



On a true value of risk

Kozin, Igor

Published in:
International Journal of General Systems

Link to article, DOI:
[10.1080/03081079.2017.1415896](https://doi.org/10.1080/03081079.2017.1415896)

Publication date:
2018

Document Version
Peer reviewed version

[Link back to DTU Orbit](#)

Citation (APA):
Kozin, I. (2018). On a true value of risk. *International Journal of General Systems*, 47(3), 263-277.
<https://doi.org/10.1080/03081079.2017.1415896>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

On a true value of risk

Igor Kozine

Technical University of Denmark

Produktionstorvet, Building 424, office 118, 2800 Lyngby, Tel. +45 45254548, igko@dtu.dk

The paper suggests looking on probabilistic risk quantities and concepts through the prism of accepting one of the views: whether a true value of risk exists or not. It is argued that discussions until now have been primarily focused on closely related topics that are different from the topic of the current paper. The paper examines operational consequences of adhering to each of the views and contrasts them. It is demonstrated that operational differences on how and what probabilistic measures can be assessed and how they can be interpreted appear tangible. In particular, this concerns prediction intervals, the use of Bayes rule, models of complete ignorance, hierarchical models of uncertainty, assignment of probabilities over possibility space and interpretation of derived probabilistic measures. Behavioural implications of favouring the either view are also briefly described.

Keywords: risk, true value, uncertainty, ignorance

1. INTRODUCTION

The topic of whether a true value of risk exists may seem to have been thoroughly scrutinised and be of only philosophical interest. However, it does have operational implications and in the view of some new developments in statistical reasoning it appears worth returning to it. As it will be seen, the conventional implicit underlying assumption is that a true value exists, which predefines how probabilistic risk quantities are derived and interpreted. Assuming the opposite (a true value does not exist) opens up a different perspective for the derivations and interpretation. The operational consequences of adhering to one of the views are examined and contrasted in the paper. It is informative to understand the differences stemming from choosing either side. It is also important to note that when referring to “a value of risk”, we mean a quantitative risk metric. That is, we

separate the concept of risk from risk metrics. More specifically, the implied in this paper risk metric is a combination of probability and severity of consequences; and, as a particular case, can be thought of as expected disutility. This is in compliance with the glossary of the Society for Risk Analysis accessible on their web page <http://www.sra.org/sra-glossary-draft>.

Answering the question on the existence of a true value of risk is not the same as, for example, answering the question of whether risk exists (see, for example, Aven [2012]) or whether it is a state of the world or a mental construct about the world (see, for example, Rosa [2010]). Neither is it the same as answering the question of whether risk can be measured objectively or only subjectively. The discussions of this type have been basically limited to the following dimensions: { ‘state of the world’ vs. ‘mental construct’; ‘objective’ vs. ‘subjective’; ‘risk as a concept’ vs. ‘risk as a metrics’ }.

When thinking of risk as a mental construct, one should further distinguish between two different cases: (1) the constructs chained either to the laws of nature or to observed events, and (2) those not having in any obvious way causal connections to nature. The concept of risk falls obviously in the first category, as we talk about the risks of events that either have been observed or can in principle be observed in future. In contrast, such abstract objects like numbers, lattices and algebras cannot be observed nor chained to the laws of nature and therefore they fall in the second category. This distinction is important to mention in order to separate true beliefs¹ about some values from the values themselves that can be measured precisely if measuring tools exist. True beliefs are possible through perceptual knowledge; while it is much harder to argue that true beliefs can exist for abstract objects (Hendriks 2006). Even though a true value of some measurable attribute can exist and it can manifest itself perceptually, true beliefs can provide wrongful answers about the true value. For example, by looking at a clock that is not going, and being unaware of this fact, a man

¹ ‘True belief’ of a subject in a proposition simply means that the subject believes that the proposition is true

may genially believe that the time at the clock is correct, while the true value of it is actually different. We can also say that something true can be derived from something false, as ‘even a stopped clock gets the right time twice a day’ (Hendriks 2006).

Something existing in the world has properties or attributes that can be measured and be assumed to have precise values at least at a certain point of time and in a certain context. Temperatures, speeds, pressures - as being properties of real objects - can be believed having true values. Risk is measured and quantified as well. So, one may think this is a sufficient prerequisite to assert that a true value of it exists, despite we may not necessarily be able to measure it precisely at the time being because of limited resources or lacking data; nevertheless, we can do so in principle, ultimately, only in the limit. However, it may be fallacious to make affirmative conclusions about the true value on this ground. An example supporting this is the Heisenberg uncertainty principle, which is the corner stone of quantum mechanics. The principle can be briefly formulated as the following: the properties of a moving item such as position, trajectory or momentum cannot be precisely defined. That is, this two-dimensional attribute – position and momentum – simply does not have a point value in the two-dimensional space; despite the position and momentum are measurable properties. ‘For that reason everything observed is a selection from a plenitude of possibilities and a limitation on what is possible in the future,’ Heisenberg wrote (cited by Furuta [2012]). While the nature of the measured quantity is very different from that of risk, this example makes it legitimate at least to test the consequences of assuming that a true quantitative risk measure does not existence.

