



A Recursive Algorithm for Computing Inferences in Imprecise Markov Chains

Natan T'Joens^(✉), Thomas Krak, Jasper De Bock, and Gert de Cooman

ELIS – FLip, Ghent University, Ghent, Belgium
{natan.tjoens, thomas.krak, jasper.debock, gert.decooman}@ugent.be

Abstract. We present an algorithm that can efficiently compute a broad class of inferences for discrete-time imprecise Markov chains, a generalised type of Markov chains that allows one to take into account partially specified probabilities and other types of model uncertainty. The class of inferences that we consider contains, as special cases, tight lower and upper bounds on expected hitting times, on hitting probabilities and on expectations of functions that are a sum or product of simpler ones. Our algorithm exploits the specific structure that is inherent in all these inferences: they admit a general recursive decomposition. This allows us to achieve a computational complexity that scales linearly in the number of time points on which the inference depends, instead of the exponential scaling that is typical for a naive approach.

Keywords: Imprecise Markov chains · Upper and lower expectations · Recursively decomposable inferences

1 Introduction

Markov chains are popular probabilistic models for describing the behaviour of dynamical systems under uncertainty. The crucial simplifying assumption in these models is that the probabilities describing the system's future behaviour are conditionally independent of its past behaviour, given that we know the current state of the system; this is the canonical *Markov property*.

It is this Markov assumption that makes the parametrisation of a Markov chain relatively straightforward—indeed, as we will discuss in Sect. 2, the uncertain dynamic behaviour is then completely characterised by a transition matrix T , whose elements $T(x_n, x_{n+1}) = P(X_{n+1} = x_{n+1} | X_n = x_n)$ describe the probabilities that the system will transition from any state x_n at time n , to any state x_{n+1} at time $n + 1$. Note that T itself is independent of the time n ; this is the additional assumption of *time homogeneity* that is often imposed implicitly in this context. An important advantage of these assumptions is that the resulting matrix T can be used to solve various important inference problems, using one of the many available efficient algorithms.

In many cases however, the numerical value of the transition matrix T may not be known exactly; that is, there may be additional (higher-order) uncertainty about the model itself. Moreover, it can be argued that simplifying assumptions like the Markov property and time homogeneity are often unrealistic in practice. It is of interest, then, to compute inferences in a manner that is *robust*; both to violations of such simplifying assumptions, and to variations in the numerical values of the transition probabilities.

The theory of *imprecise probabilities* allows us to describe such additional uncertainties by using, essentially, *sets* of traditional (“precise”) models. In particular, such a set is comprised of all the models that we deem “plausible”; for instance, we may include all Markov chains whose characterising transition matrix T is included in some given *set* \mathcal{T} of transition matrices. In this way we can also include non-homogeneous Markov chains, by simply requiring that their (now time-dependent) transition matrices remain in \mathcal{T} . Moreover, we can even include non-Markovian models in such a set. This leads to the notion of an *imprecise Markov chain*. The robust inferences that we are after, are then the tightest possible lower and upper bounds on the inferences computed for each of the included precise models.

In this work, we present an efficient algorithm for solving a large class of these inferences within imprecise Markov chains. Broadly speaking, this class consists of inferences that depend on the uncertain state of the system at a finite number of time instances, and which can be decomposed in a particular recursive form. As we will discuss, it contains as special cases the (*joint*) *probabilities* of sequences of states; the *hitting probabilities* and *expected hitting times* of subsets of the possible states; and *time averages* of functions of the state of the system.

Interestingly, existing algorithms for some of these inferences turn out to correspond to special cases of our algorithm, giving our algorithm a unifying character. Time averages, for example, were already considered in [9], and some of the results in [8]—a theoretical study of lower and upper expected hitting times and probabilities—can be interpreted as a special cases of the algorithm presented here. Readers that are familiar with recursive algorithms for credal networks under epistemic irrelevance [1, 2, 4] might also recognise some of what we do; in fact, many of the ideas behind our algorithm have previously been discussed in this more general context [2, Chapter 7].

In order to adhere to the page limit, all proofs have been relegated to the appendix of an online extended version [13].

2 Preliminaries

We denote the natural numbers, without 0, by \mathbb{N} , and let $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. The set of positive real numbers is denoted by $\mathbb{R}_{>0}$ and the set of non-negative real numbers by $\mathbb{R}_{\geq 0}$. Throughout, we let \mathbb{I}_A denote the indicator of any subset $A \subseteq \mathcal{Y}$ of a set \mathcal{Y} ; so, for any $y \in \mathcal{Y}$, $\mathbb{I}_A(y) := 1$ if $y \in A$ and $\mathbb{I}_A(y) := 0$ otherwise.

Before we can introduce the notion of an imprecise Markov chain, we first need to discuss general (non-Markovian) stochastic processes. These are arguably most commonly formalised using a measure-theoretic approach; however, the majority of our results do not require this level of generality, and so we will keep the ensuing introduction largely intuitive and informal.

