

Hitting Times and Probabilities for Imprecise Markov Chains

Thomas Krak
Natan T’Joens
Jasper De Bock

ELIS – FLip, Ghent University, Belgium

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

Abstract

We consider the problem of characterising expected hitting times and hitting probabilities for imprecise Markov chains. To this end, we consider three distinct ways in which imprecise Markov chains have been defined in the literature: as sets of homogeneous Markov chains, as sets of more general stochastic processes, and as game-theoretic probability models. Our first contribution is that all these different types of imprecise Markov chains have the same lower and upper expected hitting times, and similarly the hitting probabilities are the same for these three types. Moreover, we provide a characterisation of these quantities that directly generalises a similar characterisation for precise, homogeneous Markov chains.

Keywords: imprecise Markov chain, hitting time, hitting probability, lower and upper expectations

1. Introduction

Markov chains are mathematical models that probabilistically describe the uncertain behaviour of a dynamical system [19]. We here consider Markov chains that can only be in a finite number of states, and that can only change state at discrete steps in time. An important class of inferences for Markov chains are the so called *expected hitting times* and *hitting probabilities* for some subset A of the set of all states \mathcal{X} that the system can be in. Informally, their aim is to answer the questions “How long will it take until the system enters a state in A ?” and “What is the probability of ever visiting a state in A ?”, respectively. Under some regularity conditions, closed-form solutions to these questions are available in the literature [19, 9].

A generalisation of Markov chains that also incorporates (higher order) uncertainty about one’s knowledge of the model description itself are *imprecise Markov chains* [12, 2, 10, 22, 3, 4, 5, 11, 15]. Their theoretical foundations are based on the theory of imprecise probabilities [29, 1], and they allow one to incorporate uncertainties about the numerical model parameters as well as about structural assumptions, like history independence—the canonical Markov property—and time homogeneity.

However, the generalisation of Markov chains to their imprecise counterpart is not unambiguous [11]. There are

various ways in which this might be done, and they can lead to different conclusions for particular inferences of interest.

On the one hand we have what might be called the “sensitivity analysis” interpretation of an imprecise Markov chain. Here, one’s model essentially constitutes an entire *set* of stochastic processes that are all compatible with one’s assessments about the system’s uncertain behaviour. But there are multiple versions of this interpretation, depending on which models one chooses to include in this set; for instance, do we only include all (time-homogeneous) Markov chains that are compatible with our assessments [12, 2], or do we also include more general stochastic processes [5, 11]? Each choice has its own merits, depending on the particular situation. Regardless of the choice that one makes here, inferences for this “sensitivity analysis” interpretation always consist in computing tight lower and upper bounds on inferences for all the models that are included in the chosen set [11].

An entirely different formalisation of imprecise Markov chains is based on the game-theoretic probability framework that was popularised by Shafer and Vovk [20]. These models are not necessarily given an interpretation in terms of compatible “precise” models; rather, this theory of stochastic processes is based on rational betting behaviour in repeated games with uncertain outcomes, and naturally leads to imprecise probabilistic models [3, 15, 6]. The correspondence between this framework and the “sensitivity analysis” interpretation of imprecise Markov chains was first explored in [3, 7].

In this present work, we consider the inference problems of computing lower and upper expected hitting times and hitting probabilities for an imprecise Markov chain—regardless of the specific interpretation that one chooses for these models. In fact, the first of our main results is that these inferences are the same for all of the different types of imprecise Markov chains discussed above. Our second main result is an exact generalisation to the imprecise setting, of a well-known characterisation of these inferences for precise, time-homogeneous Markov chains.

To the best of our knowledge, this problem has never been considered in the literature at this level of generality. The most closely related work that we are aware of is that of Lopatzidis *et al.* [16, 17], who prove similar proper-

ties for imprecise Markov chains that have the structure of birth-death chains. Moreover, De Cooman *et al.* [6] previously derived a non-linear system describing expected hitting times, that is similar to our characterisation stated in Corollary 13.

Some of the lengthier proofs as well as proofs of technical lemmas had to be omitted from this work because of the page limit constraint. They are available in the appendix of an extended version of this work [14].

2. Preliminaries

Throughout, \mathbb{N} denotes the natural numbers, and we let $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. \mathbb{R} denotes the real numbers, we define $\overline{\mathbb{R}} := \mathbb{R} \cup \{+\infty\}$, that is, the reals that are closed above, and we let $\underline{\mathbb{R}} := \mathbb{R} \cup \{-\infty, +\infty\}$. The sets $\overline{\mathbb{R}}$ and $\underline{\mathbb{R}}$ are endowed with the usual (order) topology, and we adopt the convention that $0 \cdot +\infty = 0 = 0 \cdot -\infty$.

We use \mathcal{X} to denote the finite non-empty set of states that the Markov chain can be in. Without loss of generality, it can be identified with the set $\mathcal{X} = \{1, \dots, k\}$ for some $k \in \mathbb{N}$. We use $\mathcal{L}(\mathcal{X})$ to denote the set of real-valued functions on \mathcal{X} . The set $\overline{\mathcal{L}}(\mathcal{X})$ contains all functions on \mathcal{X} that take values in $\overline{\mathbb{R}}$. Since \mathcal{X} is finite, any f in $\mathcal{L}(\mathcal{X})$ or $\overline{\mathcal{L}}(\mathcal{X})$ can be identified with a vector in \mathbb{R}^k or $(\overline{\mathbb{R}})^k$, respectively. The set $\mathcal{L}(\mathcal{X})$ is endowed with the supremum norm, i.e. $\|f\| := \sup_{x \in \mathcal{X}} |f(x)|$, and the corresponding norm topology. $\overline{\mathcal{L}}(\mathcal{X})$ receives the topology of pointwise convergence.

For any $A \subset \mathcal{X}$, we consider the indicator \mathbb{I}_A of A , defined as $\mathbb{I}_A(x) := 1$ if $x \in A$ and $\mathbb{I}_A(x) := 0$ otherwise. Constant functions on \mathcal{X} are simply denoted by their constant values. Finally, point-wise multiplication of two functions (i.e. vectors) is denoted by $f \cdot g$; for example, the term $\mathbb{I}_{A^c} \cdot Th_A^P$ in Equation (13) further on denotes the pointwise multiplication of the functions \mathbb{I}_{A^c} and Th_A^P .

2.1. “Measure-Theoretic” Imprecise Markov Chains

In order to discuss the various types of imprecise Markov chains that arise from the “sensitivity analysis” interpretation of imprecise probabilities, we need a formalisation of general (non-Markovian) stochastic processes. We briefly give the measure theoretic account of this formalisation.

In this framework, the unknown—that is, uncertain—realisation of the stochastic process is a *path*, which is a function $\omega : \mathbb{N}_0 \rightarrow \mathcal{X}$. We collect all paths in the set Ω . This set Ω is endowed with a σ -algebra \mathcal{F} ¹ and augmented to a probability space (Ω, \mathcal{F}, P) with a probability measure P . A stochastic process is then a family $\{X_n\}_{n \in \mathbb{N}_0}$ of random variables on this probability space, such that

1. Specifically, we assume that \mathcal{F} is the σ -algebra generated by the cylinder sets; this guarantees that all functions that we consider are measurable.

$X_n : \omega \mapsto \omega(n)$ for all $n \in \mathbb{N}_0$. However, for ease of notation and terminology, we will often refer to the measure P as the stochastic process; different processes then correspond to different measures on the same measurable space (Ω, \mathcal{F}) .