It should also be clear that not being able to make a measurement of an attribute of an object, phenomenon or event but only to provide a subjective assessment of it does not necessarily mean that the true value of the attribute, phenomenon or event does not exist. A person can provide a judgement of his/her body temperature when he/she does not feel good or perhaps has a fever.

Nevertheless, this judgement may be wrong and different from that shown by a thermometer.

Though, it is difficult to deny that the true value of the temperature exists.

Besides the pure curiosity of attempting to understand the foundational issues of risk, we want to ultimately understand whether accepting one of the sides – ‘the true value exists’ vs. ‘it does not exist’ – will have operational implications about how risk measures can be assessed and how decision on risk acceptance can be made. As will be seen in this paper, choosing one of the sides affects indeed the way probabilistic risk measures can be assessed and, for example, whether unforeseen events in the possibility set can be framed in the analysis and whether they can influence the computed value of risk.

The following are the principle concepts and accounts underlying the remainder of this paper.

Talking about a true and no-true value of risk means talking about a risk metrics, that is the combination of probability and magnitude/severity of consequences. Specifying the consequences of some future events and making judgements about the uncertainties depend on the assessor. To measure uncertainty we use probabilities and they are subjective (grounded in beliefs), though having support in observable events. Accepting the subjective nature of probability seems justifiable, as risks concern usually rare events and making reliable assessments of the frequencies in the span of human life is often not feasible. In fact, the *principle of direct inference* can moderate and reconcile both the frequentist and subjectivist view. The *principle* says: when you know the values of probabilities assessed as frequencies, you should adopt them as your subjective probabilities (just if the experiment is long enough and considered interchangeable) and act accordingly (Walley 1991). Finally, knowledge about the probabilities can normally be only partial, meaning that it can hardly be feasible to know precisely probability masses over all a set of possibilities.

In this paper the scope is limited to exploring only the probabilities of events while considering the severity of consequences as precisely known quantities. Even with this limited focus, the topic appears rather extensive and involves a relatively large number of probabilistic concepts. It is assumed that the reader is familiar with the most of the articulated concepts or can find more explanation in the referred literature. I have basically been trying to explain the principles in which the concepts are rooted.

The adjectives ‘true’, ‘correct’ and ‘ideal’ are used interchangeably throughout the paper.

This paper unfolds as follows. Section 2 explains the principle difference between the two views and introduces the two approaches. Section 3 and 4 go through the list of operational implications that stem from whether you favour one or the other view on the existence of the true value of risk. Section 5 expands on comparing some merits of operational approaches built on the different views. The concluding Section 6 reflects on behavioural implications of favouring either view.

2. TWO VIEWS ON A TRUE VALUE OF RISK

Identifying a set of possible future events or outcomes of an activity is a first step in any probability quantification problem. In risk analysis this step is called risk identification, the objective of which is to provide an ‘exhaustive’ list of hazards and their possible consequences that may be faced in the future. In practice, the identified risks do not necessarily constitute an exhaustive set of outcomes, as some of them can be considered as having negligible probabilities of occurring and on this ground can be discarded while some other can be simply unknown. In any case, the possibility set is usually regarded closed and the probability masses are distributed over the outcomes adding up to unity or to the value that is practically indistinguishable from unity. If we had an abundance of information about the occurrence of events from a possibility set, then probability masses over the states would be known precisely; and, as a consequence, we would have to acknowledge that we are

in possession of a true value of risk. In practice, only for some mass production identical goods we may achieve the state of knowing rather precisely the values of risks.

In general, the values of risks are not known precisely and the analyst has the option to consider that convergence to a precise value of risk is possible in the limit. That is, the true value exists but due to limited time, resources or other limitations in assessing probabilities it is not known at the time being. Following this prospective, a single probability distribution over a set of possible outcomes can be chosen that is tacitly regarded as a ‘true’ or ‘ideal’ model of uncertainty. After that, computing other probabilistic risk measures of interest becomes a rather easy mathematical exercise. The choice of a single distribution is often supported by experience gained from the performance of similar systems or observed tendencies in system’s behavior. For example, if a product is exploited during the normal operation phase (wear-in period has been overcome), the failure rate can be for some products assumed constant and then, an exponential distribution of time to failure can be arguably accepted. Following this line will result in deriving point-valued risk measures based on which it is easy to make decision on risk acceptance.

In fact, accepting the true-value view does not make the adherent use only a single probability distribution as a model of uncertainty. One can introduce a class of distributions in which a particular member is considered a plausible candidate to be an ideal distribution. Then, a risk quantity of interest can be computed for each distribution-candidate and after that, lower and upper bounds can be constructed as assessments of risk. This is a robust way of compensating for ignorance that studies the sensitivity of derived numerical values to variations in probability distributions. This approach is much more time consuming and requires substantial computational resources.

