

A Particular Upper Expectation as Global Belief Model for Discrete-Time Finite-State Uncertain Processes

Natan T'Joens, Jasper De Bock, Gert de Cooman

FLip, Ghent University, Technologiepark 125, 9052 Zwijnaarde, Belgium

Abstract

To model discrete-time finite-state uncertain processes, we argue for the use of a global belief model in the form of an upper expectation that is the most conservative one under a set of basic axioms. Our motivation for these axioms, which describe how local and global belief models should be related, is based on two possible interpretations for an upper expectation: a behavioural one similar to Walley's, and an interpretation in terms of upper envelopes of linear expectations. We show that the most conservative upper expectation satisfying our axioms, that is, our model of choice, coincides with a particular version of the game-theoretic upper expectation introduced by Shafer and Vovk. This has two important implications: it guarantees that there is a unique most conservative global belief model satisfying our axioms; and it shows that Shafer and Vovk's model can be given an axiomatic characterisation and thereby provides an alternative motivation for adopting this model, even outside their game-theoretic framework. Finally, we relate our model to the upper expectation resulting from a traditional measure-theoretic approach. We show that this measure-theoretic upper expectation also satisfies the proposed axioms, which implies that it is dominated by our model or, equivalently, the game-theoretic model. Moreover, if all local models are precise, all three models coincide.

Keywords: upper expectations, uncertain processes, coherence, game-theoretic probability, measure-theoretic probability, global uncertainty model

1. Introduction

There are various ways to describe discrete-time uncertain processes, such as Markov processes, mathematically. For many, measure theory is and has been the preferred framework to describe the uncertain dynamics of such processes [1–3]. Others have used martingales or a game-theoretic approach [4–6]. However, the common starting point for all these approaches are the local belief models. They describe the dynamics of the process from one time instant to the next and they are typically learned from measurements or elicited from experts. For instance, in some specific cases, one typically may have information about ‘the probability of throwing heads on the next coin toss’, ‘the expected number of goods that are sold by a certain shop on a single day’, ‘the probability of rain tomorrow’, ... In a measure-theoretic context, such local belief models are presented in the form of probability charges or probability measures on the local state space; in a game-theoretic context, sets of allowable bets are used. The latter belong to a family of so-called imprecise probability models [7–9], e.g. upper and lower previsions (or expectations), sets of desirable gambles, credal sets... Such imprecise probability models generalise traditional precise models, in the sense that they allow for indecision and for a representation that only expresses partial knowledge about the probabilities. In order to allow for imprecision in a measure-theoretic context, we will often consider sets of probability measures (or charges) on the local state space. Then, if local state-spaces are assumed to be finite, as in our case, the game-theoretic and measure-theoretic local descriptions are mathematically equivalent. This will be clarified in Sections 2 and 3.

In practice however, we are typically interested in more general inferences such as ‘the (lower and upper) expected number of tosses until the first tails is thrown’, ‘the (lower and upper) probability of being out of stock on a given day’, ... It is less straightforward how we should directly learn about such more complicated inferences and, even if we could in principle do so, it is often not possible or feasible to gather sufficient information

because of time and budget limitations. Hence, the question arises: ‘How and to which extent, can we extend the information captured in the local models towards global information about the entire process?’

Measure theory relies on countable additivity to do this elegantly, and this results in a mathematically powerful, but rather abstract framework. The mathematical arguments are often encumbered by measurability constraints, and it is not always clear how to extend classical measure-theoretic notions to an imprecise probabilities setting. On the other hand, Shafer and Vovk’s game-theoretic framework defines global upper expectations constructively using the concept of a ‘supermartingale’. Though it allows us to directly model imprecision and does not require measurability constraints, the involved upper expectations thus far lack a characterisation in terms of mathematical properties. This impacts the generality of their model in the negative, since it can therefore only be motivated from a game-theoretic point of view.

We propose a different, third approach. Our aim is to establish a global belief model, in the form of a conditional upper expectation, that extends the information contained in the local models by using a number of mathematical properties. Notably, this model will not be bound to a single interpretation. Indeed, its characterising properties can be justified starting from a number of different interpretations. We consider and discuss two of the most significant ones in Section 4; a behavioural interpretation that is more in line with a game-theoretic approach to uncertainty, and an interpretation in terms of upper envelopes of linear expectation operators. We then define the desired global model as the most conservative—the unique model that does not contain any additional information—under this particular set of properties. In other words, it is the so-called natural extension [7, 9, 10] under this set of properties.

Care should be taken here though. While the concept of natural—that is, most conservative—extension is well-established in the field of imprecise probabilities, it usually refers to the natural extension *under coherence*. The advantage of this notion of coherence, and the related natural extension, is that it can be applied to general uncertainty models. However, the scope and properties of the resulting extension can sometimes be rather weak. A first issue with natural extension under coherence is that it traditionally only allows us to work with *bounded* variables. The problem of how to extend this method to a broader class of variables was addressed in [9, Chapter 13], yet, only for the class of *unbounded real-valued* variables, whereas we wish to consider extended real-valued variables as well. There is a second, more profound issue though, which is that the natural extension under coherence lacks some basic continuity properties that we feel are desirable. For example, in our context of uncertain processes, natural extension under coherence would lead to a global upper expectation that does not necessarily satisfy upward monotone convergence, even if we were to restrict attention to bounded variables.¹ Furthermore, and related to this, the corresponding inferences are often very conservative, or even vacuous. If we would only be interested in *bounded* variables and uncertain processes with a *finite* time horizon, all of these issues would not come to the fore, and the natural extension under coherence would then be a valuable model; see for example [11]. However, in our—vastly—more general setting, where we will be looking at *extended real-valued* variables and uncertain processes with *infinite* time horizons, we find this model unsatisfactory for the reasons described above. We will therefore not be considering the natural extension under coherence (alone); instead, we will consider the natural extension under a different set of axioms. Unlike the coherence axioms, our axioms are specifically designed with the purpose of extending local models in the context of uncertain processes (with infinite time horizons). As we will see, the resulting natural—most conservative—extension yields more satisfying results when it comes to the mathematical properties of the resulting global model.

In Section 5, we give an alternative characterisation of our most conservative model by showing that it is the unique operator that satisfies a specific set of properties. This alternative characterisation is particularly interesting from a technical point of view, since it allows us to establish our main result in Section 6: that our

¹The experienced (or interested) reader can verify this using the following example (the notation and terminology is introduced further on in the paper). Interestingly, it shows that upward monotone convergence fails even for increasing sequences of n -measurable gambles that converge to a limit that is itself a gamble. Fix any $x \in \mathcal{X}$ and let $Q_s(f) := f(x)$ for all $s \in \mathcal{X}^*$; this means that in every situation, the next state will be x with (lower) probability 1. Consider now the increasing sequence $\{f_n\}_{n \in \mathbb{N}}$ of n -measurable gambles defined, for all $n \in \mathbb{N}$, by $f_n(\omega) := 0$ if $\omega_i = x$ for all $i \in \{1, \dots, n\}$ and $f_n(\omega) := 1$ for all other $\omega \in \Omega$. Then it can be checked that the value of the natural extension under coherence is 0 for all gambles f_n , but that it is 1—and therefore vacuous—for the limit $f := \lim_{n \rightarrow +\infty} f_n$. This means that the (upper) probability of observing a state that is not x in the first n steps, for any n , is zero, but that the upper probability of ever observing such a state is 1. We are indebted to Matthias C. M. Troffaes for this example—or rather, a very similar one that inspired it. He presented it in a talk at WPMSIIP 2013 in Lugano, and recently also discussed it with us in private communication; it seems to be unpublished though.

model coincides with a version of the game-theoretic upper expectation introduced by Shafer and Vovk [4, 5]. On the one hand, this serves as an additional motivation for our model. On the other hand, it establishes a concrete axiomatisation for the game-theoretic upper expectation. Section 8 aims at strengthening the relevance of the most conservative model that results from our axioms by listing its most salient properties.

In the final section, we couch our findings in a more measure-theoretic language. As a first step, we put forward a sensible way of applying measure theory in our imprecise probabilities setting. To do this, we draw inspiration from earlier work done in our group [12]. Subsequently, we show that the measure-theoretic upper expectation this results in, satisfies the axioms presented in Section 4 and is therefore always dominated by the most conservative one satisfying these axioms—which is our model. Moreover, the fact that this model, entirely based on measure-theoretic principles, also satisfies our proposed axioms, is to be seen as an additional motivation to adopt them. Furthermore, in a context where all local models are precise, meaning that every single one of them can be represented by a single probability mass function on the local state space, both upper expectations—and hence, also the game-theoretic one—turn out to be one and the same model.

As a final remark, we want to mention that our current work extends upon an earlier conference contribution [13]. The added material is mainly gathered in the penultimate section; it concerns the study of how our model is related to a possible measure-theoretic model. Apart from that, another important difference with the conference contribution lies in how local models are extended in order to become consistent with the game-theoretic approach; see the first part of Section 6 and compare this with [13, Section 2].

2. Upper Expectations

We denote the set of all natural numbers, without 0, by \mathbb{N} , and let $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. The set of extended real numbers is denoted by $\overline{\mathbb{R}} := \mathbb{R} \cup \{+\infty, -\infty\}$ and is endowed with the usual order and the usual order topology (corresponding to the two-point compactification of \mathbb{R}). The set of positive real numbers is denoted by $\mathbb{R}_{>0}$ and the set of non-negative real numbers by $\mathbb{R}_{\geq 0}$. We also adopt the usual conventions for the addition between the reals and $+\infty$ and $-\infty$, and the conventions that $+\infty - \infty = -\infty + \infty = +\infty$ and $0 \cdot (+\infty) = 0 \cdot (-\infty) = 0$.

Informally, we consider a subject who is uncertain about the value that some variable Y assumes in a non-empty set \mathcal{Y} . More formally, we call any map on \mathcal{Y} a *variable*; our informal Y is a special case: it corresponds to the identity map on \mathcal{Y} . A subject's uncertainty about the unknown value of Y can then be represented by an *upper expectation* \overline{E} : an extended real-valued map defined on some subset \mathcal{D} of the set $\overline{\mathcal{L}}(\mathcal{Y})$ of all extended real-valued variables on \mathcal{Y} . An element f of $\overline{\mathcal{L}}(\mathcal{Y})$ is simply called an *extended real variable*. We say that f is bounded below if there is some real c such that $f(y) \geq c$ for all $y \in \mathcal{Y}$, and we say that f is bounded above if $-f$ is bounded below. We call a sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of extended real variables *uniformly bounded below* if there is a real c such that $f_n(y) \geq c$ for all $n \in \mathbb{N}_0$ and $y \in \mathcal{Y}$. An important role will be reserved for elements f of $\overline{\mathcal{L}}(\mathcal{Y})$ that are bounded, meaning that they are bounded above *and* below. Bounded real-valued variables on \mathcal{Y} are called *gambles* on \mathcal{Y} , and we use $\mathcal{L}(\mathcal{Y})$ to denote the set of all of them. The set of all bounded below elements of $\overline{\mathcal{L}}(\mathcal{Y})$ is denoted by $\overline{\mathcal{L}}_b(\mathcal{Y})$.

Consider now the special case that \overline{E} is defined on at least the set of all bounded real-valued variables; so $\mathcal{L}(\mathcal{Y}) \subseteq \mathcal{D}$. Then we call \overline{E} *boundedly coherent* [7–9] if it satisfies the following three coherence axioms:

- C1. $\overline{E}(f) \leq \sup f$ for all $f \in \mathcal{L}(\mathcal{Y})$; [upper bound]
- C2. $\overline{E}(f + g) \leq \overline{E}(f) + \overline{E}(g)$ for all $f, g \in \mathcal{L}(\mathcal{Y})$; [sub-additivity]
- C3. $\overline{E}(\lambda f) = \lambda \overline{E}(f)$ for all $\lambda \in \mathbb{R}_{\geq 0}$ and $f \in \mathcal{L}(\mathcal{Y})$. [non-negative homogeneity]

In the particular case where $\mathcal{D} = \mathcal{L}(\mathcal{Y})$, we will simply say that \overline{E} is *coherent*. For any (boundedly) coherent \overline{E} , if we let \underline{E} be the conjugate lower expectation defined by $\underline{E}(f) := -\overline{E}(-f)$ for all $f \in -\mathcal{D}$, then the following additional properties follow from C1–C3:

- C4. $f \leq g \Rightarrow \overline{E}(f) \leq \overline{E}(g)$ for all $f, g \in \mathcal{L}(\mathcal{Y})$; [monotonicity]
- C5. $\inf f \leq \underline{E}(f) \leq \overline{E}(f) \leq \sup f$ for all $f \in \mathcal{L}(\mathcal{Y})$; [bounds]

C6. $\bar{E}(f + \mu) = \bar{E}(f) + \mu$ for all $\mu \in \mathbb{R}$ and all $f \in \mathcal{L}(\mathcal{Y})$; [constant additivity]

C7. for any $f \in \mathcal{L}(\mathcal{Y})$ and any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ in $\mathcal{L}(\mathcal{Y})$: [uniform continuity]

$$\lim_{n \rightarrow +\infty} \sup |f - f_n| = 0 \Rightarrow \lim_{n \rightarrow +\infty} \bar{E}(f_n) = \bar{E}(f).$$

Proof. We only prove C5, which clearly implies that \bar{E} is real-valued on $\mathcal{L}(\mathcal{Y})$, and the remaining properties then follow from the standard argumentation in [7, Section 2.6.1.].

First, note that $\bar{E}(0) = 0$ because of C3 and our convention that $0 \cdot (+\infty) = 0 \cdot (-\infty) = 0$. Therefore, for all $f \in \mathcal{L}(\mathcal{Y})$, it follows from C2 that $0 \leq \bar{E}(f) + \bar{E}(-f)$, or equivalently, that $-\bar{E}(-f) \leq \bar{E}(f)$. Applying C1 to both sides, we find that $\inf f = -\sup(-f) \leq -\bar{E}(-f) \leq \bar{E}(f) \leq \sup f$. The result now follows readily from the definition of \bar{E} . \square

A gamble f is typically interpreted as an uncertain reward or gain that depends on the value that Y takes in \mathcal{Y} ; if Y takes the value y , the (possibly negative) gain is $f(y)$. Then, according to Walley's behavioural interpretation [7], the upper expectation $\bar{E}(f)$ of a gamble f is a subject's infimum selling price for the gamble f , implying that, for any $\alpha > \bar{E}(f)$, the subject is willing to accept the gamble $\alpha - f$.² Axioms C1–C3 are then called rationality axioms, since they ensure that these selling prices are chosen rationally. Alternatively, any coherent upper expectation on $\mathcal{L}(\mathcal{Y})$ can equivalently be represented by a set of *linear expectations* on $\mathcal{L}(\mathcal{Y})$: coherent upper expectations on $\mathcal{L}(\mathcal{Y})$ that are self-conjugate, meaning that $\bar{E}(f) = -\bar{E}(-f)$ for all $f \in \mathcal{L}(\mathcal{Y})$. More precisely, it follows from [7, Section 3.3.3] that any coherent upper expectation \bar{E} on $\mathcal{L}(\mathcal{Y})$ is the upper envelope of the set of all linear expectations on $\mathcal{L}(\mathcal{Y})$ that are dominated by \bar{E} ; so $\bar{E}(f) = \sup\{E(f) : (\forall g \in \mathcal{L}(\mathcal{Y})) E(g) \leq \bar{E}(g)\}$, where E ranges over the linear expectations on $\mathcal{L}(\mathcal{Y})$. Moreover, according to [9, Theorem 8.15], linear expectations on $\mathcal{L}(\mathcal{Y})$ are in a one-to-one relation with *probability charges* on the power set $\mathcal{P}(\mathcal{Y})$ of \mathcal{Y} , being maps $p : \mathcal{P}(\mathcal{Y}) \rightarrow \mathbb{R}_{\geq 0}$ that are finitely additive and where $p(\mathcal{Y}) = 1$ and $p(\emptyset) = 0$. Such probability charges are more general than the more conventional *probability measures*, which additionally require σ -additivity and typically also that the domain should be restricted to a σ -algebra on \mathcal{Y} . However, when \mathcal{Y} is finite, this distinction disappears, and we can simply limit ourselves to working with probability mass functions on \mathcal{Y} ; these can be seen as probability charges or measures restricted to the domain of all singletons. In that case, we have that $E(f) = \sum_{y \in \mathcal{Y}} f(y)p(y)$ for any linear expectation E on $\mathcal{L}(\mathcal{Y})$ and any $f \in \mathcal{L}(\mathcal{Y})$, where p is the unique probability mass function on \mathcal{Y} defined by $p(y) := E(\mathbf{1}_A)$ for all $y \in \mathcal{Y}$. Here, we used $\mathbf{1}_A$ to denote the *indicator* of $A \subseteq \mathcal{Y}$: the gamble on \mathcal{Y} assuming the value 1 on A and 0 elsewhere, and where we will not distinguish in notation between y and $\{y\}$. Hence, if \mathcal{Y} is finite, any coherent upper expectation \bar{E} on $\mathcal{L}(\mathcal{Y})$ can alternatively be represented by the corresponding set of probability mass functions

$$\mathbb{P}_{\bar{E}} := \left\{ p \in \mathbb{P} : (\forall f \in \mathcal{L}(\mathcal{Y})) \sum_{y \in \mathcal{Y}} f(y)p(y) \leq \bar{E}(f) \right\}, \quad (1)$$

where \mathbb{P} denotes the set of all probability mass functions on \mathcal{Y} . It then follows from [14, Theorem 2]³ that

$$\bar{E}(f) = \max \left\{ \sum_{y \in \mathcal{Y}} f(y)p(y) : p \in \mathbb{P}_{\bar{E}} \right\} \text{ for all } f \in \mathcal{L}(\mathcal{Y}). \quad (2)$$

This alternative representation of a coherent upper expectation will be our starting point in Section 9, where we aim to establish a measure-theoretic global belief model in the form of an upper expectation.

It follows from the discussion above that, in general, (boundedly coherent) upper expectations can be interpreted in at least two possible ways: in a direct behavioural manner in terms of selling prices for gambles, or as a supremum over—an upper envelope of—a set of linear expectations. As will be discussed in Section 6, the

²Strictly speaking, the interpretation as infimum selling price only guarantees that there are $\alpha > \bar{E}(f)$ *arbitrarily close to* $\bar{E}(f)$ for which our subject accepts $\alpha - f$ (by definition of the infimum). Our claim that this is true for *all* $\alpha > \bar{E}(f)$ is based on an additional implicit assumption that if a subject is willing to sell a gamble for some price, then she is also willing to sell that gamble for any higher price.

³More specifically, Equation (2) does not follow from [14, Theorem 2] itself, but rather from the last paragraph of its proof.

behavioural interpretation allows us to motivate the game-theoretic approach proposed by Shafer and Vovk. On the other hand, an interpretation in terms of upper envelopes of linear expectations seems more natural from a measure-theoretic point of view, where probability measures or mass functions are regarded as the primary objects for describing uncertainty. Linear expectations are then typically obtained by integration and limiting the domain to measurable functions. In this paper, we will not allow ourselves to be bound to any of these two interpretations. Instead, we will motivate the defining properties of our proposed global belief model in terms of either of these interpretations.

3. Upper Expectations in Discrete-Time Uncertain Processes

We consider a *discrete-time uncertain process*: a sequence $X_1, X_2, \dots, X_n, \dots$ of uncertain states, where the state X_k at each discrete time $k \in \mathbb{N}$ takes values in a fixed non-empty set \mathcal{X} , called the *state space*. We will assume that this state space \mathcal{X} is *finite*. Let a *situation* $x_{1:n}$ be any string $(x_1, \dots, x_n) \in \mathcal{X}_{1:n} := \mathcal{X}^n$ of possible state values with finite length $n \in \mathbb{N}_0$. In particular, the unique empty string $x_{1:0}$, denoted by \square , is called the *initial situation*: $\mathcal{X}_{1:0} := \{\square\}$. We denote the set of all situations by $\mathcal{X}^* := \cup_{n \in \mathbb{N}_0} \mathcal{X}_{1:n}$. We will also use the generic notations s and t to denote situations. Furthermore, when we write $x_{1:n} \in \mathcal{X}^*$, we implicitly assume that $n \in \mathbb{N}_0$.

To model our uncertainty about the *dynamics* of an uncertain process, we associate, with every situation $x_{1:n} \in \mathcal{X}^*$, a coherent upper expectation $\overline{Q}_{x_{1:n}}$ on $\mathcal{L}(\mathcal{X})$. This upper expectation expresses a subject's beliefs about the uncertain value of the next state X_{n+1} , when she has observed that $X_1 = x_1, X_2 = x_2, \dots, X_{n-1} = x_{n-1}$ and $X_n = x_n$. Hence, it gives information about how the process changes from one time instant to the next. We will therefore also refer to $\overline{Q}_{x_{1:n}}$ as the *local model* or the *local upper expectation* associated with $x_{1:n}$. We gather all local models in an *imprecise probability tree*: a function \overline{Q}_\bullet that maps any situation $s \in \mathcal{X}^*$ to its corresponding coherent upper expectation \overline{Q}_s . Hence, such an imprecise probability tree \overline{Q}_\bullet , which we will also simply denote by \overline{Q} , represents the dynamics of the uncertain process as a whole.

As discussed in the previous section, we consider two possible ways of interpreting the coherent upper expectations \overline{Q}_s : in terms of upper envelopes of linear expectations and in terms of betting behaviour. The latter is, by nature, more oriented towards a game-theoretic approach, whereas the interpretation in terms of upper envelopes of linear expectations will seem more natural for a reader with a measure-theoretic background. In fact, since \mathcal{X} is assumed to be finite, the coherent upper expectation \overline{Q}_s , for any $s \in \mathcal{X}^*$, can be represented equivalently by the set $\mathbb{P}_s := \mathbb{P}_{\overline{Q}_s}$ of probability mass functions on \mathcal{X} , related to it by Equations (1) and (2). In measure-theoretic contexts, this representation in terms of \mathbb{P}_s is then typically preferred over the representation in terms of \overline{Q}_s .

The mathematical equivalence between the different approaches—the game-theoretic, the measure-theoretic and ours—on a local level is crucial, because only then, we can meaningfully compare the differences of their associated uncertainty models on a global level. Indeed, though it is locally only a matter of interpretational differences, it is not immediately clear how these different approaches relate on a global level, where we look at the entire dynamics of the process and no longer only at the transition from one time instant to the next. Each approach will extend the local models using different concepts, assumptions and methods, resulting in possibly different global models. Our upper expectation approach will use a limited set of simple properties, that, as we will argue, are desirable for a global model to have, under both of the interpretations we consider. Before we proceed to do so, we finish this section by introducing some further notation about uncertain processes.