Let us start by considering the *realisations* of a stochastic process. At each point in time $n \in \mathbb{N}$, such a process is in a certain *state* x_n , which is an element of a *finite* non-empty state space \mathcal{X} . A realisation of the process is called a *path*, and is an infinite sequence $\omega = x_1x_2x_3 \cdots$ where, at each discrete time point $n \in \mathbb{N}$, $\omega_n := x_n \in \mathcal{X}$ is the state obtained by the process at time n , on the path ω . So, we can interpret any path as a map $\omega : \mathbb{N} \rightarrow \mathcal{X}$, allowing us to collect all paths in the set $\Omega := \mathcal{X}^{\mathbb{N}}$. Moreover, for any $\omega \in \Omega$ and any $m, n \in \mathbb{N}$ with $m \leq n$, we use the notation $\omega_{m:n}$ to denote the finite sequence of states $\omega_m \cdots \omega_n \in \mathcal{X}^{n-m+1}$.

A stochastic process is now an infinite sequence $X_1X_2X_3 \cdots$ of uncertain states where, for all $n \in \mathbb{N}$, the uncertain state at time n is a function of the form $X_n : \Omega \rightarrow \mathcal{X} : \omega \mapsto \omega_n$. Similarly, we can consider finite sequences of such states where, for all $m, n \in \mathbb{N}$ with $m \leq n$, $X_{m:n} : \Omega \rightarrow \mathcal{X}^{n-m+1} : \omega \mapsto \omega_{m:n}$. These states are uncertain in the sense that we do not know which realisation $\omega \in \Omega$ will obtain in reality; rather, we assume that we have assessments of the probabilities $P(X_{n+1} = x_{n+1} | X_{1:n} = x_{1:n})$, for any $n \in \mathbb{N}$ and any $x_{1:n} \in \mathcal{X}^n$. Probabilities of this form tell us something about which state the process might be in at time $n + 1$, given that we know that at time points 1 through n , it followed the sequence $x_{1:n}$. Moreover, we can consider probabilities of the form $P(X_1 = x_1)$ for any $x_1 \in \mathcal{X}$; this tells us something about the state that the process might start in. It is well known that, taken together, these probabilities suffice to construct a global probability model for the *entire* process $X_1X_2X_3 \cdots$, despite each assessment only being about a finite subsequence of the states; see e.g. the discussion surrounding [7, Theorem 5.16] for further details on formalising this in a proper measure-theoretic setting. We simply use P to denote this global model.

Once we have such a global model P , we can talk about *inferences* in which we are interested. In general, these are typically encoded by functions $f : \Omega \rightarrow \mathbb{R}$ of the unknown realisation $\omega \in \Omega$, and we collect all functions of this form in the set $\mathcal{L}(\Omega)$. To compute such an inference consists in evaluating the (conditional) expected value $E_P(f | C)$ of f with respect to the model P , where $C \subseteq \Omega$ is an event of the form $X_{m:n} = x_{m:n}$ with $m, n \in \mathbb{N}$ such that $m \leq n$. In particular, if P is a global model in the measure-theoretic sense, then under some regularity conditions like the measurability of f , we would be interested in computing the quantity $E_P(f | C) := \int_{\Omega} f(\omega) dP(\omega | C)$. For notational convenience, we will also use $X_{1:0} := \Omega$ as a trivial conditioning event, allowing us to regard unconditional expectations as a special case of conditional ones.

A special type of inferences that will play an important role in the remainder of this work are those for which the function f only depends on a *finite* subsequence of the path ω , thereby vastly simplifying the definition of its expectation.

In particular, if an inference only depends on the states at time points m through n , say, then it can always be represented by a function $f : \mathcal{X}^{n-m+1} \rightarrow \mathbb{R}$ evaluated in the uncertain states $X_{m:n}$; specifically, the inference is represented by the composition $f \circ X_{m:n}$, which we will denote by $f(X_{m:n})$. In the sequel, we will call a composite function of this form *finitary*. Moreover, for any $n \in \mathbb{N}$, we denote by $\mathcal{L}(\mathcal{X}^n)$ the set of all functions of the form $f : \mathcal{X}^n \rightarrow \mathbb{R}$, and we write $\mathcal{L}_{\text{fin}}(\Omega) \subset \mathcal{L}(\Omega)$ for the set of all finitary functions. For a finitary function $f(X_{1:n})$, the computation of its expected value reduces to evaluating the finite sum

$$\mathbb{E}_P(f(X_{1:n})|C) = \sum_{x_{1:n} \in \mathcal{X}^n} f(x_{1:n})P(X_{1:n} = x_{1:n}|C).$$

Let us now move on from the discussion about general uncertain processes, to the special case of Markov chains. An uncertain process P is said to satisfy the *Markov property* if, for all $n \in \mathbb{N}$ and all $x_{1:n+1} \in \mathcal{X}^{n+1}$, the aforementioned probability assessments simplify in the sense that

