Expected Utility from Multinomial Second-order Probability Distributions

David Sundgren

Abstract—**We consider the problem of maximizing expected utility when utilities and probabilities are given by discrete probability distributions so that expected utility is a discrete stochastic variable. As for discrete second-order distributions, that is probability distributions where the variables are themselves probabilities, the multinomial family is a reasonable choice at least if first-order probabilities are interpreted as relative frequencies. We suggest a decision rule that reflects the uncertainty present in distribution-based probabilities and utilities and we show an example of this rule in action with multinomial second-order distributions.**

Index Terms—**Imprecise probability. second-order probability, discrete probability distributions, multinomial distributions, expected utilty.**

I. Introduction

When computing the expected utility of a decision alternative it may not always be possible to give precise values for the utilities and probabilities of the possible outcomes. The model for imprecise probabilities used here is discrete second-order probability distributions. The term second-order probability comes from the notion that the distributions express the probability that a first-order probability has a certain value. As motivation for discrete second-order distributions we may consider updating, the form of available information and computation of expected utility.

In a continuous second-order setting, a lower bound of a probability can rarely if ever be the result of an observation. But after seeing a three-eyed dog in a kennel of ten, I know that at least one out of ten dogs in that kennel has three eyes. Outside the kennel, I cannot, based on the observation, say much more than that the probability of coming across a three-eyed dog is non-zero. Thus, in situations where the available data has the form of relative frequencies, discrete rather than continuous second-order distributions would be a suitable choice for describing imprecise probabilities. In the case of subjective probabilities, one might as well use discrete distributions unless one has the need to express the probability of particular irrational probability values. For instance that I believe that the probability of seeing another three-eyed dog in my life-time is at least $1/\pi$. That is, discrete second-order distributions are suitable for both objective and subjective probabilities while continuous distributions come into its own in subjective settings. As for computation of distribution expected utility, the fact that there are a finite number of points in a discrete distribution makes a direct computation possible. In the continuous case simulations are necessary.

There is a rich literature on imprecise probabilities, see e.g. [1], [2], [3], [4], [5]. Here second-order probabilities, see e.g. [6], [7], [8], [9], in general and discrete second-order distributions in particular are used and advocated as opposed to interval based models. The standard interval based approach to imprecise probabilities is to employ sets of probability measures, also called credal sets. A credal set is informally a set of probability distributions. Such a set is usually restricted by lower (or, equivalently, upper) bounds of probabilities, and the demand that the set is convex. The intuition appears to be that instead of choosing precise probabilities one uses the set of all probability distributions that are consistent with the beliefs of an agent. Theories of this kind include Choquet capacities [10], lower probabilities [11] and lower previsions [12]. These theories are accessibly summed up in [13]. Without going into the details of these advanced theories one might with an extreme simplification say that in the traditional models of imprecise probabilities the decision maker's or expert's knowledge is represented by intervals between the lowest and highest possible values of the probabilities.

But given some form of representation of imprecise probabilities and utilities it is often not evident what decision rule to employ. With interval based models of imprecise probabilities there are, if we choose to maximize a utility rather than minimize a cost, the rules of Γ-maximax [14], Γ-maximin [15], E-admissibility [16], [17], maximality [12] and interval dominance. [18], see [19] for a comparison of these decision rules. These rules have in common that they single out one or several decision alternatives as optimal without qualifying the ranking with regard to the uncertainty inherent in imprecise probabilities.

In contrast, with utilities and probabilities that are expressed by belief or probability distributions, second-order probabilities, one can measure the imprecision or uncertainty with e.g. variance. With this in view it would be reasonable to use a decision rule that reflects the amount of uncertainty. For example, if the probabilities for all possible utility and (first-order) probability values are known, one can compute the probability that one alternative gives higher expected utility than another alternative.

With continuous second-order probability distributions it seems to be difficult to find closed expressions for expected utility, see [20]. In practice, simulations would have to be made. Alternatively, a decision maker could use discrete distributions. Albeit closed expressions for expected utility would still be hard to find, the expected utility values can be easily computed when second-order distributions are given. Here multinomial distributions are used, but the point is the use of discrete second-order distributions. Discrete distributions offer an environment for updating that is hard to conceive of with continuous second-order distributions and at least brute-force computation of distribution of expected utility is more straight-forward. The multinomial distribution family is used here as an example because of its simplicity, future research will undoubtedly reveal distributions with more attractive properties.

