Timothy Childers
Ondrej Majer

Eliciting Uncertainties: A Two Structure Approach

Abstract. We recast subjective probabilities by rejecting behaviourist accounts of belief by explicitly distinguishing between judgements of uncertainty and expressions of those judgements. We argue that this entails rejecting that orderings of uncertainty be complete. This in turn leads naturally to several generalizations of the probability calculus. We define probability-like functions over incomplete algebras that reflect a subject's incomplete judgements of uncertainty. These functions can be further generalized to (partial) inner and outer measures that reflect approximate elicitations.

Keywords: Subjective probability, Elicitation, Foundations of probability, Partial probability functions.

1. Introduction

A calculus for degrees of belief should determine how partial beliefs can be singly and jointly well-proportioned: given a specification of degrees of belief over some propositions, the degrees of beliefs of other propositions should be calculable. The calculus would serve as a key to rationality, as it would show how to properly adjust beliefs in light of new beliefs arising from evidence. The question then arises of how to determine particular degrees of belief.

The most influential proposal comes from Ramsey: potential behaviour can be employed by taking the degree of a belief to be "a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it". Ramsey [16, p. 169]. This account of partial belief also provides a standard for rationality by placing a subject in a gambling situation: if belief leads to action, then, in certain situations, having beliefs can lead to offering publicly available betting quotients. But, of course, if degree of belief causes action, then it should be harnessed to our desired ends—"Reason is, and ought only to be the slave of the passions…". However, offering betting quotients that automatically lead to loss (technically: betting quotients that are not fair) leads to undesirable outcomes—sure losses of money. To avoid these undesirable outcomes it is necessary and sufficient, by the Ramsey-de Finetti theorem, to conform our public actions in terms of offering betting

Presented by **David Makinson**; *Received* May 25, 2017

Studia Logica (2018) 106: 615–636
https://doi.org/10.1007/s11225-017-9757-5

quotients, and hence our degrees of belief, to the probability calculus. Thus consistency of partial belief becomes consistency of action by conformation to the probability calculus.[1]

As is well known, accounts of belief based on interpreting potential action face severe difficulties. Chisholm [1] and Geach [7] famously pointed out that beliefs cannot be correlated with actions unless we know that a subject also has certain desires. Geach succinctly puts the objection thus: "But is there in fact any behaviour characteristic of a given belief? Can action be described as 'acting as if you held such-and-such a belief' unless we take for granted, or are somehow specially informed about, the needs and wants of the agent? . . . When Dr. Johnson did penance in Uttoxeter market-place, he may have begun by standing around bareheaded until the threatened shower should fall; this would not be recognizable as rain-expecting behaviour without a knowledge of Johnson's wish to do penance." Geach [7, p. 8].

This objection applies equally to behavioural accounts of partial belief. As is well known, attitudes towards gambles are notoriously sensitive to attitudes towards risk (the Allais 'paradox' being perhaps the most famous example). Responses to proposed bets change with the size of the stakes, a result of the non-linear utility of currency. The possibility of large losses lead the risk averse to hedge their bets; but if the losses are too small elicitation may be too much of a bother, and so will be inaccurate. These problems are particularly acute in the case of sets of gambles, as Schick [18] has made forcefully clear. This means that elicited betting quotients may have little to do with degrees of belief. Such a link requires knowledge of a subject's other attitudes, and so cannot be determined by behaviour alone (Christensen [2] also makes this point).

Attempts to liberalize the behavioural account face a dilemma: if we abstract too far from actual behaviour it is hard to see what force any arguments based on it could have; but if too close to actual behaviour it will face the usual objections. For example, Christensen [2] proposes a "depragmatized" version of the Dutch Book that relies on a "simple agent", one so simple it does not have interfering beliefs and dispositions, to highlight a pre-theoretic link between fair betting odds and beliefs. But it is hard to see what lessons can be learned from agents with so little in common with us, since we are hardly simple. Or perhaps a subject need not actually enter into a bet, which would allow us to bracket their other dispositions. For example, Howson and Urbach [9] consider a counterfactual reading of the Dutch Book

[1] Ramsey [16, p. 172], regards the use of bets as measures of partial belief to be "fundamentally sound", but "insufficiently general and . . . inexact".

argument where we consider a more pliable counterpart of our subject: we consider what odds she gives in the nearest possible world in which she does bet. But if the opposition to gambling forms a deeply important part of her personality, information about how she would bet at some distant world need have nothing to do with her beliefs in the actual world.[2]

It should come as no surprise that arguments linking behaviour to undesirable consequences must take into account more than publicly observable behaviour—they must also take into account a subject's other beliefs and desires. To stress: there is no getting around the Chisholm–Geach problem. The lesson we should draw from this is not that any attempt to elicit probabilities is hopeless, but that we must take into account a subject's internal judgements of uncertainty along with their public expressions of those judgements.

Cognitive science has long abandoned any attempt to exclude internal mental states, and we should as well. In what follows we take on the task of (re)describing arguments for partial belief to conform to (something like) the probability calculus: we consider a subject with a certain structure of judgements of uncertainty, a means of expressing elements of that structure publicly, and a way of relating the two. The arguments then establish the conditions under which the structure of judgements and of expressions of those judgements are in harmony.

In the next section, for the purpose of developing our formal framework, we will assume that judgements and expressions match (as they are assumed to in traditional arguments for subjective probabilities). This is, of course, a very strong condition. The following sections concern some cases when we relax the requirement that the two match. In Section 3, we lessen the burden of elicitation on the subject by requiring only partial elicitation. In Section 4 we consider approximate elicitations.

2. Eliciting Uncertainty

Standard cognitive science takes subjects to have internal structured representations over which the mind operates, belief being one of those operations. We will approach the foundations of subjective probability in this spirit, and so this paper is organized as follows: in this section we specify the structure of representations, and the subject's judgements of uncertainty with respect

[2] Eriksson and Hajek [4] offer much the same criticisms aimed at a broader range of behavioural accounts of belief.

to those representations. We then discuss how to model the elicitation of those judgements. In the following sections we explore the implications of this picture, which lead naturally to generalizations of probability measures.[3]

We begin with modelling the subject's representations: an obvious, and even standard, starting point is to model them as as propositions forming a Boolean algebra, which is essentially the same as assuming the subject uses a classical logic. A more general approach would be to leave open the possibility that the operations of the structure might be non-classical, but this would be beyond what a single paper could cover.

There are of course two broad choices for modelling a subject's representations of uncertainty. The first is attach numbers to each representation, where the numbers could be understood as e.g. betting quotients attached to propositions. The second is to use a comparative notion of likeliness. There are many reasons to prefer the latter over the former, but we need not go into these: we choose the comparative conception it allows us to more naturally develop our account. However, at the end it will become clear that we could also begin with a quantitative conception.