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

A process that satisfies this Markov property is called a *Markov chain*. Thus, for a Markov chain, the probability that it will visit state x_{n+1} at time $n + 1$ is independent of the states $X_{1:n-1}$, given that we know the state X_n at time n . If the process is moreover *homogeneous*, meaning that $P(X_{n+1} = y|X_n = x) = P(X_2 = y|X_1 = x)$ for all $x, y \in \mathcal{X}$ and all $n \in \mathbb{N}$, then the parameterisation of the process becomes exceedingly simple. Indeed, up to the initial distribution $P(X_1)$ —a probability mass function on \mathcal{X} —the process’ behaviour is then fully characterised by a single $|\mathcal{X}| \times |\mathcal{X}|$ matrix T that is called the transition matrix. It is row-stochastic (meaning that, for all $x \in \mathcal{X}$, the x -th row $T(x, \cdot)$ of T is a probability mass function on \mathcal{X}) and its entries satisfy $T(x, y) = P(X_{n+1} = y|X_n = x)$ for all $x, y \in \mathcal{X}$ and $n \in \mathbb{N}$. The usefulness of this representation comes from the fact that we can interpret T as a linear operator on the vector space $\mathcal{L}(\mathcal{X}) \simeq \mathbb{R}^{|\mathcal{X}|}$, due to the assumption that \mathcal{X} is finite. For $f \in \mathcal{L}(\mathcal{X})$, this allows us to write the conditional expectation of $f(X_{n+1})$ given X_n as a matrix-vector product: for any $x \in \mathcal{X}$, $\mathbb{E}_P(f(X_{n+1})|X_n = x)$ equals

$$\sum_{y \in \mathcal{X}} f(y)P(X_{n+1} = y|X_n = x) = \sum_{y \in \mathcal{X}} f(y)T(x, y) = [Tf](x).$$

3 Imprecise Markov Chains

Let us now move on to the discussion about *imprecise* Markov chains. Here, we additionally include uncertainty about the model specification, such as uncertainty about the numerical values of the probabilities $P(X_{n+1}|X_{1:n})$, and about the validity of structural assessments like the Markov property.

We will start this discussion by regarding the parameterisation of such an imprecise Markov chain. We first consider the (imprecise) *initial model* \mathcal{M} ;

this is simply a non-empty set of probability mass functions on \mathcal{X} that we will interpret as containing those probabilities that we deem to plausibly describe the process starting in a certain state. Next, instead of being described by a single transition matrix T , an imprecise Markov chain’s dynamic behaviour is characterised by an entire *set* \mathcal{T} of transition matrices. So, each element $T \in \mathcal{T}$ is an $|\mathcal{X}| \times |\mathcal{X}|$ matrix that is row-stochastic. In the sequel, we will take \mathcal{T} to be fixed, and assume that it is non-empty and that it has separately specified rows. This last property is instrumental in ensuring that computations can be performed efficiently, and is therefore often adopted in the literature; see e.g. [6] for further discussion. For our present purposes, it suffices to know that it means that \mathcal{T} can be completely characterised by providing, for any $x \in \mathcal{X}$, a non-empty set \mathcal{T}_x of probability mass functions on \mathcal{X} . In particular, it means that \mathcal{T} is the set of all row-stochastic $|\mathcal{X}| \times |\mathcal{X}|$ matrices T such that, for all $x \in \mathcal{X}$, the x -row $T(x, \cdot)$ is an element of \mathcal{T}_x .

Given the sets \mathcal{M} and \mathcal{T} , the corresponding imprecise Markov chain is defined as the largest *set* $\mathcal{P}_{\mathcal{M}, \mathcal{T}}$ of stochastic processes that are in a specific sense *compatible* with both \mathcal{M} and \mathcal{T} . In particular, a model P is said to be compatible with \mathcal{M} if $P(X_1) \in \mathcal{M}$, and it is said to be compatible with \mathcal{T} if, for all $n \in \mathbb{N}$ and all $x_{1:n} \in \mathcal{X}^n$, there is some $T \in \mathcal{T}$ such that

$$P(X_{n+1} = x_{n+1} | X_{1:n} = x_{1:n}) = T(x_n, x_{n+1}) \text{ for all } x_{n+1} \in \mathcal{X}.$$

Notably, therefore, $\mathcal{P}_{\mathcal{M}, \mathcal{T}}$ contains all the (precise) homogeneous Markov chains whose characterising transition matrix T is included in \mathcal{T} , and whose initial distribution $P(X_1)$ is included in \mathcal{M} . However, in general, $\mathcal{P}_{\mathcal{M}, \mathcal{T}}$ clearly also contains models that do not satisfy the Markov property, as well as Markov chains that are not homogeneous.¹

