

Global Upper Expectations for Discrete-Time Stochastic Processes: In Practice, They Are All The Same!

Natan T’Joens
Jasper De Bock

Foundations Lab for Imprecise Probabilities, ELIS, Ghent University, Belgium

NATAN.TJOENS@UGENT.BE
JASPER.DEBOCK@UGENT.BE

Abstract

We consider three different types of global uncertainty models for discrete-time stochastic processes: measure-theoretic upper expectations, game-theoretic upper expectations and axiomatic upper expectations. The last two are known to be identical. We show that they coincide with measure-theoretic upper expectations on two distinct domains: monotone pointwise limits of finitary gambles, and bounded below Borel-measurable variables. We argue that these domains cover most practical inferences, and that therefore, in practice, it does not matter which model is used.

Keywords: upper expectation, imprecise probabilities, monotone convergence, probability measure, supermartingale, capacitability

1. Introduction

To describe the dynamics of a discrete-time stochastic process, one may choose between a number of different mathematical approaches. There is of course the measure-theoretic option [2, 13, 14]—undoubtedly the most popular one—but one can also use martingales or game-theoretic principles to do so [11, 12, 22]. Each of these approaches has its own unique strengths and flaws, and each of them—rightly or not—has attracted a dedicated group of followers. Our aim here is not to argue for the use of one or the other though, but rather to study the mathematical relation between the (global) uncertainty models that arise from these approaches in a general, imprecise-probabilistic context. As we will see, they turn out to be surprisingly similar.

All the global—imprecise—uncertainty models that we will consider take the form of an upper (or lower) expectation [19, 20]; a non-linear operator that can—but need not—be interpreted as a tight upper bound on a set of expectations. They are called global because they model beliefs about the entire, uncertain path taken by the process. In that sense, they differ from—and are more general than—local uncertainty models, which only give information about how the process is likely to evolve from one time instant to the next. Such local models form the parameters of a stochastic process, whereas

the global uncertainty model that follows from it—in our case, a global upper expectation—extends the information incorporated in these local models. It is the particular way in which this extension is done that distinguishes one type of global model from the other.

We consider three global models. The first is a probabilistic model that is defined as an upper envelope over a set of measure-theoretic global expectations [10, 18]. The second is based on game-theoretic principles, and defined as an infimum over hedging prices; see Refs. [11, 12]. The last is an abstract axiomatic model, whose defining axioms we have motivated in an earlier paper [18] on the basis of both a probabilistic and a behavioural interpretation. We have already shown that the second and third of these three global upper expectations are identical [18]. In this paper, we relate the first—measure-theoretic—one to this common axiomatic/game-theoretic upper expectation.

Our contribution consists in showing that they are equal on two different domains: variables that are monotone (upward or downward) limits of finitary gambles—bounded variables that only depend on the process’ state at a finite number of time instances—and bounded below Borel-measurable variables. Upper expectations on these two types of domains cover the vast majority of inferences encountered in practice; upper and lower¹ expected hitting times, for instance, fall under the first category [9]; upper and lower expected time averages under the second [15]. Hence the title of this paper. That the three considered global upper expectations are equal on such a large domain is relevant in a number of ways. First of all, it leaves no room for discussion when it comes to choosing a global model; it simply does not matter since all of them are equal. Philosophically speaking, it is interesting that, whatever the interpretational point of view and associated system of logical reasoning is, we always end up with exactly the same object. Finally, and maybe most importantly, such a relation provides us with a large number of additional mathematical properties for the models at hand; properties that were previously only known to hold for one or two of these models, suddenly

1. Lower expectations can be derived from upper expectations using conjugacy; see Section 3.1 and Corollary 14.

hold for all three of them. We refer to Refs. [3, 9, 10] for an illustration of how properties acquired in this way have already led to important consequences.

To adhere to the page limit, we have relegated most of our proofs to the appendix of an extended online version of this paper [16].

2. Local Uncertainty Models

A *discrete-time stochastic process* is an infinite sequence $X_1, X_2, \dots, X_k, \dots$ of uncertain states, where the state X_k at each discrete time point $k \in \mathbb{N}$ 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*. Typically, when modelling the dynamics of a stochastic process, one starts off on a local level, by specifying how the process' state X_k is (likely) to evolve from one time instant to the next. In particular, we do this by attaching a so-called local uncertainty model to each possible *situation*; a finite—possibly empty—sequence $x_{1:k} := x_1 x_2 \dots x_k$ of state values that represents a possible history $X_1 = x_1, \dots, X_k = x_k$ up until some time point $k \in \mathbb{N}_0$, with $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. The local model associated with the situation $x_{1:k}$ then models beliefs about the value of the next state X_{k+1} , conditional on the history represented by $x_{1:k}$. We let $\mathcal{X}^* := \cup_{i \in \mathbb{N}_0} \mathcal{X}^i$ be the set of all situations and we denote the *initial (empty) situation* by $\square := x_{1:0} = \mathcal{X}^0$.

Among the most popular types of local uncertainty models are (probability) mass functions p on \mathcal{X} ; for any situation $x_{1:k} \in \mathcal{X}^*$, the mass function $p(\cdot | x_{1:k})$ then provides, for each $x_{k+1} \in \mathcal{X}$, the probability $p(x_{k+1} | x_{1:k})$ that the value of the state X_{k+1} will be equal to x_{k+1} . Such a family of probability mass functions is represented by a single function $p: s \in \mathcal{X}^* \mapsto p(\cdot | s)$, which we call a *precise probability tree*.² What is equivalent, but less of a popular habit, is to attach to each possible situation $x_{1:k} \in \mathcal{X}^*$ an expectation $E_{x_{1:k}}$ on the set $\mathcal{L}(\mathcal{X})$ of all real-valued functions f on \mathcal{X} . These expectations $E_{x_{1:k}}$ may then be interpreted in a measure-theoretic sense, as coming from an underlying family of mass functions $p(\cdot | x_{1:k})$, but they can also be interpreted in a direct behavioural way as a subject's fair prices, as De Finetti does [5].

Unfortunately, irrespective of one's preference between mass functions and linear expectations, both of them are rather inadequate when modelling situations where data is scarce, or when modelling the beliefs of a conservative (risk-averse) subject. In such situations, one can reach for so-called 'imprecise' probability models [19, 20, 1]. These come in many different shapes and forms (e.g. sets of desirable gambles, belief functions, credal sets,...), but, for our purpose of modelling the lo-

cal dynamics of a process, we will only consider two specific—yet wide-spread—ones; credal sets and coherent upper (and lower) expectations.

The first, *credal sets*, are closed (under the topology of pointwise convergence) convex sets of probability mass functions; see e.g. [1, Section 9.2]. If we attach to each situation $s \in \mathcal{X}^*$ a credal set \mathcal{P}_s on \mathcal{X} , then we obtain a so-called *imprecise probability tree* $\mathcal{P}: s \in \mathcal{X}^* \mapsto \mathcal{P}_s$, which we will often simply denote by \mathcal{P} . For any $s \in \mathcal{X}^*$, the associated credal set \mathcal{P}_s may then be interpreted as a set that contains all local mass functions $p(\cdot | s)$ that are deemed 'possible'. Such an imprecise probability tree \mathcal{P} parametrises the stochastic process as a whole, and clearly does so in a more general manner than the precise methods mentioned earlier; precise probability trees correspond to the special case where, for each $s \in \mathcal{X}^*$, \mathcal{P}_s consists of a single mass function $p(\cdot | s)$. We say that a precise probability tree p is *compatible* with an imprecise probability tree \mathcal{P} , and write $p \sim \mathcal{P}$, if $p(\cdot | s) \in \mathcal{P}_s$ for all $s \in \mathcal{X}^*$.