The topic of the non-existence of a true probability was raised by de Finetti (1974). He writes that “to refer to the ‘true, unknown probabilities’, is a distortion that leads to an immediate confusion of

the issues”. “Speaking of unknown probabilities must be forbidden as meaningless” (de Finetti, 1972). However, this view has not been unanimous and subject to controversy even among those who support it. There are some personal probability distributions about which we feel relatively sure as compared with others. So, many people have found the concept of uncertainty about probability – admitting the existence of a true value – to be useful in practice, see, for example, Savage (1954), Good (1962), Kaplan and Garrick (1981). Mosleh and Bier (1996) make a stronger point opposing to de Finetti: “Various theorists ... have generally failed to distinguish between different sources of uncertainty about probabilities. In particular, ... uncertainty about the underlying events, propositions, or hypotheses on which those probabilities are conditioned; and imprecision resulting from inherent cognitive limitations.” Grounding their arguments on these points, they conclude that it is meaningful and theoretically explainable to consider subjective probability as an uncertain quantity, admitting herein probability distributions over a set of personal probabilities. This admittance guarantees immediately a unique aggregated probability distribution as the ideal representative of the personal probabilities.

The completely different (the no-true-value) view is to consider again a class of probability distributions, but to regard that no member of the class can be a true model of uncertainty. Only the whole class itself - viewed as an indivisible entity - is a reasonable model for ignorance. Each single member of the class is not a reasonable model, because no single distribution can model ignorance satisfactory. This view allows computing only interval-valued measures of probability and risk, the interpretation of which is different from that based on the ideal uncertainty model. This approach opens up other ways of deriving the assessments of the bounds. It was called by Walley (1991) as imprecise statistical reasoning. The precursors of Walley’s theory were mathematical structures with relaxed axioms of probability that formed a foundation for interval-valued probabilities (Good 1962, Smith 1961, Suppes 1974, Kuznetsov 1991). Suppes (1974) advocates

the use of interval-valued probabilities stating about subjective probabilities (beliefs) that they “are rather like the leaves on a tree. They tremble and move under even a minor current of information. Surely we shall never predict in detail all of their subtle and evanescent changes”. One should not seek precision where imprecision is only possible.

The probabilistic concepts and statistical approaches articulated in the rest of the paper are well known, though, they are looked into from the angle of which view on the existence of a true value is accepted. The assumption of an ideal and correct uncertainty model is the usual (mostly implicit) underlying fundamental prerequisite for manipulating with probabilistic quantities. However, there is the opposite view, and in the following I argue that depending on this underlying fundamental assumption about the existence of a true value, the way probabilistic risk measures are derived can vary widely as can the assessment strategies and the interpretation of the results.

3. WHAT IF A TRUE VALUE EXISTS?

Let us assume that the true value of risk exists, i.e. risk can be expressed by a single value, and let us see how this view can influence what probabilistic quantities and concepts can be assessed and articulated, how this can be done and how the results can be interpreted.

Prediction intervals

When assuming that a true value of risk exists and making the assessment of it either on a set of expert judgements or a limited sample of observed failures, the analyst admits that the assessment does not necessarily hit the correct value, but more likely it lies in the vicinity of the true value.

Acknowledging this allows introducing the concept of a prediction interval that is an interval associated with an uncertain value, with a specified probability of this value lying within this interval. It is clear that prediction intervals cannot be invoked if the alternative view (the no-true-value view) is favored.

Bayes' rule

The other remarkable implication of admitting the existence of an ideal distribution as the uncertainty model is to be able to use Bayes' rule to update this distribution if some of the states (from the set of possible outcomes) have been observed. In this case, the initially expected precise distribution is referred to as *prior* while the updated is called *posterior*. Bayes' rule can in the same way be applied to a whole class of the distributions by updating each of them while the relevant observations become available. This approach is called *Bayesian sensitivity or robust analysis*.

Model of complete ignorance

Bayes' updating rule is applicable even though the analyst does not have any prior relevant information at hand about the probabilities of occurrence of events. This can be done by introducing non-informative priors that can be viewed as models of complete ignorance over the set of possible outcomes. The choice of non-informative priors is rooted in Jacob Bernoulli's and Laplace's principle of insufficient reason (also called the principle of indifference) which states that 'if we are ignorant of the ways an event can occur (and therefore have no reason to believe that one way will occur preferentially compared to another), the event will occur equally likely in any way'². If there are $n > 1$ mutually exclusive and collectively exhaustive possibilities, the principle prescribes that each possibility should be assigned a probability equal to $1/n$, which is the simplest non-informative prior.

Again, the model of complete ignorance is firmly underlain by the existence of an ideal probability model and the best candidate at this state of knowledge is a uniform distribution. As knowledge is gained, other candidates may emerge or, alternatively, by applying the updating rule, the non-informative prior will converge to the true model, possibly, in the limit.

Hierarchical uncertainty models

² <http://mathworld.wolfram.com/PrincipleofInsufficientReason.html>

Adherence to the view of an ideal uncertainty model can even have another farther reaching implication. It enables us to employ a useful strategy for constructing a precise prior “candidate” given a set of alternative distributions and uncertainty about which of them is the best candidate. The modeler’s uncertainty about the ideal candidate is then called second-order uncertainty. Models of this type are called hierarchical uncertainty models, the most common of which is the Bayesian one. Here both the first- and the second-order uncertainty are represented by probability measures. The outcome of applying a Bayesian hierarchical model is that it can always be reduced to a first-order model by integrating out the higher order parameters. Whenever a Bayesian hierarchical model is used, it is assumed that the available information is at all levels detailed enough for it to be representable by a precise probability measure (de Cooman 2002). Other hierarchical models have been studied in the literature. See, for example, Gardenfors and Sahlin (1982), Kozine and Utkin (2002), de Cooman and Walley (2002). In any case, the prerequisite for any second-order uncertainty model is the existence of a true (but unknown) first-order uncertainty model.