For such an imprecise Markov chain, we are interested in computing inferences that are in a specific sense robust with respect to variations in the set $\mathcal{P}_{\mathcal{M}, \mathcal{T}}$. Specifically, for any function of interest $f : \Omega \rightarrow \mathbb{R}$, we consider its (conditional) *lower* and *upper* expectations, which are respectively defined by

$$\underline{E}_{\mathcal{M}, \mathcal{T}}(f|C) := \inf_{P \in \mathcal{P}_{\mathcal{M}, \mathcal{T}}} E_P(f|C) \text{ and } \bar{E}_{\mathcal{M}, \mathcal{T}}(f|C) := \sup_{P \in \mathcal{P}_{\mathcal{M}, \mathcal{T}}} E_P(f|C).$$

In words, we are interested in computing the tightest possible bounds on the inferences computed for each $P \in \mathcal{P}_{\mathcal{M}, \mathcal{T}}$. These lower and upper expectations are related through conjugacy, meaning that $\underline{E}_{\mathcal{M}, \mathcal{T}}(f|C) = -\bar{E}_{\mathcal{M}, \mathcal{T}}(-f|C)$, so it suffices to consider only the upper expectations in the remaining discussion; any results for lower expectations follow analogously through this relation.

From a computational point of view, it is also useful to consider the dual representation of the set \mathcal{T} , given by the *upper transition operator* \bar{T} with respect to this set [5,6]. This is a (non-linear) operator that maps $\mathcal{L}(\mathcal{X})$ into $\mathcal{L}(\mathcal{X})$; it is defined for any $f \in \mathcal{L}(\mathcal{X})$ and any $x \in \mathcal{X}$ as

$$[\bar{T}f](x) := \sup_{T(x, \cdot) \in \mathcal{T}_x} \sum_{y \in \mathcal{X}} T(x, y)f(y).$$

¹ Within the field of imprecise probability theory, this model is called an *imprecise Markov chain under epistemic irrelevance* [5,6,9].

So, in order to evaluate $[\overline{T}f](x)$, one must solve an optimisation problem over the set \mathcal{T}_x containing the x -rows of the elements of \mathcal{T} . In many practical cases, the set \mathcal{T}_x is closed and convex and therefore, evaluating $[\overline{T}f](x)$ is relatively straightforward: for instance, if \mathcal{T}_x is described by a finite number of (in)equality constraints, then this problem reduces to a simple linear programming task, which can be solved by standard techniques. We will also make use of the conjugate *lower transition operator* \underline{T} , defined by $[\underline{T}f](x) := -[\overline{T}(-f)](x)$ for all $x \in \mathcal{X}$ and all $f \in \mathcal{L}(\mathcal{X})$. Results about upper transition operators translate to results about lower transition operators through this relation; we will focus on the former in the following discussion.

Now, the operator \overline{T} can be used for computing upper expectations in much the same way as transition matrices are used for computing expectations with respect to precise Markov chains: for any $n \in \mathbb{N}$, any finitary function $f(X_{n+1})$ and any $x_{1:n} \in \mathcal{X}^n$ it holds that

$$\overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f(X_{n+1})|X_{1:n} = x_{1:n}) = [\overline{T}f](x_n). \tag{1}$$

Observe that the right-hand side in this expression does not depend on the history $x_{1:n-1}$; this can be interpreted as saying that the model satisfies an *imprecise Markov property*, which explains why we call our model an “imprecise Markov chain”. Moreover, a slightly more general property holds that will be useful later on:

Proposition 1. *Consider the imprecise Markov chain $\mathcal{P}_{\mathcal{M},\mathcal{T}}$. For any $m, n \in \mathbb{N}$ such that $m \leq n$, any function $f \in \mathcal{L}(\mathcal{X}^{n-m+1})$ and any $x_{1:m-1} \in \mathcal{X}^{m-1}$ and $y \in \mathcal{X}$, we have that*

$$\overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f(X_{m:n})|X_{1:m-1} = x_{1:m-1}, X_m = y) = \overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f(X_{1:n-m+1})|X_1 = y).$$

Finally, we remark that, for any $m, n \in \mathbb{N}$ such that $m \leq n$, a conditional upper expectation $\overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f|X_{m:n})$ is itself a (finitary) function depending on the states $X_{m:n}$. Using this observation, we can now introduce the *law of iterated upper expectations*, which will form the basis of the algorithms developed in the following sections:

Theorem 1. *Consider the imprecise Markov chain $\mathcal{P}_{\mathcal{M},\mathcal{T}}$. For all $m \in \mathbb{N}_0$, all $k \in \mathbb{N}$ and all $f \in \mathcal{L}_{\text{fin}}(\Omega)$, we have that*

$$\overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f|X_{1:m}) = \overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}\left(\overline{\mathbb{E}}_{\mathcal{M},\mathcal{T}}(f|X_{1:m+k})\middle|X_{1:m}\right).$$