A *Markov chain* is a stochastic process that satisfies the *Markov condition*, which is a conditional independence relation between the random states that the process obtains. In particular, a process P is said to be a Markov chain if

$$P(X_{n+1} = x_{n+1} | X_{0:n} = x_{0:n}) = P(X_{n+1} = x_{n+1} | X_n = x_n),$$

for all $x_0, \dots, x_n, x_{n+1} \in \mathcal{X}$ and all $n \in \mathbb{N}_0$, where we let $X_{0:n} := (X_0, \dots, X_n)$ and similarly for $x_{0:n}$. A Markov chain is called *homogeneous* if, for all $x, y \in \mathcal{X}$ and all $n \in \mathbb{N}_0$,

$$P(X_{n+1} = y | X_n = x) = P(X_1 = y | X_0 = x). \quad (1)$$

Any homogeneous Markov chain P is uniquely characterised—up to its initial distribution $P(X_0)$ —by a *transition matrix*. A transition matrix T is simply an $|\mathcal{X}| \times |\mathcal{X}|$ matrix that is row-stochastic, meaning that for all $x \in \mathcal{X}$, $\sum_{y \in \mathcal{X}} T(x, y) = 1$ and $T(x, y) \geq 0$ for all $y \in \mathcal{X}$. Such a transition matrix identifies a homogeneous Markov chain P (up to its initial distribution) that satisfies

$$P(X_{n+1} = y | X_n = x) = T(x, y) \text{ for all } x, y \in \mathcal{X}, n \in \mathbb{N}_0. \quad (2)$$

Moreover, a transition matrix T can also be interpreted as a linear operator that maps $\mathcal{L}(\mathcal{X})$ into $\mathcal{L}(\mathcal{X})$, because we have identified $\mathcal{L}(\mathcal{X})$ with $\mathbb{R}^{|\mathcal{X}|}$. For any $f \in \mathcal{L}(\mathcal{X})$ and $x \in \mathcal{X}$ it then holds that

$$\begin{aligned} \mathbb{E}_P[f(X_{n+1}) | X_n = x] &= \sum_{y \in \mathcal{X}} P(X_{n+1} = y | X_n = x) f(y) \\ &= \sum_{y \in \mathcal{X}} T(x, y) f(y) = [Tf](x), \end{aligned}$$

so we see that T encodes the conditional expectation operator for 1 time step corresponding to a process P that satisfies (2). Moreover, T can be uniquely extended to an operator on $\overline{\mathcal{L}}(\mathcal{X})$, due to the convention that $0 \cdot +\infty = 0$.

We now move on to the characterisation of *imprecise* Markov chains. In all cases that we consider here, these are characterised by a *set* \mathcal{T} of transition matrices. In the remainder, we will assume that \mathcal{T} is non-empty, closed, convex, and that it has separately specified rows. This last condition means that if, for all $x \in \mathcal{X}$, we select any element $T_x \in \mathcal{T}$, there must be some $T \in \mathcal{T}$ such that $T(x, \cdot) = T_x(x, \cdot)$ for all $x \in \mathcal{X}$; see e.g. [11, Definition 11.6] for further discussion.

An *imprecise* Markov chain is now a set of stochastic processes that are in a specific sense “compatible” with the transition matrices in \mathcal{T} . However, there are various ways how we might construct such a set, which all lead to different types of imprecise Markov chains.

Arguably the simplest imprecise Markov chain is the set $\mathcal{P}_{\mathcal{T}}^H$, which is the set of all homogeneous Markov chains whose characterising transition matrix T is included in \mathcal{T} [12, 2]. Its corresponding lower and upper expectation operators are defined respectively as

$$\mathbb{E}_{\mathcal{T}}^H[\cdot|\cdot] := \inf_{P \in \mathcal{P}_{\mathcal{T}}^H} \mathbb{E}_P[\cdot|\cdot] \quad \text{and} \quad \overline{\mathbb{E}}_{\mathcal{T}}^H[\cdot|\cdot] := \sup_{P \in \mathcal{P}_{\mathcal{T}}^H} \mathbb{E}_P[\cdot|\cdot],$$

where, for both $\mathbb{E}_{\mathcal{T}}^H[\cdot|\cdot]$ and $\overline{\mathbb{E}}_{\mathcal{T}}^H[\cdot|\cdot]$, the first argument takes functions of the form $f : \Omega \rightarrow \mathbb{R}$, and the second is a conditioning event $X_{0:n} = x_{0:n}$ with $n \in \mathbb{N}_0$.²

Perhaps a less obvious choice is the imprecise Markov chain $\mathcal{P}_{\mathcal{T}}^I$, which is the set of all (potentially non-Markov, non-homogeneous) stochastic processes for which for all $n \in \mathbb{N}_0$ and all $x_0, \dots, x_n \in \mathcal{X}$:

$$\exists T \in \mathcal{T} : \forall y \in \mathcal{X} : P(X_{n+1} = y | X_{0:n} = x_{0:n}) = T(x_n, y). \quad (3)$$

The associated lower and upper expectation operators are

$$\mathbb{E}_{\mathcal{T}}^I[\cdot|\cdot] := \inf_{P \in \mathcal{P}_{\mathcal{T}}^I} \mathbb{E}_P[\cdot|\cdot] \quad \text{and} \quad \overline{\mathbb{E}}_{\mathcal{T}}^I[\cdot|\cdot] := \sup_{P \in \mathcal{P}_{\mathcal{T}}^I} \mathbb{E}_P[\cdot|\cdot],$$

whose domain we take to be the same as that of $\mathbb{E}_{\mathcal{T}}^H[\cdot|\cdot]$ and $\overline{\mathbb{E}}_{\mathcal{T}}^H[\cdot|\cdot]$. This type of Markov chain is often considered in the literature [5, 11], and is called an imprecise Markov chain under *epistemic irrelevance*.

Next, it will be useful to consider the dual representation(s) of the set \mathcal{T} , given by the *lower* (resp. *upper*) *transition operator* \underline{T} (resp. \overline{T}). For either domain $\mathcal{L}(\mathcal{X})$ or $\overline{\mathcal{L}}(\mathcal{X})$, these are (non-linear) operators that map these function spaces into themselves; they are respectively defined for any element $f \in \overline{\mathcal{L}}(\mathcal{X})$ and any $x \in \mathcal{X}$ as

$$[\underline{T}f](x) := \inf_{T \in \mathcal{T}} [Tf](x) \quad \text{and} \quad [\overline{T}f](x) := \sup_{T \in \mathcal{T}} [Tf](x). \quad (4)$$

Under the stated conditions on \mathcal{T} , these operators satisfy the following useful properties:

Lemma 1 *For all $f \in \overline{\mathcal{L}}(\mathcal{X})$, there exist $T, S \in \mathcal{T}$ such that*

$$Tf = \underline{T}f \quad \text{and} \quad Sf = \overline{T}f. \quad (5)$$

Moreover, \underline{T} and \overline{T} are continuous operators on $\mathcal{L}(\mathcal{X})$, and are continuous on $\overline{\mathcal{L}}(\mathcal{X})$ with respect to non-decreasing sequences.

The usefulness of these operators stems from the fact that—similar to transition matrices for homogeneous Markov chains—they encode the (1 time step) lower and upper expectation operators for $\mathcal{P}_{\mathcal{T}}^H$ and $\mathcal{P}_{\mathcal{T}}^I$. That is,

$$[\overline{T}f](x_n) = \overline{\mathbb{E}}_{\mathcal{T}}^H[f(X_{n+1}) | X_{0:n} = x_{0:n}]$$

2. We omit a technical discussion about the required measurability and integrability properties of such f , and use this definition provided that $\mathbb{E}_P[f | X_{0:n}] := \int_{\Omega} f(\omega) dP(\omega | X_{0:n})$ is well-defined; see e.g. [24] for when this is the case. For our present purposes, it suffices to know that the functions that will be of interest in this work are all non-negative and measurable, making their expectation well-defined.

$$= \overline{\mathbb{E}}_{\mathcal{T}}^I[f(X_{n+1}) | X_{0:n} = x_{0:n}], \quad (6)$$

for all $f \in \mathcal{L}(\mathcal{X})$, all $n \in \mathbb{N}_0$ and all $x_0, \dots, x_n \in \mathcal{X}$; and similarly for the lower expectations and \underline{T} . Observe that the left hand side in this expression does not depend on the states $x_{0:n-1}$, which can be interpreted as saying that the (lower and upper) expectations of $\mathcal{P}_{\mathcal{T}}^H$ and $\mathcal{P}_{\mathcal{T}}^I$ satisfy an *imprecise Markov property*. This explains in particular why we call $\mathcal{P}_{\mathcal{T}}^I$ an “imprecise Markov chain”, while it consists of processes which in general do not themselves satisfy the Markov property. Moreover, despite the above, it is worth noting that equality of $\overline{\mathbb{E}}_{\mathcal{T}}^H[f(X_{n+m}) | X_{0:n}]$ and $\overline{\mathbb{E}}_{\mathcal{T}}^I[f(X_{n+m}) | X_{0:n}]$ does not in general hold when $m > 1$; see e.g. [15, Example 10].

