

GENERALISING THE CONJUNCTION RULE FOR AGGREGATING CONFLICTING EXPERT OPINIONS

MATTHIAS C. M. TROFFAES
SYSTEMS RESEARCH GROUP, GHENT UNIVERSITY, BELGIUM

ABSTRACT. In multi-agent expert systems, the conjunction rule is commonly used to combine expert information represented by imprecise probabilities. However, it is well-known that this rule cannot be applied in case of expert conflict. In this paper, we propose to resolve expert conflict by means of a second-order imprecise probability model. The essential idea underlying the model is a notion of behavioural trust. We construct a simple linear programming algorithm for calculating the aggregate. This algorithm explains the proposed aggregation method as a generalised conjunction rule. It also provides an elegant operational interpretation of the imprecise second-order assessments, and thus overcomes the problems of interpretation that are so common in hierarchical uncertainty models.

1. INTRODUCTION

When modelling a system, one must often rely on expert information. From the modeller's perspective, one usually wants to aggregate all expert opinions into a single representative model—a “summary” of all the expert information—which will then serve as a basis for various kinds of inferences about the system, such as decision making, estimation, hypothesis testing, *etc.* The fundamental idea behind this approach is that aggregating more expert opinions eventually leads to a more reliable, and hopefully, also to a more informative representative model. There is however no agreement on how expert opinions should be aggregated. Actually, there is not even a clear agreement on how expert opinions themselves should be represented.

Recently, the use of imprecise probabilities in representing, manipulating and aggregating expert information has received an increasing amount of attention in the literature (see for instance [16, 17, 20, 18, 11, 3, 12, 5, 15, 4] and many references therein). One of the main reasons for the increasing popularity of imprecise probabilities in modelling and aggregating expert information is that they allow for a more reliable representation of expert information. Indeed, in practice we often have only limited information about probability distributions. Imprecise probabilities reliably model limited information, and do not force us to pinpoint a

E-mail address: Matthias.Troffaes@UGent.be.

single probability measure in order to represent our knowledge. Secondly, imprecise probabilities also provide a natural setting for modelling conflicting opinions, using imprecision as a means of expressing disagreement amongst different opinions.

The conjunction rule is widely used as an aggregation rule for non-conflicting expert opinions. Conjunction gains as much information as possible from each of the experts. However, it cannot be applied in case of conflicting expert opinions, and it is not entirely clear how the rule should be generalised in order to deal with conflict.

This paper aims at providing a new systematic and computationally simple way for reconciling conflicting expert opinions, in order to generalise the conjunction rule. Throughout we shall use behavioural arguments only, in particular, avoiding sure loss, coherence and natural extension—these are fundamental concepts in the behavioural theory of imprecise probabilities [19]. We define our aggregate by means of a very general imprecise second-order hierarchical model, and for a number of important special cases we derive a simple linear programming algorithm for calculating the aggregate, whose dual also provides us with an elegant operational interpretation of the second-order assessments. For some imprecise second-order models, the proposed algorithm does not work anymore—more complex techniques are needed—and hence, a *general* operational interpretation of imprecise second-order models remains an open problem.

Imprecise hierarchical models have been studied quite extensively in the literature, although these models have not always aimed at reconciling conflict in multi-agent expert systems. Our model generalises so-called lower desirability functions introduced by de Cooman [3]. Therefore, our algorithm can also be used to calculate the first order aggregate induced by such lower desirability functions (this aggregate is also called the first-order *i*-natural extension). Another very interesting and mathematically closely related aggregation algorithm was studied by Utkin [15]. The algorithm described by Utkin is more complex and more powerful—it works in all cases—but at a price: one needs to make stronger assumptions about the second-order level, and one can only use imprecise first-order expert assessments of a very specific form.

The paper is organised as follows. Section 2 introduces the basic concepts of the behavioural theory of imprecise probabilities under the form of lower previsions, and their relation to other well-known uncertainty models. In Section 3 we explain the problem of aggregating expert opinions, and we touch on the controversy surrounding it. Section 4 explains the conjunction rule. A second-order imprecise probability model is proposed and discussed in Section 5. In Section 6 the main results are presented. Section 7 gives a numerical example, and we end with a discussion in Section 8.

2. LOWER PREVISIONS

In this paper, lower previsions are taken as the fundamental imprecise probability model. Their behavioural interpretation turns out to be very convenient in describing the second-order model later on. We only introduce the most important aspects of the theory of lower previsions that are relevant to the problem at hand. More details can be found in [19].

Let us consider a subject (which can be an expert, or a modeller) who is uncertain about something, say, the outcome of some experiment. If the set of possible outcomes is Ω , then a *gamble* X is a bounded real-valued mapping on Ω , and it is interpreted as an uncertain reward: if ω turns out to be the true outcome of the experiment then the subject receives the amount $X(\omega)$, expressed in units of some linear utility. The set of all gambles on Ω is denoted by $\mathcal{L}(\Omega)$.

The information the subject has about the outcome of the experiment will lead him to accept or reject transactions whose reward depends on this outcome, and we can formulate a model for his uncertainty by looking at a specific type of transaction: the buying of gambles. The subject's *lower prevision* (or supremum acceptable buying price) $\underline{P}(X)$ for a gamble X is the highest price s such that he is disposed to buy the gamble X for any price strictly lower than s . If the subject assesses a supremum acceptable buying price for every gamble X in a subset \mathcal{K} of $\mathcal{L}(\Omega)$, the resulting mapping $\underline{P} : \mathcal{K} \rightarrow \mathbb{R}$ is called a lower prevision.

Examples of lower previsions are:

- (i) If “ ω belongs to the set $A \subseteq \Omega$ ” then $\underline{P}_A(X) = \inf_{\omega \in A} X(\omega)$: the lowest possible reward given that $\omega \in A$. We call \underline{P}_A the *vacuous lower prevision* relative to A .
- (ii) If “ ω has probability density ϕ ” we should pay $P(X) = \int_{\Omega} X(\omega)\phi(\omega)d\omega$, the expectation w.r.t. ϕ [6, 7]. This is called the *linear prevision* induced by the density ϕ .
- (iii) If “ ω has a probability density that belongs to the set Φ ” we pay at most $\underline{P}(X) = \inf_{\phi \in \Phi} \int_{\Omega} X(\omega)\phi(\omega)d\omega$.

These examples indicate that lower previsions are uncertainty *representations* that are expressive enough to capture propositional logic (example (i)), Bayesian probability theory [13] (example (ii)), and credal sets [10] (example (iii)) (credal sets are closed convex sets of probability measures). Actually, they also generalise belief functions [8, 14], possibility and necessity measures [21], Choquet capacities [2], risk measures [1], and many other uncertainty models.

\bar{P} will denote the conjugate *upper prevision* of \underline{P} . It is defined by $\bar{P}(X) = -\underline{P}(-X)$ for every $X \in -\mathcal{K}$. $\bar{P}(X)$ represents the subject's infimum acceptable selling price for the gamble X . The difference $\bar{P}(X) - \underline{P}(X)$ is a measure for the amount of imprecision in the subject's behavioural dispositions towards X .

We now introduce a method of *inference*, associated with lower previsions, that also generalises the inference methods of, for instance, classical propositional logic and Bayesian probability theory.

