



Limit Behaviour of Upper and Lower Expected Time Averages in Discrete-Time Imprecise Markov Chains

Natan T'Joens^(✉) and Jasper De Bock

FLip, Ghent University, Ghent, Belgium
{natan.tjoens, jasper.debock}@ugent.be

Abstract. We study the limit behaviour of upper and lower bounds on expected time averages in imprecise Markov chains; a generalised type of Markov chain where the local dynamics, traditionally characterised by transition probabilities, are now represented by sets of ‘plausible’ transition probabilities. Our main result is a necessary and sufficient condition under which these upper and lower bounds, called upper and lower expected time averages, will converge as time progresses towards infinity to limit values that do not depend on the process’ initial state. Remarkably, our condition is considerably weaker than those needed to establish similar results for so-called limit—or steady state—upper and lower expectations, which are often used to provide approximate information about the limit behaviour of time averages as well. We show that such an approximation is sub-optimal and that it can be significantly improved by directly using upper and lower expected time averages.

Keywords: Imprecise Markov chain · Upper expectation · Upper transition operator · Expected time average · Weak ergodicity

1 Introduction

Markov chains are probabilistic models that can be used to describe the uncertain dynamics of a large variety of stochastic processes. One of the key results within the field is the point-wise ergodic theorem. It establishes a relation between the long-term time average of a real-valued function and its limit expectation, which is guaranteed to exist if the Markov chain is ergodic.¹ For this reason, limit expectations and limit distributions have become central objects of interest. Of course, if one is interested in the long-term behaviour of time averages, one could also study the expected values of these averages directly. This is not often done though, because the limit of these expected time averages coincides with the

¹ The term ergodicity has various meanings; sometimes it refers to properties of an invariant measure, sometimes it refers to properties such as irreducibility (with or without aperiodicity), regularity, ... Our usage of the term follows conventions introduced in earlier work [2, 8] on imprecise Markov chains; see Sects. 2 and 4.

forementioned limit expectations, which can straightforwardly be obtained by solving a linear eigenproblem [10].

We here consider a generalisation of Markov chains, called imprecise Markov chains [2, 4, 9], for which the considerations above are not necessarily true. Imprecise Markov chains are sets of traditional (“precise”) probabilistic models, where the Markov property (history independence) and time-homogeneity apply to the collection of precise models as a whole, but not necessarily to the individual models themselves. Imprecise Markov chains therefore allow one to incorporate model uncertainty about the numerical values of the transition probabilities that make up a Markov chain, but also, and more importantly, about structural assumptions such as time-homogeneity and the Markov property. For such an imprecise Markov chain, one is then typically interested in obtaining tight upper and lower bounds on inferences for the individual constituting models. The operators that represent these upper and lower bounds are respectively called upper and lower expectations.

Just like traditional Markov chains can have a limit expectation, an imprecise Markov chain can have limit upper and lower expectations. There are necessary and sufficient conditions for their existence [8] as well as an imprecise variant of the point-wise ergodic theorem [2]. An important difference with traditional Markov chains however, is that upper and lower bounds on expectations of time averages—we will call these upper and lower expected time averages—may not converge to limit upper and lower expectations. Nevertheless, because they give conservative bounds [11, Lemma 57], and because they are fairly easy to compute, limit upper and lower expectations are often used as descriptors of the long-term behaviour of imprecise Markov chains, even if one is actually interested in time averages. This comes at a cost though: as we illustrate in Sect. 4, both inferences can differ greatly, with limit expectations providing far too conservative bounds.

Unfortunately, apart from some experiments in [11], little is known about the long-term behaviour of upper and lower expected time averages in imprecise Markov chains. The aim of this paper is to remedy this situation. Our main result is an accessibility condition that is necessary and sufficient for upper and lower expected time averages to converge to a limit value that does not depend on the process’ initial state; see Sect. 7. Remarkably, this condition is considerably weaker than the ones required for limit lower and upper expectations to exist.

Technical proofs are relegated to the appendix of an extended online version [12]. This is particularly true for the results in Sect. 7, where the main text provides an informal argument that aims to provide intuition.

2 Markov Chains

We consider an infinite sequence $X_0X_1X_2\cdots$ of uncertain states, where each state X_k at time $k \in \mathbb{N}_0 := \mathbb{N} \cup \{0\}$ takes values in some finite set \mathcal{X} , called the *state space*. Such a sequence $X_0X_1X_2\cdots$ will be called a (*discrete-time*) *stochastic process*. For any $k, \ell \in \mathbb{N}_0$ such that $k \leq \ell$, we use $X_{k:\ell}$ to denote the

finite subsequence $X_k \cdots X_\ell$ of states that takes values in $\mathcal{X}^{\ell-k+1}$. Moreover, for any $k, \ell \in \mathbb{N}_0$ such that $k \leq \ell$ and any $x_{k:\ell} \in \mathcal{X}^{\ell-k+1}$, we use $X_{k:\ell} = x_{k:\ell}$ to denote the event that $X_k = x_k \cdots X_\ell = x_\ell$. The uncertain dynamics of a stochastic process are then typically described by probabilities of the form $P(X_{k+1} = x_{k+1} | X_{0:k} = x_{0:k})$, for any $k \in \mathbb{N}_0$ and any $x_{0:k+1} \in \mathcal{X}^{k+2}$. They represent beliefs about which state the process will be in at time $k + 1$ given that we know that it was in the states $x_0 \cdots x_k$ at time instances 0 through k . Additionally, our beliefs about the value of the initial state X_0 can be represented by probabilities $P(X_0 = x_0)$ for all $x_0 \in \mathcal{X}$. The local probability assessments $P(X_{k+1} = x_{k+1} | X_{0:k} = x_{0:k})$ and $P(X_0 = x_0)$ can now be combined to construct a global probability model P that describes the dynamics of the process on a more general level. This can be done in various ways; one of the most common ones being a measure-theoretic approach where countable additivity plays a central role. For our purposes however, we will only require finite additivity. Regardless, once you have such a global probability model P , it can then be used to define expectations and make inferences about the uncertain behaviour of the process.

