Journal of Economic Methodology
Publication details, including instructions for authors and subscription information:
http://www.informaworld.com/smpp/title~content=t713704064

Mismeasuring the value of statistical life

Till Grüne-Yanoff[a]
[a] University of Helsinki, Finland

To cite this Article Grüne-Yanoff, Till(2009) 'Mismeasuring the value of statistical life', Journal of Economic Methodology, 16: 2, 109 — 123
To link to this Article: DOI: 10.1080/13501780902940703
URL: http://dx.doi.org/10.1080/13501780902940703

Journal of Economic Methodology,
Vol. 16, No. 2, June 2009, 109–123

Routledge
Taylor & Francis Group

Mismeasuring the value of statistical life

Till Grüne-Yanoff*

University of Helsinki, Finland

The value of a statistical life (VSL) is an important tool for cost–benefit analysis of regulatory policies that concern fatality risks. Its proponents claim that it measures people's risk preferences, and that VSL therefore is a tool of vicarious governance. This paper criticizes the revealed preference method for measuring VSL. It specifies three minimal conditions for vicarious governance: sensitivity, fairness and hypothetical compromise, and shows that the VSL measure, in its common application in policy formation and analysis, violates these conditions. It therefore concludes that the revealed preference VSL measure, in its current form, is not a tool of vicarious governance.

Keywords: cost–benefit analysis; risk; value of life; vicarious governance

1 Introduction

Governments, inevitably, make life-or-death decisions. Their policies and regulations affect the fatality risk exposure of their citizens. Statistically, the increase of fatality risk exposure will increase the number of fatalities in a population of a given size. Policy decisions affecting fatality risk therefore either save or sacrifice 'statistical lives' – the lives of yet unidentified people.

It is therefore desirable that the policy maker adheres to the strictest standards of rationality and transparency when taking these life-or-death decisions. At the very minimum, rationality requires the comparison of all policy alternatives with respect to their outcomes. In the cases of interest here, policies are compared according to their reduction of fatalities in number and probability, but also according to the costly side effects they produce. This requires making fatality risk changes commensurable with other and possibly more common goods. The loss of statistical lives is thus weighed against monetary costs, preservation of workplaces, strategic national interests, etc.

For transparency reasons, governments choosing policies that affect fatality risks must put a value on statistical life. This is a difficult task that many would want to avoid and some may find morally repugnant. But welfare economics seems to offer a way out of this conundrum. Instead of requiring the policy maker to determine the value of statistical life herself, welfare economics offers the tools to measure people's own evaluation of small changes in fatality risk and aggregate these evaluations into a single measure. A political process of evaluating life by the policy maker is thus replaced by the measurement of others' evaluations. This value-of-statistical life (VSL) measure then allows the policy maker to govern *vicariously*: it realizes that level of fatality risk that citizens value highest.

*Email: till.grune@helsinki.fi

ISSN 1350-178X print/ISSN 1469-9427 online

DOI: 10.1080/13501780902940703
http://www.informaworld.com

I argue that VSL measure does not fulfil its assigned purpose. In particular, I argue that because of the way it is measured, it does not satisfy the minimal conditions for vicarious governance: preference sensitivity, fairness and hypothetical compromise. In section 2, I specify the notion of vicarious governance, and discuss its three minimal conditions. Section 3 examines the VSL measurement process. Section 4 argues that the common usage of VSL as a transferable and universal measure violates preference sensitivity. Section 5 argues that the way VSL is currently measured violates preference sensitivity, fairness and hypothetical compromise. Section 6 concludes.

2 Vicarious governance

Every country's government is deeply involved in programmes that affect its citizens' exposure to fatal risks. The citizens cannot realize many of these policies themselves, because their realization would be inefficient (e.g. natural monopolies), because their usage cannot be individually charged (e.g. public goods) or because some institutions require by their nature a centralized institution (e.g. legal system). Examples of such risk-affecting policies range from workplace safety regulations to international diplomacy. Many governments have sought to take decisions of this sort in a transparent and justifiable manner. They seek to adhere to a scheme of rational deliberation, which identifies all available policy options, determines the consequences of each option, assigns a value and a likelihood estimate to all these consequences, and determines the 'best' option according to some rational decision rule on the basis of this model.[1]

Crucially, this procedure involves assigning a value to all of the options' consequences. For policies that affect the fatality risk exposure of a government's citizens, this involves evaluating the change in number of deaths in the population. Many governments have responded to this difficult and controversial task by (implicit) appeal to their role as *vicarious* policy makers and regulators. Instead of arriving at an evaluation by themselves, they use a measure of citizens' subjective evaluation of fatal risks as the basis of rational policy decision. This practice is based on the following principle (P):

> *Principle of vicarious risk regulation*: a government should pursue a change in fatality risk if and only if the beneficiaries value the change sufficiently to be willing to bear the costs.[2]

In the context of VSL, there is little systematic discussion of how (P) may be justified and implemented. Yet in the VSL literature, three arguments are often cited. It is claimed that a policy adhering to (P) is more *efficient*: because it strives to choose the policy on the basis of values that are constructed from people's preferences, it tends to satisfy people's preferences better than other policies that are not constructed this way. Further, it is claimed that a policy adhering to (P) minimizes violations of *citizens' sovereignty*. It is the citizens' preferences, after all, that formed the evaluative basis for the government's decision. The government takes the decision, because the citizens wanted it. Finally, in cases of heterogeneous and conflicting preferences in the population, a policy cannot satisfy all preferences, nor can it avoid going against what some people want. However, it is claimed that a policy adhering to (P) in such situations can at least be justified as adhering to a *democratic principle*, because it is based on the equal influence from all people's preferences, and will be dominated by the majority's will. In these cases, (P) may be claimed to be a

> variant of the 'one person, one vote' principle which will reflect the strength of the individual preferences, bearing in mind the constraints imposed by the overall constraints of resources. (Jones-Lee 1989, p. 11)

(P) and its justifications are controversial and have been repeatedly criticized. This paper does not aim its criticism at (P). Rather, it uses (P) as a benchmark. (P) defines the role that VSL should play in policy formation. This paper investigates whether the current way VSL is measured can satisfy this requirement.

