,\text{Prob}((A, L))}{\text{Prob}(B)}, \quad \phi(A, b, h) = 1 - \frac{(1 - H)\,\text{Prob}((B, H))}{\text{Prob}(A)},$$

$$\text{and} \quad \phi(A, b, l) = 1 - \frac{(1 - L)\,\text{Prob}((B, L))}{\text{Prob}(A)}.$$

Let us illustrate these probabilities by means of our dice example. Before the dice are tossed, nature chooses with equal probability $1/2$ whether the witness will observe high or low precision signals. If she is to observe a low precision signal and the outcome of the first toss is 4 or 5, she receives the signal (A, L) and (B, L) with equal probability $1/2$. With this signal generating process we have $\text{Prob}((A, H)) = \text{Prob}((A, L)) = 1/6$ and $\text{Prob}((B, H)) = \text{Prob}((B, L)) = 1/3$.

---

25 For $\overline{\phi} \in (0, 1)$ the perjury rule, essentially, works like a negligence rule with a “due care standard” $\overline{\phi}$. If $\overline{\phi} = 0$, the perjury rule functions like a rule of strict liability and if $\overline{\phi} = 1$ like a rule of no liability.

26 This probability distribution depends of course on the stochastic process generating the outcomes A and B. For an application see the following discussion of our dice example.

27 A more precise yet more cumbersome notation for $\phi(B, a, h)$ would be $\phi(\neg(a, h)|B)$.

We then compute $\phi(B, a, h) = 13/14$, $\phi(B, a, l) = 19/21$, $\phi(A, b, h) = 4/5$, and $\phi(A, b, l) = 11/15$.

Given these four probabilities of not having told the truth, we may now state the result that a neutral witness's testimony is misleading under the perjury rule when she has observed a high precision signal.

Proposition 5: Under the perjury rule for a neutral witness the low precision signal weakly dominates the corresponding high precision signal.

This result is easily explained. Suppose the witness has observed the signal (A, H) and she compares the honest and candid message (a, h) with the honest but misleading message (a, l). The probability that the testimony is wrong – the outcome which may trigger the perjury rule's sanction – is the same for both messages. Nevertheless, the probability of not having told the truth is higher for (a, h) than for (a, l), i.e., $\phi(B, a, h) > \phi(B, a, l)$. Consequently, the expected sanction for (a, h) is at least as great as for (a, l) and reporting (a, l) weakly dominates (a, h): If $\phi(B, a, h) < \overline{\phi}$ $[\overline{\phi} < \phi(B, a, l)]$ so that the witness is never [always] sanctioned, misleading does not hurt. If, in contrast, $\phi(B, a, l) < \overline{\phi} < \phi(B, a, h)$, being disingenuous is strictly better than being candid.[28]

Given that a neutral witness is misleading for high precision signals, the next natural question to ask is: under what conditions is her testimony under a perjury rule at least honest? Here we have the following straightforward result.

Proposition 6: If $\max[\phi(A, b, l), \phi(B, a, l)] < \overline{\phi}$ or $\min[\phi(A, b, l), \phi(B, a, l)] \geq \overline{\phi}$, the neutral witness' testimony is honest.

From Proposition 5 we know that the witness always reports a low precision signal. If the probabilities of untruthful testimony are below the legal standard for both relevant messages (b, l) and (a, l), the witness is never sanctioned and, accordingly, indifferent between the two messages. Therefore, her report is honest (though not necessarily candid). If both probabilities of being untruthful are above the legal standard, for both messages the witness is sanctioned by the amount s whenever the testimony is wrong. If the victim reports honestly, the probability of being wrong $(1 - H$ resp. $1 - L)$, is lower than for an dishonest report $(H$ resp. $L)$.

Let us finally analyze the incentives a perjury rule gives an interested witness. As can be expected, the result is negative.

Proposition 7: If the witness is interested, a perjury rule is never truth-revealing.

Incentive compatibility requires for an interested witness that the sanctions increase with the strength of the testimony. To compensate for the increasing wage, the increase in the sanctions has to take on more values than the perjury rule does where

28 See, e.g., KREPS [1990, pp. 418–421], for a discussion of the pro and cons of the concept of weak dominance. Proposition 5 is not true if, say, $\phi(B, a, h) < \phi(B, a, l)$ which, in turn, holds if $(1 - H)P(A, H) > (1 - L)P(A, L)$; high precision signals must be much more likely than low precision signals.

the sanction is either 0 or s. To put it differently: the binary perjury rule lacks the sophistication to give an interested witness proper incentives.

Let us discuss the informational requirements of our sanction system S^* and the perjury rule S_p. Our mechanism focuses essentially on the witness's wage schedule and uses this information to derive the incentive compatible sanctions. The perjury rule, in contrast, focuses on the stochastic processes generating the signals and the final outcomes to determine the probability that the witness did not tell the truth. Accordingly, both mechanisms use different pieces of information.

Being a crime, one element for perjury is the intention to do wrong (*mens rea*, guilty mind) which we have not considered so far in our rule. Here we may argue that a personal gain from lying is a necessary condition for intent. A neutral witness gains nothing from lying. Accordingly, a neutral witness should not be prosecuted for perjury. Only when the witness is interested, the perjury rule S_p is triggered. If we interpret perjury in this way, sanctions for a neutral witness are zero (as in our mechanism) and the incentive of a neutral witness to understate precision diminishes.[29] Note that with this interpretation of perjury the court uses the wage function to determine whether the perjury process is triggered.

We have looked at a simple, binary perjury rule. If we allowed for sophisticated perjury rules where the sanction varies with the probability of its application and the gain from wrongdoing, then we could also elicit the truth with a perjury rule. Notice, however, that the perjury rule then uses information about wages in much the same way as we do and in addition the perjury rule needs all the information about the stochastic processes. Actual perjury rules resemble simple rules more than sophisticated rules; see COOTER AND EMONS [2000].

6 Conclusions

Economists have devoted much effort to developing truth-revealing mechanisms, but only few of these developments have been applied to courts. SANCHIRICO [2000], [2001] investigates the role of evidence production in the regulation of private behavior via judicial and administrative process. BERNARDO, TALLEY, AND WELCH [2000] analyze how legal presumptions can mediate between costly litigation and *ex ante* incentives. DEWATRIPONT AND TIROLE [1999] and SHIN [1998] compare the adversarial with inquisitorial procedures in arbitration. DAUGHETY AND REINGANUM [2000a] model the adversarial provision of evidence as a game in which two parties engage in strategic sequential search. DAUGHETY AND REINGANUM [2000b] use axiomatic and Bayesian methods to model information and decisions in a hierarchical judicial system; axioms represent constraints that rules of evidence impose at the trial. MILLER [2001] shows that when the court has information when the witness testifies and information that surfaces after testimony, perjury rules should give greater weight to the latter. All of these papers are of different focus than ours.

---

[29] In practice the probability of prosecuting a neutral witness for perjury is close to zero; see COOTER AND EMONS [2000].

We have shown for a simple framework that existing legal practices create incentives for witnesses to give slanted or false testimony. More importantly, we have developed a mechanism that prevents slanted or false testimony.

A few qualifications are in order. First, we have looked at the incentive problem of preventing slanted testimony. We have not looked at other incentive problems, such as withholding unfavorable information.[30] Second, we did not analyze the witness's effort to gather information. The more effort a witness provides, the more precise her signal, say. If effort were observable, the court could use this information to infer the quality of the testimony. Third, we assume that the process generating the evidence confirming or disconfirming the testimony is exogenous.[31] We do not model how this evidence comes into existence and how it is brought to the attention of the court. In the inquisitorial system the court or the party against the witness has testified may create the new evidence; in the adversarial system only the latter will have an incentive to search for new evidence. Note, however, that the perjury rule also needs new evidence to be triggered. Comparing our mechanism with perjury given that new evidence pops up thus seems to be fair. Fourth, in considering implementation of truth-revelation by truth-bonds, we have not investigated the extent a market would tend to bond at the optimal level. The answer to this question determines the extent and type of regulation required in a market for truth-bonds. Fifth, POSNER [1999a] argues that reputation plays a major role in disciplining experts. Given the current discussion we have some doubts that reputation is effective; moreover, reputation does not discipline occasional witnesses. Nevertheless, it might be interesting to incorporate reputation into the analysis. These questions present interesting tasks for future research.

Many obstacles impede institutionalizing our mechanism, but a move to strict liability has the promise of significantly improving the quality of testimony in court.

Appendix

A.1 Proof of Proposition 1

(1a) and (2a) imply increasing sanctions for interested witnesses; constant sanctions for neutral witnesses follow from (1b) and (2b). *Q.E.D.*

A.2 Proof of Proposition 2

We use the first inequalities (1a)–(6a) to determine the smallest incentive compatible sanctions. Following Proposition 1, $S(A, b, l)$ and $S(B, a, l)$ are candidates to be set to zero. Yet, if we set $S(B, a, l) = 0$ and the witness is interested, (3a) implies $S(A, b, l) < 0$. Accordingly, we set $S^*(A, b, l) = 0$. (3a) then implies $S^*(B, a, l) = (w(a, l) - w(b, l))/L$. (2a) then defines $S^1(A, b, h) = (w(b, h) - w(b, l))/(1 - L)$

[30] This problem is dealt with in EMONS [2001].

[31] We ignore, e.g., the incentives of the other party to call a witness. For example, in adversarial systems competition between advocates who cannot prove every true statement can fully inform the fact-finder LIPMAN AND SEPPI [1995].

while (6a) defines

$$S^2(A, b, h) = \frac{1}{L}\left[w(b, h) - w(a, l) + \frac{1-L}{L}\,(w(a, l) - w(b, l))\right].$$

Here we have

$$S^1(A, b, h) \geq S^2(A, b, h) \quad \Leftrightarrow$$
$$[w(b, h) - w(b, l)]\left[\frac{L}{1-L} - 1\right] \geq [w(a, l) - w(b, l)]\left[\frac{1-L}{L} - 1\right],$$

which holds because $L > 1/2$. Hence, $S^*(A, b, h) = S^1(A, b, h)$.

Given this, (1a) then defines

$$S^1(B, a, h) = \frac{w(a, h) - w(a, l)}{1-L} + \frac{w(a, l) - w(b, l)}{L},$$

(5a) implies

$$S^2(B, a, h) = \frac{w(a, h) - w(b, l)}{L},$$

and (4a) defines

$$S^3(B, a, h) = \frac{w(a, h) - w(b, h)}{H} + \frac{1-H}{H}\frac{w(b, h) - w(b, l)}{1-L}.$$

While it is straightforward to see that $S^1(B, a, h) \geq S^2(B, a, h)$, proving the second inequality is more tricky. Here we have

$$\begin{aligned}
&S^1(B, a, h) \geq S^3(B, a, h)\\
&\qquad \Leftrightarrow w(a, h)\left[\frac{H}{1-L} - 1\right] + w(b, h)\left[1 - \frac{1-H}{1-L}\right]\\
&\qquad \geq w(a, l)\left[\frac{H}{1-L} - \frac{H}{L}\right] + w(b, l)\left[\frac{H}{L} - \frac{1-H}{1-L}\right]\\
&\qquad = w(a, l)\left[\frac{H}{1-L} - 1\right] + w(b, l)\left[1 - \frac{1-H}{1-L}\right]\\
&\qquad\quad + \left[\frac{H}{L} - 1\right](w(b, l) - w(a, l)),
\end{aligned}$$

which holds given $H > L$ and our assumptions on $w(\cdot)$. Consequently,

$$S(B, a, h) = \frac{w(a, h) - w(a, l)}{1-L} + \frac{w(a, l) - w(b, l)}{L}.$$

It remains to be shown that (1b)–(6b) also hold. (1b), (2b), and (3b) are obvious. Subtracting (3b) from (2b) yields $w(b, h) - w(a, l) \geq (1-H)S(A, b, h) - (1-L)S(B, a, l) \geq (1-H)S(A, b, h) - HS(B, a, l)$, implying (6b). Adding (1b) to (3b) generates $w(a, h) - w(b, l) \geq (1-H)S(B, a, h) + (H-L)S(B, a, l) > (1-H)S(B, a, h) - HS(B, a, l)$, meaning (5b) is satisfied. Last but not least, subtracting (2b) from (5b) yields $w(a, h) - w(b, h) \geq (1-H)S(B, a, h) - (1-H)S(A, b, h) \geq (1-H)S(B, a, h) - HS(A, b, h)$ which is (4b).

It is not possible to raise one sanction and at the same time lower another sanction. To see this, suppose we raise $S(A, b, l)$. (3a) then implies that $S(B, a, l)$ goes up. $S(A, b, h)$ increases, either because $S(A, b, l)$ (by (2a)) or $S(B, a, l)$ (by (6a)) is higher. $S(B, a, h)$ goes up because $S(B, a, l)$ (by (1a)), $S(A, b, l)$ (by (5a)), or $S(A, b, h)$ (by (4a)) is higher. Increasing $S(B, a, l)$, $S(A, b, h)$, or $S(B, a, h)$ does not allow us to lower the sanctions for weaker testimony; the sanctions for stronger testimony either do not change or go up, depending on the binding constraints. *Q.E.D.*

A.3 Proof of Proposition 3

If $Y = (B, L)$, $w(b, l) - (1 - L)S(A, b, l) = w(b, l) \geq 0$. If $Y = (B, H)$,

$$w(b, h) - (1 - H)S(A, b, h) = w(b, h) - \frac{1 - H}{1 - L} w(b, h) + \frac{1 - H}{1 - L} w(b, l) \geq 0,$$

since $H > L$. If $Y = (A, L)$,

$$w(a, l) - (1 - L)S(B, a, l) = w(a, l) \left[1 - \frac{1 - L}{L}\right] + \frac{1 - L}{L} w(b, l) \geq 0$$

as $L \geq 1/2$. Finally, if $Y = (A, H)$,

$$\begin{aligned} &w(a, h) - (1 - H)S(B, a, h) = \\ &w(a, h) \left[1 - \frac{1 - H}{1 - L}\right] + w(a, l) \left[\frac{1 - H}{1 - L} - \frac{1 - H}{L}\right] + \frac{1 - H}{L} w(b, l) \geq 0 \end{aligned}$$

because $H > L \geq 1/2$. *Q.E.D.*

A.4 Proof of Proposition 4

If the true state is, e.g., (B, L), reporting (b, l) must be better than (b, h) meaning $w(b, l) - (1 - L)S(A, b, l) \geq w(b, h) - (1 - L)S(A, b, h)$. If the true state is (B, H), (b, h) must be better than (b, l) meaning $w(b, h) - (1 - H)S(A, b, h) \geq w(b, l) - (1 - H)S(A, b, l)$. Rearranging gives us

$$(1 - H)[S(A, b, l) - S(A, b, h)] > w(b, l) - w(b, h) > (1 - L)[S(A, b, l) - S(A, b, h)],$$

which cannot hold since $[S(A, b, l) - S(A, b, h)]$ has to be positive and $(1 - H) < (1 - L)$. The same argument applies to the signal (A, L). *Q.E.D.*

A.5 Proof of Proposition 5

If the witness has observed the high precision signal (B, H), the expected sanction equals $(1 - H)S_p(A, b, h)$ if she is honest and candid and $(1 - H)S_p(A, b, l)$ if she is honest but not candid. Since $\phi(A, b, h) > \phi(A, b, l)$, we have $S_p(A, b, h) \geq S_p(A, b, l)$. The same argument applies to the signal (A, H). *Q.E.D.*

A.6 Proof of Proposition 6

Proposition 4 implies that the witness either reports (b, l) or (a, l). If $\min[\phi(A, b, l), \phi(B, a, l)] \geq \overline{\phi}$, the sanction is s for both messages whenever the testimony turns out to be wrong. Suppose the witness has observed (B, H). If she reports honestly (but not candidly) (b, l), the expected sanction is $(1 - H)s < Hs$ which is the expected sanction when she reports (a, l). The same reasoning applies to the other three signals.

If $\max[\phi(A, b, l); \phi(B, a, l)] < \overline{\phi}$, the expected sanction is zero for both, (a, l) and (b, l). Accordingly, the witness will provide an honest testimony. *Q.E.D.*

A.7 Proof of Proposition 7

(3a) implies that $S(B, a, l)$ and $S(A, b, l)$ cannot both be zero for an interested witness. This observation together with Proposition 1 means that for an interested witness incentive compatible sanctions have to take on at least three different values. The perjury rule S_p takes on at most two values, 0 and s. *Q.E.D.*

References

BERNARDO, A., E. TALLEY, AND I. WELCH [2000], "A Theory of Legal Presumptions," *Journal of Law, Economics, and Organization,* 16, 1–49.

COOTER, R., AND W. EMONS [2000], "Truth-Bonding and Other Truth-Revealing Mechanisms for Courts," Working Paper, Department of Economics, University of Bern, http://www.vwi.unibe.ch/theory/papers/emons/wb_0500.pdf.

DAUGHETY, A., AND J. REINGANUM [2000a], "On the Economics of Trials: Adversarial Process, Evidence, and Equilibrium Bias," *Journal of Law, Economics, and Organization,* 16, 365–394.

– – AND – – [2000b], "Appealing Judgements," *RAND Journal of Economics,* 31, 502–525.

DEWATRIPONT, M., AND J. TIROLE [1999], "Advocates," *Journal of Political Economy,* 107, 1–39.

EMONS, W. [1994], "The Provision of Environmental Protection Measures under Incomplete Information: An Introduction to the Theory of Mechanism Design," *International Review of Law and Economics*, 14, 479–491.

– – [2001], "Perjury versus Truth-Revelation: Quantity or Quality of Testimony," Working Paper, Department of Economics, University of Bern, http://www.vwi.unibe.ch/theory/papers/emons/tru_0901.pdf.

– – AND J. SOBEL [1991], "On the Effectiveness of Liability Rules when Agents are not Identical," *Review of Economic Studies*, 58, 375–390.

FLUET, C. [2003], "Enforcing Contracts: Should Courts Seek the Truth," *Journal of Institutional and Theoretical Economics*, 159, 49–64.

GROSSMAN, S., AND O. HART [1983], "An Analysis of the Principal–Agent Problem," *Econometrica*, 51, 7–45.

HOLMSTRÖM, B., AND R. B. MYERSON [1983], "Efficient and Durable Decision Rules with Incomplete Information," *Econometrica*, 51, 1799–1819.

KREPS, D. [1990], *A Course in Microeconomic Theory*, Harvester Wheatsheaf: New York.

LIPMAN, B., AND D. SEPPI [1995], "Robust Inference in Communication Games with Partial Provability," *Journal of Economic Theory,* 66, 370–405.

MANDEL, M. J. [1999], "Going for the Gold: Economists as Expert Witnesses," *Journal of Economic Perspectives,* 13, 113–120.

MILLER, J. D. [2001], "Perjury and Information Weighting," *International Review of Law and Economics*, 21, 329–341.

MYERSON, R. B. [1985], "Bayesian Equilibrium and Incentive-Compatibility: An Introduction," pp. 229–259 in: L. Hurwicz, D. Schmeidler, and H. Sonnenschein (eds.), *Social Goals and Social Organization: Essays in Memory of Elisha A. Pazner*, Cambridge University Press: Cambridge.

PETZINGER, T. [1999], *Oil & Honor: The Texaco–Pennzoil Wars*, Beard Books: Washington, D.C.

POSNER, R. A. [1999a], "The Law and Economics of the Economic Expert Witness," *Journal of Economic Perspectives,* 13, 91–99.

– – [1999b], "An Economic Approach to the Law of Evidence," *Stanford Law Review*, 51, 1477–1546.

SANCHIRICO, C. W. [2000], "Games, Information, and Evidence Production: With Application to English Legal History," *American Law and Economics Review,* 16, 342–380.

– – [2001], "Relying on the Information of Interested – and Potentially Dishonest Parties," *American Law and Economics Review,* 16, 320–357.

SHIN, H. S. [1998], "Adversarial and Inquisitorial Procedures in Arbitration," *RAND Journal of Economics,* 29, 378–405.

SLOTTJE, D. J. (ed.) [1999], *The Role of the Academic Economist in Litigation Support*, Elsevier Science: Amsterdam.

THORNTON, R., AND J. WARD [1999], "The Economist in Tort Litigation," *Journal of Economic Perspectives,* 13, 101–112.

Robert Cooter
School of Law (Boalt Hall)
University of California at Berkeley
Berkeley, CA 94720
USA
E-mail:
rdc@law.berkeley.edu

Winand Emons
Department of Economics
University of Bern
Gesellschaftsstr. 49
3012 Bern
Switzerland
E-mail:
winand.emons@vwi.unibe.ch

Valuing Mortality Risk Reductions for Environmental Policy: A White Paper

U.S. Environmental Protection Agency,
National Center for Environmental Economics

DRAFT
December 10, 2010

For consultation with the Science Advisory Board–Environmental Economics Advisory Committee

1 Introduction

The valuation of human health benefits is often a crucial, but sometimes controversial, aspect of the application of benefit-cost analysis to environmental policies. Valuing the reduced risks of mortality, in particular, poses a special set of conceptual, analytical, ethical and empirical challenges for economists and policy analysts. This white paper addresses current and recent U.S. Environmental Protection Agency (EPA) practices regarding the valuation of mortality risk reductions, focusing especially on empirical estimates of the "value of a statistical life" (VSL) from stated preference and hedonic wage studies and how they might be summarized and applied to new policy cases using some form of benefit transfer. Benefit transfer concepts will be highlighted throughout the paper, since any application of existing empirical estimates of values for health risk reductions to new policy cases is inherently a benefit transfer problem.

The main intended audience for this paper is EPA's Science Advisory Board-Environmental Economics Advisory Committee (EEAC). The main objectives of the paper are to highlight some key topics related to the valuation of mortality risks, and to describe several possible approaches for synthesizing the empirical estimates of values for mortality risk reductions from existing hedonic wage and stated preference studies for the purpose of valuing mortality risk reductions associated with future EPA policies. Some of these approaches could be implemented in the short term, but others will likely require longer term research. We are soliciting general feedback and specific recommendations from the SAB-EEAC on each of these key topics and approaches.

1.1 Key topics

We highlight several issues in this paper, offering preliminary recommendations where we feel conclusions can be supported by existing data and methods. In other cases we describe alternative methods, data and data gaps, and possible future directions, with the intention of soliciting meaningful feedback from the EEAC. The key topics addressed in this paper—loosely ordered from short- to longer-term tasks—include:

- *Improving communication by reporting value estimates in terms of risk changes rather than "statistical lives."* We fear, as do others, that the prevalence of such terms of art as "the value of a statistical life" has contributed to unnecessary confusion and consternation among decision-makers and members of the general public. We aim to ease these communication difficulties by replacing the VSL terminology with the straightforward term "value of mortality risk" (VMR). The "units" associated with the mortality risk change must be clearly delineated and in this paper we report the units in terms of willingness to pay for a reduced risk of 1/1,000,000 or a "micro-risk," following Cameron (2008) and Howard (1989). We believe that this term provides a more accurate description of the fundamental valuation concept that underlies the marginal willingness to pay for risk reduction, and that this choice of measurement unit is a more natural one considering the typically small (relative to the full suite of risks from all hazards) changes in individual-level risks resulting from most environmental policies.
- *Alternative approaches for updating EPA's best central estimate, or range of estimates, of the willingness-to-pay for mortality risk reductions for use in regulatory impact analyses.* EPA is interested in updating its guidance to better reflect the existing estimates of mortality risk reduction values in the revealed and stated preference literatures. Specifically, how can the empirical results (described below in Section 4) be used to revise EPA's mortality risk valuation guidance in the form of a revised point estimate or range or benefit transfer function?
- *Incorporating a cancer differential into mortality risk valuation guidance.* We discuss the possibility of adding a "cancer differential" (often called a "cancer premium" in the literature) to the standard (non-cancer) estimates of mortality risk reduction values, specifically for use in analyzing policies expected to reduce carcinogenic pollutants. EPA first raised the issue of a cancer premium with the

EEAC in 2000 (USEPA 2000b), but the literature has developed considerably since that time. Given its importance for the valuation of environmental health risks in particular, we review the current literature and recommend including a cancer differential in future guidance.

- *The role of altruism in valuing risk reductions.* The role of altruistic motives for improved health and safety is typically ignored in most benefit-cost analyses but may have important implications for estimating individuals' willingness to pay for environmental improvements. We review several recent studies that examine the role of altruism in benefit-cost analysis and highlight the potential relevance of these findings for the valuation of mortality risk reductions, in particular their implications for interpreting and transferring stated preference estimates of "public" versus "private" risk reductions.
- *Toward functional benefit transfer.* We discuss specific issues that we expect to arise in applying both classical and Bayesian meta-regression techniques to new datasets of stated preference and hedonic wage value estimates described in this paper, as possible approaches for developing a benefit transfer function. We also discuss the structural benefit transfer approach, which involves specifying a direct or indirect utility function, including parameters that can describe the relevant attributes of the risk to be evaluated, and then deriving analytical expressions for observable economic variables that can be used to calibrate the parameters of the preference function. Developing a valid benefit transfer function, using either meta-regression or a structural approach or some combination of these, is a longer-term task than the others mentioned above, but EEAC feedback on these issues would be very helpful in shaping EPA's research agendas in these areas.

1.2 Roadmap

The remainder of this white paper is organized as follows. Before we address our key topics in more detail, Section 2 provides background discussion that (1) describes the valuation challenge facing the Agency and the differences in the contexts underlying existing mortality risk reduction value estimates and the policy scenarios we seek to analyze; (2) briefly summarizes EPA's most recent guidelines for valuing mortality risk reductions (USEPA 2008);[1] and (3) recaps the main recommendations from several recent expert advisory committees to EPA on the valuation of human health risk reductions and the use of meta-analyses for combining estimates from different studies.