For any set A , let us write $\mathcal{L}(A)$ to denote the set of all real-valued functions on A . Throughout, for any $a \in A$, we use \mathbb{I}_a to denote the *indicator* of a : the function in $\mathcal{L}(A)$ that takes the value 1 in a and 0 otherwise. We will only be concerned with (upper and lower) expectations of *finitary functions*: functions that depend on the state of the process at a finite number of time instances. So if f is finitary, we can write $f = g(X_{0:k})$ for some $k \in \mathbb{N}_0$ and some $g \in \mathcal{L}(\mathcal{X}^{k+1})$. Note that finitary functions are bounded; this follows from their real-valuedness and the fact that \mathcal{X} is finite. The expectation of a finitary function $f(X_{0:k})$ conditional on some event $X_{0:\ell} = x_{0:\ell}$ simply reduces to a finite weighted sum:

$$E_P(f(X_{0:k}) | X_{0:\ell} = x_{0:\ell}) = \sum_{x_{\ell+1:k} \in \mathcal{X}^{k-\ell}} f(x_{0:k}) \prod_{i=\ell}^{k-1} P(X_{i+1} = x_{i+1} | X_{0:i} = x_{0:i}).$$

A particularly interesting case arises when studying stochastic processes that are described by a probability model P that satisfies

$$P(X_{k+1} = y | X_{0:k} = x_{0:k}) = P(X_{k+1} = y | X_k = x_k),$$

for all $k \in \mathbb{N}_0$, all $y \in \mathcal{X}$ and all $x_{0:k} \in \mathcal{X}^{k+1}$. This property, known as the *Markov property*, states that given the present state of the process the future behaviour of the process does not depend on its history. A process of this type is called a *Markov chain*. We moreover call it (*time*) *homogeneous* if additionally $P(X_{k+1} = y | X_k = x) = P(X_1 = y | X_0 = x)$, for all $k \in \mathbb{N}_0$ and all $x, y \in \mathcal{X}$. Hence, together with the assessments $P(X_0 = x_0)$, the dynamics of a homogeneous Markov chain are fully characterised by the probabilities $P(X_1 = y | X_0 = x)$. These probabilities are typically gathered in a *transition matrix* T ; a row-stochastic $|\mathcal{X}| \times |\mathcal{X}|$ matrix T that is defined by $T(x, y) := P(X_1 = y | X_0 = x)$ for all $x, y \in \mathcal{X}$. This matrix representation T is particularly convenient because it can be regarded as a linear operator from

$\mathcal{L}(\mathcal{X})$ to $\mathcal{L}(\mathcal{X})$, defined for any $k \in \mathbb{N}_0$, any $f \in \mathcal{L}(\mathcal{X})$ and any $x \in \mathcal{X}$ by

$$Tf(x) := \sum_{y \in \mathcal{X}} f(y)P(X_{k+1} = y | X_k = x) = E_P(f(X_{k+1}) | X_k = x).$$

More generally, we have that $E_P(f(X_{k+\ell}) | X_k = x) = T^\ell f(x)$ for all $k \in \mathbb{N}_0$, all $\ell \in \mathbb{N}_0$ and all $x \in \mathcal{X}$. Then, under some well-known accessibility conditions [8, Proposition 3], the expectation $T^\ell f(x)$ converges for increasing ℓ towards a constant $E_\infty(f)$ independently of the initial state x . If this is the case for all $f \in \mathcal{L}(\mathcal{X})$, the homogeneous Markov chain will have a steady-state distribution, represented by the limit expectation E_∞ , and we call the Markov chain *ergodic*. The expectation E_∞ is in particular also useful if we are interested in the limit behaviour of expected time averages. Indeed, let $\bar{f}_k(X_{\ell:k+k}) := 1/(k+1) \sum_{i=\ell}^{\ell+k} f(X_i)$ be the time average of some function $f \in \mathcal{L}(\mathcal{X})$ evaluated at the time instances ℓ through $k + \ell$. Then, according to [11, Theorem 38], the limit of the expected average $\lim_{k \rightarrow +\infty} E_P(\bar{f}_k(X_{0:k}))$ coincides with the limit expectation $E_\infty(f)$. One of the aims of this paper is to explore to which extent this remains true for imprecise Markov chains.

3 Imprecise Markov Chains

If the basic probabilities $P(X_{k+1}|X_{0:k} = x_{0:k})$ that describe a stochastic process are imprecise, in the sense that we only have partial information about them, then we can still model the process' dynamics by considering a set $\mathcal{T}_{x_{0:k}}$ of such probabilities, for all $k \in \mathbb{N}_0$ and all $x_{0:k} \in \mathcal{X}^{k+1}$. This set $\mathcal{T}_{x_{0:k}}$ is then interpreted as the set of all probability mass functions $P(X_{k+1}|X_{0:k} = x_{0:k})$ that we deem "plausible". We here consider the special case where the sets $\mathcal{T}_{x_{0:k}}$ satisfy a Markov property, meaning that $\mathcal{T}_{x_{0:k}} = \mathcal{T}_{x_k}$ for all $k \in \mathbb{N}_0$ and all $x_{0:k} \in \mathcal{X}^{k+1}$. Similar to the precise case, the sets \mathcal{T}_x , for all $x \in \mathcal{X}$, can be gathered into a single object: the set \mathcal{T} of all row stochastic $|\mathcal{X}| \times |\mathcal{X}|$ matrices T such that, for all $x \in \mathcal{X}$, the probability mass function $T(x, \cdot)$ is an element of \mathcal{T}_x . A set \mathcal{T} of transition matrices defined in this way is called *separately specified* [9]. For any such set \mathcal{T} , the corresponding *imprecise Markov chain under epistemic irrelevance* $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ [3] is the set of all (precise) probability models P such that $P(X_{k+1}|X_{0:k} = x_{0:k}) \in \mathcal{T}_{x_k}$ for all $k \in \mathbb{N}_0$ and all $x_{0:k} \in \mathcal{X}^{k+1}$. The values of the probabilities $P(X_0 = x_0)$ will be of no importance to us, because we will focus solely on (upper and lower) expectations conditional on the value of the initial state X_0 .

Clearly, an imprecise Markov chain $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ also contains non-homogeneous, and even non-Markovian processes. So the Markov property does in this case not apply to the individual probability assessments, but rather to the sets $\mathcal{T}_{x_{0:k}}$. The model $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ is therefore a generalisation of a traditional Markov chain where we allow for model uncertainty about, on the one hand, the mass functions $P(X_{k+1}|X_{0:k} = x_{0:k})$ and, on the other hand, about structural assumptions such as the Markov and time-homogeneity property. However, there are also types of

imprecise Markov chains that do impose some of these properties. For a given set \mathcal{T} , the *imprecise Markov chain under complete independence* $\mathcal{P}_{\mathcal{T}}^{\text{ci}}$ is the subset of $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ that contains all, possibly non-homogeneous, Markov chains in $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ [11]. The *imprecise Markov chain under repetition independence* $\mathcal{P}_{\mathcal{T}}^{\text{ri}}$ is the subset of $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$ containing all homogeneous Markov chains [11]. Henceforth, we let \mathcal{T} be some fixed, arbitrary set of transition matrices that is separately specified.