For a policy to adhere to (P), and to its justifications sketched above, the policy's formation process has to satisfy at least four conditions. First, it has to *accurately* represent individuals' preferences. It has to pick out preferences that people actually hold, and avoid both ideological biases and measurement errors. Otherwise, the policy will not be efficient, and will not respect citizens' sovereignty. Second, it has to be *sensitive* to citizens' individual preferences. It must pick up the contextual influence on people's preferences, and allow for changes in preferences due to changing environmental contexts. It must not assume an illegitimately high level of abstraction with respect to individual preferences and the contexts in which they are held. Otherwise, the policy will not be efficient, and will not respect citizens' sovereignty. Third, in cases of heterogeneous and thus possibly conflicting preferences, the formation process has to be based on an *agreed mode of compromise*. In cases of conflict, the result of the formation process will not be supported by all citizens, but the minority at least must accept the process as legitimate. In an actual democratic vote, this means that the minority registers its disagreement with the majority vote, but accepts as legitimate the rule that the majority vote determines the policy. In the measurement of values discussed here, actual voting is absent. In addition, a measurement process may lead to results significantly different from those of a voting process. For these cases, the agreed compromise condition stipulates that the process is explicitly agreed upon, and not just taken to be some default that is suggested by the form of the measurement process. Otherwise, the policy will not respect citizens' sovereignty, and will not respect democratic principles. Last, in the case of heterogeneity, the formation process has to take into account only *relevant* risk preferences. If risk preferences are context-specific, and some individuals are – through their personal characteristics like sex, age, occupation, etc. – never exposed to certain risks, then these individuals may evaluate these risks differently from those exposed to these risks. When measuring and aggregating heterogeneous risk preferences (instead of, say, opinions all things considered), the risk preferences not exposed to the risk maybe should not be taken into consideration when forming the policy. Thus, the policy formation process requires a relevance criterion; otherwise, the policy will violate efficiency and possibly will not respect democratic principles.

For vicarious policies affecting fatality risks, VSL is the tool of choice of many governments and international organizations.[3] It measures individuals' fatality risk preferences and aggregates them into a single trade-off function. In this paper, I argue that VSL does not satisfy the four minimal conditions of (P), and therefore is not a legitimate instrument of vicarious governance.

3 Measuring VSL

Two different methods exist to record people's risk preferences. The expressed preference method asks people to state their preferences for specific scenarios. The revealed preference method reconstructs preferences from choice data. National and international institutions make heavy use of both measures. Each has its own specific conceptual and methodological problems. This paper concentrates on revealed preference method alone, although some of the criticisms raised here may be applicable to the expressed preference method.

The revealed method derives people's risk preferences from their choices of risky options. To give a simple illustration, imagine a worker having a choice over three workplace options, which are identical but for their net wage and the number of workers who die per year in a population of 1 million. Let $O_1 = <\$25k,1>$, $O_2 = <\$28k,2>$, $O_3 = <\$30k,3>$. If one observes the worker to choose O_2 from these options, one may conclude (given the worker's rationality) that the worker prefers a wage increase of \$3k to decrease of fatality risk of 1/1million, but that she prefers a decrease of fatality risk of 1/1million to a wage increase of \$2k. Assuming pure self-interest, this implies that the worker values a decrease from 1/3m to 1/2m in *her own* risk of dying from a work-related accident at \$2k, and a decrease from 1/2m to 1/1m at \$3k.

It is obvious that it is difficult, or even impossible, to observe people making such 'pure' risk decisions. People choose between different levels of risks in specific contexts, where other variables inevitably also influence their decisions. To obtain a revealed risk preference measure therefore requires a theoretical analysis of these contexts, on whose basis the 'polluting' factors can be controlled for.

The most common context used for measuring the VSL is the labour market. It is assumed that it is costly for employers to provide safer working conditions, and that workers are sensitive to work risk levels, and require compensation for higher risks. Thus, workers have utility functions of the form:

$$utility = f(wage,\ risk,\ x) \tag{1}$$

where x denotes a set of unspecified but utility-relevant factors. And employers face profit functions of the form:

$$profit = revenue(output, price) - cost(wage,\ hours,\ risk) \tag{2}$$

Based on these functions, a model of the compensating wage differential was developed and has become the theoretical basis for the VSL measure (Thaler and Rosen 1975). Imagine a world with two firms and two workers. The firm's demand for labour decreases with the total costs of employing a worker. Because the cost of a worker increases with safety levels, for any given level of profits the firm must pay workers less as the safety level rises. Figure 1 depicts two firms' isoprofit curves where wage is an

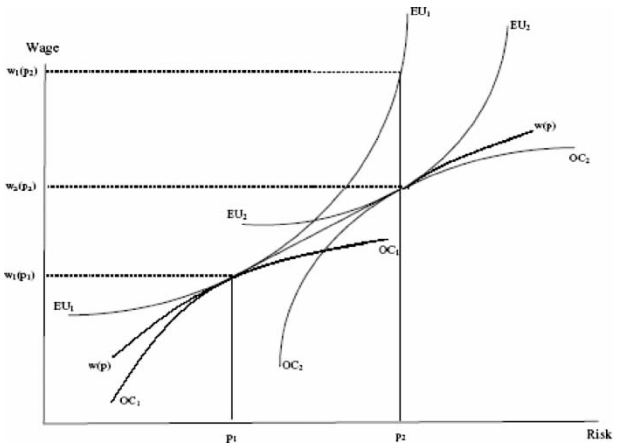

Figure 1. A two-workers/two-firm labour market.

increasing function of risk, OC_1 for firm 1 and OC_2 for firm 2. The worker's supply of labour is in part a function of the worker's preferences over wages and risk. For the worker, risk is bad. Therefore, for a given expected utility level, workers ask for higher wages as safety levels fall. Figure 1 depicts two workers' indifference curves where wage is an increasing function of risk, EU_1 for firm 1 and EU_2 for firm 2.