Finally, we note that in our definition and notation of the imprecise Markov chains above, we paid no further attention to the initial models $P(X_0)$ of their elements P . If this were to be of interest, we could specify an imprecise initial model. That is, we could consider a non-empty set \mathcal{M} of probability mass functions on \mathcal{X} , and then include P in $\mathcal{P}_{\mathcal{T}}^H$ or $\mathcal{P}_{\mathcal{T}}^I$ if and only if, in addition to its compatibility with \mathcal{T} as discussed above, it holds that $P(X_0) \in \mathcal{M}$.

However, we purposely restricted the domains of the corresponding lower and upper expectation operators to conditioning events of the form $X_{0:n} = x_{0:n}$, as this will suffice for all our results. Therefore, as one can easily see, these lower and upper expectations are invariant under any particular choice of such an initial model \mathcal{M} , which is why we have omitted any further reference to it for ease of notation and clarity of exposition.

2.2. Hitting Times and Probabilities

We next introduce the two inferences that are of interest in this work. The first of these is the expected *hitting time* of a set of states $A \subset \mathcal{X}$. The hitting time $H_A : \Omega \rightarrow \mathbb{N}_0 \cup \{+\infty\}$ for this set A is a function defined for all $\omega \in \Omega$ as $H_A(\omega) := \inf\{t \in \mathbb{N}_0 : \omega(t) \in A\}$. The vector of expected hitting times $h_A^P \in \overline{\mathcal{L}}(\mathcal{X})$ for a given stochastic process P , conditional on the starting state X_0 , is defined for all $x \in \mathcal{X}$ as

$$h_A^P(x) := \mathbb{E}_P[H_A | X_0 = x] := \int_{\Omega} H_A(\omega) dP(\omega | \omega(0) = x).$$

Thus, $h_A^P(x)$ is the expected number of steps before the process P reaches any element of A , starting from x .

For the imprecise Markov chains $\mathcal{P}_{\mathcal{T}}^H$ and $\mathcal{P}_{\mathcal{T}}^I$, the *lower* expected hitting times are defined respectively as

$$\mathbb{E}_{\mathcal{T}}^H[H_A | X_0 = x] := \inf_{P \in \mathcal{P}_{\mathcal{T}}^H} h_A^P(x), \quad (7)$$

and

$$\overline{\mathbb{E}}_{\mathcal{T}}^I[H_A | X_0 = x] := \inf_{P \in \mathcal{P}_{\mathcal{T}}^I} h_A^P(x), \quad (8)$$

with the corresponding *upper* expected hitting times defined analogously with suprema.

The second inference that we are after is the vector of conditional *hitting probabilities* $p_A^P \in \mathcal{L}(\mathcal{X})$: the probabilities that the process will eventually visit an element of A . An explicit way of encoding this inference problem uses the function $G_A : \Omega \rightarrow \{0, 1\}$, defined for all $\omega \in \Omega$ as

$$G_A(\omega) := \sup\{\mathbb{I}_A(\omega(t)) : t \in \mathbb{N}_0\}. \quad (9)$$

Thus, G_A takes the value one on a path ω if this path at some point in time passes through any of the states in A ; otherwise it takes the value zero. Therefore, clearly, for any stochastic process P , the hitting probability is given by

$$p_A^P(x) := \mathbb{E}_P[G_A | X_0 = x] := \int_{\Omega} G_A(\omega) dP(\omega | \omega(0) = x).$$

Correspondingly, the lower hitting probability for the imprecise Markov chain $\mathcal{P}_{\mathcal{F}}^H$ is given by

$$\mathbb{E}_{\mathcal{F}}^H[G_A | X_0 = x] := \inf_{P \in \mathcal{P}_{\mathcal{F}}^H} p_A^P(x), \quad (10)$$

and similarly for the upper probability and for $\mathcal{P}_{\mathcal{F}}^I$.

2.3. “Game-Theoretic” Imprecise Markov Chains

In Section 2.1 we introduced (imprecise) Markov chains using their “measure-theoretic” formalisation. An entirely different mathematical framework for describing stochastic processes—and imprecise Markov chains in particular—is the “game-theoretic” framework popularised by Shafer and Vovk [20]. For an in-depth treatise on this formalism, we refer the interested reader to [20, 28, 26]. Explicit discussions about the connection to the measure-theoretic framework can be found in references [3, 7, 15].

For our present purposes, we restrict attention to a discussion of some essential properties of the corresponding (lower or upper) expectation operators. To this end, it suffices to think of such a game-theoretic model as simply a different characterisation of the uncertain behaviour of the dynamical system of interest. And, although this characterisation is different from the measure-theoretic one, it still leads to the same inferences in a large number of cases; in fact, it is one of the aims of this present paper to show that the expected hitting times and hitting probabilities are the same for these two different characterisations.

The operators that we will consider in this section are functionals on functions on paths $\omega \in \Omega$. We will need a slight notational digression to introduce these domains. We let $\mathcal{L}(\Omega)$ be the set of all functions on Ω that take values in \mathbb{R} . The domains $\overline{\mathcal{L}}(\Omega)$ and $\underline{\mathcal{L}}(\Omega)$ contain the functions taking values in $\overline{\mathbb{R}}$ and $\underline{\mathbb{R}}$, respectively. We also need the concept of an n -measurable function: this is a function on Ω whose value $f(\omega)$ only depends on the states X_0 to X_n . For any $n \in \mathbb{N}$, we let $\mathcal{L}_n(\Omega)$ denote the set of all n -measurable functions taking values in \mathbb{R} . The sets $\overline{\mathcal{L}}_n(\Omega)$ and $\underline{\mathcal{L}}_n(\Omega)$ contain the n -measurable functions taking values in $\overline{\mathbb{R}}$ and $\underline{\mathbb{R}}$, respectively.

A *game-theoretic upper expectation operator* is now a specific $\overline{\mathbb{R}}$ -valued functional $\overline{\mathbb{E}}^V[\cdot | \cdot]$ [26, Definition 2], where the first argument takes values in $\overline{\mathcal{L}}(\Omega)$ and the second is an event of the form $X_{0:n} = x_{0:n}$.

To specify such a game-theoretic upper expectation operator, one needs to provide a family of operators $\overline{Q}_s : \overline{\mathcal{L}}(\mathcal{X}) \rightarrow \overline{\mathbb{R}}$ indexed by *situations* $s \in \mathbb{S} \cup \{\square\}$, with $\mathbb{S} := \{x_{0:n} \in \mathcal{X}^{n+1} : n \in \mathbb{N}_0\}$. Furthermore, for every situation $s \in \mathbb{S} \cup \{\square\}$, \overline{Q}_s should satisfy the following axioms:

- E1. $\overline{Q}_s(c) = c$ for all $c \in \mathbb{R}$;
- E2. $\overline{Q}_s(f + g) \leq \overline{Q}_s(f) + \overline{Q}_s(g)$ for all $f, g \in \overline{\mathcal{L}}(\mathcal{X})$;
- E3. $\overline{Q}_s(\lambda f) = \lambda \overline{Q}_s(f)$ for all positive $\lambda \in \overline{\mathbb{R}}$ and all non-negative $f \in \overline{\mathcal{L}}(\mathcal{X})$;
- E4. if $f \leq g$ then $\overline{Q}_s(f) \leq \overline{Q}_s(g)$ for all $f, g \in \overline{\mathcal{L}}(\mathcal{X})$.

Crucially, every such family leads to a unique corresponding game-theoretic upper expectation operator $\overline{\mathbb{E}}^V[\cdot | \cdot]$. Unfortunately, however, explaining how this works requires quite some technical machinery, including the notion of a supermartingale. Since we believe this would be too much of a digression, we prefer to refer the interested reader to appendix A, and here content ourselves with describing some of its properties.

A first important property is that every \overline{Q}_s can be interpreted as a *local uncertainty model* associated with the situation s . In particular, for $s = x_{0:n}$, we have that

$$\overline{\mathbb{E}}^V[f(X_{n+1}) | X_{0:n} = x_{0:n}] = \overline{Q}_{x_{0:n}}(f), \quad (11)$$

for any $f \in \overline{\mathcal{L}}(\mathcal{X})$. Similarly, the operator \overline{Q}_{\square} describes the uncertainty about the initial state. Note however that, analogous to our discussion for measure-theoretic imprecise Markov chains, we have restricted the domain of $\overline{\mathbb{E}}^V[\cdot | \cdot]$ to be conditional on $X_{0:n}$. Here too, as we explain in Appendix A, this implies that the initial model— \overline{Q}_{\square} , in this case—has no effect on our operator. For ease of notation, we will therefore make no further reference to it, and will henceforth specify $\overline{\mathbb{E}}^V[\cdot | \cdot]$ by providing a family of operators $\{\overline{Q}_s\}_{s \in \mathbb{S}}$, without \overline{Q}_{\square} .