Interpretation and assessment strategy

Another important implication concerns the interpretation of probabilistic assessments and connected to it assessment strategies.

Subjective probabilities traditionally have a behavioural interpretation (see, for example, de Finetti [1974]). According to this interpretation, probabilities are interpreted in terms of behavior such as betting behavior or other choices between actions. For each event of interest, there is some betting rate that a subject regards as fair, in the sense that he is willing to accept either side of a bet on the event at that rate (Walley 1991). Subject’s fair betting rate at which he or she is willing to bet on this event or against it is considered subject’s personal probability for the event. For example, if a subject is willing to pay a maximum price of 10€ to bet on an event A with a stake 100€, then his fair betting rate for A equals $10/100 = 0.1$. So, again, the underlying assumption is the existence of

an ideal probability of event A . The subject is attempting to provide his best guess for the true probability. Whatever information he has (even complete ignorance), he should be able to bet on or against A at any rates that are more favorable than his fair rate. It should be noted that any precise assessment under complete ignorance seems quite arbitrary.

Assignment of probabilities and possibility space

Accepting the true-value view makes the assignment of the initial probability of an event A dependent on the possibility space in which A is embedded. This may result in very different probability assignments. Let us consider the example with coloured marbles described by Walley (1996). ‘I have a closed bag containing coloured marbles. I intend to shake the bag, to reach into it and to draw out one marble. What is the probability that I will draw a red marble?’ Because there are two possible outcomes (red or non-red), by applying the principle of indifference, the probability of red must be $\frac{1}{2}$. ‘Or one might also list the ‘equipossible’ outcomes as {red, blue, green, yellow, black, white, other colours} before applying the principle of indifference, which yields probability $\frac{1}{7}$ for red’.

This consequence may be seen as of little relevance for practical risk problems, as in most cases we are not completely unaware of possible consequence states we may face in the future. However, the prior probabilities of emerging risks and very rare events may have been judged somewhat arbitrary if one counts on the existence of a true value of risk. For these risks and events the weight of prior assessments in posterior will dominate possibly in a long span of time, meaning that we may base our decisions on risk assessments that are far from being correct.

4. WHAT IF A TRUE VALUE DOES NOT EXIST?

When assuming that a true value of risk does not exist, no particular meaning can be given to an individual risk profile presented as a set of point-valued assessments in the space “probability-consequence”. As we consider the severity of consequences known precisely, then this statement

can be narrowed down and directed to a probability distribution that is not considered a reasonable model of uncertainty. Nevertheless, a class of ‘admissible’ distributions can be constructed anyway. In this class, as worded by Walley et. al (1996), we ‘do not regard the individual prior distributions as plausible or reasonable models for prior uncertainty’ and we ‘deny that any precise probability model is reasonable’. A class of prior probability distributions under some conditions is ‘not a class of reasonable priors, but a reasonable class of priors’ Walley (1996). This means that each single member of the class is not a reasonable model for ignorance, because no single distribution can model ignorance satisfactory. But the whole class is a reasonable model for ignorance. When we have little information, the upper probability of a non-trivial event should be close to one and the lower probability should be close to zero. This means that the prior probability of the event may be arbitrary from 0 to 1.

Construction of a class of admissible uncertainty models

The immediate question is how to construct such class of distributions that is not simply a collection of individual plausible candidates for an ideal uncertainty model? Walley (1991) and Kuznetsov (1991) suggested a theory based on the idea of reducing the set of all probability distributions to a set that matches partial information the analyst has at hand. It is done by solving a properly stated linear optimization problem that was called by Walley (1991) *natural extension*. To capture a wider variety of partial relevant statistical knowledge for more adequate construction of the set of admissible distributions, non-linear programs have also been suggested by Kozine and Krymsky (2009a, 2012).

By partial statistical information with regard to a specific set of possible outcomes it is meant not knowing precisely probability masses distributed over all possible events in the possibility space. Some knowledge (precise or imprecise) about probability measures (not only probability masses), which theoretically can be derived from known probability distributions (for example, mean times

to failure or mean failure rates), is also considered partial information about the probabilities of possible events.

Here I demonstrate how a set of probability distributions can be constructed from partial statistical knowledge the subject has at hand, including judgements on lower and upper probabilities; and, in the reverse order, how lower and upper probabilities can be derived from the knowledge of a permissible set of distribution, in which no one candidate is regarded a correct uncertainty model.

Assume there are three possible outcomes s_1 , s_2 and s_3 of an activity. They constitute an exhaustive set of events meaning that $P(s_1) + P(s_2) + P(s_3) = 1$, where $P(\cdot)$ stands for a probability.

Assume then that a subject has provided an upper and lower bound for s_1 . Say, the lower bound is 0.1, while the upper is 0.3. Formally, we can write $P(s_1) \in [0.1; 0.3]$. To see graphically the area these judgements cut off in the set of all probability distributions, the simplex representation can be used (Figure 1). (The interpretation of the bounds and the way of how they can be elicited is explained below in this section.)