II. Multinomial Second-order Distributions

Given that the variables of a second-order distribution are themselves probabilities, there is a normalization constraint in that probabilities must sum to one. For this reason it is hard for a decision maker to construct a second-order distribution by looking at the possible outcomes separately, and to consider all outcomes simultaneously might be untenable.

Manuscript received June 8, 2010. Manuscript accepted for publication July 25, 2010.

D. Sundgren is with the University of Gävle, Sweden (dsn@hig.se).

But the decision problem itself might imply principles that restrict the choice of second-order distributions. E.g. if it is possible to look at the probabilities of the possible outcomes as relative frequencies, and if the knowledge that restricts the lower bounds of the probabilities come from observations, multinomial distributions are natural candidates, as we shall see.

Consider an experiment with $N = \sum_{i=1}^{n} k_i$ objects of n different types. Assume that the probability to pick an object of type i is $1/n$. Since there are $\frac{N!}{\prod_{i=1}^{n} k_i!}$ permutations of the N objects, if there are k_i objects of type i, the probability that there k_i objects of type $i, i = 1, \ldots, n$ is

$$\Pr(k_1, k_2, \ldots, k_n) = \frac{N!}{n^N \prod_{i=1}^{n} k_i!}$$

The marginal distribution for a single number of objects is

$$\Pr(k) = \sum_{k_i \neq i} \Pr(k_1, k_2, \ldots, k_n) = \frac{(n-1)^{N-k}}{n^N} \binom{N}{k},$$

which is also the probability that there are k objects of a certain type among a total of N objects.

If we gain information by looking at a few of the objects and observe a_i objects of type i, the uncertainty is reduced to the remaining $N - \sum_{i=1}^{n} a_i$ objects of unknown type. So the question is how many more than the observed a_i objects there are of type i. The updated probability should be

$$\frac{(N - \sum_{i=1}^{n} a_i)!}{n^{N - \sum_{i=1}^{n} a_i} \prod_{i=1}^{n} (k_i - a_i)!}$$

But updating the prior $\frac{N!}{n^N \prod_{i=1}^{n} k_i!}$ with the hypergeometric likelihood

$$\Pr(a_i | k_i) = \frac{\prod_{i=1}^{n} \binom{k_i}{a_i}}{\binom{N}{\sum_{i=1}^{n} a_i}}$$

gives just the posterior distribution $\frac{(N - \sum_{i=1}^{n} a_i)!}{n^{N - \sum_{i=1}^{n} a_i} \prod_{i=1}^{n} (k_i - a_i)!}$ suggested above. The likelihood $\Pr(a_i | k_i)$ is the probability that one can see a_i things among the $\sum_{i=1}^{n} a_i$ observed given that there are in total k_i objects of type i.

What possibly is new here is that since we consider the frequencies k_i/N as probabilities, the multinomial distributions described here are examples of discrete second-order distributions, that is, describing the probabilities of different probability values. The advantages of discrete second-order distributions would include Bayesian updating as described above.

We must note, though, that the choice $1/n$ of underlying probabilities for the n types of objects is rather arbitrary and even questionable. For instance, if a relatively large number a_i objects of type i have been observed, it might be that the probability is larger than $1/n$. Work on more sound, alternatively less arbitrary, discrete second-order distributions is under way. For the purpose of this paper, though the multinomial distributions here are sufficient. We wish to show that known discrete probability distributions can serve as second-order distributions, to see how updating might work and suggest a decision rule based on discrete second-order probabilities. The particular distribution family used here is just an example, and the decision rule suggested here and the associated algorithms do not depend on any particular form of discrete second-order probability distribution.

TABLE I
UTILITY DISTRIBUTIONS FOR ALTERNATIVE A

Outcome	0	1	2	3	4	5	6	7	8
1	.3	.3	.2	.2	0	0	0	0	0
2	0	0	0	.3	.4	.3	0	0	0
3	0	.1	.2	.2	.2	.2	.1	0	0
4	0	0	0	0	.1	.4	.3	.2	0

We conclude this section with a remark on upper bounds of probabilities. the lower bounds are given by a_i/N, meaning that a_i objects have been observed. But then there can be at most $N - \sum_{j \neq i} a_j$ items of type i, so the upper bound of each probability is given by the lower bounds of the other probabilities.

For instance, let $N = 8$ and $n = 4$. Further, let the lower bounds be $p_1 \geq 0, p_2 \geq 1/8, p_3 \geq 3/8$ and $p_4 \geq 0$. Then we know that $p_1 \leq 1 - 1/8 - 3/8 = 1/2, p_2 \leq 1 - 3/8 = 5/8, p_3 \leq 1 - 1/8 = 7/8$ and $p_4 \leq 1 - 1/8 - 3/8 = 1/2$.