Now, for any probability model P in the imprecise Markov chain $\mathcal{P}_{\mathcal{T}}^{\text{ei}}$, we can again consider the corresponding expectation operator E_P . The *upper* and *lower expectation* are then respectively defined as the tightest upper and lower bound on this expectation:

$$\bar{E}_{\mathcal{T}}^{\text{ei}}(f|A) := \sup_{P \in \mathcal{P}_{\mathcal{T}}^{\text{ei}}} E_P(f|A) \quad \text{and} \quad \underline{E}_{\mathcal{T}}^{\text{ei}}(f|A) := \inf_{P \in \mathcal{P}_{\mathcal{T}}^{\text{ei}}} E_P(f|A),$$

for any finitary function f and any event A of the form $X_{0:k} = x_{0:k}$. The operators $\bar{E}_{\mathcal{T}}^{\text{ei}}$ and $\underline{E}_{\mathcal{T}}^{\text{ei}}$ are related by conjugacy, meaning that $\underline{E}_{\mathcal{T}}^{\text{ei}}(\cdot|\cdot) = -\bar{E}_{\mathcal{T}}^{\text{ei}}(-\cdot|\cdot)$, which allows us to focus on only one of them; upper expectations in our case. The lower expectation $\underline{E}_{\mathcal{T}}^{\text{ei}}(f|A)$ of a finitary function f can then simply be obtained by considering the upper expectation $-\bar{E}_{\mathcal{T}}^{\text{ei}}(-f|A)$.

In a similar way, we can define the upper expectations $\bar{E}_{\mathcal{T}}^{\text{ci}}$ and $\bar{E}_{\mathcal{T}}^{\text{ri}}$ and the lower expectations $\underline{E}_{\mathcal{T}}^{\text{ci}}$ and $\underline{E}_{\mathcal{T}}^{\text{ri}}$ as the tightest upper and lower bounds on the expectations corresponding to the models in $\mathcal{P}_{\mathcal{T}}^{\text{ci}}$ and $\mathcal{P}_{\mathcal{T}}^{\text{ri}}$, respectively. Since $\mathcal{P}_{\mathcal{T}}^{\text{ri}} \subseteq \mathcal{P}_{\mathcal{T}}^{\text{ci}} \subseteq \mathcal{P}_{\mathcal{T}}^{\text{ei}}$, we have that $\bar{E}_{\mathcal{T}}^{\text{ri}}(f|A) \leq \bar{E}_{\mathcal{T}}^{\text{ci}}(f|A) \leq \bar{E}_{\mathcal{T}}^{\text{ei}}(f|A)$ for any finitary function f and any event A of the form $X_{0:k} = x_{0:k}$.

As we have mentioned before, imprecise Markov chains generalise traditional Markov chains by incorporating different types of model uncertainty. The corresponding upper (and lower) expectations then allow us to make inferences that are robust with respect to this model uncertainty. For a more detailed discussion on the motivation for and interpretation behind these and other types of so-called imprecise probability models, we refer to [1, 5, 14].

Within the context of imprecise Markov chains, we will be specifically concerned with two types of inferences: the upper and lower expectation of a function at a single time instant, and the upper and lower expectation of the time average of a function. For imprecise Markov chains under epistemic irrelevance and under complete independence, both of these inferences coincide [11, Theorem 51 & Theorem 52]. For any $f \in \mathcal{L}(\mathcal{X})$ and any $x \in \mathcal{X}$, we will denote them by

$$\begin{aligned} \bar{E}_k(f|x) &= \bar{E}_{\mathcal{T}}^{\text{ei}}(f(X_k)|X_0 = x) = \bar{E}_{\mathcal{T}}^{\text{ci}}(f(X_k)|X_0 = x) \\ \text{and } \bar{E}_{\text{av},k}(f|x) &= \bar{E}_{\mathcal{T}}^{\text{ei}}(\bar{f}_k(X_{0:k})|X_0 = x) = \bar{E}_{\mathcal{T}}^{\text{ci}}(\bar{f}_k(X_{0:k})|X_0 = x), \end{aligned}$$

respectively, where the dependency on \mathcal{T} is implicit. The corresponding lower expectations can be obtained through conjugacy: $\underline{E}_k(f|x) = -\bar{E}_k(-f|x)$ and $\underline{E}_{\text{av},k}(f|x) = -\bar{E}_{\text{av},k}(-f|x)$ for all $f \in \mathcal{L}(\mathcal{X})$ and all $x \in \mathcal{X}$. In the remainder, we will omit imprecise Markov chains under repetition independence from the discussion. Generally speaking, this type of imprecise Markov chain is less studied within the field of imprecise probability because of its limited capacity to

incorporate model uncertainty. Indeed, it is simply a set of time-homogeneous precise Markov chains and therefore only allows for model uncertainty about the numerical values of the transition probabilities. Moreover, as far as we know, a characterisation for the ergodicity of such Markov chains—a central topic in this paper—is currently lacking. We therefore believe that this subject demands a separate discussion, which we defer to future work.

4 Transition Operators, Ergodicity and Weak Ergodicity

Inferences of the form $\bar{E}_k(f|x)$ were among the first ones to be thoroughly studied in imprecise Markov chains. Their study was fundamentally based on the observation that $\bar{E}_k(f|x)$ can be elegantly rewritten as the k -th iteration of the map $\bar{T}: \mathcal{L}(\mathcal{X}) \rightarrow \mathcal{L}(\mathcal{X})$ defined by

$$\bar{T}h(x) := \sup_{T \in \mathcal{T}} Th(x) = \sup_{T(x, \cdot) \in \mathcal{T}_x} \sum_{y \in \mathcal{X}} T(x, y)h(y),$$

for all $x \in \mathcal{X}$ and all $h \in \mathcal{L}(\mathcal{X})$. Concretely, $\bar{E}_k(f|x) = [\bar{T}^k f](x)$ for all $x \in \mathcal{X}$ and all $k \in \mathbb{N}_0$ [4, Theorem 3.1]. The map \bar{T} therefore plays a similar role as the transition matrix T in traditional Markov chains, which is why it is called the *upper transition operator* corresponding to the set \mathcal{T} .