Equilibrium occurs where the wage–risk trade-off is identical between workers and employers – a marginal change in risk has exactly the same value for both. In Figure 1, this equilibrium point is represented by the tangency between the constant expected utility locus EU_1 and firm 1's offer curve OC_1. Likewise, worker 2 maximizes expected utility at the tangency between EU_2 and OC_2. By increasing the number of workers and employers, one obtains a set of tangency loci. These points are represented by the wage–risk trade-off curve *w(p)*, the *compensating wage differential*. Thus, from the utility and the profit functions a relationship between wages and fatality risks is derived. This relationship is often called the worker's *willingness-to-pay* (WTP): it supposedly captures how much a worker is willing to pay (i.e. have her wage reduced) in order to gain a decrease in fatality risk. The theoretical construct of the WTP lies at the basis of the VSL measure.

This model assumes that a market-clearing equilibrium is reached in the labour market. A number of criticisms against this model point out the reasons why such an equilibrium is not reached (cf. Cookson and Dorman 2008). Because no alternative theoretical basis for VSL measurement exists, and the criticism aims at the foundations of standard economic theory, I will not pursue it further here. Instead, let me discuss the methodology of the revealed preference approach based on this theoretical model.

Risk preferences cannot be directly revealed because human agents rarely face 'pure' risk choices. The theoretical model, however, shows how risk affects wages as a separate factor. It apparently justifies the estimation of the WTP from labour market data.[4] For the case of the labour market a function of the following form is estimated:

$$w_i = \beta_1 + \beta_2 \times H_i + \beta_3 \times X_i + \beta_4 \times p_i + \beta_5 \times q_i + \varepsilon_i \tag{3}$$

w_i is the individual worker's hourly wage, H_i the worker's relevant personal characteristics (human capital measures, age, union status, etc.), X_i the respective job characteristics (blue-collar, white-collar, type of industry, physical exertion, etc.), p_i the fatality risk of the job and q_i the non-fatal injury risk of the job. Deriving (3) with respect to fatality risk, the derivative $dw/dp = \beta_4$ gives the marginal compensation rate for fatality risk changes. For a decrease of fatality risk from p^1 to p^2, an agent is willing to pay $\beta_4(p^2 - p^1)$.

Three assumptions are necessary for this analysis. First, it is assumed that fatality risks only differ in their probability. How the death comes about is considered irrelevant: whether it occurs through illness or accident, whether it is a slow or quick, consciously experienced or not. Second, it is assumed that *all* the other factors are known, and hence can be accounted for when analysing the relationship between price/wage and fatality risk. Third, it is commonly assumed that the relationship is linear.

The wage equation is estimated using labour market panel data (for an overview of results from these two estimation procedures, see Viscusi and Aldy 2003). Once the equation is estimated, VSL can be read off the equation directly. β_4 identifies the factor by which a risk change influences the compensation demands of single person. This risk change is measured as the change in probability of that person's death. VSL measures the compensation demand of a whole population for a risk change that increases the population's future death statistic by one. That is, if n is the size of the population, $\text{VSL} = n \times \beta_4/(1/n) = \beta_4$. Before addressing the real problems of this measure, a number

of distracting potential objections must be clarified. First, it may be objected that it is in principle impossible to value a human life. To meet this objection, it must be clarified that VSL does not measure the value of an individual life. Rather, it measures the value of small changes in fatality risks of low magnitude. Most people will value their personal life to have an infinite monetary price. The value of an identified, individual life can therefore not be assigned a monetary price when working under the principle of vicarious governance. However, the same does not hold for low-magnitude *risks* of death – the probability of dying under only vaguely specified circumstances within a specific time interval (for criticism of this distinction, see Broome 1978 and 1985). People deliberate and choose on the basis of values that make comparable the benefits of risk avoidance with the costs of doing so. Driving a car to work is riskier than taking public transport, but also more comfortable. Keeping a strict diet-and-exercise regime reduces the chances of a fatal heart attack, but makes life less fun. People have to make choices of this sort all the time. It seems reasonable, then, to attribute to them the ability to make 'a practical judgment – a consumer's choice ... about what it is worth to reduce the risk of death' (Schelling 1983, p. 114). Hence, from a vicarious point of view, it is justified to investigate what people are willing to accept in order to reduce risks to their lives, and to derive their preferences accordingly.

Such a measurement is limited to small changes in low-magnitude fatality risks. If risks are very high, people tend to consider the consequences as involving their near-certain death. Such scenarios will likely lead to categorical refusal, and hence cannot measured by VSL. Similarly, large changes in risk levels may lead people out of their realm of common experience, and may therefore also lead to irrational reaction of categorical refusal (or categorical acceptance, in case of a large risk reduction). Those cases cannot be measured by VSL, either. This is not a problem, however, as practical applications of VSL are generally restricted to small changes in low-level risks.[5]

Another potential objection to VSL arises from the claim that while people prefer what they choose, what they prefer is not necessarily good for them. Even if people's preferences are well considered, and free from cognitive distortions, many people's choices are determined not only by their welfare, but also by what has been called *commitments* – motives that are not connected to an agent's welfare (Sen 1977; Hausman 2005). Commitments play an important role in decisions about risk. Sons follow their fathers in their job choices; when deciding to work in the mines or not, keeping up the family tradition may be a reason to neglect their own risk preferences. Similarly, peer pressure or moral principles may lead people to ignore their own risk preferences. All these aspects may lead to an attribution of risk preferences, which do not maximize agents' welfare. However, vicarious governance and citizen sovereignty do not necessarily require vicarious *welfare* maximization. Vicarious governance attempts to deliver what people want, not what is best for them. In particular, if people consciously and voluntarily include non-welfare maximizing aspects in their deliberations, it is not clear why a vicarious government should correct this in any way.