The corresponding multinomial second-order distribution is

$$\Pr(k_1/8, k_2/8, k_3/8) = \frac{4!}{k_1!(k_2-1)!(k_3-3)!(8-k_1-k_2-k_3)!4^4}$$

Distribution of expected utility is a crucial part in the decision rule suggested here, and in the example below the distribution will be computed in a few cases. If second-order probabilities are given by the multivariate distribution $p(k_1, k_2, \ldots, k_n)$ where $\sum_{i=1}^{n} k_i = N$ (i.e. the first-order probabilities are the ratios k_i/N), and we have n independent utilities given by distributions $p_i(u_i)$, the probability mass function of expected utility is

$$h(z) = \sum_{\sum_{i=1}^{n} k_i u_i = z} p(k_1, k, \ldots, k_n) p_i(u_i)$$

For mathematical reasons concerning primality and co-primality it is hard to make general claims as to which values of k_i and u_i that are involved for every expected utility value z. Thus a brute-force technique will be employed in this paper.

III. A Decision Problem with Four Alternatives

Let us look at the four constructed decision alternatives A, B, C or D. Below we give the distributions for the utilities, and the lower bounds for probabilities that serve as parameters for our multinomial distributions. The probability and utility values are arbitrarily chosen for the sake of the example, they are not derived from any observations or any assessments of real situations. On the other hand, the model employed in the example can be used regardless of how data is collected, be it from observations or by subjective judgments.

A. Alternative A

In the first alternative there are four possible outcomes. The utilities of these are given by the distributions in table I.

The probabilities have lower bounds $p_1 \geq 0, p_2 \geq 1/8, p_3 \geq 3/8$ and $p_4 \geq 0$, so the multinomial second-order distribution is

$$\Pr(k_1/8, k_2/8, k_3/8) = \frac{4!}{k_1!(k_2-1)!(k_3-3)!(8-k_1-k_2-k_3)!4^4},$$

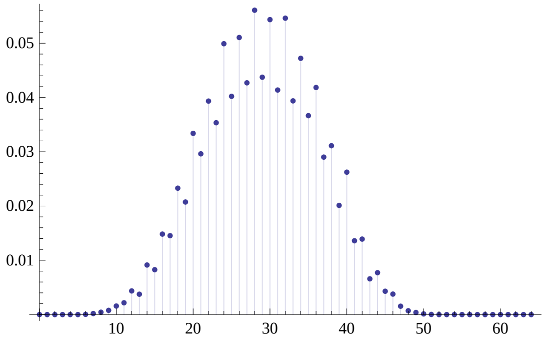

Fig. 1. Expected utility of alternative A

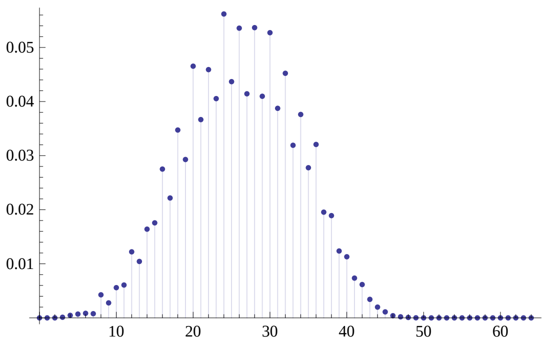

Fig. 2. Expected utility of alternative B

as in the example above at the end of Section II. These distributions for utilities and probabilities give the expected utility distributed as in Figure 1 below.

B. Alternative B

For alternative B the utilities are the same as for alternative A, the only difference between alternatives A and B is that the probability of the first outcome is higher at the expense of the probability of the second outcome; $p_1 \geq 1/8, p_2 \geq 0, p_3 \geq 3/8, p_4 \geq 0$ and the second-order probability distributions is

$$\Pr(k_1/8, k_2/8, k_3/8) = \frac{4!}{(k_1-1)!k_2!(k_3-3)!(8-k_1-k_{-}k_3)!4^4}$$

The resulting distribution of expected utility is plotted in Figure 2.

C. Alternative C

Alternative C is distinguished in that one of the possible outcomes has two possible sub-outcomes with probabilities $p_{11} \geq 1/4, p_{12} \geq 0$. The utilities of the sub-outcomes are found in Table II.

Outcome 1 then has a distribution of expected utility as plotted in Figure 3.