Figure 1. Presentation of the elicited bounds 0.1 and 0.3 for probability $P(s_1)$ on the simplex

In Figure 1 the vertexes 1, 2 and 3 correspond to the three states s_1 , s_2 and s_3 . The probability simplex is an equilateral triangle with a height of one unit, in which the probabilities assigned to the three elements are identified with perpendicular distances from the three sides of the triangle.

Adding up these three distances gives 1. Thus, each point inside of the simplex can be thought of as a precise probability distribution. Thus the elicited judgement $P(s_1) \in [0.1; 0.3]$ presented in the simplex cuts off the shaded area that is to be interpreted as the set of admissible probability distributions matching the partial statistical knowledge provided in the form of the probability interval.

Let us further assume that some more pieces of statistical evidence are available: (1) s_2 is at least two times as probable as s_3 , and (2) s_2 and s_3 is at least as probable as s_1 . Formally, these statements can be written correspondingly as follows: $P(s_2) \geq 2P(s_3)$, and $P(s_2) + P(s_3) \geq P(s_1)$. These two inequalities along with $P(s_1) \in [0.1; 0.3]$ and the condition $P(s_1) + P(s_2) + P(s_3) = 1$ define the constrained area which is shown in black in Figure 2.




Figure 2. Model produced by the additional judgements

The calculation of the upper and lower bounds for the probabilities of interest becomes a geometric or a linear programming task. The calculated values of the probabilities are $\underline{P}(s_2) = 0.466$, $\overline{P}(s_2) = 0.9$, $\underline{P}(s_3) = 0$, $\overline{P}(s_3) = 0.3$ while $\underline{P}(s_1) = 0.1$ and $\overline{P}(s_1) = 0.3$ remain unchanged. In the above two example, any underlying assumption on the existence of a true probability distribution is unnecessary. The provided partial statistical knowledge determines the admissible set

of probability distributions that as a whole is the uncertainty model. The larger the cut-off area, the larger epistemic uncertainty about the distribution of probability masses over the possible outcomes. The constrained area in the set of all probability distributions defines the lower and upper bounds of the probabilities of interest. The width of the probability intervals depends on the size of the set of admissible distributions and manifests the degree of epistemic uncertainty.

Prediction intervals

In contrast to the true-value view, prediction intervals can no longer be constructed for the probabilities, as they do not have any meaning because there is no a true value that can be confidentially covered by those intervals. Hence, the probability intervals as they have been derived above in the examples must have a completely different interpretation.

Interpretation and assessment strategy

For a behavioural interpretation of the interval-valued probabilities the fair betting rate cannot be invoked any longer, as it is simply a precise individual probability that is not regarded as a reasonable model for uncertainty under the assumption of the non-existence of a true value. Turning back to the above example, the lower and upper probabilities, $\underline{P}(s_1) = 0.1$ and $\overline{P}(s_1) = 0.3$, can be interpreted and elicited as follows.

A subject expresses the maximum rate, 0.1, at which he is prepared to bet on s_1 , and the minimum rate, 0.3, at which he is prepared to bet against s_1 . These rates are interpreted differently from the fair rates: 0.1 is the maximum price the subject is willing to pay to participate in the gamble paying off 1 unit of price if s_1 occurs and nothing otherwise; 0.3 is the minimum price for which the subject is willing to sell this gamble.

The ability to directly provide subjective lower and upper probabilities is an important operational consequence of accepting the no-true-value view.

Another example of how lower and upper probabilities can be elicited is that of the willingness of a subject to pay for reducing the chance of one mortality and the willingness to accept a premium for exposing himself to an action possibly resulting in mortality. Given the value of a human life is known, it can be divided by the willingness to pay and accept providing a lower and upper probability of a person dying due to some activity. This type of approaches are usually used to assess the value of a statistical life given the probability of people dying is known. There is a variety of those and for details the reader is referred to Hultkrantz and Svensson (2012).

Up to date and to my best knowledge, the only available interpretation of the bounds is behavioural. Nevertheless, by looking into the mechanism of deriving the bounds for probability measures, the upper bound is the supremum or maximum value of the probability measure obtained from applying different probability distributions from the set of admissible distributions. The lower bound is similarly the infimum or minimum. In different words, among all admissible distributions there is one that delivers the maximum value to the measure and there is another one that delivers its minimum value. This way of interpretation is in fact indistinguishable from that of robust Bayesian analysis.

Hierarchical uncertainty models

As far as hierarchical models of uncertainty are concerned, Kozine and Utkin (2002) proposed a very general approach that does not require the assumption of either an ideal first-order or second-order uncertainty model. The idea was to use an interval-valued second-order uncertainty model that can be interpreted as a reliability of an acquired statistical piece of evidence. For instance, if we are provided with the statistical statement that s_2 is at least two times as probable as s_3 (see the above example), then, we may not necessarily consider this evidence as 100% reliable. It is assumed that our unreliability cannot be expressed precisely but in the form of a probability interval, $[\underline{\gamma}, \bar{\gamma}]$. Now the hierarchical model appears as following: $P[P(s_2) \geq 2P(s_3)] \in [\underline{\gamma}, \bar{\gamma}]$. A

number of special cases of this model have been studied and resulted in deriving a set of rules for constructing first-order upper and lower bounds on probabilistic quantities (Kozine and Utkin, 2002).