We next remark that the axioms E1–E4 can be recognised as being analogous to familiar properties of coherent lower previsions [29, 18]. The following result essentially shows that the upper expectation operator $\overline{\mathbb{E}}^V[\cdot | \cdot]$ induced by *local* models that satisfy these properties, inherits these properties on the *global* domain $\overline{\mathcal{L}}(\Omega)$. It also provides some properties for the conjugate game-theoretic lower expectation operator, defined as $\underline{\mathbb{E}}^V[\cdot | \cdot] := -\overline{\mathbb{E}}^V[-\cdot | \cdot]$.

Proposition 2 [26, Proposition 13] *Let $\overline{\mathbb{E}}^V[\cdot | \cdot]$ be a game-theoretic upper expectation operator. Then for all $f, g \in \overline{\mathcal{L}}(\Omega)$, all $\lambda \in \mathbb{R}$ with $\lambda \geq 0$, and all $n \in \mathbb{N}_0$:*

1. $\mathbb{E}^V[f + g | X_{0:n}] \leq \mathbb{E}^V[f | X_{0:n}] + \mathbb{E}^V[g | X_{0:n}]$ ³
2. $\mathbb{E}^V[\lambda f | X_{0:n}] = \lambda \mathbb{E}^V[f | X_{0:n}]$
3. $f \leq g \Rightarrow \mathbb{E}^V[f | X_{0:n}] \leq \mathbb{E}^V[g | X_{0:n}]$

and, moreover,

4. for all $x_0, \dots, x_n \in \mathcal{X}$,

$$\begin{aligned} \inf_{\omega \in \Gamma(x_{0:n})} f(\omega) &\leq \mathbb{E}^V[f | X_{0:n} = x_{0:n}] \\ &\leq \mathbb{E}^V[f | X_{0:n} = x_{0:n}] \leq \sup_{\omega \in \Gamma(x_{0:n})} f(\omega) \end{aligned}$$

with $\Gamma(x_{0:n}) := \{\omega \in \Omega \mid \forall t \in \{0, \dots, n\} : \omega(t) = x_t\}$

5. $\mathbb{E}^V[f + \mu | X_{0:n}] = \mathbb{E}^V[f | X_{0:n}] + \mu$ for all $\mu \in \mathbb{R}$.

With the general framework of game-theoretic upper expectation operators in place, we now move on to discussing two specific kinds of such operators, that will be particularly important in the remainder of this work. The first are those that correspond to a precise stochastic process P in the measure-theoretic sense.

Proposition 3 *Let P be a stochastic process as in Section 2.1, and consider the family $\{\bar{Q}_s\}_{s \in \mathbb{S}}$ defined for all $f \in \mathcal{L}(\mathcal{X})$, all $n \in \mathbb{N}_0$, and all $x_0, \dots, x_n \in \mathcal{X}$ as*

$$\bar{Q}_{x_{0:n}}(f) := \mathbb{E}_P[f(X_{n+1}) | X_{0:n} = x_{0:n}].$$

Then the operators in $\{\bar{Q}_s\}_{s \in \mathbb{S}}$ satisfy E1–E4, and therefore determine a unique corresponding game-theoretic upper expectation operator.

We will denote this game-theoretic upper expectation operator as $\mathbb{E}_P^V[\cdot | \cdot]$, and the conjugate game-theoretic lower expectation operator as $\mathbb{E}_P^V[\cdot | \cdot]$.

Our next result establishes that these game-theoretic operators agree with the measure-theoretic expectation \mathbb{E}_P on all n -measurable real-valued functions.

Proposition 4 *Let P be a stochastic process as in Section 2.1 and let $\mathbb{E}_P^V[\cdot | \cdot]$ and $\bar{\mathbb{E}}_P^V[\cdot | \cdot]$ be its game-theoretic lower and upper expectation operators. Then for all $n \in \mathbb{N}_0$, all $x_0, \dots, x_n \in \mathcal{X}$, and all $f_m \in \mathcal{L}_m(\Omega)$ with $m \in \mathbb{N}$, it holds that*

$$\mathbb{E}_P^V[f_m | X_{0:n}] = \bar{\mathbb{E}}_P^V[f_m | X_{0:n}] = \mathbb{E}_P[f_m | X_{0:n}].$$

The second type of game-theoretic expectation operator in which we are interested, is that corresponding to an imprecise Markov chain.

3. If $f + g$ and $\mathbb{E}^V[f | X_{0:n}] + \mathbb{E}^V[g | X_{0:n}]$ are well-defined; the ambiguity of $\infty + -\infty$ makes formalising this property a bit cumbersome.

Proposition 5 *Let \mathcal{T} be a non-empty, closed, and convex set of transition matrices that has separately specified rows, let \bar{T} be the corresponding upper transition operator as in Section 2.1, and consider the family $\{\bar{Q}_s\}_{s \in \mathbb{S}}$ defined for all $f \in \mathcal{L}(\mathcal{X})$, all $n \in \mathbb{N}_0$, and all $x_0, \dots, x_n \in \mathcal{X}$ as*

$$\bar{Q}_{x_{0:n}}(f) := [\bar{T}f](x_n). \quad (12)$$

Then the operators in $\{\bar{Q}_s\}_{s \in \mathbb{S}}$ satisfy E1–E4, and therefore determine a unique corresponding game-theoretic upper expectation operator.

We will denote this game-theoretic upper expectation operator as $\mathbb{E}_{\mathcal{T}}^V[\cdot | \cdot]$, and the conjugate game-theoretic lower expectation operator as $\bar{\mathbb{E}}_{\mathcal{T}}^V[\cdot | \cdot]$.

Since the right-hand side of Equation (12) does not depend on $x_{0:(n-1)}$, it follows from Equation (11) that the induced game-theoretic upper expectation operator satisfies an imprecise Markov property that is entirely similar to that in Equation (6). It is for that reason that we call $\mathbb{E}_{\mathcal{T}}^V[\cdot | \cdot]$ and $\bar{\mathbb{E}}_{\mathcal{T}}^V[\cdot | \cdot]$ the upper and lower expectation operator of a “game-theoretic imprecise Markov chain”.

The following property shows that the operator $\mathbb{E}_{\mathcal{T}}^V[\cdot | \cdot]$ provides a lower bound for the operators $\mathbb{E}_P^V[\cdot | \cdot]$ whose characterising measure-theoretic process P is compatible with \mathcal{T} . Similarly, $\bar{\mathbb{E}}_{\mathcal{T}}^V[\cdot | \cdot]$ provides an upper bound.

Proposition 6 *For all $f \in \mathcal{L}(\Omega)$, all $n \in \mathbb{N}_0$ and all $x_0, \dots, x_n \in \mathcal{X}$, we have that*

$$\mathbb{E}_{\mathcal{T}}^V[f | X_{0:n} = x_{0:n}] \leq \inf_{P \in \mathcal{P}_{\mathcal{T}}^1} \mathbb{E}_P^V[f | X_{0:n} = x_{0:n}]$$

and

$$\sup_{P \in \mathcal{P}_{\mathcal{T}}^1} \bar{\mathbb{E}}_P^V[f | X_{0:n} = x_{0:n}] \leq \bar{\mathbb{E}}_{\mathcal{T}}^V[f | X_{0:n} = x_{0:n}].$$

Finally, we will need the following continuity property; here and in what follows, we consider $\mathcal{L}(\Omega)$ to be endowed with the topology of pointwise convergence:

Proposition 7 *Consider a non-decreasing sequence $\{f_m\}_{m \in \mathbb{N}}$ in $\mathcal{L}(\Omega)$ such that $f_m \in \mathcal{L}_m(\Omega)$ for all $m \in \mathbb{N}$ and $\lim_{m \rightarrow +\infty} f_m = f \in \mathcal{L}(\Omega)$. Then for all $n \in \mathbb{N}_0$ and all $x_0, \dots, x_n \in \mathcal{X}$ it holds that*

$$\mathbb{E}^V[f | X_{0:n} = x_{0:n}] = \lim_{m \rightarrow +\infty} \mathbb{E}^V[f_m | X_{0:n} = x_{0:n}]$$

and

$$\bar{\mathbb{E}}^V[f | X_{0:n} = x_{0:n}] = \lim_{m \rightarrow +\infty} \bar{\mathbb{E}}^V[f_m | X_{0:n} = x_{0:n}].$$

3. Characterisation and Invariance

With the various definitions of imprecise Markov chains in place, we now move on to characterising their (lower and upper) expected hitting times and probabilities, and showing that these are the same for all of the different types of models that we discussed above. We start this discussion with our result for the hitting times, in Section 3.1. The results for the hitting probabilities are largely analogous from a technical point of view, and are presented in Section 3.2.