4 A Recursive Inference Algorithm

In principle, for any function $f \in \mathcal{L}(\mathcal{X}^n)$ with $n \in \mathbb{N}$, the upper expectations of $f(X_{1:n})$ can be obtained by maximising $\mathbb{E}_P(f(X_{1:n}))$ over the set $\mathcal{P}_{\mathcal{M},\mathcal{T}}$ of all precise models P that are compatible with \mathcal{M} and \mathcal{T} . Since this will almost always be infeasible if n is large, we usually apply the law of iterated upper

expectations in combination with the Markov property in order to divide the optimisation problem into multiple smaller ones. Indeed, because of Theorem 1, we have that

$$\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})) = \bar{E}_{\mathcal{M},\mathcal{F}}\left(\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_{1:n-1})\right).$$

Using Eq. (1), one can easily show that $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_{1:n-1})$ can be computed by evaluating $[\bar{T}f(x_{1:n-1})](x_{n-1})$ for all $x_{1:n-1} \in \mathcal{X}^{n-1}$. Here, $f(x_{1:n-1}\cdot)$ is the function in $\mathcal{L}(\mathcal{X})$ that takes the value $f(x_{1:n})$ on $x_n \in \mathcal{X}$. This accounts for $|\mathcal{X}|^{n-1}$ optimisation problems to be solved. With the acquired function $f'(X_{1:n-1}) := \bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_{1:n-1})$, we can then compute the upper expectation $\bar{E}_{\mathcal{M},\mathcal{F}}(f'(X_{1:n-1})|X_{1:n-2})$ in a similar way, by solving $|\mathcal{X}|^{n-2}$ optimisation problems. Continuing in this way, we end up with a function that only depends on X_1 and for which the expectation needs to be maximised over the initial models in \mathcal{M} . Hence, in total, $\sum_{i=0}^{n-1} |\mathcal{X}|^i$ optimisation problems need to be solved in order to obtain $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n}))$. Although these optimisation problems are relatively simple and therefore feasible to solve individually, the total number of required iterations is still exponential in n , therefore making the computation of $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n}))$ intractable when n is large.

In many cases, however, $f(X_{1:n})$ can be recursively decomposed in a specific way allowing for a much more efficient computational scheme to be employed; see Theorem 2 further on. Before we present this scheme in full generality, let us first provide some intuition about its basic working principle.

So assume we are interested in $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n}))$, which, according to Theorem 1, can be obtained by maximising $E_P(\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_1))$ over $P(X_1) \in \mathcal{M}$. The problem then reduces to the question of how to compute $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_1)$ efficiently. Suppose now that $f(X_{1:n})$ takes the following form:

$$f(X_{1:n}) = g(X_1) + h(X_1)\tau(X_{2:n}), \tag{2}$$

for some $g, h \in \mathcal{L}(\mathcal{X})$ and some $\tau \in \mathcal{L}(\mathcal{X}^{n-1})$. Then, because $\bar{E}_{\mathcal{M},\mathcal{F}}$ is a supremum over linear expectations, we find that

$$\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_1) = g(X_1) + h(X_1)\bar{E}_{\mathcal{M},\mathcal{F}}(\tau(X_{2:n})|X_1),$$

where, for the sake of simplicity, we assumed that h does not take negative values. Then, by appropriately combining Proposition 1 with Theorem 1, one can express $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau(X_{2:n})|X_1)$ in terms of \bar{Y} : $\mathcal{X} \rightarrow \mathbb{R}$, defined by

$$\bar{Y}(x) := \bar{E}_{\mathcal{M},\mathcal{F}}(\tau(X_{1:n-1})|X_1 = x) \text{ for all } x \in \mathcal{X}.$$

In particular, we find that

$$\begin{aligned} \bar{E}_{\mathcal{M},\mathcal{F}}(\tau(X_{2:n})|X_1) &= \bar{E}_{\mathcal{M},\mathcal{F}}(\bar{E}_{\mathcal{M},\mathcal{F}}(\tau(X_{2:n})|X_{1:2})|X_1) = \bar{E}_{\mathcal{M},\mathcal{F}}(\bar{Y}(X_2)|X_1) \\ &= [\bar{T}\bar{Y}](X_1), \end{aligned}$$

where the equalities follow from Theorem 1, Proposition 1 and Eq. (1), respectively. So $\bar{E}_{\mathcal{M},\mathcal{F}}(f(X_{1:n})|X_1)$ can be obtained from \bar{Y} by solving a single optimisation problem, followed by a pointwise multiplication and summation.

Now, by repeating the structural assessment (2) in a recursive way, we can generate a whole class of functions for which the upper expectations can be computed using the principle illustrated above. We start with a function $\tau_1(X_1)$, with $\tau_1 \in \mathcal{L}(\mathcal{X})$, that only depends on the initial state. The upper expectation $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau_1(X_1)|X_1)$ is then trivially equal to $\tau_1(X_1)$. Next, consider $\tau_2(X_{1:2}) = g_1(X_1) + h_1(X_1)\tau_1(X_2)$ for some g_1, h_1 in $\mathcal{L}(\mathcal{X})$. $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau_2(X_{1:2})|X_1)$ is then given by $g_1(X_1) + h_1(X_1)[\bar{T}\bar{Y}_1](X_1)$, where we let $\bar{Y}_1(x) := \bar{E}_{\mathcal{M},\mathcal{F}}(\tau_1(X_1)|X_1 = x) = \tau_1(x)$ for all $x \in \mathcal{X}$ and where we (again) neglect the subtlety that h_1 can take negative values. Continuing in this way, step by step considering new functions constructed by multiplication and summation with functions that depend on an additional time instance, and no longer ignoring the fact that the functions involved can take negative values, we end up with the following result.