An infinite sequence $x_1 x_2 x_3 \dots$ of state values is called a *path*, which we denote by $\omega := x_1 x_2 x_3 \dots$. We gather all paths in the *sample space* $\Omega := \mathcal{X}^{\mathbb{N}}$. For any path $\omega := x_1 x_2 \dots \in \Omega$, the situation $x_{1:n} := x_1 x_2 \dots x_n$ that consists of its first n state values is denoted by $\omega^n \in \mathcal{X}_{1:n}$. The state value x_n at time n is denoted by $\omega_n \in \mathcal{X}$. An *event* $A \subseteq \Omega$ is a collection of paths, and in particular, the *cylinder event* $\Gamma(x_{1:n}) := \{\omega \in \Omega : \omega^n = x_{1:n}\}$ of some situation $x_{1:n} \in \mathcal{X}^*$, is the set of all paths $\omega \in \Omega$ that ‘go through’ the situation $x_{1:n}$.

A variable on Ω is called a *global variable* and we gather all extended real(-valued) global variables in the set $\overline{\mathbb{V}} := \overline{\mathcal{L}}(\Omega)$. Similarly, we let $\overline{\mathbb{V}}_b := \overline{\mathcal{L}}_b(\Omega)$ and $\mathbb{V} := \mathcal{L}(\Omega)$. For any natural $k \leq \ell$, we denote by $X_{k:\ell}$ the global variable defined by $X_{k:\ell}(\omega) := (\omega_k, \dots, \omega_\ell)$ for all $\omega \in \Omega$. As such, the state $X_k := X_{k:k}$ at time k can also be regarded as a global variable. Moreover, for any natural $k \leq \ell$ and any map $f : \mathcal{X}^{\ell-k+1} \rightarrow \overline{\mathbb{R}}$, we will write $f(X_{k:\ell})$ to denote the global extended real variable defined by $f(X_{k:\ell}) := f \circ X_{k:\ell}$. We call a global extended real variable f *n-measurable* for some $n \in \mathbb{N}_0$, if it only depends on the initial n state values; so $f(\omega_1) = f(\omega_2)$ for any two paths

ω_1 and ω_2 such that $\omega_1^n = \omega_2^n$. We will then use the notation $f(x_{1:n})$ for its constant value $f(\omega)$ on all paths $\omega \in \Gamma(x_{1:n})$. Clearly, the indicator $\mathbb{I}_{x_{1:n}} := \mathbf{1}_{\Gamma(x_{1:n})}$ of the cylinder event $\Gamma(x_{1:n})$ for $x_{1:n} \in \mathcal{X}^*$ is an n -measurable variable. Finally, we call any $f \in \overline{\mathbb{V}}$ *finitary* if it is n -measurable for some $n \in \mathbb{N}_0$. We gather all finitary *gambles* in \mathbb{V}_{fin} .

4. In Search of a Global Belief Model

Any extended real-valued map $\overline{\mathbb{E}}: \overline{\mathbb{V}} \times \mathcal{X}^* \rightarrow \overline{\mathbb{R}}: (f, s) \mapsto \overline{\mathbb{E}}(f|s)$ will be called a *global upper expectation*. The corresponding *global lower expectation* $\underline{\mathbb{E}}: \overline{\mathbb{V}} \times \mathcal{X}^* \rightarrow \overline{\mathbb{R}}: (f, s) \mapsto \underline{\mathbb{E}}(f|s)$ is then defined by conjugacy: $\underline{\mathbb{E}}(f|s) := -\overline{\mathbb{E}}(-f|s)$ for all $f \in \overline{\mathbb{V}}$ and all $s \in \mathcal{X}^*$. These lower and upper expectations also give rise to upper and lower probabilities. In particular, for any event $A \subseteq \Omega$ and any $s \in \mathcal{X}^*$, $\overline{\mathbb{P}}(A|s) := \overline{\mathbb{E}}(\mathbf{1}_A|s)$ is the *upper probability* of A conditional on s and, similarly, $\underline{\mathbb{P}}(A|s) := \underline{\mathbb{E}}(\mathbf{1}_A|s)$ is the *lower probability* of A conditional on s . Throughout this work, however, we will focus almost entirely on (global) upper expectations, and will regard (global) lower expectations and upper and lower probabilities as derived notions.

Given an imprecise probability tree $\overline{\mathbb{Q}}$ that associates a local upper expectation $\overline{\mathbb{Q}}_s$ with every situation $s \in \mathcal{X}^*$, we aim to define a global upper expectation $\overline{\mathbb{E}}$ that extends the information included in these local models in a ‘rational’ manner. To do so, we impose the following properties, where we adopt the notation that $\overline{\mathbb{E}}(f|X_{1:n}) := \overline{\mathbb{E}}(f|\cdot) \circ X_{1:n}$ for all $f \in \overline{\mathbb{V}}$ and all $n \in \mathbb{N}_0$, and where limits of variables are intended to be taken pointwise.

P1. $\overline{\mathbb{E}}(f(X_{n+1})|x_{1:n}) = \overline{\mathbb{Q}}_{x_{1:n}}(f)$ for all $f \in \mathcal{L}(\mathcal{X})$ and all $x_{1:n} \in \mathcal{X}^*$.

P2. $\overline{\mathbb{E}}(f|s) = \overline{\mathbb{E}}(f \mathbb{I}_s|s)$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $s \in \mathcal{X}^*$.

P3. $\overline{\mathbb{E}}(f|X_{1:n}) \leq \overline{\mathbb{E}}(\overline{\mathbb{E}}(f|X_{1:n+1})|X_{1:n})$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $n \in \mathbb{N}_0$.

P4. $f \leq g \Rightarrow \overline{\mathbb{E}}(f|s) \leq \overline{\mathbb{E}}(g|s)$ for all $f, g \in \overline{\mathbb{V}}$ and all $s \in \mathcal{X}^*$.

P5. For any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of finitary gambles that is uniformly bounded below and any $s \in \mathcal{X}^*$:

$$\lim_{n \rightarrow +\infty} f_n = f \Rightarrow \limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s) \geq \overline{\mathbb{E}}(f|s).$$

To motivate P1–P5, we need to attach some interpretation to $\overline{\mathbb{E}}$. We will consider two particular ones, similar to what we have done for the (boundedly) coherent upper expectations in Section 2.

We start from the interpretation of a global *gamble* $f \in \mathbb{V}$ as an uncertain reward depending on the uncertain path ω that the uncertain process takes in Ω . However, it is not clear what this means operationally if the gamble f depends on the entire length of the path. Indeed, the gamble f depends on an infinite number of subsequent state values, so there is no point in time when we can determine the actual reward linked to the gamble f . The same interpretational problem arises when considering unbounded or extended real variables (on a general set \mathcal{X} such as Ω). The simple interpretation of an uncertain reward does not suffice there because the reward itself can be unbounded or infinite, which is unrealistic—or even meaningless—in operational practice. For this reason, we prefer to only attach a direct operational meaning to the value $\overline{\mathbb{E}}(f|s)$ of a global upper expectation $\overline{\mathbb{E}}$ for a *finitary gamble* $f \in \mathbb{V}_{\text{fin}}$ conditional on a situation $s \in \mathcal{X}^*$. Such finitary gambles can be given an operationally meaningful behavioural interpretation as uncertain rewards because they take real values⁴ and only depend on the state at a finite number of time instances.

We distinguish the following two ways for interpreting the global upper expectation $\overline{\mathbb{E}}(f|s)$ of a finitary gamble $f \in \mathbb{V}_{\text{fin}}$ conditional on a situation $s \in \mathcal{X}^*$:

- **Behavioural interpretation.** It is a subject’s infimum selling price for f *contingent* on the event $\Gamma(s)$, implying that,⁵ for any $\alpha > \overline{\mathbb{E}}(f|s)$, she is willing to accept the uncertain reward associated with the gamble $\mathbb{I}_s(\alpha - f)$.

⁴One could also question the meaning of a direct behavioural interpretation for some specific (real-valued) gambles. For instance, how do we exchange money if the gamble’s value is equal to π ? In this case, we can only give an indirect interpretation in terms of simpler gambles that are ‘sufficiently’ close.

⁵The comment in Footnote 2, applies here as well, suitable adapted to the conditional setting.

- **Interpretation as an upper envelope.** It is the supremum value of $E(f|s)$, where E belongs to some given set \mathbb{E} of conditional linear expectation operators: $\bar{E}(f|s) = \sup\{E(f|s): E \in \mathbb{E}\}$.

Since P1–P3 only apply to finitary gambles, a direct justification for these axioms can be given quite easily in each of the above interpretations. Property P1 requires that the global model \bar{E} should be compatible with the local models \bar{Q}_s . The desirability of this property is self-evident, no matter which interpretation is used. Property P2 requires that the upper expectation of a finitary gamble f conditional on s should only depend on the value of f on the paths $\omega \in \Gamma(s)$. This property is clearly desirable when using the behavioural interpretation, because $\mathbb{I}_s(\alpha - f)$ only depends on the restriction of f to $\Gamma(s)$. It is also quite evident that this property is desirable for an upper envelope of conditional linear expectations, because P2 should in particular also hold for such conditional linear expectations themselves. Similarly, that P3 should hold under the upper envelope interpretation, is motivated by the fact that conditional linear expectations satisfy P3 with equality—then better known as the law of iterated expectations; an upper envelope of conditional expectations is therefore guaranteed to satisfy P3, under the assumption that P4 holds for conditional linear expectations. In order to see that property P3 is also desirable according to the behavioural interpretation, one requires a conditional version of the notion of coherence that we discussed in Section 2 [7, 9, 10].⁶ Explaining in more detail why this is the case, and what this conditional notion of coherence exactly entails, would however lead us too far.

Having attached an interpretation to finitary gambles and their conditional upper expectations, we now proceed to do the same for more general variables. We have already argued that no direct operational meaning can be given to such variables. However, this should not be taken to imply that an uncertainty model should not be able to deal with them. In fact, they can serve as useful abstract idealisations of (sequences of) variables that can be given a direct operational meaning. In particular, we will regard any extended real variable f that is bounded below and that can be written as the pointwise limit $\lim_{n \rightarrow +\infty} f_n$ of some sequence of finitary gambles $\{f_n\}_{n \in \mathbb{N}_0}$, as an abstract idealisation of f_n for large n . We gather these limits in the set

$$\bar{\mathbb{V}}_{b,\text{lim}} := \left\{ f \in \bar{\mathbb{V}}_b : f = \lim_{n \rightarrow +\infty} f_n \text{ for some sequence } \{f_n\}_{n \in \mathbb{N}_0} \text{ in } \mathbb{V}_{\text{fin}} \right\}.$$

Since P5 applies to precisely these kinds of variables, this axiom can be justified by extending the above idealisation from the variables f to their upper expectations $\bar{E}(f|s)$. Basically, since f is an abstract idealisation of f_n for large n , $\bar{E}(f|s)$ should be an abstract idealisation of $\bar{E}(f_n|s)$ for large n . The practical benefit of this is that we can then use $\bar{E}(f|s)$ to reason about $\bar{E}(f_n|s)$ for a generic large value of n , without having to specify the specific value of n . The problem, however, is that the sequence $\{\bar{E}(f_n|s)\}_{n \in \mathbb{N}_0}$ may not converge. What we then do know for sure, however, is that as n approaches infinity, $\bar{E}(f_n|s)$ will oscillate between the limit superior and inferior of the sequence $\{\bar{E}(f_n|s)\}_{n \in \mathbb{N}_0}$. Since we want $\bar{E}(f|s)$ to serve as an idealisation of $\bar{E}(f_n|s)$ for generic large values of n , $\bar{E}(f|s)$ should therefore definitely not exceed the limit superior, as this would result in an unwarranted loss of ‘information’, in the sense that $\bar{E}(f|s)$ would be too high—this will be clarified shortly. We therefore impose P5.

The final property that we impose is P4, which states that \bar{E} should be monotone. For finitary gambles f and g , this follows easily from either of our two different interpretations for $\bar{E}(f|s)$ and $\bar{E}(g|s)$. Under a behavioural interpretation, since the reward associated with g is guaranteed to be at least as high as that of f , the same should be true for a subject’s infimum selling prices for these two gambles. Under an interpretation in terms of upper envelopes of expectations, monotonicity of the envelope is implied by the monotonicity of each of the individual expectations. If f and g are more general variables—so not necessarily finitary gambles—the motivation for P4 is that, still, higher rewards—even if abstract and idealized—should correspond to higher upper expectations. It is also worth noting that the combination of P4 and P5 implies that \bar{E} is continuous with respect to non-decreasing sequences of finitary gambles. In a measure-theoretic context, this kind of continuity is usually obtained as a consequence of the assumption of σ -additivity [1, 15].

We will show in Section 6 that P1–P5 are consistent, in the sense that if the local models \bar{Q}_s are coherent, there always is at least one global upper expectation \bar{E} satisfying P1–P5. However, there may be more than one global upper expectation \bar{E} satisfying P1–P5. In that case, the best thing to do, we believe, is to choose the most

⁶There are several versions of conditional coherence [7, 9, 10], however, in the case where variables take values in a finite set, all these different versions are mathematically equivalent.

conservative model among those that satisfy P1–P5, as choosing any other would mean adding ‘information’ that is not simply implied by our axioms. We will denote this most conservative global upper expectation by \bar{E}_A and let \underline{E}_A be its conjugate lower expectation. As we will see in Section 6, \bar{E}_A is guaranteed to exist, and furthermore coincides with a particular version of the game-theoretic upper expectation defined by Shafer and Vovk [4, 16].

Of course, in order for our definition of \bar{E}_A to make sense, we need to know what it means for an upper expectation \bar{E}' to be more conservative than some other upper expectation \bar{E} . We here take this to mean that \bar{E}' is higher than \bar{E} . So higher upper expectations are more conservative, or less informative. This is why, in our motivation for P5, we said that a too high $\bar{E}(f|s)$ corresponds to an unwarranted loss of ‘information’. That it is indeed reasonable to regard higher expectations as more conservative, can again be motivated using either of the two interpretations that we considered before. Under the behavioural interpretation, higher upper expectations mean higher selling prices, which is clearly more conservative. Using an interpretation in terms of upper envelopes of expectations, higher upper expectations correspond to larger sets of expectations, so weaker constraints on expectations, which is again less informative and hence more conservative.

Having said all this, we would like to stress that—despite our extensive use of them to motivate our axioms—none of the results that we are about to develop hinge on any particular interpretation for upper expectations. Other interpretations could also be adopted, or perhaps even no interpretation at all. All that is needed is that we agree on P1–P5 and on the fact that higher upper expectations are more conservative.

5. An Axiomatisation of \bar{E}_A

For a given imprecise probability tree \bar{Q} , let $\bar{E}_{1-2}(\bar{Q})$ denote the set of all global upper expectations satisfying P1–P2, and similarly for $\bar{E}_{1-4}(\bar{Q})$ and $\bar{E}_{1-5}(\bar{Q})$. In this section, we introduce sufficient conditions for a global upper expectation \bar{E} to be the most conservative among all upper expectations in $\bar{E}_{1-5}(\bar{Q})$. We start by considering the domain of finitary gambles and then, step by step, extend the domain and introduce additional conditions on \bar{E} , in such a way that it becomes the most conservative on this extended domain.

For any situation $x_{1:n} \in \mathcal{X}^*$ and any $(n+1)$ -measurable (global) gamble f , we use $f(x_{1:n}\cdot)$ to denote the gamble on \mathcal{X} that assumes the value $f(x_{1:n+1})$ in $x_{n+1} \in \mathcal{X}$, and then use $f(x_{1:n}X_{n+1})$ as a shorthand for the global gamble $f(x_{1:n}\cdot) \circ X_{n+1}$. The following lemma establishes compatibility with the local models in a stronger way than P1 does.

Lemma 1. *Consider any $\bar{E} \in \bar{E}_{1-2}(\bar{Q})$. Then, for any situation $x_{1:n} \in \mathcal{X}^*$ and any $(n+1)$ -measurable gamble f ,*

$$\bar{E}(f|x_{1:n}) = \bar{Q}_{x_{1:n}}(f(x_{1:n}\cdot)).$$

Proof. Fix any $x_{1:n} \in \mathcal{X}^*$ and any $(n+1)$ -measurable gamble f . Note that $f(x_{1:n}X_{n+1})\mathbb{I}_{x_{1:n}} = f\mathbb{I}_{x_{1:n}}$ and hence, because of P2, $\bar{E}(f|x_{1:n}) = \bar{E}(f\mathbb{I}_{x_{1:n}}|x_{1:n}) = \bar{E}(f(x_{1:n}X_{n+1})\mathbb{I}_{x_{1:n}}|x_{1:n}) = \bar{E}(f(x_{1:n}X_{n+1})|x_{1:n})$. P1 therefore implies that, indeed, $\bar{E}(f|x_{1:n}) = \bar{Q}_{x_{1:n}}(f(x_{1:n}\cdot))$. \square

To make sure that some particular upper expectation $\bar{E} \in \bar{E}_{1-4}(\bar{Q})$ is a most conservative upper expectation on the domain of all finitary gambles, we impose on it the following property, known as *the law of iterated upper expectations* [4, 17]:

P3⁼. $\bar{E}(f|X_{1:n}) = \bar{E}(\bar{E}(f|X_{1:n+1})|X_{1:n})$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $n \in \mathbb{N}_0$.

Proposition 2. *Consider any $\bar{E} \in \bar{E}_{1-4}(\bar{Q})$ that satisfies P3⁼. Then, for any $\bar{E}' \in \bar{E}_{1-4}(\bar{Q})$, we have that*

$$\bar{E}(f|s) \geq \bar{E}'(f|s) \text{ for all } f \in \mathbb{V}_{\text{fin}} \text{ and all } s \in \mathcal{X}^*.$$

Proof. For any $p \in \mathbb{N}_0$ and any $(p+1)$ -measurable gamble g , we let $\bar{Q}_{X_{1,p}}(g)$ be the p -measurable gamble defined by $\bar{Q}_{X_{1,p}}(g)(\omega) := \bar{Q}_{\omega^p}(g(\omega^p\cdot))$ for all $\omega \in \Omega$. Note that $\bar{Q}_{X_{1,p}}(g)$ is indeed a gamble because of (bounded) coherence [C5] and the fact that g is a gamble. Then, because both the operators \bar{E} and \bar{E}' satisfy P1 and P2, it follows from Lemma 1 that

$$\bar{E}'(g|X_{1,p}) = \bar{Q}_{X_{1,p}}(g) = \bar{E}(g|X_{1,p}). \quad (3)$$

Now fix any $f \in \mathbb{V}_{\text{fin}}$ and any $x_{1:m} \in \mathcal{X}^*$. Since f is finitary, it is n -measurable for some $n \in \mathbb{N}_0$. We can assume that $m+2 < n$ without loss of generality, because f is obviously also p -measurable for every $p \geq n$. Now, it follows from Equation (3) and the fact that $\overline{\mathbb{Q}}_{X_{1:n-1}}(f)$ is an $(n-1)$ -measurable gamble, that $\overline{\mathbb{E}}'(f|X_{1:n-1}) = \overline{\mathbb{E}}(f|X_{1:n-1})$ is an $(n-1)$ -measurable gamble. Subsequently, we can apply Equation (3) once more, where $\overline{\mathbb{E}}'(f|X_{1:n-1}) = \overline{\mathbb{E}}(f|X_{1:n-1})$ now takes the role of g , in order to find that $\overline{\mathbb{E}}'(\overline{\mathbb{E}}'(f|X_{1:n-1})|X_{1:n-2}) = \overline{\mathbb{E}}(\overline{\mathbb{E}}(f|X_{1:n-1})|X_{1:n-2})$ is an $(n-2)$ -measurable gamble. Then we can continue in the same way, to finally obtain that

$$\overline{\mathbb{E}}'(\overline{\mathbb{E}}'(\dots \overline{\mathbb{E}}'(f|X_{1:n-1}) \dots |X_{1:m+1})|x_{1:m}) = \overline{\mathbb{E}}(\overline{\mathbb{E}}(\dots \overline{\mathbb{E}}(f|X_{1:n-1}) \dots |X_{1:m+1})|x_{1:m}).$$

By repeatedly applying P3 and P4, we infer that $\overline{\mathbb{E}}'(\overline{\mathbb{E}}'(\dots \overline{\mathbb{E}}'(f|X_{1:n-1}) \dots |X_{1:m+1})|x_{1:m}) \geq \dots \geq \overline{\mathbb{E}}'(\overline{\mathbb{E}}'(f|X_{1:m+1})|x_{1:m}) \geq \overline{\mathbb{E}}'(f|x_{1:m})$. Hence, plugging this back into the previous expression, we have that

$$\overline{\mathbb{E}}'(f|x_{1:m}) \leq \overline{\mathbb{E}}(\overline{\mathbb{E}}(\dots \overline{\mathbb{E}}(f|X_{1:n-1}) \dots |X_{1:m+1})|x_{1:m}).$$

On the other hand, since $\overline{\mathbb{E}}$ satisfies P3⁼, we infer that $\overline{\mathbb{E}}(\overline{\mathbb{E}}(\dots \overline{\mathbb{E}}(f|X_{1:n-1}) \dots |X_{1:m+1})|x_{1:m}) = \dots = \overline{\mathbb{E}}(\overline{\mathbb{E}}(f|X_{1:m+1})|x_{1:m}) = \overline{\mathbb{E}}(f|x_{1:m})$. So indeed, we find that $\overline{\mathbb{E}}'(f|x_{1:m}) \leq \overline{\mathbb{E}}(f|x_{1:m})$. \square

Next, we consider the domain $\overline{\mathbb{V}}_{\text{b,lim}} \subset \overline{\mathbb{V}}$ of all extended real variables that are bounded below and that can be written as the pointwise limit of a sequence of finitary gambles. The following condition, together with P3⁼, turns out to be sufficient for an upper expectation $\overline{\mathbb{E}}$ to be a most conservative one on the domain of $\overline{\mathbb{V}}_{\text{b,lim}}$ amongst all upper expectations in $\overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$.

P6. For any $f \in \overline{\mathbb{V}}_{\text{b,lim}}$ and any $s \in \mathcal{X}^*$, there is some sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of n -measurable gambles that is uniformly bounded below and that converges pointwise to f , such that moreover $\limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s) \leq \overline{\mathbb{E}}(f|s)$.

Note that, for an upper expectation $\overline{\mathbb{E}} \in \overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$, the inequality $\limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s) \leq \overline{\mathbb{E}}(f|s)$ in P6 is in fact an equality because of P5.

Proposition 3. Consider any $\overline{\mathbb{E}} \in \overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$ that satisfies P3⁼ and P6. Then, for all $\overline{\mathbb{E}}' \in \overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$, we have that

$$\overline{\mathbb{E}}(f|s) \geq \overline{\mathbb{E}}'(f|s) \text{ for all } f \in \overline{\mathbb{V}}_{\text{b,lim}} \text{ and all } s \in \mathcal{X}^*.$$

Proof. Fix any $f \in \overline{\mathbb{V}}_{\text{b,lim}}$, any $s \in \mathcal{X}^*$ and any $\overline{\mathbb{E}}' \in \overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$. According to P6, there is some sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of n -measurable gambles that is uniformly bounded below and that converges pointwise to f in such a way that $\limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s) \leq \overline{\mathbb{E}}(f|s)$ and therefore, due to P5, also that

$$\limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s) = \overline{\mathbb{E}}(f|s). \quad (4)$$

Because all f_n are finitary gambles and both $\overline{\mathbb{E}}$ and $\overline{\mathbb{E}}'$ are upper expectations in $\overline{\mathbb{E}}_{1-4}(\overline{\mathbb{Q}})$, with $\overline{\mathbb{E}}$ additionally satisfying P3⁼, we can apply Proposition 2 to find that

$$\limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}'(f_n|s) \leq \limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}(f_n|s). \quad (5)$$

Furthermore, since $\{f_n\}_{n \in \mathbb{N}_0}$ is a sequence of finitary gambles that is uniformly bounded below and that converges pointwise to f , P5 implies that $\overline{\mathbb{E}}'(f|s) \leq \limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}'(f_n|s)$. Combining this with Equations (4) and (5), we find that, indeed, $\overline{\mathbb{E}}'(f|s) \leq \overline{\mathbb{E}}(f|s)$. \square

Finally, we consider the entire domain $\overline{\mathbb{V}}$. Now, in order for an upper expectation on $\overline{\mathbb{V}}$ to be the most conservative one, it suffices to additionally impose the following property, as we will show presently.