A third objection addresses those cases where preferences are distorted by lack of information, biases and cognitive shortcomings. Here, people's preferences are unstable, and they come to regret what they chose once it is realized. One cause of such instability may be that for many people, fatality risk is *incommensurable* with other consequences of their options, and that they therefore choose risky options in a random and erratic fashion (Nussbaum 2000). In some instances, this seems a plausible possibility. But an evolutionary consideration shows that such cases must be the exception, not the rule. People are continuously faced with risky options throughout their lives. Choosing the right option is

crucial for survival. We would not be here today if we had not developed a sufficiently well-functioning practice of evaluating risky options. This holds in particular for situations that most people are well acquainted with, whose risk features are well known, and that most people make repeated choices in. Both labour markets and car markets belong into that category. At least in these areas, it does not seem to violate accuracy to use people's choices as indicators of their rational risk preferences.

Yet there is plenty of evidence for people's lack of judgment, bias and inconsistency. In practice, agencies sometimes avoid taking unstable or uninformed preferences as they find them, instead employing techniques that extract better-informed and reflected preferences from agents (for a discussion of these techniques, see Adler and Posner (2000, pp. 1128–1130)). None of these techniques will provide perfect correcting mechanisms. Yet the goal of fully informed preferences is in itself problematic (for a critical account of 'cognitive therapy' that 'confronts desires with relevant information', see Sobel (1994)). An overzealous criticism of uninformed preferences ultimately will conflict with the principle of citizens' sovereignty. After all, as Schelling pointed out, consumer and citizen sovereignty includes the 'inalienable right to make one's own mistakes' (Schelling 1983, p. 146). I therefore think that while this criticism may be serious, it affects the use of VSL more than it affects the way it is measured. In those areas where preferences are sufficiently stable and well considered, the suggested measurement process seems appropriate; in those areas where they are not, VSL should not be used for policy formation.

The objections from the non-evaluability of life, from the welfare irrelevance and the distortion of preferences therefore do not constitute reasons to alter the VSL measure. In contrast, the objections raised in the next two sections force us to reject VSL as it is currently measured and used.

4 Transferability

My first criticism of the VSL measure concerns its alleged transferability. I will illustrate this issue at the hand of the recent ruling of the Federal Drug Administration (FDA) to amend its regulation on food labelling regarding *trans* fatty acids (US Department of Health and Human Services 2003). The FDA justified its recommendation by cost–benefit analysis. The total cost of the regulation is the sum of the total testing costs, total re-labelling costs, and total reformulation costs imposed on the producer. The total health benefits of the regulation were estimated from three components: (i) the changes in trans fat intake that would result from labelling changes; (ii) the changes in health states that would result from intake changes; and (iii) the value of changes in health states in terms of life-years gained, number of cases or deaths avoided and dollar value of such benefits. The FDA measured the dollar value of cases of death avoided as 'the number of statistical deaths prevented multiplied by the willingness to pay to reduce the risk of death' (US Department of Health and Human Services 2003, p. 41488). To quantify this number, the FDA used VSL estimated from labour market data. Comparing costs and benefits, the agency estimated that the regulation in year three after the effective date would yield a total benefit between US$1 billion and US$2.5 billion (depending on discount rate and estimate method of averted deaths). On the basis of this estimate, the FDA recommended the amendment.

Two particularities of the use of VSL become apparent in this example. First, a VSL that is measured in the context of work-related risks is transferred to a context of diet-related risks to health. Second, the FDA, as a federal agency, speaks of '*the* willingness to

pay', assuming that this measure is universal across its different tasks. In this vein, other authors suggested that one unique VSL be applied across agencies: it should be attempted to 'spend up to the same marginal cost-per-life-saved amount for different agencies' (Viscusi 2000).

A valid application of VSL to a context different from that in which it was measured assumes that the difference between contexts does not matter. As discussed in section 3, VSL is measured by controlling for all other factors that have an influence on wage or price. However, various contextual factors have an influence specifically on how agents evaluate fatal risks. First, it has been shown that agents evaluate fatality risk according to the kind of death this risk may bring about: whether it occurs through illness or accident, whether it is slow or quick, consciously experienced or not.[6] This is relevant for the above FDA example. While work-related risks usually involve accidents, are sudden, grave, and concern a relatively short horizon, diet-related risks are usually slow, drawn-out diseases that will have their impact many years, if not decades, from now. Further, it has been shown that the very perception of risk (and not only, as would be obvious, its evaluation) depends on the wealth of the agents investigated (Liu and Hammitt 1999). Last, it has been shown that risk preferences depend on the labour context in more complex ways. Workers in more risky professions exhibited a kind of habituation effect, which led them to be less perceptive of risks than workers of other, less risky professions (Garen 1988). None of these context-dependencies are picked up by the current VSL measure. In fact, the concept of VSL indicates that it is constructed as a single, unique and transferable measure. Unless this framework is expanded into a set of context-dependent $VSL_{context}$ measures, VSL thus violates the sensitivity condition of vicarious governance.