TABLE II
UTILITY DISTRIBUTIONS FOR OUTCOME 1 OF ALTERNATIVE C

Sub-outcome	0	1	2	3	4	5	6	7	8
1	.2	.3	.3	.2	0	0	0	0	0
2	0	0	.2	.4	.2	.2	0	0	0

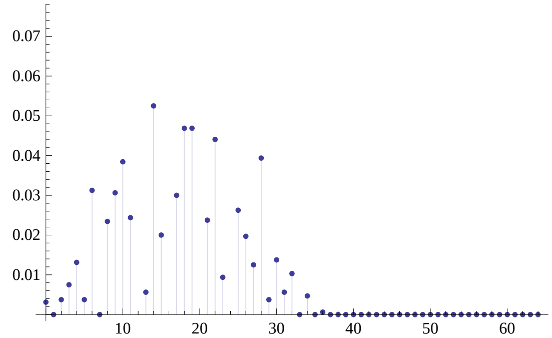

Fig. 3. Expected utility of outcome 1 in alternative C

TABLE III
UTILITY DISTRIBUTIONS FOR ALTERNATIVE C

Sub - outcome	0	1	2	3	4	5	6	7	8
1	.063	.255	.382	.262	.038	.001	0	0	0
2	0	0	.2	.5	.3	0	0	0	0
3	0	0	0	.1	.2	.3	.2	.1	.1

This expected utility distributions in turn serves as utility distribution for outcome 1. But in order to facilitate computation of expected utility of alternative C, the 65 values in Figure 3 are collected into 9 values corresponding to the scale 0 – 8 employed for the other utilities. See Table III.

The probabilities for the three main outcomes are $p_1, p_2 \geq 1/8$ and $p_3 \geq 3/8$, so the multinomial second-order distribution is

$$\Pr(k_1/8, k_2/8) = \frac{3!}{(k_1-1)!(k_2-1)!(5-k_1-k_2)!3^3},$$

resulting in the expected utility distribution plotted in figure 4.

D. Alternative D

As with alternative C there are three possible outcomes in alternative D but no sub-outcomes.

And since the lower bounds of probabilities are $p_1 \geq 0, p_2 \geq 1/8, p_3 \geq 1/8$ the second-order distribution is

$$\Pr(k_1/8, k_2/8) = \frac{6!}{k_1!(k_2-1)!(7-k_1-k_2)!3^6}$$

A plot of the distribution of expected utility is found in Figure 5

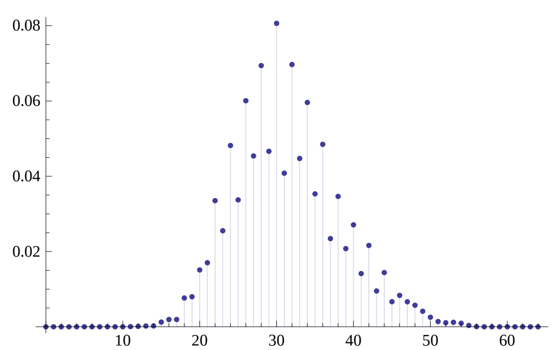

Fig. 4. Distribution of expected utility of alternative C

TABLE IV
UTILITY DISTRIBUTIONS FOR ALTERNATIVE D

Outcome	0	1	2	3	4	5	6	7	8
1	.1	.1	.2	.2	.2	.1	.1	0	0
2	.1	.1	.1	.1	.1	.2	.1	.1	.1
3	.1	.1	.1	.2	.2	.1	.1	.1	0

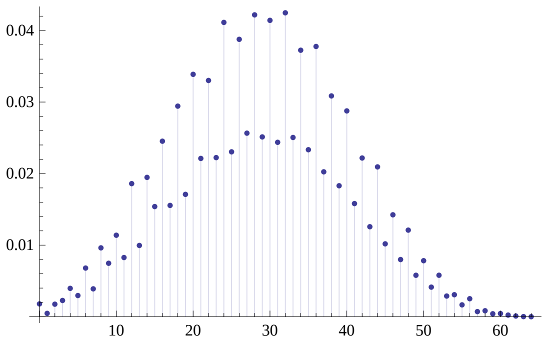

Fig. 5. Distribution of expected utility of alternative D

E. Maximizing Expected Utility

We can begin by ranking he alternatives by expectation of expected utility, $\sum_{z=0}^{64} zh(z)$, we then see that

$$\sum_{z=0}^{64} zh_A(z) = 28.9, \sum_{z=0}^{64} zh_B(z) = 26.2,$$
$$\sum_{z=0}^{64} zh_C(z) = 31.3, \sum_{z=0}^{64} zh_D(z) = 28.8\,,$$

so the ranking is C, A, D, B. Obviously, there is little difference between alternatives A and D, but what is the probability that A yields higher expected utility than D? Or that the highest ranking alternative C gives a better result than the second best A?