Another—yet equivalent—approach consists in specifying a *local coherent upper (or lower) expectation* \bar{Q}_s for each $s \in \mathcal{X}^*$ [20]: a real-valued function on $\mathcal{L}(\mathcal{X})$ that satisfies, for all $f, g \in \mathcal{L}(\mathcal{X})$ and $\lambda \in \mathbb{R}_{\geq 0}$,

- C1. $\bar{Q}_s(f) \leq \sup f$ [upper bounds];
- C2. $\bar{Q}_s(f + g) \leq \bar{Q}_s(f) + \bar{Q}_s(g)$ [sub-additivity];
- C3. $\bar{Q}_s(\lambda f) = \lambda \bar{Q}_s(f)$ [non-negative homogeneity].

Any such family $(\bar{Q}_s)_{s \in \mathcal{X}^*}$ of local coherent upper expectations will be gathered in a single *upper expectation tree* $\bar{Q}: s \in \mathcal{X}^* \mapsto \bar{Q}_s$, which we will also simply denote by \bar{Q} . For any $x_{1:k} \in \mathcal{X}^*$, the upper expectation $\bar{Q}_{x_{1:k}}$ can be interpreted as representing a subject's minimum selling prices—a generalisation of De Finetti's fair price interpretation for linear expectations.³ More concretely, this interpretation says that, given a situation $x_{1:k} \in \mathcal{X}^*$ and any $f \in \mathcal{L}(\mathcal{X})$, our subject is willing to sell the uncertain—possibly negative—payoff $f(X_{k+1})$ for any price $\alpha \geq \bar{Q}_{x_{1:k}}(f)$. Axioms C1–C3 can then be seen as rationality criteria. We refer to Walley's work [20] for a more detailed motivation and justification for coherent upper (and lower) expectations.

Mathematically speaking, it does not matter whether we use imprecise probability trees or upper expectation trees to characterise a stochastic process, because credal sets and coherent upper expectations—and therefore imprecise probability trees and upper expectation trees—are in a one-to-one relation with each other. In particular, with any imprecise probability tree \mathcal{P} , we

2. The reason why we call it a 'tree' is because it is a map on \mathcal{X}^* , which can naturally be visualised in terms of infinite (event) trees [4, Figure 1].

3. Traditionally, the behavioural interpretation of coherent upper expectations says that they represent infimum selling prices, rather than minimum selling prices; see Ref. [20]. We opt for minimum selling prices here because they fit more naturally with the supermartingales that we will introduce further on.

can associate an upper expectation tree $\bar{Q}_{s,\mathcal{P}}$ that maps each situation $s \in \mathcal{X}^*$ to the upper envelope $\bar{Q}_{s,\mathcal{P}}$ of the linear expectations corresponding to \mathcal{P}_s :

$$\bar{Q}_{s,\mathcal{P}}(f) := \sup \left\{ \sum_{x \in \mathcal{X}} f(x)p(x|s) : p(\cdot|s) \in \mathcal{P}_s \right\},$$

for all $f \in \mathcal{L}(\mathcal{X})$. That each $\bar{Q}_{s,\mathcal{P}}$ is indeed a local coherent upper expectation follows from [20, Theorem 3.6.1]. Conversely, with any upper expectation tree \bar{Q} , we can associate an imprecise probability tree $\mathcal{P}_{s,\bar{Q}}$; for any $s \in \mathcal{X}^*$, its local credal set $\mathcal{P}_{s,\bar{Q}}$ is the closed convex set of all mass functions $p(\cdot|s)$ that are dominated by \bar{Q}_s , in the sense that

$$\sum_{x \in \mathcal{X}} f(x)p(x|s) \leq \bar{Q}_s(f) \text{ for all } f \in \mathcal{L}(\mathcal{X}).$$

It follows once more from [20, Theorem 3.6.1] that this correspondence between upper expectation trees and imprecise probability trees is one-to-one; that is, the map $\mathcal{P} \mapsto \bar{Q}_{s,\mathcal{P}}$ is bijective and $\bar{Q} \mapsto \mathcal{P}_{s,\bar{Q}}$ is its inverse. We say that an imprecise probability tree \mathcal{P} and an upper expectation tree \bar{Q} *agree* if they are related through these mappings.

An important consequence of the one-to-one relation described above is that imprecise probability trees and upper expectation trees can borrow each others interpretation; local credal sets can be interpreted as representing a subject's infimum selling prices, whereas local upper expectations can be interpreted as upper envelopes of the linear expectations associated with an underlying local credal set.

3. Three Types of Global Models

Imprecise probability trees and upper expectation trees describe the dynamics of a stochastic process on a local level—how it changes from one time instant to the next—but they do not tell us anything, at least not directly, about more global features that relate to multiple time instances at once; e.g. the time it takes until the process is in a given state $x \in \mathcal{X}$. We therefore face the following question. How do we turn the local information captured by any of these trees into global information about the process as a whole? Three possible solutions are described in the current section, but we start by introducing some necessary terminology and notation.

3.1. Preliminaries

A *path* $\omega = x_1 x_2 x_3 \dots$ is an infinite sequence of state values and represents a possible evolution of the process. The *sample space* $\Omega := \mathcal{X}^{\mathbb{N}}$ denotes the set of all paths. For any $\omega = x_1 x_2 x_3 \dots \in \Omega$, we let $\omega^k := x_{1:k} \in \mathcal{X}^k$ be the finite sequence that consists of the initial k state values,

and we let $\omega_k := x_k \in \mathcal{X}$ be the k -th state value. An *event* $A \subseteq \Omega$ is a set of paths and, in particular, for any situation $x_{1:k} \in \mathcal{X}^*$, the *cylinder event* $\Gamma(x_{1:k}) := \{\omega \in \Omega : \omega^k = x_{1:k}\}$ is the set of all paths that go through the situation $x_{1:k}$.

We let $\bar{\mathbb{R}} := \mathbb{R} \cup \{+\infty, -\infty\}$ be the extended real numbers, $\bar{\mathbb{R}}_{\geq 0}$ be the subset of non-negative ones, and $\bar{\mathbb{R}}_{> 0}$ be those that are moreover real. We extend the total order relation $<$ on \mathbb{R} to $\bar{\mathbb{R}}$ by positing that $-\infty < c < +\infty$ for all $c \in \mathbb{R}$ and endow $\bar{\mathbb{R}}$ with the associated order topology.

Any extended real-valued function $f : \mathcal{Y} \rightarrow \bar{\mathbb{R}}$ on some non-empty set \mathcal{Y} will be called a *variable*. Any bounded variable—that is, a variable f for which there is a $B \in \bar{\mathbb{R}}_{\geq 0}$ such that $-B \leq f(y) \leq B$ for all $y \in \mathcal{Y}$ —will be called a *gamble*.⁴ The set of all variables will be denoted by $\bar{\mathcal{L}}(\mathcal{Y})$ and the set of all gambles by $\mathcal{L}(\mathcal{Y})$. Note that this definition is in accordance with our earlier use of $\mathcal{L}(\mathcal{X})$, where it denoted the real-valued functions on \mathcal{X} —which are automatically bounded because \mathcal{X} is finite. The elements of $\bar{\mathcal{L}}(\mathcal{X})$ and $\mathcal{L}(\mathcal{X})$ are called *local variables* and *gambles*, respectively. On the other hand, the variables in $\bar{\mathbb{V}} := \bar{\mathcal{L}}(\Omega)$ and $\mathbb{V} := \mathcal{L}(\Omega)$ are called *global variables* and *gambles*, respectively; they may depend on the entire path $\omega \in \Omega$ taken by the process. Variables that only depend on the process' state at a finite number of time instances are called *finitary*; for such a finitary variable $f \in \bar{\mathbb{V}}$, there is an $n \in \mathbb{N}$ and some $g \in \bar{\mathcal{L}}(\mathcal{X}^n)$ such that $f(\omega) = g(\omega^n)$ for all $\omega \in \Omega$. We often make this explicit by writing $f = g(X_{1:n})$, where $g(X_{1:n}) := g \circ X_{1:n}$ and where $X_{1:n}$ is the projection of $\omega \in \Omega$ on its first n state values ω^n . Sometimes, we also allow ourselves a slight abuse of notation by writing $f(x_{1:n})$ to denote the constant value of $f(\omega) = g(x_{1:n})$ on all paths $\omega \in \Omega$ such that $\omega^n = x_{1:n}$. We collect all finitary gambles in the set \mathbb{F} . A special type of global gamble is the *indicator* \mathbb{I}_A of an event A , which assumes the value 1 on A and 0 elsewhere. For any $s \in \mathcal{X}^*$, the indicator $\mathbb{I}_s := \mathbb{I}_{\Gamma(s)}$ of the cylinder event $\Gamma(s)$ is clearly a finitary gamble.