P7. For any $f \in \overline{\mathbb{V}}$ and any $s \in \mathcal{X}^*$,

$$\overline{\mathbb{E}}(f|s) = \inf \left\{ \overline{\mathbb{E}}(g|s) : g \in \overline{\mathbb{V}}_{\text{b,lim}} \text{ and } g \geq f \right\}.$$

Theorem 4. Consider any $\bar{E} \in \bar{\mathbb{E}}_{1-5}(\bar{Q})$ that satisfies P3[−], P6 and P7. Then, for all $\bar{E}' \in \bar{\mathbb{E}}_{1-5}(\bar{Q})$, we have that

$$\bar{E}(f|s) \geq \bar{E}'(f|s) \text{ for all } f \in \bar{V} \text{ and all } s \in \mathcal{X}^*.$$

Proof. Fix any $f \in \bar{V}$, any $s \in \mathcal{X}^*$ and any $\bar{E}' \in \bar{\mathbb{E}}_{1-5}(\bar{Q})$. According to P7, we have that

$$\bar{E}(f|s) = \inf \{ \bar{E}(g|s) : g \in \bar{V}_{b,\text{lim}} \text{ and } g \geq f \}.$$

Then, using Proposition 3, we get

$$\inf \{ \bar{E}'(g|s) : g \in \bar{V}_{b,\text{lim}} \text{ and } g \geq f \} \leq \inf \{ \bar{E}(g|s) : g \in \bar{V}_{b,\text{lim}} \text{ and } g \geq f \} = \bar{E}(f|s).$$

Since, for any $g \in \bar{V}_{b,\text{lim}}$ such that $f \leq g$, we have that $\bar{E}'(f|s) \leq \bar{E}'(g|s)$ because \bar{E}' satisfies P4, it follows that

$$\bar{E}'(f|s) \leq \inf \{ \bar{E}'(g|s) : g \in \bar{V}_{b,\text{lim}} \text{ and } g \geq f \} \leq \bar{E}(f|s),$$

which completes the proof. \square

We conclude that, according to Theorem 4, if there is some upper expectation $\bar{E} \in \bar{\mathbb{E}}_{1-5}(\bar{Q})$ that satisfies P3[−], P6 and P7, it must be the unique most conservative upper expectation in $\bar{\mathbb{E}}_{1-5}(\bar{Q})$, and it must therefore be the global belief model \bar{E}_Λ that we are after. So it now only remains to show that there is at least one—then necessarily unique—global upper expectation in $\bar{\mathbb{E}}_{1-5}(\bar{Q})$ that additionally satisfies P3[−], P6 and P7.

6. Game-Theoretic Upper Expectations

In this section, we show that a particular version of the game-theoretic upper expectation developed by Shafer and Vovk [4, 5] belongs to $\bar{\mathbb{E}}_{1-5}(\bar{Q})$ and furthermore satisfies P3[−], P6 and P7, which will then—because of Theorem 4—imply the existence of the unique \bar{E}_Λ we are in pursuit of. This game-theoretic upper expectation relies on the concept of a *supermartingale*,⁷ which is a capital process—the evolution of a subject’s capital—that is obtained by betting against a ‘forecasting system’. The forecasting system—called ‘Forecaster’ in Shafer and Vovk’s framework—determines for each situation a collection of allowable bets that a subject—called ‘Skeptic’ in Shafer and Vovk’s framework—can choose from. These allowable bets define Forecaster’s uncertainty model and are also captured by our notion of an imprecise probability tree. We should point out though that Shafer and Vovk also allow for more general settings where local models need not be boundedly coherent, where state spaces can be infinite and where Forecaster his commitments—that is, the allowable bets for Skeptic—in each situation need not be announced beforehand, but may also depend on previous moves by Skeptic. In that respect, our discussion here will be more particular and limited.

As before, we start from an imprecise probability tree that specifies a local coherent upper expectation \bar{Q}_s on $\mathcal{L}(\mathcal{X})$ for every situation $s \in \mathcal{X}^*$. However, in their most recent work [5, Part II], Shafer and Vovk start from local models \bar{E}_s that are defined on the extended domain $\bar{\mathcal{L}}(\mathcal{X})$ of all extended real local variables, and they require these local models \bar{E}_s to satisfy a modified version of the bounded coherence axioms C1–C3, generalised to extended real variables. Concretely, for any $s \in \mathcal{X}^*$, they impose the following properties on the upper expectation \bar{E}_s on $\bar{\mathcal{L}}(\mathcal{X})$ that models the corresponding local beliefs:

- E1. $\bar{E}_s(c) = c$ for all $c \in \mathbb{R}$;
- E2. $\bar{E}_s(f + g) \leq \bar{E}_s(f) + \bar{E}_s(g)$ for all $f, g \in \bar{\mathcal{L}}(\mathcal{X})$;
- E3. $\bar{E}_s(\lambda f) = \lambda \bar{E}_s(f)$ for all $\lambda \in \mathbb{R}_{>0}$ and all $f \in \bar{\mathcal{L}}(\mathcal{X})$;

⁷Shafer and Vovk drew inspiration from Jean Ville, whose work [18] constitutes the first major steps towards a theory of probability where (super)martingales, instead of probability measures, function as the fundamental, primary objects.

E4. $f \leq g \Rightarrow \bar{E}_s(f) \leq \bar{E}_s(g)$ for all $f, g \in \bar{\mathcal{L}}(\mathcal{X})$.

E5. $\lim_{n \rightarrow +\infty} \bar{E}_s(f_n) = \bar{E}_s(\lim_{n \rightarrow +\infty} f_n)$ for any non-decreasing and non-negative sequence $\{f_n\}_{n \in \mathbb{N}_0}$ in $\bar{\mathcal{L}}(\mathcal{X})$.

Hence, in order to place ourselves squarely in their framework, we need to extend the domain of our local models \bar{Q}_s from $\mathcal{L}(\mathcal{X})$ to $\bar{\mathcal{L}}(\mathcal{X})$ in such a way that they satisfy E1–E5. For any $f \in \bar{\mathcal{L}}(\mathcal{X})$ and any $c \in \mathbb{R}$, let $f^{\wedge c}$ and $f^{\vee c}$ be the variables in $\bar{\mathcal{L}}(\mathcal{X})$ respectively defined by $f^{\wedge c}(x) := \min\{f(x), c\}$ and $f^{\vee c}(x) := \max\{f(x), c\}$ for all $x \in \mathcal{X}$. We propose the following two continuity properties for an upper expectation \bar{Q}_s^{\uparrow} on $\bar{\mathcal{L}}(\mathcal{X})$ that extends a coherent upper expectation \bar{Q}_s on $\mathcal{L}(\mathcal{X})$:

C8. $\bar{Q}_s^{\uparrow}(f) = \lim_{c \rightarrow +\infty} \bar{Q}_s^{\uparrow}(f^{\wedge c})$ for all $f \in \bar{\mathcal{L}}_b(\mathcal{X})$;

C9. $\bar{Q}_s^{\uparrow}(f) = \lim_{c \rightarrow -\infty} \bar{Q}_s^{\uparrow}(f^{\vee c})$ for all $f \in \bar{\mathcal{L}}(\mathcal{X})$;

Properties C8 and C9 are respectively called *continuity with respect to upper* and *lower cuts* (or, alternatively, bounded above and below support). According to [19, Proposition 10 (ii)–(iii)] such an extension \bar{Q}_s^{\uparrow} always exists and is moreover unique. The upper expectation \bar{Q}_s^{\uparrow} is also an appropriate extension, in the sense that it satisfies E1–E5; this follows from [19, Proposition 10] together with [19, Proposition 1].⁸ Hence, in order to adhere to Shafer and Vovk’s approach, we can—and will—adopt this particular extension \bar{Q}_s^{\uparrow} of \bar{Q}_s for each situation $s \in \mathcal{X}^*$ in the following game-theoretic reasoning.

A map $\mathcal{M} : \mathcal{X}^* \rightarrow \mathbb{R}$ is called a *supermartingale* if it satisfies $\bar{Q}_s^{\uparrow}(\mathcal{M}(s \cdot)) \leq \mathcal{M}(s)$ for all $s \in \mathcal{X}^*$, where $\mathcal{M}(s \cdot)$ denotes the variable in $\bar{\mathcal{L}}(\mathcal{X})$ that assumes the value $\mathcal{M}(sx)$ for all $x \in \mathcal{X}$. As mentioned earlier, a supermartingale \mathcal{M} can be interpreted as a capital process that results from betting against a forecasting system. In order for this to make sense, such an interpretation will only be attached to supermartingales that are *bounded below*, meaning that there is some real c such that $\mathcal{M}(s) \geq c$ for all $s \in \mathcal{X}^*$. This represents the constraint that we can never borrow an infinite or even unbounded amount of money. To see that a bounded below supermartingale \mathcal{M} can indeed be interpreted as described above, we first need to know how a forecasting system allows us to play. In our case, the allowable bets in any given situation $s \in \mathcal{X}^*$ are those variables $f \in \bar{\mathcal{L}}_b(\mathcal{X})$ for which $\bar{Q}_s^{\uparrow}(f) \leq 0$ —we will come back to this shortly. Now, if $\mathcal{M}(s)$ is real, then the condition that $\bar{Q}_s^{\uparrow}(\mathcal{M}(s \cdot)) \leq \mathcal{M}(s)$ is equivalent to the condition that $\bar{Q}_s^{\uparrow}(\mathcal{M}(s \cdot) - \mathcal{M}(s)) \leq 0$ because \bar{Q}_s^{\uparrow} satisfies constant additivity on the domain $\bar{\mathcal{L}}_b(\mathcal{X})$; see [19, Proposition 3(E8)]. Hence, the incremental change $\mathcal{M}(s \cdot) - \mathcal{M}(s)$ of the capital process \mathcal{M} in the situation s indeed represents an allowable bet. If on the other hand $\mathcal{M}(s) = +\infty$ (the case where $\mathcal{M}(s) = -\infty$ is impossible because we assumed \mathcal{M} to be bounded below), the condition that $\bar{Q}_s^{\uparrow}(\mathcal{M}(s \cdot)) \leq \mathcal{M}(s)$ does not impose any constraints on $\mathcal{M}(s \cdot)$. The interpretation that the increment $\mathcal{M}(s \cdot) - \mathcal{M}(s)$ represents an allowable bet is, in this case, somewhat questionable. However, one could argue that, since the process \mathcal{M} has reached the maximum possible value—that is, $+\infty$ —in the situation s , it cannot increase any further, hence there is no need for any constraints on the process’ future value.

That the allowable bets f in some situation $s \in \mathcal{X}^*$ are characterised by the condition that $\bar{Q}_s^{\uparrow}(f) \leq 0$, agrees with the behavioural interpretation of (boundedly) coherent upper expectations. To see this, we limit ourselves to the domain $\mathcal{L}(\mathcal{X})$ and recall that $\bar{Q}_s^{\uparrow}(f) = \bar{Q}_s(f)$, for any $f \in \mathcal{L}(\mathcal{X})$, can then be interpreted as a subject’s infimum selling price for the gamble f . So the subject, being the forecasting system in this case, is willing to offer the gambles that have negative⁹ upper expectation. We can then take the subject up on his commitments by accepting its offer and selecting one such gamble. Moreover, since the upper expectations \bar{Q}_s are assumed to be (boundedly) coherent, the associated sets of available gambles will also satisfy certain rationality axioms. We refer to an earlier paper by one of us [11] for a more elaborate discussion of how Walley’s behavioural approach can be related (in a slightly different context) to the game-theoretic one proposed by Shafer and Vovk.

Our definition of the global game-theoretic upper expectation—and the one considered by Shafer and Vovk—will be based on bounded below supermartingales, not only because it allows for a sensible interpretation as we have explained above, but also because of more technical reasons, which we discuss in [19, Section 8]. So let

⁸To see this, note that the definition of an upper expectation in Reference [19] is different from the one we adopt here.

⁹In the interest of brevity, we ignore in our discussion of the interpretation the special case where we have zero upper expectation. See [11] for more details about this special case.

$\overline{\mathbb{M}}_b$ be the set of all such bounded below supermartingales. For any $\mathcal{M} \in \overline{\mathbb{M}}_b$, we then use $\liminf \mathcal{M}$ to denote the extended real global variable that assumes the value $\liminf_{n \rightarrow +\infty} \mathcal{M}(\omega^n)$ in each $\omega \in \Omega$. We also write, for any situation $s \in \mathcal{X}^*$ and any two f and g in $\overline{\mathbb{V}}$, that $f \geq_s g$ if $f(\omega) \geq g(\omega)$ for all paths $\omega \in \Gamma(s)$. The global game-theoretic upper expectation $\overline{\mathbb{E}}_V: \overline{\mathbb{V}} \times \mathcal{X}^* \rightarrow \overline{\mathbb{R}}$ can now be defined as

$$\overline{\mathbb{E}}_V(f|s) := \inf \left\{ \mathcal{M}(s) : \mathcal{M} \in \overline{\mathbb{M}}_b \text{ and } \liminf \mathcal{M} \geq_s f \right\} \text{ for all } f \in \overline{\mathbb{V}} \text{ and all } s \in \mathcal{X}^*.$$

The corresponding game-theoretic lower expectation $\underline{\mathbb{E}}_V$ is again defined by conjugacy; so $\underline{\mathbb{E}}_V(f|s) := -\overline{\mathbb{E}}_V(-f|s)$ for all $f \in \overline{\mathbb{V}}$ and all $s \in \mathcal{X}^*$. For a concrete and detailed motivation and interpretation for the operator $\overline{\mathbb{E}}_V$, we refer to [4, 5, 19]. Essentially, the upper expectation $\overline{\mathbb{E}}_V(f|x_{1:n})$ can once more be seen as an infimum selling price for a variable $f \in \overline{\mathbb{V}}$ conditional on the fact that we have observed some history $X_1 = x_1, X_2 = x_2, \dots, X_n = x_n$ of the process. It is based on the assumption that we should agree to sell f for any starting capital $\mathcal{M}(x_{1:n})$ such that, by appropriately betting against the forecasting system from $x_{1:n}$ onwards, we manage to end up with a higher capital than what we would receive from f , regardless of the path $\omega \in \Omega$ taken by the process. Indeed, it is reasonable to state that such a starting capital $\mathcal{M}(x_{1:n})$ is worth more than the uncertain (possibly negative) payoff corresponding to f . The global game-theoretic upper expectation $\overline{\mathbb{E}}_V(f|x_{1:n})$ is the infimum over all such starting capitals.

As we will briefly discuss in Section 7, the global game-theoretic upper expectation $\overline{\mathbb{E}}_V$ satisfies various, remarkably powerful properties. In particular, it satisfies P1–P5, P3⁼, P6 and P7, which allows us to state our main result: that $\overline{\mathbb{E}}_V$ is the unique most conservative upper expectation $\overline{\mathbb{E}}_A$ in $\overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$. Our proofs for these results are heavily based on our work in [19]. The setting there is entirely the same as the one described in this section, with the exception that we generally do not impose C9 on the local models in [19]. It will soon be clarified why we have made this distinction; for the moment, it suffices to see that the setting in [19] is more general and therefore, that the results in [19] obviously remain valid here.

Proposition 5. $\overline{\mathbb{E}}_V$ is an element of $\overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$ and furthermore satisfies P3⁼, P6 and P7.

Proof. We prove that $\overline{\mathbb{E}}_V$ satisfies P1–P5, P3⁼, P6 and P7. Axiom P1 follows immediately from [19, Proposition 14] which, in particular, says that $\overline{\mathbb{E}}(f|x_{1:n}) = \overline{\mathbb{Q}}_{x_{1:n}}(f(x_{1:n}, \cdot))$ for any situation $x_{1:n} \in \mathcal{X}^*$ and any $(n+1)$ -measurable gamble f . Indeed, for any $h \in \mathcal{L}(\mathcal{X})$, $h(X_{n+1})$ is an $(n+1)$ -measurable gamble, so it follows from [19, Proposition 14] that $\overline{\mathbb{E}}_V(h(X_{n+1})|x_{1:n}) = \overline{\mathbb{Q}}_{x_{1:n}}(h)$. To prove P2, observe that $\liminf \mathcal{M} \geq_s f$ if and only if $\liminf \mathcal{M} \geq_s f \mathbb{I}_s$ for all $f \in \overline{\mathbb{V}}$ and all $\mathcal{M} \in \overline{\mathbb{M}}_b$. The desired equality therefore follows directly from the definition of $\overline{\mathbb{E}}_V$. P3⁼, and hence also P3, follows immediately from [19, Theorem 15]. Property P4 follows from [19, Proposition 13(V4)], which says that $\overline{\mathbb{E}}_V$ satisfies the stronger property that $\overline{\mathbb{E}}(f|s) \leq \overline{\mathbb{E}}(g|s)$ if $f \leq_s g$ for all $f, g \in \overline{\mathbb{V}}$ and all $s \in \mathcal{X}^*$.

Furthermore, P6 and P7 follow from, respectively, [19, Theorem 32] and [19, Proposition 28]. Note that, in these results, the definition of $\overline{\mathbb{V}}_{b, \lim}$ seems to be slightly more general, since it there also includes pointwise limits of possibly *extended* real finitary variables. However, according to [19, Proposition 30], any such limit of extended real finitary variables is in fact also a pointwise limit of n -measurable, and therefore finitary, gambles. So both our definitions of $\overline{\mathbb{V}}_{b, \lim}$ are equivalent, therefore indeed allowing us to apply the results in [19] here.

Finally, Property P5 follows from [19, Corollary 25], which says that, for any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ in $\overline{\mathbb{V}}_b$ that is uniformly bounded below, $\overline{\mathbb{E}}_V(\liminf_{n \rightarrow +\infty} f_n|s) \leq \liminf_{n \rightarrow +\infty} \overline{\mathbb{E}}_V(f_n|s)$. Indeed, in the special case that $\{f_n\}_{n \in \mathbb{N}_0}$ is a sequence of finitary gambles that is uniformly bounded below and converges pointwise to some variable $f \in \overline{\mathbb{V}}_b$, this implies that $\overline{\mathbb{E}}_V(f|s) \leq \liminf_{n \rightarrow +\infty} \overline{\mathbb{E}}_V(f_n|s) \leq \limsup_{n \rightarrow +\infty} \overline{\mathbb{E}}_V(f_n|s)$. \square

Theorem 6. *There is a unique most conservative upper expectation $\overline{\mathbb{E}}_A$ in $\overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$ and it is the global game-theoretic upper expectation $\overline{\mathbb{E}}_V$. Furthermore, if an upper expectation $\overline{\mathbb{E}}$ in $\overline{\mathbb{E}}_{1-5}(\overline{\mathbb{Q}})$ satisfies P3⁼, P6 and P7, then it is equal to this most conservative upper expectation $\overline{\mathbb{E}}_A$, and therefore also equal to $\overline{\mathbb{E}}_V$.*

Proof. Immediately from Theorem 4 and Proposition 5. \square

7. On the Implications of Theorem 6

Shafer and Vovk define their game-theoretic upper expectations using supermartingales. Their definition has a constructive flavour and can be given a clear interpretation in terms of capital processes and betting behaviour. However, it requires that one allows unbounded and even infinite-valued bets, which we find more

questionable from a behavioural point of view. Furthermore, a complete axiomatisation of \bar{E}_V is, to the best of our knowledge, absent from the literature. Theorem 6 addresses both of these issues. First, it provides an abstract axiomatisation for \bar{E}_V using P3⁼, P6 and P7 in addition to P1–P5. Most importantly, however, Theorem 6 provides an alternative definition—and interpretation—for \bar{E}_V as the most conservative upper expectation \bar{E}_A under a limited set of intuitive properties: P1–P5. This strengthens the argument in favour of using \bar{E}_V as a global upper expectation, because it can now be motivated from both a game-theoretic point of view and from a purely axiomatic point of view. Readers need not even be familiar with the concepts of game-theoretic probability in order to use \bar{E}_V as global upper expectation. Put simply, they only have to agree on the axioms P1–P5. And even if they wish to impose additional axioms, then \bar{E}_V would still serve as a conservative upper bound for their desired (upper) expectation.

Still, recalling our assumptions about the extended local models \bar{Q}_s^\dagger , one could regret the fact that we are only considering a special case of what is generally being done by Shafer and Vovk; an upper expectation \bar{Q}_s^\dagger on $\mathcal{L}(\mathcal{X})$ satisfying E1–E5 does not necessarily have to be an extension of a coherent upper expectation \bar{Q}_s on $\mathcal{L}(\mathcal{X})$ satisfying C8 and C9, see [19, Example 1]. However, this restriction in generality actually turns out to be rather desirable. Indeed, as we show in [19], compatibility of local and global game-theoretic upper expectations—a property that we find essential—can only be obtained if we limit ourselves to extended local models \bar{Q}_s^\dagger of the form that we have described earlier, that is, boundedly coherent and satisfying C8 and C9.

That said, as mentioned before, we do not a priori impose C9 on the local models in [19]. In that paper, we mainly focus on the mathematical properties of global game-theoretic upper expectations and most of these do not require C9. Furthermore, by dropping this axiom, we can allow for more general local models than Shafer and Vovk, who impose E1–E5. Nevertheless, in this less mathematical and more philosophical contribution, we do impose C9 on the local models because, practically speaking, we always want local and global belief models to be compatible; see for example Section 4, where we considered it self-evident that P1 should be desirable. Note however that Theorem 6 actually does not depend on whether the extended local models \bar{Q}_s^\dagger satisfy C9. On the one hand, because \bar{E}_A only depends on the (boundedly) coherent local models \bar{Q}_s , and on the other hand, because \bar{E}_V only depends on what the values of the extended local models \bar{Q}_s^\dagger are on the restricted domain $\mathcal{L}_b(\mathcal{X})$. The latter can easily be seen from the fact that \bar{E}_V is defined using supermartingales that are bounded below; we refer to [19] for more details.