Theorem 2. *Consider any imprecise Markov chain $\mathcal{P}_{\mathcal{M},\mathcal{F}}$ and two sequences of functions $\{g_n\}_{n \in \mathbb{N}_0}$ and $\{h_n\}_{n \in \mathbb{N}}$ in $\mathcal{L}(\mathcal{X})$. Define $\tau_1(x_1) := g_0(x_1)$ for all $x_1 \in \mathcal{X}$, and for all $n \in \mathbb{N}$, let*

$$\tau_{n+1}(x_{1:n+1}) := h_n(x_1)\tau_n(x_{2:n+1}) + g_n(x_1) \text{ for all } x_{1:n+1} \in \mathcal{X}^{n+1}.$$

If we write $\{\bar{Y}_n\}_{n \in \mathbb{N}}$ and $\{\underline{Y}_n\}_{n \in \mathbb{N}}$ to denote the sequences of functions in $\mathcal{L}(\mathcal{X})$ that are respectively defined by $\bar{Y}_n(x) := \bar{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1 = x)$ and $\underline{Y}_n(x) := \underline{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1 = x)$ for all $x \in \mathcal{X}$ and all $n \in \mathbb{N}$, then $\{\bar{Y}_n\}_{n \in \mathbb{N}}$ and $\{\underline{Y}_n\}_{n \in \mathbb{N}}$ satisfy the following recursive expressions:

$$\begin{cases} \bar{Y}_1 = \underline{Y}_1 = g_0; \\ \bar{Y}_{n+1} = h_n \mathbb{I}_{h_n \geq 0} [\bar{T}\bar{Y}_n] + h_n \mathbb{I}_{h_n < 0} [\underline{T}\underline{Y}_n] + g_n \text{ for all } n \in \mathbb{N}; \\ \underline{Y}_{n+1} = h_n \mathbb{I}_{h_n \geq 0} [\underline{T}\underline{Y}_n] + h_n \mathbb{I}_{h_n < 0} [\bar{T}\bar{Y}_n] + g_n \text{ for all } n \in \mathbb{N}. \end{cases}$$

Here, we used $\mathbb{I}_{h_n \geq 0} \in \mathcal{L}(\mathcal{X})$ to denote the indicator of $\{x \in \mathcal{X} : h_n(x) \geq 0\}$, and similarly for $\mathbb{I}_{h_n < 0} \in \mathcal{L}(\mathcal{X})$. Note that, because we now need to evaluate both \bar{T} and \underline{T} for every iteration, we will in general need to solve $2(n-1)|\mathcal{X}|$ optimisation problems to obtain $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1)$ and $\underline{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1)$ for some $n \in \mathbb{N}$. In order to obtain the unconditional inferences $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n}))$ and $\underline{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n}))$, it then suffices to respectively maximise and minimise the expectations of $\bar{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1)$ and $\underline{E}_{\mathcal{M},\mathcal{F}}(\tau_n(X_{1:n})|X_1)$ over all initial models in \mathcal{M} .

5 Special Cases

To illustrate the practical relevance of our method, we now discuss a number of important inferences that fall within its scope. As already mentioned in the introduction section, in some of these cases, our method simplifies to a computational scheme that was already developed earlier in a more specific context. The strength of our present contribution, therefore, lays in its unifying character and the level of generality to which it extends.

Functions that depend on a single time instant. As a first, very simple inference we can consider the upper and lower expectation of a function $f(X_n)$, for some $f \in \mathcal{L}(\mathcal{X})$ and $n \in \mathbb{N}$, conditional on the initial state. The expressions for these inferences are given by $\overline{T}^{n-1}f$ and $\underline{T}^{n-1}f$, respectively [5]. For instance, for any $x \in \mathcal{X}$, $\underline{\mathbb{E}}_{\mathcal{M}, \mathcal{S}}(f(X_5) | X_1 = x) = [\underline{T}^4 f](x)$. These expressions can also easily be obtained from Theorem 2, by setting $g_0 := f$ and, for all $k \in \{1, \dots, n-1\}$, $g_k := 0$ and $h_k := 1$.