Putting these two pieces—the internal propositional structure and the internal judgements of relative likeliness of those propositions—together completes the subject's internal structure:

Definition 1. (Representational structure—qualitative) A qualitative representational structure is a couple $\langle \mathcal{F}, \precsim \rangle$ where $\mathcal{F} = (F, \mathbf{0}_{\mathcal{F}}, \mathbf{1}_{\mathcal{F}}, \vee, \wedge)$ is a Boolean algebra and $\precsim$ is a total relation on the elements of $\mathcal{F}$, i.e., either $A \precsim B$ or $B \precsim A$ or both.

The symbol $\precsim$ is a primitive relation ordering the subject's degrees of belief. The only restriction on $\precsim$ is that it is defined over all pairs of representations.[4] The representational structure constitutes the subject's (very simple) outlook on the world.

The next step is to detail how the the subject could express internal assessments of likeliness: we use the well-known device of an 'auxiliary' or 'reference experiment', as in von Neumann and Morgenstern [20] or DeGroot [3]. A reference experiment serves as a standard that the subject can employ to express their uncertainty by substituting familiar quantities such as length or area for the less familiar quantity of probability. The quantities used in the experiment depend on what the subject is comfortable using (as long as they can be normalized into probabilities).

[3]This paper thus is in the tradition of Krantz et al. [12].

[4]As usual $A \sim B$ denotes $A \precsim B$ and $B \precsim A$, and $A \prec B$ denotes $A \precsim B$ and $B \not\precsim A$.

Although our treatment will be general, the reader may wish to keep in mind a particular reference device namely, a probability wheel (one classic description of the use of the probability wheel is Spetzler and Holstein [19]). This is a perfectly balanced arrow which serves to indicate points on the circumference of the wheel. The events of the experiment are generated by landings of the arrow in an arc; probabilities are generated by a uniform distribution over the circumference.

The reference experiment provides events that can be put in correspondence to the subject's internal representations in such a way that the probabilities of the reference events correspond to her uncertainty ordering (exactly how they are put in correspondence is the subject of the next section). In other words, a reference experiment is an idealized device fully understood by the subject, in that she accepts that it can be used to express her comparisons of likelihood. The reference experiment thus provides a standard of measurement for the subject's uncertainty:

Definition 2. (Reference structure) a reference structure is a couple $\langle \mathcal{G}, p \rangle$ where $\mathcal{G} = (G, \mathbf{0}_{\mathcal{G}}, \mathbf{1}_{\mathcal{G}}, \vee_{\mathcal{G}}, \wedge_{\mathcal{G}})$ is a Boolean algebra and p is a probability distribution over $\mathcal{G}$.

When there is no risk of ambiguity we omit the subscripts. We denote by $\precsim_{\mathcal{G}}$ an ordering on the elements of $\mathcal{G}$ induced by p. In the case of the probability wheel the ordering is given by the lengths of the arcs.

The subject has some views of the world which they express via the reference experiment. That expression is successful when it is in harmony with the subject's views; elicitation is the process of harmonizing the two. Technically, elicitation is the process of identifying counterparts to propositions in a reference structure while preserving the structure and ordering of the subject's representations:

Definition 3. (Elicitation function—qualitative probability) A function e from a representational structure $\langle \mathcal{F}, \precsim \rangle$ to a reference structure $\langle \mathcal{G}, p \rangle$ is an *elicitation function* iff it satisfies the following conditions:

(i) if $A \precsim B$ then $e(A) \precsim_{\mathcal{G}} e(B)$,
if $A \prec B$ then $e(A) \prec_{\mathcal{G}} e(B)$ (monotonicity),

(ii) $e(A \wedge B) = e(A) \wedge e(B)$,
$e(A \vee B) = e(A) \vee e(B)$ (structure preservation),

(iii) $e(\mathbf{1}_{\mathcal{F}}) = \mathbf{1}_{\mathcal{G}}$,
$e(\mathbf{0}_{\mathcal{F}}) = \mathbf{0}_{\mathcal{G}}$ (scale).

Definition 3 ensures the reference structure is rich enough to provide the subject with means equal to expressing her representations. Technically, as elicitation is a mapping of the representational structure into the reference structure, it is a homomorphism, not an isomorphism.

From the definitions it follows immediately that successful elicitation yields probabilities: that is, if the subject faithfully expresses herself using the reference experiment her internal judgements of uncertainty are probabilities:

THEOREM 2.1. (Quantitative probability from elicitation) *Let* $\langle \mathcal{F}, \precsim \rangle$ *be a representational structure and* $\langle \mathcal{G}, p \rangle$ *a reference structure. If an elicitation function from* $\mathcal{F}$ *to* $\mathcal{G}$ *exists then there is probability function* f *on* $\mathcal{F}$ *preserving* $\precsim$*, i.e. a function such that for* $A, B \in \mathcal{F}$*:*

1. *if* $A \precsim B$*, then* $f(A) \leq f(B)$*,*
2. $f(A) \geq 0$*,*
3. $f(\boldsymbol{1}_{\mathcal{F}}) = 1$*,*
4. $f(A \vee B) = f(A) + f(B)$*, if* $A \wedge B = \boldsymbol{0}_{\mathcal{F}}$*.*

Moreover f *is unique for a given elicitation function* e*.*

PROOF OF THEOREM 2.1. Let e be an elicitation function from $\mathcal{F}$ to $\mathcal{G}$, and let $f(A) = p(e(A))$. Since p is fixed, f is unique for a given elicitation e.

1. By condition (i) that elicitation preserves the qualitative ordering and the fact that $\precsim_{\mathcal{G}}$ is induced by p, if $A \precsim B$, then $e(A) \precsim_{\mathcal{G}} e(B)$ and $p(e(A)) \leq p(e(B))$, i.e., $f(A) \leq f(B)$.
2. p is a probability function, and so is non-negative. Hence $0 \leq p(e(A)) = f(A)$.
3. By condition (iii), $e(\mathbf{1}_{\mathcal{F}}) = \mathbf{1}_{\mathcal{G}}$. Moreover since p is a probability function, $p(\mathbf{1}_{\mathcal{G}}) = 1 = p(e(\mathbf{1}_{\mathcal{F}})) = f(\mathbf{1}_{\mathcal{F}})$.
4. Since A, B are disjoint, $e(A), e(B)$ are as well. By (ii) and (iii): $e(\mathbf{0}_{\mathcal{F}}) = e(A \wedge B) = e(A) \wedge e(B) = \mathbf{0}_{\mathcal{G}}$. By the probability calculus, $p(e(A) \vee e(B)) = p(e(A)) + p(e(B)) = f(A) + f(B)$. ■

We have proceeded by using elicitation via a reference structure to determine that the subject's representational structure must have associated probabilities. This places the burden on the reference structure and the elicitation relation, on which representation theorems concerning the subject's internal structure will be based. Taking the subject's judgements and their expression of those judgements to be two separate structures related by elicitation is in opposition to the usual approach, where the aim is to impose conditions on the subject's internal representations to ensure that their orderings will have a unique corresponding probability. An example is the use of an auxiliary experiment (e.g., von Neumann and Morgenstern [20]) or a flat reference distribution (e.g., DeGroot [3] or Savage [17]) to supplement a subject's internal structure with events that serve as a measurement standard. This provides a formally simple framework with a single structure and a single qualitative ordering. When the subject considers a reference experiment she can be interpreted as considering a set of events to which she is willing to assign a uniform probability distribution. As the probability ordering is total, the supplemental events can be compared to every other event. But these new events have the special property that they have an associated quantitative probability distribution. It can then be proved that there is a unique ordering over the subject's internal structure, since the ordering of the supplemental events is a probability ordering. That is, the ordering of the events, in conjunction with full comparability, imposes a total probability ordering on the elements of the subject's internal structure (this is the approach of DeGroot [3]—it is rejected by French [6]).