A *global upper expectation*, finally, is a map $\bar{E} : \bar{\mathbb{V}} \times \mathcal{X}^* \rightarrow \bar{\mathbb{R}}$; it maps global variables $f \in \bar{\mathbb{V}}$ and situations $s \in \mathcal{X}^*$ to a corresponding (conditional) upper expectation $\bar{E}(f|s)$. As we will see, such maps can play the role of a global uncertainty model, in the sense that they can represent beliefs or knowledge about the path ω taken by the process, or about the value attained by a global variable f . Apart from global upper expectations, one can also consider *global lower expectations* $\underline{E} : \bar{\mathbb{V}} \times \mathcal{X}^* \rightarrow \bar{\mathbb{R}}$; for each of the models that we will consider, these are conjugate to the corresponding global upper expectation \bar{E} , in the sense that $\underline{E}(f|s) = -\bar{E}(-f|s)$ for all $f \in \bar{\mathbb{V}}$ and $s \in \mathcal{X}^*$. It therefore suffices to focus on only one of them; our theoretical developments focus on \bar{E} , leaving the implications for \underline{E} for Section 6.

4. This choice of terminology is due to Walley [20]. However, for us, the mathematical object of a gamble is not necessarily bound to the interpretation as an uncertain payoff.

3.2. Measure-Theoretic Global Upper Expectations

We start by presenting a traditional measure-theoretic approach, where global upper expectations are defined as upper envelopes of sets of (linear) expectations, and where each of these (linear) expectations on its turn is derived from a different probability measure on Ω .

Consider an imprecise probability tree \mathcal{P} and let $p \sim \mathcal{P}$ be any precise probability tree that is compatible with \mathcal{P} . With each $x_{1:k} \in \mathcal{X}^*$, we associate a probability measure $P_p(\cdot|x_{1:k})$ on the σ -algebra \mathcal{F} generated by all cylinder events as follows. First, for any $\ell \in \mathbb{N}_0$ and any $C \subseteq \mathcal{X}^\ell$, let

$$P_p(C|x_{1:k}) := P_p(\cup_{z_{1:\ell} \in C} \Gamma(z_{1:\ell})|x_{1:k}) := \sum_{z_{1:\ell} \in C} P_p(z_{1:\ell}|x_{1:k}),$$

where $P_p(z_{1:\ell}|x_{1:k}) :=$ (1)

$$\begin{cases} \prod_{i=k}^{\ell-1} p(z_{i+1}|z_{1:i}) & \text{if } k < \ell \text{ and } z_{1:k} = x_{1:k} \\ 1 & \text{if } k \geq \ell \text{ and } z_{1:\ell} = x_{1:\ell} \\ 0 & \text{otherwise.} \end{cases}$$

It is then easy to see that, on the algebra generated by the cylinder events, $P_p(\cdot|x_{1:k})$ forms a finitely additive probability [10, Chapter 3]. Hence, by [2, Theorem 2.3], it is also a countably additive probability—that is, a probability measure—on this algebra and so, by Carathéodory's extension theorem [22, Theorem 1.7], $P_p(\cdot|x_{1:k})$ can be uniquely extended to a probability measure on \mathcal{F} .

In accordance with standard practices, we then associate with every probability measure $P_p(\cdot|s)$ an expectation $E_p(\cdot|s)$ using Lebesgue integration. That is, we let $E_p(f|s) := \int_{\Omega} f dP_p(\cdot|s)$ for all $f \in \bar{\mathbb{V}}$ for which $\int_{\Omega} f dP_p(\cdot|s)$ exists, which is guaranteed if f is \mathcal{F} -measurable and bounded below (or bounded above). For general $f \in \bar{\mathbb{V}}$, we adopt an upper integral $\bar{E}_p(f|s)$ defined by

$$\bar{E}_p(f|s) := \inf\{E_p(g|s) : g \in \bar{\mathbb{V}}_{\sigma,b} \text{ and } g \geq f\}, \quad (2)$$

where $\bar{\mathbb{V}}_{\sigma,b}$ is the set of all bounded below \mathcal{F} -measurable variables in $\bar{\mathbb{V}}$. It follows from [18, Proposition 12] that $\bar{E}_p(\cdot|s)$ coincides with $E_p(\cdot|s)$ on the entire domain where $E_p(\cdot|s)$ is well-defined—that is, where the Lebesgue integral with respect to $P_p(\cdot|s)$ exists—and hence, that $\bar{E}_p(\cdot|s)$ is an extension of $E_p(\cdot|s)$.

Finally, the global upper expectation $\bar{E}_{\mathcal{P}}$ corresponding to the imprecise probability tree \mathcal{P} is defined as the upper envelope of the upper integrals \bar{E}_p corresponding to each of the precise trees $p \sim \mathcal{P}$. That is, for each $f \in \bar{\mathbb{V}}$ and $s \in \mathcal{X}^*$,

$$\bar{E}_{\mathcal{P}}(f|s) := \sup\{\bar{E}_p(f|s) : p \sim \mathcal{P}\}.$$

This definition is in line with the sensitivity analysis interpretation for imprecise probability models [20, Section 1.1.5], which regards them as resulting from a lack of knowledge about a single ideal precise model.

The approach set out above should look familiar to anyone with a measure-theoretic background, and we therefore omit an in-depth conceptual discussion; we instead refer to [18, Section 9] for more details. One aspect, however, that we feel is worth pointing out is the difference between our way of conditioning and what is usually done in measure-theory. Usually, conditional expectations (and probabilities) are derived from a single unconditional probability measure through the Radon-Nikodym derivative [13, Section 2.7.2]. We, on the other hand, associate with each situation $s \in \mathcal{X}^*$ a separate—in the traditional sense, unconditional—probability measure $P_p(\cdot|s)$ and use this probability measure $P_p(\cdot|s)$ to define the expectation $E_p(\cdot|s)$. The reason why we do so is because, unlike the traditional approach, it allows us to condition—in a meaningful way—on (cylinder) events with probability zero; again, we refer to [18, Section 9] for more details.

3.3. Game-Theoretic Global Upper Expectations

The second global model that we will consider is the game-theoretic upper expectation introduced and, for the most part, developed by Shafer and Vovk [11, 12]. This operator is defined in terms of infimum hedging prices; starting capitals that allow a gambler to cover—or hedge—the costs or gains of a given global gamble. These hedging prices—and hence, these game-theoretic upper expectations—are determined using the notion of a supermartingale; a function that describes the possible evolution of a gambler's capital as he gambles in a way that is in accordance with the local models \bar{Q}_s .

Formally, for any upper expectation tree Q , a *supermartingale* \mathcal{M} is a real-valued function on \mathcal{X}^* that satisfies $\bar{Q}_s(\mathcal{M}(s \cdot)) \leq \mathcal{M}(s)$ for all $s \in \mathcal{X}^*$, where $\mathcal{M}(s \cdot) \in \mathcal{L}(\mathcal{X})$ denotes the local gamble that takes the value $\mathcal{M}(sx)$ in $x \in \mathcal{X}$. How can such a supermartingale be interpreted in the way described above? Consider any situation $x_{1:k} \in \mathcal{X}^*$ and a gambler—called 'Skeptic' in Shafer and Vovk's framework—whose current capital equals $\mathcal{M}(x_{1:k})$. Then, recalling our interpretation for the local model $\bar{Q}_{x_{1:k}}$ as representing a subject's minimum selling prices, the condition that $\bar{Q}_{x_{1:k}}(\mathcal{M}(x_{1:k} \cdot)) \leq \mathcal{M}(x_{1:k})$ implies that Skeptic can use his capital $\mathcal{M}(x_{1:k})$ to buy the uncertain reward $\mathcal{M}(x_{1:k} X_{k+1})$ from this subject—called 'Forecaster' in Shafer and Vovk's framework. If Skeptic chooses to commit to such a transaction, he is actually gambling *against* Forecaster, which explains why these players are called Skeptic and Forecaster. So we see that a supermartingale describes the evolution of Skeptic's capital if he chooses, in each situation, to buy a gamble that Forecaster is willing to sell.