Sums of functions. One can also use our method to compute upper and lower expectations of sums $\sum_{k=1}^n f_k(X_k)$ of functions $f_k \in \mathcal{L}(\mathcal{X})$. Then we would have to set $g_0 := f_n$ and, for all $k \in \{1, \dots, n-1\}$, $g_k := f_{n-k}$ and $h_k := 1$. Although we allow the functions f_k to depend on k , it is worth noting that, if we set them all equal to the same function f , our method can also be employed to compute the upper and lower expectation of the time average $1/n \sum_{k=1}^n f(X_k)$ of f over the time interval n . The subtlety of the constant factor $1/n$ does not raise a problem here, because upper and lower expectations are homogeneous with respect to non-negative scaling.

Product of functions. Another interesting class of inferences are those that can be represented by a product $\prod_{k=1}^n f_k(X_k)$ of functions $f_k \in \mathcal{L}(\mathcal{X})$. To compute upper and lower expectations of such functions, it suffices to set $g_0 := f_n$ and, for all $k \in \{1, \dots, n-1\}$, $g_k := 0$ and $h_k := f_{n-k}$. A typical example of an inference than can be described in this way is the probability that the state will be in a set $A \subseteq \mathcal{X}$ during a certain time interval. For instance, the upper expectation of the function $\mathbb{I}_A(X_1)\mathbb{I}_A(X_2)$ gives us a tight upper bound on the probability that the state will be in A during the first two time instances.

Hitting probabilities. The hitting probability of some set $A \subseteq \mathcal{X}$ over a finite time interval n is the probability that the state X_k will be in A somewhere within the first n time instances. The upper and lower bounds on such a hitting probability are equal to the upper and lower expectation of the function $f(X_{1:n}) := \mathbb{I}_{A'_n} \in \mathcal{L}(\Omega)$, where $A'_n := \{\omega \in \Omega : (\exists k \leq n) \omega_k \in A\}$. Note that $f(X_{1:n})$ can be decomposed in the following way:

$$f(X_{1:n}) = \mathbb{I}_A(X_1) + \mathbb{I}_A(X_2)\mathbb{I}_{A^c}(X_1) + \dots + \mathbb{I}_A(X_n) \prod_{k=1}^{n-1} \mathbb{I}_{A^c}(X_k)$$

Hence, these inferences can be obtained using Theorem 2 if we let $g_0 := \mathbb{I}_A$ and, for all $k \in \{1, \dots, n-1\}$, $g_k := \mathbb{I}_A$ and $f_k := \mathbb{I}_{A^c}$. Additionally, one could also be interested in the probability that the state X_k will *ever* be in A . Upper and lower bounds on this probability are given by the upper and lower expectation of the function $f := \mathbb{I}_{A'} \in \mathcal{L}(\Omega)$ where $A' := \{\omega \in \Omega : (\exists k \in \mathbb{N}) \omega_k \in A\}$. Since the function f is non-finitary, we are unable to apply our method in a direct way. However, it is shown in [8, Proposition 16] that, if the set \mathcal{S} is convex and closed, the upper and lower bounds on the hitting probability over a finite time interval converge to the upper and lower bounds on the hitting probability over an infinite time interval, therefore allowing us to approximate these inferences by choosing n sufficiently large.

Hitting times. The *hitting time* of some set $A \subseteq \mathcal{X}$ is defined as the time τ until the state is in A for the first time; so $\tau(\omega) := \inf\{k \in \mathbb{N}_0 : \omega_k \in A\}$ for all $\omega \in \Omega$. Once more, the function τ is non-finitary, necessitating an indirect approach to the computation of its upper and lower expectation. This can be done in a similar way as we did for the case of hitting probabilities, now considering the finitary functions $\tau_n(X_{1:n})$, where $\tau_n \in \mathcal{L}(\mathcal{X}^n)$ is defined by $\tau_n(x_{1:n}) := \inf\{k \in \mathbb{N} : x_k \in A\}$ if $\{k \in \mathbb{N} : x_k \in A\}$ is non-empty, and $\tau_n(x_{1:n}) := n+1$ otherwise, for all $n \in \mathbb{N}$ and all $x_{1:n} \in \mathcal{X}^n$. These functions correspond to choosing $g_0 := \mathbb{I}_{A^c}$ and, for all $k \in \{1, \dots, n-1\}$, $g_k := \mathbb{I}_{A^c}$ and $f_k := \mathbb{I}_A$. If the set \mathcal{T} is convex and closed, the upper and lower expectations of these functions for large n will then approximate those of the non-finitary hitting time [8, Proposition 10].

6 Discussion

The main contribution of this paper is a single, unified method to efficiently compute a wide variety of inferences for imprecise Markov chains; see Theorem 2. The set of functions describing these inferences is however restricted to the finitary type, and therefore a general approach for inferences characterised by non-finitary functions is still lacking. In some cases, however, as we already mentioned in our discussion of hitting probabilities and hitting times, this issue can be addressed by relying on a continuity argument.