3.1. Lower and Upper Expected Hitting Times

The starting point for our results in this section is the following well-known characterisation of the expected hitting times of a (precise) homogeneous Markov chain:

Lemma 8 ([19] Theorem 1.3.5) *Consider a homogeneous Markov chain P with transition matrix T . Its vector of expected hitting times $h_A^P \in \overline{\mathcal{L}}(\mathcal{X})$ is the minimal non-negative solution to*

$$h_A^P = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot T h_A^P, \quad (13)$$

where $A^c = \mathcal{X} \setminus A$, and minimality means that $h_A^P(x) \leq g(x)$ for all $x \in \mathcal{X}$, for any non-negative $g \in \overline{\mathcal{L}}(\mathcal{X})$ that also satisfies (13).

Inspired by this result, we introduce a recursive scheme that essentially iterates an imprecise version of Equation (13). To this end, let $\underline{h}_A^{(0)} := \overline{h}_A^{(0)} := \mathbb{I}_{A^c}$ and, for all $n \in \mathbb{N}_0$, define

$$\underline{h}_A^{(n+1)} := \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T} \underline{h}_A^{(n)} \quad (14)$$

and

$$\overline{h}_A^{(n+1)} := \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \overline{T} \overline{h}_A^{(n)}. \quad (15)$$

We will see in Lemma 9 below that these functions can be given a clear interpretation. To this end, for all $n \in \mathbb{N}_0$, let $H_A^{(n)} : \Omega \rightarrow \{0, \dots, n+1\}$ be defined for all $\omega \in \Omega$ as

$$H_A^{(n)}(\omega) := \begin{cases} H_A(\omega) & \text{if } H_A(\omega) \leq n, \text{ and} \\ n+1 & \text{otherwise.} \end{cases}$$

Thus, $H_A^{(n)}(\omega)$ is the number of steps until A was visited on the path ω , provided that this happened in at most n steps; otherwise its value is fixed to be $n+1$. The aforementioned interpretation now goes as follows:

Lemma 9 *For all $n \in \mathbb{N}_0$ it holds that*

$$\underline{h}_A^{(n)} = \mathbb{E}_{\mathcal{F}}^V[H_A^{(n)} | X_0] \quad \text{and} \quad \overline{h}_A^{(n)} = \mathbb{E}_{\mathcal{F}}^V[H_A^{(n)} | X_0].$$

Moreover, it clearly holds that $\lim_{n \rightarrow +\infty} H_A^{(n)} = H_A$. The next result tells us that the equalities in Lemma 9 continue to hold as we pass to this limit; therefore, we can use the above recursive scheme to compute the (lower and upper) expected hitting times for a game-theoretic imprecise Markov chain:

Proposition 10 $\mathbb{E}_{\mathcal{F}}^V[H_A | X_0] = \underline{h}_A^* := \lim_{n \rightarrow +\infty} \underline{h}_A^{(n)}$ and $\mathbb{E}_{\mathcal{F}}^V[H_A | X_0] = \overline{h}_A^* := \lim_{n \rightarrow +\infty} \overline{h}_A^{(n)}$.

Proof Each $H_A^{(n)}$ is n -measurable and the sequence $H_A^{(n)}$ is non-decreasing. Therefore, using Lemma 9 and Proposition 7, the limit \underline{h}_A^* exists and equals $\mathbb{E}_{\mathcal{F}}^V[H_A | X_0]$. The proof for \overline{h}_A^* is completely analogous. ■

In a similar manner, we can use these functions $H_A^{(n)}$ to establish that the game-theoretic hitting times corresponding to a (precise) stochastic process P , agree with the measure-theoretic expected hitting times of this process; this property allows us to relate the quantities obtained under the two different frameworks that we are using.

Lemma 11 *Let P be any measure-theoretic stochastic process. Then $\mathbb{E}_P^V[H_A | X_0] = \overline{\mathbb{E}}_P^V[H_A | X_0] = \mathbb{E}_P[H_A | X_0]$.*

Proof Note that each $H_A^{(n)}$ is n -measurable, that the sequence $H_A^{(n)}$ is non-decreasing and non-negative, and that $\lim_{n \rightarrow +\infty} H_A^{(n)} = H_A$. Hence, using Proposition 7, Proposition 4, and the continuity of $\mathbb{E}_P[\cdot | \cdot]$ with respect to pointwise converging non-decreasing non-negative sequences (Lebesgue's monotone convergence theorem), we find that

$$\begin{aligned} \mathbb{E}_P^V[H_A | X_0] &= \lim_{n \rightarrow +\infty} \mathbb{E}_P^V[H_A^{(n)} | X_0] \\ &= \lim_{n \rightarrow +\infty} \mathbb{E}_P[H_A^{(n)} | X_0] = \mathbb{E}_P[H_A | X_0]. \end{aligned}$$

The proof for $\overline{\mathbb{E}}_P^V[H_A | X_0]$ is completely analogous. ■

We now need one more property before we can state our first main result. Since the sequence $H_A^{(n)}$ is non-decreasing, it follows from Proposition 2 together with Lemma 9 that the sequences $\underline{h}_A^{(n)}$ and $\overline{h}_A^{(n)}$ are also non-decreasing. Hence, using the continuity of \underline{T} with respect to non-decreasing sequences in $\overline{\mathcal{L}}(\mathcal{X})$ —see Lemma 1—we find that

$$\underline{h}_A^* = \lim_{n \rightarrow +\infty} (\mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T} \underline{h}_A^{(n)}) = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T} \underline{h}_A^*. \quad (16)$$

So, \underline{h}_A^* is a fixed-point of the iterative scheme (14). Similarly, \overline{h}_A^* is a fixed-point of (15).

By combining Equation (16) with the properties of game-theoretic expectation operators and the known characterisation for precise, homogeneous Markov chains in Lemma 8, we can now derive the following remarkable consequence; it states that the (lower and upper) expected hitting time for any type of imprecise Markov chain is obtained by a homogeneous Markov chain that is compatible with it. Consequently, the (lower and upper) expected hitting times are the same for all types of imprecise Markov chains!

Theorem 12 *There exists a $P \in \mathcal{P}_{\mathcal{F}}^H$ such that*

$$\mathbb{E}_{\mathcal{F}}^V[H_A | X_0] = \mathbb{E}_{\mathcal{F}}^I[H_A | X_0] = \mathbb{E}_{\mathcal{F}}^H[H_A | X_0] = \mathbb{E}_P[H_A | X_0]. \quad (17)$$

Moreover, there exists a $P \in \mathcal{P}_{\mathcal{F}}^H$, such that

$$\overline{\mathbb{E}}_{\mathcal{F}}^V[H_A | X_0] = \overline{\mathbb{E}}_{\mathcal{F}}^I[H_A | X_0] = \overline{\mathbb{E}}_{\mathcal{F}}^H[H_A | X_0] = \mathbb{E}_P[H_A | X_0]. \quad (18)$$

Proof From the fixed-point claim (16) and the reachability property (5), we find a $T \in \mathcal{T}$ such that

$$\underline{h}_A^* = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot T \underline{h}_A^*. \quad (19)$$

Using (2), we find a homogeneous Markov chain P with transition matrix T , and clearly $P \in \mathcal{P}_{\mathcal{F}}^H$. It remains to show that (17) holds for this P .