A *hedging price* $\alpha \in \mathbb{R}$ for any $f \in \bar{\mathbb{V}}$ is now a real number for which there is a bounded below supermartingale \mathcal{M} that starts in $\mathcal{M}(\square) = \alpha$ and such that $\liminf \mathcal{M}(\omega) :=$

$\liminf_{k \rightarrow +\infty} \mathcal{M}(\omega^k) \geq f(\omega)$ for all $\omega \in \Omega$. A hedging price α for f is therefore worth more to Skeptic than the global gamble f , because he is always able to eventually turn the initial capital $\mathcal{M}(\square) = \alpha$ into a capital that is higher than the uncertain payoff corresponding to f , simply by choosing the right gambles from the ones Forecaster is offering. That \mathcal{M} should be bounded below, represents the condition that Skeptic can borrow at most a finite amount.

For any $f \in \mathbb{V}$, the infimum over all the hedging prices α is then what defines the (unconditional) global game-theoretic upper expectation $\bar{E}_{G, \bar{Q}}(f)$ of f . More generally, the *global game-theoretic upper expectation* of any $f \in \mathbb{V}$ conditional on any $s \in \mathcal{X}^*$, is defined as

$$\bar{E}_{G, \bar{Q}}(f|s) := \inf \{ \mathcal{M}(s) : \mathcal{M} \in \mathbb{M}_b(\bar{Q}), (\forall \omega \in \Gamma(s)) \liminf \mathcal{M}(\omega) \geq f(\omega) \}, \quad (3)$$

where $\mathbb{M}_b(\bar{Q})$ denotes the set of all bounded below supermartingales. The unconditional case corresponds to $s = \square$; so, $\bar{E}_{G, \bar{Q}}(f) := \bar{E}_{G, \bar{Q}}(f|\square)$.

As the attentive reader may have noticed, the definition above only applies to global gambles. So why not to general variables $f \in \bar{\mathbb{V}}$? The reason is that, on this extended domain, the formula presented above would yield an upper expectation with rather weak continuity properties [17, section 8]. A simple solution is to use continuity with respect to so-called upper and lower cuts to extend the domain from $\mathbb{V} \times \mathcal{X}^*$ to $\bar{\mathbb{V}} \times \mathcal{X}^*$.⁵ To do so, for any $f \in \bar{\mathbb{V}}$ and any $c \in \mathbb{R}$, let $f^{\wedge c}$ be defined by $f^{\wedge c}(x) := \min\{f(x), c\}$ for all $x \in \mathcal{X}$, and let $f^{\vee c}$ be defined analogously, as a pointwise maximum. Then we henceforth let $\bar{E}_{G, \bar{Q}}: \bar{\mathbb{V}} \times \mathcal{X}^* \rightarrow \bar{\mathbb{R}}$ be defined by Equation (3) on $\mathbb{V} \times \mathcal{X}^*$, and furthermore impose, for any $s \in \mathcal{X}^*$, that

$$G1. \quad \bar{E}_{G, \bar{Q}}(f|s) = \lim_{c \rightarrow +\infty} \bar{E}_{G, \bar{Q}}(f^{\wedge c}|s) \text{ for all } f \in \bar{\mathbb{V}}_b;$$

$$G2. \quad \bar{E}_{G, \bar{Q}}(f|s) = \lim_{c \rightarrow -\infty} \bar{E}_{G, \bar{Q}}(f^{\vee c}|s) \text{ for all } f \in \bar{\mathbb{V}}.$$

Properties G1 and G2 together clearly imply that $\bar{E}_{G, \bar{Q}}$ is uniquely determined by its values on $\mathbb{V} \times \mathcal{X}^*$. Hence, since $\bar{E}_{G, \bar{Q}}$ on this domain is described by Equation (3), it follows that $\bar{E}_{G, \bar{Q}}$ is uniquely defined on all of $\bar{\mathbb{V}} \times \mathcal{X}^*$.

This way of extending a global game-theoretic upper expectation is not that common, though. A technique that is used more often consists in directly applying Equation (3) to the entire domain $\bar{\mathbb{V}} \times \mathcal{X}^*$, but with the real-valued supermartingales replaced by extended real-valued ones [12, 18, 17]. This of course first requires an extension \bar{Q}_s^1 of the local models \bar{Q}_s to the domain $\mathcal{L}(\mathcal{X})$, which can be done in a way similar to what we have done with $\bar{E}_{G, \bar{Q}}$, by imposing continuity with respect

⁵ This is similar to how [19, Chapter 15] extends the notion of coherence from gambles to unbounded real-valued variables.

to upper and lower cuts.⁶ An extended real-valued supermartingale $\mathcal{M}: \mathcal{X}^* \rightarrow \bar{\mathbb{R}}$ is then characterised by the condition that $\bar{Q}_s^1(\mathcal{M}(s)) \leq \mathcal{M}(s)$ for all $s \in \mathcal{X}^*$. Remarkably enough, the global game-theoretic upper expectation that results from this ‘extended supermartingale’-approach is identical to the operator $\bar{E}_{G, \bar{Q}}$ we have defined above, using Properties G1 and G2; see for example the end of [17, Section 8]. We favor our approach, though, because the use of extended real-valued supermartingales undermines what we think is a key strength of the game-theoretic approach: that supermartingales—and hence the resulting game-theoretic upper expectations—can be given a clear behavioural meaning in terms of betting.

3.4. Axiomatic Global Upper Expectations

Instead of relying on measure-theoretic or game-theoretic principles, one can also simply adopt an abstract global model \bar{E} that is completely characterised by a number of axioms. In particular, starting from any given upper expectation tree \bar{Q} , we suggest to impose the following list of axioms:

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

$$P2. \quad \bar{E}(f|s) = \bar{E}(f \mathbb{I}_s|s) \text{ for all } f \in \mathbb{F} \text{ and all } s \in \mathcal{X}^*.$$

$$P3. \quad \bar{E}(f|X_{1:k}) \leq \bar{E}(\bar{E}(f|X_{1:k+1})|X_{1:k}) \text{ for all } f \in \mathbb{F} \text{ and all } k \in \mathbb{N}_0.$$

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

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

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

Here, as well as further on, we call a sequence $(f_n)_{n \in \mathbb{N}}$ of variables uniformly bounded below if there is some $B \in \mathbb{R}$ such that $f_n(\omega) \geq B$ for all $n \in \mathbb{N}$ and $\omega \in \Omega$. Furthermore, the limit $\lim_{n \rightarrow +\infty} f_n$, as well as all others in this paper, are intended to be taken pointwise.

Axioms P1–P5 are put forward here because, as we argue in [18, Section 4], they can be motivated on the basis of two different interpretations for a global upper expectation; a direct behavioural interpretation in terms of minimum selling prices, or a probabilistic interpretation in terms of sets of linear expectations (or probability measures). Basically, we find Axioms P1–P4

⁶ The choice of extending the local models \bar{Q}_s in this particular way, by imposing continuity with respect to upper and lower cuts, is motivated in [17, Sections 2 and 8] and [18, section 6], and is, as far as the resulting global game-theoretic upper expectation—with extended real-valued supermartingales—is concerned, completely equivalent with how Shafer and Vovk axiomatise their local models in [12, Part II].

straightforward and believe them to be almost unquestionable, regardless of the adopted interpretation. Axiom P5, which imposes a form of continuity, is perhaps more disputable. Nonetheless, compared to other well-known continuity properties, such as dominated convergence or monotone convergence, Property P5 is rather weak because it only applies to sequences of finitary gambles. Note that, in general, finitary gambles play a central role in our axiomatisation; with the exception of monotonicity (Axiom P4), all our axioms exclusively apply to finitary gambles (and their limits). We find this important because, as explained in [18, Section 4], they are the only global variables that we feel can be given a direct operational meaning, and hence, the only global variables for which axioms can be motivated directly. More general global variables in $\bar{\mathbb{V}}$, on the other hand, that depend on an infinite number of state values, or are unbounded or even infinite-valued, should be regarded as abstract idealisations.