2.1. Inference. Through a procedure called natural extension, we are able to derive from the assessments embodied in \underline{P} , a supremum buying price $\underline{E}(X)$ for each gamble X in $\mathcal{L}(\Omega)$; we want to find the point-wise smallest (and therefore most conservative) lower prevision \underline{E} that satisfies for any gambles X and Y

- $\underline{E}(X) \geq \inf[X]$ (accepting sure gain)
- $\underline{E}(\lambda X) = \lambda \underline{E}(X)$ whenever $\lambda > 0$ (scale independence)
- $\underline{E}(X + Y) \geq \underline{E}(X) + \underline{E}(Y)$ (super-additivity)
- $\underline{E}(X) \geq \underline{P}(X)$ (compatibility)

If such \underline{E} exists, \underline{P} is said to *avoid sure loss*. It can be easily shown that a lower prevision avoids sure loss if and only if no “Dutch book” argument can be made against \underline{P} , that is, if and only if there is no combination of transactions—buying gambles for acceptable buying prices—that leads to a sure loss. Mathematically, this means that $\sup_{\omega \in \Omega} [\sum_{i=1}^m [X_i(\omega) - \underline{P}(X_i)]] \geq 0$ must hold for any $m \in \mathbb{N}$ and any $X_1, \dots, X_m \in \mathcal{K}$.

In case \underline{P} avoids sure loss, \underline{E} exists and is called the *natural extension* of \underline{P} . For any gamble X , the natural extension $\underline{E}(X)$ can be easily calculated: assuming \mathcal{K} to be finite, it is equal to the supremum α^* achieved by the free variable α subject to

$$(1) \quad X(\omega) - \alpha \geq \sum_{Y \in \mathcal{K}} \lambda_Y (Y(\omega) - \underline{P}(Y))$$

for each $\omega \in \Omega$, with variables $\lambda_Y \geq 0$ for each $Y \in \mathcal{K}$ —if also Ω is finite, which happens quite often in practice, then this is a linear program. The fact that \underline{P} avoids sure loss guarantees that the problem has a solution $\alpha^* \in \mathbb{R}$.

Of course, we may not know whether \underline{P} avoids sure loss or not. For arbitrary \underline{P} , if the supremum α^* happens to be $+\infty$ for some gamble X , then it will be $+\infty$ for any gamble X and the natural extension does not exist. Hence, in such a case \underline{P} incurs sure loss; this identifies a conflict in the assessments. Thus, the linear program used to calculate natural extension can also be used to detect a “Dutch book” simply by solving it for an arbitrary gamble X .

If \underline{E} and \underline{P} coincide on \mathcal{K} , then \underline{P} is called *coherent*. It is easy to see that natural extension is always coherent, and in fact, by its definition it is the point-wise smallest coherent lower prevision on $\mathcal{L}(\Omega)$ that is compatible with \underline{P} .

3. AGGREGATION: A SHORT REVIEW

We now shortly review some different ways to tackle the problem of aggregating expert opinions. Basically, there are two ways to approach the problem: *axiomatic* (also called normative), and *ad hoc*. No rule is ever purely axiomatic, or purely *ad*

hoc. Many rules can be given an axiomatic as well as an *ad hoc* explanation (such as the conjunction rule and the unanimity rule described below).

Axiomatic approaches aim at deriving a preferably unique rule of aggregation from axioms or properties that this rule should satisfy. Typical axioms are requirements of commutativity of the rule with respect to some other action, such as updating (external Bayesianity), marginalisation, permutation of experts (symmetry) *etc.* They can also refer to some other property of the rule, such as unanimity-preservation (if all experts agree, then the aggregate should also agree with all experts), invariance with respect to non-informative expert opinions, independence preservation, *etc.*

Especially among Bayesians (see for instance [9] for an overview, and references therein), where expert opinions and the aggregate are to be represented by probability measures, there still is a lot of controversy about these axioms. Indeed, imposing even only a few axioms easily leads to contradictions or undesirable aggregation rules such as so-called dictatorship rules. What counts is how the rule will eventually be used. From this perspective, it is not always clear what axioms should be imposed.

In imprecise probability theory, the axiomatic approach is somewhat less problematic (see [16] for a discussion). Yet, it is still not clear how to define a unique aggregation rule under this uncertainty model. The *conjunction rule* is defined as the smallest (and therefore most conservative) coherent lower prevision that dominates each of the experts' lower previsions. Conjunction aims at gaining as much information as possible from each of the experts: the aggregate is at least as informative as each of the experts' lower previsions, and it can only become more informative as more experts enter the scene. The conjunction however does not always exist, in particular when different experts make conflicting statements. On the other hand, the *unanimity rule*, defined as the lower envelope of the experts' lower previsions, is guaranteed to exist. It aims at reconciling the experts' assessments. As a result however, it may lead to extremely imprecise results: the aggregate will be at least as imprecise as the most imprecise expert, and its imprecision can only increase as more experts enter the scene. Unanimity certainly leads to a very reliable aggregate. However, it fails completely to produce also a more informative aggregate as more expert assessments become available.

One imprecise probability aggregation rule could consist of using the conjunction rule if the conjunction exists, and the unanimity rule if the conjunction does not exist. The problem of this rule is that it is far from stable: a small variation of an expert's lower prevision may yield huge differences in the aggregate lower prevision.

Ad hoc approaches are not as much concerned with axioms: one simply proposes or derives a mathematical formula, together with some form of justification. (Afterwards of course, it is usually investigated which of the axioms it satisfies. This usually provides the *ad hoc* rule with an additional source for motivation or

criticism.) They generally divide into three sub-categories: hierarchical models, weighting schemes, and consensus methods. Consensus methods are based on expert interaction: before an aggregate is constructed, the experts are allowed to interact with each other (see [12] for an excellent discussion of a consensus method using imprecise probabilities). Weighting rules, the linear opinion pool in Bayesian aggregation being maybe the most prevailing example, try to take each expert's expertise into account (a feature lacking most of the purely axiomatic approaches). The same holds for hierarchical models, and in fact, hierarchical models may be seen as one attempt to motivate, and generalise, some of the existing weighting schemes (many weighting rules are however not instances of hierarchical models).

Using probability measures, the most reasonable approach seems to be linear pooling; taking a convex combination of expert probability measures. It is very easy to implement, and gives quite good results in practice. Subject of debate is of course how one should assign the weights. However, it still seems a bit ridiculous to be precise about the weights, if one is not even sure about the first order level. An imprecise pooling method can resolve this, but it is not clear at first sight how this can be done. We note that the method proposed in this paper could be somehow interpreted as an imprecise pooling method—in Section 6.1 we shall derive precise weights on the set of all expert conjunctions from imprecise weights on the original expert models only. It is surprising that we can do this by natural extension only.

Concluding, besides theoretical and practical problems associated with each of these methods separately, any method using single probability distributions for both the experts and the aggregate fails to model conflict among experts, and forces experts to pinpoint a single probability, even for those events of which he does not have much expertise. Imprecise probabilities address both these problems, because they allow for experts to assess their expertise using a closed convex set of probability measures (also called credal sets), a lower prevision, a set of desirable gambles, an ordering on gambles, a possibility measure, *etc.*—rather than forcing them to choose a single probability measure. Consequently, it is also easier to avoid conflict when combining imprecise probabilities because, roughly speaking, experts are not forced to give precise probabilities on events of which they have only little knowledge—they can simply say they don't know. And should there be conflict anyway, imprecision can be used to reflect it (for instance, using the unanimity rule). These characteristics are the main motivation for introducing the second-order imprecise probability model in Section 5.

4. THE CONJUNCTION RULE

Suppose there are n (male) subjects, called *experts*. The set of all experts is denoted by $N = \{1, \dots, n\}$. Assume that the parameter ω of interest assumes values in a *finite* set of possible values Ω . Each expert $k \in N$ expresses his beliefs about ω through a lower prevision \underline{P}_k on some *finite* subset \mathcal{K}_k of $\mathcal{L}(\Omega)$. We assume

that each \underline{P}_k has a natural extension (i.e., avoids sure loss). The natural extension of \underline{P}_k will be denoted by \underline{E}_k . Throughout, we shall assume all \underline{E}_k to be different, i.e., for any $k, \ell \in N$, $k \neq \ell$, we assume there is at least one gamble X such that $\underline{E}_k(X) \neq \underline{E}_\ell(X)$. This simplifies the analysis and notation used further on, does not essentially change any of the results, and is satisfied in most cases of interest. Now, how can the lower previsions $\underline{P}_1, \dots, \underline{P}_n$ be combined into an aggregate—a single coherent lower prevision defined on the set of all gambles $\mathcal{L}(\Omega)$?