Generalised Bayes rule

Having a mechanism for revising statistical evidence as some events in the possibility space have been observed is of extreme importance for statistical reasoning. As individual probability distributions are not endowed with any meaning under the no-true-value view, Bayes updating rule is not applicable here. Nevertheless, what is possible to do is to update directly upper and lower bounds of probability measures or, which is technically the same, to update extreme points of the region cut off in the set of all probability distributions based on statistical evidence. For example, the black tetragon in Figure 2 defines the set of probability distributions matching the provided statistical information. As soon as one out of the three events s_1 , s_2 and s_3 has been observed, the so-called generalized Bayes rule developed by Walley (1991) can be applied. It updates the vertices of the tetragon and, by doing so, the rule produces a new updated region.

Another approach suggests using either the beta-Bernoulli (Walley 1991) or the imprecise Dirichlet model (Walley 1996) for updating directly upper and lower probability bounds. Technically, the use of these two models is rather simple and this is why they have been applied in different domains (see for example, Abellan [2006], Utkin [2004], Utkin [2006], Utkin and Kozine [2010])

Assignment of probabilities and possibility space

Accepting the no-true-value view admits a different type of procedure for assigning initial probabilities over a possibility space about which we have very limited knowledge. The assignment of the initial probability of an event A can be made independent of the possibility space in which A is embedded. Let us consider the example about the coloured marbles from the previous section. If we are completely ignorant about the content of the bag, ‘it is quite plausible that all the marbles in

the bag are red, or that none of them red, or that the proportion of red marbles is anywhere between 0 and 1. That suggests the *vacuous probability* model: take the lower probability of red to be 0 and the upper probability to be 1. [...] If you have no information about the chance of drawing a red marble, you have no compelling reason to bet on or against red, whatever the odds. [...] The upper and lower probabilities assigned to red should not depend on whether the possibility space is {red, non-red}, or {red, blue, green, other colours} ...' (Walley 1996).

The vacuous probability is a model of complete ignorance under the assumption of no-true-value that is the counterpart of non-informative priors under the opposite assumption. It is clear that the vacuous probability is a trivial model that cannot generate useful inferences. However, as a starting point for a learning process it is completely reliable and will not change if unexpected events are witnessed and the possibility space becomes populated with more events.

5. SOME COMPARATIVE NOTES

When contrasting the operational implications of choosing either the true-value view or no-true-value, the advantages of one concept or approach over the other is far from obvious. The modeler has to simply choose the view he supports. In some cases, though, computational complexity can become a barrier hampering the use of a mathematical tool. Complexity is often attributed to the tools based on the no-true-value view. In some other cases, the underlying conceptual and mathematical foundation can be regarded as inadequate to produce a satisfactory result. For example, providing too wide intervals for probabilistic measures is often practically useless, and introducing some arbitrariness by accepting simplifying assumptions may appear justifiable to be determinate when taking decision. In this regard, the true-value view is definitely advantageous over the no-true-value view, as it has in its toolkit rather simple mathematical mechanisms to narrow down epistemic uncertainty up to precise probabilities. The opponents, though, can consider this way of tackling uncertainty as unjustifiably reckless.

As it was mentioned above, hierarchical uncertainty models are rather common in statistical reasoning whenever there is an ideal first-order probability distribution but the modeler is uncertain what it is. The uncertainty of the modeler is then called second-order uncertainty. These models are convenient devices for reducing the problem to a single probability distribution (first-order uncertainty) that is considered the best compounded candidate for the ideal distribution over a possibility set. The common criticism of this approach concerns the two points: (1) deriving a precise second-order uncertainty model seems unrealistic and, in practice, there is considerable arbitrariness in the choice of it; and (2) even though the second-order uncertainty is somehow derived, it is unclear how to interpret it, and in particular how to give it a particular interpretation in terms of behavior (de Cooman and Walley, 2002).

To address these issues possibilistic hierarchical models have been developed, which represent second-order uncertainty by a possibility model (Walley 1997; de Cooman and Walley 2002) and the related models of Gardenfors and Sahlin (1982), Levi (1980) and Nau (1992). They relax the constraint of having a precise second-order uncertainty model and suggest alternative ways of deriving it, for example, by constructing a possibility distribution. These are a kind of hybrid models accepting an ideal (but not known) first-order uncertainty model, though not accepting this assumption for the second-order uncertainty. The outcome of applying the hybrid hierarchical models will be first-order lower and upper bounds on the probability measures. Although, these modifications do not require a precise second-order uncertainty model, the issue of interpretation of both the second-order uncertainty and derived interval-valued probability measures stays unresolved.

Mathematically, it is possible to define a very general hierarchical model that does not require the assumption of either an ideal first-order or second-order uncertainty model (Kozine and Utkin, 2002). However, such model, except for some special cases, is rather complex to apply than the

Bayesian and hybrid hierarchical model, and at present we do not have a behavioural foundation for it, nor a clear understanding of how to elicit the interval-valued second-order probabilities.