With this context in mind, in Section 3 we describe and discuss three of the key topics of this whitepaper: terminology and metrics, cancer risk valuation, and altruism. In Section 4, we review the empirical mortality risk value estimates from the stated preference and hedonic wage literatures, including recent meta-analyses of these literatures. The discussion of the stated preference literature includes a newly assembled database of stated preference estimates of mortality risk reduction values in anticipation of an updated meta-analysis. We also review and extract value estimates and other attributes from hedonic wage studies that have provided estimates of the VSL, with selected studies spanning 1974 to the present. We discuss strengths and weaknesses of these studies for application to environmental policies.

In Section 5 we discuss alternative approaches for synthesizing the estimates from these literatures as a necessary step for updating EPA guidance. A longer term goal is to develop a benefit transfer function for valuing mortality risk reductions, rather than relying on the current practice of transferring a single central point estimate. We discuss two basic approaches for developing such a benefit transfer function: meta-analysis and structural benefit transfer. Meta-analysis uses statistical regression techniques to quantify the influence of study, policy, demographic, and possibly other variables on the willingness to pay for health risk reductions. The structural benefit transfer approach involves specifying a direct or indirect utility function and then deriving analytical expressions for observable economic variables that can be used to calibrate the parameters of the preference function.

[1] These are reflected in EPA's revised *Guidelines for Preparing Economic Analyses* (2008).

Section 6 concludes with summaries of the key topics and needs for both short-term guidance and longer-term research.

2 Background

2.1 The valuation challenge

Benefit cost analysis is a useful tool that provides detailed information on a wide variety of consequences associated with environmental policies. Benefits are based on what individuals would be willing to pay for risk reductions or for other improvements from pollution reduction. Costs are determined using the value of the resources directed to pollution reduction. As safeguarding human health is among the EPA's primary goals, to develop more complete and more accurate benefit-cost analyses of its policies, EPA must estimate individuals' willingness to pay for reductions in health risks from environmental harms. Ideally, benefit-cost analysis of policies that reduce health risks would account for all of the factors that may cause willingness to pay to vary across different types of policies and individual characteristics and circumstances. The literature has indicated that these factors may include the sources of risk affected by the policy (e.g., hazardous air pollutants, water contamination, etc.), the resulting health conditions (e.g., cancer, cardio-respiratory diseases, gastro-intestinal diseases, etc.), how the policy affects the timing of morbidity and mortality risks across each individuals' life span (i.e., how it shifts the "survival curve"), the income and other personal characteristics of the affected individuals, and how the changes in risks are perceived by those individuals. While addressing all of these factors simultaneously is currently empirically infeasible, there are three challenges that we highlight for their direct relevance to EPA.

First, fundamental to this valuation challenge is that the risk reductions provided by EPA policies are inherently public in nature, unlike, for example, private purchase decisions. The distinction is important because individuals may reasonably value risk reductions from public policies differently than those from private actions even if their own mortality risks are affected in a quantitatively identical manner. Such differences could be due to differences in "controllability," "dread," or other tangible or intangible factors (e.g., Slovic 1987, Savage 1993, Chilton et al. 2006). Furthermore, public policies raise issues about altruistic values for risk reductions to others, something that may be of particular relevance for environmental risks. EPA would like to use the existing literature to evaluate the extent and nature of altruistic values and consider how to formulate mortality risk valuation guidance accordingly. We address altruism in greater detail in Section 6.3.

A second major challenge for the valuation of mortality risk reductions for environmental policies is the intertwined nature of morbidity and mortality risks. Environmental policies generally do not reduce the risks of fatal workplace or automobile accidents, for example, which provide the context for many of the mortality valuation estimates in the literature and generally have little or no accompanying morbidity or period of illness. Ideally, we would use an integrated model that could estimate willingness to pay for mortality and associated morbidity risk reductions simultaneously. Developing such a model is beyond the scope of this white paper and current guidance development effort, and is near the frontier of the empirical valuation literature. Nevertheless, to the extent possible with currently available data and models, we would like to account for how individuals consider morbidity in existing estimates of mortality risk reduction values when they always occur together. It also is important to capture some related losses that may not be reflected in willingness to pay estimates, depending on context in which they were estimated. For example, reduced health from illness preceding death is certainly a loss to an individual and his or her quality of life, but may not be reflected in VSL estimates from the hedonic wage literature, which are based on the risks of workplace injuries that lead to death. Society also is worse off because of the illness due to the individual's lost productivity, something that may not be reflected in revealed or stated willingness to pay estimates, depending upon the type of insurance held by the individual and possibly the scenario description.

This issue is of particular relevance to EPA when addressing reductions in cancer risks since many EPA policies focus on reducing exposure to carcinogens. Ten years ago EPA reviewed the economic literature on valuing fatal cancer risk reductions and discussed a number of risk characteristics that may influence people's values, including but not limited to the timing of the risks (USEPA 2000b,c). The committee recognized many of the issues reviewed by EPA as theoretically valid but empirically ambiguous, and therefore recommended that "the only risk characteristic for which adjustments to the VSL can be made is the timing of the risk" (USEPA 2000c p 1). In particular, this recommendation advised against the application of any differential to reflect preferences for reducing cancer risks relative to other types of risk because of dread or other factors. With an additional decade of valuation literature to draw upon, EPA is seeking to re-examine this question using data from the stated and revealed preference studies described below, as well as other relevant empirical results. We will discuss cancer valuation in more detail in Section 6.4.

Finally, the empirical literature may allow us to account for the extent to which individuals value different categories of risks differently in a systematic transfer of benefits. For example, if environmental risk reductions are valued differently from workplace or auto accidents, regardless of whether the mitigation is from private or public actions, our guidance should reflect this difference.

It is important to keep the overarching valuation challenge in mind as we begin discussing recent studies and value estimates. Each study reflects an attempt to measure the value of a reduction in mortality risk from a specific cause (or small set of causes), in a specific context, among a specific population. By now there is ample theoretical and empirical evidence to indicate that values for health risk reductions are not "one-size-fits-all"—that is, they are "individuated" (e.g., Sunstein 2004, Evans and Smith 2008, Scotton and Taylor 2009). For this reason, we believe that there is great scope for improving upon the point value benefit transfer approach that has traditionally been applied to mortality risk reductions based on a central estimate of the VSL. Therefore, we ultimately are seeking both short-term recommendations as well as advice on a longer-term research agenda on how these heterogeneous studies can best be synthesized for systematic benefit transfers to improve the application of benefit-cost analysis to future environmental policies.

2.2 Existing EPA Guidance

EPA's draft *Guidelines for Preparing Economic Analyses* (2008) (hereafter, the draft *Guidelines*) retains the recommendation from the 2000 version, a default central VSL value $4.8 million in 1990 real dollars. This estimate, after adjusting for inflation and real income growth, is to be applied to mortality risk reductions for all types of policies, no matter the source of the risk.[2] The estimate is based on the mean of a probability distribution fit to twenty-six published VSL estimates. The draft *Guidelines* also indicates that the distribution itself can be used for formal uncertainty analysis. The underlying studies, the probability distribution parameters, and other useful information are available in Appendix B of the draft *Guidelines* (USEPA 2008).

The draft *Guidelines* also retains the 2000 version recommendation that the VSL for mortality risk reductions should not be adjusted for differences in sources of risk or population characteristics—rather, these factors should be examined qualitatively. In some cases, the analysis may include a quantitative sensitivity analysis. Analysts should account for timing when valuing mortality risk reductions, and should discount the benefits of future risk reductions at the same rate used to discount other costs and benefits. Because the VSL represents the marginal willingness to pay for contemporaneous risk reductions, this is typically done by estimating the lag between reduced exposure and reduced mortality risks, calculating willingness to pay in all future periods when mortality risks are reduced, and discounting back to the present.

Finally, EPA's draft *Guidelines* also recommends accounting for increases over time in average income. This is done by using projections of real GDP per capita and applying an income elasticity

[2] We report all estimates in 2009 US dollars unless otherwise noted.

estimate. The resulting future (real) VSL will therefore reflect the idea that health risk reductions are normal goods and so willingness to pay will increase with income.

2.3 Recommendations from prior expert committees

This white paper is one stage in a detailed process that EPA has undertaken with the SAB-EEAC to improve the Agency's ability to value health risk reductions. Since its review of EPA's *Guidelines for Preparing Economic Analyses* (USEPA 2000a) the SAB has offered several specific sets of recommendations on valuing risk reductions, particularly for mortality risks.

In July 2000 the SAB-EEAC released an advisory report in response to EPA's white paper, *Valuing the Benefits of Fatal Cancer Risk Reduction*, which focused on benefit transfer issues associated with using existing mortality risk values to estimate the benefits of EPA actions on carcinogens, including potential adjustments that could be made to existing risk values to account for this category of benefits (USEPA 2000b). As noted earlier, after reviewing the white paper and current economics literature, the SAB concluded that, while many of the issues raised in the white paper were theoretically valid and potentially important, the empirical literature supported only accounting for latency and for income growth over time. The SAB-EEAC did not consider other adjustments to EPA's default mortality risk value to be appropriate for the Agency's primary analyses, but could be addressed separately using sensitivity analysis.

An August 2001 SAB report, *Arsenic Rule Benefits Analysis: An SAB Review* (USEPA 2001), generally supported EPA's estimate of the marginal willingness to pay for mortality risk reductions. The SAB also offered additional recommendations to account for the time between reduced exposure and reduced mortality risks. This report coined the term "cessation lag" for this concept and offered specific recommendations for estimating cessation lags based on the types of risk data available. The SAB review also clarified that reductions in exposure to carcinogens—that is, exposure *per se*, aside from the increased cancer risks that the exposure causes—are not a separate benefit category under a damage function approach to valuing reduced risks. The board noted that it is possible that there is an existence value for protected drinking water; however, without sufficient empirical evidence to estimate the magnitude of this value, it cannot be included in the quantitative benefits analysis. Finally, the report indicated that it is appropriate to add the costs of illness to the willingness to pay for mortality risk reductions when estimating the benefits of reduced cancer mortality.

EPA further consulted with the SAB-EEAC on additional mortality risk valuation issues in 2004, developing a strategy to gather additional information on meta-analysis to inform both the SAB-EEAC and EPA (USEPA 2004b). In 2006, EPA returned to the SAB-EEAC with two documents for formal review: a white paper addressing how remaining life expectancy affects willingness to pay for mortality risk reductions, and an expert report on the use of meta-analysis for combining existing mortality risk value estimates. A 2007 report, *SAB Advisory on EPA's Issues in Valuing Mortality Risk Reduction*, responded to both topics (USEPA 2007).

On the subject of life expectancy, the SAB-EEAC noted that there was theoretical ambiguity on how willingness to pay might change with age (and, hence, remaining life expectancy). The committee concluded that the existing economics literature does not provide clear theoretical or empirical support for using different values for mortality risk reductions for differently-aged adults or a constant "value of statistical life year" (VSLY). Thus, the SAB-EEAC recommended that EPA continue using its traditional assumption of an age-independent willingness to pay for mortality risk reductions.

To address meta-analysis, EPA assembled a work group of expert statisticians in December 2005 to discuss the meta-analysis of VSL estimates and to examine three existing meta-analyses: Mrozek and Taylor (2002), Viscusi and Aldy (2003), and Kochi et al. (2006). While the expert workgroup did not endorse any one of these studies, the panel did encourage the use of meta-analytic techniques for the analysis of the existing literature on VSL. The workgroup recommended analyzing stated preference and hedonic wage data separately, and offered a set of principles that should be followed in conducting such an analysis (USEPA 2007).

The SAB-EEAC review of the Meta-analysis workgroup's report stated that meta-regression is "a useful statistical technique for identifying various aspects of study design or population characteristics that are associated with differences in VSL," but concluded that meta-regression is "not appropriate [for] combin[ing] VSL estimates" into a summary measure (USEPA 2007 p i). Rather, the SAB-EEAC suggested using meta-regression to examine how study design characteristics influence the VSL estimates and relying on other statistical techniques to determine a central estimate or range of estimates for use in benefit transfer to new policy cases.

Based on these expert recommendations and other considerations, we believe that updated reviews and meta-analyses of the stated preference and hedonic wage literatures could help refine the Agency's central estimate(s) or range of estimates of the marginal willingness to pay for mortality risk reductions. Studies have shown that values for health risk reductions may depend on differences among policies and the affected individuals. These factors include the sources of risk affected by the policy (e.g., hazardous air pollutants, water contamination, etc.), the resulting health conditions (e.g., cancer, cardio-respiratory diseases, gastro-intestinal diseases, etc.), as well as how the policy affects the timing of morbidity and mortality risks across each individuals' life span (i.e., how it shifts the "survival curve"). Therefore, as is widely recognized in most other contexts where some form of benefit transfer is used for policy analysis, we believe a functional benefit transfer approach should be more accurate than a single point estimate applied in all circumstances. Consequently, we are interested in exploring approaches for developing benefit transfer functions that can account for some or all of these factors.

3 Key Issues for EPA

3.1 Fundamental Concepts and Recommended Terminology Changes

3.1.1 Fundamental Valuation Concept

We begin by identifying the fundamental valuation concept that economists aim to estimate using non-market valuation methods and apply in benefit-cost analyses of policies that reduce human health risks. Consider a general utility function for an individual *i* with income Y_i and some health risk R_i among the arguments: $U_i = U(Y_i, R_i, \mathbf{Z}_i)$. The vector $\mathbf{Z}_i$ is included to emphasize that, in addition to income and risk, the individual's utility (and therefore the willingness to pay for health risk reductions) also may be influenced by many other factors specific to the case at hand. We will highlight several of these factors throughout this white paper. The individual's *marginal rate of substitution* between income and risk is:

$$dU_i = \frac{\partial U}{\partial Y_i} dY_i + \frac{\partial U}{\partial R_i} dR_i = 0 \quad \Rightarrow \quad \frac{dY_i}{dR_i} = -\frac{\partial U / \partial R_i}{\partial U / \partial Y_i}.$$

This marginal rate of substitution, dY_i / dR_i, also can be interpreted as the individual's *marginal willingness to pay* (*wtp*) for a change in risk—that is, the amount of money the individual would be willing to swap for a small change in risk on the margin.[3] This is the fundamental value concept that must be estimated for use in benefit-cost analyses of policies that may improve human health. With estimates of these quantities, conditioned as necessary on possibly many observable characteristics of the policy and the affected individuals, it is straightforward to calculate the total willingness to pay for the risk reductions that are expected to be produced by the policy: $\sum_i wtp_i \times \Delta R_i$, where *i* indexes all individuals affected by the policy, and wtp_i and ΔR_i are the estimated marginal willingness to pay and risk

[3] Throughout this white paper, we will use "*wtp*" to refer to marginal willingness to pay, which will have units of $/change in risk, and we will use "*WTP*" to refer to discrete willingness to pay amounts, which will have units of $.

reduction for individual *i*, both of which may depend on individual-level characteristics and circumstances.[4]

It is important to emphasize that this is a *marginal* value concept—a dollar value *per unit change in risk*. These values should be thought of as the slope of a curve at a point, rather than the height of the curve.[5] For practical purposes, the units used to report estimates of these slope values are of no consequence. They could be reported as dollars per nano-risk (10^{-9}), or micro-risk (10^{-6}), or mili-risk (10^{-3}), etc. As long as the measurement units are known, then the risk changes to be valued can be expressed in the same units and the correct total value can be calculated. The conventional measurement units used for reporting these slope estimates are (effectively) "dollars per mortality" risk changes, usually simply written as "$," where "per mortality" is understood (or misunderstood, depending on the audience). This quantity was often referred to as the "value of life" in the early literature on the subject (e.g., Rice and Cooper 1967). While the terminology varies, the quantity is now typically called the "value of a statistical life," or VSL, where "statistical" has been added to emphasize that valuation is based on changes in risk rather than the loss of life with certainty.[6]

3.1.2 Change in metric and terminology

Despite its widespread usage, this particular selection of measurement units for the denominator of the marginal rate of substitution between income and risk, and the VSL label that has been attached to

---

[4] For ease of exposition we ignore the time dimension here. We will allude to some of the complications that arise in the more realistic dynamic case, using a life-cycle model, in Section 6.2.2 and Appendix A.

[5] Also note that if the risk changes to be valued are large, then the slope of the willingness to pay function may change over the relevant range and so the marginal willingness to pay ´ the change in risk may not give an accurate estimate of total willingness to pay. For the most part in this white paper we will ignore this complication, though we do come back to it in an illustrative example in Section 5.2.1.

[6] A common way of explaining the meaning of the VSL is based on a population's aggregate willingness to pay for an aggregate risk reduction. For example, suppose in a town of 1,000 people a policy is enacted that reduces each person's risk of dying by 1 in 1,000 in a year. Then the expected number of avoided deaths (lives saved) by the policy for the year would be equal to one—a so-called "statistical life." Suppose further that we know (from a survey or other study) that the average amount that people in the town would be willing to pay for the risk reduction of 1 in 1,000 was $8,000. We then know that the aggregate willingness to pay is $8,000,000 for saving the one statistical life, so the "value of a statistical life" would be $8,000,000.

it, have caused or contributed to needless confusion and controversy, especially among non-economists (Cameron 2009). Most economists recognize that the "units" associated with the VSL reflect the aggregation of the small risk reductions across many individuals until that aggregate reflects a total of 1.0, or one statistical life. However, for non-specialists this potentially subtle point is often lost; the addition of the word "statistical" to the terminology does not seem sufficient to clarify the concept.[7]

To help reduce the misconceptions that seem to be inspired or aggravated by the VSL terminology, we propose a change in EPA standard practice such that estimates of health values will be referred to as the "value of mortality risk" (VMR), and report the associated units using standard metric prefixes to indicate the size of the risk change and the associated time scale, e.g., \$/μr/person/yr (dollars per micro[10^{-6}]-risk per person per year) (Howard 1989, Cameron 2009).[8]

As noted earlier the choice of risk increment for aggregating and reporting risk changes is mainly one of convenience. However, we believe that explicitly labeling the units of the VMR in this way more clearly emphasizes that these values refer to small changes in individual-level risks over a definite time span rather than how much money any single individual or group would be willing to pay to prevent the certain death of any particular person. It also should be emphasized that the use of a standardized

[7] A recent example of the confusion surrounding this concept in the popular press can be found in an AP story titled, "American Life Worth Less Today" (Bornstein 2008) that opened by saying "[EPA] has decided that an American life isn't worth what it used to be." The story was referring to an alternate analysis in some air regulatory impact analyses that used a more recent review of the literature to report a lower VSL than is reflected in EPA's 2000 Guidelines. This story quickly spread throughout the media even appearing on the Colbert Report as EPA's efforts to "devalue life." Video clip at http://www.colbertnation.com/the-colbert-report-videos/176175/july-14-2008/the-word---priceless (04:06) Posted on 7/14/2008.

[8] Other alternatives to the VSL to better describe marginal wealth-risk tradeoffs have been used or proposed as well. For example, the UK government uses the term "value of prevented fatality (VPF)," but as described by Wolfe (2007) this designation confronts the same misinterpretations as VSL. Cameron (2009) suggests a greater departure from standard terminology not only to communicate that "lives" are not being valued, but also to clarify that "value" itself should be understood in terms of opportunity costs. After considering several alternatives, the term suggested is "willingness to swap (WTS) other goods and services for a micro-risk reduction," abbreviated WTS (μr). In recent empirical work, Cameron and DeShazo (2008) report results in terms of micro-risk reductions. Scotton and Taylor (2009) use the term "value of a risk reduction" (VRR), noting that "explicit consideration of the heterogeneous values for heterogeneous risks underscores the importance of moving the policy discussion from 'a VSL' to valuation of marginal changes in fatality risks specific to the type of the risk affected by the policy" (p 23).

measurement unit for reporting values for health risk reductions should neither be taken to imply that the values themselves are invariant across individuals or contexts, nor that these marginal values will be constant across the full range of relevant risk changes.

For the remainder of this paper we will use the general term "value of mortality risk" whenever possible. We will report estimates as VMRs, as defined above, to the extent possible, using the VSL terminology only as necessary in discussing the previous literature.

3.2 Altruism and willingness to pay for mortality risk reductions

We now turn to an overarching conceptual issue that may affect the conduct of benefit-cost analysis more generally: altruism. The default assumption for most applications of revealed and stated preference methods for non-market valuation is that individuals' (or households') well-being depends on their own consumption (interpreted broadly to include market and non-market goods and services) and is not directly influenced by the consumption or well-being of others. If this assumption is invalid, we may be concerned that our standard methods of estimating willingness to pay assuming "atomistic" individuals or households may give misleading results in benefit-cost analysis.

There are at least two ways that altruism may be relevant for the valuation of mortality risk reductions. First, some stated preference studies are based on surveys that make a distinction between "public" and "private" risk reductions.[9] The difference, if any, between WTP for public versus private risk reductions may be partly due to altruism, but other factors could be at work as well. For example, a distrust of government may lead some respondents to express a lower WTP for public risk reductions provided through government programs compared to those provided through private initiatives. While stated preference studies may in principle be able to distinguish altruistic preferences from other

[9] Few studies explicitly address the public versus private issue. However, for most of the studies it is possible to assign the estimates to one category: estimates that accrue to an individual only, such as an individual health risk reduction or the decision to wear a seatbelt or purchase a health care treatment, are "private" and estimates that can accrue to the individual and others, such as reductions in highway safety-related deaths, are "public." See section 6.1 for more details on the stated preference studies.

confounding factors, it is difficult to draw clear conclusions from the existing literature because most studies that have been conducted to date were not designed to examine altruism per se.[10] Therefore, the proper application of the results of these stated preference studies may depend in part on how altruism should be treated in benefit-cost analyses. Second, since hedonic wage studies are focused on compensation received by individual workers for taking on private, job-related risk, the mortality risk values from hedonic wage studies do not incorporate altruism. Therefore, if (some forms of) altruistic preferences should be included in benefit-cost analysis, then hedonic wage-based estimates of mortality risk values may need to be supplemented with separate value estimates that capture altruistic preferences alone. On the other hand, if (some forms of) altruistic preferences should be excluded from benefit-cost analyses, then this may influence whether (or how) some stated preference studies should be used for benefit transfers.

EPA's *Guidelines for Preparing Economic Analyses* (USEPA 2000a) discussed the role of altruism in estimating the total benefits of public actions, and noted the key distinctions between paternalistic (or "safety focused") and non-paternalistic (or "preference respecting") forms of altruism.[11] If altruistic motives are non-paternalistic, then individuals care not only about the benefits others receive, but also the costs they bear, and most economists who have studied this issue have concluded that it is generally inappropriate to add these altruistic values for benefits others receive to total willingness to pay. Doing so could lead to "double-counting" some of the benefits and/or costs. Paternalistic altruism, on the other hand, should be included in the calculation of total benefits. EPA's *Guidelines* (USEPA 2000a p 61) describes the issue as follows:

[10] Stated preference studies and the treatment of altruism also may hold promise for identifying preferences related to equity or environmental justice (EJ) concerns. For example, preferences for reductions in risks for others, particularly those who may be disproportionately exposed to pollutants (which are often low income and minority groups typically associated with EJ) could be identified through a well designed stated preference study.

[11] Formally, the utility function of non-paternalistic altruists includes others' utility, while the utility function of paternalistic altruists includes others' consumption of one or more types of private or public goods or services.

While benefits are generally calculated by summing each individual's WTP for his or her own welfare, there are conditions under which it is appropriate to include altruistic values, or individuals' WTP for the welfare of others. Economic theory concludes that if one cares about a neighbor but respects the neighbor's preferences, and if the neighbor would have to pay for the policy action being analyzed, then altruistic benefits should not be counted in a benefit-cost analysis. The intuition behind this result is that, if one respects the neighbor's preferences, one cares about both the benefits and the costs the neighbor faces. It is therefore inappropriate to add the value one attaches to the neighbor's benefits without considering the cost implications of doing so. Comparing individual benefits and costs in this case is the appropriate decision rule.

Altruistic benefits may be counted either when altruism toward one's neighbor is paternalistic or when one will in fact bear the costs of the project but the neighbor will not. In the first case (paternalistic altruism), one cares about the benefits the neighbor will enjoy, e.g., from a health or safety project, but not about the costs the project will impose on him. An example of the second case would be a project whose costs are borne entirely by the current generation; i.e., the project imposes no costs on future generations. In this case, altruism toward future generations by the current generation could legitimately be counted as a benefit.

The conclusions in the *Guidelines* were based largely on Bergstrom (1982) and McConnell (1997) who demonstrated that the optimal provision of public goods based upon selfish preferences is a necessary and sufficient condition for the optimal provision based on social preferences (including altruistic preferences). However, since the publication of the *Guidelines*, Flores (2002) has challenged the conventional wisdom that (non-paternalistic) altruism should be excluded from benefit-cost analysis. Flores showed that passing a private values benefit-cost test is a sufficient but not a necessary condition for non-marginal policies to be potentially Pareto improving, except under special circumstances. That is, even if all altruism is non-paternalistic, failure to include altruistic values may lead to the rejection of policies that are potentially Pareto improving. Flores concluded that "benefit-cost analysis with altruism cannot simply be conducted independent of who pays." In a more recent study, Bergstrom (2006) concluded that "The assumptions under which the private values benefit-cost test is necessary for potential Pareto improvements need not always be satisfied;" nevertheless, "Despite these qualifications… for a broad class of economies, a comparison of the sum of private values to the cost of a project is the appropriate test for determining whether it can lead to a Pareto improvement" (p 348-349).

Bergstrom's conclusion seems to summarize the prevailing view regarding non-paternalistic altruism in benefit-cost analysis, especially for policies that would cause marginal changes in environmental quality (since Flores' counter-examples involved non-marginal changes). Therefore, the main relevance of altruism for mortality risk valuation lies in the distinction between the paternalistic and non-paternalistic forms. Including the former but excluding the latter may require supplementing revealed preference estimates of health risk valuations with a careful selection of results from previous stated preference studies. Stated preference surveys that elicit only private willingness to pay would exclude both forms of altruism. One way to include paternalistic but exclude non-paternalistic altruism would be to design a survey that would inform respondents about health improvements that others would experience from the policy, but also ask each respondent to assume that all others would be taxed an amount equal to their private willingness to pay for the policy, so that their utility remains unchanged (Johansson 1994). It is not clear which if any of existing stated preference studies (many of which are reviewed below in Section 6.1) were designed this way, so the current body of empirical results cannot support the separation of paternalistic from non-paternalistic altruism. We recommend additional research in this area to help estimate paternalistic willingness to pay for environmental policies that reduce health risks. Additional examination of existing studies may shed light on this issue in the relative short-term, and we are interested in feedback on this issue.