Consider therefore a new (female) subject, called the *modeller*. She wishes to aggregate the expert assessments to a single coherent lower prevision defined on $\mathcal{L}(\Omega)$. Let us first introduce a notion of behavioural trust.

Definition 1. *Let α and β be two subjects. Assume that each of the subjects models his/her knowledge about $\omega \in \Omega$ through a coherent lower prevision \underline{P}_α resp. \underline{P}_β on \mathcal{K}_α resp. \mathcal{K}_β . Let \underline{E}_α resp. \underline{E}_β denote their natural extension. The following conditions are equivalent; if any (hence all) of them are satisfied, we say that α trusts β .*

- (A) *α is willing to accept every decision β makes concerning buying gambles on Ω , that is, for each gamble $X \in \mathcal{L}(\Omega)$, α is willing to accept β 's price $s < \underline{E}_\beta(X)$ for buying X as his/her price for buying X .*
- (B) *\underline{E}_α point-wise dominates \underline{E}_β on $\mathcal{L}(\Omega)$, that is, $\underline{E}_\alpha(X) \geq \underline{E}_\beta(X)$ for any gamble $X \in \mathcal{L}(\Omega)$.*

The point of the first part of the definition is that any behavioural theory of uncertainty inherently has a notion of trust in a multi-agent environment and hence, as we show now, also notions of conjunction and consistency, which can be derived from behavioural trust in a straightforward way.

Definition 2. *If there is a point-wise smallest, and hence most conservative, coherent lower prevision on $\mathcal{L}(\Omega)$ the modeller can have such that she still trusts each of the experts of $\underline{P}_1, \dots, \underline{P}_n$, then this lower prevision is called the conjunction of $\underline{P}_1, \dots, \underline{P}_n$. If this conjunction exists, then the experts are said to be consistent, otherwise they are said to be conflicting.*

By Definition 1(B), the conjunction is simply the (point-wise) smallest coherent lower prevision that dominates all the experts' natural extensions \underline{E}_k . The conjunction of $\underline{P}_1, \dots, \underline{P}_n$ is denoted by $\sqcap_{k \in N} \underline{P}_k$; the conjunction of two consistent coherent lower previsions \underline{P}_1 and \underline{P}_2 is also denoted by $\underline{P}_1 \sqcap \underline{P}_2$. It is easy to show that \sqcap is an associative and commutative operator on coherent lower previsions, but the result is only defined in case of consistency (see [4] for many more properties). Conjunction can be calculated through linear programming in a similar way as natural extension.

Proposition 1. Consider the maximum α^* achieved by the free variable α subject to the linear constraints

$$X(\omega) - \alpha \geq \sum_{k \in N} \sum_{Y \in \mathcal{K}_k} \lambda_{k,Y} (Y(\omega) - \underline{P}_k(Y))$$

for each $\omega \in \Omega$, with variables $\lambda_{k,Y} \geq 0$ for each $k \in N$ and $Y \in \mathcal{K}_k$. If α^* is finite then the expert assessments $\underline{P}_1, \dots, \underline{P}_n$ are consistent and $\bigcap_{k \in N} \underline{P}_k(X) = \alpha^*$. If $\alpha^* = +\infty$, then the conjunction does not exist: in such a case the assessments are conflicting.

Again, note that if α^* is $+\infty$ for some gamble X then it will be $+\infty$ for any gamble X , as in Section 2.1. If the assessments are conflicting—if no conjunction exists—then there is no coherent way to accept *every* decision of *every* expert, since the modeller incurs a sure loss if she would do so. It is easily established that in case of inconsistency there are gambles $X_k \in \mathcal{L}(\Omega)$ such that

$$(2) \quad \sup_{\omega \in \Omega} \left[\sum_{k \in N} [X_k(\omega) - \underline{E}_k(X_k)] \right] < 0,$$

i.e., the combination of the transactions in which the gambles X_k are bought for a price $\underline{E}_k(X_k)$ leads to a loss, whatever the actual value of the parameter ω . Blindly accepting decisions of all experts is clearly unacceptable in case of inconsistency. The modeller is therefore certain that some of the experts' assessments cannot be trusted. But she does not necessarily know which ones.

One solution in case of conflict is to use the unanimity rule. This consists in choosing the modeller's lower prevision, the aggregate, such that each of the experts trust the modeller. This means that each of the experts agrees with the modeller's behavioural dispositions (hence the name of the rule). But as we have already noted before, the resulting aggregate may be too imprecise to be useful.

It may however happen that the modeller may have actual information about which of the experts are to be trusted more than others. In the next section we propose a second-order hierarchical imprecise probability model that aims at modelling such knowledge. Its interpretation is based on the notion of behavioural trust.

5. A GENERAL SECOND-ORDER IMPRECISE PROBABILITY MODEL

The modeller wishes to recover information regarding ω using the information revealed by the experts, taking into account that some experts are more trustworthy than others. We describe how behavioural trust can be used to aggregate information revealed by experts.

The modeller first assumes the existence of a so-called *true coherent lower prevision* \underline{P}_T on $\mathcal{L}(\Omega)$, but she is not sure about what it is. \underline{P}_T could refer to the behaviour of a hypothetical “representative” expert, an operational procedure designed to measure uncertainty such as an imprecise Dirichlet (or other) model

updated through a contingency table, or even a real system that behaves just like an expert. The modeller is interested in what the hypothetical expert knows about ω , or what the result of the operational procedure will be about ω , or how the system behaves with respect to ω , but, she is only able to infer information about ω through $\underline{P}_1, \dots, \underline{P}_n$. She cannot talk to the hypothetical representative expert, cannot perform the operational procedure, has no access to the system of interest: it may be too expensive, or she might not have the necessary means. Her uncertainty thus regards the random variable \underline{P}_T which we assume to take all values in the set $\underline{\mathcal{P}}(\Omega)$ of coherent lower previsions on $\mathcal{L}(\Omega)$. Her possibility space $\underline{\mathcal{P}}(\Omega)$ is also called the second-order possibility space.¹

Often, even with imprecise hierarchical models, the second-order possibility space is restricted to the set of all linear previsions (see for instance [15]). It is well-known that restricting the second-order possibility space to linear previsions may lead to different results: the so-called *precision-imprecision equivalence* does not always hold (see [3] for a discussion and an example where the equivalence fails). Our main motivation for *not* restricting to linear previsions is that we should not expect experts to be able to pinpoint a single probability measure. We want experts to be honest about their information, so if there really is uncertainty, we sure want them to be able to tell us. This should hold as well for the “real” experts as for the hypothetical representative expert.

5.1. Trust and Dual Trust. In terms of events on the modeller’s second-order possibility space, we may consider the event that the true behavioural dispositions implied by \underline{P}_T include at least expert k ’s behavioural dispositions implied by \underline{E}_k (remember that \underline{E}_k is the natural extension of \underline{P}_k). This event obtains when \underline{P}_T belongs to the set

$$\underline{\mathcal{M}}(\underline{E}_k) = \{\underline{P} \in \underline{\mathcal{P}}(\Omega) : (\forall X \in \mathcal{L}(\Omega))(\underline{E}_k(X) \leq \underline{P}(X))\}.$$

The modeller is unsure about \underline{P}_T , but we assume she can assess a supremum buying price \underline{t}_k for the gamble $I_{\underline{\mathcal{M}}(\underline{E}_k)}$ that returns a unit gain if the event $\underline{P}_T \in \underline{\mathcal{M}}(\underline{E}_k)$ obtains and nothing otherwise.² Likewise, she assesses a supremum buying price $1 - \bar{t}_k$ for the gamble $I_{\underline{\mathcal{C}}\underline{\mathcal{M}}(\underline{E}_k)} = 1 - I_{\underline{\mathcal{M}}(\underline{E}_k)}$ that returns a unit gain if the event $\underline{P}_T \in \underline{\mathcal{M}}(\underline{E}_k)$ does not obtain and nothing if it does obtain. If she is completely ignorant about $\underline{P}_T \in \underline{\mathcal{M}}(\underline{E}_k)$ or about its complement, she should choose supremum buying price zero.