There are two ways to interpret this procedure. First, for any event A in the (original) representational structure there is a reference event $e(A)$ in the enriched representational structure that is equally likely, that is $e(A) \sim A$. Since the reference event $e(A)$ has by definition some quantitative probability r, we can conclude that the event A has probability r, that is $p(A) = r$. Second, for any representational proposition A, there is always an equivalent reference event $e(A)$, and so we can effectively ignore the representational events. Given the structure of the reference experiment this implies that there is an r such that $p(e(A)) = p(A) = r$, where p is a normalization of the quantities of the reference experiment.

The former is an introspectionist reading which takes the reference structure to be subsumed into the representational structure. This is very strong, requiring the subject be able to contemplate a full probability measure over all propositions. The latter, more mainstream behaviourial reading, does the opposite. The subject's internal representational structure is made entirely public, with representational events identified as observable reference events. This is also a very strong reading—it reduces introspection to expression, and so equates belief with action in exactly the same way that vitiates the Dutch Book argument.

However, note that our formal framework is compatible with both of these interpretations: the requirement of completeness of the subject's uncertainty ordering in Definition 1 ensures that whatever is said about one structure can be said about the other. This clearly shows that the requirement of completeness obscures the distinction between the two structures. Moreover, imposing completeness is against the spirit of our programme, as we aim to place all the requirements on elicitation using a reference structure, leaving the internal uncertainty ordering completely unspecified.

In the next sections we will show what happens when the two structures are kept separate. In so doing we provide natural generalizations of the probability function, while also allowing for the use of other more general representation theorems of qualitative probability orderings without completeness. Having developed our approach fully we will be in a position to see that there are other varieties of elicitation, which will be explored in Section 4.

3. The Partial Subject

In this section we consider a subject upon whom no prior constraints, other than willingness to conform to a reference experiment, are imposed. Obviously, this requires changing Definition 1 to exclude totality.

Definition 4. (Representational structure—partial qualitative ordering) A representational structure with a partial qualitative ordering is a couple $\langle \mathcal{F}, \precsim \rangle$ such that $\mathcal{F}$ is a Boolean algebra and $\precsim$ is a relation on the elements of $\mathcal{F}$ which satisfies $A \precsim A$, $\mathbf{0} \precsim A \precsim \mathbf{1}$ and $\mathbf{0} \prec \mathbf{1}$.

Definition 4 effectively models a subject who does not, for whatever reason, possess judgements of likeliness for all pairs of propositions. This of course leads to significant differences with a total ordering: When the subject provides judgements for *all* pairs of propositions it is a simple matter to determine what properties the internal ordering must have given faithful elicitation using a reference experiment. But if the subject provides only some judgements of uncertainty adjustments will be needed to infer the structure of her representations. As an illustration take the axioms of qualitative probability (as in e.g., Savage [17], Section 3.2):

Definition 5. (Qualitative probabilities) A relation $\precsim$ is a qualitative probability associated with a Boolean algebra $\mathcal{F}$ iff for any $A, B, C \in \mathcal{F}$:

(Pre) Preorder (reflexivity, transitivity)
$A \precsim A$;
If $A \precsim B$ and $B \precsim C$ then $A \precsim C$,

(Tot) Totality
Either $A \precsim B$ or $B \precsim A$ or both,

(Non) Nontriviality
$\mathbf{0} \precsim A$, moreover $\mathbf{0} \prec \mathbf{1}$,

(Add) Additivity (Independence of disjoint events)
If $A \precsim B$ then $A \vee C \precsim B \vee C$, given $A \wedge C = \mathbf{0} = B \wedge C$.

Since elicitation guarantees harmony between the reference and representational expressions, totality ensures all properties of the reference structure are also present in the representational structure. Technically, this follows from the fact that elicitation is a homomorphism and both structures are total. Hence:

LEMMA 3.1. (Qualitative probability from elicitation) *Let $\langle \mathcal{F}, \precsim \rangle$ be a qualitative representational structure. If there is an elicitation function e from $\langle \mathcal{F}, \precsim \rangle$ to a reference structure $\langle \mathcal{G}, p \rangle$ such that e satisfies $(i) - (iii)$ from Definition 3, then $\precsim$ is a qualitative probability.*

PROOF OF LEMMA 3.1

Tot Totality is assumed.

Pre Reflexivity trivially follows from **Tot**. Transitivity: assume for contradiction that $A \precsim B \precsim C$ but not $A \precsim C$, that is $C \prec A$, which by monotonicity of elicitation is equivalent to $e(A) \precsim_{\mathcal{G}} e(B) \precsim_{\mathcal{G}} e(C)$ and $e(C) \prec_{\mathcal{G}} e(A)$, which contradicts the transitivity of $\precsim_{\mathcal{G}}$. Hence it cannot be the case that $C \prec A$, and totality gives us $A \precsim C$ as required.

Non Both $\mathbf{0} \precsim A$ and $\mathbf{0} \prec \mathbf{1}$ follow from **Tot** and (iii).

Add We need to show that if $A \wedge C = \mathbf{0}_{\mathcal{F}} = B \wedge C$, and $A \precsim B$ then $A \vee C \precsim B \vee C$. Assume for contradiction that $B \vee C \prec A \vee C$. Then according to (i) $e(B \vee C) \prec_{\mathcal{G}} e(A \vee C)$ and from (ii) we get $e(A) \wedge e(C) = \mathbf{0}_{\mathcal{G}} = e(B) \wedge e(C)$ and $e(B) \vee e(C) \prec_{\mathcal{G}} e(A) \vee e(C)$. As the ordering $\prec_{\mathcal{G}}$ on the reference structure satisfies the qualitative probability axioms, in particular additivity, we get $e(B) \prec_{\mathcal{G}} e(A)$. But we assumed $A \precsim B$ which implies $e(A) \precsim_{\mathcal{G}} e(B)$. Contradiction. ■

Yet even if the the subject does not provide judgements for all pairs we can still proceed (almost) as in the previous section. The reference structure remains as above. Elicitation is obtained by substituting the representational structure from Definition 4 into Definition 3. The result of Theorem 2.1 then applies as totality is used neither in its formulation or proof.