3.3 Valuing cancer risks

As noted in our description of EPA's valuation challenge, willingness to pay for cancer risk reductions may be systematically different than that for workplace or auto accidents or other risks not associated with a lengthy and painful illness. This difference is sometimes referred to as a a "cancer premium," but we will use the more general term "cancer differential." While not often defined precisely, the differential is posited as capturing elements of dread and fear of cancer, as well as the pain and suffering from the period of illness preceding death. It might also include income and household productivity losses over this period of morbidity.

Several authors have recommended accounting for this differential in benefit-cost analysis of policies that reduce exposure to carcinogens (e.g., Revesz 1999, Sunstein 2004). To the extent that existing policy guidance on valuing mortality risk reductions is based on non-cancer risk-wealth tradeoffs, this would involve an "adjustment" to the default (generic, non-cancer) mortality risk reduction value. Governmental analyses in the UK have adopted this approach, applying a 100% differential for cancer risks (HM Treasury 2003).[12] In addition, the European Commission has recommended a 50% differential for carcinogenic pollutants over its default value of preventing a fatality (European Commission 2000).

For the purpose of developing guidance, we are interested in assessing the valuation literature on cancer risks and any cancer risk differential, both in the short-term and the longer term. Ultimately, this literature could inform the development of a benefit transfer function, in combination with the stated preference and hedonic wage estimates described in greater detail below. While such longer-term research is being conducted, we believe it is reasonable that evidence of systematically different preferences for cancer risk reductions be part of any recommended short-term guidance.

To inform this discussion, this section contains a somewhat more detailed assessment of the empirical literature on cancer risk valuation, with a particular emphasis on studies that examine risks in both cancer and non-cancer contexts. These studies are described in Table 1 in the following categories:

- studies comparing values for cancer and non-cancer fatal risk reductions
  - stated preference studies that estimate willingness to pay
  - risk-risk studies
- stated preference studies of cancer risks without internal comparisons, and
- related hedonic property and hedonic wage studies.

[12] Specifically, this adjustment is applied for the benefits from asbestos proposals by the UK Health and Safety Executive (HSE).

The first of these categories contains the most direct evidence on any cancer differentials.

Note on Cessation Lag and Latency

Reduced exposure to carcinogens results in reduced cancer incidence after a period of time that EPA has referred to as "cessation lag," a term originally coined by the SAB in its review of the Agency's arsenic in drinking water benefits analysis. Cessation lag addresses only reduced risks from reduced exposure and thus applies best to populations currently at risk. The time between initial exposure and increased cancer incidence is referred to as "latency" in recent EPA analyses, but it is often used in the literature in a broader sense to refer to the time difference between a change in exposure and a change in risk.

Prior SAB-EEAC advice and agency practice has been to estimate cessation lag and latency from available epidemiologic data, apply a value of statistical life estimate at the time at which cancer mortality reductions occur, and discount this value back to the present at the rates prescribed in Agency guidance. The practice has generally been supported by research findings suggesting that individuals discount over these lag times at rates generally consistent with market rates, although some recent stated preference studies find near-zero discount rates over latency periods (Hammitt and Haninger, 2010; Alberini and Scasny, 2010a).

An important issue in estimating a cancer differential is the potential need to consider differences in the time profile of mortality risks between cancer and non-cancer cases. Earlier studies were often silent on the issue, but more recent ones have attempted to address it explicitly. Our focus in this section is on a potential cancer differential that captures the difference in marginal willingness to pay for reduction of cancer mortality risks relative to that of a non-specific mortality risk holding timing equal. That is, the differential, in principle, compares a contemporaneous non-cancer risk reduction with a contemporaneous cancer risk reduction. We recognize that timing may be intertwined with how people perceive and value risk reductions, something that should be considered more fully in any rigorous, systematic benefit-transfer exercise as we develop guidance.

Stated Preference studies including cancer and non-cancer risks

Several stated preference studies have estimated willingness to pay for both cancer and non-cancer risks, in large part to examine a possible cancer differential. A few studies have focused only on cancer risk reductions without an internal comparison to other types of risk. The results of these studies are somewhat mixed—some have found evidence of a cancer differential (Hammitt and Liu 2004, Tsuge et al. 2005, Alberini and Scasny 2010a, and Alberini and Scasny 2010b), while a few others found no such evidence (Hammitt and Haninger 2010, Adamowicz et al. 2008) when looking at whole-household or public risks. Cameron and DeShazo (2008) found evidence of a differential for some cancers (breast and prostate) over other cancers (colon, lung, and skin), but not over other health endpoints (heart attacks and disease).

There have been two risk-risk tradeoff studies specifically examining how preferences for cancer risk reduction compare to those for automobile accident risk reductions. By asking respondents to choose among different bundles of risks, these simplified choice experiments aim to estimate the relative values of various types of risk reductions. They do not, however, provide a willingness to pay for either risk type and therefore are not included in our reviews of the willingness to pay literature above. Van Houtven et al. (2008) found a strong preference for avoiding cancer risks relative to automobile accidents even after controlling for latency and morbidity periods. With a 5-year latency, values for reductions in fatal cancer risk were approximately three times larger than those for immediate accident risks, declining to fifty percent larger for a 25-year latency. By contrast, in a study by Magat et al. (1996), the median respondent was indifferent between fatality risk from auto accidents and lymphoma, suggesting that cancer mortality is no more ‘dreaded’ than accidental mortality. It is difficult to draw firm conclusions, however, because the study did not specify the timing of the risks, and, in particular, any latency associated with cancer. Therefore, if respondents assumed that cancer risks would be realized after a latency period then the results suggest that any preference for cancer reductions was approximately offset by discounting future risks.

Three additional stated preference studies focus on WTP for cancer risks without direct comparisons to other risks. These do not internally address the question of how cancer risks are valued differently from non-cancer risks, but may be combined with the results from other studies to address this question. Focusing on cancer risks from hazardous waste sites Alberini, et al. (2010) estimated a cancer VSL of approximately $5.6 million (2009 dollars) using the results of choice experiments in Italy. Carson and Mitchell (2006) examined willingness to pay for installing a water filtration system to remove trihalomethanes (THM) in public drinking water. Estimated values depend upon an assumed latency and discount rate, as well as the specific risk reduction, but generally range from $3.4 to $8.0 at the smallest risk changes for a 25-year latency. Buzby et al. (1995) used a telephone-mail survey to examine the value of reduced fatal cancer risk from exposure to pesticides in grapefruit, and estimated a value of statistical cancer fatality at $6.99 million based on exposure assumptions.

Related Hedonic Property and Wage Studies

There are a small number of studies that have estimated WTP for reduced cancer risks using revealed preference approaches. The results have generally shown that the value of a statistical cancer case is similar to prevailing VSL estimates from hedonic wage studies. Direct comparison, however, is difficult without additional assumptions about latency or cessation lag and cancer fatality rates, as noted for each study.

In the context of hazardous waste, Gayer et al. (2000) and Gayer et al. (2002) employed a hedonic property framework to estimate the implicit value of a statistical cancer case from surrounding Superfund sites. In the first study, the value of a statistical cancer case was approximately $5.5 million, but did not include any assumptions or information on latency or fatality. The 2002 study calculated estimates under a variety of latency and discounting assumptions with results ranging from $5.2 million to $10.0 with no latency, and from $6.2 to $11.8 million using a 3% discount rate and 10-year latency period.

Davis (2004) used housing price responses to an observed cancer cluster in Nevada to estimate marginal willingness to pay for a change in lifetime pediatric leukemia risk ranging from $3.7 million to $11.1 million, which is generally consistent with the Gayer et al. studies, although the leukemia values are specific to children. Ho and Hite (2008) included risks from air toxics and hazardous waste sites in a hedonic property model and estimated the implicit value of cancer mortality to be $6.0 million. Finally, Lott and Manning (2000) explored the presence of compensating wage differentials for carcinogenic exposures in the workplace using the hedonic wage framework, finding that workers were being compensated for carcinogenic exposures. By making assumptions about the proportion of cancer deaths that arise from occupational exposures they calculated a cancer-specific VSL of $12.4 million.[13]

Because reducing environmental cancer risk is an important part of EPA's mission to protect human health, a key question is how the results from the empirical literature summarized here, along with other literature described in this report, can be systematically synthesized to account for individuals' preferences for reducing cancer risks relative to other types of health risks. As a first-cut, the simple average of the central estimates of the cancer differential from the subset of studies in Table 1 that reported values for both cancer and non-cancer risks is 52%.[14] This is a preliminary estimate and should be refined or replaced with a more systematic synthesis of the literature, possibly incorporating results

[13] As stated earlier, all figures have been updated to 2009 dollars using the Consumer Price Index, unless otherwise noted.

[14] Specifically, the summary point estimates that we drew from each of the nine studies in Table 1 that reported results pertaining directly to the cancer differential (i.e., VSLcancer / VSLnon-cancer - 1) are: 0 (Hammit & Hanninger 2010), 0.5 (Alberini & Scansy 2010a), 0.85 (Alberini & Scansy 2010b), -0.15 (Adamowiz et al. 2008), 0 (Cameron & Deshazo 2008), 0.2 (Tsuge et al. 2005), 0.3 (Hammitt & Liu 2004), 3 (Van Houtven et al. 2008), and 0 (Magat et al. 1996). The average of these figures is 0.52.

from other relevant studies. In the meantime, a cancer differential of 50% might be a reasonable placeholder value for use in upcoming RIAs.[15]

4 Review of stated preference and hedonic wage studies

Our reviews of the literature in the sections that follow focus on results from stated preference and hedonic wage and studies. This reflects where the majority of potentially relevant empirical estimates are found and is consistent with prior consultations and advisory reports. The hedonic wage approach is well-established and vetted and remains influential in informing guidance across the federal government. However, the approach is limited to work-related risks and the associated risk characteristics, many of which differ from EPA policy scenarios, as has been detailed many times in the economics literature.

There has been a tremendous growth in the number of stated preference studies to estimate values for mortality risk reductions in recent years; certainly there is now a far larger and more sophisticated body of literature to draw upon than was available at the time of EPA's last revision of its guidance. These developments potentially allow for an examination of important valuation dimensions including risk source (e.g., environmental, traffic-related); type of illness (e.g., any cancer differential or associated morbidity); and altruism. Our review of the empirical literature and how it can be synthesized attempts to address these issues.

However, additional studies exist that may supplement the reviews of the stated preference and hedonic wage literatures below. First, some stated preference studies do not seek to estimate willingness to pay or accept, but rather relative preferences for different types of mortality risk reduction. Two examples addressing cancer risks are described more completely above (Magat et al. 1996 and Van

[15] Another possible way to represent the cancer differential would be to estimate the absolute (rather than fractional) increment of the cancer mortality risk values over the values for non-cancer risks (i.e., VSLcancer - VSLnon-cancer). This would require an additional step of estimating the income elasticity of this absolute cancer differential. Estimating the fractional cancer differential implicitly assumes that the income elasticity of the absolute cancer differential equals that for the non-cancer VSL.

Houtven et al. 2008). The study results do not estimate willingness to pay, but it may be possible to combine the estimates from the studies on relative tradeoffs with the willingness to pay literature to refine our benefit transfers.

Another segment of the literature that we do not examine in detail here includes studies that evaluate only public preferences for risk reducing policies. Examples from this literature include Cropper et al. (1994) and Subramanian and Cropper (2000), who used survey methods to examine how respondents would allocate a given public budget to public programs for lifesaving and risk reduction; and Bosworth et al. (2009) who assessed community-level preferences for public programs to improve health and safety. The SAB previously concluded that these studies can be informative in their own right, but cannot be directly related to individual willingness to pay and used directly for benefit-cost analysis (USEPA 2001). EPA is open to suggestions on whether and how this literature may be effectively and appropriately synthesized with the results of other studies for the development of guidance on mortality risk valuation.

The hedonic property method has been used to estimate the value of environmental amenities and disamenities including mortality risks. A major challenge has been to limit the analysis to risk reduction rather than more comprehensive measures or indicators of environmental quality, such as air quality (e.g., Chay and Greenstone 2005) or the presence of or distance to hazardous waste sites (e.g., Greenstone and Gallagher 2008). These studies can be useful for evaluating some policies directly, such as the remediation of hazardous sites, but cannot be directly informative for mortality risk valuation. Willingness to pay for reduced mortality risks have been estimated in hedonic property studies, as first described and demonstrated in Portney (1981), who examined the relationship between housing prices and mortality risks from air quality. Four other studies, described more completely above in this paper, estimate marginal willingness to pay for cancer risk (Gayer et al. 2000, 2002; Davis 2004; and Ho and Hite 2008).

Finally, implicit values for risk reductions can be estimated in "averting behavior" studies, wherein an individual or household uses the good as an input into the production of health or safety. Blomquist (2004) conducted an extensive review of this literature and concluded, with some caveats, that the findings are broadly similar to hedonic wage estimates. Recent additions to the literature are generally consistent with this conclusion (e.g., Andersson 2005, 2008 (automobile risks); Hakes and Viscusi 2007 (seatbelt use)). Key concerns about averting behavior studies include issues of risk perception and the separability of joint benefits and costs (USEPA 2000b). Viscusi (1992) explicitly excluded these studies from consideration in his meta-analysis of VSL estimates. Further, the lack of available studies on environmentally-related risks limits the usefulness of this class of studies for the present purpose of developing guidance for mortality risk valuation.[16]

4.1 Stated preference studies

Stated preference (SP) is a survey-based method for estimating willingness to pay or accept for non-market goods or services. SP methods are widely used to value environmental amenities or improvements in human health endpoints that may be difficult or impossible to estimate using revealed preference methods because of long lag times, unclear causality, or other factors. For example, SP studies have been used to elicit willingness to pay for reductions in the risks of dying from cancer and cardio-vascular disease. SP studies vary widely in terms of the types of risk considered, payment vehicles, latency periods, mode of survey administration, etc. The number of and variation among existing SP studies is now large enough that the variation in their results can be analyzed statistically, although this involves a number of data collection and model estimation challenges.

[16] Note that there are some studies that relate averting behaviors to environmental quality or even related risks (e.g., Dickie and Gerking, 2009; Um, Kwak, and Kim, 2002), but, as documented in Blomquist, 2006, relatively few studies estimate WTP for reduced mortality risks in an environmental context.

4.1.1 Recent meta-analyses of SP studies

Three recent meta-analyses examined the stated preference literature using statistical methods. Kochi et al. (2006) used both stated and revealed preference studies in an empirical Bayes framework. Dekker et al. (2008) focused exclusively on stated preference studies, also with Bayesian methods. Braathen et al. (2009) conducted a meta-regression analysis of a wide variety of stated preference studies using classical econometric tools. Each of these studies is discussed in more detail below.

Kochi et al. (2006) used an empirical Bayes estimation method to generate predicted VSL estimates using multiple estimates from both stated preference and hedonic wage studies. Here we focus on the analysis and results for the stated preference data in their study. Study selection criteria were similar to those used by Viscusi (1992), including the use of studies for the general population and those conducted in high income countries only, and a minimum sample size.[17] Another important criterion was the use of estimates for immediate risk reductions; specifically, estimates for risks involving a latency period were excluded.

Kochi et al. analyzed 45 VSL estimates drawn from 14 stated preference studies. The authors recorded all estimates from each study and then separated them into "homogeneous subsets." Specifically, they grouped estimates by lead study author and used a Q-test for homogeneity to determine whether the estimates within a group are homogenoeous. After completing the separation of the estimates into homogenous subsets, they recalculated the VSL for the subset to create a unique VSL for that author. The recalculated mean reflects a weighted VSL of the estimates in the homogeneous subset, where the weights are based on the standard errors for the estimates.[18] This technique is intended to address the troubling issue of choosing among multiple estimates from each study when those

[17] Viscusi (1992) excluded two studies with sample sizes of around 30. Kochi et al. (2006) chose a minimum sample size of 100 for their analysis.

[18] Another implicit selection criterion in this study was the use of estimates with reported standard errors. In the assembly of our new meta-analysis dataset, described in Section 4.1.2 below, we find that this may be a highly constraining selection criterion.

estimates may be based on overlapping samples. The process of creating homogeneous subsets resulted in 18 stated preference VSL estimates with a mean of $3.5 million and a standard error of $0.67 million (in 2009 dollars).

Dekker et al. (2008) examined the influence of risk context (i.e., deaths from automobile-related accidents, air pollution, and all causes) on willingness to pay estimates from SP studies. The authors discussed the benefits transfer challenge associated with applying estimates from one context (e.g., auto risks) to another (e.g., air pollution), particularly when there is limited empirical evidence on the size and direction of the effects. Employing Bayesian techniques in a meta-regression, they compared willingness to pay or accept estimates in three different risk contexts—air pollution, traffic safety, and environment/general—while attempting to control for the size of the risk change and other respondent and study characteristics. Several study design decisions by Dekker et al. were based on recommendations from the EPA meta-analysis work group (USEPA 2006).

The authors used existing meta-analyses and additional literature searches to identify stated preference studies for auto, air pollution, or context-free (unspecified) mortality risk reductions. After searching the literature and applying screening criteria, a final database was assembled containing 98 VSL estimates from 27 studies, including three studies from the U.S. Seventy-one of the estimates were based on studies of road safety, seven on studies of air pollution, and twenty on studies of "general mortality" (presumably deaths from all, or unspecified, causes). The authors drew multiple estimates from each study, although it appears that they attempted to ensure that those estimates were from non-overlapping subsamples. Because of the small sample size that results from this approach they use Bayesian techniques suitable for these situations.

The analysis by Dekker et al. focused on explaining variation in willingness to pay for discrete changes in mortality risk reductions rather than the VSL and therefore includes as an independent variable the magnitude of the risk change associated with each estimate. They found that willingness to pay estimates are lower when the commodity is described as a public good and that there is a premium for risk reductions from air/general context over automobile risks.

Braathen et al. (2009) reviewed and conducted a meta-analysis of 75 studies with 900 estimates from developed and developing countries. The authors recorded a variety of attributes for each estimate: type of risk, country, survey mode, type of study, etc. The purpose of the study was to examine how these attributes influence the resulting VSL estimates. Using classical econometric techniques, their results show that methodological variables (i.e., type of payment questions, survey mode) explain 70 percent of the variation in the estimates. Of particular relevance to EPA, the authors found that health risks are valued lower than traffic and environmental risks, in contrast to the results of Dekker et al. However, risks to individuals are valued higher than risks to the public, similar to the results of Dekker et al. (2008). The work of Braathen et al. still is preliminary and, like the Dekker et al. meta-analysis, it includes studies from both developed and developing countries.

4.1.2 A new meta-analysis dataset

In an effort to both update the estimate or range of estimates used by EPA, we have constructed a new dataset containing information from a set of studies reflecting the current literature appropriate for application to U.S. environmental policy.[19] We used EconLit, conference proceedings, published and unpublished meta-analyses, working paper series, and personal contacts to identify and generate a comprehensive list of stated preference mortality risk valuation studies from 1974 and later.[20]

Each study was screened to ensure that it provided empirical estimates of the value of mortality risk reductions (i.e., purely theoretical studies and those that only examined morbidity were not included). Following the advice from the SAB-EEAC (USEPA 2007), we established a set of selection

[19] There is substantial overlap between our data set and those reflected in the meta-analyses reviewed in this section. Differences are due to different selection criteria and new studies that have appeared since the other meta-analysis studies were conducted.

[20] The earliest study that forms the basis of the recommendations of the existing EPA *Guidelines* (2000a) was conducted in 1974. Therefore, we limited our search for relevant literature to this starting date, assuming that the earlier literature had been vetted and judged to be obsolete prior to the release of the 2000 *Guidelines*.

criteria that determined which studies to include in our final data set. These criteria are based on information from other meta-analyses, as well as our own best judgment regarding study features necessary for application to valuing mortality risk reductions when analyzing U.S. environmental policies. The criteria we applied are as follows:

- minimum sample size of 100,
- sample frame based on general population,
- conducted in a high-income country,[21]
- results based on exclusive dataset,
- written in English,
- provides enough information to calculate a WTP estimate if one is not reported in the paper,
- provides estimates for willingness to pay (willingness to accept estimates were not included),[22] and
- provides estimates for willingness to pay for risk reductions to adults (estimates for risk reductions to children are not included).

We focus on studies with a sample size of at least 100 because smaller samples tend to suffer from small sample size problems (e.g., less precision) and are less likely to be representative of the general population. Because the purpose of this exercise is to determine an estimate or range of estimates for use in environmental policy, we limit our studies to those of the general population as opposed to specialized subgroups, like students or business owners. In addition, because our focus is on U.S. environmental policy we choose to limit our studies to those conducted in high-income countries. Socio-economic and cultural differences between the U.S. and most developing countries may be too large for reliable

[21] High-income countries are defined as having a gross national income per-capita of $11,906 (2008 US dollars) according to the World Bank reports (www.worldbank.org). The most recent World Bank data is for 2008.

[22] Three studies report willingness to accept estimates. These studies also report WTP estimates so we do not reject any study based solely on this criterion.

transfers of value estimates. Our own language limitations required that we restrict ourselves to studies written in English. Finally, we limit our investigation to willingness to pay estimates for adults only.

Thirty-three studies published between 1988 and 2009[23] meet the selection criteria described above, yielding nearly 450 willingness to pay estimates. For each of the studies we recorded all willingness to pay and value of statistical life estimates that were reported in the study, as well as those we could calculate based on information available in the study.[24] The meta-analyses using stated preference studies we described earlier draw multiple estimates from each study, and each has a different way to address the fact that these estimates are almost always drawn from overlapping samples (e.g., authors report multiple results from different estimation exercises or sub-samples within their data). However, we believe that the issues associated with using multiple estimates from each study are sufficiently problematic to warrant selection of independent estimates from each study.[25] Table 3 reports selected data for each study with detailed footnotes to describe the decisions to support the selected estimates.[26] This exercise results in 40 independent estimates. We report select characteristics for each estimates along with the willingness to pay and standard errors (reported in \$/μr). The willingness to pay for micro-risks are either directly extracted from the underlying studies (when the information was reported in the papers) or calculated by dividing the VSL estimates by 10^{-6} when the WTP estimates are not reported.

All estimates were recorded in the currency and dollar year presented in the study. If the dollar year was not noted or could not be gleaned from other information in the study then we assumed that it

[23] While we set a start date of 1974 for inclusion in our data set, only studies published after 1988 met our selection criteria.

[24] For the most part, all possible estimates were calculated or recorded for each study. We did not, however, record or calculate estimates for various levels of confidence respondents had in their responses, passing/failing quizzes about risk, and various forms of scenario rejection. We felt that these estimates were designed mainly to test the validity of the survey instrument and not to produce central estimates of mortality risk valuations per se.

[25] Later we discuss in detail the various issues associated with using multiple estimates and how this can be addressed econometrically.

[26] In general we opted for the estimate(s) that were the most inclusive of all the data in the study. Alternatively, we could select more estimates from each study – for example, by including estimates by age group – if this was determined to be an important dimension to the analysis.

was the year prior to the release or publication of the paper. All estimates are for individuals; when it was clear that an estimate reflected a household willingness to pay, we divided those estimates by the average household size for the country and year when the study was conducted. We then converted all estimates to U.S. dollars using the Purchasing Power Parity Index for the dollar year of the estimates. Next, all estimates were converted to 2009 dollars using the Consumer Price Index (CPI) and adjusted for income growth over time assuming an income elasticity of 0.5.

In addition to the willingness to pay estimates and standard errors (when available), we quantified and recorded as much information as we could for each study. Our data set includes whether or not the study was published in a peer-reviewed journal, the year it was conducted and published or released, the country where the study was conducted, sample characteristics, risk reduction information (e.g., magnitude, type of risk), scope tests, public versus private risk reductions, etc. See Table 2 for a description of many of the variables in our data set. Much of this information is only available for a subset of studies, particularly information on the demographic characteristics of the sample.

Twenty-two studies were published in journals, with 13 published in the *Journal of Risk and Uncertainty.* Six of the remaining studies are unpublished reports or working papers and five are book chapters. We identified nine different sources of mortality risk represented in the studies, including automobile accidents, air pollution, drinking water, hazardous waste sites, and food. The studies were predominantly conducted in the U.S. and Europe. Other countries represented in the data include Canada, Japan, Taiwan, and New Zealand.

Most of the studies are contingent valuation studies where the choice question involves stating a response (e.g., yes/no to a dichotomous choice question, open-ended response) to a scenario with a fixed set of attributes. Several studies are choice experiments in which respondents choose one option from several in which the attributes, including the magnitude of the risk reductions and the cost, vary across the options.

The average sample size for the estimates is 814 observations with a range of 13 to over 2,000.[27] Most studies were conducted with a self-administered mode via web-TV or a centralized computer facility. The second most common mode is an in-person survey. Other modes represented in the data include mail, telephone, and a combination of the two. A scope test was performed or calculated for about half of the estimates, and of those about 90 percent passed a weak form of the test (i.e., willingness to pay estimates exhibited a statistically significant increase with the size of the risk reduction, but was not necessarily proportional). Fifteen percent passed a strong form of the scope test (i.e., willingness to pay was proportional or nearly proportional to the size of the risk reduction).

4.2 Hedonic wage studies

In their recommendations to EPA, the SAB-EEAC and the Meta-Analysis workgroup clearly stated that both revealed hedonic wage and stated preference studies should be considered when deriving estimates of mortality risk values (USEPA 2006, 2007). Both groups also recommended that the two segments of the literature be analyzed separately. In this section we focus on the hedonic wage literature.