The interval $[\underline{t}_k, \bar{t}_k]$ can be interpreted as a probability interval for the event $\underline{P}_T \in \underline{\mathcal{M}}(\underline{E}_k)$. Consider the case in which the modeller believes with certainty that $\underline{P}_T \in \underline{\mathcal{M}}(\underline{E}_k)$, i.e., $\underline{t}_k = \bar{t}_k = 1$. In that case, the modeller is sure that any

¹The first-order possibility space is Ω , and $\underline{P}_1, \dots, \underline{P}_n$ and \underline{P}_T are called first-order models.

²The gamble $I_{\underline{\mathcal{M}}(\underline{E}_k)}$ is nothing but the indicator function of the set $\underline{\mathcal{M}}(\underline{E}_k)$: $I_{\underline{\mathcal{M}}(\underline{E}_k)}(\underline{P}) = 1$ if $\underline{P} \in \underline{\mathcal{M}}(\underline{E}_k)$, and $I_{\underline{\mathcal{M}}(\underline{E}_k)}(\underline{P}) = 0$ otherwise. Lower previsions on gambles that are indicator functions are a common (but not the most general) method for representing imprecise probabilities.

buying price of expert k for a gamble X is also a buying price of the representative expert for X , and hence, she should fully accept expert k 's decisions (regarding the buying of gambles) as hers (remember that ideally, she wants to behave as \underline{P}_T). According to Definition 1 she “trusts” expert k whenever $\underline{t}_k = \bar{t}_k = 1$. Therefore, we shall call \underline{t}_k and \bar{t}_k the modeller’s *lower and upper trust* assigned to expert k .

Dually, it may happen that the behavioural dispositions implied by the expert k include at least the true behavioural dispositions: the expert’s assessments may to be too precise (for instance, he might be a Bayesian restricting to linear previsions as in Example (ii) of Section 2), but not necessarily contradicting \underline{P}_T . This obtains when \underline{P}_T belongs to the set

$$\mathcal{N}(\underline{E}_k) = \{\underline{P} \in \mathcal{P}(\Omega) : (\forall X \in \mathcal{L}(\Omega))(P(X) \leq \underline{E}_k(X))\}.$$

(note that $\mathcal{N}(\underline{E}_k) \neq \mathcal{CM}(\underline{E}_k)$). Again, the modeller can assess a supremum buying price \underline{t}'_k for the gamble $I_{\mathcal{N}(\underline{E}_k)}$, and likewise, a supremum buying price $1 - \bar{t}'_k$ for the gamble $1 - I_{\mathcal{N}(\underline{E}_k)}$.

The interval $[\underline{t}'_k, \bar{t}'_k]$ can be interpreted as a probability interval for the event $\underline{P}_T \in \mathcal{N}(\underline{E}_k)$. In case $\underline{t}'_k = \bar{t}'_k = 1$, the modeller is sure about $\underline{P}_T \in \mathcal{N}(\underline{E}_k)$, and this means that any price for a gamble X that is not acceptable as a buying price for expert k will also not be an acceptable as a buying price for the representative expert \underline{P}_T . She should fully reject any behaviour that is not included in the behaviour of expert k (regarding the buying of gambles). Therefore, we shall call \underline{t}'_k and \bar{t}'_k the modeller’s *lower and upper dual trust* assigned to expert k .

An important issue in these definitions of lower and upper trust, and lower and upper dual trust, is that the events $\underline{P}_T \in \mathcal{M}(\underline{E}_k)$ and $\underline{P}_T \in \mathcal{N}(\underline{E}_k)$ (and their complements) are not observable in general. Indeed, the behaviour described by \underline{P}_T refers to a hypothetical representative expert, and in practice it is far from clear how to set up an objective method for measuring this “representative” behaviour \underline{P}_T . Stated as such, the model described here belongs to the realm of fantasy. However, being stubborn and investigating this at first sight useless model a little further, we shall see in Section 6.1 that, most surprisingly, it is possible to give lower trust, upper trust, and lower dual trust an operational interpretation whenever all upper dual trust is one (i.e., vacuous).

Remark that if an expert has different expertise in different domains we may wish to assign higher (and perhaps also more precise) trust intervals to those domains in which he’s more experienced. A natural solution to this is to consider such an expert as a set of “sub-experts”, each member corresponding to a particular domain of expertise, and then to assign different (dual) trust intervals to each sub-expert. We then treat the sub-experts simply as regular experts and proceed with the aggregation as usual. In this way we can improve the overall precision of the result.

We defined \underline{t}_k , \bar{t}_k , \underline{t}'_k and \bar{t}'_k as specifications of buying prices on particular gambles on the second order possibility space $\mathcal{P}(\Omega)$. In terms of a lower prevision

\underline{Q} on the second order possibility space, we have for $k \in N$:

$$\begin{aligned}\underline{Q}(I_{\mathcal{M}(E_k)}) &= \underline{\mathfrak{t}}_k, & \underline{Q}(1 - I_{\mathcal{M}(E_k)}) &= 1 - \bar{\mathfrak{t}}_k, \\ \underline{Q}(I_{\mathcal{N}(E_k)}) &= \underline{\mathfrak{t}}'_k, & \underline{Q}(1 - I_{\mathcal{N}(E_k)}) &= 1 - \bar{\mathfrak{t}}'_k.\end{aligned}$$

Since we assumed all \underline{E}_k to be distinct, all gambles $I_{\mathcal{M}(E_k)}$, $1 - I_{\mathcal{M}(E_k)}$, $I_{\mathcal{N}(E_k)}$ and $1 - I_{\mathcal{N}(E_k)}$ are also distinct, so \underline{Q} is well-defined.³ With the following notation for coherent lower previsions \underline{P}_1 and \underline{P}_2 on $\mathcal{L}(\Omega)$:

$$I_{\underline{P}_1 \geq \underline{P}_2} = \begin{cases} 1, & \text{if for each } X \in \mathcal{L}(\Omega): \underline{P}_1(X) \geq \underline{P}_2(X), \\ 0, & \text{otherwise,} \end{cases}$$

we can write $I_{\mathcal{M}(E_k)}(\underline{P})$ as $I_{\underline{P} \geq \underline{E}_k}$ and $I_{\mathcal{N}(E_k)}(\underline{P})$ as $I_{\underline{E}_k \geq \underline{P}}$.

5.2. A first-order aggregate through natural extension. If \underline{Q} avoids sure loss, that is, if

$$\sup_{\underline{P} \in \underline{\mathcal{P}}(\Omega)} \left\{ \sum_{k \in N} \kappa_k \left(I_{\underline{P} \geq \underline{E}_k} - \underline{\mathfrak{t}}_k \right) + \lambda_k \left(\bar{\mathfrak{t}}_k - I_{\underline{P} \geq \underline{E}_k} \right) \right. \\ \left. + \mu_k \left(I_{\underline{E}_k \geq \underline{P}} - \underline{\mathfrak{t}}'_k \right) + \nu_k \left(\bar{\mathfrak{t}}'_k - I_{\underline{E}_k \geq \underline{P}} \right) \right\} \geq 0,$$

for every $\kappa_k \geq 0$, $\lambda_k \geq 0$, $\mu_k \geq 0$ and $\nu_k \geq 0$, then, as explained in Section 2.1, the natural extension \underline{E} of \underline{Q} exists and is a coherent lower prevision on $\mathcal{L}(\underline{\mathcal{P}}(\Omega))$. In such a case we say that there is *second-order consistency*. If \underline{Q} does not avoid sure loss, then we say that there is *second-order conflict*, and \underline{Q} has no natural extension. In such a case, the modeller should revise her (dual) trust assignments. Typically, she can do this by decreasing the lower (dual) trust and increasing the upper (dual) trust assigned to some of the experts, until \underline{Q} avoids sure loss.