The no-true-value view on risk has not found strong support among risk and reliability practitioners.

The reasons for it in my opinion are: rather complex models and often too imprecise derived probabilistic measures that make them impracticable for decision making. Root causes of high imprecision in interval-valued probabilistic quantities are illuminated in Kozine and Krymsky (2009b), from where the reader can also judge on the complexity of the mathematical derivations.

6. CONCLUDING NOTES

The operational implications of adhering to the different views on the existence of the true value of risk described in this paper are not exhaustive. Some other that have not been discussed but known to the author are the case of multiple agents and the influence on decision making. The case of multiple agents is the topic tightly linked to the combination rules of statistical evidence obtained from different agents or, more generally, from different information sources. The combination rules are drastically different depending on the view the modeler favours. As far as the influence on decision making is concerned, the consequences of favouring either view are still unclear.

In conclusion I want to mention few other points that do not have any operational implication of favouring either view discussed in the paper; however, they may have behavioural implications influencing the reliable operation of complex systems.

The topic of whether a true value of risk exists is motivated by acknowledging epistemic uncertainty in our risk-related knowledge. Sound scepticism is a necessary trait of an epistemologist as well as of a risk analyst. As ‘error is a starting point of scepticism’ (Hendriks 2006), failure is starting point of high reliability organisations. And continuing this chain, I can add that uncertainty is a starting point of reliable decision. Preoccupation with uncertainty – like preoccupation with failure in high reliability organisations – may contribute to making a positive difference in high risk

occupations and environments. It may contribute to the avoidance of wrongful decisions and, as a consequence, to lessening the risk of unwanted events.

A similar point of view is promoted by Sahlin (2012). His stand is that in the presence of epistemic risk, possibly complemented by value uncertainty, we need to take a Socratic approach to risk analysis and risk management. To explain what it means he cites Plato's *Apology*: 'Probably neither of us knows anything really worth knowing; but whereas this man imagines he knows, without really knowing, I, knowing nothing, do not even suppose I know. On this one point, at any rate, I appear to be a little wiser than he, because I do not even think I know things about which I know nothing.' Except for the call to pay proper attention to epistemic uncertainty when it comes to risk assessment and management, Sahlin argues that 'a Socratic approach allows for rational decision making.' On the opposite, 'we cannot be rational if we do not take our complete state of knowledge into account.'

Mentioning the theory of *epistemic reliabilism* (Goldman 1979) seems to me relevant too. As it is considered that probabilities in relation to risk should be regarded as subjective quantities grounded in beliefs; beliefs in turn cannot always be qualified as knowledge, that is, they can fail to output the truth. Nevertheless, we should acknowledge that some belief-generating and belief-representing methods may appear more reliable than some other. Reliability in this context "consists in tendency to produce beliefs that are rather true than false" (Goldman 1979). More reliable methods produce more often true statistical statements, meaning that the frequency with which the beliefs turn up true is greater than the frequency of false. The frequency can be substituted by propensity to produce more reliable judgements. 'A method is reliable as it has the propensity to produce true beliefs' (Hendriks 2006).

With regard to the reliability and propensity my belief is that the no-true-value view is more reliable than its counterpart because it has a higher propensity to produce true statements. I ground it on the

following hypothesis that appears to me a sensible statement: A subject does not have any disposition to favour an option if it is not a feasible option among alternatives.

For the no-true-value view, there is not a feasible option to reduce an uncertainty interval except for acquiring more relevant knowledge; while for the true-value view, there are feasible options to reduce the interval up to a precise value by simply employing some extant assessment strategy or computational mechanism. An assessment strategy may not necessarily be substantiated by relevant knowledge. Neither is the computational mechanism necessarily adequate. They might simply be used because these options exist and one may be tempted to reduce uncertainty in this mechanistic way which is in general arbitrary. As arbitrariness is the counterpart to reliability when making choice, the true-value view has the propensity to produce less reliable decisions compared to the no-true-value view.

It is appealing to be able to derive a precise (point-valued) risk measure of interest. Though, it should not be done at the cost of lower reliability of statements. Or, perhaps, a proper balance should be sought. The more precisely we are trying to distribute probability masses over future (complex) events, the more unreliable our derived conclusions tend to be. This is like the Heisenberg principle of uncertainty: ‘... the more closely we try to pin down some location for the electron on its way toward the detection screen, the more it ... will be likely to miss that target.’³

References

- Abellan, J. (2006). Uncertainty measures on probability intervals from the imprecise Dirichlet model. *International Journal of General Systems*, 35(5), 509–528.
- Aven, T. (2012). Probability does not exist, but does risk exist? *ESRA Newsletter*, pp. 2–3.
- Aven, T., & Renn, O. (2009). On risk defined as an event where the outcome is uncertain. *Journal of Risk Research*, 12(1), 1–11. <http://doi.org/10.1080/13669870802488883>