Hedonic pricing models use statistical methods to measure the contribution of a good's characteristics to its price. As applied to the labor market, hedonic wage studies (also known as compensating wage studies) are based on the premise that heterogeneous goods and services can be viewed as "bundles" of attributes and are differentiated from each other by the quantity and quality of these attributes. Fatal and nonfatal risks are among the many attributes that differ across jobs. All else equal, we would expect riskier jobs to pay higher wages. Therefore, it should be possible to estimate the value associated with reduced occupational fatality risk using data on wage and risk differentials among

[27] This is the sample size for the recorded estimates. Most studies used a subset of the data when recording different estimates (e.g., males only, younger respondents only). All studies meet the criteria of a minimum sample size of 100 respondents.

jobs, controlling for other factors that might influence the wage. A standard regression equation in the hedonic wage literature is

$$\ln w_i = \mathbf{X}_i \boldsymbol{\beta} + \phi p_i + u_i ,$$

where w_i is the wage for individual *i*, $\mathbf{X}_i$ is a vector of explanatory variables including various characteristics for the individual and her job, p_i is the probability of dying on the job, and $\boldsymbol{\beta}$ and ϕ are parameters to be estimated. If the prevailing wages are the result of a market equilibrium in which individuals have sorted themselves among jobs to optimize their individual-level trade offs between wages and risks, then the slope of the hedonic wage function with respect to the risk variable, $\partial w_i / \partial p_i$, will equal the individuals' marginal willingness to swap wages for risks.

4.2.1 Data sources

Some of the principal differences between hedonic wage studies arise from the data sources used to characterize workers and the job risks they face (Bellavance et al. 2009). Since no large data sets exist that contain both worker and risk information, researchers must match observations from various sources, which requires judgments on how best to combine data that are often reported at different levels of aggregation. Most hedonic wage studies conducted in the U.S. rely on one of two datasets for information on wages, other job characteristics, and worker characteristics: the Panel Study of Income Dynamics (PSID) and the Current Population Survey (CPS). Until recently, most studies had relied on two primary sources of risk characteristics: the Bureau of Labor Statistics (BLS) Survey of Working Conditions and the National Institute of Occupational Safety and Health (NIOSH) National Traumatic Occupational Fatality Survey. The BLS data are reported as annual counts of deaths by three-digit occupation or industry while the NIOSH data provide rates of death, averaged over five years, by one-digit occupation or industry by state. Users of these data necessarily consider risks by broad industry classifications (assigning all occupations within an industry the same risk) or by broad occupational classification (ignoring potential differences within an occupation across industries).

A number of recent studies, however, have turned to the Bureau of Labor Statistics' Census of Fatal Occupational Injuries (CFOI) as the source for workplace risk characteristics. The CFOI data are considered the most comprehensive data on workplace fatalities available (Viscusi 2004), compiling detailed information since 1992 from all states and the District of Columbia. Not only are the counts of these fatal events reported by 3-digit occupation and 4-digit industry classifications, but the circumstances of the fatal events as well as other characteristics of the workers involved (e.g., age, gender, race) also are recorded.[28] To ensure the veracity and completeness of the reported data, multiple sources are consulted and cross-referenced, including death certificates, workers' compensation reports and Federal and State administration reports. To form a complete dataset for estimation, these data still must be paired with worker samples drawn from another source (often the Current Population Survey) and fatality rates still must be constructed by the researcher using estimates of the number of workers, as with the other BLS data.

4.2.2 Estimation issues

Recently, EPA funded a study to examine the hedonic wage methodology and to provide a quantitative assessment of the robustness of the resulting value estimates for mortality risk reductions. The results of this research are summarized in Black et al. (2003) and were subsequently published in Black and Kniesner (2003). These studies examined the roles of the functional form of the estimating equation, measurement error, and unobservable characteristics using various commonly used data sets. Their findings highlighted a number of potential problems with previous hedonic wage studies. First, they found that estimates of the value of risk reductions can be very sensitive to seemingly minor changes in the specification of the regression equation. In fact, many specifications lead to negative estimates,

[28] More information on the CFOI data is available at: http://www.bls.gov/iif/oshfat1.htm.

which would suggest that people would be willing to accept *lower* wages for jobs with *higher* risks. They were unable to alleviate this problem using more flexible functional forms, so they concluded that this instability is not due to equation mis-specification. Instead, they found strong evidence that the job risk estimates contain considerable measurement error.

Black and Kniesner (2003) examined both the BLS SWC and NIOSH data sets (the CFOI dataset had not been widely used by that time). Their results indicate that, while both datasets have advantages and disadvantages, they both also are subject to considerable measurement error. They identified three sources of measurement error in the two data sets:

- sampling variation within industry and occupation cells given the small size of some of the cells (in recognition of this problem, BLS and NIOSH suppress data when the number of fatalities is low),
- heterogeneity in job risks and non-random assignment of those risks within occupations (e.g., late night convenience store clerks tend to be male and older), and
- industry and occupation are not measured accurately, especially at the three-digit level.

Moreover, they found that the measurement error is correlated with covariates commonly used in the wage equations and is likely correlated with the regression error as well. They concluded that studies that do not control for measurement error suffer from attenuation bias, resulting in under-estimates of mortality risk values. They also concluded that the NIOSH data produce results most consistent with economic theory.

4.2.3 Recent meta-analyses of hedonic wage studies

In addition to the methodological assessment conducted by Black and others, several meta-analyses of the hedonic wage literature have been conducted in recent years. We focus here on four recent studies, three of which were reviewed by the Meta-analysis workgroup convened by EPA. The fourth was published after their deliberations.

Mrozek and Taylor (2002) used multiple observations from 47 hedonic wage studies. Variables included in their meta-regressions were of three types: (1) those which may influence wage/risk tradeoffs (e.g., mean hourly earnings), (2) those describing the sample, and (3) methodological choices of the original researchers (e.g., if a risk-squared term was included in the estimating equation).

The authors used weighted least squares where the weights were the number of estimates provided by the study. This ensured that each study was weighted equally, regardless of the number of observations drawn from it. Four meta-regression models were estimated, each using log(VSL) as the dependent variable. All four models indicated a positive and significant relationship between the mean risk and VSL. The authors used the meta-analysis results to "predict" the VSL as if the original studies had all followed a set of "best practice" assumptions. The predicted values range from $1.78 million to $15.4 million (2009 dollars). Those assuming the use of National Institute for Occupational Safety and Health (NIOSH) data are higher than those assuming use of Bureau of Labor Statistics (BLS) data. The authors concluded that the evidence best supports an estimate of $2.69 million at the average occupational risk level of 0.5 per 10,000 (2009 dollars).

While this study provides a comprehensive overview of the hedonic wage literature, it includes studies using older (and possibly unreliable) occupational risk data. In addition, the authors excluded estimates in original studies that were statistically insignificant or negative.

Viscusi and Aldy (2003) conducted a review of more than 60 hedonic wage studies of values for mortality risk reductions across 10 countries (including 52 from the U.S.), examining a number of econometric issues, the effects of unionization on risk premiums, and the effects of age and income on VSL estimates. No studies were eliminated from the sample, and no attempt was made to modify the original VSL estimates. Point estimates extracted from each study were those based on the "whole sample" and the original authors' preferred model specification. Viscusi and Aldy generated summary VSL estimates by using the estimated coefficients from the meta-analysis to predict the natural logarithm of VSL for each original study, then study-specific predicted-VSLs were averaged to produce an overall mean estimate. Predicted U.S. mean values were constructed based on regression samples using all countries, but with averaging across U.S. studies only. The predicted values in the study for the U.S. range from $6.85 million to $9.47 million (2009 dollars), and the median predicted values were generally very close to the means.

Kochi et al. (2006) used an empirical Bayes estimation method to generate predicted VSL estimates based on previous hedonic and stated preference studies. Here we focus on the analysis and results for the hedonic wage data. Using selection criteria similar to those from Viscusi (1992), the analysis included 162 VSL estimates from 31 hedonic wage studies. All possible VSL estimates and associated standard errors for each included study were re-estimated based on information provided in each original study. Estimates without standard errors were not included. The homogeneous subsetting method described earlier also was applied to the hedonic wage estimates (the hedonic and stated preference data were analyzed together), resulting in 42 VSL estimates from hedonic wage studies with a mean of $11.96 million and a standard error of $0.62 million (2009 dollars). Because of the subsetting technique employed to pool the estimates, Kochi et al. could not explicitly account for study design and population characteristics in their analysis.

Bellavance et al. (2009) is the most recent meta-analysis of the hedonic wage literature. The authors' principle objective was to better understand the variability in VSL estimates from hedonic wage studies, which is described as ranging from $0.5 to $50 million. Thirty-nine VSL estimates from 37 studies were assembled based on those used in prior meta-analyses and further searches of several economics databases. The resulting dataset contains sixteen studies from the U.S., seven from Canada, and three or fewer from each of several other countries. The earliest study is from 1974 and the most recent is Viscusi (2004).

The authors draw only one VSL estimate from each study. Standard errors were recorded or computed for 32 of the 39 estimates. Criteria were established to chose the specification within each study, including: (1) no interaction terms between the probability of death and other explanatory variables (in order to more easily compute the standard error), (2) similarity of specification to other included studies, (3) larger samples with characteristics most similar to other studies, and (4) the recommendations of authors of prior meta-analyses. Bellavance et al. acknowledged that the source for U.S. risk data varies and has evolved over time from early BLS surveys to NIOSH to BLS' Census of Fatal Occupational Injuries (CFOI). However, their analysis did not control for the data source other than for the use of Society of Actuaries (SOA) data, which was found to have a significant impact on the estimated VSL. Sensitivity analyses were conducted with and without studies using SOA risk data.

Using a mixed effects model (random intercept with fixed effects for study characteristics), the authors regressed the VSL estimates on average income, probability of death, and several study design variables. The mean weighted average VSL is approximately $7.23 million (2009 dollars). Other key findings include that the VSL is significantly higher for studies that treat risk as endogenous, and there is some evidence that the VSL declines with the baseline risk.

4.2.4 A new meta-analysis of hedonic wage studies

Using Appendix 1 from Bellevance et al. as a starting point, we constructed a new data set of hedonic wage studies, augmenting the information contained therein with data from Kochi et al. (2006) and Viscusi and Aldy (2003). We also conducted a full text search in JSTOR for "Census of Fatal Occupational Injuries" and "CFOI" in order to develop a comprehensive list of studies using these data. A total of 14 CFOI studies were reviewed, with those actually using the CFOI data in an original, hedonic wage analysis retained for further assessment. These seven studies were further augmented with an unpublished manuscript using the CFOI data, for a total of eight additional studies.

Additional searches were conducted in JSTOR for studies published in 2000 or later using the key words "hedonic wage" and "compensating wage." We also conducted a search in the Social Science Citations Index for studies citing Viscusi (2004), a paper that derives mortality risk valuation estimates controlling for occupation and industry using the CFOI data.

In constructing our data set, we generally employed the same selection criteria used in Bellavance et al. (2009), with some exceptions based on our own judgment and to ensure consistency with the criteria used for the stated preference data set. First, we limited our data to those studies with a sample size of 100 or more. We also retained only those studies conducted in a high-income country as defined by the World Bank. Third, we omitted studies that rely on Society of Actuaries data as the source of risk information as these data are thought to reflect broader risks than those experienced on the job (Viscusi 1992, Kochi et al. 2006). We further limited our data by excluding those studies that focus on extremely dangerous jobs (e.g., police officers), since the risk preferences of individuals who take these jobs may differ substantially from those of the general public. We do, however, apply the other selection criteria employed by Bellavance et al., including retaining only those studies using a model specification similar to that given near the beginning of this section, excluding studies based on specific causes of death, excluding studies using the same samples as other studies, and excluding studies failing to report enough information to calculate the value of mortality risk reductions and/or the average probability of death. Applying all of these criteria resulted in the selection of 37 studies.

For each of our selected studies we recorded the following key variables: year of publication, the country in which the sample was drawn, sample size, average income, average annual probability of death, source of risk information, the estimated coefficient on the risk variable, whether the sample was exclusively male, manufacturing, blue collar and white, as well as whether the regression controlled for nonfatal risks, union status, and worker compensation. We calculated VMRs for each study by deriving the VSL and dividing these estimates by 10^6. As with our stated preference data, all estimates are reported in 2009 dollars after adjusting for inflation using CPI and accounting for income growth over time assuming an income elasticity of 0.5.

Similar to the stated preference data, we capture only one specification per study in our database, following the criteria established by Bellavance et al.[29] Because the hedonic studies are more homogeneous in their design than the stated preference studies, we are able to be more selective in which specifications to include. Although one motivation here is to minimize the influence of each individual study, it does not necessarily rid us entirely of the problem of overlapping subsamples as many of the studies draw their samples from the same source.

Table 4 lists key characteristics for our selected studies.[30] A total of 24 studies out of 37 were conducted in the U.S. with 3 using NIOSH data, 13 using BLS data and 8 using CFOI data as the source of occupational risk. Seven of these twenty-four studies rely on the Panel Study of Income Dynamics (PSID) as a source for worker characteristics with another 11 using CPS data. Twenty-six studies included women in their samples and 7 focused on blue collar workers only. Three studies restricted their samples to union members only. Average sample size across studies was 17,741, and the average income was $40,508 per year (2009 dollars). The mean probability of occupational death across studies was 0.00014.

5 Income Elasticity Considerations

EPA first attempted to address the income elasticity of VSL issue in its analysis of *The Benefits and Costs of the Clean Air Act, 1990 to 2010* (US EPA, 1999), which made a distinction between application of income adjustments for longitudinal changes in income over time and cross-sectional income differences for benefit transfer. The report applied a range of VSL income elasticities in a sensitivity analysis to project the value of reduced mortality risks in the year 2010.

[29] Note that some hedonic studies report results for multiple non-overlapping subsamples (e.g., male vs. female, union vs. non-union) within the study. Rather than capture these multiple observations, we have elected to implement the selection criteria used by Bellavance et al.

[30] Information reported in the table was adapted from Bellavance et al. (2009).

The issue was further developed in EPA's White Paper *Valuing the Benefits of Fatal Cancer Risk Reductions*, where income was one of the many benefit transfer issues to be addressed. The SAB-EEAC review of the White Paper concluded: "With regard to population characteristics, the Committee believes that it is appropriate to adjust the value of the projected statistical lives saved in future years to reflect higher incomes in those years, but not for cross-sectional differences in income, because of the sensitivity of making such distinctions."[31] The SAB-EEAC recommended that any appropriate adjustments for income growth should be part of the Agency's main analysis.

Based on a review of the empirical literature on the cross-sectional income elasticity of VSL literature originally developed for use in *The Benefits and Costs of the Clean Air Act, 1990 to 2010* report, EPA analyses have typically applied a range of estimates with a low end of 0.08, a central value of 0.4, and a high end of 1.0. Many analyses characterize this range with a triangular distribution with a resulting mean estimate of approximately 0.48. Income elasticity is then typically paired with projections of growth in real US GDP per capita.

More recent information on the income elasticity of VSL has come primarily from meta-analyses of hedonic wage studies. The results in Mrozek and Taylor (2002) suggest income elasticities ranging from 0.37 to 0.49, although the authors note that these results should be interpreted with caution because of measurement error in the income variable and the functional form used by many hedonic wage studies included in their meta-analysis. As described earlier in this paper, more recent work from Viscusi and Aldy (2003) estimates the income elasticity of the VSL in the range of 0.5 to 0.6, slightly higher than the mean value used in many EPA analyses. None of the 95 percent confidence bounds on the Viscusi and Aldy estimates include a VSL income elasticity as high as 1.0. The Bellevance *et al*. (2009) meta-analysis,

[31] "An SAB Report on EPA's White Paper *Valuing the Benefits of Fatal Cancer Risk Reduction*," US EPA, 2000, page 7. A 2007 SAB review also noted the empirical difficulties of accounting for differences in real income and wealth across populations due, in part, to "uncertainty about the value(s) of income elasticity and very little empirical evidence concerning the relationship between wealth and mortality valuation." US EPA 2007, page D-7.

also described earlier, predicts somewhat higher elasticity estimates ranging from 0.84 to 1.08 depending upon the model.

Some recent theoretical research has examined the relationship between the income elasticity of the VSL and the coefficient of relative risk aversion and noted that these two quantities should be very close in magnitude. This can be seen most easily in a simple two-period model. Let "lifetime" utility be the expected discounted sum of utility in both periods: $U = u_1 + p\beta u_2$, where u_t is utility in period t, p is the probability of survival between periods 1 and 2, and is the utility discount factor. Also assume that u_t depends on income in period t and takes the standard "constant relative risk aversion" (CRRA) form: $u_t = \left(y_t^{1-\eta}\right)/(1-\eta)$, where η is the coefficient of relative risk aversion. The VSL is the marginal rate of substitution between the individual's first period income and her probability of survival to the second period, i.e., $VSL \equiv \left(\partial U/\partial P\right)/\left(\partial U/\partial y_1\right) = y_1^{\eta}\beta\left(y_2^{1-\eta}\right)/(1-\eta)$, and so the income elasticity of the VSL is $\left(dVSL/dy_1\right)\times\left(y_1/VSL\right) = \eta$. Kaplow (2005) examined a more realistic version of this model by allowing for self-defensive expenditures that could increase the individual's survival probability. Using that elaborated model, Kaplow showed that the income elasticity of the VSL should be at least as large as 1 when $0 \leq \eta < 1$, and at least as large as η when $\eta \geq 1$.

Empirical estimates of the coefficient of relative risk aversion span a wide range—from around 0.5 to 1 at the lower end (e.g., Shepard and Zeckhauser 1984, Eeckhoudt and Hammit 2001, Chetty 2006) to 10 or more at the high end (e.g., Kocherlakota 1990)—but most estimated or assumed values for η seem to fall in the range of 1 to 3. For example, Hall and Jones (2008) and Hall (2010) estimated η to be around 2, based on the recent trend of income growth and the more rapid growth in health care expenditures in the United States. Szpiro (1986), Feldstein and Ranguelova (2001), Barro (2006), and Layard *et al*. (2008), among others, also estimate or use values of η in this range. And in the contemporary climate change economics literature, the most commonly used values of η are 2 to 3 (e.g., Arrow 2007; Nordhaus 2008; Dasgupta 2008; Weitzman 2009, 2010a,b).

The theoretical considerations combined with (most of) the empirical estimates of relative risk aversion cited above are at odds with the early estimates of the income elasticity of the VSL in the neighborhood of 0.5 cited above. In a more recent study, Kneisner *et al*. (2009) applied a quantile regression approach to a dataset assembled from the Panel Study of Income Dynamics (PSID) and the Census of Fatal Occupational Injuries (CFOI). Their preferred regression model produced estimates of the income elasticity of the VSL between 1.23, for the lowest quantile, to 2.24, for the highest quantile. Kneisner *et al*. note that "Our estimates of a large income elasticity of VSL are consistent with the simple theoretical models that have been developed [by Kaplow (2005)]," and "With recent estimates of the coefficient of relative risk aversion being around 2 based on the labor supply analysis of Chetty (2006) and the consumption analysis of Kneisner and Ziliak (2002), one would expect the VSL to be income elastic, which is what the results above indicate."

Based on theoretical considerations such as those examined by Kaplow (2005) and the new empirical results of Kneisner *et al*. (2009), EPA believes that its recommended estimate of the income elasticity of the VSL appears to be on the low end of the range of estimates and may need to be updated to a higher value or range of values.

6 Methods for Combining Data

The values for mortality risk reductions estimated in the stated preference and hedonic wage studies described above constitute a current empirical summary of the literature, which can be used to inform the revision of EPA's mortality risk valuation guidance. These studies could be combined or synthesized in a number of ways, from a simple point estimate to range, distribution, or systematically combined in a more rigorous meta-analysis. Our objective in this section is to outline analytical options that can be implemented in the longer term for updating the estimate or range of estimates used by EPA in our guidance on valuing morality risk reductions. We begin with meta-analysis methods, including methods similar to those used in our current guidance and extending to more rigorous application of meta-regression techniques. This is followed by the structural benefit transfer approach, which involves calibrating a direct or indirect utility function so that it is consistent with summary estimates of values for health risk. Our goal is to provide enough information on the analytical options and key issues to receive clear recommendations from the SAB-EEAC on an approach to implement for updating our guidance and on future research directions.

6.1 Meta-analysis

There are several options for obtaining simple summary statistics or ranges from the existing data. We outline these options and key issues in order of increasing complexity.

6.1.1 Parametric distribution

EPA's current guidance took one best estimate from each of five stated preference and twenty-one hedonic wage studies and then fit a parametric distribution to the values. The resulting mean and distribution has become EPA's default estimate for valuing mortality risk reductions. To replicate this approach we could use the databases of SP and HW studies discussed above and then separately characterize the resulting distributions in a curve-fitting exercise. Based on these distributions we could define a range of default values for the value of mortality risk for EPA policies. Key choices and principles are:

- *Use all "independent estimates" from the studies rather than one estimate per study*. Because many studies provide estimates for different subpopulations or other treatments, we can often include multiple study estimates without gross violations of independence. An alternative is to rely upon a single estimate per study, which has been done for several meta-analysis.

- *Update all study estimates to a common year, including the effect of real income (GDP per capita) growth over time and the estimates income elasticity of the VSL.* The review of the literature in the prior section already includes this update.
- *Limit SP study estimates to those that are non-cancer and non-latent.* In so doing, we will produce a "base value" that should be more consistent with estimates stemming from the hedonic wage literature. We will attempt to address any systematic difference in value between reduced cancer risks and other types of risk separately. In part, this is simply recognizing that EPA policies affect both cancer and non-cancer mortality risks and different values for each may be appropriate. Similarly, EPA policies address risk reductions varying from the near-immediate to those delayed over many years, a benefit-transfer aspect that we address by discounting over estimated latency periods. Including latent risks in this simple aggregation would double-count the effects of timing on value.
- *Include public-risk studies or rely only on private-risk SP studies*. Most EPA regulations result in public risk reductions. To avoid under-counting benefits, we would want to err toward inclusion, basing guidance on the full set of relevant studies including those that incorporate altruism even if we cannot distinguish whether it is paternalistic or non-paternalistic. On the other hand, to avoid double-counting of benefits we would want to use only those studies that capture private willingness to pay for mortality risk reduction. Clear recommendations from the EEAC on this issue in particular would be very helpful.

6.1.2 Classical econometrics

A second approach to combining the information from multiple studies—to determine the characteristics of the studies that influence the value estimates or to generate a benefit transfer function—is to perform a meta-regression using classical econometrics. Two issues arise when considering this approach. First, the analyst must decide which observations to include in the analysis. Some previous meta-regression studies have used all relevant observations in the analysis (e.g., Nelson and Kennedy

2008, Braathen et al. 2009, Mrozek and Taylor 2002). This approach incorporates all available information, but runs the risk of including estimates from overlapping samples (and therefore non-independent observations). For example, the same individual(s) may be represented multiple times in the data when a paper reports multiple estimates using different modeling assumptions. Restricting the data to non-overlapping samples is a non-trivial exercise because choosing the most appropriate estimate(s) from each study involves subjective judgment. In addition, small sample size problems—already a hurdle in meta-analysis—are exacerbated when the sample is limited in this way. The stated preference and hedonic wage meta-analysis datasets described in Section 4 draw independent samples based on procedures outlined above. However, a very strict interpretation of the requirement for non-overlapping subsamples for the hedonic wage studies could result in just a handful of estimates for use in a meta-analysis given the reliance by authors on the same sources of data.

Second, there are econometric issues to consider when analyzing these data. Nelson and Kennedy (2008) discuss “factual” versus “methodological heterogeneity.” Factual heterogeneity arises because of real differences in what the primary studies are measuring. For example, the *wtp* for auto risks may factually differ from that for cancer risks. Similarly, the *wtp* for occupational risks for male blue-collar workers may factually differ from that estimated for a more inclusive sample. Methodological heterogeneity arises because of different study design choices, such as the use of different models to estimate willingness to pay. When these sources of heterogeneity are unobserved, errors may be correlated. It also is likely that estimates produced by different surveys and designed by different authors have different variances, making heteroskedasticity a concern. Classical econometrics provides several approaches for dealing with correlated errors and heteroskedasticity. A fixed effects model assumes that the unobserved heterogeneity among studies can be captured with an intercept shift. By including a dummy variable for all but one of the studies, the intercept shift is estimated directly. This approach can result in low degrees of freedom if each study contributes a small number of estimates. An alternative approach that does not require a new independent variable for each study is the random effects model. Using the "composite error" exposition of the random effects model, the estimating equation is

$$y_{ij} = \mathbf{x}_{ij}\boldsymbol{\beta} + \varepsilon_{ij},$$

where y_{ij} is *WTP* estimate *j* from study *i,* $\mathbf{x}_{ij}$ is the row of data for that estimate, and $\boldsymbol{\beta}$ is a vector of coefficients. The error term ε_{ij} has the following structure

$$\varepsilon_{ij} = u_i + \nu_{ij}, \text{ where } u \sim N\left(0, \sigma_u^2\right) \text{ and } \nu \sim N\left(0, \sigma_\nu^2\right).$$

ε is a composite error term with components *u,* which can vary between studies but has the same value within studies, and ν, which can vary within and across studies.

If heteroskedasticity also is a concern, there will be two potential violations of the classical assumption of spherical errors. In this setting, coefficient estimates will be consistent but inefficient and standard error estimates may be inconsistent. One solution to this problem is to estimate the model in the traditional way but calculate standard errors that are robust to heteroskedasticity. White standard errors are robust to heteroskedasticity (Greene 1997 p 503-505), and there is a class of robust standard errors that imposes the panel structure on the calculation and is thus robust to correlation within clusters as well. Statistical packages such as Stata and SAS are able to produce "cluster-robust" standard errors.