The natural extension, if it exists, is given by

$$(3) \quad \underline{E}(Z) = \sup \left\{ \alpha \in \mathbb{R} : (\exists \kappa_k \geq 0, \lambda_k \geq 0, \mu_k \geq 0, \nu_k \geq 0) (\forall \underline{P} \in \underline{\mathcal{P}}(\Omega)) \right. \\ \left. Z(\underline{P}) - \alpha \geq \sum_{k \in N} \kappa_k \left(I_{\underline{P} \geq \underline{E}_k} - \underline{\mathfrak{t}}_k \right) + \lambda_k \left(\bar{\mathfrak{t}}_k - I_{\underline{P} \geq \underline{E}_k} \right) \right. \\ \left. + \mu_k \left(I_{\underline{E}_k \geq \underline{P}} - \underline{\mathfrak{t}}'_k \right) + \nu_k \left(\bar{\mathfrak{t}}'_k - I_{\underline{E}_k \geq \underline{P}} \right) \right\}$$

for any (second-order) gamble $Z \in \mathcal{L}(\underline{\mathcal{P}}(\Omega))$. The lower prevision \underline{E} represents the modeller's knowledge about the representative expert's knowledge as inferred from

³The general case in which (some of) the \underline{E}_k are allowed to be equal does not pose any conceptual difficulties, it only requires the introduction of a more complicated notation which we would like to avoid for the clarity of exposition.

the experts' judgements $(\underline{P}_1, \dots, \underline{P}_n)$ and the modeller's second-order judgements $(\underline{t}_1, \bar{t}_1, \underline{t}'_1, \bar{t}'_1, \dots, \underline{t}_n, \bar{t}_n, \underline{t}'_n, \bar{t}'_n)$.

From the natural extension \underline{E} , we can, theoretically, deduce lower and upper trust, and lower and upper dual trust of any coherent lower prevision \underline{P} on $\mathcal{L}(\Omega)$:

$$\begin{aligned} \underline{t}(\underline{P}) &= \underline{E}(I_{\mathcal{M}(\underline{P})}), & \bar{t}(\underline{P}) &= 1 - \underline{E}(1 - I_{\mathcal{M}(\underline{P})}) = \bar{E}(I_{\mathcal{M}(\underline{P})}), \\ \underline{t}'(\underline{P}) &= \underline{E}(I_{\mathcal{N}(\underline{P})}), & \bar{t}'(\underline{P}) &= 1 - \underline{E}(1 - I_{\mathcal{N}(\underline{P})}) = \bar{E}(I_{\mathcal{N}(\underline{P})}). \end{aligned}$$

Note that these natural extensions will agree with the original second-order judgements, that is, $\underline{t}(\underline{E}_k) = \underline{t}_k$, $\bar{t}(\underline{E}_k) = \bar{t}_k$, $\underline{t}'(\underline{E}_k) = \underline{t}'_k$ and $\bar{t}'(\underline{E}_k) = \bar{t}'_k$, exactly when \underline{Q} is coherent.

It is perhaps of more interest that we can also infer supremum buying prices and infimum selling prices for the supremum buying price and infimum selling price of a gamble X with respect to the representative expert:

$$\begin{aligned} \underline{E}_*(X) &= \underline{E}(X_*), & \underline{E}^*(X) &= \underline{E}(X^*), \\ \bar{E}_*(X) &= \bar{E}(X_*), & \bar{E}^*(X) &= \bar{E}(X^*), \end{aligned}$$

where X_* is the lower and X^* is the upper evaluation map corresponding to X , defined by

$$X_* : \mathcal{P}(\Omega) \rightarrow \mathbb{R}; \underline{P} \mapsto \underline{P}(X), \quad X^* : \mathcal{P}(\Omega) \rightarrow \mathbb{R}; \underline{P} \mapsto \bar{P}(X).$$

We can interpret $[\underline{E}_*(X), \bar{E}_*(X)]$ as an interval estimate for the ‘‘true’’ lower prevision $\underline{P}_T(X)$, and $[\underline{E}^*(X), \bar{E}^*(X)]$ as an interval estimate for the ‘‘true’’ upper prevision $\bar{P}_T(X)$. The next proposition shows that it makes sense to take the most conservative point-estimates $\underline{E}_*(X)$ and $\bar{E}^*(X)$ as a first order aggregate.

Proposition 2. *\underline{E}_* is a coherent lower prevision and \underline{E}^* is a coherent upper prevision. Moreover, for every gamble X it holds that $\underline{E}_*(X) = -\bar{E}^*(-X)$, that is, \bar{E}^* is the conjugate of \underline{E}_* .*

An alternative and theoretically perhaps more appealing interpretation of \underline{E}_* can be obtained as follows. We can interpret X_* as a coherent conditional lower prevision on $\mathcal{L}(\Omega)$,

$$\underline{F}(X|\underline{P}) = X_*(\underline{P}).$$

By the marginal extension theorem [19], \underline{E}_* is then exactly the natural extension of $\underline{F}(\cdot|\cdot)$ and $\underline{E}(\cdot)$:

$$\underline{E}_*(X) = \underline{E}(\underline{F}(X|\underline{P})).$$

This line of reasoning also proves the coherence of \underline{E}_* .

6. MAIN RESULT

A problem with Eq. (3) to calculate \underline{E} , and hence, also to calculate \underline{E}_* , $\underline{\mathfrak{t}}$, $\bar{\mathfrak{t}}$, $\underline{\mathfrak{t}'}$ and $\bar{\mathfrak{t}'}$, is that it involves a linear inequality for each lower prevision $\underline{P} \in \underline{\mathcal{P}}(\Omega)$. Hence, to calculate \underline{E} we need to solve a linear program with an infinite number of inequalities. The following theorem establishes that in case of upper dual trust equal to one, it takes a linear program with only a finite number of inequalities to calculate $\underline{E}_*(X)$ for every gamble X , and $\underline{\mathfrak{t}}(\underline{P})$ and $\bar{\mathfrak{t}}'(\underline{P})$ for every coherent lower prevision \underline{P} (by calculating $\underline{E}(Z)$ and taking for Z either X_* , $I_{\mathcal{M}(\underline{P})}$ or $1 - I_{\mathcal{N}(\underline{P})}$). Indeed, for those cases the collection of all constraints in Eq. (3) are implied by only a finite subset of them.

The reason why the second-order gambles X_* , $I_{\mathcal{M}(\underline{P})}$ and $1 - I_{\mathcal{N}(\underline{P})}$ are so special is that they are *monotonically increasing*. Note that a second-order gamble Z is said to be monotonically increasing if for any two coherent lower previsions $\underline{P}_1, \underline{P}_2 \in \underline{\mathcal{P}}(\Omega)$ it holds that

$$\underline{P}_1(X) \geq \underline{P}_2(X) \text{ for all gambles } X \in \mathcal{L}(\Omega) \implies Z(\underline{P}_1) \geq Z(\underline{P}_2).$$

It is convenient to introduce the following notation. We denote the conjunction of a subset $J \subseteq N$ of agents by

$$\underline{E}_J(X) = (\bigcap_{k \in J} \underline{P}_k)(X).$$

Of course, in case the assessments \underline{P}_k for $k \in J$ are conflicting, \underline{E}_J does not exist. We also define $\underline{E}_\emptyset(X) = \inf_{\omega \in \Omega} X(\omega)$, the vacuous lower prevision on Ω . We have the following result.