We can then explicate the properties of the subject's internal ordering (as in Lemma 3.1). Since the order is incomplete it cannot be directly proven that the subject's internal ordering satisfies the axioms of qualitative probability, since it might not exist for some pairs. But it can be shown that the internal ordering *does not violate* the axioms:

LEMMA 3.2. (Partial qualitative probability from elicitation) $\langle \mathcal{F}, \precsim \rangle$ *be a representational structure with a partial qualitative ordering from the previous definition. If there is a function e from* $\langle \mathcal{F}, \precsim \rangle$ *to a reference structure* $\langle \mathcal{G}, p \rangle$*, such that e satisfies conditions* $(i) - (iii)$ *from the definition of an elicitation function, then it holds for any* $A, B, C \in \mathcal{F}$*:*

(Pre′) *Preorder (reflexivity, transitivity)*

$A \precsim A$*;*

If $A \precsim B$ *and* $B \precsim C$ *then not* $C \prec A$*,*

(Non′) *Nontriviality*

$\mathbf{0} \precsim A \precsim \mathbf{1}$*, moreover* $\mathbf{0} \prec \mathbf{1}$*,*

(Add′) *Additivity (independence of disjoint events)*

If $A \precsim B$ *then not* $B \vee C \prec A \vee C$*, given* $A \wedge C = \mathbf{0} = B \wedge C$*.*

PROOF. Reflexivity and nontriviality replace totality in the definition of the partial representational structure. The proofs of the other properties differ only trivially from the proofs in Lemma 3.1. ∎

We stress that while we use qualitative probabilities as illustration, other representation theorems can of course be employed.

We have so far explored the consequences of a partial uncertainty ordering arising from removing all constraints on the subject's uncertainty ordering. There is, however, another constraint we can remove, namely, that the subject's *representations* be complete. This allows us to consider a subject who might have 'isolate' beliefs about tangerines and about tsunamis, but not about tangerines *and* tsunamis. A standard, and obvious, way of modelling this would use sub-algebras of the representational structure.[5] However, the sub-algebra approach imposes much more structure than is necessary. It would be better to allow a subject to have individual views that they are unwilling to combine. To this end, we introduce a much more general approach using structures where the operations are partial, i.e., they may not be defined for all elements of the structure.

[5] As in Koopman [11].

Not just any partial operations will do, of course: they should be compatible with their complete counterparts. We can secure this by requiring that the partial structure be in principal completable, i.e. there is a mapping of the partial structure into a corresponding full structure respecting the partial operations (in this paper we consider only finite cases to avoid unnecessary technical complications):

DEFINITION 6. (Boolean partial algebra) A Boolean partial algebra is a partial structure $\mathcal{F} = (F, \mathbf{1}', \mathbf{0}', \vee', \wedge', ^{c'})$ such that there is a Boolean algebra $\mathcal{S} = (S, \mathbf{1}, \mathbf{0}, \vee, \wedge, ^{c})$ which is a completion of $\mathcal{F}$. A structure $\mathcal{S} = (S, o_1, \ldots, o_n)$ is a *completion* of a partial structure $\mathcal{F} = (F, q_1, \ldots, q_n)$ iff for each i, q_i is an operation of arity k_i on F, o_i is a (complete) operation of arity k_i on S and there is a homomorphism h from $\mathcal{F}$ to $\mathcal{S}$, i.e., for any $x_1, \ldots, x_{k_i} \in \mathcal{F}$ and any q_i, $h(q_i(x_1, \ldots, x_{k_i})) = o_i(h(x_1), \ldots, h(x_{k_i}))$ whenever $q_i(x_1, \ldots, x_{k_i})$ defined.

To give an example of a Boolean partial algebra, consider the structure $\mathcal{F}$ where the only defined operations are for each $A_i, A_i \vee \mathbf{1} = \mathbf{1}, A_i \wedge \mathbf{0} = \mathbf{0}$. This partial algebra is indeed very partial: none of the A_i have complements, everything is disconnected, excepting the top and the bottom elements. Clearly, however, the partial algebra can be completed as, for example, the Boolean algebra generated by taking the A_i to be disjoint (obviously there are many more possible completions).

The internal uncertainty ordering is as in Definition 4, as it must be, since imposing a total order on a partial Boolean structure would constrain the previously undefined operations. For example, suppose that $A \vee B$ is undefined and that the representational structure is totally ordered, e.g. $A \precsim B$. Since the likelihood ordering must respect the algebraic structure, it cannot be the case that $A > B$, because than it must hold $B \prec A$, contrary to what we assumed. This implies that $A \vee B \neq A$ and hence the operation $A \vee B$ is undefined but not unconstrained.

These considerations give us a representational structure in keeping with our programme:

DEFINITION 7. (Partial representational structure) A partial representational structure is a couple $\langle \mathcal{F}, \precsim \rangle$ such that $\mathcal{F}$ is a partial Boolean algebra and $\precsim$ is a relation on the elements of $\mathcal{F}$ satisfying: $A \precsim A$, $\mathbf{0} \precsim A \precsim \mathbf{1}$ and $\mathbf{0} \prec \mathbf{1}$.

The reference structure, Definition 2, remains the same, as it has all the resources needed for expression (we will consider the case where it is not rich enough in the next section). The elicitation function requires the minor modification that elicitation preserves each operation in the representational structure:

Definition 8. (Partial elicitation function) A function e from a partial representational structure $\langle \mathcal{F}, \precsim \rangle$ to a reference structure $\langle \mathcal{G}, p \rangle$ is a *partial elicitation function* iff it satisfies the following conditions:

(i) If $A \precsim B$ then $e(A) \precsim_{\mathcal{G}} e(B)$,
if $A \prec B$ then $e(A) \prec_{\mathcal{G}} e(B)$ (monotonicity),

(ii) $e(A \wedge B) = e(A) \wedge e(B)$,
$e(A \vee B) = e(A) \vee e(B)$,
$e(A^c) = e(A)^c$ if the corresponding operations are defined (structure preservation),

(iii) $e(\mathbf{1}_{\mathcal{F}}) = \mathbf{1}_{\mathcal{G}}$,
$e(\mathbf{0}_{\mathcal{F}}) = \mathbf{0}_{\mathcal{G}}$ (scale).

A substantial generalization of Theorem 2.1 follows:

Theorem 3.1. (Existence of partial probability functions) *Let $\langle \mathcal{F}, \precsim \rangle$ be a partial representational structure, $\langle \mathcal{G}, p \rangle$ a reference structure, and e a partial elicitation function from $\mathcal{F}$ to $\mathcal{G}$. Then there is a* partial probability function *f from $\mathcal{F}$ to* $[0, 1]$ *which is unique for a given e, i.e. a function such that for $A, B \in \mathcal{F}$:*

1. *if $A \precsim B$, then $f(A) \leq f(B)$, if $A \prec B$, then $f(A) < f(B)$,*
2. *$f(A) \geq 0$,*
3. *$f(\mathbf{1}_{\mathcal{F}}) = 1$,*
4. *$f(A \vee B) = f(A) + f(B)$, if both $A \vee B, A \wedge B$ exist and $A \wedge B = \mathbf{0}_{\mathcal{F}}$.*

Proof of Theorem 3.1. The proof is almost the same as that of Theorem 2.1. We again put $f(A) = p(e(A))$. Clearly f is unique.