Indeed, consider any function $f = \lim_{n \rightarrow +\infty} \tau_n(X_{1:n})$ that is the pointwise limit of a sequence $\{\tau_n(X_{1:n})\}_{n \in \mathbb{N}}$ of finitary functions, defined recursively as in Theorem 2. If $\bar{E}_{\mathcal{M}, \mathcal{T}}$ is continuous with respect to $\{\tau_n(X_{1:n})\}_{n \in \mathbb{N}}$, meaning that $\lim_{n \rightarrow +\infty} \bar{E}_{\mathcal{M}, \mathcal{T}}(\tau_n(X_{1:n})) = \bar{E}_{\mathcal{M}, \mathcal{T}}(f)$, the inference $\bar{E}_{\mathcal{M}, \mathcal{T}}(f)$ can then be approximated by $\bar{E}_{\mathcal{M}, \mathcal{T}}(\tau_n(X_{1:n}))$ for sufficiently large n . Since we can recursively compute $\bar{E}_{\mathcal{M}, \mathcal{T}}(\tau_n(X_{1:n}))$ for any $n \in \mathbb{N}$ using the methods discussed at the end of Sect. 4, this yields an efficient way of approximating $\bar{E}_{\mathcal{M}, \mathcal{T}}(f)$. A completely analogous argument can be used for the lower expectation $\underline{E}_{\mathcal{M}, \mathcal{T}}(f)$. This begs the question whether the upper and lower expectations $\bar{E}_{\mathcal{M}, \mathcal{T}}$ and $\underline{E}_{\mathcal{M}, \mathcal{T}}$ satisfy the appropriate continuity properties for this to work.

Unfortunately, results about the continuity properties of these operators are rather scarce —especially compared to their precise counterparts—and depend on the formalism that is adopted. In this paper, for didactical reasons, we have considered one formalism: we have introduced imprecise Markov chains as being *sets* of “precise” models that are in a specific sense compatible with the given set \mathcal{T} . It is however important to realise that there is also an entirely different formalisation of imprecise Markov chains that is instead based on the game-theoretic probability framework that was popularised by Shafer and Vovk; we refer to [10, 11] for details. It is well known that the inferences produced under these two different frameworks agree for finitary functions [3, 9], so the method described by Theorem 2 is also applicable when working in a game-theoretic framework. The continuity properties of the game-theoretic upper and lower expectations, however, are not necessarily the same as those of the measure-theoretic operators that we considered here. So far, the continuity properties

of game-theoretic upper and lower expectations are better understood [10–12], making these operators more suitable if we plan to employ the continuity argument above.

Acknowledgments. The work in this paper was partially supported by H2020-MSCA-ITN-2016 UTOPIAE, grant agreement 722734.

References

1. De Bock, J.: Credal networks under epistemic irrelevance. *Int. J. Approximate Reasoning* **85**, 107–138 (2017)
2. De Bock, J.: Credal networks under epistemic irrelevance: theory and algorithms. Ph.D. thesis, Ghent University (2015)
3. de Cooman, G., Hermans, F.: Imprecise probability trees: bridging two theories of imprecise probability. *Artif. Intell.* **172**(11), 1400–1427 (2008)
4. de Cooman, G., Hermans, F., Antonucci, A., Zaffalon, M.: Epistemic irrelevance in credal nets: the case of imprecise Markov trees. *Int. J. Approximate Reasoning* **51**(9), 1029–1052 (2010)
5. de Cooman, G., Hermans, F., Quaeghebeur, E.: Imprecise Markov chains and their limit behaviour. *Probab. Eng. Inf. Sci.* **23**(4), 597–635 (2009)
6. Hermans, F., Škulj, D.: Stochastic processes. In: Augustin, T., Coolen, F.P.A., De Cooman, G., Troffaes, M.C.M. (eds.) *Introduction to Imprecise Probabilities*. Wiley (2014). Chapter 11
7. Kallenberg, O.: *Foundations of Modern Probability*. Springer Science & Business Media, New York (2006)
8. Krak, T’Joens, N., De Bock, J.: Hitting times and probabilities for imprecise Markov chains (2019). Accepted for publication in the conference proceedings of ISIPTA 2019. A preprint can be found at [arXiv:1905.08781](https://arxiv.org/abs/1905.08781)
9. Lopatzidis, S.: *Robust Modelling and Optimisation in Stochastic Processes using Imprecise Probabilities, with an Application to Queueing Theory*. Ph.D. thesis, Ghent University (2017)
10. Shafer, G., Vovk, V.: *Probability and Finance: It’s Only a Game!*. Wiley, New York (2001)
11. T’Joens, N., De Bock, J., de Cooman, G.: In search of a global belief model for discrete-time uncertain processes (2019). Accepted for publication in the conference proceedings of ISIPTA 2019
12. T’Joens, N., De Bock, J., de Cooman, G.: Continuity properties of game-theoretic upper expectations. [arXiv:1902.09406](https://arxiv.org/abs/1902.09406) (2019)
13. T’Joens, N., Krak, T., De Bock, J., de Cooman, G.: A Recursive Algorithm for Computing Inferences in Imprecise Markov Chains. [arXiv:1905.12968](https://arxiv.org/abs/1905.12968) (2019)