A second way to address heteroskedasticity is via weighted least squares, where the weights are inversely proportional to the variance of the willingness to pay estimate. However, since a number of the studies in our meta-analysis datasets do not provide standard errors of the estimates, we can use the number of estimates drawn from a single study (assuming we draw multiple estimates per study) as a proxy for variance. The rationale for this proxy is that studies reporting a large number of estimates are more likely to report all possible willingness to pay estimates based on different characteristics of the sample, versions of the survey, etc., and these estimates may be less precise than those from a study that presents a few, select estimates. Mrozek and Taylor (2002) used this approach, as discussed above. This insures that each study is given equal weight, as opposed to each estimate. The sample size for each estimate also could be used to generate weights. Observations that arise from larger samples should be more precise, all else equal. However, sample sizes are not available for all observations in our meta-analysis datasets. Mrozek and Taylor (2002) used the level of significance of the VSL estimate to create a *t*-statistic weight in an appendix to their paper. The estimating equation for this approach is:

$$\frac{1}{n_i} y_{ij} = \frac{1}{n_i} x_{ij}\beta + \varepsilon_{ij},$$

where n_i is the number of estimates or sample size from the i^{th} study. This technique provides more efficient estimates than unweighted estimation of the analogous model.

Considering the data issues common to meta-analyses of willingness to pay estimates for mortality risk reductions, we propose two classical approaches meant to address both heteroskedasticity and correlated errors arising from unobserved study heterogeneity when multiple estimates are drawn from each study. Weighted least squares estimation, as discussed above, can correct for heteroskedasticity. However, relevant statistics may not be reported to construct the ideal weights. If weighted least squares is used, we suggest testing for heteroskedasticity and using standard errors that are robust to clustering. Alternatively, one could estimate a study-level panel model to account for unobserved heterogeneity and calculate standard errors that are robust to heteroskedasticity. Since many studies provide just a few estimates, a fixed effects model may not be feasible while a random effects model would preserve degrees of freedom. We are particularly interested in EEAC comments on these alternatives.

6.1.3 Bayesian estimation

In the previous section we discussed how classical estimation techniques could be used to estimate a meta-regression of values for reductions in mortality risks while addressing heteroskedasticity and correlated errors. However, if we use data sets with non-overlapping estimates—as has been recommended by the Meta-analysis workgroup, and as is reflected in the summary of stated preference and hedonic wage estimates in Tables 3 and 4—our data selection criteria leave us with relatively small samples for meta-regression. The combination of small sample size and non-spherical errors presents a particular problem for classical approaches to estimation. Specification tests, including those for heteroskedasticity, and calculations of robust standard errors rely on asymptotic relationships and therefore may not be reliable when the sample size is small (Moeltner and Woodward 2009). Bayesian estimation has desirable small sample properties and can more easily accommodate general error structures.

Bayesian analogs to the classical approaches discussed above have been developed and can be used to estimate a meta-regression model to improve value estimates and provide richer inference into the results. Koop (2003 p 124-129) presented a Bayesian pooled regression model with an error structure general enough to be robust to correlated errors and heteroskedasticity even when the form of heteroskedasticity is unknown. Moeltner and Woodward (2009) use this model to estimate a meta-regression of wetland valuation estimates from a sample of just 12 values from 9 studies. They use Gibbs sampling to estimate the model

$$y_j = \mathbf{x}_j\boldsymbol{\beta} + \varepsilon_j \quad \text{with} \quad \varepsilon_j \sim N\left(0, \sigma^2\omega_j\right), \quad \text{and} \quad \omega_j \sim IG\left(\frac{\nu}{2}, \frac{\nu}{2}\right),$$

where y_j is WTP reported in study j, $\mathbf{x}_j$ is a row vector of population and other characteristics associated with study j, $\boldsymbol{\beta}$ is a vector of regression coefficients, ε_j is a zero mean regression error with variance $\sigma^2\omega_j$, and *IG* denotes the inverse-gamma distribution. This approach allows the authors to estimate study-specific variances by estimating a single parameter ν and drawing ω_j in a data augmentation step. Moeltner and Woodward (2009) showed that Bayesian estimation can be used to conduct meta-regression on small heteroskedastic samples and produce consistent and efficient parameter estimates.

A Bayesian analogue to the study-level panel model is also developed by Koop (2003 p 149-157). Bayesian estimation of a study-level panel model with a *non-hierarchical prior* is analogous to the fixed effects model in classical econometrics because the unobserved heterogeneity between studies is attributed to a constant (intercept shift) for each study. If the number of studies is large relative to the number of estimates from each study then, just as would be the case under classical assumptions, the high-dimensional parameter space can be problematic. In these cases it may be beneficial to use a *hierarchical* prior which places more structure on the unobserved heterogeneity by assuming the study-level effects can be drawn from a distribution, thus only the parameters of that distribution, and not the individual effects themselves, need to be estimated. Bayesian estimation of a panel model with a hierarchical prior is analogous to the classical random effects panel model. In both cases the error structure imposed on the model is general enough to be robust to non-spherical errors due to correlation within studies and heteroskedasticity.

6.2 Structural benefit transfer

Thus far we have discussed meta-analysis, including classical and Bayesian approaches to estimating a meta-regression model, which then could be used for functional benefit transfers. Using meta-regression, the form of the estimating equation, and therefore the transfer function, typically would be based on a combination of statistical tests and qualitative theorizing about the important variables to include in the model. The resulting function can be viewed as a low-order Taylor series approximation to the "true" preference function within the range of the data used to estimate it.

In contrast to the meta-regression approach, structural benefit-transfer (also known as preference calibration) involves first specifying a direct or indirect utility function for a representative individual, then deriving analytical expressions for observable economic outcomes from the utility function (Smith et al. 2002, 2006). Such observable outcomes could include labor-leisure tradeoffs, demand for related market commodities, equilibrium wage schedules for jobs with differing risk or other characteristics, responses to stated preference survey questions, etc. The parameters of the utility function are calibrated using data on such outcomes, and the calibrated model then can be used to predict willingness to pay or accept for any policy changes that can be described by variations in one or more of the parameters that appear in the calibrated preference function.

The key advantages of the structural benefit transfer approach are that it provides a means of combining estimates from separate studies that use different benefit concepts (e.g., marginal or non-marginal willingness to pay or accept, consumer surplus, compensating or equivalent variation, etc.), and it assures the economic consistency of transfers (Smith et al. 2002, 2006). In this context "economic consistency" means, for example, that estimated willingness to pay will never exceed income, that value estimates will always be responsive to scope (the size of the postulated change in quantity or quality), that *WTP* and *WTA* will always stand in the proper relationship to each other, and so forth. The way that such consistency is achieved is through the ex ante imposition of a specific form for the utility function, from which all subsequent value estimates and behavioral responses are then derived. One way to think about the contrast between meta-regression and structural benefit transfer is that the former uses relatively more data and fewer theoretical assumptions, while the latter uses relatively fewer data (or more highly aggregated data) and stronger theoretical assumptions. Therefore, the meta-regression approach may give more accurate value estimates within the range of the data used to estimate the function, while the structural benefit transfer approach may be more accurate in out-of-sample transfers. Thus, the choice of one approach over the other may depend in part on whether the policy case(s) to be examined fall largely within or largely outside of the range of data available for a meta-regression transfer function.

6.2.1 Static preference functions

A simplistic example may help clarify the structural benefit transfer approach. Here we follow Smith et al. (2003, 2006) and use a static model of the tradeoff between income and survival. (In the next sub-section we will consider a more general dynamic life-cycle model.) Assume that utility conditional on survival is proportional to the log of scaled income, so expected utility is $U = p\ln(aY)$, where *p* is the individual's survival probability. Using this functional form, the marginal willingness to pay for an increase in the probability of survival is $wtp = (\partial U/\partial p)/(\partial U/\partial Y) = Y\ln(aY)/p$. Next suppose that, based on a comprehensive review of the hedonic wage literature, *wtp* is estimated to be \$8/μr (i.e., the VSL is \$8,000,000) for individuals with average annual income 35,000 \$/yr and average annual survival probability p = 0.984. This allows calibration of the single unknown parameter of the utility function: $\ln a = pVSL/Y - \ln Y$ = 214.5, which gives a function that can be transferred to individuals with different background mortality risk levels. This function could vary by age and other personal and environmental characteristics, and/or different income levels by adjusting *p* and/or *Y*, respectively. Using this functional form, *wtp* is inversely proportional to the baseline survival probability (and therefore increases with the background mortality risk) and is (nearly) proportional to income.

We also can use the calibrated utility function to calculate willingness to pay for changes in mortality risks of any magnitude, rather than relying on the first-order approximation represented by the *wtp*. In this case the willingness to pay function is $WTP = Y - \exp\left[p\ln(aY)/(p+\Delta p)\right]/a$. Note that for large enough Δp's the marginal approximation may exceed total income while the actual *WTP* cannot.[32] As noted by Smith et al. (2006), this is one of the key advantages of a structural benefit transfer

[32] Letting Δp go to its maximum value $1-p$ gives $WTP = Y\left[1 - 1/(aY)^{1-p}\right]$, which is necessarily less than *Y*. Also note that, in this model, the smaller is *p* the larger is *WTP*, approaching *Y* as *p* goes to zero. This gives a simple illustration of the "dead-anyway effect" (Pratt and Zeckhauser 1996).

approach: it can produce more realistic predictions of *WTP* well outside of the range of data used to estimate marginal willingness to pay. (Additional numerical examples are provided in Appendix A.)

Another advantage of the structural approach is that it can help to account for potential behavioral responses. We can illustrate this by extending the simple model given above. Again following the hedonic wage literature, suppose that wages, *W*, are an increasing function of job-related mortality risk, *m*. Specifically, suppose that $W = W_0 + \alpha m^{\beta}$. Total income is comprised of wages plus non-wage income, *y*. With this extension, expected utility is $U = \ p_0 - m \ \ln \ a \ \ y + W_0 + \alpha m^{\beta}$, where p_0 is the background (non-job related) survival probability. Now suppose that after careful examination of the hedonic wage literature we estimate that, for a sample of individuals of prime working age (say, around 40 years old), *y* = 5,000 \$/yr, *W* = 30,000 \$/yr, p_0 = 0.99, *m* = 0.006, and $wtp = \partial W / \partial m$ = \$8/μr. So, for example, if β = 0.5, then $\partial W / \partial m = \beta \alpha m^{\beta - 1} \Rightarrow \ \alpha = \ \ 8 / 10^{-6} \ \ / 2 / 0.006 = \ 6.67 \times 10^{8}$ and $W_0 = W - \alpha m^{\beta} = -8.6 \times 10^{9}$. (Note that with two estimates of *wtp* at two levels of job risk, we could calibrate α and β simultaneously.) Now recall the standard assumption underlying the hedonic wage literature that the individual has chosen her job-risk level optimally, and assume she is able to adjust that level to re-optimize her expected utility after a policy intervention changes p_0 by some amount Δp. To determine the maximum willingness to pay for an exogenous change in mortality risk, we must solve the two-equation system comprised of (1) the equality between expected utility with and without the policy, and (2) the first-order condition for maximized expected utility with respect to job-risk with the policy and a reduction in income equal to *WTP*.

Results from some simple numerical experiments with this model are given in Appendix A. The main lesson from these examples is that if individuals are able to adjust their job risk level, then *WTP* generally will be higher and the total number of "statistical lives saved" will be lower than otherwise predicted under the assumption of no behavioral response. The numerical examples in Appendix A are not intended to represent any specific real-world case; nevertheless, they clearly illustrate that the structural benefit transfer approach is able to capture these effects.

6.2.2 Life-cycle preference functions

The structural benefit transfer function illustrated above was based on the simplifying assumption that the representative individual looks ahead only one period at a time—that is, utility depends only on the probability of survival to the next period and expected consumption in the next period. A more realistic framework would account for expectations of survival and consumption in all future periods. This brings us to the life-cycle consumption modeling approach. A life-cycle consumption model represents consumption-versus-saving (and possibly other) choices by an individual over the course of her lifetime. Life-cycle models are inherently dynamic, with age-specific mortality probabilities included as key parameters. Individuals are assumed to maximize the expected present value of discounted utility, where the expectation is conditional on the probabilities of living to all possible future ages (e.g., Yaari 1964, Shepard and Zeckhauser 1984, Rosen 1988, Cropper and Sussman 1990, Ehrlich 2000, Johansson 2002, Aldy and Smyth 2006, Murphy and Topel 2006, Hall and Jones 2007, USEPA 2007 p. 14-16).

A life-cycle consumption modeling framework could be used as the basis for a generalized structural benefit transfer function. Such a transfer function would allow calculation of willingness to pay for any marginal or non-marginal changes in the individual's mortality profile (i.e., "survival curve") at any point in the life cycle. As emphasized by Hammit (2007 p. 232), "the survival curve and how it shifts are the fundamental concepts; the number of life-years saved and lives saved in a specified time period are the alternative and partial summary measures of the shift." The life-cycle consumption framework is tailor-made to account for shifts in the survival curve, and it can easily account for the age and lifetime income profile of the individual and the latency and cessation lag characteristics of the policy.

As in any structural benefit transfer application, it may be necessary to calibrate the parameters of a life-cycle consumption model using only a few aggregate data—for example, summary statistics on labor-leisure tradeoffs, average rates of saving over a representative individual's life span, average market wage differentials for more versus less risky jobs, summary results from stated preference surveys on risk tradeoffs, etc. Thus, like other structural-benefit transfer functions, one based on the life-cycle consumption framework would necessarily sacrifice statistical sophistication for theoretical consistency, so many of the advantages and disadvantages of structural benefit-transfer functions discussed by Smith et al. (2002, 2006) will apply to life-cycle models as well.

An important potential advantage of using a life-cycle consumption framework for structural benefit transfers is that it could help to avoid the transfer errors that may arise from using a single VSL point estimate for all varieties of mortality risk reductions. As shown in Appendix A, the life-cycle framework allows calculation of the marginal willingness to pay at any age *a* for risk reductions at any later age *b*, $wtp_{a,b}$. VSL estimates from hedonic wage studies may be most plausibly interpreted as the marginal willingness to pay for contemporaneous mortality risk reductions for adults of prime working age, e.g., $wtp_{40,40}$. It may be inaccurate to use such estimates to calculate the willingness to pay for, say, a 20 year-old who will experience mortality risk reductions at ages 55 through the end of life. In contrast, a schedule of $wtp_{a,b}$ estimates based on a calibrated life-cycle consumption model would give a ready means of calculating total willingness to pay for any exogenous shift in the survival curve for individuals of any age. Furthermore, this approach can properly account for all latency and cessation lag effects associated with the specific pattern of mortality risk changes caused by the policy, without the need for possibly inaccurate transfers of a VSL point estimate to earlier and later ages and across individuals with different levels of wealth and income.

Implementing such a structural life-cycle benefit transfer function would be challenging. Estimating or calibrating such a model would require specifying or solving for the life-cycle pattern of consumption, and specifying a functional form for the utility function as well as calibrating or estimating its parameters. Any structural benefit transfer approach—whether based on a life-cycle consumption framework or something else—would represent a significant departure from the traditional point estimate transfer approach typically used for mortality risk valuations, mainly based on the VSL. To accelerate the development of such an approach, we recommend conducting additional case studies applying existing structural benefit transfer functions (e.g., Smith et al. 2002, 2003, 2006) to a wider range of illustrative policy scenarios, and additional research aimed at expanding and refining the calibration of existing benefit transfer functions or developing new ones for potential use in future policy analyses. The scholarly research on structural benefit transfer methods is still in an early stage, so we are especially interested in EEAC recommendations in this area.

7 Conclusions

EPA continually strives to improve the quality of its economic analyses of proposed environmental policies. This is especially important in the area of human health valuation, in particular the value of mortality risk reductions, since such a large fraction of the (monetized) benefits of EPA rules are based on this category of impacts. This white paper represents the latest round of literature review and study by EPA's National Center for Environmental Economics on this topic, submitted to the SAB-EEAC for feedback. Advice from the committee will be carefully considered as EPA updates its *Guidelines for Preparing Economic Analyses*.

7.1 Addressing key issues: terminology, altruism, cancer valuation

EPA plans to change its metric and terminology for mortality risk valuation in benefit-cost analysis to better reflect the risk-dollar tradeoffs faced by individuals as evaluated in the economics literature, and risk reductions provided by environmental policies. As detailed in section 3.1.2 of this white paper, for valuation purposes we will report changes in risk reductions valued in terms of the value of mortality risk (VMR), scaled to micro-risk reductions. This is consistent with recent suggestions in the economics literature and is aimed at reducing confusion about how mortality risks are evaluated in benefit-cost analysis.

A second key issue for EPA is the valuation of cancer risk reductions and how these risks are valued systematically differently from the more immediate risks typically considered in WTP studies. Our review of the cancer literature, while not conclusive, suggests a “cancer differential” of roughly 50% over immediate accidental or “generic” risk valuation estimates. We recommend including a differential of this general magnitude as part of Agency benefits analyses for reduced cancer risks. Specific guidance on the application of this differential will be developed by the Agency at a later date.

7.2 Longer term analytical directions

In the longer term, EPA plans to perform analysis to better and more rigorously synthesize the existing mortality risk valuation literature. Two key directions include meta-analysis and structural benefit-transfer.

7.2.1 Meta-analysis

Section 5.1 described simplified approaches to aggregating the existing empirical valuation data, along with some key issues to consider in this process. These include whether to (i) use multiple estimates from studies, (ii) update all studies to a common year accounting for real income growth, and (iii) limit SP studies to avoid double-counting the effects of cancer risks and latency or cessation lag. The suggested approach evaluates the RP and SP studies separately, from which EPA would develop a range of default values for the value of mortality risk (VMR).

Alternatively, a new addition to the discussion of mortality risk meta-analysis with the SAB-EEAC is the potential for Bayesian meta-regression, and we are particularly interested in the SAB-EEAC comments on the potential advantages and disadvantages of this approach. Another key question to consider is how the results of any meta-regression would be used to inform guidance, and the merits of developing a statistical benefit-transfer function from these results.

7.2.2 Structural Benefit Transfer

An alternative to meta-regression and other largely statistical approaches to synthesizing literature results for policy, is to impose more structure on the benefit-transfer problem and then calibrate a preference function based on a specified utility function and data on observable outcomes. This is a relatively new approach that has been developed and demonstrated in only a few previous studies. We recommended conducting additional scoping studies and further research to develop structural benefit transfer functions, possibly based on a life-cycle consumption framework, suitable for application in benefit-cost analyses of future EPA policies.

7.3 Other research directions

We see three other areas where more research would be valuable in developing guidance for mortality risk valuation, and we welcome SAB-EEAC comment on these (as requested in the accompanying charge questions).

First, additional applied research on the altruistic components of WTP for public risk reductions would be a valuable contribution, potentially allowing EPA to rigorously include theoretically-appropriate altruistic values and better reflect the public value of environmental policies. We acknowledge that this is a difficult task. The economics literature on the proper treatment of altruism in benefit-cost analysis is well-developed, enumerating the conditions under which it is appropriate to include altruistic values in evaluating the benefits of public programs. EPA programs are inherently public and ideally should include paternalistic altruism. However, while the empirical literature has been able to capture some altruistic values for public risk reductions, it has not generally been able to distinguish among types of altruism sufficiently well for the values to be included neatly in applied analysis.

Second, more and more research reflects the general understanding that value of reducing mortality risks is not “one-size-fits-all.” Rather, these values are heterogeneous, or “individuated,” and depend upon a wide array of individual and risk characteristics. More detailed research in this area also will provide data needed for developing more general and more accurate benefit-transfer functions.

Third and finally, most of the valuation literature, and many theoretical frameworks, have treated mortality and morbidity risks separately, focusing on just one of these endpoints at a time. However, some recent work also suggests that changes in health risks may be best framed as changes in health risk profiles that include both mortality and morbidity. Individuals may value different combinations of changes in risk or illness and risk of death in complex ways. Systematic empirical work to evaluate these relationships could lead to much more robust and complete benefits analysis.

References

Adamowicz W, Dupont D, Krupnick A, Zhang J. 2008. Valuation of cancer and microbial disease risk reductions in municipal drinking water: an analysis of risk context using multiple valuation methods. *Working Paper*.

Alberini A, Scasny M. 2010a. Labels and perceptions in mortality risk reduction valuations. *Working Paper.*

Alberini A, Scasny M. 2010b. Context and the VSL: Evidence from a Stated Preference Study in Italy and the Czech Republic. *Working Paper.*

Alberini A, Scasny, M, Guignet D, Tonin S. 2010. The benefits of contaminated site cleanup revisted: the case of Naples and Caserta, Italy. *Working Paper.*

Alberini A, Chiabai A. 2007. Urban environmental health and sensitive populations: how much are the Italians willing to pay to reduce their risks? *Regional Science and Urban Economics* 37: 239-258.

Alberini A, Cropper M, Krupnick A, Simon NB. 2004. Does the value of a statistical life vary with age and health status? Evidence from the US and Canada. *Journal of Environmental Economics and Management* 48(1): 769-792.

Alberini A, Cropper M, Krupnick A, Simon NB. 2006a. Willingness to pay for mortality risk reductions: does latency matter? *Journal of Risk and Uncertainty* 32: 231-245**.**

Alberini A, Hunt A, Markandya A. 2006b. Willingness to pay to reduce mortality risks: evidence from a three-country contingent valuation study. *Environmental and Resource Economics* 33:251-264.

Alberini A, Scasny M, Kohlova MB, Melichar J. 2006c. The value of a statistical life in the Czech Republic: evidence from a contingent valuation study. *In* Menne B, Ebi K, eds. *Climate Change Adaptation Strategies for Europe*. Darmstadt, Germany: Springer.

Alberini A, Tonin S, Turvani M, Chiabai A. 2007. Paying for permanence: public preferences for contaminated site cleanup. *Journal of Risk and Uncertainty* 34(2):155-178.

Aldy JE, Smyth SJ. 2006. A numerical analysis of the value of life. Working paper, http://cbey.research.yale.edu/uploads/File/aldy_35.pdf

Aldy JE, Viscusi WK. 2007. Age differences in the value of statistical life: revealed preference evidence. Resources for the Future Discussion Paper, http://www.rff.org/documents/RFF-DP-07-05.pdf

Aldy JE, Viscusi WK. 2008. Adjusting the value of a statistical life for age and cohort effects. *The Review of Economics and Statistics* 90(3):573-581.

Andersson H, Lindberg G. 2009. Benevolence and the value of road Safety. *Accident Analysis and Prevention* 41:286-293.

Andersson H. 2005. The value of safety as revealed in the Swedish car market: an application of the hedonic pricing approach. *Journal of Risk and Uncertainty* 30(3):211-239.

Andersson H. 2008. Willingness to pay for car safety: evidence from Sweden. *Environmental and Resource Economics* 42(4):579-594.

Arabsheibani RG, Marin A. 2000. Stability of estimates of the compensation for danger. *Journal of Risk and Uncertainty* 20(3):247-269.

Arrow KJ. 2007. Global climate change: a challenge to policy. *The Economist's Voice* June 2007 p 1-5.

Ashenfelter O, Greenstone M. 2004. Estimating the value of a statistical life: the importance of omitted variables and publication bias. *The American Economic Review* 94(2):454–460.

Ashenfelter O. 2006. Measuring the value of a statistical life: problems and prospects. *The Economic Journal* 116(510):C10-C23.

Barro RJ. 2006. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics* 121(3):823-866.

Bellavance F, Dionne G, Lebeau M. 2009. The value of statistical life: a meta-analysis with a mixed effects regression model. *Journal of Health Economics* 28:444-464.

Berger MC, Gabriel PE. 1991. Risk aversion and the earnings of U.S. immigrants and natives. *Applied Economics* 23: 311-318.

Bergstrom TC. 1982. When is a man's life worth more than his human capital? In: Jones-Lee MW, ed. *The Value of Life and Safety: Proceedings from a Conference held by the Geneva Association*. Amsterdam: North Holland Press, p 3-25.

Bergstrom TC. 2006. Benefit-cost in a benevolent society. *American Economic Review* 96(1):339-351.

Bergstrom JC, Taylor LO. 2006. Using meta-analysis for benefits transfer: theory and practice. *Ecological Economics* 60:351-360.

Black D, Galdo J, Liu L. 2003. *How Robust are Hedonic Wage Estimates of the Price of Risk?* Final Report to the USEPA R 829-43-001.

Black DA, Kniesner TJ. 2003. On the measurement of job risk in hedonic wage models. *Journal of Risk and Uncertainty* 27(3)205-220.

Blomquist GC. 2004. Self-protection and averting behavior, values of statistical lives, and benefit-cost analysis of environmental policy. *Review of Economics and the Household* 2(1):89-110.

Bornstein, Seth. "American Life Worth Less Today." Associated Press 10 July 2008.

Bosworth R, Cameron TA, DeShazo JR. 2009. Demand for environmental policies to improve health: evaluating community-level policy scenarios. *Journal of Environmental Economics and Management* 57(3):293-308.

Braathen NA, Lindhjem H, Navrud S. 2009. Valuing lives saved from environmental, transport and health policies: a meta-analysis of stated preference studies. http://www.oecd.org/dataoecd/20/48/43809818.pdf

Buzby JC, Ready RC, Skees JR. 1995. Contingent valuation in food policy analysis: a case study of a pesticide residue risk reduction program. *Journal of Agriculture and Applied Economics* 27(2):613-625.

Cameron TA. 2009. The value of a statistical life: [they] do not think it means what [we] think it means. University of Oregon Working Paper, http://www.uoregon.edu/~cameron/vita/ REEP_VSL_102509.pdf

Cameron TA, DeShazo JR 2008. Demand for health risk reductions. Working Paper, http://www.uoregon.edu/~cameron/vita/Manuscript_20051269R2.pdf

Cameron TA, DeShazo JR, Johnson EH. 2008. Willingness to pay for health risk reductions: differences by type of illness. Working paper, http://www.aere.org/meetings/documents/ 0608Erica_Johnson.pdf

Carson RT, Mitchell RC. 2006. Public preferences toward environmental risks: the case of trihalomethanes. *In* Alberini A, Bjornstad D, Kahn JR, eds., *Handbook of Contingent Valuation.* Northampton, MA: Edward Elgar.