Theorem 1. *Suppose that $Z \in \mathcal{L}(\underline{\mathcal{P}}(\Omega))$ is monotonically increasing and $\bar{\mathfrak{t}}'_k = 1$ for every $k \in N$. Then*

$$\begin{aligned} \underline{E}(Z) = \sup \left\{ \alpha \in \mathbb{R} : (\exists \kappa_k \geq 0, \lambda_k \geq 0, \mu_k \geq 0)(\forall J \subseteq N) \right. \\ \left. \underline{E}_J \text{ exists} \implies Z(\underline{E}_J) - \alpha \geq \sum_{k \in N} \kappa_k \left(I_{\underline{E}_J \geq \underline{E}_k} - \underline{\mathfrak{t}}_k \right) + \lambda_k \left(\bar{\mathfrak{t}}_k - I_{\underline{E}_J \geq \underline{E}_k} \right) \right. \\ \left. + \mu_k \left(I_{\underline{E}_k \geq \underline{E}_J} - \underline{\mathfrak{t}}'_k \right) \right\} \end{aligned}$$

Proof. If $\bar{\mathfrak{t}}'_k = 1$ for all k then the supremum in Eq. (3) will be achieved for $\nu_k = 0$, and hence, we may omit these terms.

Next, we show that for any $\underline{P} \in \underline{\mathcal{P}}(\Omega)$ we can find a $J \subseteq N$ such that \underline{E}_J exists and

$$Z(\underline{E}_J) \leq Z(\underline{P}), \quad I_{\underline{E}_J \geq \underline{E}_k} = I_{\underline{P} \geq \underline{E}_k}, \quad \text{and} \quad I_{\underline{E}_k \geq \underline{E}_J} \geq I_{\underline{E}_k \geq \underline{P}},$$

for all $k \in N$. In such a case the inequality for \underline{P} in Eq. (3) is implied by the inequality for \underline{E}_J in Eq. (3), and hence, we may ‘replace’ \underline{P} by \underline{E}_J in Eq. (3), establishing the proof.

Choose $J = \{k: \underline{E}_k \leq \underline{P}\}$. Observe that $\underline{E}_J = \bigcap_{k \in J} \underline{P}_k$ always exists—if $J = \emptyset$ then \underline{E}_J is the vacuous lower prevision.

Also observe that $\underline{E}_J \leq \underline{P}$, and hence, it immediately follows that $Z(\underline{E}_J) \leq Z(\underline{P})$ since Z is monotone, and $I_{\underline{E}_J \geq \underline{E}_k} \leq I_{\underline{P} \geq \underline{E}_k}$ and $I_{\underline{E}_k \geq \underline{E}_J} \geq I_{\underline{E}_k \geq \underline{P}}$ for every $k \in N$, since

$$\begin{aligned} \underline{E}_J \geq \underline{E}_k &\implies \underline{P} \geq \underline{E}_k, \\ \underline{E}_k \geq \underline{P} &\implies \underline{E}_k \geq \underline{E}_J. \end{aligned}$$

We are left to show that

$$\underline{P} \geq \underline{E}_k \implies \underline{E}_J \geq \underline{E}_k,$$

which would establish $I_{\underline{E}_J \geq \underline{E}_k} = I_{\underline{P} \geq \underline{E}_k}$. Indeed, suppose that $\underline{P} \geq \underline{E}_k$. This means that $k \in J$. Since $\underline{E}_J \geq \underline{E}_j$ for all $j \in J$ by definition of \underline{E}_J , we indeed find that $\underline{E}_J \geq \underline{E}_k$. \square

We must require that $\bar{\mathfrak{t}}_k$ is 1 for all $k \in N$ because in general it is impossible to establish that for every \underline{P} there is a \underline{E}_J such that

$$Z(\underline{E}_J) \leq Z(\underline{P}), \quad I_{\underline{E}_J \geq \underline{E}_k} = I_{\underline{P} \geq \underline{E}_k}, \quad I_{\underline{E}_k \geq \underline{E}_J} = I_{\underline{E}_k \geq \underline{P}},$$

for all $k \in N$. For example, take $n = 1$ and any $\underline{P} \not\leq \underline{P}_1$. For every choice of J , that is, $\underline{E}_J = \underline{P}_1$ or $\underline{E}_J = \inf_{\omega \in \Omega}$, it *cannot* hold that

$$\underline{P}_1 \geq \underline{E}_J \implies \underline{P}_1 \geq \underline{P}.$$

There does not seem to exist an efficient method for calculating neither the upper trust $\bar{\mathfrak{t}}(\underline{P})$ nor the lower dual trust $\mathfrak{t}'(\underline{P})$ associated with a general first-order model \underline{P} —since the second-order gambles involved are not monotonically increasing Theorem 1 does not apply. Perhaps a different choice of constraints might solve the problem (*e.g.*, see [15] for a solution in case of a precise second-order possibility space).

6.1. A generalised conjunction rule and an operational interpretation.

Theorem 1 shows that in some cases the natural extension can be calculated by solving a finite linear program (an infinite number of linear inequalities reduces to a finite number of linear inequalities). In its dual form, this linear program has a very nice form and provides us with an operational interpretation of lower and upper trust, and lower dual trust. It is given in the next theorem.

Theorem 2 (Dual form). *Suppose that $Z \in \mathcal{L}(\mathcal{P}(\Omega))$ is monotonically increasing and $\bar{\mathfrak{t}}'_k = 1$ for every $k \in N$. Define for any $J \subseteq N$ such that \underline{E}_J exists a non-negative variable α_J . Then*

$$(4) \quad \underline{E}(Z) = \min \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists}}} \alpha_J Z(\underline{E}_J),$$

where the variables α_J are subject to

$$\begin{aligned} \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists}}} \alpha_J &= 1, & \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists} \\ \underline{E}_k \geq \underline{E}_J}} \alpha_J &\geq \underline{\mathfrak{t}}'_k, \\ \underline{\mathfrak{t}}_k &\leq \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists} \\ \underline{E}_J \geq \underline{E}_k}} \alpha_J &\leq \bar{\mathfrak{t}}_k, \end{aligned}$$

for all $k \in N$. (Notice that the constraints do not depend on the gamble Z). There is second-order consistency if and only if the above system of constraints has a feasible solution.

For calculating for instance $\underline{E}_*(X)$ for some gamble X , we should minimise

$$(5) \quad \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists}}} \alpha_J \underline{E}_J(X)$$

subject to the above constraints, and for $\bar{E}^*(X)$ we should maximise

$$(6) \quad \sum_{\substack{J \subseteq N \\ \underline{E}_J \text{ exists}}} \alpha_J \bar{E}_J(X),$$

subject to the same constraints. But, these equations tell us that the first-order aggregate is a convex combination of all possible conjunctions \underline{E}_J , and the coefficients α_J of this convex combination—which may depend on the first-order gamble X —can be interpreted as frequencies at which we are willing to choose the conjunction \underline{E}_J among all possible conjunctions. If we are willing to buy X for a price $\underline{E}_J(X)$ —whenever it exists—at rate α_J , then we should be willing to buy X for a price given by Eq. (5). Hence, the coefficients α_J have a simple operational interpretation.

This leads naturally to an operational interpretation of lower and upper trust, and lower dual trust. Looking at the constraints in Theorem 2 we see that the lower trust $\underline{\mathfrak{t}}_k$ is a lower bound for the sum of frequencies α_J for which \underline{E}_J dominates \underline{E}_k , that is, for which any decision of expert k is also a decision of at least one expert in J . Thus, through the frequencies α_J , we have an operational interpretation for lower trust. A similar argument shows that also upper trust and lower dual trust can be given an operational meaning.