1. The monotonicity of partial elicitation in Definition 8 is the same as Definition 3, and so the partiality of operations plays no role. Since the elicitation function must preserve the qualitative probability ordering, if $A \precsim B$, then $e(A) \precsim_{\mathcal{G}} e(B)$ and $p(e(A)) \leq p(e(B))$, i.e., $f(A) \leq f(B)$. Similarly in the case $A \prec B$.
2. p is a probability function, and so is non-negative. Hence $0 \leq p(e(A)) = f(A)$.
3. By condition (iii) of Definition 8, $e(\mathbf{1}_{\mathcal{F}}) = \mathbf{1}_{\mathcal{G}}$. Moreover since p is a probability function, $p(\mathbf{1}_{\mathcal{G}}) = 1 = p(e(\mathbf{1}_{\mathcal{F}})) = f(\mathbf{1}_{\mathcal{F}})$.

4. As e commutes with both $\vee$ and $\wedge$, we have $e(A \vee B)) = e(A) \vee e(B)$ and $e(A \wedge B) = e(A) \wedge e(B)$. Hence $f(A \vee B) = p(e(A \vee B)) = p(e(A) \vee e(B))$. Since we assume $A \wedge B = \mathbf{0}_{\mathcal{F}}$, $e(A)$ and $e(B)$ must be disjoint as well: $e(\mathbf{0}_{\mathcal{F}}) = e(A \wedge B) = e(A) \wedge e(B) = \mathbf{0}_{\mathcal{G}}$. By the probability calculus, $p(e(A) \vee e(B)) = p(e(A)) + p(e(B)) = f(A) + f(B)$. ■

Theorem 3.1 establishes that f is a partial version of a standard probability function (and that the representation theorem is a partial version of the standard representation theorems of, e.g., Narens [15, p. 36], Krantz et al. [12, p. 432]). At one extreme, if all operations are defined over the elements of the representational structure, f is a standard probability function. At the other extreme, f is still recognizably a probability-like function, as the following lemma illustrates:

LEMMA 3.3. *Let $\langle \mathcal{F}, \precsim \rangle, \langle \mathcal{G}, p \rangle, e$ be as before and f a partial probability function from Theorem 3.1. Then for $A, B \in \mathcal{F}$:*

1. $f(\boldsymbol{0}_{\mathcal{F}}) = 0$,
2. $f(A \vee B) \leq f(A) + f(B)$ *if* $A \vee B$ *is defined,*
3. $f(A^c) = 1 - f(A)$, *if* A^c *is defined.*

PROOF OF LEMMA 3.3

1. Using the scale condition, $0 = p(\mathbf{0}_{\mathcal{G}}) = p(e(\mathbf{0}_{\mathcal{F}})) = f(\mathbf{0}_{\mathcal{F}})$.
2. If $A \vee B$ is defined then according to condition (ii) for elicitation $e(A \vee B) = e(A) \vee e(B)$. By the probability calculus, $f(A \vee B) = p(e(A \vee B)) = p(e(A) \vee e(B)) \leq p(e(A)) + p(e(B)) = f(A) + f(B)$.
3. By condition (ii) of the definition of partial elicitation function $e(A^c) = e(A)^c$ if A^c is defined. Then it easily follows that $f(A^c) = p(e(A^c)) = p(e(A)^c) = 1 - p(e(A)) = 1 - f(A)$. ■

As before, we can explicate the properties of the subject's uncertainty ordering. For example, as in the previous section, we will now show (the partial versions of) the axioms of qualitative probabilities hold.

By definition, partial representational structures possess Preorder and Nontriviality. Of course, Additivity cannot be proved, so we shall concentrate on a partial version of Additivity:

(Add*) Additivity for partial structures

For any $A, B, C \in \mathcal{F}$ such that $A \wedge C = \mathbf{0}_{\mathcal{F}} = B \wedge C$, if $A \precsim B$ and both $A \vee C, B \vee C$ are defined then it is not the case that $B \vee C \prec A \vee C$.

This change is significant. In the full case the compatibility of the uncertainty ordering with the Boolean is guaranteed by Additivity.[6] This guarantee does not hold in the partial case.[7] We therefore need to establish that the following condition holds:

(Com*) Compatibility for partial structures

For each $A, B \in \mathcal{F}$:

1. if $A \vee B$ is defined, then neither $A \vee B \prec A$ nor $A \vee B \prec B$,
2. if $A \wedge B$ is defined, then neither $A \prec A \wedge B$ nor $B \prec A \wedge B$.
3. if A^c is defined, then not $A \vee A^c \prec \mathbf{1}_{\mathcal{F}}$ and $A \wedge A^c \prec \mathbf{0}_{\mathcal{F}}$

To put it differently, in the partial case Additivity splits into Partial Additivity and Compatibility.

We can now prove that the subject's uncertainty ordering is a partial version of qualitative probability:

LEMMA 3.4. (Partial Qualitative probability from partial elicitation) *Let $\langle \mathcal{F}, \precsim \rangle$ be a partial representational structure from Definition 7. If there is a partial elicitation function e from $\langle \mathcal{F}, \precsim \rangle$ to a reference structure $\langle \mathcal{G}, p \rangle$ such that e satisfies (i)–(iii) from Definition 8 then $\precsim$ satisfies **Pre′**, **Non′**, **Add***, and **Com*** i.e., is a partial qualitative probability.*

PROOF OF LEMMA 3.4. **Com** and **Non′** are assumed in the definition of partial representational structures.

Pre′ As reflexivity is assumed we need only to establish transitivity, the proof of which is the same as in Lemma 3.2.

Add* The proof is again the same as in Lemma 3.2, with the assumption that all operations are defined.

Com* By the properties of the elicitation function, if $A \vee B$ is defined then $e(A \vee B) = e(A) \vee e(B)$. Assume for contradiction that $A \vee B \prec A$. Then by monotonicity $e(A \vee B) \prec_{\mathcal{G}} e(A)$, and so $e(A) \vee e(B) \prec_{\mathcal{G}} e(A)$. But the qualitative ordering on $\mathcal{G}$ respects the Boolean operations—contradiction. The proof for $\wedge$ is completely analogous. ■

---

[6]We show e.g. $A \precsim A \vee B$. Denote $B - A = B \wedge A^c$. By **Non**, $\mathbf{0}_{\mathcal{F}} \precsim B - A$. By **Add**, $\mathbf{0}_{\mathcal{F}} \vee A \precsim (B - A) \vee A$. Finally we have $A = \mathbf{0}_{\mathcal{F}} \vee A \precsim (B - A) \vee A = A \vee B$.