Chay KY, Greenstone M. 2005. Does air quality matter? Evidence from the housing market. *Journal of Political Economy* 113(2):376-424.

Chetty R. 2006. A new method of estimating risk aversion. *The American Economic Review* 96(5):1821-1834.

Chilton S, Jones-Lee M, Kiraly F, Metcalf H, Pang W. 2006. Dread risks. *Journal of Risk and Uncertainty* 33:165-182.

Corso PS, Hammitt JK, Graham JD. 2001. Valuing mortality-risk reduction: using visual aids to improve the validity of contingent valuation. *Journal of Risk and Uncertainty* 23(2):165-184.

Cropper ML, Sussman FG. 1990. Valuing future risks to life. *Journal of Environmental Economics and Management* 19:160-174.

Cropper M, Aydede SK, Portney PR. 1994. Preferences for lifesaving programs – how the public discounts time and age. *Journal of Risk and Uncertainty* 8(3):243-265.

Cousineau J-M, Lacroix R, Girard A-M. 1992. Occupational hazard and wage compensating differentials. *Review of Economics and Statistics* 74(1):166-169.

Dasgupta P. 2008. Discounting climate change. *Journal of Risk and Uncertainty* 37:141-169.

Davis LW. 2004. The effect of health risk on housing values: evidence from a cancer cluster. *The American Economic Review* 94(5):1693-1704.

Dekker T, Brouwer R, Hofkes M, Moeltner K. 2008. The effect of risk context on the value of a statistical life: a Bayesian meta-model. IVM Working Paper.

Desaigues B, Rabl A. 1995. Reference values for human life: an econometric analysis of a contingent valuation in France. In Schwabe Christie NG, Soguel NC, *eds.*, *Contingent Valuation, Transport Safety and the Value of Life*. Boston, MA: Kluwer Academic Publishers.

Dickie M, Gerking S. 2009. Family behavior: implications for health benefits transfer from adults to children. *Environmental and Resource Economics* 43:31-43.

Dillingham AE. 1985. The influence of risk variable definition on value of life estimates. *Economic Inquiry* 24:277-294.

Dillingham AE, Smith RS. 1984. Union effects on the valuation of fatal risk. In *Proceedings of the Industrial Relations Research Association 36th annual meeting*, San Francisco, CA.

Dorsey S, Norman W. 1983. Workers' compensation, job hazards and wages. *Industrial and Labor Relations Review* 36(4): 642-654.

Ehrlich I. 2000. Uncertain lifetime, life protection, and the value of life saving. *Journal of Health Economics* 19:341-367.

European Commission. 2000. *Recommended interim values for the value of preventing a fatality in DG environment cost benefit analysis.* http://ec.europa.eu/environment/enveco/others/pdf/ recommended_interim_values.pdf

Evans MF, Smith VK. 2006. Do we really understand the age-VSL relationship? *Resource and Energy Economics* 28(3):242-261.

Evans MF, Smith VK. 2008. Complementarity and the measurement of individual risk tradeoffs: accounting for quantity and quality of life effects. NBER working paper 13722. http://www.nber.org/papers/w13722

Federal Register. 1993. Executive Order 12866: Regulatory Planning and Review. *Federal Register* 58(190):51735-51744.

Feldstein M, Ranguelova E. 2001. Individual risk in an investment-based social security system. *The American Economic Review* 91(4):1116-1125.

Flores NE. 2002. Non-paternalistic altruism and welfare economics. *Journal of Public Economics* 83(2):293-305.

Garen J. 1988. Compensating wage differentials and the endogeneity of job riskiness. *The Review of Economics and Statistics* 70(1):9-16.

Gayer T, Hamilton JT, Viscusi WK. 2000. Private values of risk tradeoffs at Superfund sites: housing market evidence on learning about risk. *Review of Economics and Statistics* 82:439-451.

Gayer T, Hamilton JT, Viscusi WK. 2002. The market value of reducing cancer risk: hedonic housing prices with changing information. *Southern Economic Journal* 69:266-289.

Gegax D, Gerking S, Schulze W. 1991. Perceived risk and the marginal value of safety. *Review of Economics and Statistics* 73(4):589-596.

Gerking S, De Haan M, Schulze W. 1988. The marginal value of job safety: a contingent valuation study. *Journal of Risk and Uncertainty* 1(2):185-199.

Greene WH. 1997. *Econometric Analysis*, Third Edition. Upper Saddle River, NJ: Prentice-Hall.

Greenstone M, Gallagher J. 2008. Does hazardous waste matter? Evidence from the housing market and the Superfund program. *Quarterly Journal of Economics* 123(3):951-1003.

Gunderson M, Hyatt D. 2001. Workplace risks and wages: Canadian evidence from alternative models. *The Canadian Journal of Economics* 34(2):377-395.

Gyrd-Hansen D, Halvorsen PA, Kristiansen IS. 2007. Willingness-to-pay for a statistical life in the times of a pandemic. *Health Economics* 17(1):55-66.

Hakes JK, Viscusi KW. 2007. Automobile seatbelt usage and the value of statistical life. *Southern Economics Journal* 73(3):659-676.

Hall RE, Jones CI. 2007. The value of life and the rise in health spending. *The Quarterly Journal of Economics* 122(1):39-72.

Hammitt JK. 2007. Valuing changes in mortality risk: lives saved versus life years saved. *Review of Environmental Economics and Policy* 1(2):228-240.

Hammitt JK, Graham JD. 1999. Willingness to pay for health protection: inadequate sensitivity to probability. *Journal of Risk and Uncertainty* 8:33-62.

Hammitt JK, Haninger K. 2007. Willingness to pay for food safety: sensitivity to duration and severity of illness. *American Journal of Agricultural Economics* 89(5):1170-1175.

Hammitt JK, Haninger K. 2010. Valuing fatal risks to children and adults: effects of disease, latency, and risk aversion. *Journal of Risk and Uncertainty* 40:57-83.

Hammitt JK and Liu J-T. 2004. Effects of disease type and latency on the value of mortality risk. *Journal of Risk and Uncertainty* 28:73-95.

HM Treasury. 2003. *The green book: Appraisal and evaluation in Central Government*. http://www.hm-treasury.gov.uk/d/green_book_complete.pdf

Ho C-S, Hite D. 2008. The benefit of environmental improvement in the Southeastern United States: evidence from a simultaneous model of cancer mortality, toxic chemical releases and house values. *Papers in Regional Science* 87(4):589-604.

Howard RA. 1989. Microrisks for medical decision analysis. *International Journal of Technology Assessment in Health Care* 5:357-370.

Hultkrantz L, Lindberg G, Andersson C. 2006. The value of improved road safety. *Journal of Risk and Uncertainty* 32:151-170.

Itaoka K, Krupnick A, Akai M, Alberini A, Cropper M, Simon N. 2007. Age, health and willingness to pay for mortality risk reductions: a contingent valuation survey of Shizuoka, Japan residents. *Environmental Economics and Policy Studies* 8:211-237.

Johannesson M, Johansson P-O, O'Connor RM. 1996. The value of private safety versus the value of public safety. *Journal of Risk and Uncertainty* 13:263-175.

Johannesson M, Johansson P-O, Lofgren K-G. 1997. On the value of changes in life expectancy: blips versus parametric changes. *Journal of Risk and Uncertainty* 15:221-239.

Johansson P-O. 2002. On the definition and age-dependency of the value of a statistical life. *The Journal of Risk and Uncertainty* 25(3):251-263.

Kaplow L. 2005. The value of a statistical life and the coefficient of relative risk aversion. *The Journal of Risk and Uncertainty* 31(1):23-34.

Kidholm K. 1995. Assessing the value of traffic safety using the contingent valuation technique: the Danish survey. *In* Schwabe NG, Soguel NC, eds. *Contingent Valuation, Transport Safety and the Value of Life: Studies in Risk and Uncertainty*. Boston, MA: Kluwer.

Kim S-W, Fishback PV. 1999. The impact of institutional change on compensating wage differentials for accident risk: South Korea, 1984-1990. *Journal of Risk and Uncertainty* 18:231-248.

Kniesner TJ, Leeth JD. 1991. Compensating wage differentials for fatal injury risk in Australia, Japan, and the United States. *Journal of Risk and Uncertainty* 4(1):75-90.

Kniesner, TJ, Viscusi WK, Ziliak JP. 2006. Life-cycle consumption and the age-adjusted value of life. *Contributions to Economic Analysis and Policy* 5(1):Article 4. http://www.bepress.com/bejeap/contributions/vol5/iss1/art4.

Kocherlakota NR. 1996. The equity premium: it's still a puzzle. *Journal of Economic Literature* 34:42-71.

Kochi I, Hubbell B, Kramer R. 2006. An empirical Bayes approach to combining and comparing estimates of the value of a statistical life for environmental policy analysis. *Environmental & Resource Economics* 34:385-406.

Koop, G. 2003. *Bayesian Econometrics*. John Wiley and Sons Ltd. West Sussex England.

Lanoie P, Pedero C, Latour R. 1995. The value of a statistical life: a comparison of two approaches. *Journal of Risk and Uncertainty* 10:235-257.

Layard R, Nickell S, Mayraz G. 2008. The marginal utility of income. *Journal of Public Economics* 92:1846-1857.

Leeth, J. D. and J. Ruser. 2003. Compensating wage differentials for fatal and nonfatal injury risk by gender and race. *Journal of Risk and Uncertainty* 27(3): 257-77.

Leigh JP. 1995. Compensating wages, value of a statistical life, and inter-industry differentials. *Journal of Environmental Economics and Management* 28(1):83-97.

Leigh JP. 1991. No evidence of compensating wages for occupational fatalities. *Industrial Relations* 30(3):382-395.

Leigh PJ, Folsom RN. 1984. Estimates of the value of accident avoidance at the job depend on concavity of the equalizing differences curve. *The Quarterly Review of Economics and Business* 24(1):55-56.

Lott, John R., Jr., and Richard L. Manning. 2000. Have changing liability rules compensated workers twice for occupational hazards? *The Journal of Legal Studies* 29(1): 99-130.

Magat WA, Viscusi WK, Huber J. 1996. A reference lottery metric for valuing health, *Management Science* 42(8):1118-11130.

Marin A, Psacharopoulos G. 1982. The reward for risk in the labor market: evidence from the United Kingdom and reconciliation with other studies. *Journal of Political Economy* 90(4):827-853.

Martinello F, Meng R. 1992. Workplace risks and the value of hazard avoidance. *Canadian Journal of Economics* 25(2):333-345.

McConnell KE. 1997. Does altruism undermine existence value? *Journal of Environmental Economics and Management* 32:22-37.

Meng R. 1989. Compensating differences in the Canadian labour market. *Canadian Journal of Economics* 12(2):413-424.

Meng R, Smith D. 1990. The valuation of risk of death in public sector decision-making. *Canadian Public Policy* 16(2):137-144.

Meng RA, Smith DA. 1999. The impact of workers' compensation on wage premiums for job hazards. *Applied Economics* 31:1101–1108.

Miller P, Mulvey C, Norris K. 1997. Compensating differentials for risk of death in Australia. *Economic Record* 73(223):363-372.

Miller T, Guria J. 1991. The value of statistical life in New Zealand. *Report to the New Zealand Ministry of Transport*, Land Transportation Division, #0-477-05255-X.

Moeltner K, Woodward R. 2009. Meta-functional benefit transfer for wetland valuation: making the most of small samples. *Environmental & Resource Economics* 42:89-108.

Moore CC, Phaneuf DJ, Thurman WN. 2009. A Bayesian bioeconometric model of invasive species control: the case of the Hemlock Wooly Adelgid. Working paper, http://www. ncsu.edu/cenrep/research/documents/Moore_Phaneuf_Thurman_invasivepaper-Aug09.pdf

Moore MJ, Viscusi KW. 1988. Doubling the estimated value of life: results using new occupational fatality data. *Journal of Policy Analysis and Management* 7(3):476-490.

Morris J, Hammitt JK. 2001. Using life expectancy to communicate benefits of health care programs in contingent valuation. *Medical Decision Making* 21:468-478.

Morzek J, Taylor LO. 2002. What determines the value of life? A meta-analysis. *Journal of Policy Analysis and Management* 21(2):253-270.

Murphy KM, Topel RH. 2006. The value of health and longevity. *Journal of Political Economy* 114(5):871-904.

Nelson JP, Kennedy PE. 2008. The use (and abuse) of meta-analysis in environmental and natural resource economics: an assessment. http://papers.ssrn.com/sol3/papers. cfm?abstract_id=1117490

Nordhaus W. 2008. *A Question of Balance: Weighing the Options on Global Warming Policies*. New Haven, CT: Yale University Press.

Olson CA. 1981. An analysis of wage differentials received by workers on dangerous jobs. *Journal of Human Resources* 16(2):167-185.

OMB (Office of Management and Budget). 2003. Circular A-4: Rgulatory Analysis. Executive Office of the President, Washington, DC. http://www.whitehouse.gov/omb/circulars

Persson U, Norinder A, Hjalte K, Gralen K. 2001. The value of a statistical life in transport: findings from a new contingent valuation study in Sweden. *Journal of Risk and Uncertainty* 23(2):121-134.

Philips PR, Russell IT, Jones-Lee MW. 1989. The empirical estimation of individual valuation of safety: results of a national sample survey. *In* Jones-Lee MW, ed. *The Economics of Safety and Physical Risk*. Oxford, UK: Basil-Blackwell.

Portney PR. 1981. Housing prices, health effects, and valuing reductions in risk of death. *Journal of Environmental Economics and Management* 8:72-78.

Pratt JW, Zeckhauser RJ. 1996. Willingness to pay and the distribution of risk and wealth. *Journal of Political Economy* 1(4):747-763.

Revesz RL. 1999. Environmental regulation, cost-benefit analysis, and the discounting of human lives. *Columbia Law Review* 99(4):941-1017.

Rice DP, Cooper BS. 1967. The economic value of human life. *American Journal of Public Health* 57(11):1954-1966.

Rosen S. 1988. The value of changes in life expectancy. *Journal of Risk and Uncertainty* 1:285-304.

Sandy R, Elliott RF. 1996. Unions and risks: their impact on the level of compensation for fatal risk. *Economica* 63:291-310.

Savage I. 1993. An empirical investigation into the effect of psychological perceptions on the willingness-to-pay to reduce risk. *Journal of Risk and Uncertainty* 6(1):75-90.

Scotton CR, Taylor LO. 2009. Valuing risk reductions: incorporating risk heterogeneity into a revealed preference framework. Working paper, http://www.ncsu.edu/cenrep/research/ documents/ riskhet_REE_2009.pdf

Shepard DS, Zeckhauser RJ. 1984. Survival versus consumption. *Management Science* 30(4):423-439.

Siebert SW, Wei X. 1994. Compensating wage differentials at workplace accidents: evidence for union and nonunion workers in the UK. *Journal of Risk and Uncertainty* 9(1):61-76.

Slovic P. 1987. Perception of risk. *Science* 30(4):423-439.

Smith RS. 1974. The feasibility of an 'injury tax' approach to occupational safety. *Law and Contemporary Problems* 38(4):730-744.

Smith et al. 2004. Do the Near Elderly Value Mortality Risks Differently? *The Review of Economics and Statistics* 86(1): 423-429.

Smith VK, Van Houtven G, Pattanayak SK. 2002. Benefit transfer via preference calibration: “prudential algebra” for policy. *Land Economics* 78(1):132-152.

Smith VK, Pattanayak SK, Van Houtven G. 2003. Using preference calibration for VSL estimation. RTI Working Paper 03_03. http://www.rti.org/pubs/rtipaper_03_03.pdf

Smith VK, Pattanayak SK, Van Houtven GL. 2006. Structural benefit transfer: an example using VSL estimates. *Ecological Economics* 60:361-371.

Smith TC, Spiegelhalter DJ, Thomas A. 1995. Bayesian approaches to random-effects meta-analysis: a comparative study. *Statistics in Medicine* 14(24): 2685-2699.

Strand. 2002. Public- and private-good values of statistical life: results from a combined choice experiment and contingent valuation survey. *Working Paper* 2002:2, University of Oslo, Health Economics Research Program.

Subramanian U, Cropper M. 2000. Public choices between life saving programs: the tradeoff between qualitative factors and lives saved. *Journal of Risk and Uncertainty* 21(1):117-149.

Sunstein CR. 2004. Valuing life: a plea for disaggregation. *Duke Law Journal* 54:385-445.

Szpiro GG. 1986. Measuring risk aversion: an alternative approach. *The Review of Economics and Statistics* 68(1):156-159.

Tsuge T, Kishimot A, Takeuchi K. 2005. A choice experiment approach to the valuation of mortality. *Journal of Risk and Uncertainty* 31(1):73-95.

Um M, Kwak S, Kim T. 2002. Estimating willingness to pay for improved drinking water quality using averting behavior method with perception measure. *Environmental and Resource Economics* 21: 287-302.

US EPA. 1999. *The Benefits and Costs of the Clean Air Act, 1990 to 2010.*

USEPA. 2000a. *Guidelines for Preparing Economic Analyses*. EPA-240-R-00-003. http://yosemite.epa.gov/ee/epa/eed.nsf/pages/Guidelines.html

USEPA. 2000b. Valuing fatal cancer risk reductions. Science Advisory Board review draft.

USEPA. 2000c. *An SAB Report on EPA's White Paper "Valuing the Benefits of Fatal Cancer Risk Reduction*." EPA-SAB-EEAC-00-013. http://www.hks.harvard.edu/fs/rstavins/ Monographs_&_Reports/SAB_Report_on_Fatal_Cancer.pdf

USEPA. 2001. *Arsenic Rule Benefits Analysis: An SAB Review*. EPA-SAB-EC-01-008. http://www.epa.gov/ogwdw000/arsenic/pdfs/ec01008.pdf

USEPA. 2003. *Children's Health Valuation Handbook*. EPA-100-R-03-003. http://yosemite.epa.gov/EE/epa/eed.nsf/pages/HandbookChildrensHealthValuation.html

USEPA. 2004a. Value of statistical life analysis and environmental policy: a white paper. http://yosemite.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0483-01.pdf/$file/EE-0483-01.pdf

USEPA. 2004b. *The Valuation of Mortality Risk Reduction: A Science Advisory Board Notification of a Consultation*. EPA-SAB-CON-04-004. http://yosemite.epa.gov/sab/SABPRODUCT.nsf/0/ 4CAC87DC95C91646852571AA00668C74/$File/sab_con_04004.pdf

USEPA. 2006. *Report of the EPA Work Group on VSL Meta-analyses*. http://yosemite.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0494-01.pdf/$file/EE-0494-01.pdf

USEPA. 2007. *SAB Advisory on EPA's Issues in Valuing Mortality Risk Reduction*. (EPA-SAB-08-001) http://nepis.epa.gov/Adobe/PDF/P10007U3.PDF

USEPA. 2008. *Guidelines for Preparing Economic Analyses*. External Review Draft. http://yosemite1.epa.gov/ee/epa/eerm.nsf/vwAN/EE-0516-01.pdf/$file/EE-0516-01.pdf

Van Houtven G, Sullivan MB, Dockins C. 2008. Cancer premiums and latency effects: a risk tradeoff approach for valuing reductions in fatal cancer risks. *Journal of Risk and Uncertainty* 36(2):179-199.

Viscusi WK. 1978. Labor market valuations of life and limb: empirical evidence and policy implications. *Public Policy* 26(3):359-386.

Viscusi, WK. 1981. Occupational Safety and Health Regulation: Its Impact and Policy Alternatives. *Research in Public Policy Analysis and Management* 2: 281-299.

Viscusi WK. 1992. *Fatal Tradeoffs: Public and Private Responsibilities or Risk*. New York, NY: Oxford University Press.

Viscusi WK. 2003. Racial differences in labor market values of statistical life. *Journal of Risk and Uncertainty* 27(3):239-256.

Viscusi WK. 2004. The value of life: estimates with risks by occupation and industry. *Economic Inquiry* 42(1):29-48.

Viscusi WK, Aldy JE. 2003. The value of a statistical life: a critical review of market estimates throughout the world. *Journal of Risk and Uncertainty* 27(1):5-76.

Viscusi WK, Aldy JE. 2007. Labor market estimates of the senior discount for the value of statistical life. *Journal of Environmental Economics and Management* 53(3):377-392.

Weiss P, Maier G, Gerking S. 1986. The economic evaluation of job safety: a methodological survey and some estimates for Austria. *Empirica* 13(1):53-67.

Weitzman ML. 2010a. What is the "damages function" for global warming—and what difference might it make? *Climate Change Economics* 1(1):57-69.

Weitzman ML. 2010b. GHG targets as insurance against catastrophic climate damage. Working paper.

Wolfe J. 2007. What is the value of preventing a fatality? *In* Lewens, Tim, ed. *Risk: Philosophical Perspectives.* New York, NY: Routledge.

Yaari ME. 1964. On the consumer's lifetime allocation process. *International Economic Review* 5(3):304-317.

Zhang J, Adamowicz W, Krupnick A, Dupont D. 2009. The role of altruism in the valuation of community drinking water risks. *Working Paper*.

Tables and figures

Table 1: Cancer Valuation Literature Summary

Study	Timing	Treatment of Morbidity	Dread	Risk context and characteristics	Affected Pop.	Other health effects	Findings / Notes
Hammitt & Haninger (2010) *Choice experiment*	Latency periods of 1, 10, 20 years Implied discount rates not stat diff from zero (-1.2 to 3.9)	½ sample: no symptom descriptions; ½ sample: 150-200 word descriptions Self-assessed severity based on EQ-5D and visual analog scales	Not separately treated	Pesticide risks from food. Safer food from Pesticide Safety System (not organic) Auto accident for whole family simultaneously from "product" on next car purchase.	Adult selves Adult others Children	Organ: brain, liver, bladder, lymphocytes Mortality only (no non-fatal outcomes)	No statistical difference by cancer/non-cancer; target organ; auto/other risks (for protecting whole family simultaneously). Child VSL=1.8 Adult Other VSL=1.15 Self Insensitivity to number of people for the "whole household" question.
Alberini & Scasny (2010a) "Labels & Perceptions . . ." *Choice experiment*	Latency periods of 0, 2, 5, 10 years Discount rate = zero (may reflect changes in future baseline risk)	Description of morbidity or illness Focus of study is mortality risk	Rated subjectively by respondents for each type of risk	"cancer" designation varied independently from dread Private good v. nationwide public program Other independent variables: salience("familiarity") exposure, sensitivity to illness, beliefs in prevalence	Adults Children (Italy)	Respiratory fatality Cancer fatality Auto fatality	EC Cancer differential of 50% relative to "general" VSL is consistent with findings. Auto accident risks valued less than respiratory or cancer VSL higher for public v. private (if public programs are effective) "cancer" designation effect persists after controlling for other risk characteristics.