Of course, even if we agree upon Axioms P1–P5, it does not necessarily provide us with a global upper expectation because there may be multiple—or, worse, no—global upper expectations satisfying these axioms. The following result shows that there is at least one model that satisfies P1–P5, and that among all the ones that satisfy them, there is a unique most conservative—that is, largest—one. We denote this model by $\bar{E}_{A,\bar{Q}}$.

Theorem 1 ([18, Theorem 6]) *For any upper expectation tree Q , there is a unique most conservative global upper expectation $\bar{E}_{A,\bar{Q}}$ that satisfies P1–P5.*

4. An Equality for Monotone Limits of Finitary Gambles

Having introduced all three global upper expectations, we can finally turn to the central problem of this paper: how are these upper expectations related to each other? More specifically, we ask ourselves the following. If the parameters of a stochastic process are equivalent—that is, if the trees \mathcal{P} and \bar{Q} agree—are the global models $\bar{E}_{\mathcal{P}}$, $\bar{E}_{G,\bar{Q}}$ and $\bar{E}_{A,\bar{Q}}$ then equal? In a recent paper [18], we have shown that the answer is affirmative for the latter two models.

Theorem 2 ([18, Theorem 6]) *The global upper expectations $\bar{E}_{A,\bar{Q}}$ and $\bar{E}_{G,\bar{Q}}$ are equal.*

So it only remains to study the relationship between the measure-theoretic upper expectation $\bar{E}_{\mathcal{P}}$ and the common upper expectation $\bar{E}_{\bar{Q}} := \bar{E}_{A,\bar{Q}} = \bar{E}_{G,\bar{Q}}$. To do so, we will build on two earlier results, gathered from that same paper [18]; the first one [18, Theorem 14] says that $\bar{E}_{\mathcal{P}}$ coincides with $\bar{E}_{\bar{Q}}$ if the tree \mathcal{P} is a precise probability tree p (and \bar{Q} is the agreeing (upper) expectation

tree); the second one [18, Proposition 21] says that they are also equal for general imprecise probability trees \mathcal{P} , provided that we limit ourselves to finitary gambles. Our main results extend this equality for general imprecise probability trees in two ways: to variables that are monotone limits of finitary gambles and to bounded below \mathcal{F} -measurable variables. In the current section, we work towards establishing the first extension. Our approach is straightforward; we will prove that $\bar{E}_{\mathcal{P}}$ and $\bar{E}_{\bar{Q}}$ are both continuous with respect to monotone sequences of finitary gambles. Since they coincide on finitary gambles, this directly implies the desired equality.

We start by showing that $\bar{E}_{\mathcal{P}}$ and $\bar{E}_{\bar{Q}}$ are both continuous with respect to non-decreasing sequences in $\bar{\mathbb{V}}_b$ —and hence definitely in \mathbb{F} .

Proposition 3 *For any \mathcal{P} and \bar{Q} , any $s \in \mathcal{X}^*$ and any non-decreasing sequence $(f_n)_{n \in \mathbb{N}}$ in $\bar{\mathbb{V}}_b$, we have that*

$$\lim_{n \rightarrow +\infty} \bar{E}_{\mathcal{P}}(f_n | s) = \bar{E}_{\mathcal{P}}(f | s), \text{ with } f = \sup_{n \in \mathbb{N}} f_n = \lim_{n \rightarrow +\infty} f_n,$$

and similarly for $\bar{E}_{\bar{Q}}$.

Proof That the statement holds for $\bar{E}_{\bar{Q}}$ follows immediately from [18, Theorem 9(i)]. To prove the statement for $\bar{E}_{\mathcal{P}}$, recall [18, Theorem 14], which says that, for any precise probability tree p and the agreeing (upper) expectation tree \bar{Q}_p , we have that $\bar{E}_p(g | s) = \bar{E}_{\bar{Q}_p}(g | s)$ for all $g \in \bar{\mathbb{V}}$. Then, since $\bar{E}_{\bar{Q}_p}$ is continuous with respect to non-decreasing sequences in $\bar{\mathbb{V}}_b$ [18, Theorem 9(i)], we have, for any precise probability tree p , that $\lim_{n \rightarrow +\infty} \bar{E}_p(f_n | s) = \bar{E}_p(f | s)$. Hence, it follows that

$$\sup_{p \sim \mathcal{P}} \lim_{n \rightarrow +\infty} \bar{E}_p(f_n | s) = \sup_{p \sim \mathcal{P}} \bar{E}_p(f | s) = \bar{E}_{\mathcal{P}}(f | s).$$

On the other hand, we also have that

$$\begin{aligned} \sup_{p \sim \mathcal{P}} \lim_{n \rightarrow +\infty} \bar{E}_p(f_n | s) &\leq \lim_{n \rightarrow +\infty} \sup_{p \sim \mathcal{P}} \bar{E}_p(f_n | s) \\ &= \lim_{n \rightarrow +\infty} \bar{E}_{\mathcal{P}}(f_n | s), \end{aligned}$$

where the two last limits exist because $(f_n)_{n \in \mathbb{N}}$ is non-decreasing and \bar{E}_p —and therefore also $\bar{E}_{\mathcal{P}}$ —is monotone; see e.g. [16, Lemma 18]. So we obtain that $\bar{E}_{\mathcal{P}}(f | s) \leq \lim_{n \rightarrow +\infty} \bar{E}_{\mathcal{P}}(f_n | s)$. The converse inequality follows from the fact that $f_n \leq f$ for all $n \in \mathbb{N}$ and the monotonicity of $\bar{E}_{\mathcal{P}}$. ■

Next, we prove that $\bar{E}_{\mathcal{P}}$ is also continuous with respect to non-increasing sequences in \mathbb{F} —that $\bar{E}_{\bar{Q}}$ satisfies this type of continuity was already established in [18, Theorem 9(ii)]. The proof is less straightforward, though, and first requires us to establish the following two topological lemmas concerning probability trees. We will say that a sequence $(p_i)_{i \in \mathbb{N}}$ of precise probability trees converges if there is some limit tree p such that, for each $s \in \mathcal{X}^*$, the mass functions $(p_i(\cdot | s))_{i \in \mathbb{N}}$ converge (pointwise) to the mass function $p(\cdot | s)$.

Lemma 4 Consider any imprecise probability tree \mathcal{P} . Then any sequence $(p_i)_{i \in \mathbb{N}}$ of precise probability trees that are compatible with \mathcal{P} has a converging subsequence whose limit is compatible with \mathcal{P} .

Lemma 5 Consider any sequence $(p_i)_{i \in \mathbb{N}}$ of precise probability trees that converges to some limit tree p . Then, for any $g \in \mathbb{F}$ and any $s \in \mathcal{X}^*$, $\lim_{i \rightarrow +\infty} E_{p_i}(g|s) = E_p(g|s)$.

Combined, the two lemmas above suffice to prove the continuity of $\bar{E}_{\mathcal{P}}$ with respect to non-increasing sequences in \mathbb{F} .

Proposition 6 For any \mathcal{P} and \bar{Q} , any $s \in \mathcal{X}^*$ and any non-increasing sequence $(f_n)_{n \in \mathbb{N}}$ of finitary gambles, we have that

$$\lim_{n \rightarrow +\infty} \bar{E}_{\mathcal{P}}(f_n|s) = \bar{E}_{\mathcal{P}}(f|s), \text{ with } f = \inf_{n \in \mathbb{N}} f_n = \lim_{n \rightarrow +\infty} f_n,$$

and similarly for $\bar{E}_{\bar{Q}}$.