[7]For example $A \vee B$ might be defined in $\mathcal{F}$, but not $B \wedge A^c$, which is needed for the previous derivation.

We now have completed the framework for exploring a subject's internal judgements of uncertainty with respect to elicitation via a reference experiment, and have already seen that it provides a novel characterization of probabilities. In the next section we will further demonstrate the power of this approach by exploring different forms of elicitation.

4. Approximate Elicitation

In the previous sections we began with a representational structure and a minimal characterization of the subject's judgements of uncertainty, deriving the properties of a subject's judgements via elicitation, where elicitation provided perfect matching. This section provides a further generalization: we explore imperfect elicitation of partial structures. We consider two cases in which the elicitation procedure can be imperfect: first, in which the reference experiment is not rich enough to fully express the subject's judgements; second, in which a subject might be unable or unwilling to provide exact matches in the reference experiment.

4.1. Imperfect Reference Experiments

The case of a coarse-grained reference experiment is the easier of the two. In this case, the subject would be able to give an exact match between a proposition and an element were a fine-grained reference experiment available. But if the only available reference experiment is coarse-grained, the subject can only provide approximations. We can give a simple example:

Example 1. Suppose a subject considers two propositions A and B, and uses a twelve-sided die, that is, a reference structure $\mathcal{G}$ generated by outcomes $\{1, \dots, 12\}$ and a flat distribution, for elicitation. Further suppose our subject sets $e(A) = \{1, \dots, 7\}$ and $e(B) = \{1, \dots, 5\}$. Now imagine that in elicitation only a six-sided die is available, that is, a reference experiment $\mathcal{G}'$ generated in the obvious way. Clearly, $\mathcal{G}$ can be mapped onto $\mathcal{G}'$, by, say, pairing $\{1\}$ in $\mathcal{G}'$ with $\{1, 7\}$ in $\mathcal{G}$, $\{2\}$ to $\{2, 8\}$ and so on (it is important to keep in mind the numbers are simply labels for events in the reference experiment, nothing more). Let $e^*(A)$ denote the elicitation of 'it is at most as likely as...that A', and $e_*(A)$ denote the elicitation of 'it is at least as likely as...that A'. The subject assigns $e^*(A) = \{1, 2, 3, 4\}, e_*(A) = \{1, 2, 3\}, e^*(B) = \{2, 3, 4\}, e_*(B) = \{2, 3\}$ as an approximation.

Formally we can describe the situation of an imprecise reference experiment as follows: we have a reference experiment $\mathcal{G}'$ (say, a 6 sided die) which could be extended to a more fine-grained reference experiment $\mathcal{G}$ (say, a 12 sided die), e.g., $\mathcal{G}' \subset \mathcal{G}$. Although a subject might be in principle able to find an elicitation x in the fine-grained experiment $\mathcal{G}$ that exactly corresponds to a proposition A, there may be no such corresponding elicitation in the coarse-grained experiment $\mathcal{G}'$, i.e., the element x might be in $\mathcal{G} - \mathcal{G}'$. In this case the subject must find an approximation of x in $\mathcal{G}'$, which will be the approximate elicitation of A. Consider the case where the subject gives a best approximation from below by giving the greatest lower bound. Formally, the subject aims to find an $x_* \in \mathcal{G}'$ such that if $\mathcal{G}'$ is extended to $\mathcal{G}$, then $x_* \leq_{\mathcal{G}} x$ (it is a lower bound) and for all $y \in \mathcal{G}', y \leq_{\mathcal{G}} x$, then $y \leq_{\mathcal{G}} x_*$ (it is a greatest lower bound).

Requiring the subject give a best approximation to individual elements of the reference experiment imposes the further condition of giving a best approximation to their combinations. Consider the elicitations $x, y, x \wedge y \in \mathcal{G}$ of some propositions $A, B, A \wedge B$. Let the operator $_*$ denote the greatest lower bound of an element from $\mathcal{G}$ in $\mathcal{G}'$. It is clear that $(x \wedge y)_*$ is a lower approximation of *both* elements x and y: $(x \wedge y)_* \leq_{\mathcal{G}} x \wedge y \leq_{\mathcal{G}} x, y$. But since x_*, y_* are the *best* lower approximations, it must be $(x \wedge y)_* \leq_{\mathcal{G}} x_*, y_*$. But that entails that $(x \wedge y)_* \leq_{\mathcal{G}} x_* \wedge y_*$. In the other direction, if the subject chooses $(x \wedge y)_* < x_* \wedge y_*$, then since $x_* \leq x, y_* \leq y$ we get $x_* \wedge y_* \leq x \wedge y$, i.e. $(x_* \wedge y_*)$ is a lower approximation of $x \wedge y$. We have $(x \wedge y)_* < x_* \wedge y_* \leq x \wedge y$ so there is a closer approximation of $x \wedge y$ than $(x \wedge y)_*$, contrary to our assumption.

EXAMPLE 2. To continue our earlier example, suppose the subject gives $(A \wedge B)_* = \{2\}$. But $(A)_* \wedge (B)_* = \{2, 3\}$, and so is a closer approximation to $A \wedge B$, contrary to the assumption that $(A \wedge B)_*$ is the best approximation. A similar argument argument can be made if the subject chooses, say, $(A \vee B)^* = \{1, \ldots, 5\}$.

It follows that if the subject wants to provide a consistent closest approximation from below, her lower elicitation function must obey the condition: if $A \wedge B$ exists, then $(A \wedge B)_* = (A)_* \wedge (B)_*$, i.e. the lower elicitation must commute with meets. By a similar line of reasoning the upper elicitation must commute with joins: if $A \vee B$ exists, then $(A \vee B)^* = (A)^* \vee (B)^*$. Hereafter we will use the notation e_*, e^* when we wish to stress that the function is to be taken as an elicitation.

Together with these two consistency conditions we obviously want the elicitation to respect the qualitative ordering. This requirement is indeed very natural: if $A \precsim B$ and $e_*(A)$ is a lower elicitation of A then the closest lower elicitation $e_*(B)$ of B should not be below it, i.e. $e_*(A) \precsim_{\mathcal{G}} e_*(B)$. Similarly the closest upper elicitation $e^*(A)$ of A should not be above the upper elicitation $e^*(B)$ of B, i.e., $e^*(A) \precsim_{\mathcal{G}} e^*(B)$. These considerations result in the following definition:

DEFINITION 9. (Approximate elicitation function) The functions e_* and e^* from a partial representational structure $\mathcal{F}$ to a reference structure $\mathcal{G}$, are lower and upper elicitation functions iff for any $A, B \in \mathcal{F}$:

(i) if $A \precsim B$, then $e_*(A) \precsim_{\mathcal{G}} e_*(B)$ and $e^*(A) \precsim_{\mathcal{G}} e^*(B)$,

(ii) if $A \wedge B$ exists, then $e_*(A \wedge B) = e_*(A) \wedge e_*(B)$,

(iii) if $A \vee B$ exists, then $e^*(A \vee B) = e^*(A) \vee e^*(B)$.