It may be argued that these context-dependencies do not necessarily present a disturbance of the measurement process when performed within the same context. Problems arise, however, when VSL measurements are transferred from one context to the next. Whether the problems are serious depends on whether the contexts are similar enough in all relevant factors. This is currently difficult to determine, as many of these factors are only insufficiently investigated. What is clear, however, is that without giving good reasons for the similarity of contexts, applying VSL measured in one context to another may lead to a distorted representation of the actual preferences people hold in that second context. To the extent that such distortions cannot be ruled out, the current practice of transferring VSL violates the accuracy condition of vicarious governance.

These considerations certainly do not rule out the use of VSL altogether. First of all, VSL can be applied to contexts that are sufficiently similar to its context of measurement. Further, with some ingenuity, new contexts of measurement can often be found that are sufficiently similar to a relevant context of application. In assessing clean-water policies, to give a negative example, the US Environmental Protection Agency (EPA) neglected the opportunity to use an existing similar context to measure the relevant VSL:

> Because a thriving market for bottled water exists and because a desire for safe and healthy drinking water drives much of this market, the EPA's choice to use wages for risky jobs as a proxy for the value of clean drinking water is problematic. There is a commercial market in the very item the EPA was valuing, but the EPA ignored it. (Heinzerling 2002, p. 2324)

Last, if such similar contexts cannot be found, the contingent method may be the best alternative. Asking people what they really prefer within a specified policy context, regardless of all the problems and costs associated with this method, is better than imposing risk preferences that were measured in a context that has little to do with the current situation.

5 Aggregation

My second criticism focuses on the way the VSL measure aggregates different people's risk preferences. As discussed in section 3, the compensating wage differential is estimated from panel data from a large and often heterogeneous population. Each individual is present in the dataset with one choice locus. The wage–risk trade-off function $w(p)$ shown in Figure 1 does not capture the risk preferences of any individual worker. This much even partisans of the VSL admit.

> The estimated trade-off curve w(p) does not imply how a particular worker must be compensated for non-marginal changes in risk. (Viscusi and Aldy 2003, p. 9)

The question then is how this trade-off curve $w(p)$ can be interpreted. What does it measure? How does it relate to any characteristics of the individual workers? Why does it support vicarious governance? Two answers have been given. The first interprets VSL as a *group compensation measure* (Bergstrom 1982; Dehez and Dreze 1982). Such an interpretation does not clarify much. It does not tell us what group property is measured. Nor does it tell us why the group should be compensated for a risk change. Finally, it does not make clear what individual properties $w(p)$ measures. But one thing is clear: $w(p)$ cannot be the willingness-to-pay expressing a risk preference of the group. For that, a social welfare function is needed, which maps individual preferences into a single group preference. The measurement process of VSL certainly does not provide that. Hence, this interpretation leaves the nature of $w(p)$ blurred, and one may suspect that a violation of the accuracy condition of vicarious governance hides behind this fuzziness.

The alternative answer interprets VSL as 'an *average trade-off rate* across different levels of risk' (Viscusi and Aldy 2003, p. 9). This may be understood as the claim that VSL is an average of *individual* trade-off rates. An individual trade-off rate captures the individual's response to a change of risk exposure – by e.g. the introduction of new legislation, new technology, etc. – other factors held equal. In particular, it captures her change in job choices in response to the changed risk, and hence her marginal rate of substitution of income for a change in probability of her death. It is therefore called her individual wage–risk trade-off rate.

Can VSL be interpreted as the average of these individual trade-off rates? I think not. The derivative of the hedonic wage equation with respect to fatality risk is *not* identical to the aggregate of individual marginal rates of substitutions. This becomes clear when considering the following: The standard measurement procedure fits a curve to individuals' choice loci and derives the marginal rate of substitution from this curve. This does not capture any individual's rate of substitution. As Figure 2 illustrates, each of the individuals may have a rate of substitution that is considerably flatter than the estimated curve $w(p)$.

Figure 2 depicts the case where all risk–wage indifference curves are linear. This is a simplified version of the convex indifference curves of Figure 1. In Figure 2(a), the little squares represent individuals' risk–wage choices: given that the job they are on has a certain fatality risk, they require a certain wage. The VSL measurement procedure fits a line $w(p)$ to all the choice loci. The result is the dashed line in Figure 2(a). But each individual worker has a personal risk–wage trade-off line. They are depicted in Figure 2(b). Because the measurement procedure takes only one choice locus from each individual, it does not capture these individual trade-off lines at all. In fact, it would be consistent if all these lines were either flatter than the estimated $w(p)$ – as depicted in Figure 2b – or steeper. Therefore, *w(p)* does *not* represent the average of the individual trade-off lines.

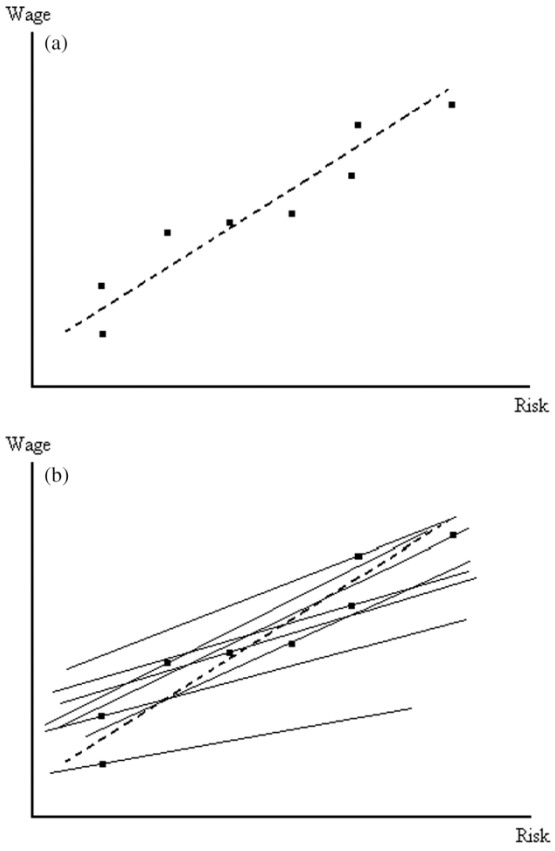

Figure 2. Divergence of *w(p)* and individual risk–wage trade-offs.