Proof The statement for $\bar{E}_{\bar{Q}}$ follows from [18, Theorem 9(ii)]. To prove the statement for $\bar{E}_{\mathcal{P}}$ first note that, since $(f_n)_{n \in \mathbb{N}}$ is non-increasing and all f_n are gambles, the variable f is bounded above. Moreover, f is \mathcal{F} -measurable because it is a pointwise limit of finitary—and therefore certainly \mathcal{F} -measurable—gambles [13, Theorem II.4.2]. Taking both facts into account, we deduce that, for any $p \sim \mathcal{P}$, the expectation $E_p(f|s)$ exists and hence, because \bar{E}_p is an extension of E_p (see Section 3.2), that $\bar{E}_p(f|s) = E_p(f|s)$. Since this obviously also holds for each f_n —because they are finitary and bounded—the desired statement follows if we manage to show that

$$\lim_{n \rightarrow +\infty} \sup\{E_p(f_n|s) : p \sim \mathcal{P}\} = \sup\{E_p(f|s) : p \sim \mathcal{P}\}.$$

The ‘ \geq ’-inequality follows immediately from the fact that $f_m \geq \inf_{n \in \mathbb{N}} f_n = f$ for all $m \in \mathbb{N}$ and the monotonicity of E_p . It remains to prove the converse inequality.

Fix any $\epsilon > 0$ and let $(p_i)_{i \in \mathbb{N}}$ be a sequence of precise probability trees such that $p_i \sim \mathcal{P}$ and

$$E_{p_i}(f_i|s) + \epsilon \geq \sup\{E_p(f_i|s) : p \sim \mathcal{P}\} \text{ for all } i \in \mathbb{N}.$$

Note that this is indeed possible because, for all $i \in \mathbb{N}$, $\sup\{E_p(f_i|s) : p \sim \mathcal{P}\} \leq \sup f_i$ and, since f_i is a gamble, $\sup f_i \in \mathbb{R}$. Then Lemma 4 guarantees that $(p_i)_{i \in \mathbb{N}}$ has a convergent subsequence $(p_{i(k)})_{k \in \mathbb{N}}$ whose limit p^* is compatible with \mathcal{P} . Since E_{p^*} satisfies continuity with respect to non-increasing sequences [18, Property M9] (the required conditions are obviously satisfied because f_n is finitary and $f_n \leq f_1 \leq \sup f_1 \in \mathbb{R}$ for all $n \in \mathbb{N}$), there is, for any real $a > E_{p^*}(f|s)$, some $n^* \in \mathbb{N}$ such that $a \geq E_{p^*}(f_{n^*}|s)$. Furthermore, since f_{n^*} is finitary, and since $E_{p^*}(f_{n^*}|s) \in \mathbb{R}$ because f_{n^*} is a gamble, Lemma 5 implies that there is some $k^* \in \mathbb{N}$ such

that $E_{p_{i(k)}}(f_{n^*}|s) - \epsilon \leq E_{p^*}(f_{n^*}|s)$ for all $k \geq k^*$. We therefore get that

$$a \geq E_{p_{i(k)}}(f_{n^*}|s) - \epsilon \text{ for all } k \geq k^*. \quad (4)$$

Now consider any $k \geq k^*$ such that $i(k) \geq n^*$, which is possible because $(i(k))_{k \in \mathbb{N}}$ is increasing. Then, since $(f_n)_{n \in \mathbb{N}}$ is non-increasing, and $E_{p_{i(k)}}$ is monotone, Equation (4) implies that $a \geq E_{p_{i(k)}}(f_{i(k)}|s) - \epsilon$. Since the tree $p_{i(k)}$ was chosen in such a way that $E_{p_{i(k)}}(f_{i(k)}|s) + \epsilon \geq \sup\{E_p(f_{i(k)}|s) : p \sim \mathcal{P}\}$, this implies that $a \geq \sup\{E_p(f_{i(k)}|s) : p \sim \mathcal{P}\} - 2\epsilon$. Because this holds for any $k \geq k^*$ such that $i(k) \geq n^*$, we find that

$$\begin{aligned} a &\geq \lim_{k \rightarrow +\infty} \sup\{E_p(f_{i(k)}|s) : p \sim \mathcal{P}\} - 2\epsilon \\ &= \lim_{n \rightarrow +\infty} \sup\{E_p(f_n|s) : p \sim \mathcal{P}\} - 2\epsilon, \end{aligned}$$

where the equality follows from the fact that $(f_n)_{n \in \mathbb{N}}$ is non-increasing and the monotonicity of E_p . Since this holds for any real $a > E_{p^*}(f|s)$, it follows that $E_{p^*}(f|s) \geq \lim_{n \rightarrow +\infty} \sup\{E_p(f_n|s) : p \sim \mathcal{P}\} - 2\epsilon$. Finally, it suffices to recall that $p^* \sim \mathcal{P}$, to see that

$$\sup\{E_p(f|s) : p \sim \mathcal{P}\} \geq \lim_{n \rightarrow +\infty} \sup\{E_p(f_n|s) : p \sim \mathcal{P}\} - 2\epsilon,$$

which, since $\epsilon > 0$ was arbitrary, concludes the proof. ■

It now remains to combine the two types of continuity with the fact that $\bar{E}_{\mathcal{P}}$ and $\bar{E}_{\bar{Q}}$ coincide on $\mathbb{F} \times \mathcal{X}^*$ [18, Proposition 21] to arrive at our first main result.

Theorem 7 Consider any \mathcal{P} and \bar{Q} that agree, any $s \in \mathcal{X}^*$ and any $f \in \bar{\mathbb{V}}$ that is the pointwise limit of a non-decreasing or non-increasing sequence of finitary gambles. Then we have that $\bar{E}_{\mathcal{P}}(f|s) = \bar{E}_{\bar{Q}}(f|s)$.

5. An Equality for \mathcal{F} -Measurable Variables

In order to prove our second main result—that $\bar{E}_{\mathcal{P}}$ coincides with $\bar{E}_{\bar{Q}}$ on bounded below \mathcal{F} -measurable variables—we require the notions of upper and lower semicontinuity.

Let Ω be endowed with the topology generated by the cylinder events $\{\Gamma(s) : s \in \mathcal{X}^*\}$. As we show in [16, Appendix A.2], this topology is metrisable and compact, and coincides with the product topology on $\Omega = \mathcal{X}^{\mathbb{N}}$. For any topological space \mathcal{Y} —and hence for Ω in particular—a function $f : \mathcal{Y} \rightarrow \bar{\mathbb{R}}$ is called *upper semicontinuous* (u.s.c.) if $\{y \in \mathcal{Y} : f(y) < a\}$ is an open subset of \mathcal{Y} for each $a \in \mathbb{R}$; see [8, Section 11.C and 23.F] or [21, Section 3.7.K]. A function $f : \mathcal{Y} \rightarrow \bar{\mathbb{R}}$ is called *lower semicontinuous* (l.s.c.) if $-f$ is u.s.c. and it is called *continuous* if it is both u.s.c. and l.s.c. In general, semicontinuous functions can always be written as pointwise limits of monotone sequences of continuous real-valued functions (see e.g. [8, Theorem 23.19]). In our case, though, where $\mathcal{Y} = \Omega$, a stronger property holds.

Lemma 8 Any $f \in \overline{\mathbb{V}}$ is u.s.c. (l.s.c.) if and only if it is the pointwise limit of a non-increasing (resp. non-decreasing) sequence $(f_n)_{n \in \mathbb{N}}$ of extended real variables, each of which is finitary and bounded below (resp. bounded above). Moreover, f is both u.s.c. (l.s.c.) and bounded above (resp. bounded below) if and only if it is the pointwise limit of a non-increasing (resp. non-decreasing) sequence $(f_n)_{n \in \mathbb{N}}$ of finitary gambles.