We defer discussion of any resulting probability function until the next section.

4.2. Imperfect Elicitation

We have shown that if the subject is able to provide best approximations to some 'ideal' elicitation, the corresponding upper/lower elicitation must satisfy the conditions of Definition 9. The converse, however does not in general hold—an elicitation function can satisfy the conditions without being the best approximation in the sense discussed. It might therefore seem that we should include the condition of *best* approximation into the definition of approximate elicitation.

This solution however, would only cover the case of imprecise elicitation due to a coarse-grained reference experiment. It would not be applicable in the case where the reference experiment $\mathcal{G}$ is rich enough, but the subject is not able or willing to determine precise values: there simply would be no better elicitation to be had, even via a richer reference experiment. We can address this case by replacing the notion of best approximation in the sense of closest lower or upper bound with a weaker notion: that the upper approximation cannot be improved using the lower approximation and vice versa. This condition gives a relative notion of best approximation as opposed to the absolute notion given by the assumption of an ideal elicitation.

We can explicate this notion as a 'no improvement' condition. Formally, take an $x \in \mathcal{G}$. By definition, $(x^c)^* \geq_{\mathcal{G}} x^c$, since * is an *upper* approximation. By complementation $((x^c)^*)^c \leq_{\mathcal{G}} x$. From the assumption that x_* is a lower approximation of x that cannot be improved it follows that for any $z \leq_{\mathcal{G}} x$ it must hold that $z \leq_{\mathcal{G}} x_*$, hence $((x^c)^*)^c \leq_{\mathcal{G}} x_*$. Assume the last inequality is a strict inequality: $((x^c)^*)^c <_{\mathcal{G}} x_*$. Then by complementation $(x^c)^* >_{\mathcal{G}} (x_*)^c$

and similarly from $x_* \leq_{\mathcal{G}} x$ we get $(x_*)^c \geq_{\mathcal{G}} x^c$. Hence $(x^c)^* >_{\mathcal{G}} (x_*)^c \geq_{\mathcal{G}} x^c$, and so $(x_*)^c$ would be an improvement of the upper approximation $(x^c)^*$, contrary to assumption. We should therefore require $((x^c)^*)^c = x_*$, and for the same reason $((x^c)_*)^c = x^*$.

Example 3. Continuing with the previous examples, the subject assigns an upper estimate of A of $\{1, 2, 3, 4\}$. By the nature of complementation in the reference experiment, this is a lower estimate of A^c, namely $\{5, 6\}$. If the subject were to offer $(A^c)_* = \{6\}$, we could point out that she has already in effect given a better estimate.

We term pairs of elicitation functions meeting this consistency criterion *regular*:

Definition 10. (Regular elicitation functions) A pair of elicitation functions e_*, e^* is *regular* iff $e_*(A) = (e^*(A^c))^c$ and $e^*(A) = (e_*(A^c))^c$ whenever A^c defined.

Regularity also holds for best approximations with a coarse-grained reference experiment. If $e_*(A)$ is the best lower approximation of A, it cannot be improved—in particular it cannot be improved using its upper approximation e^*.

These conditions lead to a probability-like function:

Theorem 4.1. (Approximate quantitative probabilities) *Let $\langle \mathcal{F}, \precsim \rangle$ be a partial representational structure, $\langle \mathcal{G}, p \rangle$ a reference structure and e_*, e^* a lower (upper) elicitation function. Then there are functions f_*, f^* such that:*

1. *if $A \precsim B$, then $f_*(A) \leq f_*(B)$ and $f^*(A) \leq f^*(B)$*
2. *$f_*(\mathbf{0}_{\mathcal{F}}) = 0 = f^*(\mathbf{0}_{\mathcal{F}}), f_*(\mathbf{1}_{\mathcal{F}}) = f^*(\mathbf{1}_{\mathcal{F}})$*
3. *$f_*(A \vee B) \geq f_*(A) + f_*(B)$ if both $A \vee B, A \wedge B$ are defined and $A \wedge B = \mathbf{0}_{\mathcal{F}}$*
4. *$f^*(A \vee B) \leq f^*(A) + f^*(B)$ if both $A \vee B, A \wedge B$ are defined and $A \wedge B = \mathbf{0}_{\mathcal{F}}$*
5. *$f_*(A) = 1 - f^*(A^c)$ if e_*, e^* is a regular pair.*
6. *Let us denote $\mathcal{F}' = \{A \in \mathcal{F}, e_*(A) = e^*(A)\}$. Then the function $f(A) = f_*(A) = f^*(A)$ from $\mathcal{F}'$ to $\mathcal{G}$ is a partial probability function satisfying conditions 1–5 from Theorem 3.1.*

Proof of Theorem 4.1. We define $f_*(A) = p(e_*(A))$ and $f^*(A) = p(e^*(A))$ and show that 1.–5. follow from the properties of upper and lower elicitations. For $A \in \mathcal{F}$ we write A_*, A^* instead of $e_*(A), e^*(A)$

1. From the requirement that e_*, e^* preserve the qualitative ordering.
2. From the preservation of the top and the bottom. For 3. and 4. assume $A \vee B, A \wedge B$ are defined and $A \wedge B = \mathbf{0}_{\mathcal{F}}$. Then:
3. The lower elicitation e_* preserves meets, so $(A \wedge B)_* = A_* \wedge B_* = \mathbf{0}_{\mathcal{G}}$. Since A_*, B_* are disjoint and p is a probability, $p(A_* \vee B_*) = p(A_*) + p(B_*)$. Also, $p(A \vee B)_* \geq p(A_* \vee B_*)$, since $(A \vee B)_* \geq_{\mathcal{G}} A_* \vee B_*$, as otherwise $A_* \vee B_*$ would be a better approximation of $A \vee B$ than $(A \vee B)_*$. Putting these together $f_*(A \vee B) = p((A \vee B)_*) \geq p(A_* \vee B_*) = p(A_*) + p(B_*) = f_*(A) + f_*(B)$. Hence f_* is superadditive.
4. Since e^* preserves joins $f^*(A \vee B) = p((A \vee B)^*) = p(A^* \vee B^*)$. By the probability calculus $p(A^* \vee B^*) = p(A^*) + p(B^*) - p(A^* \wedge B^*)$, as well as $p(A^* \wedge B^*) \geq 0$. So we can deduce that $f^*(A \vee B) = p(A^* \vee B^*) \leq p(A^*) + p(B^*) = f^*(A) + f^*(B)$. Thus f^* is subadditive.
5. As e_*, e^* are regular, $f_*(A) = p(e_*(A)) = p((e^*(A^c))^c) = 1 - p(e^*(A^c)) = 1 - f^*(A^c)$.
6. By Defintion 9, e_* preserves meets and e^* preserves joins, and are moreover both monotonic with respect to $\precsim_{\mathcal{F}}$. $e(A) = e_*(A) = e^*(A)$ is clearly an elicitation function in the sense of Definition 8. Hence by Theorem 3.1 there exists a partial probability function satisfying conditions 1–5. ■

If the elicitation functions are regular, f_* and f^* are fully analogous to the partial probability function of Section 3, as they are partial inner and outer measures (as made clear by 5.). If, however, the elicitation functions are not regular, then we obtain a generalization of inner and outer measures, since while 5. no longer holds f^* and f_* still have the sub- and superadditivity characteristic of lower and upper measures (e.g., 3. and 4. still hold). At the other extreme, 6. shows that we can see precise elicitation as a special case of approximate elicitation (again, we only deal with the finite case).