Instead, it only measures something that may be called a cross-individual rate of substitution: starting from some individual's wage-risk locus, there is somebody else who is willing to take more risk for a higher wage. This measure cannot support vicarious risk regulation. As discussed above, the choice locus of a (rational) worker in this graph is the tangency between her indifference curve and a firm's offer curve. The shape and position of the firm's offer curve is determined by factors irrelevant to the risk–wage trade-off: by its production capacities, market demand, and the qualities of the worker. The worker's choice locus is crucially influenced by these factors. Fitting a curve across different workers' choice loci will capture the influence of these factors, but will not measure the workers' indifference curves between risk and wage. The described measurement process therefore violates the accuracy condition of vicarious governance.

It is conceivable that this problem may be solved with an improved measurement technique. For example, it may be possible to use time series data in which individuals are identified so that the individual's choices can be observed before and after the change in fatality risk. I do not know, however, of any such research. In any case, the current measurement procedure violates the sensitivity conditions of vicarious governance: it aggregates choice loci in a curve that does not reflect an average of the individual risk–wage trade-off curves at all.

But even if the problem could be technically solved, the question arises whether an average value can be a good instrument of vicarious governance. Imagine that an improved measurement technique can estimate individual hedonic wage functions and derive from them individual wage–risk trade-off coefficients $w_i(p)$. Let the function $V = \Sigma\ w_i(p_i)$ be the average of these individual functions. Commonly, it is assumed that V assigns equal weights to each w_i. This is justified with the claim that every individual in the population is equally exposed to the fatality risk, and hence his or her risk preference should count as much as those of others. In that case, a policy that increases the death statistic by one in the population changes the fatality risk for every individual in the population (with n members) by $\Delta p_i = 1/n$. With w_i linear, V then simplifies to the arithmetic mean of all w_i:

$$V = \sum w_i(p_i) = \sum w_i \times \Delta p_i = 1/n \sum w_i \tag{4}$$

The aggregate VSL measure is the slope of the function V. Figure 3 shows a possible distribution of risk–wage trade-off functions in a population of five individuals. The function V is constructed as the equally weighed average of the individual trade-off functions w_i.

Now, depending on how widely w_i is distributed, any policy evaluated by this mean might be far off the mark.[7] It will provide too little security for the wealthy (who were willing to pay much more for a reduction of risk, but whose willingness to pay was offset by the presence of poorer members of the group), and it will provide too much security at too high a price for the poor (who were not willing to pay such a high price, but whose low bid was counterbalanced by the presence of the wealthy in the group).[8]

Forming the average may seem a fair aggregation procedure. As an extension of the 'one citizen, one vote' principle, VSL might realize a 'one risk-preference, one voice' principle. But the purported analogy hides a crucial difference: majority-voting rules guarantee that those in the majority get what they voted for. However, a majority of risk-takers, for example, may not get what they want from VSL-based policies if there is a significant minority of risk-avoiders, and if both groups are clustered around two distinct risk–wage trade-offs. VSL, by averaging risk preferences, will come down somewhere close to the middle. Thus, a democratic referendum on a particular risk policy in such circumstances would have given a radically different result from basing the policy

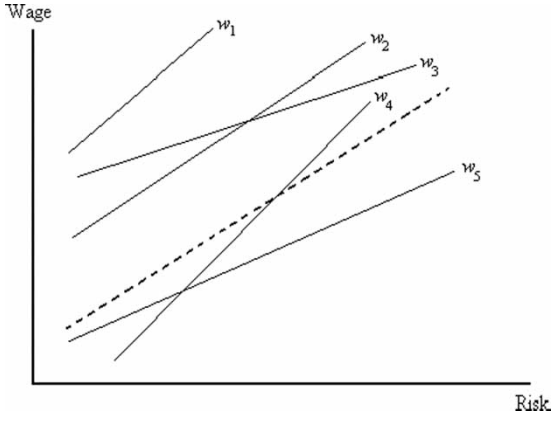

Figure 3. Possible distribution of individual risk–wage trade-offs.

decision on VSL. The referendum would have led to the instalment of the risk-prone policy; the VSL-based policy, based on the average of all risk preferences, would be much more cautious. Such an average might well be a sub-optimal outcome. Risk-takers may not be willing (or able) to pay for the extra safety provided, and risk-avoiders may find safety levels too low. One or even both groups may reasonably prefer a policy that heeds the others' preferences (with appropriate compensation) to a policy that takes the average.

The crucial point, however, is that citizens may not consent to the averaging built into the VSL measure. As it in some cases differs from a majority-voting principle, it cannot be legitimized as an extension of it. Without explicit consent, un-weighted averaging is just a technical default, but not a legitimate procedure. Thus, without such consent, VSL violates the agreed compromise condition of vicarious governance. Because consent may be difficult to reach on this issue, it may be advisable to abstain from using averaging preference measured for policy decisions in such cases of strong preference heterogeneity.