8. Additional Properties

Of course, there is more to an uncertainty model than an axiomatisation or an interpretation, however compelling they may be. In order to be practically useful, its mathematical properties should also be sufficiently powerful such that inferences can be computed or approximated efficiently. For instance, the popularity of the Lebesgue integral as a tool for defining expected values, is in part due to its strong continuity properties (e.g. the Dominated Convergence Theorem [1, 15]). What may perhaps be somewhat surprising, is that despite the simplicity of our axioms P1–P5, our most conservative model \bar{E}_A also scores well on this count. For example, it satisfies the following generalisation of bounded coherence. Its proof, as well as the proofs of the other results in this section, follows immediately from the fact that \bar{E}_A coincides with \bar{E}_V and the fact that these properties hold for \bar{E}_V , as is shown in [19].

Proposition 7. [19, Proposition 13 and Theorem 15] *For all $f, g \in \bar{V}$, all $\lambda \in \mathbb{R}_{\geq 0}$, all $\mu \in \mathbb{R}$, all $n \in \mathbb{N}_0$ and all $s \in \mathcal{X}^*$, \bar{E}_A satisfies*

$$V1. \inf_{\omega \in \Gamma(s)} f(\omega) \leq \underline{E}_A(f|s) \leq \bar{E}_A(f|s) \leq \sup_{\omega \in \Gamma(s)} f(\omega);$$

$$V2. \bar{E}_A(f + g|s) \leq \bar{E}_A(f|s) + \bar{E}_A(g|s);$$

$$V3. \bar{E}_A(\lambda f|s) = \lambda \bar{E}_A(f|s);$$

$$V4. \bar{E}_A(f + \mu|s) = \bar{E}_A(f|s) + \mu.$$

$$V5. \bar{E}_A(f|X_{1:n}) = \bar{E}_A(\bar{E}_A(f|X_{1:n+1})|X_{1:n}).$$

Another important result, better known as Fatou's Lemma, shows that the upper bound imposed by property P5 for a global upper expectation $\bar{E}(f|s)$ of a variable $f \in \bar{\mathcal{V}}_{b,\text{lim}}$ need not be attained by the most conservative upper expectation \bar{E}_A in $\bar{\mathcal{E}}_{1-5}(\bar{\mathcal{Q}})$. Interestingly enough, it can be replaced by a generally tighter one, which is established in the following lemma.

Lemma 8. [19, Corollary 25] *For any $s \in \mathcal{X}^*$ and any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ in $\bar{\mathcal{V}}_b$ that is uniformly bounded below, $\bar{E}_A(\liminf_{n \rightarrow +\infty} f_n|s) \leq \liminf_{n \rightarrow +\infty} \bar{E}_A(f_n|s)$.*

The following theorem states that \bar{E}_A satisfies continuity with respect to non-decreasing sequences of bounded below variables and with respect to non-increasing sequences of finitary, bounded above variables.

Theorem 9. [19, Theorem 23 and Theorem 31]

- i. *For any $s \in \mathcal{X}^*$ and any non-decreasing sequence $\{f_n\}_{n \in \mathbb{N}_0}$ in $\bar{\mathcal{V}}_b$ that converges point-wise to a variable $f \in \bar{\mathcal{V}}_b$, we have that $\bar{E}_A(f|s) = \lim_{n \rightarrow +\infty} \bar{E}_A(f_n|s)$.*
- ii. *For any $s \in \mathcal{X}^*$ and any non-increasing sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of finitary, bounded above variables that converges point-wise to a variable $f \in \bar{\mathcal{V}}$, we have that $\bar{E}_A(f|s) = \lim_{n \rightarrow +\infty} \bar{E}_A(f_n|s)$.*

Apart from the brief overview above, $\bar{E}_V = \bar{E}_A$ also satisfies a weak and a strong law of large numbers, a law of the iterated logarithm, Lévy's zero-one law, and many more surprisingly strong properties; we refer the interested reader to Shafer and Vovk's work [5, 16] and to ours [19, 20].

9. Measure-Theoretic Belief Models

In contradistinction with the game-theoretic approach, where the central concept is that of a supermartingale, the measure-theoretic approach relies on σ -additive probability measures in order to describe the global dynamics of an uncertain process. In a precise probabilities context, this global measure is usually obtained by extending the information that is included in the (precise) local models, now represented by probability mass functions on \mathcal{X} , using Ionescu Tulcea's extension theorem [1, Theorem 2.9.2]. This consists of two consecutive steps. First, a probability measure on the algebra generated by all cylinder events is constructed in a straightforward manner, by combining local (conditional) probabilities using the well-known product rule. Subsequently, this probability measure is extended to a probability measure on the σ -algebra \mathcal{F} generated by all cylinder events, which, according to Carathéodory's extension theorem [6, Theorem 1.7], can be done uniquely. Linear expectations are then typically obtained by Lebesgue integration and by restricting the domain to (random) variables that are measurable with respect to the σ -algebra \mathcal{F} .

This traditional measure-theoretic approach has a few shortcomings however. First of all, it only considers the special case where each local model is represented by a single probability mass function on \mathcal{X} . Hence, we cannot directly apply it here in our more general, imprecise probabilities setting. Secondly, the domain of the expectation operators is restricted to measurable (random) variables and therefore, no explicit information is given about non-measurable variables. Thirdly, by condensing all local probability mass functions into a single global probability measure, we disregard part of the information that is incorporated in the local probability mass functions. More specifically, the part that describes the process' dynamics conditional on the fact that its path will lie in a set of probability zero (with respect to the global probability measure). Indeed, conditional probabilities and expectations are typically derived from the global probability measure using Bayes' rule, yet, for those cases where the conditioning event has probability zero, the conditional probabilities and expectations become ill-defined or are chosen arbitrarily; see Appendix A and [1, 3].

In this section, we aim to adapt the traditional measure-theoretic approach to our setting and deal with these problems in a satisfactory manner. As a first step, we will limit ourselves to the precise case and try to resolve the last two issues mentioned above. The global probability measure on \mathcal{F} , which is the starting point for the traditional approach, will be replaced by a 'conditional probability measure' in the sense of [12, Definition 6]. Essentially, such a conditional measure specifies a probability measure on \mathcal{F} for each situation $s \in \mathcal{X}^*$, which describes the global dynamics of the uncertain process if we are sure that the path ω taken by the process will pass through this situation s . This will enable us to also meaningfully define linear expectations conditional

on situations that have probability zero. Subsequently, we will extend the domain of these conditional linear expectation operators to all global variables $\bar{\mathbb{V}}$ by taking a particular upper integral leading to an upper expectation. As a final step, we allow for imprecision in the local models by considering a set \mathbb{P}_s of probability mass functions on \mathcal{X} in each situation s . We can then apply the foregoing course of reasoning for each possible selection $p: s \in \mathcal{X}^* \mapsto p(\cdot|s)$ of local probability mass functions, where, in each situation s , the selected probability mass function $p(\cdot|s)$ belongs to \mathbb{P}_s . Hence, instead of a single upper expectation operator, we now obtain a set of them. The global measure-theoretic upper expectation \bar{E}_{meas} will then be defined as the upper envelope of all such ‘compatible’ upper expectation operators, similar to what is done on a local level.

The resulting operator \bar{E}_{meas} will satisfy P1–P5, which implies by Theorem 6 that the upper expectation \bar{E}_Λ —and therefore also $\bar{E}_\mathbb{V}$ —is an upper bound for \bar{E}_{meas} . Moreover, if all local models are precise, all upper expectation operators \bar{E}_Λ , $\bar{E}_\mathbb{V}$ and \bar{E}_{meas} will be shown to coincide.

Precise Probability Trees

In this subsection, we consider a *precise probability tree* p , that is, a map $p: s \in \mathcal{X}^* \mapsto p(\cdot|s)$ where each $p(\cdot|s)$ is a probability mass function on \mathcal{X} . Recalling our considerations in Section 2, each mass function $p(\cdot|s)$ is in a one-to-one correspondence with a linear expectation \bar{Q}_s on $\mathcal{L}(\mathcal{X})$, so we could just as well have started from an imprecise probability tree \bar{Q} where all local models \bar{Q}_s are linear on $\mathcal{L}(\mathcal{X})$. However, we prefer to start with p because it is in accordance with the traditional measure-theoretic view towards uncertainty, where probability measures or charges—and therefore also probability mass functions—are considered to be primary objects.

As mentioned above, we aim to construct a conditional probability measure that allows us to condition on any situation $s \in \mathcal{X}^*$ meaningfully. To do so, we follow the same reasoning as proposed in [12, Chapter 3]. First, we define, for any $x_{1:n} \in \mathcal{X}^*$, a probability charge $P(\cdot|x_{1:n})$ on the algebra generated by all cylinder events. This can be done in a simple and intuitive manner [12, Lemma 14]: for any $m \in \mathbb{N}_0$ and any $C \subseteq \mathcal{X}^m$, we let

$$P(C|x_{1:n}) := \sum_{z_{1:m} \in C} P(z_{1:m}|x_{1:n}), \text{ where } P(z_{1:m}|x_{1:n}) := \begin{cases} \prod_{i=n}^{m-1} p(z_{i+1}|z_{1:i}) & \text{if } n < m \text{ and } z_{1:n} = x_{1:n} \\ 1 & \text{if } n \geq m \text{ and } z_{1:m} = x_{1:m} \\ 0 & \text{otherwise,} \end{cases} \quad (6)$$

where we use the shorthand notation $P(C|x_{1:n})$ for any $C \subseteq \mathcal{X}^m$ to mean $P(\cup_{z_{1:m} \in C} \Gamma(z_{1:m})|x_{1:n})$, and where we did not make any particular distinction between the notations $z_{1:m}$ and $\{z_{1:m}\}$. Since each $P(\cdot|x_{1:n})$ defined in this way is a (finitely additive) probability charge on the algebra generated by all cylinder events [12, Chapter 3], and therefore automatically a σ -additive probability measure on this algebra [3, Theorem 2.3], we can use Carathéodory’s extension theorem [6, Theorem 1.7] to extend each $P(\cdot|x_{1:n})$ to a unique σ -additive measure on the σ -algebra \mathcal{F} generated by all cylinder events. Then, according to [12, Theorem 15], we can gather all such probability measures $P(\cdot|x_{1:n})$ to obtain a single *conditional probability measure* $P: \mathcal{F} \times \mathcal{X}^* \rightarrow \mathbb{R}$ [12, Definition 6], which is *coherent* in the sense of [12, Definition 5]. This notion of coherence is related to the bounded coherence condition in Section 2; both notions represent rational behaviour, but, the version used in [12, Definition 5] is adapted to probability measures and involves conditioning on events.

Before defining our global measure-theoretic (upper) expectation, we first recall the following concepts about measurability and integration; we refer to Appendix A for a more elaborate introduction. We say that a global variable $f \in \bar{\mathbb{V}}$ is \mathcal{F} -measurable if, for all $c \in \mathbb{R}$, the set $\{\omega \in \Omega: f(\omega) \leq c\}$ is \mathcal{F} -measurable, meaning that it is an element of the σ -algebra \mathcal{F} . We will commonly use the properties that \mathcal{F} -measurability of variables is preserved under taking maxima or minima (this follows immediately from the fact that the class of all \mathcal{F} -measurable sets is closed under countable unions and intersections [21, Section II.4]) and that \mathcal{F} -measurability is preserved under pointwise convergence of variables [21, Theorem II.4.2]. Moreover, recall that the Lebesgue integral $\int_\Omega f dP_u$ [1, 3, 15] with respect to any (unconditional) probability measure P_u on \mathcal{F} always exists if f is \mathcal{F} -measurable and either bounded or non-negative. However, for a general \mathcal{F} -measurable variable $f \in \bar{\mathbb{V}}$, the integral only exists if $\min\{\int_\Omega f^+ dP_u, \int_\Omega f^- dP_u\} < +\infty$, where $f^+ := f \vee 0$ and $f^- := -f \wedge 0$. In that case, we let $\int_\Omega f dP_u := \int_\Omega f^+ dP_u - \int_\Omega f^- dP_u$. Then, because the Lebesgue integral is trivially real-valued for a bounded \mathcal{F} -measurable variable (this can for instance be deduced from M1 and M2 in Appendix A), it can easily be seen that $\int_\Omega f dP_u$ also always exists for an \mathcal{F} -measurable variable $f \in \bar{\mathbb{V}}_b$.

We now adopt an approach that is similar to what is done in a traditional context, defining expectations as Lebesgue integrals. An important difference, however, is that we integrate with respect to a probability measure that depends on the conditioning event. Concretely, we define the *measure-theoretic expectation* $E_{\text{meas},p}(f|s)$ of an \mathcal{F} -measurable variable $f \in \overline{\mathbb{V}}$ conditional on the situation $s \in \mathcal{X}^*$ by the Lebesgue integral $\int_{\Omega} f dP|_s$ with respect to the (unconditional) probability measure $P|_s := P(\cdot|s)$, provided that $\int_{\Omega} f dP|_s$ exists. The expectation $E_{\text{meas},p}(\cdot|\square)$ then corresponds to the unconditional expectation as introduced in traditional measure theory. The subscript p in $E_{\text{meas},p}$ is used to remind us that $E_{\text{meas},p}$ depends on the precise probability tree p . Since $E_{\text{meas},p}$ is defined in terms of the Lebesgue integral, it inherits several properties of this integral; again, we refer to Appendix A for more information. The following is a subset of these properties—simplified and applied to $E_{\text{meas},p}$ —that we will need in the main text. We take $s \in \mathcal{X}^*$ to be any situation and implicitly assume that the considered expectations exist.

M1. $E_{\text{meas},p}(af + b|s) = aE_{\text{meas},p}(f|s) + b$ for all \mathcal{F} -measurable variables $f \in \overline{\mathbb{V}}$ and all $a, b \in \mathbb{R}$.

M2. $f \leq g \Rightarrow E_{\text{meas},p}(f|s) \leq E_{\text{meas},p}(g|s)$ for all \mathcal{F} -measurable variables $f, g \in \overline{\mathbb{V}}$.

M8. Consider any non-decreasing sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of \mathcal{F} -measurable variables in $\overline{\mathbb{V}}$. If there is an \mathcal{F} -measurable function f^* such that $E_{\text{meas},p}(f^*|s) > -\infty$ and $f_n \geq f^*$ for all $n \in \mathbb{N}_0$, then

$$\lim_{n \rightarrow +\infty} E_{\text{meas},p}(f_n|s) = E_{\text{meas},p}(f|s) \text{ where } \lim_{n \rightarrow +\infty} f_n = f.$$

M9. Consider any non-increasing sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of \mathcal{F} -measurable variables in $\overline{\mathbb{V}}$. If there is an \mathcal{F} -measurable function f^* such that $E_{\text{meas},p}(f^*|s) < +\infty$ and $f_n \leq f^*$ for all $n \in \mathbb{N}_0$, then

$$\lim_{n \rightarrow +\infty} E_{\text{meas},p}(f_n|s) = E_{\text{meas},p}(f|s) \text{ where } \lim_{n \rightarrow +\infty} f_n = f.$$

Now that our measure-theoretic definitions are in place, we aim at relating the expectation $E_{\text{meas},p}$ to our upper expectation operator \overline{E}_A discussed in the earlier sections of this paper. For the sake of clarity, we will use $\overline{E}_{A,p}$ to denote our upper expectation \overline{E}_A corresponding to the imprecise probability tree \overline{Q} associated with p ; so each \overline{Q}_s is defined by $\overline{Q}_s(f) := \sum_{x \in \mathcal{X}} f(x)p(x|s)$ for all $f \in \mathcal{L}(\mathcal{X})$. We first focus on the domain of all \mathcal{F} -measurable variables in $\overline{\mathbb{V}}_b$. This seems appropriate because, though we cannot guarantee existence of the integral $\int_{\Omega} f dP|_s$ —and therefore also $E_{\text{meas},p}(f|s)$ —for a general \mathcal{F} -measurable variable $f \in \overline{\mathbb{V}}$, it can always be guaranteed for an \mathcal{F} -measurable variable $f \in \overline{\mathbb{V}}_b$, as we have explained earlier. It will turn out that both operators $E_{\text{meas},p}$ and $\overline{E}_{A,p}$ coincide on the domain of all \mathcal{F} -measurable variables in $\overline{\mathbb{V}}_b$.

To prove this, we base ourselves on [5, Theorem 9.3]; a result by Shafer and Vovk—yet the essential idea originates from Jean Ville—that shows that game-theoretic and measure-theoretic (upper) expectations coincide on \mathcal{F} -measurable gambles, in case we are considering precise probability trees. Proposition 10 below establishes the equality of $\overline{E}_{A,p}$ and $E_{\text{meas},p}$ on this domain and therefore, due to Theorem 6, also of \overline{E}_V and $E_{\text{meas},p}$. Though the proof of Proposition 10 is based on the same principles and ideas as those that underly the proof of [5, Theorem 9.3], they are stated in a somewhat different setting and, for an unexperienced reader, it may therefore not be entirely clear how both are related to each other. Moreover, our result here also applies to conditional expectations, whereas [5, Theorem 9.3] only considers the unconditional case. We have therefore decided to include a self-contained proof of Proposition 10 in Appendix B, starting from Ville's martingale theorem.

Proposition 10. *For any \mathcal{F} -measurable $f \in \overline{\mathbb{V}}$ and any $s \in \mathcal{X}^*$, we have that $E_{\text{meas},p}(f|s) = \overline{E}_{A,p}(f|s)$.*

The next result extends the domain of the equality to all \mathcal{F} -measurable variables in $\overline{\mathbb{V}}_b$.

Theorem 11. *For any \mathcal{F} -measurable $f \in \overline{\mathbb{V}}_b$ and any $s \in \mathcal{X}^*$, we have that $E_{\text{meas},p}(f|s) = \overline{E}_{A,p}(f|s)$.*

Proof. Consider any $f \in \overline{\mathbb{V}}_b$ that is \mathcal{F} -measurable and any $s \in \mathcal{X}^*$. Because f is bounded below and \mathcal{F} -measurable, $E_{\text{meas},p}(f|s)$ exists. We can moreover assume that f is non-negative without loss of generality because it is bounded below and both $E_{\text{meas},p}$ and $\overline{E}_{A,p}$ are constant additive with respect to real constants; see M1 and V4. Consider now the non-decreasing sequence $\{f^{\wedge n}\}_{n \in \mathbb{N}_0}$ of upper cuts and note that each $f^{\wedge n}$ is bounded and

\mathcal{F} -measurable; this follows from the fact that \mathcal{F} -measurability is preserved under taking minima. Using M8—which we are allowed to use because $E_{\text{meas},p}(f^{\wedge 0}|s)$ is real (since $f^{\wedge 0}$ is bounded) and because $\{f^{\wedge n}\}_{n \in \mathbb{N}_0}$ is non-decreasing—we have that $E_{\text{meas},p}(f|s) = \lim_{n \rightarrow +\infty} E_{\text{meas},p}(f^{\wedge n}|s)$. As a consequence,

$$E_{\text{meas},p}(f|s) = \lim_{n \rightarrow +\infty} E_{\text{meas},p}(f^{\wedge n}|s) = \lim_{n \rightarrow +\infty} \bar{E}_{\Lambda,p}(f^{\wedge n}|s) = \bar{E}_{\Lambda,p}(f|s),$$

where the second equality follows from Proposition 10 and the third one from Theorem 9(i). \square

Next, we want to drop the condition of \mathcal{F} -measurability and generalise towards the domain $\bar{\mathbb{V}}$ of all global extended real variables. So we are looking for an upper expectation operator that appropriately extends $E_{\text{meas},p}$ to the entire domain $\bar{\mathbb{V}} \times \mathcal{X}^*$. Inspired by Theorem 11, we simply suggest the following upper integral:

$$\bar{E}_{\text{meas},p}^1(f|s) := \inf\{E_{\text{meas},p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \text{ for all } f \in \bar{\mathbb{V}} \text{ and } s \in \mathcal{X}^*.$$

That it is valid to write $E_{\text{meas},p}(g|s)$ for any $g \in \bar{\mathbb{V}}_b$ that is \mathcal{F} -measurable, follows from the fact that $E_{\text{meas},p}(g|s) = \int_{\Omega} g dP|_s$ is guaranteed to exist for such a variable. Let us now show that $\bar{E}_{\text{meas},p}^1$ is indeed an extension of $E_{\text{meas},p}$.

Proposition 12. *For any $f \in \bar{\mathbb{V}}$ and any $s \in \mathcal{X}^*$, we have that $\bar{E}_{\text{meas},p}^1(f|s) = E_{\text{meas},p}(f|s)$ if $E_{\text{meas},p}(f|s)$ exists.*

Proof. Suppose that $E_{\text{meas},p}(f|s)$ exists, which means that the variable f is \mathcal{F} -measurable and that $E_{\text{meas},p}(f^+|s)$ or $E_{\text{meas},p}(f^-|s)$ is real-valued. By the monotonicity [M2] of $E_{\text{meas},p}$, we immediately have that

$$\bar{E}_{\text{meas},p}^1(f|s) = \inf\{E_{\text{meas},p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \geq E_{\text{meas},p}(f|s). \quad (7)$$

If $E_{\text{meas},p}(f^+|s) = +\infty$ and therefore $E_{\text{meas},p}(f^-|s) \in \mathbb{R}$ and $E_{\text{meas},p}(f|s) = +\infty$, the desired equality follows trivially from Equation (7). We proceed to show that this is also the case if $E_{\text{meas},p}(f^+|s) < +\infty$.

Consider the non-increasing sequence $\{f^{\vee-n}\}_{n \in \mathbb{N}_0}$ of lower cuts of f . Then, for any $n \in \mathbb{N}_0$, the global variable $f^{\vee-n}$ is bounded below and \mathcal{F} -measurable. Hence,

$$\begin{aligned} \bar{E}_{\text{meas},p}^1(f|s) &= \inf\{E_{\text{meas},p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \\ &\leq \inf_{n \in \mathbb{N}_0} E_{\text{meas},p}(f^{\vee-n}|s) = \lim_{n \rightarrow +\infty} E_{\text{meas},p}(f^{\vee-n}|s) = E_{\text{meas},p}(f|s), \end{aligned}$$

where the second equality follows from the non-increasing character of the sequence $\{f^{\vee-n}\}_{n \in \mathbb{N}_0}$ and the monotonicity [M2] of $E_{\text{meas},p}$, and the final equality follows from M9, which we can use because $f^{\vee-n} \leq f^{\vee 0} = f^+$ for all $n \in \mathbb{N}_0$ and $E_{\text{meas},p}(f^+|s) < +\infty$ by assumption. Combining this with Equation (7), the desired equality again follows. \square

Our next result shows that $\bar{E}_{\text{meas},p}^1$ is in fact the most conservative extension of $E_{\text{meas},p}$ that satisfies P4.