Since \underline{h}_A^* satisfies (19), it clearly is a solution to (13). Hence, by Lemma 8 and Proposition 10, it holds that

$$\mathbb{E}_P[H_A | X_0] \leq \underline{h}_A^* = \mathbb{E}_{\mathcal{F}}^V[H_A | X_0].$$

Conversely, we infer from Proposition 6 and Lemma 11 that

$$\begin{aligned} \mathbb{E}_{\mathcal{F}}^V[H_A | X_0] &\leq \inf_{Q \in \mathcal{P}_{\mathcal{F}}^V} \mathbb{E}_Q^V[H_A | X_0] \\ &= \inf_{Q \in \mathcal{P}_{\mathcal{F}}^V} \mathbb{E}_Q[H_A | X_0] \\ &= \mathbb{E}_{\mathcal{F}}^I[H_A | X_0] \\ &\leq \mathbb{E}_{\mathcal{F}}^H[H_A | X_0] \leq \mathbb{E}_P[H_A | X_0], \end{aligned}$$

where the last two inequalities hold since $P \in \mathcal{P}_{\mathcal{F}}^H \subseteq \mathcal{P}_{\mathcal{F}}^I$.

The proof for the upper expected hitting time is far more tedious; it can be found in Reference [14]. ■

We want to stress how powerful this result is: no matter what kind of imprecise generalisation of a Markov chain one wishes to use, the corresponding expected hitting time will always be the same (provided the regularity conditions of the set \mathcal{F} are satisfied). This is not only powerful from a theoretical point of view; numerically, it allows one to use algorithms for computing (lower and upper) expectations of, say, a game-theoretic model, even when the model that one is using is a set of homogeneous Markov chains.

We conclude this section with the following characterisation of the lower and upper expected hitting times of an arbitrary imprecise Markov chain; note that this is a direct generalisation of Lemma 8.

Corollary 13 *Consider an imprecise Markov chain with lower transition operator \underline{T} and upper transition operator \overline{T} . Its vector of lower expected hitting times $\underline{h}_A \in \mathcal{L}(\mathcal{X})$ is the minimal non-negative solution to*

$$\underline{h}_A = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \underline{T} \underline{h}_A, \quad (20)$$

and its vector of upper expected hitting times $\overline{h}_A \in \mathcal{L}(\mathcal{X})$ is the minimal non-negative solution to

$$\overline{h}_A = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot \overline{T} \overline{h}_A. \quad (21)$$

Proof Due to Theorem 12, the lower expected hitting time is the same for every type of imprecise Markov chain; let $\underline{h}_A := \mathbb{E}_{\mathcal{F}}^H[H_A | X_0] = \mathbb{E}_{\mathcal{F}}^V[H_A | X_0]$ be this lower expected hitting time. That \underline{h}_A satisfies (20) is immediate from Proposition 10 and (16). That it is non-negative follows from the non-negativity of H_A .

It remains to show that it is the minimal solution. So, let $\underline{g}_A \in \mathcal{L}(\mathcal{X})$ be any non-negative solution of (20), and suppose *ex absurdo* that $\underline{g}_A(x) < \underline{h}_A(x)$ for some $x \in \mathcal{X}$.

From (20) and (5), we find a $T \in \mathcal{F}$ such that

$$\underline{g}_A = \mathbb{I}_{A^c} + \mathbb{I}_{A^c} \cdot T \underline{g}_A. \quad (22)$$

Using (2), we find a homogeneous Markov chain P with transition matrix T , and clearly $P \in \mathcal{P}_{\mathcal{F}}^H$. By Lemma 8 and (22) we conclude that

$$h_A^P(x) \leq \underline{g}_A(x) < \underline{h}_A(x) = \mathbb{E}_{\mathcal{F}}^H[H_A | X_0 = x],$$

which yields a contradiction using (7).

The proof of the corresponding statement for the upper expected hitting time is a bit more tedious; it can be found in Reference [14]. ■

3.2. Lower and Upper Hitting Probabilities

We now move on to study the (lower and upper) hitting probabilities of an imprecise Markov chain. Both the discussion and the technical results largely mirror that of the hitting times in Section 3.1. We again start with a well-known characterisation for (precise) homogeneous Markov chains:

Lemma 14 ([19] Theorem 1.3.2) *Consider a homogeneous Markov chain P with transition matrix T . Its vector of hitting probabilities $p_A^P \in \mathcal{L}(\mathcal{X})$ is the minimal non-negative solution to*

$$p_A^P = \mathbb{I}_A + \mathbb{I}_{A^c} \cdot T p_A^P. \quad (23)$$

Once more, we proceed by defining a recursive scheme that is inspired by this characterisation: we let $\underline{p}_A^{(0)} := \overline{p}_A^{(0)} := \mathbb{I}_A$ and, for all $n \in \mathbb{N}_0$, we define

$$\underline{p}_A^{(n+1)} := \mathbb{I}_A + \mathbb{I}_{A^c} \cdot \underline{T} \underline{p}_A^{(n)} \quad (24)$$

and

$$\overline{p}_A^{(n+1)} := \mathbb{I}_A + \mathbb{I}_{A^c} \cdot \overline{T} \overline{p}_A^{(n)}. \quad (25)$$

In order to give these functions a clear interpretation, we require some auxiliary functions. For all $n \in \mathbb{N}_0$, we let $G_A^{(n)} : \Omega \rightarrow \{0, 1\}$ be defined for all $\omega \in \Omega$ as

$$G_A^{(n)}(\omega) := \sup\{\mathbb{I}_A(\omega(t)) : t \in \{0, \dots, n\}\}.$$

Thus $G_A^{(n)}$ takes the value one on ω if ω visits A in the first n steps; otherwise it takes the value zero. The aforementioned interpretation now goes as follows:

Lemma 15 *For all $n \in \mathbb{N}_0$ it holds that*

$$\underline{p}_A^{(n)} = \mathbb{E}_{\mathcal{F}}^V[G_A^{(n)} | X_0] \quad \text{and} \quad \overline{p}_A^{(n)} = \mathbb{E}_{\mathcal{F}}^V[G_A^{(n)} | X_0].$$

Moreover, we again clearly have that $\lim_{n \rightarrow +\infty} G_A^{(n)} = G_A$. As the following result tells us, the equalities in Lemma 15 continue to hold as we pass to this limit; so, the above recursive scheme can be used to compute the (lower and upper) hitting probabilities for a game-theoretic imprecise Markov chain:

Proposition 16 $\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0] = \underline{p}_A^* := \lim_{n \rightarrow +\infty} \underline{p}_A^{(n)}$ and $\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0] = \bar{p}_A^* := \lim_{n \rightarrow +\infty} \bar{p}_A^{(n)}$.

Proof First, we remark that each $G_A^{(n)}$ is n -measurable and that the sequence $G_A^{(n)}$ is non-decreasing. Therefore, using Lemma 15 and Proposition 7, the limit \underline{p}_A^* exists and equals $\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0]$. The proof for \bar{p}_A^* is completely analogous. ■

We can also use these functions $G_A^{(n)}$ to establish a connection between the game-theoretic hitting probabilities corresponding to a (precise) stochastic process P , and the measure-theoretic hitting probabilities of this process:

Lemma 17 *Let P be any measure-theoretic stochastic process. Then $\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0] = \mathbb{E}_P^{\mathbb{V}}[G_A | X_0] = \mathbb{E}_P[G_A | X_0]$.*

Proof Completely analogous to the proof of Lemma 11. ■

Since each $\underline{p}_A^{(n)}, \bar{p}_A^{(n)} \in \mathcal{L}(\mathcal{X})$ is bounded—this follows from the boundedness of G_A and each $G_A^{(n)}$, together with Lemma 15, Proposition 16, and Proposition 2—continuity of \underline{T} on $\mathcal{L}(\mathcal{X})$ immediately yields the fixed-point property of the iterative scheme (24):

$$\underline{p}_A^* = \lim_{n \rightarrow +\infty} (\mathbb{I}_A + \mathbb{I}_{A^c} \cdot \underline{T} \underline{p}_A^{(n)}) = \mathbb{I}_A + \mathbb{I}_{A^c} \cdot \underline{T} \underline{p}_A^*. \quad (26)$$

Similarly, \bar{p}_A^* is a fixed-point of the scheme (25). We again conclude that the lower and upper hitting probabilities are the same for all types of imprecise Markov chains:

Theorem 18 *There exists a $P \in \mathcal{P}_{\mathcal{T}}^{\mathbb{H}}$ such that*

$$\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0] = \mathbb{E}_{\mathcal{T}}^{\mathbb{I}}[G_A | X_0] = \mathbb{E}_{\mathcal{T}}^{\mathbb{H}}[G_A | X_0] = \mathbb{E}_P[G_A | X_0]. \quad (27)$$

Moreover, it holds that⁴

$$\mathbb{E}_{\mathcal{T}}^{\mathbb{V}}[G_A | X_0] = \mathbb{E}_{\mathcal{T}}^{\mathbb{I}}[G_A | X_0] = \mathbb{E}_{\mathcal{T}}^{\mathbb{H}}[G_A | X_0]. \quad (28)$$

Proof The proof for the lower hitting probability is completely analogous to the proof of the lower expected hitting time in Theorem 12, only relying on Lemma 14 instead of Lemma 8; on the fixed-point property (26) instead of (16); on Proposition 16 instead of Proposition 10; and on Lemma 17 instead of Lemma 11. The proof for the upper hitting probability is again far more tedious and can be found in Reference [14]. ■

We close with a characterisation of the lower and upper hitting probabilities for an arbitrary imprecise Markov chain, that directly generalises Lemma 14.