³ http://simple.wikipedia.org/wiki/Uncertainty_principle

- De Cooman, G. (2002). Precision–imprecision equivalence in a broad class of imprecise hierarchical uncertainty models. *Journal of Statistical Planning and Inference*, 105, 175–198.
- De Cooman, G., & Walley, P. (2002). A possibilistic hierarchical model for behaviour under uncertainty. *Theory and Decision*, 52, 327–374.
- De Finetti, B. (1972). *Probability, Induction and Statistics: The Art of Guessing*. New York: Wiley.
- De Finetti, B. (1974). *Theory of Probability*. New York: John Wiley and Sons.
- Furuta, A. (2012). *One Thing Is Certain: Heisenberg’s Uncertainty Principle Is Not Dead*. Scientific American.
- Gardenfors, P., & Sahlin, N.-E. (1982). Unreliable probabilities, risk taking, and decision making. *Synthese*, 53, 361–386.
- Goldman, A. (1979). What is Justified Belief? In G. S. Pappas (Ed.), *Justification and Knowledge* (pp. 1–23). Dordrecht: D. Reidel Publishing Company.
- Good, I.J. (1962) Subjective probability as the measure of a non-measurable set. In *Logic, Methodology and Philosophy of Science*. Eds. E. Nagel, P. Suppes, and A. Tarski. Stanford: Stanford University Press.
- Hendriks, V. F. (2006). *Mainstream and Formal Epistemology*. New York: Press, Cambridge University.
- Hultkrantz, L., & Svensson, M. (2012). The value of a statistical life in Sweden: A review of the empirical literature. *Health Policy*, 108(2-3), 302–310. Retrieved from [dx.doi.org/10.1016/j.healthpol.2012.09.007](https://doi.org/10.1016/j.healthpol.2012.09.007)
- Kaplan, S. and Garrick B.J. (1981). On the Quantitative Definition of Risk. *Risk Analysis*, Vol. 1, No. 1, p. 11-27
- Kozine, I., & Utkin, L. V. (2002). Processing Unreliable Judgements with an Imprecise Hierarchical Model. *Risk Decision and Policy*, 7(3), 325–339.
- Kozine, I., & Krymsky, V. (2009a). Bounded Densities and Their Derivatives: Extension to other Domains. In P. Coolen-Schrijner, F. Coolen, M. Troffaes, T. Augustin, & S. Gupta (Eds.), *Imprecision in Statistical Theory and Practice* (pp. 29–42). Greensboro: Grace Scientific Publishing. <http://doi.org/10.1080/15598608.2009.10411909>
- Kozine, I., & Krymsky, V. (2009b). Computing interval-valued statistical characteristics: what is the stumbling block for reliability applications? *International Journal of General Systems*, 38(5), 547–565. <http://doi.org/10.1080/03081070802187798>

- Kozine, I., & Krymsky, V. (2012). An interval-valued reliability model with bounded failure rates. *International Journal of General Systems*, 41(8), 760–773.
<http://doi.org/10.1080/03081079.2012.721201>
- Kuznetsov, V. (1991). *Interval statistical models*. Moscow: Radio and Sviaz (in Russian).
- Levi, I. (1980). *The Enterprise of Knowledge*. London: MIT Press.
- Mosleh, A. and Bier, V.M. (1996). Uncertainty About Probability: A Reconciliation with the Subjective Viewpoint. *IEEE Transactions on Systems, Man, and Cybernetics – Part A: Systems and Humans*, Vol. 26, No. 3, p. 303-310.
- Nau, R. F. (1992). Indeterminate probabilities on finite sets. *Annals of Statistics*, 20, 1737–1767.
- Rosa, E. a. (2010). The logical status of risk – to burnish or to dull. *Journal of Risk Research*, 13(3), 239–253. <http://doi.org/10.1080/13669870903484351>
- Sahlin, N.-E. (2012). Unreliable Probabilities, Paradoxes, and Epistemic Risks. In *Handbook of Risk Theory* (pp. 477–498). Dordrecht: Springer Netherlands. http://doi.org/10.1007/978-94-007-1433-5_18
- Savage, L. J. (1954). *The Foundations of Statistics*. New York: Wiley.
- Smith, C.A.B. (1961). Consistency in statistical inference and decision. *J. Royal Stat. Soc. Ser. B*, vol. 23, pp. 1-37.
- Suppes, P. (1974). The measurement of belief. *J. Royal Stat. Soc. Ser. B*, vol. 36, pp. 160-175.
- Utkin, L. (2004). Probabilities of judgments provided by unknown experts by using the imprecise Dirichlet model. *Risk, Decision and Policy*, 9(4), 391–4000.
- Utkin, L. (2006). Cautious reliability analysis of multi-state and continuum-state systems based on the imprecise Dirichlet model. *International Journal of Reliability, Quality and Safety Engineering*, 13(5), 433–453.
- Utkin, L., & Kozine, I. (2010). On new cautious structural reliability models in the framework of imprecise probabilities. *Structural Safety*, 32(6), 411–416.
doi:10.1016/j.strusafe.2010.08.004
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall.
- Walley, P. (1996). Inferences from multinomial data: Learning about a bag of marbles. *Journal of the Royal Statistical Society*, 58(1), 3–57.
- Walley, P., Gurrin, L., & Burton, P. (1996). Analysis of clinical data using imprecise prior probabilities. *The Statistician*, 45(4), 457–485.