Finally, the interpretations given here show that for instance lower trust assigned to an expert is not a numerical property of this expert only. Rather, it is a property of the expert within a given group of experts.

6.2. Special case: lower trust only. Restricting to lower trust only, it is easy to obtain the following results (the lower trust assignments are assumed to be non-negative):

- (i) If lower trust is equal to one for all experts, then the first order aggregate is equal to the conjunction of all the experts, and second-order consistency is equivalent with consistency. We thus obtain the conjunction rule as a special case.
- (ii) If a lower trust model is second-order consistent, then it will remain so for any lower assignment of lower trust.
- (iii) A lower trust assignment such that $\sum_{k \in N} \mathbf{t}_k \leq 1$ is always second-order consistent: in that case, a first order aggregate always exists.
- (iv) If all experts are pair-wise conflicting, that is, if there are no conjunctions except for the trivial ones, then any lower trust assignment such that $\sum_{k \in N} \mathbf{t}_k > 1$ is second-order conflicting. If $\sum_{k \in N} \mathbf{t}_k \leq 1$, then there is second-order consistency, and the aggregate is given by

$$\underline{E}_*(X) = \sum_{k \in N} \mathbf{t}_k \underline{E}_k(X) + \left(1 - \sum_{k \in N} \mathbf{t}_k\right) \inf_{\omega \in \Omega} X(\omega).$$

Thus in case of total conflict (which is quite common if all experts use a single probability measure to represent their knowledge), the model produces a linear opinion pool mixed with a vacuous lower prevision.

These results show that the highest possible assignments for lower trust measure the amount of conflict between the expert assessments. If they can be chosen maximal, all equal to one, then no conflict is present. If they cannot even be chosen such that their sum is larger than one, then there is total conflict.

7. A NUMERICAL EXAMPLE

Let's now illustrate our ideas with a simple example. Suppose we have a system that can switch between 20 possible states ω . We could think for instance of the twenty lowest quantum states of a hydrogen atom (we assume that higher states do not occur):

$$\Omega = \{1, 2, 3, \dots, 20\}$$

The natural number ω is called the principle quantum number. If the hydrogen atom has principle quantum number ω , its energy level is given by $X(\omega) = -\frac{13.6 \text{ eV}}{\omega^2}$: the energy level X of the hydrogen atom is a gamble on Ω (expressed in units of eV). Assume we have three experts, 1, 2 and 3, judging over the quantum state of the hydrogen atom. They make the following, very weak assessments:

1. The probability that $\omega \leq 8$ is at least 0.4.

2. The probability that $\omega \geq 4$ is at most 0.7.
3. The mean energy is at least -0.24 eV.

Many probability measures are compatible with each one of these statements separately. But, intuitively it is clear that no probability measure is compatible with all of them. In particular expert 2 and expert 3 appear to be conflicting. Indeed, principle quantum numbers strictly less than 4, to which a probability mass of at least 0.3 is assigned by the second expert, correspond to energy levels less than -1.5 eV; this is considerably lower than the value -0.24 eV assessed by the third expert. Nevertheless, from these assessments, we would like to find lower and upper bounds on the mean value and the standard deviation of the energy level X and the principle quantum number ω (we could also calculate bounds for other moments). Before performing any calculation, we should perhaps note that we do not expect the aggregate to have very tight bounds, because of the weakness of each of the assessments (the aggregate may quite precise of course, but we can't know this a priori). In more realistic examples, where each expert provides for instance bounds on a larger set of events or a larger set of gambles, we may expect tighter bounds. In terms of lower previsions, we have:

$$\begin{aligned}\underline{P}_1(I_{\{\omega \leq 8\}}) &= 0.4 \\ \underline{P}_2(-I_{\{\omega \geq 4\}}) &= -0.7 \\ \underline{P}_3(X) &= -0.24\end{aligned}$$

Note that all of these lower previsions are coherent.

We now prove that the second and the third expert are conflicting. Perhaps the easiest way to see the conflict is by solving the linear program described in Section 2.1 with $\underline{P} = \underline{P}_2$ in Eq. (1). We find that the second expert is willing to sell X for any price strictly larger than $\bar{E}_2(X) \approx -0.48$, for instance, for $s = -0.47$. But since $\underline{P}_3(X) = \underline{E}_3(X) = -0.24$, this selling price s is lower than for instance the buying price $b = -0.25$ of the third expert for X . Combining these two experts we can create a money pump (or a ‘‘Dutch book’’) by repeatedly buying X for $b = -0.25$ and selling it again for a lower price $s = -0.47$ yielding a net sure loss of $b - s = 0.22$ in each transaction. This points to a conflict between the second and the third expert. Mathematically, the conflict follows from Eq. (2) for $N = \{2, 3\}$, $X_2 = -X$ and $X_3 = X$. Obviously, this also implies conflict for $N = \{1, 2, 3\}$.

These are the only conflicts. Indeed, solving the linear program described in Proposition 1 for $N = \{1, 2\}$ and $N = \{1, 3\}$ with for instance X equal to the zero gamble, it is easily shown that the first expert is consistent both with the second and the third expert: in all cases the maximum α^* is finite (and equal to zero)—there are no money pumps.

After further inspection, assume we assign lower and upper trust as follows:

$$[\underline{t}_1, \bar{t}_1] = [0.3, 0.5]$$

$$[\underline{t}_2, \bar{t}_2] = [0.4, 0.6]$$

$$[\underline{t}_3, \bar{t}_3] = [0.5, 1.0]$$

We know nothing about dual trust: $\underline{t}'_k = 0$ and $\bar{t}'_k = 1$ for all $k \in \{1, 2, 3\}$. This imprecise second-order hierarchical model has the following solution for the lower bound of the mean energy level (Theorem 2):

$$\begin{aligned} \underline{E}_*(X) = \min \left(\alpha_\emptyset \inf_{\omega \in \Omega} X(\omega) + \alpha_{\{1\}} \underline{E}_1(X) + \alpha_{\{2\}} \underline{E}_2(X) + \alpha_{\{3\}} \underline{E}_3(X) \right. \\ \left. + \alpha_{\{1,2\}} (\underline{E}_1 \sqcap \underline{E}_2)(X) + \alpha_{\{1,3\}} (\underline{E}_1 \sqcap \underline{E}_3)(X) \right), \end{aligned}$$

where the non-negative variables α_J are subject to the constraints

$$\begin{aligned} \alpha_\emptyset + \alpha_{\{1\}} + \alpha_{\{2\}} + \alpha_{\{3\}} + \alpha_{\{1,2\}} + \alpha_{\{1,3\}} &= 1, \\ 0.3 \leq \alpha_{\{1\}} + \alpha_{\{1,2\}} + \alpha_{\{1,3\}} &\leq 0.5, \\ 0.4 \leq \alpha_{\{2\}} + \alpha_{\{1,2\}} &\leq 0.6, \\ 0.5 \leq \alpha_{\{3\}} + \alpha_{\{1,3\}} &\leq 1.0. \end{aligned}$$

Obviously, $\inf_{\omega \in \Omega} X(\omega) = -13.6$. Through Proposition 1, we find

$$\begin{aligned} \underline{E}_1(X) &= \underline{E}_2(X) = -13.6, \\ \underline{E}_3(X) &= -0.24, \\ (\underline{E}_1 \sqcap \underline{E}_2)(X) &= -13.6, \\ (\underline{E}_1 \sqcap \underline{E}_3)(X) &= -0.24. \end{aligned}$$

The linear program has a solution, so there is second-order consistency. The solution is $\underline{E}_*(X) = -6.92$. For the upper bound, we solve a similar linear program, but now maximising and using the upper previsions. Obviously, $\bar{E}_\emptyset(X) = \sup_{\omega \in \Omega} X(\omega) = -0.034$, and

$$\begin{aligned} \bar{E}_1(X) &= -\underline{E}_1(-X) \approx -0.11, \\ \bar{E}_2(X) &= -\underline{E}_2(-X) \approx -0.48, \\ \bar{E}_3(X) &= -\underline{E}_3(-X) \approx -0.034, \\ (\bar{E}_1 \sqcap \bar{E}_2)(X) &= -(\underline{E}_1 \sqcap \underline{E}_2)(-X) \approx -0.49, \\ (\bar{E}_1 \sqcap \bar{E}_3)(X) &= -(\underline{E}_1 \sqcap \underline{E}_3)(-X) \approx -0.11. \end{aligned}$$

The solution is $\bar{E}^*(X) \approx -0.22$.