							VSL increases with dread;
Alberini & Scasny (2010b) "Context and the VSL . . ." *Choice experiment* * The Italy sample appears to overlap with Alberini et al. (2010a)	Yes (0, 2, 5, 10 years) Implicit discount rates from 0.3 to 7.4%	Description of morbidity or illness Focus of study is mortality risk	No	Public & Private programs Perceived "effectiveness" of the program Finds that "risk characteristics and mode of delivery primarily drive heterogeneity in VSL"	Adults Children (CzechRep. & Italy)	Respiratory illness Cancer Road-Traffic accidents	"Evidence of cancer premium": ~ 1.25x (Italy-children) ~ 1.90x (Italy-adults) ~ 1.75x (Czech-children) ~ 2.5 x (Czech-adults) Premium for public programs Any premium for reduced children's risk is modest (small for cancer risks, larger for other causes).
Adamowicz et al. (2008) *CVM and CE*	Risks described as community deaths over a 35-year time period for microbials and carcinogens	Symptoms described for microbial illness and for bladder cancer	Not addressed	Risk reductions are strictly public Describes tradeoffs between reduced microbials in DW and reduced carcinogens	Households (Canada)	Microbial illness Microbial fatality Bladder cancer - fatal - nonfatal	Modest cancer "discount" (for mortality) Cancer VSL = .85*Microbial VSL Cancer illness = 20-50% of cancer mortality
Cameron & DeShazo (2008) *Choice experiment*	Illness profile over specific, varying times Results support lower values for longer latency. No implicit	Health states defined as: - Current health - Sickness - Remission years - Lost life-years Illness characterized	Not addressed	Intervention is generally a screening and treatment program to prevent the given risk profile. For auto accidents it is a safety program.	Adults	12 major common risks, including: Heart disease, heart attack, stroke, respiratory disease, diabetes, Alzheimer's	Difficult to draw general conclusions. Heart attacks & heart disease risks valued similarly to some cancers (and more than others).

	discount rate estimated.	by length and severity (pain, disability)				Cancer (5 types) Auto accidents	
Tsuge et al. (2005) *Choice experiment*	Latency periods of 0, 5, 10 years Implied discount rate = 20%	Unclear, but does not appear to be detailed. Focus is on mortality.	Subjective perceptions of voluntariness, controllability, dread(pain), dread(fear), severity, exposure	--	Adults (Japan)	Accidents Generalized cancer Heart disease Non-specific	Unique formulation of "quantity-based" VSL distinguishing WTP for opportunities for risk reduction Depends on model specifications, but perhaps 20% differential over "general" risks; reduced cancer risks preferred to reduced heart attack risk;
Hammitt & Liu (2004) *CVM*	Latent: 20 years to onset of symptoms Acute: symptoms "within a few months" Implied discount rates of 1.5% (with up to 3% plausible)	brief description of symptoms progressive severity over time from mild to bedridden and unable to care for themselves lasting 2-3 years before mortality	Not addressed directly	All symptoms held the same except for "cancer" designation: Lung cancer v. bronchitis (from pollution from factories) Liver cancer v. liver failure (from drinking water contaminants)	Adults (Taiwan)	Liver (failure v. cancer) Lung (bronchitis v. cancer)	~30% differential for cancer relative to identical non-cancer degenerative disease (marginal significance) Environmental context. No "trauma" or "accident" alternative for comparison.
Philips et al. (1989)	No	No	No	No	Adults (U.K.)	Motor vehicles Heart disease Fatal & Nonfatal	Mean estimates higher for cancer; median estimates are not
The following two studies are risk-risk studies							
Van Houtven et al. (2008)	Latency periods of 5, 15, 25-year periods specified	Symptoms described for three types of cancer;	Not treated separately	Organ-specific cancer risks vs. auto-accident risks.	Adults	Fatal cancer (stomach, liver, brain);	Significant cancer differential (3x over auto accidents at 5-year

Risk-Risk Survey	Morbidity varied from 2 or 5 years	morbidity duration varied separately from latency				Fatal auto accident	latency; 1.5x at 25 years) Differential declines with length of latency; Latency would need to be 30+ years for indifference
Magat et al. (1996) *risk-risk survey*	Not addressed explicitly	Symptoms described for lymphoma and nerve disease	No	Included separate treatments for non-fatal lymphoma and nerve disease Respondents told not to consider out of pocket medical costs	Adults	Fatal lymphoma, non-fatal lymphoma; fatal auto accidents nerve disease	No evidence of differential for cancer fatality (ratio of fatal cancer: fatal auto is 1:1.) Ratio of non-fatal cancer to auto is ~.58.
The following studies examine cancer only (without comparison to other risks)							
Carson & Mitchell (2006) *Open-ended CVM*	Not in survey; VSL estimates assume 25 years	No	No	Public/social decision Cancer risks from THM in drinking water	Adults (Household?)	No	Cancer VSL depends upon assumptions about latency and discount rate. Also sensitive to risk reduction. Assuming 0.4/100,000 reduction, 25-yr latency, results range from $3.4m at 3% to $8.8m at 7%
Alberini et al. (2010)	0, 2, 5, 10 years Employed a zero discount rate for estimation based on prior work	Extent unclear	Unclear	Cancer risks from hazardous waste in Italy	Adults	Fatal Cancer (type unspecified)	New estimates of the cancer VSL using data from 2008 survey in Milan Cancer VSL of ~$5.6m

							About 20% higher if delivered via public program (if public programs are considered "effective").
Buzby et al. (1995)	No	No	No	Exposure to pesticides in grapefruit	Grapefruit customers (Adults)	No	Makes assumptions about lifetime exposure to estimate VSL=$6.99m

Revealed Preference Cancer Valuation Studies

Study	Timing	Treatment of Morbidity	Dread	Risk context and characteristics	Affected Pop.	Other health effects	Findings / Notes
Gayer et al. (2000) *Hedonic Property*	No	No	No	No; just "cancer" w/o distinction between fatal and non-fatal cancers	Adults / Household near Superfund sites	No	cancer risk reductions valued similarly to workplace fatal risks
Gayer et al. (2002) *Hedonic Property*	No	No	No	No; just "cancer" w/o distinction between fatal and non-fatal cancers	Adults / Household near Superfund sites	No	$5.2m to $10.0m cancer VSL with no latency (and 100% fatality.) With 10-year latency: - $6.2 to $11.7 at 3% - $10.2 to $19.8 at 7%
Davis (2004) *Hedonic Property*	Unclear	No	No	Pediatric leukemia from cancer cluster; no distinction between fatal and non-fatal	Children	No	value of prevented pediatric leukemia ranges from $4.1m to $11.5m depending on model used
Ho and Hite (2008)	No	No	No	cancer mortality only (didn't include non-	Adults	No	Hedonic property with $6.0m Value of statistical

Hedonic Property				fatal)			cancer fatality (without latency treatment or assumptions).
Lott & Manning (2000) *Hedonic wage*	No	No	No	Cancer	Workers	No	Hedonic Wage Cancer VSL = $12.4 million

Table 2. Select variables included in the stated preference meta-analysis dataset

Variable Name	Description
STUDY	Study identifier
PUBYEAR	Year study was published or released
PUBLISH	0=unpublished or working paper; 1=published in a peer-review outlet (includes book chapters)
JRU	0=does not appear in *Journal of Risk and Uncertainty (JRU);* 1=published in *JRU*
ALBSERIES	0=not part of the Alberini, Krupnick, Cropper and Simon series of studies; 1=part of this series
AUTO	0=non-auto/traffic risk; 1=auto/traffic risk
ENVIRONMENTAL	0=non-environmental risk source; 1=environmental-related risk (i.e., air pollution, drinking water, hazardous waste site, or unspecified general death risk)
PUBLIC	0=risk affects individual only; 1=risk affects public
CANCER	0=non-cancer death; 1=cancer death
ESTIMATES	Number of estimates reported or calculated from study
WTP	Willingness to pay for risk reduction (2009 US dollars)
WTP_SE	Standard error for WTP
VSL	VSL in millions, adjusted for inflation and income growth (2009 dollars)
SE	Standard error in millions of VSL estimate
MEAN	0=WTP/VSL is based on median WTP; 1=WTP/VSL is based on mean WTP
YEARCONDUCT	Year study was conducted
US	0=non-US study; 1=US study
CV	0=choice experiment; 1=contingent valuation
BASE	Baseline risk presented to survey respondents
REDUCE	Size of risk reduction presented to respondents
PCTREDUCE	Percent reduction in risk
TIMING	0=immediate risk reduction, 1=latent risk reduction
LENGTH	Length of latency period in years (0=immediate risk reduction)
SIZE	Sample size used to calculate WTP/VS: estimate
MALE	0=female, 1=male
AGE	Average age
RACE	Percent white
INCOME	Annual mean household income (thousands, 2007 US dollars)
HEALTH	Percent reporting exceptional or very good health, no reported disease or illness, or non-smoker
NSCENARIO	Number of scenarios each respondent was asked to value
MODE	0= self administered survey mode, 1=survey administered with an interviewer (e.g., in-person, telephone
DOTS	0=ladder, bar chart used for visual aid; 1=grid used for visual aid
SCOPE	0=no scope test performed or calculated, 1=scope test performed or calculated
WEAK	0=does not pass a weak scope test, 1=passes a weak test, but WTP is less than proportional to the size of the risk reduction

Variable Name	Description
STRONG	0=does not pass a strong scope test, 1=passes a strong test; WTP is proportional to the size of the risk reduction

Table 3. Stated preference dataset

Study	Country	Sample Size	Risk Characteristics: Cancer	Public	Latency (yrs)	Auto risk	Env. risk	Unspec. Source	Other risk type	Risk reduction	WTP (2009$)*	SE
Adamowicz et al. (2008)	USA	366	0	1	0	0	1	0	0	0.0000029	6.65 (1)	0.91
Adamowicz et al. (2008)	USA	366	1	1	10	0	1	0	0	0.0000029	6.03 (1)	0.75
Alberini and Chiabai (2007)	Italy	756	0	0	0	0	1	0	0	0.0001	6.03 (2)	.
Alberini et al. (2007)	Italy	782	0	1	0	0	1	0	0	0.000001	6.96 (1)	.
Alberini et al. (2004)	USA	548	0	0	0	0	0	1	0	0.0001	6.59 (3)	1.00
Alberini et al. (2004)	Canada	292	0	0	0	0	0	1	0	0.0001	5.05 (3)	0.66
Alberini et al. (2006a)	USA	403	0	0	10	0	0	1	0	0.0005	0.95 (4)	0.44
Alberini et al. (2006a)	Canada	589	0	0	10	0	0	1	0	0.0005	1.42 (4)	0.26
Alberini et al. (2006b)	France, Italy, UK	.	0	0	0	0	0	1	0	0.0005	3.22 (4)	0.57
Alberini and Scasny (2010)	Italy	1906	1	1	4.25	1	1	0	0	0.000425	4.68 (16)	0.30
Alberini and Scasny (2010)	Czech Republic	1506	1	1	4.25	1	1	0	0	0.000425	1.27 (16)	0.14
Alberini et al. (2006c)	Czech Republic	954	0	0	0	0	1	0	0	0.0003	3.11 (4)	0.21
Andersson and Lindberg (2009)	Sweden	216	0	0	0	1	0	0	0	0.0002	13.02 (5)	.
Andersson and Lindberg (2009)	Sweden	222	0	1	0	1	0	0	0	0.0002	7.45 (5)	.
Buzby et al. (1995)	USA	512	1	0	75	0	0	0	1	0.00000066	6.99 (6)	.
Cameron et al. (2008)	USA	1619	1	0	10	0	0	1	0	0.000001	0.86 (7)	.
Carson and Mitchell (2006)	USA	121	1	1	25	0	1	0	0	0.0000004	8.64 (8)	.
Corso et al. (2001)	USA	275	0	0	0	1	0	0	0	0.00005	4.29 (9)	.
Desaigues and Rabl (1995)	France	1000	0	1	0	1	0	0	0	0.000046	1.64 (1)	.
Gerking et al. (1988)	USA	861	0	0	0	0	0	0	1	0.00025	6.86 (6)	.
Gyrd-Hansen et al. (2007)	Norway	1168	0	0	0	0	0	0	1	0.0028	0.04 (1)	.
Hakes and Viscusi (2007)	USA	465	0	0	0	1	0	0	0	0.0001	7.22 (10)	.
Hammitt and Graham (1999)	USA	992	0	0	0	1	0	0	0	0.00005	2.96 (11)	0.32
Hammitt and Graham (1999)	USA	978	0	0	0	0	0	0	1	0.000073	2.72 (11)	0.56
Hammitt and Haninger (2010)	USA	1997	0.5	0	1	0	0	0	1	0.00015	6.77 (12)	1.24
Hammitt and Liu (2004)	Taiwan	1248	1	0	20	0	1	0	0	0.00005	1.94 (13)	.
Hultkrantz et al. (2006)	Sweden	225	0	0	0	1	0	0	0	0.000165	6.40 (14)	.
Itaoka et al. (2007)	Japan	248	0	0	0	0	0	1	0	0.001	2.92 (17)	0.76
Johannesson et al. (1997)	Sweden	2029	0	0	22.5	0			0	0.0002	5.13 (10)	.
Johannesson et al. (1996)	Sweden	389	0	0	0	1	0	0	0	0.000162	4.49 (18)	0.48

Study	Country	Sample Size	Risk Characteristics							Risk reduction	WTP (2009$)*	SE
			Cancer	Public	Latency (yrs)	Auto risk	Env. risk	Unspec. Source	Other risk type			
Johannesson et al. (1996)	Sweden	410	0	1	0	1	0	0	0	0.000162	3.73 (18)	0.48
Kidholm (1995)	Denmark	908	0	0	0	1	0	0	0	0.000022	2.38 (19)	.
Lanoie et al. (1995)	Canada	162	0	0	0	1	0	0	0	0.0002	2.92 (10)	.
Miller and Guria ((1991)	New Zealand	629	0	0	0	1	0	0	0	.	1.59 (21)	.
Morris and Hammitt (2001)	USA	167	0	0	20	0	0	0	1	0.046	0.19 (20)	.
Persson et al. (2001)	Sweden	675	0	0	0	1	0	0	0	0.00003	3.59 (1)	.
Philips et al. (1989)	U.K.	1563	1	0	0	1	1	0	0	.	6.90 (1)	.
Strand (2002)	Norway	.	0	0	1	0	1	0	0	.	0.57 (22)	
Tsuge et al. (2005)	Japan	400	1	0	5	0	1	1		0.0001	3.62 (15)	.
Zhang, et al. (2009)	Canada	366	1	0	.	0	0	0	0	.	12.69 (23)	.

* The WTP and SE estimates reported in this table are adjusted for inflation (using the CPI) and income growth (using an elasticity of 0.5).

(1) author's preferred

(2) healthy 30-49 year old, based on mean and smaller risk reduction (from Table 7 in paper)

(3) based on mean and smaller risk reduction (Table 6 in paper)

(4) based on mean

(5) based on parametric estimation (Table 7 in paper)

(6) only estimate reported in paper

(7) 45 year old who is diagnosed with lung cancer 10 years after exposure, is sick for 5 years and then dies; estimate is chosen because it most closely matches many EPA policy scenarios (Table 3 in paper)

(8) based on corrected mean for the smallest risk reduction (Table 19.2 in paper) (note: We could also obtain other independent estimates for different risk reductions)

(9) from a model with co-variates for the smaller risk reduction using dots for a visual aid (Table 3 in paper)

(10) based on the full sample

(11) based on median (mean not reported) for the smallest risk reduction (Table 5 and 7 in paper)

(12) based on model of WTP for reductions in risk to self, which is based on median WTP, one year latency and cancer set to 0.5 and affected organs set to 0.25 (options are brain, bladder, liver and lymphocytes) (Table 2 in paper)

(13) based on latent lung cancer from model with full set of co-variates (Table 3 in paper)

(14) based on private risk reduction (there is also an estimate for a public risk reduction, but they are not independent)

(15) only estimate in paper; reflects the idea that wtp is independent of the source of risk; CE asks about cancer, accidents, heart disease, over different latency periods
(16) based on pooled model (Table 5 in paper)
(17) based on smaller risk reduction with no latency from wave 2 (where smaller risk reduction was presented first (Table 7 in paper)
(18) based on standard estimates (Table 2 in paper)
(19) based on mean estimate for risk reduction provided through an air bag (assumed to be a private risk reduction) using the maximum WTP results (Table 2 in paper)
(20) based on WTP for vaccine at age 60 (Table 3 in paper)
(21) based on WTP for a safer car (Table 3 in paper)
(22) based on WTP for private reductions in risk from environmental causes (Table 10 in paper)
(23) based on WTP for private cancer risk reductions assuming no treatment or purchase of bottled water (Table 9 in paper)

Table 4: Hedonic Wage dataset

Study	Country	Sample	Sample Size	Sample Characteristics: Union	White	Male	Manual/Mfg	Blue Collar	Risk Variable	Mean Risk	Nonfatal Risk Included (1=Yes)	Workers' Comp Included? (1=Yes)	WTP (2009$)	SE
Smith (1974)	USA	CPS 1967; Census of Manufactures 1960; Employment and Earnings 1963	3183	0	1	1	0	0	BLS 1966, 1967	0.000125	1	0	14.06	5.87
Viscusi (1978)	USA	SWC 1969-1970	496	0	0	0	0	1	BLS, subjective risk of job (SWC)	0.000118	1	0	3.72	2.15
Olson (1981)	USA	CPS 1978	5993	0	0	0	0	0	BLS 1973	0.0001	1	0	18.15	7.30
Viscusi (1981)	USA	PSID 1976	3977	0	0	0	0	0	BLS 1973-1976	0.000104	1	0	12.33	2.13
Marin and Psacharopoulos (1982)	UK	General Household Survey 1975	5509	0	0	0	0	0	OPCS Occupational Mortality Decenniel Survey 1970-1972	0.00009	0	0	9.09	2.01
Dorsey and Walzer (1983)	USA	CPS May 1978	1697	1	0	0	0	1	BLS 1976	0.000058	1	1	17.27	7.29
Dillingham and Smith (1984)	USA	CPS May 1979	879	0	1	0	0	0	BLS industry data 1976, 1979	0.00012	1	0	4.81	2.30
Leigh and Folsom (1984)	USA	PSID 1974, QES 1977	1529	0	1	0	0	0	BLS	0.00014	1	0	15.12	6.40
Dillingham (1985)	USA	QES 1977	514	0	0	0	0	0	BLS 1976; NY workers' compensation data 1970	0.00014	0	0	6.21	3.47

Study	Country	Sample	Sample Size	Sample Characteristics: Union	White	Male	Manual/Mfg	Blue Collar	Risk Variable	Mean Risk	Nonfatal Risk Included (1=Yes)	Workers' Comp Included? (1=Yes)	WTP (2009$)	SE
Weiss et al. (1986)	Austria	Austrian Microcensus File of Central Bureau of Statistics 1981	4225	0	0	0	0	0	Austrian Social Insurance Data on Job-related Accidents 1977-1984	0.00013	1	0	12.23	5.03
Moore and Viscusi (1988)	USA	PSID 1982	1349	0	1	0	0	0	BLS 1972-1982, NIOSH National NTOF Survey 1980-85	0.00008	0	1	13.15	5.21
Garen (1988)	USA	PSID 1981-1982	2863	0	0	0	0	1	BLS 1980, 1981	0.000108	1	0	24.08	5.17
Meng (1989)	Canada	National Survey of Class Structure and Labour Process 1981	718	0	0	0	0	0	Labour Canada and Quebec Occupational Health and Safety Board 1981	0.00019	0	0	6.85	3.99
Meng and Smith (1990)	Canada	National Election Survey	777	0	0	0	1	0	Labour Canada and Quebec Occupational Health and Safety Board 1981-83	0.00012	0	0	1.78	3.28
Berger and Gabriel (1991)	USA	1980 Census	22837	0	0	1	0	0	BLS 1979	0.000097	0	0	11.17	1.95
Leigh (1991)	USA	PSID 1974, 1981	1502	0	0	1	0	1	BLS 1979	0.000134	0	0	10.74	3.23
Kniesner and Leeth (1991)	USA	CPS 1978	8868	0	0	0	1	0	NIOSH NTOF Survey 1980-1985	0.000436	1	1	0.67	0.46
Gegax (1991)	USA	Authors' mail survey 1984	228	1	0	0	0	0	Workers' assessed fatality risk at work 1984	0.00086	0	0	3.92	1.99
Martinello and Meng (1992)	Canada	Labor Market Activity Survey 1986	4352	0	0	0	1	0	Labor Canada and Statistics Canada 1986	0.00025	1	0	4.45	1.34

Study	Country	Sample	Sample Size	Sample Characteristics: Union	White	Male	Manual/Mfg	Blue Collar	Risk Variable	Mean Risk	Nonfatal Risk Included (1=Yes)	Workers' Comp Included? (1=Yes)	WTP (2009$)	SE
Cousineau et al. (1992)	Canada	Labor Canada Survey 1979	32713	0	0	0	1	0	Quebec Compensation Board	7.64E-05	1	0	7.01	0.67
Siebert and Wei (1994)	UK	General Household Survey 1983	1353	1	0	1	1	0	Health and Safety Executive 1986-88	3.32E-05	1	0	20.70	9.85
Leigh (1995)	USA	PSID 1981	1528	0	0	1	0	1	NIOSH 1980-85	0.00011	0	0	16.23	3.04
Sandy and Elliot (1996)	UK	Social Change and Economic Life Initiative Survey 1986	440	0	0	1	1	0	OPCS Occupational Mortality Decenniel Survey 1979/80-1982/83	4.52E-05	0	0	76.00	32.55
Milleret al. (1997)	Australia	Australian Census of Population and Housing 1991	18,850	0	0	1	0	0	Worksafe Australia, National Occupational Health and Safety Commission 1992-93	0.000068	0	0	23.86	1.82
Meng and Smith (1999)	Canada	Labor Market Activity Survey 1986	1503	0	0	0	0	0	Ontario Workers' Compensation Board	0.00018	1	1	3.33	0.86
Kim and Fishback (1999)	South Korea	Ministry of Labor's Report on Monthly Labor Survey and Survey on Basic Statistics for the Wage Structures	321	0	0	1	0	0	Ministry of Labor's Analysis for Industrial Accidents	0.000485	1	1	2.20	0.45
Arabsheibani and Marin (2000)	UK	General Household Survey (1980s)	3608	0	0	1	0	0	OPCS Occupational Mortality Decennial Survey 1979-80	0.00005	1	0	43.88	8.82

Study	Country	Sample	Sample Size	Sample Characteristics: Union	White	Male	Manual/Mfg	Blue Collar	Risk Variable	Mean Risk	Nonfatal Risk Included (1=Yes)	Workers' Comp Included? (1=Yes)	WTP (2009$)	SE
Gunderson and Hyatt (2001)	Canada	Survey of Ontario Workers with Perma- nent Impairment	2014	0	0	0	0	1	Ontario Workers' Compensation Board	0.000167	1		34.03	4.83
Viscusi (2003)	USA	CPS MORG 1997	83625	0	1	0	0	0	CFOI 1992-1997	3.62E-05	1	1	21.45	2.01
Leeth and Ruser (2003)	USA	CPS ORG 1996-98	45001	0	0	1	0	1	CFOI 1996-1998	9.76E-05	1	1	3.61	0.80
Smith et al. (2004)	USA	Health & Retirement Survey (Wave 1)	3632	0	0	0	0	0	BLS 1993	5.8E-05	0	0	7.97	
Viscusi (2004)	USA	CPS MORG 1997	99033	0	0	0	0	0	CFOI 1992-1997	4.02E-05	1	1	6.79	0.80
Kniesner et al. (2006)	USA	PSID 1997	1875	0	0	1	0	0	CFOI 1992-1997	0.00004	0	0	29.59	
Viscusi and Aldy (2007)	USA	CPS MORG 1992-1997	120,008	0	0	0	0	0	CFOI 1992-1997	0.00004	1	1	12.23	
Aldy and Viscusi (2008)	USA	CPS MORG 1993-1997	123,439	0	0	0	0	0	CFOI 1992-2000		1	1	13.09	
Evans and Smith (2008)	USA	Health & Retirement Survey	2,708	0	0	0	0	0	CFOI	0.000064	0	0	13.06	
Scotton and Taylor (2009)	USA	CPS MORG 1996-1998	43,261	0	0	0	0	0	CFOI 1992-1997	4.895E-05	1	0	6.16	1.89

Appendix A

This appendix gives some illustrative numerical examples using the simple static (single-period) structural benefit transfer function from Section 5.2.1, and a more formal exposition of the life-cycle modeling framework discussed in Section 5.2.2. Table B1 shows willingness to pay values for a range of mortality risk reductions using the static model in Section 5.2.1. The first three columns in the table show the difference between the marginal approximation and the exact *WTP* [$] for a range of changes in baseline risks Δp [yr^{-1}]. The final six columns in the table show *WTP* [$] and m^{**} [yr^{-1}] (explained below) for a range of Δp's and three possible values of β, accounting for the behavioral response described in Section 5.2.1. To determine the maximum willingness to pay for an exogenous change in background mortality risks, we must solve the two-equation system comprised of the equality between expected utility with and without the policy,

$$p_0 - m^* \;\ln\; a\; y + W_0 + \alpha m^{*\beta} \;=\; p_0 + \Delta p - m^{**} \;\ln\; a\; y + W_0 + \alpha m^{**\beta} - WTP \quad ,$$

and the first-order condition for maximized expected utility with respect to job-risk with the policy <u>and</u> a reduction in income equal to *WTP*, i.e.,

$$\beta\alpha m^{**\beta-1} \;\; p_0 + \Delta p - m^{**} \;\; / \;\; y + W_0 + \alpha m^{**\beta} - WTP \;\; - \ln\; a\; y + W_0 + \alpha m^{**\beta} - WTP \;\; = 0 \,,$$

where m^{**} is the job-risk level that the individual would choose if her baseline survival probability were increased by Δp <u>and</u> if she were charged the amount *WTP* for this change. The level of *m* that she would actually choose after the policy is implemented would depend on the actual cost of the policy to her.

The main lesson from these examples is that—when preferences for consumption and risk are not separable, as in this example—if individuals are able to freely adjust their job risk level, then *WTP* generally will be higher and the total number of "statistical lives saved" will be lower than otherwise predicted under the assumption of no behavioral response. In fact, if $\beta = 1$ and if each individual were charged their maximum *WTP* for the change, then the individuals' behavioral responses would fully offset the changes in their baseline mortality risk. In this extreme case, *WTP* would exactly equal $wtp \cdot \Delta p$ and, if each individual had to pay this full amount to fund the policy, then the number of "lives saved" would be zero. If the full costs of the policy were less than the aggregate *WTP*, then both the net social benefits and the number of statistical lives saved would be positive, though the latter still would be less than $\Delta p \times N$. If the full costs of the policy were greater than the aggregate *WTP*, then of course the net social benefits would be negative, but also note that the number of statistical lives "saved" would be negative as well—that is, even though environmental risks were reduced, the policy would *increase* overall mortality rates since people's behavioral responses to the increased costs would involve shifting to jobs with higher mortality risks. The numerical results in Table B1 are not necessarily intended to be realistic, especially considering that they involve mortality risk reductions that are much larger than those we would typically expect from most environmental regulations, but they nevertheless highlight the importance of calculating benefits and costs simultaneously for non-marginal policies when behavioral adjustments are expected.

Next, a brief exposition of a generalized life-cycle (multi-period) model may help to describe the potential usefulness of this framework as a basis for structural benefit transfers of mortality risk reductions. Suppose that the value function for a representative individual is given by $V_a = \sum_{t=a}^{T} u(c_t, h_t, t)\, s_{a,t} e^{-\rho(t-a)}$, where $u(c_t, h_t, t)$ is utility in period *t* (assumed here to depend on consumption c_t, health status h_t, and possibly age *t*), s_τ is the probability of surviving to the beginning of age $\tau + 1$ given that the individual is alive at the beginning of age τ, $s_{a,t} = \prod_{\tau=a}^{t} s_\tau$, and *T* is the individual's maximum possible lifespan. Marginal willingness to pay at age *a* for mortality risk reductions (or, equivalently, an increase in survival probability) at age *b* ($\geq a$) is $wtp_{a,b} \equiv \frac{dc_a}{ds_b} = \frac{\partial V_a / \partial s_b}{\partial V_a / \partial c_a}$.

To help interpret this willingness to pay measure, we can break the value function into two parts at some future age $t = b$, $V_a = \sum_{t=a}^{b-1} u(c_t, h_t, t) s_{a,t} e^{-\rho(t-a)} + \sum_{t=b}^{T} u(c_t, h_t, t) s_{a,t} e^{-\rho(t-a)}$, then re-write second term on the right hand side of this equation in terms of the value function at age *b*, $V_a = \sum_{t=a}^{b-1} u(c_t, h_t, t) s_{a,t} e^{-\rho(t-a)} + V_b s_{a,b} e^{-\rho(b-a)}$, which means $\frac{\partial V_a}{\partial s_b} = V_b s_{a,b-1} e^{-\rho(b-a)}$.[33] Thus, the marginal willingness to pay at age *a* for a reduction in mortality risk at some future age *b* is

$$wtp_{a,b} = \frac{V_b s_{a,b-1} e^{-\rho(b-a)}}{\partial u(c_a, h_a, a) / \partial c_a}.$$[34]

This is the expected remaining lifetime utility at the beginning of age *b*, discounted by the survival probability and the pure rate of time preference between ages *a* and *b*, and then monetized by the marginal utility of consumption at age *a*.