4. But note that the upper hitting probability is not necessarily reached by any $P \in \mathcal{P}_{\mathcal{T}}^{\mathbb{H}}$.

Corollary 19 *Consider an imprecise Markov chain with lower transition operator \underline{T} and upper transition operator \bar{T} . Its vector of lower hitting probabilities $\underline{p}_A \in \mathcal{L}(\mathcal{X})$ is the minimal non-negative solution to*

$$\underline{p}_A = \mathbb{I}_A + \mathbb{I}_{A^c} \cdot \underline{T} \underline{p}_A, \quad (29)$$

and its vector of upper hitting probabilities $\bar{p}_A \in \mathcal{L}(\mathcal{X})$ is the minimal non-negative solution to

$$\bar{p}_A = \mathbb{I}_A + \mathbb{I}_{A^c} \cdot \bar{T} \bar{p}_A. \quad (30)$$

Proof The proof for \underline{p}_A is completely analogous to the proof of \underline{h}_A in Corollary 13. The proof for \bar{p}_A is again different, and can be found in Reference [14]. ■

4. Summary and Discussion

We have studied lower and upper expected hitting times and probabilities for imprecise Markov chains. To this end, we considered three different ways in which an imprecise Markov chain might be defined: as a set of precise, homogeneous Markov chains; as a set of precise but general (*non-Markovian*) stochastic processes; and as a game-theoretic model with imprecise local models. We have shown that these quantities of interest are the same for all these types of imprecise Markov chains. Moreover, we have presented characterisations of these quantities that are direct generalisations of their well-known counterparts for precise homogeneous Markov chains.

One unexplored line of research would be to investigate the connections of these results to the theory of Markov Decision Processes (MDPs) [8]. In an MDP, the aim is to choose, at each point in time $n \in \mathbb{N}_0$, an *action* a_n from an admissible action set $\mathbb{A}_n(x_n)$ that determines the transition probabilities $P(X_{n+1} = x_{n+1} | X_n = x_n)$. If we interpret the choice of these actions in our current context as a selection $T \in \mathcal{T}$, then the connection between MDPs and the theory of imprecise Markov chains becomes intuitively clear. It is worth mentioning that this connection has been known for a while—see, e.g., the introductions of [27, 13]—yet an important semantic difference has always been the goal with which actions are selected. In an imprecise Markov chain, we optimise over \mathcal{T} in order to compute bounds on inferential quantities of interest; the goal is the quantification of uncertainty. In contrast, in an MDP, the intended outputs are typically the optimal actions themselves, which are selected to optimise a given utility function that is typically interpreted as an operational reward. However, as it pertains to the results in this current work, in Corollaries 13 and 19 the characterising equations that we have derived are very reminiscent of the equations of optimality that one often encounters in the theory of MDPs, and it would be very interesting to see if this connection could be formalised.

Finally, we hope in future work to derive efficient algorithms for numerically computing the inferences that we have discussed, and aim to also extend our results to the setting of imprecise continuous-time Markov chains [23, 13].

Appendix A. Introduction to Game-theoretic Upper Expectations

For readers that would like to have a better understanding of game-theoretic upper expectations, and how they are derived from the local models \bar{Q}_s , this appendix provides a brief introduction to the game-theoretic probability framework of Shafer and Vovk [20, 21].

Rather than using transition probabilities to describe the uncertain behaviour of a process, they assume that, for every situation $s \in \mathbb{S} \cup \{\square\}$, we are given an operator $\bar{Q}_s: \mathcal{L}(\mathcal{X}) \rightarrow \bar{\mathbb{R}}$ that satisfies E1–E4. As we have already mentioned in Section 2.3, these operators can be thought of as the local uncertainty models of the process. Quite conveniently, as we have for example seen in Propositions 3 and 5, such a local model can be a linear expectation operator corresponding to a probability mass function on \mathcal{X} , or an upper envelope of a set of such linear expectation operators (provided the corresponding set of probability mass functions is closed and convex). In the game-theoretic framework, however, these local models will typically be interpreted as representing the bets that a subject is willing to offer to others. To do this, a function $f(X_{n+1})$, with $f \in \mathcal{L}(\mathcal{X})$ and $n \in \mathbb{N}_0$, is regarded as a bet that yields a (possibly negative) uncertain reward $f(x)$ if $X_{n+1} = x$.⁵ The adopted interpretation for the local model $\bar{Q}_{x_{0:n}}$ is then that conditional on the fact that he observed $X_{0:n} = x_{0:n}$, the subject is willing to offer the bet $f(X_{n+1})$, for any $f \in \mathcal{L}(\mathcal{X})$ such that $\bar{Q}_{x_{0:n}}(f) \leq 0$. Axioms E1–E4 can then be regarded as constraints on what it means to offer bets rationally. In the same way, the operator $\bar{Q}_{\square}(\cdot): \mathcal{L}(\mathcal{X}) \rightarrow \bar{\mathbb{R}}$ represents bets on the initial state X_0 that the subject is willing to offer.

The idea is now to combine all these local bets to obtain a global uncertainty model $\bar{\mathbb{E}}^V[\cdot|\cdot]$ that extends the information that is gathered in the local models. This is achieved using the concept of a *supermartingale*.

Formally, we define a supermartingale \mathcal{M} to be an extended real-valued map on $\mathbb{S} \cup \{\square\}$ that is uniformly bounded below, i.e. there is a real c such that $\mathcal{M}(s) \geq c$ for all $s \in \mathbb{S} \cup \{\square\}$, and that satisfies $\bar{Q}_s(\mathcal{M}(s\cdot)) \leq \mathcal{M}(s)$ for all $s \in \mathbb{S} \cup \{\square\}$. Here, we used $\mathcal{M}(s\cdot)$ to denote the function in $\mathcal{L}(\mathcal{X})$ that takes the value $\mathcal{M}(sx)$ for each $x \in \mathcal{X}$. Indeed, \mathcal{M} is uniformly bounded below, so $\mathcal{M}(s\cdot)$ will only take values in $\bar{\mathbb{R}}$. The key property here is that a supermartingale \mathcal{M} should satisfy $\bar{Q}_s(\mathcal{M}(s\cdot)) \leq \mathcal{M}(s)$ for all $s \in \mathbb{S} \cup \{\square\}$, which essentially states that \mathcal{M} represents a

5. Note that the reward associated with these bets may also be equal to $+\infty$. In that case, it is not immediately clear how we should interpret these bets. This topic is for instance discussed in [25].

possible way to take the subject up on the bets that he is offering. Indeed, if, for the sake of simplicity, we assume that $\mathcal{M}(s)$ is real, then it follows from the constant additivity of \bar{Q}_s —see [26, Proposition 1]—that $\bar{Q}_s(\mathcal{M}(s\cdot)) \leq \mathcal{M}(s)$ is equivalent to $\bar{Q}_s(\mathcal{M}(s\cdot) - \mathcal{M}(s)) \leq 0$, which implies that our subject is willing to offer the bet $\mathcal{M}(s\cdot) - \mathcal{M}(s)$. In this way, it becomes clear that \mathcal{M} describes a possible evolution of a person’s capital when he is gambling according to the bets offered by our subject. For a given choice of local models $\{\bar{Q}_s\}_{s \in \mathbb{S}}$ and \bar{Q}_{\square} , we will use $\bar{\mathbb{M}}_b$ to denote the corresponding set of all such supermartingales.