Proposition 13. *$\bar{E}_{\text{meas},p}^1$ is the most conservative extension of $E_{\text{meas},p}$ that satisfies P4.*

Proof. Fix any $f \in \bar{\mathbb{V}}$, any $s \in \mathcal{X}^*$ and any $\bar{E}'_p : \bar{\mathbb{V}} \times \mathcal{X}^* \rightarrow \bar{\mathbb{R}}$ that satisfies P4 and that coincides with $E_{\text{meas},p}$ on its domain. Then,

$$\begin{aligned} \bar{E}_{\text{meas},p}^1(f|s) &= \inf\{E_{\text{meas},p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \\ &= \inf\{\bar{E}'_p(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \geq \bar{E}'_p(f|s), \end{aligned}$$

where the second equality follows from the fact that $E_{\text{meas},p}(g|s)$ exists for any \mathcal{F} -measurable variable $g \in \bar{\mathbb{V}}_b$, and the inequality follows from the fact that \bar{E}'_p satisfies monotonicity [P4]. Hence, the proposition now follows immediately from the fact that, due to Proposition 12, $\bar{E}_{\text{meas},p}^1$ itself is an extension of $E_{\text{meas},p}$ and the fact that $\bar{E}_{\text{meas},p}^1$ satisfies P4 because $E_{\text{meas},p}$ is monotone [M2]. \square

In Section 4, we argued that P1–P5 are desirable for a global upper expectation and, under those axioms, we put forward the most conservative upper expectation as our model of choice. Now, Proposition 13 guarantees that $\bar{E}_{\text{meas},p}^1$ is the most conservative extension of $E_{\text{meas},p}$ only under P4. So it may very well be that $\bar{E}_{\text{meas},p}^1$ does not necessarily satisfy the axioms P1–P5. Since we consider these axioms to be desirable, that would mean that $\bar{E}_{\text{meas},p}^1$ is a too conservative extension of $E_{\text{meas},p}$, in the sense that it results in an unwarranted loss of information. Our following result assures that this is not the case; the extension $\bar{E}_{\text{meas},p}^1$ coincides with $\bar{E}_{\Lambda,p}$ and hence, $\bar{E}_{\text{meas},p}^1$ is guaranteed to satisfy P1–P5.

Theorem 14. *For any $f \in \bar{\mathbb{V}}$ and any $s \in \mathcal{X}^*$, we have that $\bar{E}_{\text{meas},p}^1(f|s) = \bar{E}_{\Lambda,p}(f|s)$.*

Proof. For any $f \in \bar{\mathbb{V}}$ and any $s \in \mathcal{X}^*$, we immediately find that

$$\begin{aligned} \bar{E}_{\text{meas},p}^1(f|s) &= \inf\{E_{\text{meas},p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \\ &= \inf\{\bar{E}_{\Lambda,p}(g|s) : g \in \bar{\mathbb{V}}_b, g \text{ is } \mathcal{F}\text{-measurable and } g \geq f\} \geq \bar{E}_{\Lambda,p}(f|s), \end{aligned}$$

where the second equality follows from Theorem 11 and the inequality from the monotonicity [P4] of $\bar{E}_{\Lambda,p}$. To show that the converse inequality holds, we will use the fact that $\bar{E}_{\Lambda,p}$ satisfies P7 (this follows from Proposition 5 and Theorem 6).

Consider any $g \in \bar{\mathbb{V}}_{b,\text{lim}}$. Since g is the pointwise limit of a sequence $\{g_n\}_{n \in \mathbb{N}_0}$ of finitary gambles $g_n \in \mathbb{V}_{\text{fin}}$ and since any finitary gamble is clearly \mathcal{F} -measurable, g is the pointwise limit of a sequence of \mathcal{F} -measurable global gambles. Then it follows that g itself is also \mathcal{F} -measurable (\mathcal{F} -measurability is preserved under pointwise convergence). Furthermore, by the definition of $\bar{\mathbb{V}}_{b,\text{lim}}$, g is also bounded below. Hence, $E_{\text{meas},p}(g|s)$ exists and due to Theorem 11, we have that $\bar{E}_{\Lambda,p}(g|s) = E_{\text{meas},p}(g|s)$. Since $\bar{E}_{\text{meas},p}^1$ is an extension of $E_{\text{meas},p}$ by Proposition 12, this implies that $\bar{E}_{\Lambda,p}(g|s) = \bar{E}_{\text{meas},p}^1(g|s)$. We conclude that

$$\bar{E}_{\Lambda,p}(f|s) = \inf\{\bar{E}_{\Lambda,p}(g|s) : g \in \bar{\mathbb{V}}_{b,\text{lim}} \text{ and } g \geq f\} = \inf\{\bar{E}_{\text{meas},p}^1(g|s) : g \in \bar{\mathbb{V}}_{b,\text{lim}} \text{ and } g \geq f\} \geq \bar{E}_{\text{meas},p}^1(f|s),$$

where the first step follows from P7 and the last step follows from the monotonicity [P4] of $\bar{E}_{\text{meas},p}^1$ (as established by Proposition 13). \square

Theorem 14 does not only show us that $\bar{E}_{\text{meas},p}^1$ satisfies P1–P5, but also, and more importantly, that it is the most conservative global upper expectation satisfying P1–P5. Furthermore, if we combine Theorem 14 with Proposition 13, it can also easily be seen that it is the most conservative extension of $E_{\text{meas},p}$ satisfying P1–P5.

Corollary 15. *$\bar{E}_{\text{meas},p}^1$ is the most conservative extension of $E_{\text{meas},p}$ that satisfies P1–P5.*

Proof. We know from Proposition 13 that $\bar{E}_{\text{meas},p}^1$ is the most conservative extension of $E_{\text{meas},p}$ that satisfies P4. This implies in particular that $\bar{E}_{\text{meas},p}^1$ is more (or at least as) conservative than any extension of $E_{\text{meas},p}$ that satisfies P1–P5. Since $\bar{E}_{\text{meas},p}^1$ itself satisfies P1–P5 due to Theorem 14, it follows that $\bar{E}_{\text{meas},p}^1$ is indeed the most conservative extension of $E_{\text{meas},p}$ satisfying P1–P5. \square

Another interesting consequence of Theorem 14 and Proposition 13 is that $\bar{E}_{\Lambda,p}$ turns out to coincide with $E_{\text{meas},p}$ on the entire domain where it is defined, or, in other words, our model $\bar{E}_{\Lambda,p}$ is an extension of the traditional measure-theoretic expectation $E_{\text{meas},p}$.

Corollary 16. *For any $f \in \bar{\mathbb{V}}$ and any $s \in \mathcal{X}^*$, we have that $\bar{E}_{\Lambda,p}(f|s) = E_{\text{meas},p}(f|s)$ if $E_{\text{meas},p}(f|s)$ exists.*

Proof. Immediate consequence of Theorem 14 and Proposition 13. \square

Now, before we continue to consider the more general case of imprecise probability trees, we briefly want to come back to our choice of the extension $\bar{E}_{\text{meas},p}^1$. Recall that we defined it as an upper approximation—called an upper integral—with respect to all \mathcal{F} -measurable variables in $\bar{\mathbb{V}}_b$. However, one could rightfully wonder why we approximate with respect to this particular domain. It turned out to be a suitable choice, as Corollary 15

shows, but there was a priori no concrete reason for doing so. We could just as well have adopted the following alternative upper integral:¹⁰

$$\bar{E}_{\text{meas},p}^2(f|s) := \inf \left\{ E_{\text{meas},p}(g|s) : g \in \bar{\mathcal{V}}, E_{\text{meas},p}(g|s) \text{ exists and } g \geq f \right\} \text{ for all } f \in \bar{\mathcal{V}} \text{ and } s \in \mathcal{X}^*.$$

The upper expectation $\bar{E}_{\text{meas},p}^2$ is an extension of $E_{\text{meas},p}$ because $E_{\text{meas},p}$ is monotone [M2]. Moreover, it is also clear that $\bar{E}_{\text{meas},p}^2(f|s)$ is dominated by $\bar{E}_{\text{meas},p}^1(f|s)$ for all $f \in \bar{\mathcal{V}}$ and all $s \in \mathcal{X}^*$ because the infimum in $\bar{E}_{\text{meas},p}^2$ is taken over an equal or larger set than the set over which the infimum is taken in $\bar{E}_{\text{meas},p}^1$. Hence, one could be tempted to adopt the upper integral $\bar{E}_{\text{meas},p}^2$ instead of $\bar{E}_{\text{meas},p}^1$ with the aim of obtaining a more informative extension—recall that we regard higher upper expectations as less informative. However, it can straightforwardly be shown that this is only vain hope.

Proposition 17. *We have that $\bar{E}_{\text{meas},p}^1(f|s) = \bar{E}_{\text{meas},p}^2(f|s)$ for all $f \in \bar{\mathcal{V}}$ and all $s \in \mathcal{X}^*$.*

Proof. Fix any $f \in \bar{\mathcal{V}}$ and any $s \in \mathcal{X}^*$. That $\bar{E}_{\text{meas},p}^1(f|s) \geq \bar{E}_{\text{meas},p}^2(f|s)$ follows immediately from the fact that $E_{\text{meas},p}(g|s)$ exists for each \mathcal{F} -measurable variable $g \in \bar{\mathcal{V}}_b$ and therefore, that the infimum in $\bar{E}_{\text{meas},p}^2(f|s)$ is taken over a set that is at least as large as the set over which the infimum is taken in $\bar{E}_{\text{meas},p}^1(f|s)$. On the other hand, we have that

$$\begin{aligned} \bar{E}_{\text{meas},p}^2(f|s) &= \inf \left\{ E_{\text{meas},p}(g|s) : g \in \bar{\mathcal{V}}, E_{\text{meas},p}(g|s) \text{ exists and } g \geq f \right\} \\ &= \inf \left\{ \bar{E}_{\text{meas},p}^1(g|s) : g \in \bar{\mathcal{V}}, E_{\text{meas},p}(g|s) \text{ exists and } g \geq f \right\} \geq \bar{E}_{\text{meas},p}^1(f|s), \end{aligned}$$

where the second equality follows from Proposition 12 and the inequality follows from the monotonicity of $\bar{E}_{\text{meas},p}^1$ (see Proposition 13). \square

Hence, both measure-theoretic extensions $\bar{E}_{\text{meas},p}^1$ and $\bar{E}_{\text{meas},p}^2$ are equal and therefore, by Theorem 14, both of them coincide with our model $\bar{E}_{\Lambda,p}$ which, on its turn, coincides with the game-theoretic upper expectation $\bar{E}_{\mathcal{V}}$ according to Theorem 6. In summary, we conclude that, if the local models are precise, all of the three possible approaches—the game-theoretic, the measure-theoretic and ours—are equivalent.

Imprecise Probability Trees

Let us now turn to the general case where we consider an imprecise probability tree $\bar{\mathcal{Q}}$ for which the local models $\bar{\mathcal{Q}}_s$ are (boundedly) coherent upper expectations that are not necessarily self-conjugate. As we have explained earlier, these local models can equivalently be represented by sets \mathbb{P}_s of probability mass functions on \mathcal{X} , where $\mathbb{P}_s := \mathbb{P}_{\bar{\mathcal{Q}}_s}$ is related to $\bar{\mathcal{Q}}_s$ according to Equations (1) and (2) for all $s \in \mathcal{X}^*$. We call a precise probability tree $p : s \in \mathcal{X}^* \mapsto p(\cdot|s)$ compatible with the imprecise probability tree $\bar{\mathcal{Q}}$ and write that $p \sim \bar{\mathcal{Q}}$ if $p(\cdot|s) \in \mathbb{P}_s$ for all $s \in \mathcal{X}^*$. In other words, a compatible precise probability tree $p \sim \bar{\mathcal{Q}}$ corresponds to a possible selection, where a local probability mass function $p(\cdot|s)$ is chosen from the set \mathbb{P}_s for each situation $s \in \mathcal{X}^*$. Now, for any such compatible tree $p \sim \bar{\mathcal{Q}}$, we can consider the global upper expectation $\bar{E}_{\text{meas},p} := \bar{E}_{\text{meas},p}^1 = \bar{E}_{\text{meas},p}^2$ as defined in the previous section. Hence, instead of a single upper expectation, we now obtain a set of compatible global upper expectations. We define the *global measure-theoretic upper expectation* \bar{E}_{meas} as the upper envelope over all such compatible upper expectations:

$$\bar{E}_{\text{meas}}(f|s) := \sup \left\{ \bar{E}_{\text{meas},p}(f|s) : p \sim \bar{\mathcal{Q}} \right\} \text{ for all } f \in \bar{\mathcal{V}} \text{ and } s \in \mathcal{X}^*.$$

This definition is in correspondence with the usual measure-theoretic view towards imprecision where it is regarded as partial information about a single, unknown probability measure. Upper (and lower) expectations are

¹⁰This upper integral seems to be the one used by Shafer and Vovk in [5, Chapter 9].

then considered as upper (and lower) bounds on the set of all ‘possible’ linear expectations. A similar, yet unconditional version of this measure-theoretic upper expectation is also adopted by Shafer and Vovk in [5, Chapter 9].

As we will now show, the measure-theoretic upper expectation \bar{E}_{meas} satisfies our axioms P1–P5. This implies, by Theorem 6, that \bar{E}_{meas} is always dominated by the global upper expectation \bar{E}_A , and therefore also by \bar{E}_V . Moreover, that the operator \bar{E}_{meas} is entirely based on a measure-theoretic view towards uncertainty, yet turns out to satisfy our axioms P1–P5, can only be seen as an additional motivation for the use of these axioms.

Theorem 18. *The global measure-theoretic upper expectation \bar{E}_{meas} is an element of $\bar{E}_{1-5}(\bar{Q})$.*

Proof. To show that P1 holds, consider any $f \in \mathcal{L}(\mathcal{X})$ and any $x_{1:n} \in \mathcal{X}^*$. Since $f(X_{n+1})$ is bounded and clearly \mathcal{F} -measurable, we have, for any $p \sim \bar{Q}$, that $E_{\text{meas},p}(f(X_{n+1})|x_{1:n})$ exists and therefore that $\bar{E}_{\text{meas},p}(f(X_{n+1})|x_{1:n}) = E_{\text{meas},p}(f(X_{n+1})|x_{1:n})$, due to Proposition 12. Moreover, using $P_{|x_{1:n}}(x_{n+1})$ to denote $P_{|x_{1:n}}(\cup_{z_{1:n} \in \mathcal{X}^n} \Gamma(z_{1:n}, x_{n+1}))$ for any $x_{n+1} \in \mathcal{X}$, it is clear that $E_{\text{meas},p}(f(X_{n+1})|x_{1:n}) = \sum_{x_{n+1} \in \mathcal{X}} f(x_{n+1})P_{|x_{1:n}}(x_{n+1})$. It also follows immediately from the construction (6) of $P_{|x_{1:n}} = P(\cdot|x_{1:n})$ that $P_{|x_{1:n}}(x_{n+1}) = P_{|x_{1:n}}(x_{1:n+1}) = p(x_{n+1}|x_{1:n})$ for all $x_{n+1} \in \mathcal{X}$ and $p \sim \bar{Q}$. Hence,

$$\begin{aligned} \bar{E}_{\text{meas}}(f(X_{n+1})|x_{1:n}) &= \sup\{\bar{E}_{\text{meas},p}(f(X_{n+1})|x_{1:n}): p \sim \bar{Q}\} \\ &= \sup\{E_{\text{meas},p}(f(X_{n+1})|x_{1:n}): p \sim \bar{Q}\} \\ &= \sup\left\{\sum_{x_{n+1} \in \mathcal{X}} f(x_{n+1})p(x_{n+1}|x_{1:n}): p \sim \bar{Q}\right\} \\ &= \sup\left\{\sum_{x_{n+1} \in \mathcal{X}} f(x_{n+1})p(x_{n+1}|x_{1:n}): p(\cdot|x_{1:n}) \in \mathbb{P}_{x_{1:n}}\right\} = \bar{Q}_{x_{1:n}}(f), \end{aligned}$$

where the last step follows from the relation between $\mathbb{P}_{x_{1:n}}$ and $\bar{Q}_{x_{1:n}}$ which is given by Equation (2).

The remaining properties P2–P5 can be easily derived from the definition of \bar{E}_{meas} and the fact that, due to Corollary 15, these properties are satisfied by each compatible $\bar{E}_{\text{meas},p}$. For Properties P2 and P4 this is immediate. To see that P3 holds, consider any $f \in \mathbb{V}_{\text{fin}}$ and any $n \in \mathbb{N}_0$. Note that, for any compatible tree $p \sim \bar{Q}$ and any $x_{1:n} \in \mathcal{X}^n$, we have that $\bar{E}_{\text{meas},p}(f|x_{1:n}) \leq \bar{E}_{\text{meas},p}(\bar{E}_{\text{meas},p}(f|X_{1:n+1})|x_{1:n})$ because $\bar{E}_{\text{meas},p}$ satisfies P3. Hence,

$$\begin{aligned} \bar{E}_{\text{meas}}(f|x_{1:n}) &= \sup\{\bar{E}_{\text{meas},p}(f|x_{1:n}): p \sim \bar{Q}\} \leq \sup\{\bar{E}_{\text{meas},p}(\bar{E}_{\text{meas},p}(f|X_{1:n+1})|x_{1:n}): p \sim \bar{Q}\} \\ &\leq \sup\{\bar{E}_{\text{meas},p}(\bar{E}_{\text{meas}}(f|X_{1:n+1})|x_{1:n}): p \sim \bar{Q}\} \\ &= \bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:n+1})|x_{1:n}), \end{aligned}$$

where the second inequality follows from the monotonicity [P4] of $\bar{E}_{\text{meas},p}$ and the fact that $\bar{E}_{\text{meas},p}(f|z_{1:n+1}) \leq \bar{E}_{\text{meas}}(f|z_{1:n+1})$ for all $p \sim \bar{Q}$ and all $z_{1:n+1} \in \mathcal{X}^{n+1}$. The inequality above holds for any $x_{1:n} \in \mathcal{X}^n$, so we indeed find that $\bar{E}_{\text{meas}}(f|X_{1:n}) \leq \bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:n+1})|X_{1:n})$. To prove P5, we consider any $s \in \mathcal{X}^*$ and any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of finitary gambles that is uniformly bounded below and that converges pointwise to some variable $f \in \bar{V}$. For any compatible tree $p \sim \bar{Q}$, $\bar{E}_{\text{meas},p}$ satisfies P5 and therefore, $\bar{E}_{\text{meas},p}(f|s) \leq \limsup_{n \rightarrow +\infty} \bar{E}_{\text{meas},p}(f_n|s)$. Then we also have that

$$\begin{aligned} \bar{E}_{\text{meas}}(f|s) &= \sup\{\bar{E}_{\text{meas},p}(f|s): p \sim \bar{Q}\} \leq \sup\left\{\limsup_{n \rightarrow +\infty} \bar{E}_{\text{meas},p}(f_n|s): p \sim \bar{Q}\right\} \\ &\leq \sup\left\{\limsup_{n \rightarrow +\infty} \left[\sup\{\bar{E}_{\text{meas},p'}(f_n|s): p' \sim \bar{Q}\}\right]: p \sim \bar{Q}\right\} \\ &= \limsup_{n \rightarrow +\infty} \left[\sup\{\bar{E}_{\text{meas},p'}(f_n|s): p' \sim \bar{Q}\}\right] = \limsup_{n \rightarrow +\infty} \bar{E}_{\text{meas}}(f_n|s). \quad \square \end{aligned}$$

As mentioned before, the result above implies, in combination with Theorem 6, that \bar{E}_{meas} is dominated by our model \bar{E}_A .

Corollary 19. *We have that $\bar{E}_{\text{meas}}(f|s) \leq \bar{E}_A(f|s) = \bar{E}_V(f|s)$ for all $f \in \bar{V}$ and all $s \in \mathcal{X}^*$.*

Proof. Immediate consequence of Theorem 6 and Theorem 18. \square

Corollary 19 extends upon an earlier result by Lopatatzidis et al. [12, Theorem 33], which says that $\bar{E}_{\text{meas}}(f|s) \leq \bar{E}_V(f|s)$ for all $s \in \mathcal{X}^*$ and some specific variables f in $\bar{V}_{\text{b,lim}}$. The definition of \bar{E}_V in [12] slightly differs from ours though. The version in [12] is always larger or equal than—so at least as conservative as—ours because they work with supermartingales that can only take real values; see [19, section 8] for more details. Hence, Corollary 19 certainly implies [12, Theorem 33]. Another existing result that is strongly related to Corollary 19 is due to Shafer and Vovk [5, Proposition 9.8], which seems to say that $\bar{E}_{\text{meas}}(f|\square) \leq \bar{E}_V(f|\square)$ for all bounded below variables $f \in \bar{V}_{\text{b}}$. Since our result applies to all $f \in \bar{V}$ and all $s \in \mathcal{X}^*$, our result would in that sense be stronger than theirs. However, it is not immediately clear whether we can actually draw such conclusions. Unlike in our case, Shafer and Vovk [5, Section 9.2] do not require that, for each possible situation, a local model should be specified beforehand. To put it in their terms: Forecaster need not declare to Skeptic beforehand—that is, before the game starts—which local upper expectations he is going to choose. Instead, Shafer and Vovk only start off with a set \mathcal{Q} of local upper expectations from which, at each situation, Forecaster can choose. Since this set \mathcal{Q} is moreover required to be constant throughout the whole game, this crucially implies that we, as subjects outside the game, are unable to specify Forecaster’s moves—and hence, the choice of the local upper expectation at each situation. The global game-theoretic upper expectation will therefore not depend on Forecaster’s moves, but rather only on the set \mathcal{Q} . Moreover, both of their global models—their global game-theoretic upper expectation and their global measure-theoretic upper expectation—are defined on variables with a domain that differs from Ω . Besides the fact that we find it far from trivial how this whole framework relates to ours here, they also impose some relatively strong topological conditions on the set \mathcal{Q} . Hence, all together, in order to draw a fair comparison, we feel like a closer look is required. We hope to clarify this in our future work.

Moving on, the inevitable question that comes to mind when looking at Corollary 19 is whether the inequality can be turned into an equality. We are aware of two partial answers. On the one hand, Shafer and Vovk seem to establish in [5, Theorem 9.7] that $\bar{E}_{\text{meas}}(f|\square) = \bar{E}_V(f|\square)$ for all Suslin gambles $f \in \mathbb{V}$; this includes all \mathcal{F} -measurable gambles and is therefore a very strong result. However, here too, as before, their framework differs substantially from ours and the result requires additional topological assumptions. It is therefore again not immediately clear if, and how, it translates to our context; we leave this for future work. On the other hand, we also have a result by Lopatatzidis et al. [12, Theorem 29] that shows that $\bar{E}_{\text{meas}}(f|s) = \bar{E}_V(f|s)$ for all $s \in \mathcal{X}^*$ and all n -measurable gambles f —and therefore all finitary gambles f in \mathbb{V}_{fin} . Here too, however, the result is not immediately applicable because—as explained directly after Corollary 19—the version of \bar{E}_V that is used in [12] differs from ours. We therefore give a self-contained proof instead. Nevertheless, our ideas and techniques are essentially the same as the ones that were used to obtain [12, Theorem 29].

We start by showing that \bar{E}_{meas} satisfies $\text{P3}^=$; the desired equivalence result will then follow as an immediate consequence.