In an analogous way, inferences of the form $\bar{E}_{av,k}(f|x)$ can be obtained as the k -th iteration of the map $\bar{T}_f: \mathcal{L}(\mathcal{X}) \rightarrow \mathcal{L}(\mathcal{X})$ defined by $\bar{T}_f h := f + \bar{T}h$ for all $h \in \mathcal{L}(\mathcal{X})$. In particular, if we let $\tilde{m}_{f,0} := f = \bar{T}_f(0)$ and

$$\tilde{m}_{f,k} := f + \bar{T}\tilde{m}_{f,k-1} = \bar{T}_f\tilde{m}_{f,k-1} \text{ for all } k \in \mathbb{N}, \tag{1}$$

then it follows from [11, Lemma 41] that $\bar{E}_{av,k}(f|x) = \frac{1}{k+1}\tilde{m}_{f,k}(x)$ for all $x \in \mathcal{X}$ and all $k \in \mathbb{N}_0$. Applying Eq. (1) repeatedly, we find that for all $x \in \mathcal{X}$:

$$\bar{E}_{av,k}(f|x) = \frac{1}{k+1}\tilde{m}_{f,k}(x) = \frac{1}{k+1}[\bar{T}_f^k \tilde{m}_{f,0}](x) = \frac{1}{k+1}[\bar{T}_f^{k+1}(0)](x). \tag{2}$$

The same formula can also be obtained as a special case of the results in [13].

These expressions for $\bar{E}_k(f|x)$ and $\bar{E}_{av,k}(f|x)$ in terms of the respective operators \bar{T} and \bar{T}_f are particularly useful when we aim to characterise the limit behaviour of these inferences. As will be elaborated on in the next section, there are conditions on \bar{T} that are necessary and sufficient for $\bar{E}_k(f|x)$ to converge to a limit value that does not depend on the process' initial state $x \in \mathcal{X}$. If this is the case for all $f \in \mathcal{L}(\mathcal{X})$, the imprecise Markov chain is called *ergodic* and we then denote the constant limit value by $\bar{E}_\infty(f) := \lim_{k \rightarrow +\infty} \bar{E}_k(f|x)$. Similarly, we call an imprecise Markov chain *weakly ergodic* if, for all $f \in \mathcal{L}(\mathcal{X})$, $\lim_{k \rightarrow +\infty} \bar{E}_{av,k}(f|x)$ exists and does not depend on the initial state x . For a weakly ergodic imprecise Markov chain, we denote the common limit value by $\bar{E}_{av,\infty}(f) := \lim_{k \rightarrow +\infty} \bar{E}_{av,k}(f|x)$. In contrast with standard ergodicity, weak ergodicity and, more generally, the limit behaviour of $\bar{E}_{av,k}(f|x)$, is almost entirely unexplored. The aim of this paper is to remedy this situation. The main

contribution will be a necessary and sufficient condition for an imprecise Markov chain to be weakly ergodic. As we will see, this condition is weaker than those needed for standard ergodicity, hence our choice of terminology. The following example shows that this difference already becomes apparent in the precise case.

Example 1. Let $\mathcal{X} = \{a, b\}$, consider any function $f = \begin{bmatrix} f_a \\ f_b \end{bmatrix} \in \mathcal{L}(\mathcal{X})$ and assume that \mathcal{T} consists of a single matrix $T = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Clearly, \bar{T} is not ergodic because $\bar{T}^{(2\ell+1)}f = T^{(2\ell+1)}f = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}f = \begin{bmatrix} f_b \\ f_a \end{bmatrix}$ and $\bar{T}^{(2\ell)}f = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}f = \begin{bmatrix} f_a \\ f_b \end{bmatrix}$ for all $\ell \in \mathbb{N}_0$. \bar{T} is weakly ergodic though, because

$$\bar{T}_f^{(2\ell)}(0) = \ell \begin{bmatrix} f_a + f_b \\ f_a + f_b \end{bmatrix} \quad \text{and} \quad \bar{T}_f^{(2\ell+1)}(0) = f + \bar{T} \bar{T}_f^{(2\ell)}(0) = f + \ell \begin{bmatrix} f_a + f_b \\ f_a + f_b \end{bmatrix},$$

for all $\ell \in \mathbb{N}_0$, which implies that $\bar{E}_{\text{av},\infty}(f) := \lim_{k \rightarrow +\infty} \bar{T}_f^k(0)/k = (f_a + f_b)/2$ exists. \diamond

Notably, even if an imprecise Markov chain is ergodic (and hence also weakly ergodic) and therefore both $\bar{E}_\infty(f)$ and $\bar{E}_{\text{av},\infty}(f)$ exist, these inferences will not necessarily coincide. This was first observed in an experimental setting [11, Section 7.6], but the differences that were observed there were marginal. The following example shows that these differences can in fact be very substantial.

Example 2. Let $\mathcal{X} = \{a, b\}$, let \mathcal{T}_a be the set of all probability mass functions on \mathcal{X} and let $\mathcal{T}_b := \{p\}$ for the probability mass function $p = (p_a, p_b) = (1, 0)$ that puts all mass in a . Then, for any $f = \begin{bmatrix} f_a \\ f_b \end{bmatrix} \in \mathcal{L}(\mathcal{X})$, we have that

$$\bar{T}f(x) = \begin{cases} \max f & \text{if } x = a; \\ f_a & \text{if } x = b, \end{cases} \quad \text{and} \quad \bar{T}^2f(x) = \begin{cases} \max \bar{T}f = \max f & \text{if } x = a; \\ \bar{T}f(a) = \max f & \text{if } x = b. \end{cases}$$

It follows that $\bar{T}^k f = \max f$ for all $k \geq 2$, so the limit upper expectation $\bar{E}_\infty(f)$ exists and is equal to $\max f$ for all $f \in \mathcal{L}(\mathcal{X})$. In particular, we have that $\bar{E}_\infty(\mathbb{I}_b) = 1$. On the other hand, we find that $\bar{T}_{\mathbb{I}_b}^{(2\ell)}(0) = \ell$ and $\bar{T}_{\mathbb{I}_b}^{(2\ell+1)}(0) = \mathbb{I}_b + \bar{T} \bar{T}_{\mathbb{I}_b}^{(2\ell)}(0) = \begin{bmatrix} \ell \\ \ell+1 \end{bmatrix}$ for all $\ell \in \mathbb{N}_0$. This implies that the upper expectation $\bar{E}_{\text{av},\infty}(\mathbb{I}_b) := \lim_{k \rightarrow +\infty} \bar{T}_{\mathbb{I}_b}^k(0)/k$ exists and is equal to $1/2$. This value differs significantly from the limit upper expectation $\bar{E}_\infty(\mathbb{I}_b) = 1$.