To have an idea of what the assessments tell us about the statistical variance of the energy level ω , we calculate the so-called *lower and upper variances* $\underline{\sigma}^2$ and $\bar{\sigma}^2$

of X under \underline{E}_* [19, Appendix G]:

$$\underline{\sigma}^2(X) = \min_{\mu \in \mathbb{R}} \underline{E}_*((X - \mu)^2), \quad \bar{\sigma}^2(X) = \min_{\mu \in \mathbb{R}} \bar{E}^*((X - \mu)^2)$$

The lower variance $\underline{\sigma}^2(X)$ is the supremum buying price that we are willing to pay for all gambles $(X - \mu)^2$, and the upper variance $\bar{\sigma}^2(X)$ is the infimum selling price we are willing to sell some gamble $(X - \mu)^2$ for. In fact, these bounds coincide with the minimal and maximal variance σ^2 of X under all probability measures that are compatible with the aggregate \underline{E}_* . After some calculations similar to the ones above, we find for the lower variance $\underline{\sigma}^2(X) \approx 0.16$ and the upper variance $\bar{\sigma}^2(X) \approx 46.0$.

In conclusion, from the expert information and the trust assignments it follows that the mean energy level of the hydrogen atom is approximately between -6.92 eV and -0.22 eV, with standard deviation approximately between 0.39 eV and 6.8 eV. Similar calculations can be done for the principle quantum number ω , but let's simply give the result: the expected value of the principle quantum number is approximately between 4.3 and 17.6 , with standard deviation approximately between 1.5 and 9.5 . Note that all these bounds are rather imprecise, but this was to be expected due the weakness of the expert judgements \underline{P}_1 , \underline{P}_2 and \underline{P}_3 . In general, the method cannot reduce imprecision that is present in all of the expert judgements (and arguably, this is how we should prefer it).

8. DISCUSSION AND CONCLUSION

A second-order imprecise probability model was proposed based on a behavioural notion of trust. As most second-order hierarchical models, the interpretation of this model relies on the existence of a hypothetical “representative” expert. Unluckily, at first sight this leads to philosophical as well as to practical problems. The second-order gambles that were used to derive the aggregate are defined on a possibility space that cannot always be sampled in a meaningful way. How should one deduce the second-order lower and upper trust and dual trust values? One could argue that the model should only be used in applications where the representative expert can be identified (for instance, one choice could be identifying the representative expert with the modeller itself).

One important conclusion of this paper is that, by expressing the aggregation algorithm in its dual form (Theorem 2), the proposed aggregation method can be explained as a generalised conjunction rule. Moreover, in doing this, we *do* find an elegant operational interpretation of the imprecise second-order assessments, and thus overcome the above-mentioned problems of interpretation that are so common in many hierarchical uncertainty models.

The method is both mathematically simple, quite general, and practically appealing, especially if only a limited number of expert opinions need to be aggregated. In the general case however, it cannot be excluded that the size of the

linear program to be solved will grow exponentially in the number of experts, limiting the applicability of the model. However, this will only occur if the number of subsets $J \subseteq N$ for which \underline{E}_J exists grows exponentially too. In the extreme case where there is always total conflict, the size of the linear program grows only linearly in the number of experts. Secondly, we note that the linear inequalities in Theorem 2 contain mostly zeros and ones, so perhaps there is an efficient way to deal with linear programs of this type (we have not investigated this further). Finally, extending the aggregation algorithm to upper dual trust less than one and calculating $\bar{\mathfrak{t}}(\underline{P})$ and $\underline{\mathfrak{t}}(\underline{P})$ remain the subject of further research.

Acknowledgements. The author is indebted to Gert de Cooman, Lev Utkin and Igor Kozine for stimulating discussions and constructive comments, and two anonymous referees for valuable remarks that greatly improved the readability of this paper. Finally, the author also wishes to thank Lyudmila Mihaylova, Sofie Troffaes and Hektor Letraublou for reading and helping with an earlier draft of this paper.

This paper presents research results of project G.0139.01 of the Fund for Scientific Research, Flanders (Belgium), and of the Belgian Programme on Inter-university Poles of Attraction initiated by the Belgian state, Prime Minister's Office for Science, Technology and Culture. The scientific responsibility rests with the author.

REFERENCES

- [1] P. Artzner, F. Delbaen, J. M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.
- [2] G. Choquet. Theory of capacities. *Annales de l'Institut Fourier*, 5:131–295, 1953–54.
- [3] Gert de Cooman. Precision-imprecision equivalence in a broad class of imprecise hierarchical models. *Journal of Statistical Planning and Inference*, pages 175–198, June 2002.
- [4] Gert de Cooman and Matthias C. M. Troffaes. Coherent lower previsions in systems modelling: products and aggregation rules. *Reliability Engineering and System Safety*, 85:113–134, 2004.
- [5] Gert de Cooman and Peter Walley. A possibilistic hierarchical model for behaviour under uncertainty. *Theory and Decision*, 52(4):327–374, 2002.
- [6] Bruno de Finetti. Sul significato soggettivo della probabilità. *Fundamenta Mathematicae*, 17:298–329, 1931.
- [7] Bruno De Finetti. *Theory of Probability: A Critical Introductory Treatment*. Wiley, New York, 1974–5. Two volumes.
- [8] A. P. Dempster. Upper and lower probabilities induced by a multivalued mapping. *Ann. Math. Statist.*, 38:325–339, 1967.
- [9] Christian Genest and James V. Zidek. Combining probability distributions: A critique and an annotated bibliography. *Statistical Science*, 1(1):114–148, 1986.
- [10] Isaac Levi. *The Enterprise of Knowledge. An Essay on Knowledge, Credal Probability, and Chance*. MIT Press, Cambridge, 1983.
- [11] S. Moral and J. del Sagrado. Aggregation of imprecise probabilities. In B. Bouchon-Meunier, editor, *Aggregation and Fusion of Imperfect Information*, pages 162–188. Physica-Verlag, New York, 1998.

- [12] Robert F. Nau. The aggregation of imprecise probabilities. *Journal of Statistical Planning and Inference*, 105(1):265–282, June 2002.
- [13] Christian P. Robert. *The Bayesian Choice*. Springer, 1994.
- [14] G. Shafer. *A Mathematical Theory of Evidence*. Princeton University Press, 1976.
- [15] Lev V. Utkin. Imprecise second-order hierarchical uncertainty model. *Uncertainty, Fuzziness and Knowledge-Based Systems*, 11(2):301–317, June 2003.
- [16] P. Walley. The elicitation and aggregation of beliefs. Technical report, University of Warwick, Coventry, 1982. Statistics Research Report 23.
- [17] P. Walley. Measures of uncertainty in expert systems. *Artificial Intelligence*, 83:1–58, 1996.
- [18] P. Walley. Statistical inferences based on a second-order possibility distribution. *International Journal of General Systems*, 26:337–383, 1997.
- [19] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
- [20] N. Wilson and S. Moral. Fast Markov chain algorithms for calculating Dempster-Shafer belief. In W. Wahlster, editor, *Proceedings of the 12th European Conference on Artificial Intelligence*, pages 672–676. J. Wiley, 1996.
- [21] L. A. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets Syst.*, 1:3–28, 1978.