Developing a usable structural benefit-transfer function based on a lifecycle framework would be challenging. Estimating or calibrating such a model would require specifying or solving for the life-cycle pattern of consumption, calibrating or estimating the pure rate of time preference, and specifying a

---

[33] Throughout this section we treat the path of consumption over the life cycle as exogenous; that is, we ignore any behavioral responses to changes in mortality risks that would adjust the levels of consumption in future periods. This simplification will be strictly valid only under some special conditions—namely, that that the individual can never be a net borrower (Cropper and Sussman 1990, USEPA 2007 p D-15)—but it should provide a close approximation for small changes in exogenous mortality risks. More specifically, we would expect it to provide a close lower bound on willingness to pay in most cases of interest—a lower bound because it assumes that the individual is constrained to maintain the same consumption path after the change, and a close approximation because we would expected any adjustments in future consumption levels to be very small for reasonably small changes in mortality risks.

[34] Direct inspection of this equation suggests some simple comparative static results: (1) $wtp_{a,b}$ decreases with the latency period $b - a$ because all elements of the numerator— V_{b+1}, $s_{a,b-1}$, and $e^{-\rho(b-a)}$ —decrease and the denominator does not change. (2) $wtp_{a,a}$ could increase or decrease with *a* because, while V_{a+1} and $s_{a,a-1}$ decrease with *a*, the denominator could decrease or increase with *a* depending on the pattern of consumption and health status over the life cycle (USEPA 2007 p D-16). If the pattern of consumption were perfectly flat over the life cycle, and if utility depended only on consumption and not health status or age per se, then $wtp_{a,a}$ would unambiguously decrease with age. However, observed consumption patterns generally are not flat; consumption typically is low in the early (adult) years, high in middle age, and lower again in later years, which, all else equal, would tend to increase then decrease $wtp_{a,a}$.

functional form for the period utility function $u(c_t, h_t, t)$ and calibrating or estimating its parameters.

The simplest reasonable implementation of such an approach might proceed as follows:

1.) Specify the lifetime pattern of consumption for a "representative" individual as the pattern of average consumption levels for a random sample of individuals of various ages from the population of interest. Alternatively, multiple representative life-cycle consumption patterns could be generated based on average consumption levels for sub-samples of the population, e.g., by gender, race, geographic region, etc., as appropriate for the exposed sub-population relevant for the policy to be examined.

2.) Set ρ equal to a suitable central value from a relevant set of revealed or stated preference studies (presumably somewhere between, say, 0% and 5% per year).

3.) Assume the utility function is of the standard CRRA form with a lower bound on utility: $u_t = (c_t^{1-\eta} - d^{1-\eta}) / (1-\eta)$. Then either

    a. set η equal to a suitable central value from a relevant set of revealed or stated preference studies (presumably somewhere between, say, 0.5 and 3), and use at least one valid estimate of willingness to pay for well-specified mortality risk changes from the revealed or stated preference literature to calibrate *d*, or

    b. use at least two valid estimates of marginal willingness to pay from the RP or SP literature to calibrate η and *d* simultaneously.

Such a calibrated life-cycle model then could be used to calculate $wtp_{a,b}$ for all combinations of *a* and *b* for each representative individual identified in step 1. These estimates then could be transferred to any pattern of mortality risk changes that are projected for one or more policies under consideration. More sophisticated versions of this approach could specify u_t as a function of age and/or health status, which might facilitate a link to the QALY literature.

<u>Table A1</u>. Maximum willingness to pay for a range of changes in survival probabilities, Δp, based on a marginal approximation ($wtp \cdot \Delta p$) and direct calculation (WTP), with and without a behavioral response. Baseline job risk is $m = 0.006$. Estimates of the adjusted job risk with a behavioral response (m^{**}) assume that the individual's income is simultaneously reduced by *WTP* (that is, expected utility without the policy is equal to that with the policy combined with the charge *WTP*).

	No behavioral response		With behavioral response					
			$\beta = 0.33$		$\beta = 0.67$		$\beta = 1$	
Δp [yr^{-1}]	$wtp \cdot \Delta p$ [$]	WTP [$]	WTP [$]	m^{**} [yr^{-1}]	WTP [$]	m^{**} [yr^{-1}]	WTP [$]	m^{**} [yr^{-1}]
0.000005	40.0	25.0	40.0	0.0060034	400	0.0060040	40.0	0.0060050
0.00005	400.0	397.7	399.2	0.0060337	399.6	0.0060404	400.0	0.0060500
0.0005	4,000.0	3,778.1	3,926.6	0.0063399	3,956.7	0.0064062	4,000.0	0.0065000
0.005	40,000.0	23,773.6	33,977.6	0.0096250	36,431.3	0.0102351	40,000.0	0.0110000

Oxford Review of Economic Policy, Volume 36, Number S1, 2020, pp. S56–S63

What is the meaning of (statistical) life? Benefit–cost analysis in the time of COVID-19

Jonathan Colmer*

Abstract Efforts to support public policy decisions need to be conducted carefully and thoughtfully. Recent efforts to estimate the social benefits of reductions in mortality risks associated with COVID-19 interventions are likely understated. There are large uncertainties over how much larger the social benefits could be. This raises questions about how helpful conventional approaches to valuing mortality and morbidity risks for benefit–cost analyses can be in contexts such as the current crisis.

Keywords: value of a statistical life, benefit-cost analysis, COVID-19

JEL classification: A13, D61, I18, J17

I. Introduction

Are the economic costs of policy interventions to limit the spread of COVID-19 worth the potential health benefits? How long should social distancing and shelter-in-place rules be in place for? These are the multi-trillion dollar questions that economists, public health experts, politicians, and your neighbour are currently debating.

But how do we answer these questions? When considering any policy intervention, it is important to get a sense of the costs and benefits of different courses of action. In many settings, including the COVID-19 pandemic, this is particularly challenging. While many of the costs, such as forgone income and production, are easily measured in your currency of choice, measuring the benefits is a lot harder.[1]

In a recent interview the Governor for New York, Andrew Cuomo declared 'you cannot put a value on a human life.' (Cuomo, 2020). He is correct. Researchers instead estimate what is known as the value of a statistical life (VSL) (Drèze,1962; Schelling, 1968; Jones-Lee, 1974; Viscusi, 1992; Viscusi and Aldy, 2003; Banzhaf, 2014).

*Department of Economics, University of Virginia; e-mail: j.colmer@virginia.edu

I thank David Bradford, Trudy Cameron, Andres Clarens, Scott Doney, Leora Friedberg, Cameron Hepburn, Michael Livermore, Jacquelyn Pless, Jay Shimshack, Bill Shobe, Alex Teytelboym, and Kip Viscusi for helpful comments and suggestions.

[1] In the context of COVID-19, some costs are also very difficult to quantify, e.g. forgone opportunities to enjoy time with friends and family, mental and physical health risks from sedentary behaviour and loneliness, etc.

doi:10.1093/oxrep/graa022

© The Author(s) 2020. Published by Oxford University Press.
For permissions please e-mail: journals.permissions@oup.com

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

VSL estimates do not measure the value of life. The VSL represents aggregate demand for wide-spread, but individually very small, reductions in mortality risk, i.e. how much individuals are willing to pay for a very small reduction in the probability of death, paid for by forgoing the consumption of other goods and services. For example, if a policy reduced the risk of death by 0.00001 per person, 1 in 100,000, then it would take 100,000 people to accumulate a *collective* risk reduction of one 'statistical life'. If, on average, each individual is willing to pay $100 per year to reduce the probability of dying by 0.00001, then collectively the group would be willing-to-pay $10m per year to prevent the loss of one 'statistical life'. This is the value of a statistical life. If the policy applied to a large fraction of the country's population, say 100m people (a thousand times as many), then the expected total number of lives saved would be 1,000 'statistical lives', with a collective willingness to pay of $10 billion.

VSL estimates play a very large role in determining benefit–cost assessments for many government policies in the areas of health, transportation, and the environment. In a recent review of 115 major federal regulations in the United States, 70 per cent of the total benefits were directly attributable to the monetized value of reducing early mortality (OMB, 2014).

A considerable amount of work and effort goes into credibly estimating the value of a statistical life and there are a number of challenging issues that researchers need to account for when estimating and interpreting VSL estimates (see Viscusi and Aldy (2003) and Viscusi (2018) for an overview of estimates from the literature and Ashenfelter (2006), Cameron (2010), and Viscusi (2011) for a discussion of the empirical and conceptual issues). A broader issue is how any given VSL estimate is used in benefit–cost analysis. Almost all benefit–cost analyses apply VSL estimates to new populations. As such, it is important that VSL estimates are chosen carefully, and that the assumptions that give VSL estimates a meaningful interpretation are plausible, when applied to new contexts. These considerations are of particular importance in the context of COVID-19.

The COVID-19 crisis has prompted a slew of benefit–cost analyses using VSL estimates. This article calls for introspection. I argue that recent efforts have, in many cases, lacked clarity on some of these issues and may consequently have underestimated the social benefits of policy interventions to mitigate the spread of COVID-19. Ultimately, more questions are raised than answered. However, one conclusion is clear: when benefit–cost analysis is implemented, we need to ensure that it is implemented carefully and thoughtfully to most effectively serve and support public policy decisions.

II. Large vs small reductions in risk

VSL estimates are only valid when based on very small (marginal) changes in risk. When there are large changes in risk, VSL estimates may be too small. An important consideration in the face of large changes in risk is the distinction between willingness to pay (WTP) for a reduction in risk and willingness to accept (WTA) compensation to forgo reductions in mortality risk. When risks are small these distinctions don't matter, as WTP ≈ WTA. However, when considering larger reductions in risk, an individual's

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

WTA compensation to forgo a reduction in mortality risk will be larger than their WTP for the same reduction in mortality risk, WTA > WTP.[2]

Consider the following example. Figure 1 shows an individual's willingness to accept (*a*) and willingness to pay (*b*) for a large increase/decrease in risk. Moving from p to p_1, which reduces survival probabilities, we observe that the individual would need an increase in income, moving from $w(p)$ to $w(p_1)$. By contrast, reducing risk by the same amount, from p to p_2, we observe that the individual would be willing to reduce their income, moving from $w(p)$ to $w(p_2)$. The willingness to accept the large increase in risk, $w(p_1)–w(p)$, is much larger than the willingness to pay for a reduction in risk of the same magnitude, $w(p)–w(p_2)$. There are two additional insights that can be drawn from Figure 1. First, the VSL (represented by the red and blue lines, which captures the marginal rate of substitution between income and mortality risk, $–dw(p)/dp$) will be higher for people with higher baseline risk, represented by the steeper gradient. Second, in response to large changes in risk, the VSL based on willingness to accept/pay (represented by the red lines in Figure 1) will be larger/smaller than the VSL for small changes in risk (the blue line).

Whether willingness to pay or willingness to accept is the more appropriate measure relates to property rights. WTP is appropriate when the individual does not have property rights—we want to know what the individual would pay for something to which they currently have no legal right. WTA is appropriate when the individual does have property rights—we are compensating the individual for the loss of something to which they do have a legal right. The right to life is conventionally regarded as the primary natural or human right (Garnsey, 2007). As such, when considering large reductions in mortality risk, WTA is arguably the more appropriate measure.

Figure 1: WTA vs WTP for large changes in risk

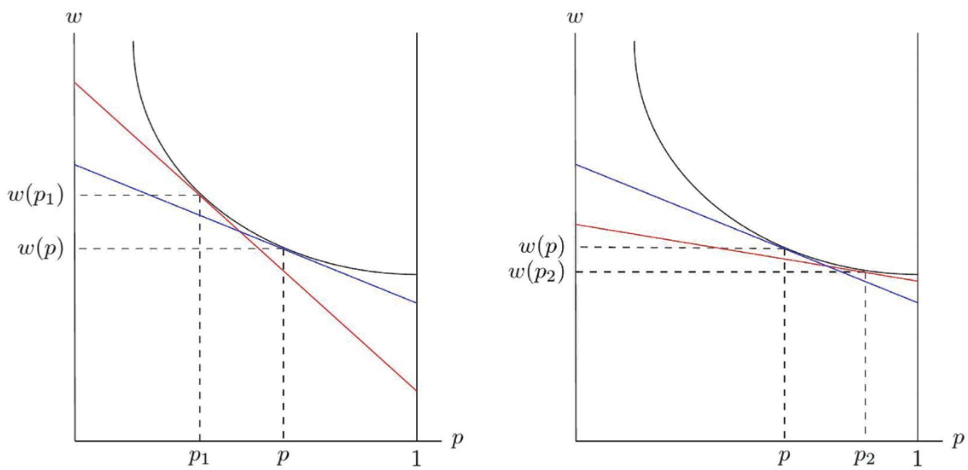

(*a*) WTA an increase in risk p to p_1. (*b*) WTP for a reduction in risk p to p_2

[2] The extrapolation of VSL estimates based on very small changes in risks to larger changes implicitly assumes that people are risk neutral. This assumption is implausible.

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

The COVID-19 pandemic reflects a context in which increases in mortality risk are large, at least for some groups of the population. VSL estimates are most often derived from hedonic wage models—a revealed preference approach. Researchers estimate the additional wages paid to workers in riskier jobs, capturing the compensating wage differentials associated with greater mortality risk. As such, VSL estimates are usually derived from contexts in which changes in risk and wages are relatively small. The wage–risk trade-offs upon which we base benefit–cost analyses for typically sized risk-reductions cannot confidently be extrapolated to the risks presented by COVID-19. In technical terms, the magnitude of the mortality risks avoided by social distancing and shelter-in-place policies is outside the domain of the risk reductions that have been used to estimate the VSL. Use of VSL estimates derived from very small changes in mortality risk could substantially underestimate the social benefits of policy interventions that seek to reduce risk of exposure. Any benefit–cost analysis of interventions that induce non-marginal reductions in mortality risk should be explicit about this issue.

III. Assumption: everyone is a middle-aged, male, blue-collar worker . . .

It is very easy, and tempting, to apply VSL estimates to new settings. Policy-makers crave a single number that they can employ in any context, and it would indeed be much easier if we could be confident that there was just *one* number that was always the right number to use when measuring the social benefits of risk reduction. However, the value of a statistical life is not a universal constant.

Willingness to pay for changes in mortality risk will be different for different people. Such differences may arise for a number of reasons, including: differences in how informed individuals are about risk reductions; differences in life expectancy; differences in risk aversion; differences in cultural or theological beliefs (e.g. some groups may believe in fatalism, the belief that fate governs major life outcomes, or theological determinism, the belief that all events are pre-ordained); differences in income. In practice, VSL estimates reflect the risk preferences of middle-aged, male, blue-collar workers. There is less credible evidence about the risk preferences of people that are not in this demographic, including the elderly, those with pre-existing health conditions that prevent them from working, carers, students, and children below working age. In the context of COVID-19, the people who face the greatest risks (the elderly, those with pre-existing conditions, and service workers) are not well represented in the populations that inform existing VSL estimates.

Willingness to pay for risk reductions depends on the nature of the risk in question. The context for many VSL estimates is sudden deaths due to industrial accidents. The type of risk presented by COVID-19 does not match the profile of risk used to support existing VSL estimates. COVID-19 presents unfamiliar and poorly understood risks that can involve weeks of severe illness, hospitalization, ventilators, and isolation from family and friends, before a lonely death among strangers. When a victim recovers, we do not yet know the extent of recovery, whether the disease will recur, or whether there are any long-term compromises to health. Cameron and DeShazo (2013) present suggestive evidence that willingness to pay is smaller when considering the risk of sudden death compared to the risk of death following illness.

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

IV. Should we use age-specific VSLs?

One of the most controversial issues associated with the use of VSL in benefit–cost analysis is whether VSL estimates should be differentiated by age. The controversy emerges because individuals think it unfair to value the lives of some people more than others. But this is not what the VSL represents. It represents individual willingness to pay for reductions in mortality risk. That said, the relationship between age and willingness to pay is theoretically ambiguous and the empirical evidence is mixed (Hammitt, 2007; Krupnick, 2007; Aldy and Viscusi, 2007). One might expect willingness to pay for an extension of life to fall as remaining life expectancy falls, but economic resources and exposure to risk also change over the life cycle and so willingness to pay could increase or remain constant as age increases. However, existing VSL estimates treat every risk reduction in the same way. There is no distinction between reducing the risk of someone who is likely to die in the next few years or reducing the risk of someone who is younger with dependants.

One strategy for constructing age-specific VSL estimates is to convert a given VSL estimate into the value of a statistical life year (VSLY). This is calculated using the following equation,

$$VSLY = \frac{rVSL}{1-(1+r)^{-L_{\bar{a}}}}$$

where $L_{\bar{a}}$ represents the average number of remaining life years for the average person in the sample and r represents the real discount rate, commonly, but arbitrarily, assumed to be 3 per cent. If the average VSL estimate in a sample was \$10.9m (\$2020) and the average individual had 41 expected remaining life years, we would calculate the value of a statistical life year to be \$465,565.25 (\$2020). Age-specific VSL estimates can then be calculated for different ages, using the following equation,

$$VSL_a = \sum_{t=\ell}^{L} \frac{VSLY}{(1+r)^t}$$

where t is the current expected life-years remaining in each age group and r is the real discount rate. For example, if 20-year-olds had 60 expected life-years remaining, VSL_{20} would be \$12.884m. If 80-year-olds had 8 expected life-years remaining, VSL_{80} would be \$3.268m.

There are a number of issues with this approach. First, this approach doesn't reflect underlying heterogeneity in willingness to pay by age—it is based solely on differences in average life expectancy. It completely ignores the complexity of the underlying relationship between willingness to pay for a risk reduction and age. As discussed, one theoretical basis for using declining age-profile is that consumption and income decline with age, but one should model this explicitly for the population under study. Murphy and Topel (2006) provide a framework for this that can be used to construct population-specific age-adjusted VSL estimates. However, neither approach accounts for how willingness to pay varies with baseline risk. An individual facing a high probability of death has little incentive to limit their spending on risk reductions as the probability of survival is low. As such, age-specific VSL estimates may understate the willingness to pay for risk reductions. The net effect of these competing forces is unclear.

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

In the context of COVID-19 the benefits of social distancing are likely to accrue to older populations and so it is reasonable that VSL estimates reflect that. However, whether a declining age-profile is appropriate is unclear. Even if a declining age-profile is justified in normal circumstances, it is unclear whether it is appropriate if exposure to risk increases with age. The current situation is one in which increases in mortality risk are quite large for older populations. As discussed above, this would imply that willingness to pay based on marginal changes in risk is likely understated. If so, the true age-profile for the VSL in the context of COVID-19 could be flat, or even increasing with age if WTP/WTA was high enough. It seems likely that VSLY-derived estimates with a declining age-profile represent a lower bound on the social benefits associated with COVID-19 interventions.

V. Decision-making when lives are identifiable

As is hopefully clear, the use of VSL estimates is only applicable in settings where interventions consider a probabilistic loss of life—statistical lives. Using the value of a statistical life to aid decision-making is completely inappropriate when the lives at risk are identified. Identifiable lives reflect specific people who need help now, for example, individuals who are in need of rescuing.

When considering whether to save the life of a given individual, society's choice cannot be based on the individual's willingness to pay. Their willingness to pay doesn't come into it. The tools of constrained optimization that economists use, and that provide the foundations of benefit–cost analysis, are not appropriate in these circumstances because the identified individual is unable to make a trade-off. Sadly, the COVID-19 pandemic is likely to put decision-makers in a position in which they have to make choices about identifiable lives. Society's choice as to whether to give an ICU bed to one patient or another has nothing to do with the patients' willingness to pay or the willingness to pay of their families. The choice has to be made on some other basis.

VI. The social value of life

The value of a statistical life is an analytical tool used in benefit–cost analysis, but its scope is limited. The VSL reflects private willingness to pay for a very small reduction in mortality risk. It does not capture the value of a person's life to the rest of society. As such, the total benefits of interventions to reduce mortality risk will be understated as the social benefits are not accounted for. This is not a critique of VSL. VSL does not claim to measure such considerations. However, it is important to acknowledge that as an analytical tool the VSL does not provide a complete representation of the total benefits associated with reductions in mortality risk. This is true in any context, including the COVID-19 pandemic.

Henry Moseley was a physicist at Oxford in the early twentieth century, who died during the First World War in August 1915 at the age of 27. His contributions to physics prior to his death were monumental. The social cost of his death in terms of the further contributions he could have made is immeasurable. The social benefit to society

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

associated with the fact that Isaac Newton did not die during the bubonic plague, or that C. S. Lewis and J. R. R. Tolkien did not lose their lives during the First World War, is equally immeasurable. Yet, we are cognizant of their existence. What of the scientists, artists, and innovators who lost their lives before new ideas were even conceived? We are not aware of these forgone benefits, but we are nevertheless worse off because they do not exist.

Society also bears the cost of grief and loss of companionship that is experienced by the family and friends of the deceased (Posner and Sunstein, 2015; Beckerman, 2019). As in the case of identifiable lives, it is not possible to credibly assign a monetary value to such considerations.

Measuring the value of a statistical life in a credible way is hard. Measuring the social costs of the loss of life in a credible way is arguably impossible. Inevitably, this means that such considerations are not included in benefit–cost analyses. However, not including something doesn't mean that we shouldn't acknowledge its existence. The social cost can be acknowledged, even if it is to state that the social benefits of intervention are understated due to our inability to measure the social costs associated with loss of life. Just because something can't be measured, doesn't mean that it has no value.

VII. Conclusion

The value of a statistical life (VSL) can be a very useful tool when used appropriately. However, it is often misunderstood and used inappropriately. In the context of COVID-19 there are three main issues. First, the mortality risks presented by COVID-19 are larger than the risks that typically underpin VSL estimates. When there are large changes in mortality risk, existing VSL estimates likely underestimate the social benefits of risk reductions. Second, the type of risk presented by COVID-19 does not match the profile of risks that underpin existing VSL estimates. Third, the populations at risk from COVID-19 do not match the populations used to support existing VSL estimates. Researchers should be clear to highlight differences and be explicit about the possible implications.

Ultimately, it is incredibly challenging to determine a credible and relevant measure of VSL that is appropriate for benefit–cost analysis in the time of COVID-19. This isn't a problem with VSL *per se*. I argue that existing estimates likely represent a lower bound on the social benefits of reductions in mortality risk; however, there are large uncertainties over how much larger the social benefits could be. This raises questions about how helpful conventional benefit–cost analyses can be in the current crisis.

References

Aldy, J. and Viscusi, W. K. (2007), 'Age Differences in the Value of Statistical Life: Revealed Preference Evidence', *Review of Environmental Economics and Policy*, **1**(2), 241–60.

Ashenfelter, O. (2006), 'Measuring the Value of Statistical Life: Problems and Prospects', *The Economic Journal*, **116**(510), 10–23.

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025

Banzhaf, H. (2014), 'The Cold-war Origins of the Value of Statistical Life', *Journal of Economic Perspectives*, **28**(4), 213–26.
Beckerman, W. (2019), 'The Valuation of Life: A Kantian Perspective', *mimeo*.
Cameron, T. (2010), 'Euthanizing the Value of Statistical Life', *Review of Environmental Economics and Policy*, **4**(2), 161–78.
— DeShazo, J. (2013), 'Demand for Health Risk Reductions', *Journal of Environmental Economics and Management*, **65**, 87–109.
Cuomo, A. (2020), https://www.governor.ny.gov/news/audio-rush-transcript-governor-cuomo-guest-cnns-cuomo-prime-time-1#
Drèze, J. (1962), 'L'utilité sociale d'une vie humaine', *Revue Française de Recherche Opérationnelle*, **6**, 93–118
Garnsey, P. (2007), *Thinking about Property: From Antiquity to the Age of Revolution*. Cambridge, Cambridge University Press.
Hammitt, J. (2007), 'Valuing Changes in Mortality Risk: Lives Saved versus Life Years Saved', *Review of Environmental Economics and Policy*, **1**(2), 228–40.
Jones-Lee, M. (1974), 'The Value of Changes in the Probability of Death or Injury', *Journal of Political Economy*, **82**(4), 835–49.
Krupnick, A. (2007), 'Mortality-risk Valuation and Age: Stated Preference Evidence', *Review of Environmental Economics and Policy*, **1**(2), 261–82.
Murphy, K., and Topel, R. (2006), 'The Value of Health and Longevity', *Journal of Political Economy*, **114**(5), 871–904.
OMB (2014), *2013 Report to Congress on the Benefits and Costs of Federal Regulations and Unfunded Mandate on State, Local, and Tribal Entities*, Washington, DC, US Office of Management and Budget.
Posner, E., and Sunstein, C. R. (2015), 'Dollars and Death', *University of Chicago Law Review*, **72**(2), 537–98.
Schelling, T. (1968), 'The Life You Save May Be Your Own', in J. Samuel, B. Chase (ed.), *Problems of Public Expenditure Analysis*, Washington, DC, Brookings Institution, 127–62.
Viscusi, W. K. (1992), *Fatal Tradeoffs*, New York, Oxford University Press.
— (2011), 'What's to Know? Puzzles in the Literature on the Value of Statistical Life', *Journal of Economic Surveys*, **26**(5), 763–8.
— (2018), 'Best Estimate Selection Bias of Estimates of the Value of a Statistical Life', *Journal of Benefit–Cost Analysis*, **9**(2), 205–46.
— Aldy, J. (2003), 'The Value of a Statistical Life: A Critical Review Of Market Estimates throughout the World', *Journal of Risk and Uncertainty*, **27**(1), 5–76.

Downloaded from https://academic.oup.com/oxrep/article/36/Supplement_1/S56/5899010 by guest on 19 April 2025