In fact, this result could have been expected simply by taking a closer look at the dynamics that correspond to \mathcal{T} . Indeed, it follows directly from \mathcal{T} that, if the system is in state b at some instant, then it will surely be in a at the next time instant. Hence, the system can only reside in state b for maximally half of the time, resulting in an upper expected average that converges to $1/2$. These underlying dynamics have little effect on the limit upper expectation $\bar{E}_\infty(\mathbb{I}_b)$ though, because it is only concerned with the upper expectation of \mathbb{I}_b evaluated at a single time instant. \diamond

Although we have used sets \mathcal{T} of transition matrices to define imprecise Markov chains, it should at this point be clear that, if we are interested in the inferences $\bar{E}_k(f|x)$ and $\bar{E}_{\text{av},k}(f|x)$ and their limit values, then it suffices to

specify \bar{T} . In fact, we will henceforth forget about \mathcal{T} and will assume that \bar{T} is a *coherent* upper transition operator on $\mathcal{L}(\mathcal{X})$, meaning that it is an operator from $\mathcal{L}(\mathcal{X})$ to $\mathcal{L}(\mathcal{X})$ that satisfies

- C1. $\min h \leq \bar{T}h \leq \max h$ [boundedness];
- C2. $\bar{T}(h + g) \leq \bar{T}h + \bar{T}g$ [sub-additivity];
- C3. $\bar{T}(\lambda h) = \lambda \bar{T}h$ [non-negative homogeneity],

for all $h, g \in \mathcal{L}(\mathcal{X})$ and all real $\lambda \geq 0$ [5, 14, 15], and we will regard $\bar{E}_k(f|x)$ and $\bar{E}_{\text{av},k}(f|x)$ as objects that correspond to \bar{T} . Our results and proofs will never rely on the fact that \bar{T} is derived from a set \mathcal{T} of transition matrices, but will only make use of C1–C3 and the following two properties that are implied by them [14, Section 2.6.1]:

- C4. $\bar{T}(\mu + h) = \mu + \bar{T}h$ [constant additivity];
- C5. if $h \leq g$ then $\bar{T}h \leq \bar{T}g$ [monotonicity],

for all $h, g \in \mathcal{L}(\mathcal{X})$ and all real μ . This can be done without loss of generality because an upper transition operator \bar{T} that is defined as an upper envelope of a set \mathcal{T} of transition matrices—as we did in Sect. 4—is always coherent [14, Theorem 2.6.3]. Since properties such as ergodicity and weak ergodicity can be completely characterised in terms of \bar{T} , we will henceforth simply say that \bar{T} itself is (weakly) ergodic, instead of saying that the corresponding imprecise Markov chain is.

5 Accessibility Relations and Topical Maps

To characterise ergodicity and weak ergodicity, we will make use of some well-known graph-theoretic concepts, suitably adapted to the imprecise Markov chain setting; we recall the following from [4] and [8]. The *upper accessibility graph* $\mathcal{G}(\bar{T})$ corresponding to \bar{T} is defined as the graph with vertices $x_1 \cdots x_n \in \mathcal{X}$, where $n := |\mathcal{X}|$, with an edge from x_i to x_j if $\bar{T}\mathbb{1}_{x_j}(x_i) > 0$. For any two vertices x_i and x_j , we say that x_j is *accessible* from x_i , denoted by $x_i \rightarrow x_j$, if $x_i = x_j$ or if there is a directed path from x_i to x_j , which means that there is a sequence $x_i = x'_0, x'_1, \dots, x'_m = x_j$ of vertices, with $m \in \mathbb{N}$, such that there is an edge from $x'_{\ell-1}$ to x'_ℓ for all $\ell \in \{1, \dots, m\}$. We say that two vertices x_i and x_j *communicate* and write $x_i \leftrightarrow x_j$ if both $x_i \rightarrow x_j$ and $x_j \rightarrow x_i$. The relation \leftrightarrow is an equivalence relation (reflexive, symmetric and transitive) and the equivalence classes are called *communication classes*. We call the graph $\mathcal{G}(\bar{T})$ *strongly connected* if any two vertices x_i and x_j in $\mathcal{G}(\bar{T})$ communicate, or equivalently, if \mathcal{X} itself is a communication class. Furthermore, we say that \bar{T} (or $\mathcal{G}(\bar{T})$) has a *top class* \mathcal{R} if

$$\mathcal{R} := \{x \in \mathcal{X} : y \rightarrow x \text{ for all } y \in \mathcal{X}\} \neq \emptyset.$$

So, if \bar{T} has a top class \mathcal{R} , then \mathcal{R} is accessible from any vertex in the graph $\mathcal{G}(\bar{T})$. As a fairly immediate consequence, it follows that \mathcal{R} is a communication

class that is *maximal* or *undominated*, meaning that $x \not\prec y$ for all $x \in \mathcal{R}$ and all $y \in \mathcal{R}^c$. In fact, it is the only such maximal communication class.

Having a top class is necessary for \bar{T} to be ergodic, but it is not sufficient. Sufficiency additionally requires that the top class \mathcal{R} satisfies [8, Proposition 3]:

- E1. $(\forall x \in \mathcal{R})(\exists k^* \in \mathbb{N})(\forall k \geq k^*) \min \bar{T}^k \mathbb{1}_x > 0$ [Regularity];
- E2. $(\forall x \in \mathcal{R}^c)(\exists k \in \mathbb{N}) \bar{T}^k \mathbb{1}_{\mathcal{R}^c}(x) < 1$ [Absorbing].

We will say that \bar{T} is *top class regular* (TCR) if it has a top class that is regular, and analogously for *top class absorbing* (TCA). Top class regularity represents aperiodic behaviour: it demands that there is some time instant $k^* \in \mathbb{N}$ such that all of the elements in the top class \mathcal{R} are accessible from each other in k steps, for any $k \geq k^*$. In the case of traditional Markov chains, top class regularity suffices as a necessary and sufficient condition for ergodicity [4, 10]. However, in the imprecise case, we need the additional condition of being top class absorbing, which ensures that the top class will eventually be reached. It requires that, if the process starts from any state $x \in \mathcal{R}^c$, the lower probability that it will ever transition to \mathcal{R} is strictly positive. We refer to [4] for more details. From a practical point of view, an important feature of both of these accessibility conditions is that they can be easily checked in practice [8].