Lemma 20. *The upper expectation \bar{E}_{meas} satisfies $\text{P3}^=$.*

Proof. Consider any $f \in \mathbb{V}_{\text{fin}}$. Let $n \in \mathbb{N}_0$ be such that f is n -measurable—there is at least one such an n because f is finitary. Then, to establish the fact that \bar{E}_{meas} satisfies $\text{P3}^=$, it suffices to prove that, for all $m \in \mathbb{N}_0$,

$$\text{there is some } p \sim \bar{Q} \text{ such that } \bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:m+1})|X_{1:m}) = \bar{E}_{\text{meas},p}(f|X_{1:m}). \quad (8)$$

Indeed, by definition of \bar{E}_{meas} , $\bar{E}_{\text{meas},p}(f|X_{1:m}) \leq \bar{E}_{\text{meas}}(f|X_{1:m})$ for all $p \sim \bar{Q}$ and all $m \in \mathbb{N}_0$, and therefore, expression (8) implies that $\bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:m+1})|X_{1:m}) \leq \bar{E}_{\text{meas}}(f|X_{1:m})$ for all $m \in \mathbb{N}_0$. The converse inequality then follows from Theorem 18, which says that \bar{E}_{meas} satisfies P3 .

Let us first show that expression (8) holds for any $m \geq n$. Consider any $\ell \geq n$ and any $x_{1:\ell} \in \mathcal{X}^\ell$. Since f is n -measurable and $\ell \geq n$, f takes a constant value $f(x_{1:n})$ on $\Gamma(x_{1:\ell})$, and therefore, we can write that $f\mathbb{I}_{x_{1:\ell}} = f(x_{1:n})\mathbb{I}_{x_{1:\ell}}$. Hence, for any $p \sim \bar{Q}$, since $\bar{E}_{\text{meas},p}$ satisfies P2 due to Theorem 14, we have that

$$\bar{E}_{\text{meas},p}(f|x_{1:\ell}) \stackrel{\text{P2}}{=} \bar{E}_{\text{meas},p}(f\mathbb{I}_{x_{1:\ell}}|x_{1:\ell}) = \bar{E}_{\text{meas},p}(f(x_{1:n})\mathbb{I}_{x_{1:\ell}}|x_{1:\ell}) \stackrel{\text{P2}}{=} \bar{E}_{\text{meas},p}(f(x_{1:n})|x_{1:\ell}) = E_{\text{meas},p}(f(x_{1:n})|x_{1:\ell}).$$

where the last step follows from Proposition 12 and the fact that $E_{\text{meas},p}(f(x_{1:n})|x_{1:\ell})$ exists because $f(x_{1:n})$ is a constant. Furthermore, since $f(x_{1:n})$ is real-valued [because f is a gamble], M1 implies that $E_{\text{meas},p}(f(x_{1:n})|x_{1:\ell}) =$

$f(x_{1:n})$. Hence, we have that

$$\bar{E}_{\text{meas},p}(f|x_{1:\ell}) = f(x_{1:n}) \text{ for all } p \sim \bar{Q}, \text{ all } \ell \geq n \text{ and all } x_{1:\ell} \in \mathcal{X}^\ell. \quad (9)$$

So, since $m \geq n$, the right hand side of Equation (8) is equal to $f(X_{1:n}) = f$, whatever the tree $p \sim \bar{Q}$. Moreover, note that Equation (9) also implies that $\bar{E}_{\text{meas}}(f|x_{1:\ell}) = f(x_{1:n})$ for all $\ell \geq n$ and all $x_{1:\ell} \in \mathcal{X}^\ell$. Then, since $m \geq n$, and therefore certainly $m+1 \geq n$, the left hand side of Equation (8) is equal to

$$\bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:m+1})|X_{1:m}) = \bar{E}_{\text{meas}}(f(X_{1:n})|X_{1:m}) = \bar{E}_{\text{meas}}(f|X_{1:m}) = f,$$

where the last step follows once more from the fact that $m \geq n$, and that $\bar{E}_{\text{meas}}(f|x_{1:\ell}) = f(x_{1:n})$ for all $\ell \geq n$ and all $x_{1:\ell} \in \mathcal{X}^\ell$. Hence, irrespectively of the tree $p \sim \bar{Q}$, both sides of Equation (8) are equal to f —and therefore equal to each other—for the case that $m \geq n$.

For the case that $m < n$, we will establish expression (8) using a decreasing induction argument in m . The case where $m = n$ will serve as an induction base; that expression (8) holds for this case was just proved above. Now suppose that expression (8) is satisfied with $m = \ell$ where $1 \leq \ell \leq n$, meaning that there is some $p \sim \bar{Q}$ such that $\bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:\ell+1})|X_{1:\ell}) = \bar{E}_{\text{meas},p}(f|X_{1:\ell})$. First note that this implies that $\bar{E}_{\text{meas}}(f|X_{1:\ell}) = \bar{E}_{\text{meas},p}(f|X_{1:\ell})$ because, on the one hand, $\bar{E}_{\text{meas}}(f|X_{1:\ell}) \leq \bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:\ell+1})|X_{1:\ell})$ due to Theorem 18 which says that \bar{E}_{meas} satisfies P3, and on the other hand, $\bar{E}_{\text{meas},p}(f|X_{1:\ell}) \leq \bar{E}_{\text{meas}}(f|X_{1:\ell})$ due to the definition of \bar{E}_{meas} . Furthermore, because f is a finitary gamble, and therefore an \mathcal{F} -measurable gamble, $E_{\text{meas},p}(f|X_{1:\ell})$ exists and therefore, by Proposition 12,

$$\bar{E}_{\text{meas}}(f|X_{1:\ell}) = \bar{E}_{\text{meas},p}(f|X_{1:\ell}) = E_{\text{meas},p}(f|X_{1:\ell}).$$

Moreover, using M1 and M2, and the fact that $\inf f$ and $\sup f$ are both real (because f is a gamble), it can easily be inferred that $\inf f \leq E_{\text{meas},p}(f|X_{1:\ell}) \leq \sup f$. So $\bar{E}_{\text{meas}}(f|X_{1:\ell}) = E_{\text{meas},p}(f|X_{1:\ell})$ is a gamble, and more specifically an ℓ -measurable gamble. Then, fix any $x_{1:\ell-1} \in \mathcal{X}^{\ell-1}$ and note that, since \bar{E}_{meas} satisfies P1 and P2 by Theorem 18, we can apply Lemma 1 to obtain that

$$\bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:\ell})|x_{1:\ell-1}) = \bar{E}_{\text{meas}}(E_{\text{meas},p}(f|X_{1:\ell})|x_{1:\ell-1}) = \bar{Q}_{x_{1:\ell-1}}(E_{\text{meas},p}(f|x_{1:\ell-1} \cdot)).$$

Then, due to the fact that $\mathbb{P}_{x_{1:\ell-1}}$ is related with $\bar{Q}_{x_{1:\ell-1}}$ through Equation (2), there is some $p'(\cdot|x_{1:\ell-1}) \in \mathbb{P}_{x_{1:\ell-1}}$ such that

$$\bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:\ell})|x_{1:\ell-1}) = \bar{Q}_{x_{1:\ell-1}}(E_{\text{meas},p}(f|x_{1:\ell-1} \cdot)) = \sum_{x_\ell \in \mathcal{X}} p'(x_\ell|x_{1:\ell-1}) E_{\text{meas},p}(f|x_{1:\ell}). \quad (10)$$

Moreover, since f is n -measurable, we can write that $f = \sum_{z_{1:n} \in \mathcal{X}^n} f(z_{1:n}) \mathbb{I}_{z_{1:n}}$, and therefore, the Lebesgue integral $E_{\text{meas},p}(f|x_{1:\ell})$ for any $x_\ell \in \mathcal{X}$, reduces—by definition; see [1, Section 2.6.1]—to the finite sum

$$E_{\text{meas},p}(f|x_{1:\ell}) = \sum_{z_{1:n} \in \mathcal{X}^n} f(z_{1:n}) P_{|x_{1:\ell}}(z_{1:n}) = \sum_{z_{1:\ell} \in \mathcal{X}^\ell} \sum_{x_{\ell+1:n} \in \mathcal{X}^{n-\ell}} f(z_{1:\ell} x_{\ell+1:n}) P_{|x_{1:\ell}}(z_{1:\ell} x_{\ell+1:n}),$$

where the last step simply is a convenient change of notation. Due to Equation (6), we have, for all $z_{1:\ell} \in \mathcal{X}^\ell$ and all $x_{\ell+1:n} \in \mathcal{X}^{n-\ell}$, that $P_{|x_{1:\ell}}(z_{1:\ell} x_{\ell+1:n}) = 0$ if $z_{1:\ell} \neq x_{1:\ell}$ and $P_{|x_{1:\ell}}(z_{1:\ell} x_{\ell+1:n}) = \prod_{i=\ell}^{n-1} p(x_{i+1}|x_{1:i})$ if $z_{1:\ell} = x_{1:\ell}$. Substituting these expressions in the equality above, gives us

$$E_{\text{meas},p}(f|x_{1:\ell}) = \sum_{x_{\ell+1:n} \in \mathcal{X}^{n-\ell}} f(x_{1:n}) P_{|x_{1:\ell}}(x_{1:n}) = \sum_{x_{\ell+1:n} \in \mathcal{X}^{n-\ell}} f(x_{1:n}) \prod_{i=\ell}^{n-1} p(x_{i+1}|x_{1:i}).$$

Then, plugging this back into Equation (10), we obtain that

$$\begin{aligned} \bar{E}_{\text{meas}}(\bar{E}_{\text{meas}}(f|X_{1:\ell})|x_{1:\ell-1}) &= \sum_{x_\ell \in \mathcal{X}} p'(x_\ell|x_{1:\ell-1}) \sum_{x_{\ell+1:n} \in \mathcal{X}^{n-\ell}} f(x_{1:n}) \prod_{i=\ell}^{n-1} p(x_{i+1}|x_{1:i}) \\ &= \sum_{x_{\ell:n} \in \mathcal{X}^{n-\ell+1}} f(x_{1:n}) p'(x_\ell|x_{1:\ell-1}) \prod_{i=\ell}^{n-1} p(x_{i+1}|x_{1:i}). \end{aligned} \quad (11)$$

Now let $p^* \sim \bar{Q}$ be any compatible precise probability tree such that $p^*(\cdot|x_{1:\ell-1}) = p'(\cdot|x_{1:\ell-1})$ and $p^*(\cdot|x_{1:i}) = p(\cdot|x_{1:i})$ for all $\ell \leq i \leq n-1$ and all $x_{\ell:i} \in \mathcal{X}^{i-\ell+1}$; it is always possible to choose such a compatible tree p^* because $p \sim \bar{Q}$ and $p'(\cdot|x_{1:\ell-1}) \in \mathbb{P}_{x_{1:\ell-1}}$. Then note that, for any $x_{\ell:n} \in \mathcal{X}^{n-\ell+1}$, the corresponding probability $\mathbb{P}_{x_{1:\ell-1}}^*(x_{1:n})$, related with p^* through Equation (6), is equal to $p'(x_\ell|x_{1:\ell-1}) \prod_{i=\ell}^{n-1} p(x_{i+1}|x_{1:i})$. Hence, recalling Equation (11), we infer that

$$\begin{aligned} \bar{\mathbb{E}}_{\text{meas}}(\bar{\mathbb{E}}_{\text{meas}}(f|X_{1:\ell})|x_{1:\ell-1}) &= \sum_{x_{\ell:n} \in \mathcal{X}^{n-\ell+1}} f(x_{1:n}) \mathbb{P}_{x_{1:\ell-1}}^*(x_{1:n}) = \sum_{z_{1:\ell-1} \in \mathcal{X}^{\ell-1}} \sum_{x_{\ell:n} \in \mathcal{X}^{n-\ell+1}} f(z_{1:\ell-1} x_{\ell:n}) \mathbb{P}_{x_{1:\ell-1}}^*(z_{1:\ell-1} x_{\ell:n}) \\ &= \mathbb{E}_{\text{meas}, p^*}(f|x_{1:\ell-1}) = \bar{\mathbb{E}}_{\text{meas}, p^*}(f|x_{1:\ell-1}), \end{aligned}$$

where the second step follows from the fact that $\mathbb{P}_{x_{1:\ell-1}}^*(z_{1:\ell-1} x_{\ell:n}) = 0$ if $z_{1:\ell-1} \neq x_{1:\ell-1}$ due to Equation (6), where the penultimate step follows from $f = \sum_{z_{1:n} \in \mathcal{X}^n} f(z_{1:n}) \mathbb{1}_{z_{1:n}}$ and the definition of the Lebesgue integral [1, Section 2.6.1], and where the last step follows from Proposition 12. Hence, since the equation above holds for any $x_{1:\ell-1} \in \mathcal{X}^{\ell-1}$, we have established expression (8) for the case where $m = \ell - 1$, concluding our proof of the induction step. \square

That $\bar{\mathbb{E}}_{\text{meas}}$, $\bar{\mathbb{E}}_A$ and $\bar{\mathbb{E}}_V$ are all equal on the domain \mathbb{V}_{fin} now follows straightforwardly by combining Lemma 20 with Corollary 19, Theorem 18 and Proposition 2.

Proposition 21. *We have that $\bar{\mathbb{E}}_{\text{meas}}(f|s) = \bar{\mathbb{E}}_A(f|s) = \bar{\mathbb{E}}_V(f|s)$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $s \in \mathcal{X}^{**}$.*

Proof. Due to Corollary 19, it suffices to prove that $\bar{\mathbb{E}}_{\text{meas}}(f|s) \geq \bar{\mathbb{E}}_A(f|s)$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $s \in \mathcal{X}^{**}$. This can easily be done by combining Theorem 18, Lemma 20 and Proposition 2. Indeed, Theorem 18 says that $\bar{\mathbb{E}}_{\text{meas}}$ is an element of $\bar{\mathbb{E}}_{1-5}(\bar{Q}) \subseteq \bar{\mathbb{E}}_{1-4}(\bar{Q})$, and Lemma 20 says that $\bar{\mathbb{E}}_{\text{meas}}$ satisfies P3[−], so Proposition 2 implies that $\bar{\mathbb{E}}_{\text{meas}}(f|s) \geq \bar{\mathbb{E}}'(f|s)$ for all $f \in \mathbb{V}_{\text{fin}}$, all $s \in \mathcal{X}^{**}$ and all $\bar{\mathbb{E}}' \in \bar{\mathbb{E}}_{1-4}(\bar{Q})$. In particular, this implies that $\bar{\mathbb{E}}_{\text{meas}}(f|s) \geq \bar{\mathbb{E}}_A(f|s)$ for all $f \in \mathbb{V}_{\text{fin}}$ and all $s \in \mathcal{X}^{**}$ because, by definition, $\bar{\mathbb{E}}_A \in \bar{\mathbb{E}}_{1-5}(\bar{Q}) \subseteq \bar{\mathbb{E}}_{1-4}(\bar{Q})$. \square

Whether the equality with $\bar{\mathbb{E}}_{\text{meas}}$ can be extended to a more general domain, remains an open question that we would like to pursue in our future research. We expect that there is much to learn from the results in [5, Section 9.2], provided that they can be translated to our setting here.

10. Conclusion

We have put forward a small set of simple axioms P1–P5 and argued that they are desirable for a global upper expectation in the context of discrete-time finite-state uncertain processes. We have established the existence of a unique most conservative model under these axioms, and have in addition given sufficient conditions to uniquely characterise this most conservative model. Using this characterisation, we have shown that our most conservative upper expectation coincides with a version of the game-theoretic upper expectation used by Shafer and Vovk, and therefore has particularly powerful mathematical properties, despite the simplicity of our defining axioms.

We also considered an alternative, more traditional measure-theoretic approach, where we defined a global belief model using an upper envelope of upper integrals corresponding to compatible precise probability trees. We have shown that this model satisfies axioms P1–P5 and therefore, that our most conservative model gives guaranteed upper bounds on the value of this measure-theoretic upper expectation. For precise probability trees, both models, and consequently also the game-theoretic model, coincide. In the imprecise case, equality was only established for the domain of finitary gambles, and it remains to be seen whether this can be extended to a more general domain.

Given our current findings, we believe there are a number of reasons why our global model $\bar{\mathbb{E}}_A$ candidates as an excellent choice when it comes to modelling discrete-time finite-state uncertain processes. First and foremost, its definition is based on the axioms P1–P5; a set of properties which, in our opinion, are essential for a global belief model to have. Our model $\bar{\mathbb{E}}_A$ is taken to be the most conservative or, equivalently, the least informative model under these axioms because we do not want to include any further information. In this way, we obtain a definition that is both convincing and interpretationally clear. Moreover, if one desires her global belief model

to also have other properties, additional to P1–P5, then our model \bar{E}_A would still serve as a conservative upper bound for her model. As a second argument, we recall from Section 8 that our model \bar{E}_A possesses rather strong continuity properties, as well as several fundamental probabilistic laws. Thirdly and lastly, our model \bar{E}_A coincides with the game-theoretic upper expectation \bar{E}_V and dominates the measure-theoretic upper expectation \bar{E}_{meas} . Hence, our model \bar{E}_A plays an important role to anyone: for a person with a game-theoretic background, the equivalent models \bar{E}_A and \bar{E}_V can be used interchangeably and can therefore benefit from each others properties; for a person with a measure-theoretic background, our model \bar{E}_A always gives a conservative upper bound for the model \bar{E}_{meas} .

One counterargument that could be given to the plea above, is whether axiom P5 is really desirable or even justified for a global belief model. Our justification for it is based on the fact that we consider upper expectations of non-finitary variables to be abstract idealisations of upper expectations of finitary variables. Still, this could be seen as a rather artificial and unnecessary assumption, similar to the countable additivity property in measure theory. However, compared to other common continuity properties, e.g. monotone convergence, property P5 is fairly weak since it only applies to sequences of finitary gambles and only imposes an inequality on the value of the global upper expectations. Moreover, that both the game-theoretic and the measure-theoretic model satisfy property P5, can only be seen as additional motivation to adopt it.

Acknowledgements

The research of Jasper De Bock and Gert de Cooman was partially funded by the European Commission's H2020 programme, through the UTOPIAE Marie Curie Innovative Training Network, H2020-MSCA-ITN-2016, Grant Agreement number 722734, as well as by FWO (Research Foundation - Flanders), through project number 3GO28919, entitled "Efficient inference in large-scale queueing models using imprecise continuous-time Markov chains". We are also indebted two anonymous reviewers that provided helpful feedback and pointed us to some relevant literature.

References

- [1] A. N. Shiryaev, *Probability-1*, Graduate Texts in Mathematics, third ed., Springer, New York, 2016.
- [2] A. N. Shiryaev, *Probability-2*, Graduate Texts in Mathematics, third ed., Springer, New York, 2019.
- [3] P. Billingsley, *Probability and Measure*, Wiley Series in Probability and Statistics, Wiley, 1995.
- [4] G. Shafer, V. Vovk, *Probability and Finance: It's Only a Game!*, Wiley, New York, 2001.
- [5] G. Shafer, V. Vovk, *Game-Theoretic Foundations for Probability and Finance*, Wiley, Hoboken, 2019.
- [6] D. Williams, *Probability with Martingales*, Cambridge mathematical textbooks, Cambridge University Press, 1991.
- [7] P. Walley, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, London, 1991.
- [8] T. Augustin, F. P. Coolen, G. de Cooman, M. C. Troffaes, *Introduction to Imprecise Probabilities*, Wiley, Chichester, 2014.
- [9] M. C. Troffaes, G. de Cooman, *Lower Previsions*, Wiley, Chichester, 2014.
- [10] P. M. Williams, Notes on conditional previsions, Technical Report, School of Mathematical and Physical Science, University of Sussex, 1975. For a revised journal version, see [14].
- [11] G. de Cooman, F. Hermans, Imprecise probability trees: Bridging two theories of imprecise probability, *Artificial Intelligence* 172 (2008) 1400–1427.
- [12] S. Lopatatzidis, *Robust Modelling and Optimisation in Stochastic Processes using Imprecise Probabilities*, with an Application to Queueing Theory, Ph.D. thesis, Ghent University, 2017.
- [13] N. T'Joens, J. De Bock, G. de Cooman, In search of a global belief model for discrete-time uncertain processes, in: *Proceedings of the 11th International Symposium on Imprecise Probabilities: Theories and Applications*, volume 103, 2019, pp. 377–385.
- [14] P. M. Williams, Notes on conditional previsions, *International Journal of Approximate Reasoning* 44 (2007) 366–383.
- [15] H. L. Royden, P. M. Fitzpatrick, *Real Analysis*, Prentice Hall, New Jersey, 2010.
- [16] G. Shafer, V. Vovk, A. Takemura, Lévy's zero–one law in game-theoretic probability, *Journal of Theoretical Probability* 25 (2012) 1–24.
- [17] G. de Cooman, J. De Bock, S. Lopatatzidis, Imprecise stochastic processes in discrete time: global models, imprecise Markov chains, and ergodic theorems, *International Journal of Approximate Reasoning* 76 (2016) 18–46.
- [18] J. Ville, *Étude critique de la notion de collectif*, Monographies des probabilités, Gauthier-Villars, Paris, 1939.
- [19] N. T'Joens, J. De Bock, G. de Cooman, Game-theoretic upper expectations for discrete-time finite-state uncertain processes, 2020. Submitted for publication; see arXiv:2008.03133 for a preliminary online version.
- [20] N. T'Joens, G. de Cooman, J. De Bock, Continuity of the Shafer-Vovk-Ville operator, in: *Proceedings of the 9th International Conference on Soft Methods in Probability and Statistics*, volume 832, 2018, pp. 200–207.
- [21] A. N. Shiryaev, *Probability*, Graduate Texts in Mathematics, second ed., Springer, New York, 1996.

Appendix A. Basic Measure-Theoretic Concepts

A *measurable space* $(\mathcal{Y}, \mathcal{F})$ is a couple, where \mathcal{Y} is a non-empty set and \mathcal{F} is a σ -algebra on \mathcal{Y} . We say that $A \subseteq \mathcal{Y}$ is \mathcal{F} -*measurable* if $A \in \mathcal{F}$. Such a set $A \in \mathcal{F}$ is also called an (\mathcal{F} -measurable) *event*. We say that an extended real-valued function $f: \mathcal{Y} \rightarrow \overline{\mathbb{R}}$ is \mathcal{F} -*measurable* if the set $f^{-1}(A) := \{y \in \mathcal{Y} : f(y) \in A\}$ is \mathcal{F} -measurable for every $A \in \mathcal{B}(\overline{\mathbb{R}})$. Here, $\mathcal{B}(\overline{\mathbb{R}})$ denotes the Borel algebra on $\overline{\mathbb{R}}$, being the σ -algebra generated by all open—or equivalently, closed—sets in $\overline{\mathbb{R}}$. Recall that we consider the usual order topology on $\overline{\mathbb{R}}$ and therefore, the Borel algebra $\mathcal{B}(\overline{\mathbb{R}})$ can alternatively be generated from the sets $\{x \in \overline{\mathbb{R}} : x \leq c\}$ where $c \in \mathbb{R}$; see for instance [1, Section 2.2.2]. Then, as we have done in the main text, we can alternatively characterise the \mathcal{F} -measurable functions as those functions $f: \mathcal{Y} \rightarrow \overline{\mathbb{R}}$ such that $\{y \in \mathcal{Y} : f(y) \leq c\}$ is \mathcal{F} -measurable for all $c \in \mathbb{R}$ [3, Theorem 13.1. (i)]. Typically, an \mathcal{F} -measurable (extended) real-valued function f is also called a *random variable*.