Example 4. We can now give probabilities. In the fine-grained experiment the subject gives the probabilities $p(A) = \frac{7}{12}$, $p(B) = \frac{5}{12}$, $p(A \vee B) = \frac{7}{12}$, $p(A \wedge B) = \frac{5}{12}$. In the coarse-grained experiment the subject gives the inner and outer probabilities $p_*(A) = \frac{3}{6} = \frac{1}{2}$, $p^*(A) = \frac{4}{6} = \frac{2}{3}$, $p_*(B) = \frac{2}{6} = \frac{1}{3}$, $p^*(B) = \frac{3}{6} = \frac{1}{2}$. Also, $p^*(A \vee B) = \frac{4}{6} = \frac{2}{3}$ and $p_*(A \wedge B) = \frac{2}{6} = \frac{1}{3}$.

We should stress that we are not providing imprecise probabilities—that is, families of probabilities. This would impose significant requirements on the representational structure that are not in the spirit of this section.

Instead, we have used upper- and lower-probabilities to provide a characterization of faulty elicitation, although, of course, our framework can be used to explicate imprecise probabilities as well.

5. Conclusion

One issue we have not discussed is the relation of our approach to standard approaches employing gambling devices. We note that our approach can be used to explicate depragmatized arguments of the kind in Christensen [2].[8] Elicitation imposes harmony between the subject's judgements and expressions of uncertainty, with no need for penalties. Instead, a clash between those judgements and expressions is a logico-epistemic flaw: incoherent orderings cannot be matched with a reference experiment. Pragmatic considerations could be added, of course by, for example, attaching penalties to the reference experiment. Or, elicitation could take the form of some distance function, yielding a version of a non-pragmatic argument of the kind in Joyce [10].

Considering explicitly the relation between judgement and expression opens up many other possibilities, such as subjects with non-Boolean judgements of uncertainty, subjects with multiple judgements of uncertainty, and perhaps, by considering families of different elicitations, subjects who are ambivalent, or who are not being completely truthful. The generality of our account makes clear the specificity of arguments for probability. Probabilistic reference experiments yield probabilities; non-probabilistic reference experiments will yield non-probabilistic calculi. In the case of subjects with a representational structure suitable for fuzzy logics (e.g. an MV algebra), who express themselves via a fuzzy reference experiment (as in Marra [13]) we obtain constraints on uncertainty in the form of a many-valued logic. Or, there could be subjects with mixed internal representational structures, as well as subjects with mixed reference experiments. For example, a subject's representations could form a fuzzy logical structure while the reference experiment remains Boolean (allowing Mundici [14] results to be applied).

Thus, the justification for probabilities then turns into the justification for a particular kind of reference experiment. This points to further research. We have so far stayed with a very standard model: that of a subject in a laboratory providing responses to an experiment. An ecological approach

[8] As well as with Gillies [8], who envisages a willing subject engaging a psychologist to determine her probabilities.

removing the subject from the laboratory would provide a much more general (and realistic) model. Exploring these possibilities can lead to much future research.

Acknowledgements. The authors wish to thank Colin Howson, Peter Milne, Samuel Fletcher, Jeff Paris and two anonymous reviewers for helpful comments. Work on this paper was supported by grant no. GA16-15621S of the Czech Science Foundation.

References

[1] Chisholm, R., *Perceiving: A Philosophical Study*, Cornell University Press, Ithaca, NY, 1957.

[2] Christensen, D., *Putting Logic in its Place*, Oxford University Press, Oxford, 2004.

[3] DeGroot, M., *Optimal Statistical Decisions*, Wiley, Hoboken, NJ, 2004.

[4] Eriksson, L., and A. Hájek, What are degrees of belief. *Studia Logica* 86:183–213, 2007.

[5] Fishburn, P., The axioms of subjective probability. *Statistical Science* 1:335–345, 1986.

[6] French, S., On the axiomatisation of subjective probabilities. *Theory and Decision* 14:19–33, 1982.

[7] Geach, P., *Mental Acts: Their Content and Their Objects*, Routledge & Kegan Paul, London, 1957.

[8] Gillies, D., *Philosophical Theories of Probability*, Routledge, London, 2000.

[9] Howson, C., and P. Urbach, *Scientific Reasoning: The Bayesian Approach*, Open Court, La Salle, IL, 1993.

[10] Joyce, J. M., Accuracy and coherence: Prospects for an alethic epistemology of partial belief, in F. Huber and C. Schmidt-Petri (eds.), *Degrees of Belief*, Springer, Berlin, 2009, pp. 263–300.

[11] Koopman, B. O., The bases of probability. *Bulletin of the American Mathematical Society* 46:763–774, 1940.

[12] Krantz, D. H., R. D. Luce, P. Suppes, and A. Tversky, *Foundations of Measurement: Vol 1: Additive and Polynomial Representations*, Academic Press, New York, 1971.

[13] Marra, V., The problem of artificial precision in theories of Vagueness a note on the Rôle of maximal consistency. *Erkenntnis* 79:1015–1026, 2014.

[14] Mundici, D., Bookmaking over infinite-valued events. *International Journal of Approximate Reasoning* 43:223–240, 2006.

[15] Narens, L., *Theories in Probability: An Examination of Logical and Qualitative Foundations*, World Scientific Publishing Co., Inc., River Edge, NJ, 2007.

[16] Ramsey, F. P., Truth and probability, in R. B. Braithwaite (ed.), *The Foundations of Mathematics and Other Logical Essays*, Routledge and Kegan Paul, London, 1931.

[17] Savage, L. J., *The Foundations of Statistics*, 2nd ed., Dover Publications, New York, 1972.

[18] Schick, F., Dutch bookies and money pumps. *The Journal of Philosophy* 83:112–119, 1986.

[19] Spetzler, C. S., and C.-A. S. Stael Von Holstein, Probability encoding in decision analysis. *Management Science* 22:340–358, 1975.

[20] Von Neumann, J., and O. Morgenstern, *Theory of Games and Economic Behavior*, 2nd rev. ed, Princeton University Press, Princeton, 1947.

T. Childers, O. Majer
Institute of Philosophy
Czech Academy of Sciences
Jilska 1
110 00 Prague
Czech Republic
majer@flu.cas.cz

T. Childers
childers@flu.cas.cz