The characterisation of ergodicity using (TCR) and (TCA) was strongly inspired by the observation that upper transition operators are part of a specific collection of order-preserving maps, called *topical maps*. These are maps $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ that satisfy

- T1. $F(\mu + h) = \mu + Fh$ [constant additivity];
- T2. if $h \leq g$ then $F(h) \leq F(g)$ [monotonicity],

for all $h, g \in \mathbb{R}^n$ and all $\mu \in \mathbb{R}$. To show this, we identify $\mathcal{L}(\mathcal{X})$ with the finite-dimensional linear space \mathbb{R}^n , with $n = |\mathcal{X}|$; this is clearly possible because both are isomorph. That every coherent upper transition operator is topical now follows trivially from C4 and C5. What is perhaps less obvious, but can be derived in an equally trivial way, is that the operator \bar{T}_f is also topical. This allows us to apply results for topical maps to \bar{T}_f in order to find necessary and sufficient conditions for weak ergodicity.

6 A Sufficient Condition for Weak Ergodicity

As a first step, we aim to find sufficient conditions for the existence of $\bar{E}_{av,\infty}(f)$. To that end, recall from Sect. 4 that if $\bar{E}_{av,\infty}(f)$ exists, it is equal to the limit $\lim_{k \rightarrow +\infty} \bar{T}_f^k(0)/k$. Then, since \bar{T}_f is topical, the following lemma implies that it is also equal to $\lim_{k \rightarrow +\infty} \bar{T}_f^k h/k$ for any $h \in \mathcal{L}(\mathcal{X})$.

Lemma 1 [7, Lemma 3.1]. *Consider any topical map $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$. If the limit $\lim_{k \rightarrow +\infty} F^k h/k$ exists for some $h \in \mathbb{R}^n$, then the limit exists for all $h \in \mathbb{R}^n$ and they are all equal.*

Hence, if $\lim_{k \rightarrow +\infty} \overline{T}_f^k h/k$ converges to a constant vector μ for some $h \in \mathcal{L}(\mathcal{X})$, then $\overline{E}_{av,\infty}(f)$ exists and is equal to μ . This condition is clearly satisfied if the map \overline{T}_f has an (additive) eigenvector $h \in \mathcal{L}(\mathcal{X})$, meaning that $\overline{T}_f^k h = h + k\mu$ for some $\mu \in \mathbb{R}$ and all $k \in \mathbb{N}_0$. In that case, we have that $\overline{E}_{av,\infty}(f) = \mu$, where μ is called the eigenvalue corresponding to h .

To find conditions that guarantee the existence of an eigenvector of \overline{T}_f , we will make use of results from [6] and [7]. There, accessibility graphs are defined in a slightly different way: for any topical map $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$, they let $\mathcal{G}'(F)$ be the graph with vertices v_1, \dots, v_n and an edge from v_i to v_j if $\lim_{\alpha \rightarrow +\infty} [F(\alpha \mathbb{I}_{v_j})](v_i) = +\infty$. Subsequently, for such a graph $\mathcal{G}'(F)$, the accessibility relation $\cdot \rightarrow \cdot$ and corresponding notions (e.g. ‘strongly connected’, ‘top class’, ...) are defined as in Sect. 5. If we identify the vertices v_1, \dots, v_n in $\mathcal{G}'(\overline{T})$ and $\mathcal{G}'(\overline{T}_f)$ with the different states x_1, \dots, x_n in \mathcal{X} , this can in particular be done for the topical maps \overline{T} and \overline{T}_f . The following results show that the resulting graphs coincide with the one defined in Sect. 5.

Lemma 2. *For any two vertices x and y in $\mathcal{G}'(\overline{T})$, there is an edge from x to y in $\mathcal{G}'(\overline{T})$ if and only if there is an edge from x to y in $\mathcal{G}'(\overline{T}_f)$.*

Proof. Consider any two vertices x and y in the graph $\mathcal{G}'(\overline{T})$. Then there is an edge from x to y if $\lim_{\alpha \rightarrow +\infty} [\overline{T}(\alpha \mathbb{I}_y)](x) = +\infty$. By non-negative homogeneity [C3], this is equivalent to the condition that $\lim_{\alpha \rightarrow +\infty} \alpha [\overline{T} \mathbb{I}_y](x) = +\infty$. Since moreover $0 \leq \overline{T} \mathbb{I}_y \leq 1$ by C1, this condition reduces to $\overline{T} \mathbb{I}_y(x) > 0$. \square

Corollary 1. *The graphs $\mathcal{G}'(\overline{T}_f)$, $\mathcal{G}'(\overline{T})$ and $\mathcal{G}(\overline{T})$ are identical.*

Proof. Lemma 2 implies that $\mathcal{G}'(\overline{T})$ and $\mathcal{G}'(\overline{T}_f)$ are identical. Moreover, that $\mathcal{G}'(\overline{T}_f)$ is equal to $\mathcal{G}'(\overline{T})$, follows straightforwardly from the definition of \overline{T}_f . \square

In principle, we could use this result to directly obtain the desired condition for the existence of an eigenvector from [6, Theorem 2]. However, [6, Theorem 2] is given in a multiplicative framework and would need to be reformulated in an additive framework in order to be applicable to the map \overline{T}_f ; see [6, Section 2.1]. This can be achieved with a bijective transformation, but we prefer to not do so because it would require too much extra terminology and notation. Instead, we will derive an additive variant of [6, Theorem 2] directly from [6, Theorem 9] and [6, Theorem 10].

The first result establishes that the existence of an eigenvector is equivalent to the fact that trajectories are bounded with respect to the Hilbert semi-norm $\|\cdot\|_H$, defined by $\|h\|_H := \max h - \min h$ for all $h \in \mathbb{R}^n$.

Theorem 1 [6, Theorem 9]. *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a topical map. Then F has an eigenvector in \mathbb{R}^n if and only if $\{\|F^k h\|_H : k \in \mathbb{N}\}$ is bounded for some (and hence all) $h \in \mathbb{R}^n$.*

That the boundedness of a single trajectory indeed implies the boundedness of all trajectories follows from the non-expansiveness of a topical map with respect

to the Hilbert semi-norm [6]. The second result that we need uses the notion of a *super-eigenspace*, defined for any topical map F and any $\mu \in \mathbb{R}$ as the set $S^\mu(F) := \{h \in \mathbb{R}^n : Fh \leq h + \mu\}$.

Theorem 2 [6, Theorem 10]. *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a topical map such that the associated graph $\mathcal{G}'(F)$ is strongly connected. Then all of the super-eigenspaces are bounded in the Hilbert semi-norm.*

Together, these theorems imply that any topical map $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ for which the graph $\mathcal{G}'(F)$ is strongly connected, has an eigenvector. The connection between both is provided by the fact that trajectories cannot leave an eigenspace. The following result formalises this.