Lemma 8 leads us to two important intermediate results, the first of which being that $\overline{E}_{\mathcal{P}}$ and $\overline{E}_{\overline{\mathcal{Q}}}$ coincide on the domain of all u.s.c. variables that are bounded above and all l.s.c. variables that are bounded below. The result can simply be seen as a restatement of Theorem 7.

Corollary 9 For any \mathcal{P} and $\overline{\mathcal{Q}}$ that agree, any $s \in \mathcal{X}^*$ and any variable $f \in \overline{\mathbb{V}}$ that is u.s.c. and bounded above, or l.s.c. and bounded below, we have that $\overline{E}_{\mathcal{P}}(f|s) = \overline{E}_{\overline{\mathcal{Q}}}(f|s)$.

On the other hand, Lemma 8 also implies that continuity with respect to non-increasing sequences of (bounded above) u.s.c. variables is actually not stronger than continuity with respect to non-increasing sequences of finitary gambles; see [16, Lemma 17]. Since both $\overline{E}_{\mathcal{P}}$ and $\overline{E}_{\overline{\mathcal{Q}}}$ satisfy the latter type of continuity, we immediately obtain the following result.

Proposition 10 Consider any \mathcal{P} and $\overline{\mathcal{Q}}$, any $s \in \mathcal{X}^*$ and any non-increasing sequence $(f_n)_{n \in \mathbb{N}}$ of u.s.c. variables that are bounded above. Then we have that $\lim_{n \rightarrow +\infty} \overline{E}_{\mathcal{P}}(f_n|s) = \overline{E}_{\mathcal{P}}(f|s)$ for $f = \lim_{n \rightarrow +\infty} f_n$, and similarly for $\overline{E}_{\overline{\mathcal{Q}}}$.

Proof This follows from [16, Lemma 17], Proposition 6 and the fact that $\overline{E}_{\mathcal{P}}$ and $\overline{E}_{\overline{\mathcal{Q}}}$ are clearly both monotonous (see [16, Lemma 18]). ■

Note that, conversely, $\overline{E}_{\mathcal{P}}$ and $\overline{E}_{\overline{\mathcal{Q}}}$ are also continuous with respect to non-decreasing sequences of l.s.c. variables that are bounded below, simply because, due to Proposition 3, they satisfy continuity with respect to any non-decreasing (bounded below) sequence.

As a final step towards establishing our desired result, we will use what is called Choquet's capacitability theorem. This theorem can be found in many different textbooks, but we will make use of the specific version of Dellacherie [6]. We do this because Dellacherie's notion of a capacity can directly be applied to an extended real-valued functional—such as $\overline{E}_{\mathcal{P}}$ and $\overline{E}_{\overline{\mathcal{Q}}}$ —whereas most other sources restrict capacities to take the form of set-functions. Let us start by introducing some key concepts and terminology regarding capacitability and analytic functions.

Let $\overline{\mathbb{V}}_{\geq 0}$ be the set of all variables taking values in $\overline{\mathbb{R}}_{\geq 0}$ and $\mathbb{V}_{\geq 0}^u$ the set of all (possibly unbounded) variables taking values in $\mathbb{R}_{\geq 0}$. A functional $F: \overline{\mathbb{V}}_{\geq 0} \rightarrow \overline{\mathbb{R}}_{\geq 0}$ is called

a Ω -capacity if it satisfies the following three properties [6, Section II.1.1]:

- CA1. $f \leq g \Rightarrow F(f) \leq F(g)$ for all $f, g \in \overline{\mathbb{V}}_{\geq 0}$;
- CA2. $\lim_{n \rightarrow +\infty} F(f_n) = F(\lim_{n \rightarrow +\infty} f_n)$ for any non-decreasing sequence $(f_n)_{n \in \mathbb{N}}$ in $\overline{\mathbb{V}}_{\geq 0}$;
- CA3. $\lim_{n \rightarrow +\infty} F(f_n) = F(\lim_{n \rightarrow +\infty} f_n)$ for any non-increasing sequence $(f_n)_{n \in \mathbb{N}}$ of u.s.c. variables in $\mathbb{V}_{\geq 0}^u$.

Recall from the beginning of this section that Ω is compact and metrisable, which is in line with Dellacherie's assumption about the set 'E' in [6, Section II.1.1]; see [6, Introduction, Paragraph 2]. Furthermore, observe that CA3 only applies to sequences in $\mathbb{V}_{\geq 0}^u$ instead of sequences in $\overline{\mathbb{V}}_{\geq 0}$; this too corresponds to the definition given in [6, Section II.1.1] because Dellacherie always considers u.s.c. functions to be real-valued [6, Introduction, Paragraph 2]. In fact, one could restate CA3 so as to only apply to sequences that are uniformly bounded above; this follows immediately from the non-increasing character and the following lemma.

Lemma 11 Any u.s.c. variable $f \in \mathbb{V}_{\geq 0}^u$ is bounded above.

For any Ω -capacity F , we say that a variable $f \in \overline{\mathbb{V}}_{\geq 0}$ is F -capacitable if

$$F(f) = \sup\{F(g); g \in \mathbb{V}_{\geq 0}^u, g \text{ is u.s.c. and } f \geq g\}. \quad (5)$$

A variable $f \in \overline{\mathbb{V}}_{\geq 0}$ is called *universally capacitable* if it is F -capacitable for all Ω -capacities F . Now, Choquet's capacitability theorem [6, Theorem II.2.5] states that any analytic variable is universally capacitable. The definition of an analytic variable can be found in [6, 8]; we do not explicitly give it here, because it is a rather abstract concept that, in practice, can often be replaced by the simpler and better-known notion of a Borel-measurable variable. Indeed, according to [6, Section I.2.6], each Borel-measurable variable in $\overline{\mathbb{V}}_{\geq 0}$ is analytic. Moreover, by [16, Corollary 16], the Borel σ -algebra on Ω coincides with the σ -algebra \mathcal{F} generated by all cylinder events, so the notions of Borel-measurability and \mathcal{F} -measurability are equivalent. Combined with [6, Theorem II.2.5], this allows us to state the following weaker version of Choquet's capacitability theorem:

Theorem 12 (Choquet's capacitability light) Any \mathcal{F} -measurable variable $f \in \overline{\mathbb{V}}_{\geq 0}$ is universally capacitable.

As an almost immediate consequence of Proposition 3, Proposition 6 and Lemma 11, it can be shown that, for any $s \in \mathcal{X}^*$, the restrictions of both $\overline{E}_{\mathcal{P}}(\cdot|s)$ and $\overline{E}_{\overline{\mathcal{Q}}}(\cdot|s)$ to $\overline{\mathbb{V}}_{\geq 0}$ are Ω -capacities [16, Appendix A.2]. Therefore, and because these upper expectations coincide on the u.s.c. variables in $\mathbb{V}_{\geq 0}^u$ —due to Corollary 9 and Lemma 11

above—the desired equality for \mathcal{F} -measurable variables in $\bar{\mathbb{V}}_{\geq 0}$ follows from Equation (5) and Theorem 12. We can moreover replace $\bar{\mathbb{V}}_{\geq 0}$ by $\bar{\mathbb{V}}_b$, simply because $\bar{E}_{\mathcal{P}}$ and $\bar{E}_{\bar{Q}}$ are linear with respect to adding constants [16, Lemma 18]. This leads to our second main result.

Theorem 13 *For any \mathcal{P} and \bar{Q} that agree, any $s \in \mathcal{X}^*$ and any \mathcal{F} -measurable variable $f \in \bar{\mathbb{V}}_b$ that is bounded below, we have that $\bar{E}_{\mathcal{P}}(f|s) = \bar{E}_{\bar{Q}}(f|s)$.*

6. Relation with Shafer and Vovk’s Work

Before we conclude this paper, it seems appropriate to say a few words about how our work here compares to that of Shafer and Vovk. As readers that are familiar with their work may have noticed, the idea to use Choquet’s capacitability theorem to extend the domain of the equality to \mathcal{F} -measurable (or analytic) variables already appears in [12, Chapter 9]. Another part that strongly builds on ideas from [12, Chapter 9] is the proof of Proposition 6; some key steps there were inspired by the proof of [12, Lemma 9.10]. So it is fair to say that [12, Chapter 9] served as an important inspiration for our work. In fact, to the untrained eye, it might perhaps even seem as if our results do not differ much from those in [12, Chapter 9]; but take a closer look.