A *probability space* $(\mathcal{Y}, \mathcal{F}, P)$ is a measurable space $(\mathcal{Y}, \mathcal{F})$ equipped with a σ -additive probability measure P on \mathcal{F} . We say that an event $A \in \mathcal{F}$ is P -*null* if $P(A) = 0$. We will also say that a property about the elements in \mathcal{Y} holds P -*almost surely* (P -a.s.) if there is a P -null event $A \in \mathcal{F}$ such that the property holds for all $y \in \mathcal{Y} \setminus A$. Note that the intersection $A \cap B \in \mathcal{F}$ of two P -almost sure events $A, B \in \mathcal{F}$, is itself also P -almost sure. Now, for any probability space $(\mathcal{Y}, \mathcal{F}, P)$ and any extended real-valued function $f: \mathcal{Y} \rightarrow \overline{\mathbb{R}}$ that is \mathcal{F} -measurable, the (measure-theoretic) expectation $E(f)$ of f will be defined using the Lebesgue integral [1, 3, 15]: $E(f) := \int_{\mathcal{Y}} f dP$. As we already mentioned in the main text, this integral is definitely well-defined if f is non-negative and is otherwise only defined if $\min\{\int_{\mathcal{Y}} f^+ dP, \int_{\mathcal{Y}} f^- dP\} < +\infty$, with $f^+ = f \vee 0$ and $f^- = -f \wedge 0$, in which case we let $\int_{\mathcal{Y}} f dP := \int_{\mathcal{Y}} f^+ dP - \int_{\mathcal{Y}} f^- dP$. When we say that ‘ $E(f)$ exists’, we take this to mean that its defining integral $\int_{\mathcal{Y}} f dP$ is well-defined. Furthermore, an extended real-valued function $f: \mathcal{Y} \rightarrow \overline{\mathbb{R}}$ is called P -*integrable* if it is \mathcal{F} -measurable and $\int_{\mathcal{Y}} |f| dP = \int_{\mathcal{Y}} f^+ dP + \int_{\mathcal{Y}} f^- dP < +\infty$.

The (measure-theoretic) conditional expectation $E(f|\mathcal{G})$ of a non-negative \mathcal{F} -measurable extended real-valued function f with respect to a σ -algebra $\mathcal{G} \subseteq \mathcal{F}$, is any extended real-valued function on \mathcal{Y} that is \mathcal{G} -measurable and that satisfies $\int_A f dP = \int_A E(f|\mathcal{G}) dP$ or, equivalently, $\int_{\mathcal{Y}} f \mathbf{1}_A dP = \int_{\mathcal{Y}} E(f|\mathcal{G}) \mathbf{1}_A dP$ for all $A \in \mathcal{G}$. If f is \mathcal{F} -measurable and non-negative, the conditional expectation $E(f|\mathcal{G})$ exists and is unique up to a P -null set [1, Section 2.7]. The value of $E(f|\mathcal{G})$ on a P -null set can be chosen arbitrarily since it will not change the value of the integral $\int_A E(f|\mathcal{G}) dP = \int_{\mathcal{Y}} E(f|\mathcal{G}) \mathbf{1}_A dP$; see Property M3 below. Similarly to the unconditional case, the measure-theoretic conditional expectation $E(f|\mathcal{G})$ of a general \mathcal{F} -measurable extended real-valued function f with respect to a σ -algebra $\mathcal{G} \subseteq \mathcal{F}$ is only defined if $\min\{E(f^+|\mathcal{G}), E(f^-|\mathcal{G})\} < +\infty$ P -almost surely, and then $E(f|\mathcal{G}) := E(f^+|\mathcal{G}) - E(f^-|\mathcal{G})$ up to a P -null set. If f is P -integrable, such a conditional expectation $E(f|\mathcal{G})$ always exists; see also [3, Section 34]. We moreover have the following convenient properties.

Lemma 22. *For any probability space $(\mathcal{Y}, \mathcal{F}, P)$ and any two extended real-valued functions f and g that are \mathcal{F} -measurable, the following properties hold, provided that the considered unconditional or conditional expectations exist.*

M1. $E(af + b) = aE(f) + b$ for all $a, b \in \mathbb{R}$;

M2. $f \leq g \Rightarrow E(f) \leq E(g)$;

M3. $f = g$ P -almost surely $\Rightarrow E(f) = E(g)$;

M4. $E(f|\mathcal{F}) = f$ P -almost surely;

M5. $E(f|\mathcal{G}^*) = E(f)$ where $\mathcal{G}^* = \{\emptyset, \mathcal{Y}\}$;

M6. $E(E(f|\mathcal{G}_2)|\mathcal{G}_1) = E(f|\mathcal{G}_1)$ P -almost surely for every two σ -algebras $\mathcal{G}_1 \subseteq \mathcal{G}_2 \subseteq \mathcal{F}$.

M7. $\inf f \leq E(f|\mathcal{G})$ P -almost surely and $E(f|\mathcal{G}) \leq \sup f$ P -almost surely for any σ -algebra $\mathcal{G} \subseteq \mathcal{F}$;

Proof. Let us first show that M5 follows from [1, Section 2.7.4(E*)]. There, it is stated that $E(f|\mathcal{G}^*) = E(f)$ holds P -almost surely. However, since $\mathcal{G}^* = \{\emptyset, \mathcal{Y}\}$ and $E(f|\mathcal{G}^*)$ is \mathcal{G}^* -measurable, we find that $E(f|\mathcal{G}^*)$ must be constant. So if $E(f|\mathcal{G}^*) = E(f)$ holds P -almost surely, it also holds on the entire domain \mathcal{Y} .

To prove Property M1, consider [1, Section 2.7.4(D*)] which in particular states that $E(af + b|\mathcal{G}^*) = aE(f|\mathcal{G}^*) + E(b|\mathcal{G}^*)$ P -almost surely. Since $E(b|\mathcal{G}^*) = b$ P -almost surely by [1, Section 2.7.4(A*)], we have that $E(af + b|\mathcal{G}^*) =$

$aE(f|\mathcal{G}^*) + b$ P-almost surely, which by M5 implies that $E(af + b) = aE(f) + b$ P-almost surely. Since both sides are constants, M1 follows.

Property M2 follows from [1, Section 2.7.4(B*)], which states that $E(f|\mathcal{G}^*) \leq E(g|\mathcal{G}^*)$ P-almost surely if $f \leq g$ P-almost surely. Indeed, if $f \leq g$, then certainly $f \leq g$ P-almost surely and then [1, Section 2.7.4(B*)] together with M5 implies that $E(f) \leq E(g)$ P-almost surely; since both sides are constants, M2 follows.

To prove M3, note that $f = g$ P-almost surely implies that both $f \leq g$ P-almost surely and $f \geq g$ P-almost surely. Then M3 follows by applying the previous reasoning to both of these inequalities. Furthermore, Properties M4 and M6 are taken directly from [1, Section 2.7.4].

To see that property M7 holds, consider [1, Section 2.7.4(A*)], which in particular says that $E(\sup f|\mathcal{G}) = \sup f$ P-almost surely. Combining this with monotonicity [1, Section 2.7.4(B*)], and taking into account that the intersection of two P-almost sure events is itself also P-almost sure, we indeed find that $E(f|\mathcal{G}) \leq \sup f$ P-almost surely. The fact that $\inf f \leq E(f|\mathcal{G})$ P-almost surely can then easily be obtained using the linearity of $E(\cdot|\mathcal{G})$ [1, Section 2.7.4(D*)]. \square

Lemma 23. *Consider any probability space $(\mathcal{Y}, \mathcal{F}, P)$, any σ -algebra $\mathcal{G} \subseteq \mathcal{F}$ and any sequence $\{f_n\}_{n \in \mathbb{N}_0}$ of extended real-valued functions such that f_n is \mathcal{F} -measurable for all $n \in \mathbb{N}_0$.*

M8. *If $\{f_n\}_{n \in \mathbb{N}_0}$ is non-decreasing and there is an \mathcal{F} -measurable function f^* such that $E(f^*) > -\infty$ and $f_n \geq f^*$ for all $n \in \mathbb{N}_0$, then*

$$\lim_{n \rightarrow +\infty} E(f_n) = E(f) \text{ where } \lim_{n \rightarrow +\infty} f_n = f.$$

M9. *If $\{f_n\}_{n \in \mathbb{N}_0}$ is non-increasing and there is an \mathcal{F} -measurable function f^* such that $E(f^*) < +\infty$ and $f_n \leq f^*$ for all $n \in \mathbb{N}_0$, then*

$$\lim_{n \rightarrow +\infty} E(f_n) = E(f) \text{ where } \lim_{n \rightarrow +\infty} f_n = f.$$

Proof. Both properties can directly be obtained by applying M5 to the properties stated in [1, Theorem 2.7.2] and observing that both sides of the acquired equations will be constants (allowing us to infer that the equations always hold, instead of P-almost surely as is the case for the more general versions in [1, Theorem 2.7.2]). \square

A (discrete) *filtration* $\{\mathcal{F}_n\}_{n \in \mathbb{N}_0}$ on a measurable space $(\mathcal{Y}, \mathcal{F})$ is a sequence of increasing σ -algebras in \mathcal{F} ; so $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}$. We will then also use \mathcal{F}_∞ to denote the smallest σ -algebra $\sigma(\cup_{n \in \mathbb{N}_0} \mathcal{F}_n)$ generated by the σ -algebras \mathcal{F}_n . We say that $(\mathcal{Y}, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0})$ is a *filtered measurable space* if $(\mathcal{Y}, \mathcal{F})$ is equipped with a filtration $\{\mathcal{F}_n\}_{n \in \mathbb{N}_0}$, and moreover say that $(\mathcal{Y}, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0}, P)$ is a *filtered probability space* if it additionally has a σ -additive measure P on \mathcal{F} . A sequence $\{\mathcal{N}_n\}_{n \in \mathbb{N}_0}$ of extended real-valued functions on \mathcal{Y} is called a *measure-theoretic process* in a filtered probability space $(\mathcal{Y}, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0}, P)$ if \mathcal{N}_n is \mathcal{F}_n -measurable for all $n \in \mathbb{N}_0$. It is moreover called a *measure-theoretic martingale* if \mathcal{N}_n is real-valued and $E(\mathcal{N}_{n+1}|\mathcal{F}_n) = \mathcal{N}_n$ P-almost surely for all $n \in \mathbb{N}_0$, where the existence of $E(\mathcal{N}_{n+1}|\mathcal{F}_n)$ is an implicit condition. The following two results are fundamental in establishing a relation between the measure-theoretic and the game-theoretic framework.

Proposition 24 (Lévy's zero-one law; [2, Theorem 7.4.3]). *For any filtered probability space $(\mathcal{Y}, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0}, P)$ and any P-integrable function f , we have that*

$$\lim_{n \rightarrow +\infty} E(f|\mathcal{F}_n) = E(f|\mathcal{F}_\infty) \text{ P-almost surely}$$

Proposition 25 (Ville's theorem; [4, Proposition 8.14]). *For any filtered probability space $(\mathcal{Y}, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0}, P)$ and any $A \in \mathcal{F}_\infty$, we have that $P(A) = 0$ if and only if there is a non-negative measure-theoretic martingale that converges to $+\infty$ on A .*

Appendix B. Proof of Proposition 10

Proposition 10 will be shown to hold by establishing the equivalence between the global measure-theoretic and global game-theoretic (upper) expectation on a particular restricted domain and then applying Theorem 6. Before we relate both global models, however, we first establish the following alternative characterisation for the unique extension \overline{Q}_s^1 of the local linear (upper) expectation \overline{Q}_s corresponding to any probability mass function p_s on \mathcal{X} :

Lemma 26. *Consider any probability mass function p_s on \mathcal{X} and let \overline{Q}_s be the unique linear expectation on $\mathcal{L}(\mathcal{X})$ corresponding to p_s . Then we have that $\overline{Q}_s^1(f) = \sum_{x \in \mathcal{X}} f(x)p_s(x)$ for all $f \in \overline{\mathcal{L}}(\mathcal{X})$.*

Proof. Let $\overline{E}_s: \overline{\mathcal{L}}(\mathcal{X}) \rightarrow \overline{\mathbb{R}}$ be defined by $\overline{E}_s(f) := \sum_{x \in \mathcal{X}} f(x)p_s(x)$ for all $f \in \overline{\mathcal{L}}(\mathcal{X})$. We show that \overline{Q}_s^1 coincides with \overline{E}_s . It is clear that this is the case on the domain $\mathcal{L}(\mathcal{X})$ of all gambles on \mathcal{X} ; this follows immediately from the fact that \overline{Q}_s^1 is an extension of \overline{Q}_s together with the fact that $\overline{E}_s(f) = \sum_{x \in \mathcal{X}} f(x)p_s(x) = \overline{Q}_s(f)$ for all $f \in \mathcal{L}(\mathcal{X})$. To see that both coincide on $\overline{\mathcal{L}}_b(\mathcal{X})$, consider any $f \in \overline{\mathcal{L}}_b(\mathcal{X})$ and observe that

$$\begin{aligned} \overline{Q}_s^1(f) &\stackrel{\text{C9}}{=} \lim_{c \rightarrow +\infty} \overline{Q}_s^1(f^{\wedge c}) = \lim_{c \rightarrow +\infty} \sum_{x \in \mathcal{X}} f^{\wedge c}(x)p_s(x) = \sum_{x \in \mathcal{X}} \lim_{c \rightarrow +\infty} f^{\wedge c}(x)p_s(x) = \sum_{x \in \mathcal{X}} p_s(x) \lim_{c \rightarrow +\infty} f^{\wedge c}(x) \\ &= \sum_{x \in \mathcal{X}} p_s(x)f(x) = \overline{E}_s(f), \end{aligned}$$

where the second equality follows from the fact that $f^{\wedge c} \in \mathcal{L}(\mathcal{X})$ for any $c \in \mathbb{R}$ and, as we have just shown, \overline{Q}_s^1 and \overline{E}_s coincide on $\mathcal{L}(\mathcal{X})$; where the third equality follows from the fact that \mathcal{X} is finite and that, for all $x \in \mathcal{X}$, $f^{\wedge c}(x)p_s(x)$ is real and non-decreasing in c ; where the fourth equality follows from the fact that $p_s(x) \geq 0$ for all $x \in \mathcal{X}$ and our convention that $0 \cdot (+\infty) = 0$; and where the last equality follows from the definition of \overline{E}_s . In an analogous way, we show that \overline{Q}_s^1 and \overline{E}_s coincide on their entire domain $\overline{\mathcal{L}}(\mathcal{X})$: consider any $f \in \overline{\mathcal{L}}(\mathcal{X})$ and note that

$$\begin{aligned} \overline{Q}_s^1(f) &\stackrel{\text{C9}}{=} \lim_{c \rightarrow -\infty} \overline{Q}_s^1(f^{\vee c}) = \lim_{c \rightarrow -\infty} \sum_{x \in \mathcal{X}} f^{\vee c}(x)p_s(x) = \sum_{x \in \mathcal{X}} \lim_{c \rightarrow -\infty} f^{\vee c}(x)p_s(x) = \sum_{x \in \mathcal{X}} p_s(x) \lim_{c \rightarrow -\infty} f^{\vee c}(x) \\ &= \sum_{x \in \mathcal{X}} p_s(x)f(x) = \overline{E}_s(f), \end{aligned}$$

where the second equality follows from the fact that $f^{\vee c} \in \overline{\mathcal{L}}_b(\mathcal{X})$ for any $c \in \mathbb{R}$ and, as we have just shown, \overline{Q}_s^1 and \overline{E}_s coincide on $\overline{\mathcal{L}}_b(\mathcal{X})$; where the third equality follows from the fact that \mathcal{X} is finite, our convention that $+\infty - \infty = -\infty + \infty = +\infty$ and that, for all $x \in \mathcal{X}$, $f^{\vee c}(x)p_s(x)$ is in $\mathbb{R} \cup \{+\infty\}$ and non-decreasing in c ; the fourth equality follows from the fact that $p_s(x) \geq 0$ for all $x \in \mathcal{X}$ and our convention that $0 \cdot (-\infty) = 0$; and the last equality follows from the definition of \overline{E}_s . \square

Henceforth, for any precise probability tree p and any $s \in \mathcal{X}^*$, we will always let \overline{Q}_s^1 be this unique extension of the linear expectation \overline{Q}_s corresponding to the probability mass function $p(\cdot|s)$. We will then say that a map $\mathcal{M}: \mathcal{X}^* \rightarrow \overline{\mathbb{R}}$ is a game-theoretic supermartingale *with respect to a precise probability tree p* if \mathcal{M} is a game-theoretic supermartingale with respect to the local upper expectations \overline{Q}_s^1 .

Furthermore, for any precise probability tree $p: s \in \mathcal{X}^* \mapsto p(\cdot|s)$, we let $P: \mathcal{F} \times \mathcal{X}^* \rightarrow \mathbb{R}$ be the corresponding conditional probability measure as discussed in Section 9, where \mathcal{F} is the σ -algebra generated by all cylinder events. Recall that, for any $s \in \mathcal{X}^*$, the map $P(\cdot|s) = P_s$ is then a probability measure on \mathcal{F} . This allows us to apply the concepts and results in Appendix A here, by considering the probability space $(\Omega, \mathcal{F}, P_s)$, for any $s \in \mathcal{X}^*$. For notational convenience, we will let $E_{|s}(f)$ be the Lebesgue integral $\int_{\Omega} f dP_s$ with respect to the measure P_s for all \mathcal{F} -measurable $f \in \overline{\mathbb{V}}$ and all $s \in \mathcal{X}^*$ such that $\int_{\Omega} f dP_s$ exists. So $E_{|s}(f)$ is an alternative notation for $E_{\text{meas}, p}(f|s)$. We introduce this notation because it allows us to write $E_{|s}(f|\mathcal{G})$ to denote a \mathcal{G} -measurable function representing the measure-theoretic expectation of f conditional on a σ -algebra $\mathcal{G} \subseteq \mathcal{F}$, as defined in Appendix A. Furthermore, we equip the measurable space (Ω, \mathcal{F}) with the filtration $\{\mathcal{F}_n\}_{n \in \mathbb{N}_0}$ where, for any $n \in \mathbb{N}_0$, \mathcal{F}_n is the σ -algebra generated by all cylinder events $\Gamma(x_{1:n})$ where $x_{1:n} \in \mathcal{X}^n$. Note that, for any $n \in \mathbb{N}_0$, an

\mathcal{F}_n -measurable function is then an n -measurable variable because the cylinder events $\Gamma(x_{1:n})$ form the atoms of \mathcal{F}_n and \mathcal{X} is finite. Hence, any measure-theoretic process $\{\mathcal{N}_n\}_{n \in \mathbb{N}_0}$ in $(\Omega, \mathcal{F}, \{\mathcal{F}_n\}_{n \in \mathbb{N}_0})$ is a sequence of n -measurable variables. This allows us to write $\mathcal{N}_n(x_{1:n})$ for any $n \in \mathbb{N}_0$ and any $x_{1:n} \in \mathcal{X}^n$ to mean the constant value of \mathcal{N}_n on all $\omega \in \Gamma(x_{1:n})$.

Lemma 27. *Consider any $x_{1:n} \in \mathcal{X}^*$, any precise probability tree p and let \mathbb{P} be the corresponding conditional probability measure. Then, for any non-negative measure-theoretic martingale $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ in the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_m\}_{m \in \mathbb{N}_0}, \mathbb{P}_{|x_{1:n}})$, there is a non-negative game-theoretic (super)martingale¹¹ \mathcal{M} with respect to the tree p such that $\liminf_{i \rightarrow +\infty} \mathcal{M} \geq_{x_{1:n}} \liminf_{i \rightarrow +\infty} \mathcal{N}_i$ and moreover $\mathcal{M}(x_{1:n}) = \mathcal{N}_0(\square)$.*

Proof. Fix any non-negative measure-theoretic martingale $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ and let $\mathcal{M} : \mathcal{X}^* \rightarrow \overline{\mathbb{R}}$ be defined by

$$\mathcal{M}(z_{1:m}) := \begin{cases} \mathcal{N}_m(z_{1:m}) & \text{if } n \leq m, z_{1:n} = x_{1:n} \text{ and } \mathbb{P}_{|x_{1:n}}(z_{1:m}) > 0; \\ +\infty & \text{if } n \leq m, z_{1:n} = x_{1:n} \text{ and } \mathbb{P}_{|x_{1:n}}(z_{1:m}) = 0; \\ \mathcal{N}_n(x_{1:n}) & \text{otherwise,} \end{cases} \quad \text{for all } z_{1:m} \in \mathcal{X}^*.$$

We show that \mathcal{M} is a non-negative game-theoretic supermartingale such that $\mathcal{M}(x_{1:n}) = \mathcal{N}_0(\square)$ and $\liminf_{i \rightarrow +\infty} \mathcal{M} \geq_{x_{1:n}} \liminf_{i \rightarrow +\infty} \mathcal{N}_i$.