A final problem for VSL arises from the possibility that individuals' risk preferences are systematically correlated with the degree of risk they are exposed to.[9] The improved VSL measure proposed in Equation (4) rests on the assumption that all agents are equally exposed to a given risk. The standard regression analyses use industry or occupation specific fatality frequencies to approximate these risk levels. In the context of fatality risk, there are at least four reasons to believe that risk exposure and risk preferences are correlated. First, affluent people tend to prefer safety more than poor people. White-collar jobs, in which most affluent people work, tend to be safer than blue-collar jobs. Hence risk preferences tend to be correlated with risk exposure. Also, the wealthy have access to safety devices to protect themselves, which the poor do not have. Further, people who value safety more tend to behave more cautiously than those who value safety less (Jones-Lee 1989, chap. 1, n. 15), maybe because their preferences make them more sensitive to risk, and hence more apprehensive of dangerous situations. However, in other contexts, this relation may be reversed: people, because they know that they are cautious and hence not as much exposed to a particular risk as others, may value a reduction in risk less than others who are exposed more. Crudely put, HIV infects the promiscuous, sharp crash barriers hurt motorcyclists, and high ozone levels endanger the elderly. Using those contexts to measure VSL may lead to distorting results, as the unaffected majority may not be concerned with the relevant risks at all, and hold risk preferences accordingly.

All four cases show why it is reasonable to expect a strong correlation between risk preferences and risk exposure.[10] The simplifying assumption made for Equation (4), namely that $\Delta p_i = 1/n$, is not justified for cases of such correlations. Instead, V must be computed as the weighted average of individual trade-off functions, where the weights are the individual risk exposure.

$$V = \sum m_i{}^{*} \Delta p_i \tag{5}$$

Such a weighting is *not* part of standard VSL measurement in use today. This lack affects policy formation with VSL in the following way. The assumption of a uniform p_i in measuring VSL makes one underestimate the risk–wage trade-off of those who have a lower risk exposure than average, but a higher than average risk–wage trade-off. Those agents may have chosen a safe job because they put a very high value on safety, despite their lower-than average risk exposure. Using a uniform risk level to estimate their risk preferences will attribute their choices too much to the assumed risk exposure, while underestimating their particularly high risk preferences. In cases of correlation between risk preferences and risk exposure, using the standard VSL measurement therefore violates

the relevance condition of vicarious governance. To counter this deficiency, VSL should be weighed according to risk exposure.

6 Conclusion

In this paper, I proposed a basic principle (P) for vicarious governance. This principle and its four conditions accuracy, sensitivity, agreed mode of compromise and relevance were meant to capture – in a very simplified way – the often implicit arguments for forming policy with the help of the VSL measure. I employed this principle to appraise the way VSL is currently measured.

The result is that the measurement of VSL is insufficient, given the role it is supposed to play in policy formation: it violates the minimal requirements of vicarious governance. First, construing VSL as a unique and transferable measure violates the sensitivity condition. Second, the practice of transferring VSL between widely differing contexts violates the accuracy condition. Third, the basic measurement procedure makes it difficult to interpret the resulting risk–wage trade-off function; the commonly proposed interpretation of an average trade-off rate violates the sensitivity condition. Fourth, in cases of heterogeneous risk preferences, averaging across different risk preferences in some cases yields results substantially different from a majority voting procedure, and without consent violates the agreed compromise condition. Last, in the case of correlation between risk preferences and risk exposure, the VSL violates the relevance condition.

All these arguments involved specific measurement practices and particular contexts of applications. Hence, my arguments do not rule out the use of VSL as an instrument of vicarious governance in general. Rather, they point to various potential improvements of the measurement of VSL. Absent such improvements, the VSL measure in its current form may still be applicable, albeit in strictly circumscribed areas that avoid the pitfalls of transfer across contexts and of strong heterogeneity. In this way, this paper does not advocate an all-out rejection of VSL, but rather hopes to contribute to its improvement, or at least to its more adequate use.

Notes

1. Currently leading in implementation of this approach is the US federal government. By executive order, all significant US Federal regulations are subject to economic impact analysis. See US Environmental Protection Agency (1993).
2. This principle does not require that beneficiaries have the same preferences, but does require that they value the benefits higher than the costs. This intentionally leaves open the question whether each beneficiary must individually value the benefits higher than the costs, or whether overall benefits must be higher than overall costs. In the following discussion of VSL, I will come back to this issue.
3. In the US, the Office of Management and Budget, responsible for overseeing and coordinating the review of regulatory impact analysis, recommends the use of VSL as part of economic impact analysis. See US Office of Management and Budget (2003, pp. 30–31). This recommendation is not legally binding, nor is any particular measurement procedure or numerical value suggested; yet the practice of using VSL in federal regulatory analysis has become standard. On the international level, the intergovernmental panel on climate change uses VSL in its efforts to cost the effects of global warming (cf. Bruce, Lee, and Haites 1995).
4. Similar arguments can be found for the derivation of the WTP for risk reductions from consumer market data (e.g. housing or cars) or traffic behaviour data.
5. To put a number to this, the relevant risks usually are in the range from 1/10,000 to 1/100,000. See e.g. Viscusi (2004), who discusses fatality risks ranking from 1/100,000 to 45/100,000.

6. For a discussion, see Slovic, Fischoff, and Lichtenstein (1981), Mendeloff and Kaplan (1989). Chilton et al. (2002) measure specific VSL for rail accidents, domestic fires and fires in public places *relative* to VSL for road accidents. Chilton et al. (2006) measure the differences in 'dread' about various causes of death.
7. Of course, this argument does not apply if w_i is the same for all individuals. Indeed, some proponents of VSL make exactly that assumption. But on the national or international level, this assumption is empirically questionable. Further, if that assumption were true, why engage in such a complex measurement procedure at all? Why not just take somebody's w_i – e.g. the policy maker's – and proceed with that?
8. This argument is commonly accepted for international comparisons. Wealth differences between nations make it plausible that these nations also have different VSLs. 'A poor nation would do well to adopt a lower VSL than a wealthier nation; for China or India, it would be disastrous to use a VSL equivalent to that of the United States or Canada. But this point should not be taken to support the ludicrous position that for donor institutions, both public and private, risk reduction in a wealthy nation deserves more attention than equivalent risk reduction in a poor nation' (Sunstein 2004, p. 395). I agree that it would be wrong to dismiss VSL because it depends on wealth levels. Rather, it is a question how VSL measures can be used appropriately; and it should be clear that they must not be used to compare groups with widely divergent wealth. But if not groups, then also not individuals: aggregating the marginal rate of compensation across a population with pronouncedly unequal wealth distribution is wrong by the same argument as comparing poor and wealthy countries.
9. As an anonymous referee suggests, the correlation between preferences and exposures provides some reason to believe that the possibility depicted in Figure 2 should be taken seriously.
10. From contingent evaluation measurement of VSL, Jones-Lee (1989) tries to get a feel for the extent to which ignoring the covariance of m_i and Δp_i may result in an overestimate of the VSL. He concludes that ignoring it overestimates it by at most 20%.