With probability

$$\sum_{z=0}^{64}\sum_{x=0}^{z} h_C(z)h_A(x) = .6046$$

alternative C would give at least as high expected utility as alternative A. Given that anything less than .5 would mean that C is worse than A, the superiority of C over A is not very impressive. The point is that the uncertainties inherent in the second-order probabilities and utilities carry over to uncertainty of expected utility and that given the probability of 60 % one can give a cautious recommendation for alternative C.

Further, the probability that A gives at least as high expected utility as D is a mere .52, as slight change in the underlying data could reverse the verdict in favor of D. And with a probability of .59 alternative D is at least as good as B.

The supports for distributions of expected utility, or interval-based expected utilities, are as follows,

$$6 \leq \mathrm{EU}_A \leq 51, 3 \leq \mathrm{EU}_B \leq 49, 11 \leq \mathrm{EU}_C \leq 57, 0 \leq \mathrm{EU}_D \leq 63$$

Overlapping of intervals makes it hard to draw conclusions in terms of which is the better alternative without applying a maximax or maximin rule. The highest lower bound of expected utility is 11 for alternative C but the highest upper bound is 63 for D. In such an interval-based decision analysis there would not be room for estimation of how probable these extreme values are.

IV. TIME COMPLEXITY

The relevant parameters for the complexity of computing discrete expected utility are n, the number of possible outcomes, and N, the number of points in the discrete distributions. The number of points are not necessarily the same for probabilities and utilities, or even between probabilities or utilities, but for simplicity we assume that they are, as in the example above. In any case the different numbers of distributions points would hardly differ by magnitudes.

Assuming $N+1$ first-order probability values k_i/N ranging from 0 to 1 and $N+1$ utility values, there are N^1+1 different possible values of expected utility, from 0 to N^2. Given expected utility z, we must collect all allowable probability and utility vectors that produce the value z of expected utility. But since there is not as yet an expression for the minimal solution x to a diophantine equation $ax+by=c$, (the formula $x = a^{-1}c(\mathrm{mod} b)$ depends on a and b being co-prime), we cannot use a closed expression for directly computing the distribution of discrete expected utility. The raggedness of the plots give some indication that discrete distributions of expected utility might not have simple expressions independent of primality and co-primality. Instead we have to go through all $(N+1)^{2n-1}$ possible choices of $n-1$ probabilities and n utilities and see whether they produce the expected utility z. Each such check costs $\mathcal{O}(\log^2 N)$ arithmetic operations. In total we have $\mathcal{O}(N^{2n}\log^2 N)$ operations for computing expected utility.

If we consider the number of possible consequences n of a decision alternative as a constant inherent in the problem, the time complexity is polynomial in N, meaning that increasing the granularity of the distributions is not prohibitively expensive. However, it is also possible to imagine that deeper investigation of the decision problem leads to splitting of consequences, thus increasing n. Than again, such a situation would rather lead to deeper levels of the decision tree as in alternative B as in our example in Section III. Then n does not change, but computing the distribution of utility for the sub-events makes for another $\mathcal{O}(N'^{2n'}\log N')$ operations.

We have suggested a decision rule based on expectation of expected utility and probability assessments of the probability that one alternative yields higher expected utility than another.

Expectation of expected utility is computed by N^2+1 multiplications and N^2 additions, $\mathcal{O}(N^2\log N)$. And the suggested comparison of alternatives, the probability that, say alternative A has at least as high expected utility as alternative B, $\sum_{z=0}^{N^2}\sum_{x=0} zh_A(z)h_B(x)$ means

$$\frac{N^2(N^2+1)(2n^2+1)}{12} - \frac{N^2(N-1)}{4} \in \mathcal{O}(N^6)$$

multiplications, or $\mathcal{O}(N^6\log^2 N)$ elementary operations.

V. CONCLUSIONS

With second-order probabilities it is possible to express any consistent beliefs about the probabilities of an event. In fact, it is not hard to imagine that a decision maker might shy away from all the possibilities and express his or her knowledge through e.g. intervals. but it may be that the nature of the decision problem makes some second-order distributions more suitable than others. Such differentiation will be a matter for future research. Also, given a certain distribution family, consistency constraints might limit the choice to the lower bounds of the probabilities. Indeed such local information might be all that the decision maker has access to. Here we have looked at multinomial distributions for the purpose of expressing second-order probabilities.