First of all—and most importantly—the setting in which we define game-theoretic upper expectations differs considerably from theirs. More specifically, they consider supermartingales under the prequential principle, which says that Forecaster’s moves—the specification of the local models \bar{Q}_s (or \bar{Q}_s^\dagger)—are not necessarily known beforehand for each situation $s \in \mathcal{X}^*$, but instead are allowed to also depend on previous moves by Skeptic; see [12, Theorem 7.5] for more details. While this assumption allows them to remain more general—though, in many practical cases, it does not make much of a difference—the benefit that we gain from dropping it is remarkable; it allows us to replace [12, Lemma 9.10] and [12, Theorem 9.7], which require strong topological conditions on the parametrisation of the local models, with respectively Theorem 7 and Theorem 13, which are similar, but do not need any topological conditions at all.

A second notable difference is that our results involve a larger domain; Theorem 7, or equivalently, Corollary 9, applies to both u.s.c. variables that are bounded above and l.s.c. variables that are bounded below, whereas [12, Lemma 9.10] only applies to bounded u.s.c. variables; Theorem 13 applies to bounded below (\mathcal{F} -measurable) variables, whereas [12, Theorem 9.7] only applies to bounded (analytic) variables.⁷ Our results also allow conditioning on situations; theirs only apply to unconditional upper expectations. The fact that this extension in

domain is relevant in practice becomes clear when we also take a look at lower expectations. Indeed, in (more) practical situations, we are usually not only interested in the upper expectation of a variable, but also, and simultaneously, in its lower expectation [9, 10]. Our results can be easily extended to this two-sided setting, by combining the conjugacy relation between global upper and lower expectations with our two main results.

Corollary 14 *Consider any \mathcal{P} and \bar{Q} that agree, any $s \in \mathcal{X}^*$ and any $f \in \bar{\mathbb{V}}$ that is (a) the pointwise limit of a monotone sequence of finitary gambles or (b) an \mathcal{F} -measurable gamble. Then we have that*

$$\bar{E}_{\mathcal{P}}(f|s) = \bar{E}_{\bar{Q}}(f|s) \text{ and } \underline{E}_{\mathcal{P}}(f|s) = \underline{E}_{\bar{Q}}(f|s).$$

Note that many practically relevant inferences—e.g. hitting times [9]—fall under category (a) but not under category (b), simply because they are not bounded. Yet, it is exactly this class of variables that is missing in Shafer and Vovk’s main result [12, Theorem 9.7].

Finally, recall that our results relate $\bar{E}_{\mathcal{P}}$ to $\bar{E}_{\bar{Q}}$, where the latter represents, apart from the game-theoretic upper expectation $\bar{E}_{G, \bar{Q}}$, also the axiomatic upper expectation $\bar{E}_{A, \bar{Q}}$. Shafer and Vovk, on the other hand, only relate $\bar{E}_{\mathcal{P}}$ to the game-theoretic upper expectation $\bar{E}_{G, \bar{Q}}$.

7. Conclusion

Our main results, Theorem 7 and Theorem 13, show that measure-theoretic, game-theoretic and axiomatic upper expectations are equal on a large domain of variables; it contains all variables that are the limit of a monotone sequence of finitary gambles, and all variables that are bounded below and \mathcal{F} -measurable. It remains to be seen whether we can extend this equivalence even further, to all variables; so far, we have yet to find a counterexample showing that this is not possible.

We would also like to investigate the relation between our models and the Daniell-Stone type of (global) upper expectations described in [7]. Comparing Theorem 13 and [7, Theorem 3.10], and taking into account their use of Choquet’s capacitability theorem, it seems that a close connection must exist, at least for bounded measurable variables. A more thorough study is required though before we can make accurate statements.

Acknowledgments

The research of Natan T’Joens was supported and funded by the Special Research Fund (BOF) of Ghent University (reference number: 356). The research of Jasper De Bock was partially funded by project number 3GO28919 of the FWO (Research Foundation - Flanders). We thank the reviewers for their thorough reading of our manuscript.

⁷ Recall that we could just as well have stated Theorem 13 for analytic variables instead of \mathcal{F} -measurable variables.

References

- [1] Thomas Augustin, Frank P.A. Coolen, Gert de Cooman, and Matthias C.M. Troffaes. *Introduction to Imprecise Probabilities*. Wiley, Chichester, 2014.
- [2] Patrick Billingsley. *Probability and Measure*. Wiley Series in Probability and Statistics. Wiley, third edition, 1995.
- [3] Jasper De Bock and Natan T'Joens. Average behaviour of imprecise Markov chains: A single pointwise ergodic theorem for six different models. Accepted for publication in the Proceedings of ISIPTA 2021.
- [4] Gert de Cooman, Jasper De Bock, and Stavros Lopotatzidis. Imprecise stochastic processes in discrete time: global models, imprecise Markov chains, and ergodic theorems. *International Journal of Approximate Reasoning*, 76:18–46, 2016.
- [5] Bruno De Finetti. *Theory of Probability*. John Wiley & Sons, New York, 1974.
- [6] Claude Dellacherie. *Ensembles Analytiques, Capacités, Mesures de Hausdorff*, volume 295 of *Lecture Notes in Mathematics*. Springer-Verlag, Berlin, 1972.
- [7] Robert Denk, Michael Kupper, and Max Nendel. Kolmogorov-type and general extension results for nonlinear expectations. *Banach Journal of Mathematical Analysis*, 12(3):515–540, 2018.
- [8] Alexander Kechris. *Classical Descriptive Set Theory*, volume 156 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1995.
- [9] Thomas Krak, Natan T'Joens, and Jasper De Bock. Hitting times and probabilities for imprecise Markov chains. In *International Symposium on Imprecise Probabilities : Theories and Applications, ISIPTA 2019, Proceedings*, volume 103, pages 265–275. PMLR, 2019.
- [10] Stavros Lopotatzidis. *Robust Modelling and Optimisation in Stochastic Processes using Imprecise Probabilities, with an Application to Queueing Theory*. PhD thesis, Ghent University, 2017.
- [11] Glenn Shafer and Vladimir Vovk. *Probability and Finance: It's Only a Game!* Wiley, New York, 2001.
- [12] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, Hoboken, 2019.
- [13] Albert N. Shiryaev. *Probability-1*. Graduate Texts in Mathematics. Springer, New York, third edition, 2016.
- [14] Albert N. Shiryaev. *Probability-2*. Graduate Texts in Mathematics. Springer, New York, third edition, 2019.
- [15] Natan T'Joens and Jasper De Bock. Average behaviour in discrete-time imprecise Markov chains: A study of weak ergodicity. *International Journal of Approximate Reasoning*, 132:181–205, 2021.
- [16] Natan T'Joens and Jasper De Bock. Global upper expectations for discrete-time stochastic processes: In practice, they are all the same! Extended version, see ArXiv:2102.13075.
- [17] Natan T'Joens, Jasper De Bock, and Gert de Cooman. Game-theoretic upper expectations for discrete-time finite-state uncertain processes. Submitted for publication; see arXiv:2008.03133 for a preliminary online version, 2020.
- [18] Natan T'Joens, Jasper De Bock, and Gert de Cooman. A particular upper expectation as global belief model for discrete-time finite-state uncertain processes. *International Journal of Approximate Reasoning*, 131:30–55, 2021.
- [19] Matthias C.M. Troffaes and Gert de Cooman. *Lower Previsions*. Wiley, Chichester, 2014.
- [20] Peter Walley. *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London, 1991.
- [21] Stephen Willard. *General Topology*. Addison-Wesley series in mathematics. Dover Publications, 2004.
- [22] David Williams. *Probability with Martingales*. Cambridge mathematical textbooks. Cambridge University Press, 1991.