Theorem 3. *Let $F: \mathbb{R}^n \rightarrow \mathbb{R}^n$ be a topical map such that the associated graph $\mathcal{G}'(F)$ is strongly connected. Then F has an eigenvector in \mathbb{R}^n .*

Proof. Consider any $h \in \mathbb{R}^n$ and any $\mu \in \mathbb{R}$ such that $\max(Fh - h) \leq \mu$. Then $Fh \leq h + \mu$, so $h \in S^\mu(F)$. Now notice that $F(Fh) \leq F(h + \mu) = Fh + \mu$ because of **T1** and **T2**, which implies that also $Fh \in S^\mu(F)$. In the same way, we can also deduce that $F^2h \in S^\mu(F)$ and, by repeating this argument, that the whole trajectory corresponding to h remains in $S^\mu(F)$. This trajectory is bounded because of Theorem 2, which by Theorem 1 guarantees the existence of an eigenvector. □

In particular, if $\mathcal{G}'(\overline{T}_f)$ is strongly connected then \overline{T}_f has an eigenvector, which on its turn implies the existence of $\overline{E}_{\text{av},\infty}(f)$ as explained earlier. If we combine this observation with Corollary 1, we obtain the following result.

Proposition 1. *An upper transition operator \overline{T} is weakly ergodic if the associated graph $\mathcal{G}(\overline{T})$ is strongly connected.*

Proof. Suppose that $\mathcal{G}(\overline{T})$ is strongly connected. Then, by Corollary 1, $\mathcal{G}'(\overline{T}_f)$ is also strongly connected. Hence, since \overline{T}_f is a topical map, Theorem 3 guarantees the existence of an eigenvector of \overline{T}_f . As explained in the beginning of this section, this implies by Lemma 1 that $\overline{E}_{\text{av},\infty}(f)$ exists, so we indeed find that \overline{T} is weakly ergodic. □

In the remainder of this paper, we will use the fact that \overline{T} is coherent—so not just topical—to strengthen this result. In particular, we will show that the condition of being strongly connected can be replaced by a weaker one: being top class absorbing. It will moreover turn out that this property is not only sufficient, but also necessary for weak ergodicity.

7 Necessary and Sufficient Condition for Weak Ergodicity

In order to gain some intuition about how to obtain a more general sufficient condition for weak ergodicity, consider the case where \overline{T} has a top class \mathcal{R} and the process’ initial state x is in \mathcal{R} . Since \mathcal{R} is a maximal communication class, the

process surely remains in \mathcal{R} and hence, it is to be expected that the time average of f will not be affected by the dynamics of the process outside \mathcal{R} . Moreover, the communication class \mathcal{R} is a strongly connected component, so one would expect that, due to Proposition 1, the upper expected time average $\bar{E}_{av,k}(f|x)$ converges to a constant that does not depend on the state $x \in \mathcal{R}$. Our intuition is formalised by the following proposition. Its proof, as well as those of the other statements in this section, are available in the appendix of [12].

Proposition 2. *For any maximal communication class \mathcal{S} and any $x \in \mathcal{S}$, the upper expectation $\bar{E}_{av,k}(f|x)$ is equal to $\bar{E}_{av,k}(f\mathbb{I}_{\mathcal{S}}|x)$ and converges to a limit value. This limit value is furthermore the same for all $x \in \mathcal{S}$.*

As a next step, we want to extend the domain of convergence of $\bar{E}_{av,k}(f|x)$ to all states $x \in \mathcal{X}$. To do so, we will impose the additional property of being top class absorbing (TCA), which, as explained in Sect. 5, demands that there is a strictly positive (lower) probability to reach the top class \mathcal{R} in a finite time period. Once in \mathcal{R} , the process can never escape \mathcal{R} though. One would therefore expect that as time progresses—as more of these finite time periods go by—this lower probability increases, implying that the process will eventually be in \mathcal{R} with practical certainty. Furthermore, if the process transitions from $x \in \mathcal{R}^c$ to a state $y \in \mathcal{R}$, then Proposition 2 guarantees that $\bar{E}_{av,k}(f|y)$ converges to a limit and that this limit value does not depend on the state y . Finally, since the average is taken over a growing time interval, the initial finite number of time steps that it took for the process to transition from x to y will not influence the time average of f in the limit. This leads us to suspect that $\bar{E}_{av,k}(f|x)$ converges to the same limit as $\bar{E}_{av,k}(f|y)$. Since this argument applies to any $x \in \mathcal{R}^c$, we are led to believe that \bar{T} is weakly ergodic. The following result confirms this.

Proposition 3. *Any \bar{T} that satisfies (TCA) is weakly ergodic.*

Conversely, suppose that \bar{T} does not satisfy (TCA). Then there are two possibilities: either there is no top class or there is a top class but it is not absorbing. If there is no top class, then it can be easily deduced that there are at least two maximal communication classes \mathcal{S}_1 and \mathcal{S}_2 . As discussed earlier, the process cannot escape the classes \mathcal{S}_1 and \mathcal{S}_2 once it has reached them. So if it starts in one of these communication classes, the process’ dynamics outside this class are irrelevant for the behaviour of the resulting time average. In particular, if we let f be the function that takes the constant value c_1 in \mathcal{S}_1 and c_2 in \mathcal{S}_2 , with $c_1 \neq c_2$, then we would expect that $\bar{E}_{av,k}(f|x) = c_1$ and $\bar{E}_{av,k}(f|y) = c_2$ for all $k \in \mathbb{N}_0$, any $x \in \mathcal{S}_1$ and any $y \in \mathcal{S}_2$. In fact, this can easily be formalised by means of Proposition 2. Hence, $\bar{E}_{av,\infty}(f|x) = c_1 \neq c_2 = \bar{E}_{av,\infty}(f|y)$, so the upper transition operator \bar{T} cannot be weakly ergodic.

Proposition 4. *Any weakly ergodic \bar{T} has a top class.*

Finally, suppose that there is a top class \mathcal{R} , but that it is not absorbing. This implies that there is an $x \in \mathcal{R}^c$ and a compatible precise model such that the process is guaranteed to remain in \mathcal{R}^c given that it started in x . If we now let

$f = \mathbb{I}_{\mathcal{R}^c}$, then conditional on the fact that $X_0 = x$, the expected time average of f corresponding to this precise model is equal to 1. Furthermore, since $f \leq 1$, no other process can yield a higher expected time average. The upper expected time average $\bar{E}_{av,k}(f|x)$ is therefore equal to 1 for all $k \in \mathbb{N}_0$. However, using Proposition 2, we can also show that $\bar{E}_{av,k}(f|y) = 0$ for any $y \in \mathcal{R}$ and all $k \in \mathbb{N}_0$. Hence, $\bar{E}_{av,\infty}(f|x) = 1 \neq 0 = \bar{E}_{av,\infty}(f|y)$, which precludes \bar{T} from being weakly ergodic.