It has been demonstrated that discrete second-order probability distributions allow for updating through observations in a way that continuous distributions would not. And such discrete distributions lend themselves naturally to probability viewed as relative frequency. Furthermore, even when relative frequencies are inappropriate or unavailable, or simply when subjective probabilities are desired or required, continuous second-order distributions offer more possibilities than their discrete counterparts only in rather contrived examples.

In the second-order model, uncertainty is in a manner of speaking made precise. For utilities, probabilities and expected utility, variances may be computed, or the probability that the probability of an outcome is lower than a given value, or the probability that a decision alternative has a higher expected utility than another alternative. Such computations are however for the foreseeable future impossible to conduct save by simulation when using continuous distributions.

We have shown an example of second-order decision making with discrete distributions and shown that the necessary calculations are computationally costly, but far from intractable.

References

[1] A. P. Dempster, “Upper and lower probabilities induced by a multivalued mapping,” *Annals of Mathematical Statistics, vol xxxviii*, pp. 325–399, 1967.

[2] G. Shafer, *A Mathematical theory of evidence*. Princeton University Press, 1976.

[3] P. Huber, “The case of choquet capacities in statistics,” *Bulletin of the International Statistical Institute, 45*, pp. 181–188, 1973.

[4] P. Huber and V. Strassen, “Minimax tests and the neyman-pearsons lemma for capacities,” *Annals of Statistics, 1*, pp. 251–263, 1973.

[5] I. Good, “Subjective probability as the measure of a non-measurable set,” in *Logic, Methodology, and the Philosophy of Science*, P. Suppes, B. Nagel, and A. Tarski, Eds. Stanford University Press, 1962, pp. 319–329.

[6] R. F. Nau, “Uncertainty aversion with second-order utilities and probabilities,” *Management Science*, vol. 52, no. 1, pp. 136–145, 2006.

[7] L. Ekenberg and J. Thorbiörnson, “Second-order decision analysis,” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, vol. 9, No 1*, vol. 9, no. 1, pp. 13–38, 2001.

[8] L. V. Utkin and T. Augustin, “Decision making with imprecise second-order probabilities,” in *ISIPTA ’03 - Proceedings of the Third International Symposium on Imprecise Probabilities and Their Applications*, 2003, pp. 547–561.

[9] L. V. Utkin, “Imprecise second-order hierarchical uncertainty model,” *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems, Volume 11:3*, pp. 301–317, 2003.

[10] G. Choquet, “Theory of capacities,” *Ann. Inst. Fourier*, pp. 5:131–295, 1954.

[11] C. A. B. Smith, “Consistency in statistical inference and decision,” *Journal of the Royal Statistical Society, Series B, xxiii*, pp. 1–25, 1961.

[12] P. Walley, *Statistical reasoning with Imprecise Probabilities*. Chapman and Hall, 1991.

[13] F. Cozman, “A brief introduction to the theory of sets of probability measures,” cMU Tech. report CMU-RI-TR 97-24.

[14] J. K. Satia and R. E. L. Jr., “Markovian decision processes with uncertain transition probabilities,” *Operations Research*, vol. 21, pp. 728–740, 1973.

[15] H. Robbins, “Asymptotically sub-minimax solutions to compound statistical decision problems,” *Proc. Second Berkerly Symp. Math. Stat. Probab.*, vol. 1, 1951.

[16] I. J. Good, “Rational decisions,” *Journal of the Royal Statistical Society, Series B*, vol. 14 (1), pp. 107–114, 1952.

[17] I. Levi, *The Enterprise of Knowledge*. MIT Press, 1980.

[18] M. Zaffalon, K. Wesnes, and O. Petrini, “Reliable diagnoses of dementia by the naive credal classifier inferred from incomplete cognitive data,” *Artificial Intelligence in Medicine*, vol. 29, no. 1-2, pp. 61–79, 2003.

[19] M. C. Troffaes, “Decision making under uncertainty using imprecise probabilities,” *International Journal of Approximate Reasoning*, vol. 45, no. 1, pp. 17 – 29, 2007.

[20] D. Sundgren, L. Ekenberg, and M. Danielson, “Some properties of aggregated distributions over expected values,” in *MICAI 2008: Advances in Artificial Intelligence*, ser. Lecture Notes in Computer Science, no. 5317. Springer, 2008, pp. 699–709.