Consider now any $f \in \mathcal{L}(\Omega)$ and $s = x_{0:n} \in \mathbb{S}$. The global (game-theoretic) upper expectation of f conditional on s is then defined by

$$\bar{\mathbb{E}}^V[f|s] := \inf\{\mathcal{M}(s) : \mathcal{M} \in \bar{\mathbb{M}}_b \text{ and } \liminf \mathcal{M} \geq_s f\},$$

where $\liminf \mathcal{M} \geq_s f$ is taken to mean that for every path $\omega = z_0 z_1 \dots z_n \dots \in \Omega$ such that $z_{0:n} = x_{0:n}$, $\liminf_{m \rightarrow +\infty} \mathcal{M}(z_{0:m}) \geq f(\omega)$.

An intuitive meaning can be given to these upper expectations if we interpret $f \in \mathcal{L}(\Omega)$ as an uncertain reward that depends on the path $\omega \in \Omega$ that is taken by the process. In particular, the game-theoretic upper expectation $\bar{\mathbb{E}}^V[f|s]$ can then be interpreted as the infimum starting capital $\mathcal{M}(s)$ that is needed in order to guarantee that, by starting in the situation $s = x_{0:n}$ and then gambling against the subject in an appropriate (and allowed) way, we can be *sure* to (eventually) end up with a capital that is larger than the reward that is associated with f , in the sense that this will be true for every path $\omega = z_0 z_1 \dots z_n \dots \in \Omega$ such that $z_{0:n} = x_{0:n}$. In other words, if we are in the situation $X_{0:n} = x_{0:n}$, then any capital α larger than $\bar{\mathbb{E}}^V[f|s]$ should be worth more to us than the uncertain reward f , because we can bet with α to (eventually) obtain a reward that is guaranteed to be higher than f . We can therefore regard $\bar{\mathbb{E}}^V[f|s]$ as a lower bound on these capitals α .

As can be expected from this interpretation, the upper expectation $\bar{\mathbb{E}}^V[f|s]$ does not depend on the chosen initial model \bar{Q}_{\square} , thereby justifying our claim in Section 2.3. The following result formalizes this.

Proposition 20 *For any $f \in \mathcal{L}(\mathcal{X})$ and $s \in \mathbb{S}$, the upper expectation $\bar{\mathbb{E}}^V[f|s]$ does not depend on the choice of \bar{Q}_{\square} .*

Finally, we want to add that the global game-theoretic upper expectation operator $\bar{\mathbb{E}}^V[\cdot|\cdot]$ can also be characterised in a completely different way, without the use of game-theoretic concepts such as supermartingales. Indeed, in [25], it is shown that this operator is the most conservative—so least informative—upper expectation that is consistent with the local uncertainty models and satisfies a number of basic rationality axioms.

For more information on the subject of game-theoretic probabilities, we refer the interested reader to the textbooks of Shafer and Vovk [20, 21].

Acknowledgments

The work in this paper was partially supported by H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734. We are very grateful to three anonymous reviewers for their helpful suggestions and constructive comments.

References

- [1] Thomas Augustin, Frank P.A. Coolen, Gert De Cooman, and Matthias C.M. Troffaes. *Introduction to Imprecise Probabilities*. Wiley, 2014.
- [2] Marcilia Andrade Campos, Graçaliz Pereira Dimuro, Antônio Carlos da Rocha Costa, and Vladik Kreinovich. Computing 2-step predictions for interval-valued finite stationary Markov chains. 2003.
- [3] Gert De Cooman and Filip Hermans. Imprecise probability trees: Bridging two theories of imprecise probability. *Artificial Intelligence*, 172(11):1400–1427, 2008.
- [4] Gert De Cooman, Filip Hermans, and Erik Quaeghebeur. Imprecise Markov chains and their limit behavior. *Probability in the Engineering and Informational Sciences*, 23(4):597–635, 2009.
- [5] Gert De Cooman, Filip Hermans, Alessandro Antonucci, and Marco Zaffalon. Epistemic irrelevance in credal nets: the case of imprecise Markov trees. *International Journal of Approximate Reasoning*, 51(9):1029–1052, 2010.
- [6] Gert De Cooman, Jasper De Bock, and Stavros Lopatzidis. Imprecise stochastic processes in discrete time: global models, imprecise Markov chains, and ergodic theorems. *International Journal of Approximate Reasoning*, 76:18–46, 2016.
- [7] Sébastien Destercke and Gert De Cooman. Relating epistemic irrelevance to event trees. In *Soft Methods for Handling Variability and Imprecision*, pages 66–73. Springer, 2008.
- [8] Eugene A Feinberg and Adam Shwartz. *Handbook of Markov decision processes: methods and applications*, volume 40. Springer Science & Business Media, 2012.
- [9] Charles Miller Grinstead and James Laurie Snell. *Introduction to probability*. American Mathematical Soc., 2012.
- [10] Darald J Hartfiel. *Markov set-chains*. Springer, 2006.
- [11] Filip Hermans and Damjan Škulj. Stochastic processes. In Thomas Augustin, Frank P.A. Coolen, Gert De Cooman, and Matthias C.M. Troffaes, editors, *Introduction to Imprecise Probabilities*, chapter 11. Wiley, 2014.
- [12] Igor O Kozine and Lev V Utkin. Interval-valued finite Markov chains. *Reliable computing*, 8(2):97–113, 2002.
- [13] Thomas Krak, Jasper De Bock, and Arno Siebes. Imprecise continuous-time Markov chains. *International Journal of Approximate Reasoning*, 88:452–528, 2017.
- [14] Thomas Krak, Natan T’Joens, and Jasper De Bock. Hitting Times and Probabilities for Imprecise Markov Chains. <https://arxiv.org/abs/1905.08781>, 2019.
- [15] Stavros Lopatzidis. *Robust Modelling and Optimization in Stochastic Processes Using Imprecise Probabilities, with an Application to Queueing Theory*. PhD thesis, Ghent University, 2016.
- [16] Stavros Lopatzidis, Jasper De Bock, and Gert De Cooman. Calculating bounds on expected return and first passage times in finite-state imprecise birth-death chains. In *ISIPTA*, volume 15, pages 177–186, 2015.
- [17] Stavros Lopatzidis, Jasper De Bock, and Gert De Cooman. Computing lower and upper expected first-passage and return times in imprecise birth–death chains. *International Journal of Approximate Reasoning*, 80:137–173, 2017.
- [18] Enrique Miranda and Gert de Cooman. Lower previsions. In Thomas Augustin, Frank P.A. Coolen, Gert De Cooman, and Matthias C.M. Troffaes, editors, *Introduction to Imprecise Probabilities*, chapter 2. Wiley, 2014.
- [19] J.R. Norris. *Markov Chains*. Cambridge University Press, 1997.
- [20] Glenn Shafer and Vladimir Vovk. *Probability and finance: it’s only a game!* Wiley, 2001.
- [21] Glenn Shafer and Vladimir Vovk. *Game-Theoretic Foundations for Probability and Finance*. Wiley, 2019.
- [22] Damjan Škulj. Finite discrete time Markov chains with interval probabilities. In *Soft Methods for Integrated Uncertainty Modelling*, pages 299–306. Springer, 2006.
- [23] Damjan Škulj. Efficient computation of the bounds of continuous time imprecise Markov chains. *Applied mathematics and computation*, 250:165–180, 2015.

- [24] Terence Tao. *An introduction to measure theory*. American Mathematical Society Providence, 2011.
- [25] Natan T'Joens, Jasper De Bock, and Gert de Cooman. In Search of a Global Belief Model for Discrete-Time Uncertain Processes. Accepted for publication in the conference proceedings of ISIPTA 2019.
- [26] Natan T'Joens, Jasper De Bock, and Gert de Cooman. Continuity Properties of Game-Theoretic Upper Expectations. <https://arxiv.org/abs/1902.09406>, 2019.
- [27] Matthias Troffaes and Damjan Skulj. Model checking for imprecise Markov chains. Society for Imprecise Probability: Theories and Applications (SIPTA), 2013.
- [28] Vladimir Vovk and Glenn Shafer. Game-theoretic probability. In Thomas Augustin, Frank P.A. Coolen, Gert De Cooman, and Matthias C.M. Troffaes, editors, *Introduction to Imprecise Probabilities*, chapter 6. Wiley, 2014.
- [29] Peter Walley. *Statistical reasoning with imprecise probabilities*, 1991.