Proposition 5. *Any weakly ergodic \bar{T} that has a top class satisfies (TCA).*

Together with Propositions 3 and 4, this allows us to conclude that (TCA) is a necessary and sufficient condition for weak ergodicity.

Theorem 4. *\bar{T} is weakly ergodic if and only if it is top class absorbing.*

8 Conclusion

The most important conclusion of our study of upper and lower expected time averages is its final result: that being top class absorbing is necessary and sufficient for weak ergodicity; a property that guarantees upper and lower expected time averages to converge to a limit value that does not depend on the process’ initial state. In comparison with standard ergodicity, which guarantees the existence of a limit upper and lower expectation, weak ergodicity thus requires less stringent conditions to be satisfied. We illustrated this difference in Example 1, where we considered a(n imprecise) Markov chain that satisfies (TCA) but not (TCR).

Apart from the fact that their existence is guaranteed under weaker conditions, the inferences $\bar{E}_{av,\infty}(f)$ are also able to provide us with more information about how time averages might behave, compared to limit expectations. To see why, recall Example 2, where the inference $\bar{E}_{av,\infty}(\mathbb{I}_b) = 1/2$ significantly differed from $\bar{E}_\infty(\mathbb{I}_b) = 1$. Clearly, the former was more representative for the limit behaviour of the time average of \mathbb{I}_b . As a consequence of [11, Lemma 57], a similar statement holds for general functions. In particular, it implies that $\bar{E}_{av,\infty}(f) \leq \bar{E}_\infty(f)$ for any function $f \in \mathcal{L}(\mathcal{X})$. Since both inferences are upper bounds, $\bar{E}_{av,\infty}(f)$ is therefore at least as informative as $\bar{E}_\infty(f)$.

In summary then, when it comes to characterising long-term time averages, there are two advantages that (limits of) upper and lower expected time averages have over conventional limit upper and lower expectations: they exist under weaker conditions and they are at least as (and sometimes much more) informative.

That said, there is also one important feature that limit upper and lower expectations have, but that is currently still lacking for upper and lower expected time averages: an (imprecise) point-wise ergodic theorem [2, Theorem 32]. For the limit upper and lower expectations of an ergodic imprecise Markov chain, this result states that

$$\underline{E}_\infty(f) \leq \liminf_{k \rightarrow +\infty} \bar{f}_k(X_{0:k}) \leq \limsup_{k \rightarrow +\infty} \bar{f}_k(X_{0:k}) \leq \bar{E}_\infty(f),$$

with lower probability one. In order for limit upper and lower expected time averages to be the undisputed quantities of interest when studying long-term time averages, a similar result would need to be obtained for weak ergodicity, where the role of $\bar{E}_\infty(f)$ and $\underline{E}_\infty(f) := -\bar{E}_\infty(-f)$ is taken over by $\bar{E}_{av,\infty}(f)$ and $\underline{E}_{av,\infty}(f) := -\bar{E}_{av,\infty}(-f)$, respectively. If such a result would hold, it would provide us with (strictly almost sure) bounds on the limit values attained by time averages that are not only more informative as the current ones, but also guaranteed to exist under weaker conditions. Whether such a result indeed holds is an open problem that we would like to address in our future work.

A second line of future research that we would like to pursue consists in studying the convergence of $\bar{E}_{av,k}(f|x)$ in general, without imposing that the limit value should not depend on x . We suspect that this kind of convergence will require no conditions at all.

References

1. Augustin, T., Coolen, F.P., de Cooman, G., Troffaes, M.C.: *Introduction to Imprecise Probabilities*. Wiley, Chichester (2014)
2. de Cooman, G., De Bock, J., Lopatzidis, S.: Imprecise stochastic processes in discrete time: global models, imprecise Markov chains, and ergodic theorems. *Int. J. Approximate Reasoning* **76**, 18–46 (2016)
3. 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)
4. de Cooman, G., Hermans, F., Quaeghebeur, E.: Imprecise Markov chains and their limit behaviour. *Prob. Eng. Inf. Sci.* **23**(4), 597–635 (2009)
5. de Cooman, G., Troffaes, M.C.: *Lower Previsions*. Wiley, Chichester (2014)
6. Gaubert, S., Gunawardena, J.: The Perron-Frobenius theorem for homogeneous, monotone functions. *Trans. Am. Math. Soc.* **356**(12), 4931–4950 (2004)
7. Gunawardena, J.: From max-plus algebra to nonexpansive mappings: a nonlinear theory for discrete event systems. *Theor. Comput. Sci.* **293**(1), 141–167 (2003)
8. Hermans, F., de Cooman, G.: Characterisation of ergodic upper transition operators. *Int. J. Approximate Reasoning* **53**(4), 573–583 (2012)
9. Hermans, F., Škulj, D.: Stochastic processes. In: Augustin, T., Coolen, F.P., de Cooman, G., Troffaes, M.C. (eds.) *Introduction to Imprecise Probabilities*, pp. 258–278. Wiley, Chichester (2014)
10. Kemeny, J.G., Snell, J.L.: *Finite Markov chains*. Undergraduate Text in Mathematics. Springer, New York (1976)
11. Lopatzidis, S.: *Robust Modelling and Optimisation in Stochastic Processes using Imprecise Probabilities, with an Application to Queueing Theory*. Ph.D. thesis, Ghent University (2017)
12. T’Joens, N., De Bock, J.: Limit behaviour of upper and lower expected time averages in discrete-time imprecise markov chains (2020). <http://arxiv.org/abs/2002.05661>
13. T’Joens, N., Krak, T., Bock, J.D., Cooman, G.: A Recursive Algorithm for Computing Inferences in Imprecise Markov Chains. In: Kern-Isberner, G., Ognjanović, Z. (eds.) *ECSQARU 2019. LNCS (LNAI)*, vol. 11726, pp. 455–465. Springer, Cham (2019). https://doi.org/10.1007/978-3-030-29765-7_38

14. Walley, P.: *Statistical Reasoning with Imprecise Probabilities*. Chapman and Hall, London (1991)
15. Williams, P.M.: Notes on conditional previsions. *Int. J. Approximate Reasoning* **44**(3), 366–383 (2007)