References

Adler, E.A., and Posner, M.D. (2000), 'Implementing Cost-Benefit Analysis When Preferences Are Distorted', *Journal of Legal Studies*, 29(2), 1105–1147.

Bergstrom, T.C. (1982), 'Is a Man's Life Worth More Than His Human Capital?', in *The Value of Life and Safety*, ed. M. Jones-Lee, New York: North-Holland, pp. 3–26.

Broome, J. (1978), 'Trying to Value a Life', *Journal of Public Economics*, 9, 91–100.

——— (1985), 'The Economic Value of Life', *Economica*, 52, 281–294.

Bruce J.P., Lee H., and Haites E.F., (eds.) (1995), *Climate Change 1995: Economic and Social Dimensions of Climate Change*, Cambridge: Cambridge University Press.

Chilton, S., Covey, J., Hopkins, L., Jones-Lee, M.W., Loomes, G., Pidgeon, N., and Spencer, A. (2002), 'Public Perceptions of Risk and Preference-Based Values of Safety', *Journal of Risk and Uncertainty*, 25(3), 211–232.

Chilton, S., Jones-Lee, M., Kiraly, F., Metcalf, H., and Pang, W. (2006), 'Dread Risks', *Journal of Risk and Uncertainty*, 33(3), 165–182.

Cookson, R., and Dorman, P. (2008), 'Lessons From the Literature on Valuing Reductions in Physical Risks', in *Economic Evaluation of Interventions for Occupational Health and Safety: Developing Good Practice*, eds. E. Tompa, A.J. Culver and R. Dolinschi, Oxford: Oxford University Press, pp. 71–92.

Dehez, P., and Dreze, J. (1982), 'State-Dependent Utility: The Demand for Insurance and the Value of Safety', in *The Value of Life and Safety*, ed. M. Jones-Lee, New York: North-Holland, pp. 41–65.

Garen, J. (1988), 'Compensating Wage Differentials and the Endogeneity of Job Riskiness', *Review of Economics and Statistics*, 70(1), 9–16.

Hausman, D. (2005), 'Sympathy Commitment, and Preference', *Economics and Philosophy*, 21, 33–50.

Heinzerling, L. (2002), 'Markets for Arsenic', *Georgetown Law Journal*, 90, 2311–2339.

Jones-Lee, M. (1989), *The Economics of Safety and Physical Risk*, Oxford: Basil Blackwell.

Liu, J.-T., and Hammitt, J.K. (1999), 'Perceived Risk and Value of Workplace Safety in a Developing Country', *Journal of Risk Research*, 2(3), 263–275.

Mendeloff, J., and Kaplan, R.M. (1989), 'Are Large Differences in "Lifesaving" Costs Justified? A Psychometric Study of the Relative Value Placed on Preventing Deaths', *Risk Analysis*, 9(3), 349–363.

Nussbaum, M. (2000), 'The Cost of Tragedy', *Journal of Legal Studies*, 29, 1005–1036.

Schelling, T. (1983), 'The Life You Save May Be Your Own', in *Choices and Consequences*, Cambridge, MA: Harvard University Press, ch. 5.

Sen, A. (1977), 'Rational Fools', *Philosophy and Public Affairs*, 6, 317–344.

Slovic, P.B., Fischoff, B., and Lichtenstein, S. (1981), 'Perceived Risk: Psychological Factors and Social Implications', in *The Assessment and Perception of Risk*, eds. F. Warner and D.H. Slater, London: Proceedings of the Royal Society, Vol. 376, pp. 17–34.

Sobel, D. (1994), 'Full Information Accounts of Well-Being', *Ethics*, 104(4), 784–810.

Sunstein, C. (2004), 'Valuing Life: A Plea for Disaggregation', *Duke Law Journal*, 54, 358–445.

Thaler, R., and Rosen, S. (1975), 'The Value of Saving a Life: Evidence From the Labor Market', in *Household Production and Consumption*, ed. N.E. Terleckyj, New York: Columbia University Press, pp. 265–300.

US Department of Health and Human Services (2003), 'Food Labeling: *Trans* Fatty Acids in Nutrition Labeling, Nutrient Content Claims, and Health Claims', *Federal Register*, 68(133), 41433–41506.

US Environmental Protection Agency (1993), 'Executive Order 12866 of September 30, 1993 – Regulatory Planning and Review', *Federal Register*, 58, 51735, *http://www.epa.gov/fedrgstr/eo/eo12866.htm*.

US Office of Management and Budget (2003), *Regulatory Analysis* (OMB Circular A-4), Washington, DC: The White House, *http://www.whitehouse.gov/omb/circulars/a004/a-4.pdf*.

Viscusi, K. (2000), 'Risk Equity', *Journal of Legal Studies*, 29, 843–871.

——— (2004), 'The Value of Life: Estimates With Risk by Occupation and Industry', *Economic Inquiry*, 42(1), 29–48.

Viscusi, K., and Aldy, J.E. (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.