Let us first show that $\mathcal{M}(z_{1:m}) \geq \overline{\mathbb{Q}}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$ for all $z_{1:m} \in \mathcal{X}^*$ and therefore, that \mathcal{M} is a game-theoretic supermartingale. Recall that, because $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ is a non-negative measure-theoretic martingale, we have that $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ is a sequence of non-negative real i -measurable variables such that $\mathbb{E}_{|x_{1:n}}(\mathcal{N}_{i+1} | \mathcal{F}_i) = \mathcal{N}_i \mathbb{P}_{|x_{1:n}}$ -almost surely for all $i \in \mathbb{N}_0$ (note that the considered expectations exist because $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ is non-negative). For all $i \in \mathbb{N}_0$ and all $A \in \mathcal{F}_i$, since $\mathbb{E}_{|x_{1:n}}(\mathcal{N}_{i+1} | \mathcal{F}_i)$ and \mathcal{N}_i only differ on a $\mathbb{P}_{|x_{1:n}}$ -null set, we have that $\int_A \mathbb{E}_{|x_{1:n}}(\mathcal{N}_{i+1} | \mathcal{F}_i) d\mathbb{P}_{|x_{1:n}} = \int_A \mathcal{N}_i d\mathbb{P}_{|x_{1:n}}$ because of M3. In particular, this implies that

$$\int_{\Gamma(z_{1:m})} \mathbb{E}_{|x_{1:n}}(\mathcal{N}_{m+1} | \mathcal{F}_m) d\mathbb{P}_{|x_{1:n}} = \int_{\Gamma(z_{1:m})} \mathcal{N}_m d\mathbb{P}_{|x_{1:n}} \quad \text{for any } m \in \mathbb{N}_0 \text{ and any } z_{1:m} \in \mathcal{X}^m.$$

Moreover, since $\Gamma(z_{1:m}) \in \mathcal{F}_m$, it follows from the definition of the measure-theoretic conditional expectation that $\int_{\Gamma(z_{1:m})} \mathcal{N}_{m+1} d\mathbb{P}_{|x_{1:n}} = \int_{\Gamma(z_{1:m})} \mathbb{E}_{|x_{1:n}}(\mathcal{N}_{m+1} | \mathcal{F}_m) d\mathbb{P}_{|x_{1:n}}$. Combining both equalities, we find that

$$\int_{\Gamma(z_{1:m})} \mathcal{N}_{m+1} d\mathbb{P}_{|x_{1:n}} = \int_{\Gamma(z_{1:m})} \mathcal{N}_m d\mathbb{P}_{|x_{1:n}} = \mathcal{N}_m(z_{1:m}) \int_{\Gamma(z_{1:m})} d\mathbb{P}_{|x_{1:n}} = \mathcal{N}_m(z_{1:m}) \mathbb{P}_{|x_{1:n}}(z_{1:m}),$$

where the second equality follows from M1 and the fact that \mathcal{N}_m is constant and real on the cylinder event $\Gamma(z_{1:m})$. Moreover, because \mathcal{N}_{m+1} is $(m+1)$ -measurable and real-valued, the term on the left hand side of the equation above reduces to the finite sum $\sum_{z_{m+1} \in \mathcal{X}} \mathcal{N}_{m+1}(z_{1:m+1}) \mathbb{P}_{|x_{1:n}}(z_{1:m+1})$, allowing us to conclude that

$$\mathcal{N}_m(z_{1:m}) \mathbb{P}_{|x_{1:n}}(z_{1:m}) = \sum_{z_{m+1} \in \mathcal{X}} \mathcal{N}_{m+1}(z_{1:m+1}) \mathbb{P}_{|x_{1:n}}(z_{1:m+1}) \quad \text{for any } m \in \mathbb{N}_0 \text{ and any } z_{1:m} \in \mathcal{X}^m. \quad (\text{B.1})$$

Consider now first any $m \geq n$ and any $z_{1:m} \in \mathcal{X}^m$ such that $z_{1:n} = x_{1:n}$ and $\mathbb{P}_{|x_{1:n}}(z_{1:m}) > 0$. It then follows from the definition of \mathcal{M} that

$$\begin{aligned} \mathcal{M}(z_{1:m}) \mathbb{P}_{|x_{1:n}}(z_{1:m}) &= \mathcal{N}_m(z_{1:m}) \mathbb{P}_{|x_{1:n}}(z_{1:m}) \stackrel{(\text{B.1})}{=} \sum_{z_{m+1} \in \mathcal{X}} \mathcal{N}_{m+1}(z_{1:m+1}) \mathbb{P}_{|x_{1:n}}(z_{1:m+1}) \\ &= \sum_{z_{m+1} \in \mathcal{X}} \mathcal{M}(z_{1:m+1}) \mathbb{P}_{|x_{1:n}}(z_{1:m+1}) \\ &= \sum_{z_{m+1} \in \mathcal{X}} \mathcal{M}(z_{1:m+1}) \mathbb{P}_{|x_{1:n}}(z_{1:m}) p(z_{m+1} | z_{1:m}) \\ &= \mathbb{P}_{|x_{1:n}}(z_{1:m}) \sum_{z_{m+1} \in \mathcal{X}} \mathcal{M}(z_{1:m+1}) p(z_{m+1} | z_{1:m}) = \mathbb{P}_{|x_{1:n}}(z_{1:m}) \overline{\mathbb{Q}}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot)), \end{aligned}$$

¹¹As shown in the proof, the inequality $\mathcal{M}(z_{1:m}) \geq \overline{\mathbb{Q}}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$ is actually an equality, for all $z_{1:m} \in \mathcal{X}^*$. Combining this with the non-negativity of \mathcal{M} and the characterisation of the upper expectations $\overline{\mathbb{Q}}_{z_{1:m}}^\dagger$ that we established in Lemma 26, it can then be deduced that $-\mathcal{M}$ is also a game-theoretic supermartingale. In this case, we can therefore actually call \mathcal{M} a *game-theoretic martingale*.

where the third equality follows because $\mathcal{N}_{m+1}(z_{1:m+1})$ only differs from $\mathcal{M}(z_{1:m+1})$ if $P_{|x_{1:n}}(z_{1:m+1}) = 0$ [because $n \leq m+1$ and $z_{1:n} = x_{1:n}$] and our convention that $0 \cdot (+\infty) = 0$, where the fourth equality follows from Equation (6) together with the fact that $n < m+1$ and $z_{1:n} = x_{1:n}$, and where the last equality follows from the expression for $\bar{Q}_{z_{1:m}}^\dagger$ that we established in Lemma 26. Recall that $P_{|x_{1:n}}(z_{1:m}) > 0$, so we can divide both sides by $P_{|x_{1:n}}(z_{1:m})$ to obtain that $\mathcal{M}(z_{1:m}) = \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$. As a consequence, the condition that $\mathcal{M}(z_{1:m}) \geq \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$ is satisfied for all $z_{1:m} \in \mathcal{X}^*$ such that $m \geq n$, $z_{1:n} = x_{1:n}$ and $P_{|x_{1:n}}(z_{1:m}) > 0$.

Second, consider any $z_{1:m} \in \mathcal{X}^*$ such that $m \geq n$, $z_{1:n} = x_{1:n}$ and $P_{|x_{1:n}}(z_{1:m}) = 0$. Then we also have that $P_{|x_{1:n}}(z_{1:m+1}) = 0$ for any $z_{m+1} \in \mathcal{X}$ because $\Gamma(z_{1:m+1}) \subset \Gamma(z_{1:m})$, and it therefore follows from the definition of \mathcal{M} that $\mathcal{M}(z_{1:m}) = \mathcal{M}(z_{1:m+1}) = +\infty$ for all $z_{m+1} \in \mathcal{X}$. Hence, using Lemma 26, we find that

$$\bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot)) = \sum_{z_{m+1} \in \mathcal{X}} \mathcal{M}(z_{1:m+1}) p(z_{m+1}|z_{1:m}) = +\infty = \mathcal{M}(z_{1:m}),$$

which establishes that indeed $\mathcal{M}(z_{1:m}) \geq \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$. So we have that $\mathcal{M}(z_{1:m}) \geq \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$ for all $z_{1:m} \in \mathcal{X}^*$ such that $m \geq n$ and $z_{1:n} = x_{1:n}$.

Thirdly, consider any other situation $z_{1:m} \in \mathcal{X}^*$. That is, consider any $z_{1:m} \in \mathcal{X}^*$ such that $n > m$ or $z_{1:n} \neq x_{1:n}$. It then follows from the definition of \mathcal{M} that $\mathcal{M}(z_{1:m}) = \mathcal{N}_n(x_{1:n})$ and, as we will now show, that also $\mathcal{M}(z_{1:m+1}) = \mathcal{N}_n(x_{1:n})$ for all $z_{m+1} \in \mathcal{X}$. Our assumption about $z_{1:m}$ tells us that there are two cases: either $n \leq m$ and $z_{1:n} \neq x_{1:n}$, either $n > m$. If $n \leq m$ and $z_{1:n} \neq x_{1:n}$, then also $n \leq m+1$ and therefore, due to the definition of \mathcal{M} , we find that $\mathcal{M}(z_{1:m+1}) = \mathcal{N}_n(x_{1:n})$ for all $z_{m+1} \in \mathcal{X}$. If $n > m$, then either $n > m+1$ or $n = m+1$. If $n > m+1$ or, equivalently, $n \not\leq m+1$, then it is once more clear from the definition of \mathcal{M} that $\mathcal{M}(z_{1:m+1}) = \mathcal{N}_n(x_{1:n})$ for all $z_{m+1} \in \mathcal{X}$. If $n = m+1$, then for any $z_{m+1} \in \mathcal{X}$, we either have that $z_{1:m+1} = x_{1:n}$ or that $z_{1:m+1} \neq x_{1:n}$. In the latter case, it is clear, again due to the definition of \mathcal{M} , that $\mathcal{M}(z_{1:m+1}) = \mathcal{N}_n(x_{1:n})$. Otherwise, so if $z_{1:m+1} = x_{1:n}$, the definition of \mathcal{M} also implies that $\mathcal{M}(z_{1:m+1}) = \mathcal{M}(x_{1:n}) = \mathcal{N}_n(x_{1:n})$ because $P_{|x_{1:n}}(x_{1:n}) = 1 > 0$ according to Equation (6). We conclude that, indeed, $\mathcal{M}(z_{1:m}) = \mathcal{N}_n(x_{1:n})$ and $\mathcal{M}(z_{1:m+1}) = \mathcal{N}_n(x_{1:n})$ for all $z_{m+1} \in \mathcal{X}$. It now follows trivially from Lemma 26 that $\mathcal{M}(z_{1:m}) = \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$, and therefore definitely that $\mathcal{M}(z_{1:m}) \geq \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$.

So we have proved that $\mathcal{M}(z_{1:m}) \geq \bar{Q}_{z_{1:m}}^\dagger(\mathcal{M}(z_{1:m}\cdot))$ for all $z_{1:m} \in \mathcal{X}^*$ and therefore, that \mathcal{M} is indeed a game-theoretic supermartingale with respect to the tree p . Moreover, \mathcal{M} is non-negative because $\{\mathcal{N}_i\}_{i \in \mathbb{N}_0}$ is non-negative. That $\liminf_{i \rightarrow +\infty} \mathcal{M}(\omega^i) \geq \liminf_{i \rightarrow +\infty} \mathcal{N}_i(\omega)$ holds for all $\omega \in \Gamma(x_{1:n})$, follows immediately from the fact that $\mathcal{M}(\omega^i) \geq \mathcal{N}_i(\omega^i) = \mathcal{N}_i(\omega)$ for any $\omega \in \Gamma(x_{1:n})$ and any $i \geq n$, where the inequality is simply implied by the definition of \mathcal{M} together with the fact that $\omega^n = x_{1:n}$ and $i \geq n$. It still remains to show that $\mathcal{M}(x_{1:n}) = \mathcal{N}_0(\square)$.

We only need to prove that $\mathcal{N}_0(\square) = \mathcal{N}_n(x_{1:n})$, since the desired equality then trivially follows from the definition of \mathcal{M} and the fact that $P_{|x_{1:n}}(x_{1:n}) = 1 > 0$ due to Equation (6). To do so, consider any $z_{1:m} \in \mathcal{X}^*$ such that $m < n$ and $z_{1:m} = x_{1:m}$. Note that, by Equation (6), we then have that $P_{|x_{1:n}}(z_{1:m}) = 1$. Moreover, it should also be clear from Equation (6) that $P_{|x_{1:n}}(z_{1:m+1}) = 1$ if $z_{m+1} = x_{m+1}$ [because then $z_{1:m+1} = x_{1:m+1}$ and $m+1 \leq n$] and that $P_{|x_{1:n}}(z_{1:m+1}) = 0$ otherwise [because then $z_{1:m+1} \neq x_{1:m+1}$]. Hence, plugging this back into (B.1), we find that $\mathcal{N}_m(z_{1:m}) = \mathcal{N}_{m+1}(z_{1:m} x_{m+1})$. Since this holds for any $z_{1:m} \in \mathcal{X}^*$ such that $m < n$ and $z_{1:m} = x_{1:m}$, it follows that $\mathcal{N}_m(x_{1:m}) = \mathcal{N}_{m+1}(x_{1:m+1})$ for all $m < n$. This clearly implies that $\mathcal{N}_0(\square) = \mathcal{N}_1(x_1) = \dots = \mathcal{N}_n(x_{1:n})$. \square

In the following proof, we will write, for any two $f, g \in \bar{\mathbb{V}}$, any $s \in \mathcal{X}^*$ and any (unconditional) probability measure P' on \mathcal{F} , that $f =_s g$ P' -almost surely—and similarly for \geq_s and \leq_s —if the event $\{\omega \in \Gamma(s) : f(\omega) \neq g(\omega)\}$ is P' -null. Note that then $f =_s g$ P' -almost surely if $f = g$ P' -almost surely.

Proof of Proposition 10. Fix any \mathcal{F} -measurable $f' \in \bar{\mathbb{V}}$ and any $x_{1:n} \in \mathcal{X}^*$. It suffices to show that $E_{\text{meas},p}(f'|x_{1:n}) = \bar{E}_V(f'|x_{1:n})$; the desired equality is then automatically implied by Theorem 6. First observe that, because f' is bounded and \mathcal{F} -measurable, $E_{\text{meas},p}(f'|x_{1:n}) = E_{|x_{1:n}}(f')$ exists. We will now prove that $E_{\text{meas},p}(f'|x_{1:n}) = \bar{E}_V(f'|x_{1:n})$ for the non-negative \mathcal{F} -measurable gamble $f := f' - \inf f'$ (the variable f is indeed a gamble because f' is a gamble and therefore $\inf f' \in \mathbb{R}$), which then implies that $E_{\text{meas},p}(f'|x_{1:n}) = \bar{E}_V(f'|x_{1:n})$ because $E_{\text{meas},p}(\cdot|x_{1:n})$ and \bar{E}_V both satisfy the constant additivity property; see M1 and V4.

We first show that $\bar{E}_V(f|x_{1:n}) \leq E_{\text{meas},p}(f|x_{1:n})$. To do so, we will prove that there is some $c \in \mathbb{R}$ such that, for all $\epsilon > 0$, there is a bounded below game-theoretic supermartingale \mathcal{M}_ϵ with respect to the tree p such that

$\mathcal{M}_\epsilon(x_{1:n}) = E_{\text{meas},p}(f|x_{1:n}) + \epsilon c$ and $\liminf \mathcal{M}_\epsilon \geq_{x_{1:n}} f$. Indeed, the desired inequality then follows immediately from the definition of \bar{E}_V .

Consider the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_m\}_{m \in \mathbb{N}_0}, P_{|x_{1:n}})$ and the corresponding measure-theoretic expectation $E_{|x_{1:n}} = E_{\text{meas},p}(\cdot|x_{1:n})$. Since f is bounded and \mathcal{F} -measurable, it is surely $P_{|x_{1:n}}$ -integrable (the Lebesgue integral of a bounded \mathcal{F} -measurable function is real-valued), and therefore, by Proposition 24 in Appendix A, we have that

$$\lim_{m \rightarrow +\infty} E_{|x_{1:n}}(f|\mathcal{F}_m) = E_{|x_{1:n}}(f|\mathcal{F}_\infty) \quad P_{|x_{1:n}}\text{-almost surely.}$$

Note that $\mathcal{F}_\infty = \sigma(\cup_{m \in \mathbb{N}_0} \mathcal{F}_m) = \mathcal{F}$ is the smallest σ -algebra generated by all cylinder events, which, by M4, implies that $E_{|x_{1:n}}(f|\mathcal{F}_\infty) = f$, $P_{|x_{1:n}}$ -almost surely. Hence, since the intersection of two $P_{|x_{1:n}}$ -almost sure events is itself also $P_{|x_{1:n}}$ -almost sure, we have that $\lim_{m \rightarrow +\infty} E_{|x_{1:n}}(f|\mathcal{F}_m) = f$ $P_{|x_{1:n}}$ -almost surely. Moreover, note that the conditional expectations $E_{|x_{1:n}}(f|\mathcal{F}_m)$ can be chosen in such a way that the sequence $\{E_{|x_{1:n}}(f|\mathcal{F}_m)\}_{m \in \mathbb{N}_0}$ is a non-negative measure-theoretic martingale in the filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_m\}_{m \in \mathbb{N}_0}, P_{|x_{1:n}})$. Indeed, each $E_{|x_{1:n}}(f|\mathcal{F}_m)$ is real-valued and non-negative $P_{|x_{1:n}}$ -almost surely because of M7 and the fact that f is bounded and non-negative. So, since the value of the measure-theoretic conditional expectation can be chosen arbitrarily on a null set, $E_{|x_{1:n}}(f|\mathcal{F}_m)$ can be chosen such that it is real-valued and non-negative everywhere. Moreover, because of M6, we have that $E_{|x_{1:n}}(E_{|x_{1:n}}(f|\mathcal{F}_{m+1})|\mathcal{F}_m) = E_{|x_{1:n}}(f|\mathcal{F}_m)$ $P_{|x_{1:n}}$ -almost surely for all $m \in \mathbb{N}_0$, where $E_{|x_{1:n}}(E_{|x_{1:n}}(f|\mathcal{F}_{m+1})|\mathcal{F}_m)$ exists because $E_{|x_{1:n}}(f|\mathcal{F}_{m+1})$ is \mathcal{F} -measurable and non-negative. So let $\{E_{|x_{1:n}}(f|\mathcal{F}_m)\}_{m \in \mathbb{N}_0}$ be a version of the conditional expectations that forms a non-negative measure-theoretic martingale. We can then use Lemma 27 to infer the existence of a non-negative game-theoretic supermartingale \mathcal{M}_0 with respect to the tree p such that $\liminf \mathcal{M}_0 \geq_{x_{1:n}} \liminf_{m \rightarrow +\infty} E_{|x_{1:n}}(f|\mathcal{F}_m)$ and $\mathcal{M}_0(x_{1:n}) = E_{|x_{1:n}}(f|\mathcal{F}_0)(\square)$. Then, because $\lim_{m \rightarrow +\infty} E_{|x_{1:n}}(f|\mathcal{F}_m) = f$ $P_{|x_{1:n}}$ -almost surely and therefore also $\lim_{m \rightarrow +\infty} E_{|x_{1:n}}(f|\mathcal{F}_m) =_{x_{1:n}} f$ $P_{|x_{1:n}}$ -almost surely, we have that $\liminf \mathcal{M}_0 \geq_{x_{1:n}} f$ $P_{|x_{1:n}}$ -almost surely. Moreover, we have that $\mathcal{M}_0(x_{1:n}) = E_{|x_{1:n}}(f|\mathcal{F}_0)(\square) = E_{|x_{1:n}}(f)$, due to the fact that $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and property M5. So, we can conclude that \mathcal{M}_0 is a non-negative game-theoretic supermartingale with respect to p such that $\mathcal{M}_0(x_{1:n}) = E_{|x_{1:n}}(f) = E_{\text{meas},p}(f|x_{1:n})$ and $\liminf \mathcal{M}_0 \geq_{x_{1:n}} f$ $P_{|x_{1:n}}$ -almost surely.

As our final step towards obtaining \mathcal{M}_ϵ , consider Proposition 25 and note that it ensures that there is a non-negative measure-theoretic supermartingale $\{\mathcal{N}_m\}_{m \in \mathbb{N}_0}$ in $(\Omega, \mathcal{F}, \{\mathcal{F}_m\}_{m \in \mathbb{N}_0}, P_{|x_{1:n}})$ that converges to $+\infty$ on all paths $\omega \in \Gamma(x_{1:n})$ such that $\liminf \mathcal{M}_0(\omega) < f(\omega)$. Indeed, the set of all such paths ω has probability zero because $\liminf \mathcal{M}_0 \geq_{x_{1:n}} f$ $P_{|x_{1:n}}$ -almost surely. Let $c := \mathcal{N}_0(\square)$, which is real-valued because $\{\mathcal{N}_m\}_{m \in \mathbb{N}_0}$ is a measure-theoretic supermartingale. By Lemma 27, we find that there is a non-negative game-theoretic supermartingale \mathcal{M}' with respect to p such that $\liminf \mathcal{M}' \geq_{x_{1:n}} \liminf_{m \rightarrow +\infty} \mathcal{N}_m$ and $\mathcal{M}'(x_{1:n}) = \mathcal{N}_0(\square) = c$. Since $\{\mathcal{N}_m\}_{m \in \mathbb{N}_0}$ converges to $+\infty$ on all paths $\omega \in \Gamma(x_{1:n})$ such that $\liminf \mathcal{M}_0(\omega) < f(\omega)$, it follows that \mathcal{M}' also converges to $+\infty$ on all such paths ω . Consider now any $\epsilon > 0$ and let \mathcal{M}_ϵ be the process defined by $\mathcal{M}_\epsilon(s) := \mathcal{M}_0(s) + \epsilon \mathcal{M}'(s)$ for all $s \in \mathcal{X}^*$. Then \mathcal{M}_ϵ is clearly non-negative—and therefore bounded below—and it is a game-theoretic supermartingale because of [19, Lemma 12].¹² Furthermore, note that $\liminf \mathcal{M}_\epsilon(\omega) \geq f(\omega)$ for all $\omega \in \Gamma(x_{1:n})$. Indeed, if $\liminf \mathcal{M}_0(\omega) \geq f(\omega)$ for some $\omega \in \Gamma(x_{1:n})$, then also $\liminf \mathcal{M}_\epsilon(\omega) \geq f(\omega)$ because ϵ and \mathcal{M}' are non-negative. If $\liminf \mathcal{M}_0(\omega) < f(\omega)$ for some $\omega \in \Gamma(x_{1:n})$, then \mathcal{M}' , and therefore also $\epsilon \mathcal{M}'$, converges to $+\infty$, which, together with the non-negativity of \mathcal{M}_0 , implies that \mathcal{M}_ϵ converges to $+\infty$ on ω . Hence, also in this case, we have that $\liminf \mathcal{M}_\epsilon(\omega) \geq f(\omega)$ so we can conclude that $\liminf \mathcal{M}_\epsilon \geq_{x_{1:n}} f$. Moreover, recall that $\mathcal{M}'(x_{1:n}) = c \in \mathbb{R}$ and that $\mathcal{M}_0(x_{1:n}) = E_{|x_{1:n}}(f)$, so we have that $\mathcal{M}_\epsilon(x_{1:n}) = \mathcal{M}_0(x_{1:n}) + \epsilon \mathcal{M}'(x_{1:n}) = E_{|x_{1:n}}(f) + \epsilon c = E_{\text{meas},p}(f|x_{1:n}) + \epsilon c$. Hence, \mathcal{M}_ϵ satisfies all the desired conditions and we conclude that indeed $\bar{E}_V(f|x_{1:n}) \leq E_{\text{meas},p}(f|x_{1:n})$.

Then we are left to show the remaining inequality $\bar{E}_V(f|x_{1:n}) \geq E_{\text{meas},p}(f|x_{1:n})$. However, this can be easily deduced from the already obtained inequality and the self-conjugacy of $E_{\text{meas},p}$. Indeed, $E_{|x_{1:n}}(-f)$ exists because $-f$ is \mathcal{F} -measurable and bounded, so we can apply M1 to find that $E_{\text{meas},p}(f|x_{1:n}) = E_{|x_{1:n}}(f) = -E_{|x_{1:n}}(-f) = -E_{\text{meas},p}(-f|x_{1:n})$. Since we have already shown that $\bar{E}_V(g|x_{1:n}) \leq E_{\text{meas},p}(g|x_{1:n})$ for all \mathcal{F} -measurable $g \in \mathbb{V}$, we have in particular that $\bar{E}_V(-f|x_{1:n}) \leq E_{\text{meas},p}(-f|x_{1:n})$, which implies that $E_{\text{meas},p}(f|x_{1:n}) = -E_{\text{meas},p}(-f|x_{1:n}) \leq$

¹²Alternatively, instead of using [19, Lemma 12], one could also easily deduce this using the alternative expression for the local models \bar{Q}_s^1 that we established in Lemma 26.

$-\bar{E}_V(-f|x_{1:n}) = \underline{E}_V(f|x_{1:n}) \leq \bar{E}_V(f|x_{1:n})$, where the last inequality follows from V1 and the fact that $\bar{E}_\Lambda = \bar{E}_V$ —and therefore also $\underline{E}_\Lambda = \underline{E}_V$ —because of Theorem 6. \square