Queueing Syst (2016) 82:75–101
DOI 10.1007/s11134-015-9458-6

Robust queueing theory: an initial study using imprecise probabilities

Stavros Lopatatzidis[1] · **Jasper De Bock**[1] · **Gert de Cooman**[1] · **Stijn De Vuyst**[2] · **Joris Walraevens**[3]

Received: 24 November 2014 / Revised: 17 August 2015 / Published online: 28 August 2015

Abstract We study the robustness of performance predictions of discrete-time finite-capacity queues by applying the framework of imprecise probabilities. More concretely, we consider the *Geo/Geo/1/L* model with probabilities of arrival and departure that are no longer fixed, but are allowed to vary within given intervals. We distinguish between two concepts of independence in this framework, namely repetition independence and epistemic irrelevance. In the first approach, we assume the existence of time-homogeneous probabilities for arrival and departure, which leads us to consider a collection of stationary queues. In the second, the stationarity assumption is dropped and we allow the arrival and departure probabilities to vary from time point to time point; they may even depend on the complete history of queue lengths. We calculate bounds on the expected queue length, the probability of a particular queue length and the probability of turning on the server. For the expected queue length, both approaches coincide. For the other performance measures, we observe and dis-

✉ Stavros Lopatatzidis
stavros.lopatatzidis@UGent.be

Jasper De Bock
jasper.debock@UGent.be

Gert de Cooman
gert.decooman@UGent.be

Stijn De Vuyst
stijn.devuyst@UGent.be

Joris Walraevens
jw@telin.UGent.be

1 Ghent University, Technologiepark–Zwijnaarde 914, 9052 Zwijnaarde, Belgium

2 Ghent University, Technologiepark–Zwijnaarde 903, 9052 Zwijnaarde, Belgium

3 Ghent University, Sint-Pietersnieuwstraat 41, 9000 Ghent, Belgium

cuss various differences between the bounds obtained for these two approaches. One of our observations is that ergodicity may break down due to imprecision: bounds on expected time averages of certain functions on the state space are not necessarily equal to the bounds on the expectation of that function at random instants in a steady-state queue.

Keywords *Geo/Geo/1/L* · Imprecise probabilities · Time-homogeneous · Robustness · Performance measures · Discrete-time queueing

Mathematics Subject Classification 60K25 · 90B22 · 37A50 · 60-08 · 60G20 · 68M20

1 Introduction

The Achilles' heel of queueing models for the use of decision support, prediction and dimensioning in practical applications is that, usually, it is difficult to specify the parameters of these models in a way that is both exact and reliable. A stochastic model $\mathbf{Y} = f(\mathbf{X})$ assumes variables that are subject to randomness, and the convention is to capture the extent of this randomness by assigning precise probabilities to all relevant events. In particular, a specific probability distribution is put forward for the independent variables $\mathbf{X}$ of the model, after which the challenge is to obtain the probability distribution for $\mathbf{Y}$, either analytically, algorithmically, by estimation or otherwise. It is clear that the confidence we can put in the precise probability distribution for $\mathbf{Y}$ depends both on how skilful the model f is and on how confident we are in the probability distribution for $\mathbf{X}$. The goal of queueing theory in the last hundred years has been to analyse increasingly generalised, complicated and intricate models $f(\cdot)$. Although successful on its own, this approach often inspires overconfidence in its results, as the second part, gaining confidence in the probability model for $\mathbf{X}$, is not treated with the same attention and is frequently neglected altogether. Once numbers are produced, their dependence on the probability model for $\mathbf{X}$ is often all too easily forgotten.

In this paper, we do not question the relevance of the queueing model itself but rather discuss the consequences of representing the stochastic quantities by a set of distributions—a so-called *credal set*—rather than a single precise distribution. The variables can have any of the distributions in the set, without preference and without further specification. This allows for our beliefs about the input variables to be expressed more robustly. For example, when the distribution of the number of arrivals to a queue per time unit needs to be estimated from measured data, we put forward, instead of a single distribution (chosen based on some criterion such as, for instance, maximum likelihood), a credal set around this single distribution. This is a more prudent and arguably more honest approach which, to some extent, allows us to account for the unreliability of the estimation of the input variables in the queueing model.

This approach—expressing beliefs with credal sets—falls squarely within the theory of imprecise probabilities, a theory that was brought to a synthesis by Williams [33–35] and Walley [27,29], but goes back to Boole [4] and Keynes [16], with crucial contributions by quite a number of statisticians and philosophers [17,21,24]. It looks at (conservative) probabilistic inference in the following way: how can we calculate as efficiently as possible the consequences—in the sense of most conservative tightest bounds—of making certain probability assessments. Since Walley's [27] seminal work, significant advances have been made in the field.[1] On the side of model calibration, important work was done on coherently translating expert opinions or measurements of data into imprecise probability models (elicitation) [3,18,28]. Stochastic processes with imprecise models can also be described, in particular Markov chains with finite state spaces and imprecise local models (i.e. imprecise transition probabilities). Algorithms were constructed to calculate upper and lower expectations for functions of the state after n steps [11]. Generalisations of using (Viterbi) filters for state estimation in hidden Markov models exist as well [12]. In all of these applications, the crux of the matter is that for imprecise probabilities the concept of independence of two random variables is no longer self-evident and splits up into several notions. Amongst these, we explicitly mention *repetition independence* and *epistemic irrelevance*, both of which will be explained and used further on.

In this contribution, our aim is to demonstrate the practical use of imprecise probabilities for queueing analysis. For the sake of demonstration, the queueing model itself was admittedly kept as simple as possible: a discrete-time finite-capacity queue with Bernoulli arrivals and departures. In the classic framework of precise probabilities, such a model poses no problems and all performance indicators can be obtained easily. Even so, this simplicity serves all the better to exhibit the implications of imprecision in the local models. For the Bernoulli variables involved, we specify a credal set by stating a lower and upper probability for both the event of an arrival and for the event of a departure during a slot. Performance of the queue is assessed in terms of (transient and stationary) queue length and the server switching from idle to busy. Our main conclusions are that

1. upper and lower expectations of the performance highly depend on the assumed notion of independence;
2. ergodicity may break down under imprecision, and expected time averages are no longer equivalent to ensemble averages.

For queueing theory, the second conclusion is particularly important. In the conventional case of precise Markov chains, if the queue is irreducible, aperiodic and positive recurrent, it is called ergodic and the idea of 'mean value' of a function on the state space is unambiguous. The mean can either be seen as the expected time average over a very long sample path, or it can be the expected value of that function at an arbitrary time point. In the case of imprecision, these may be entirely different performance metrics.

The imprecise probability framework is not the only attempt at a more robust description of uncertainty. There is, for instance, also the Dempster–Shafer theory [22], interval probabilities [30], fuzzy set theory [37] and others. To some extent, one can also characterise the influence of parameter uncertainty on performance predictions strictly from within the framework of precise probabilities, for example, by studying

[1] A good overview can be obtained by perusing the proceedings of the biennial ISIPTA conferences at www.sipta.org.

perturbations of the Markov chain underlying the queueing model [2,6]. For small enough ε, the stationary vector $\boldsymbol{\pi}_\varepsilon$ of a perturbated chain with transition matrix (or generator) $\mathbf{Q}_\varepsilon = \mathbf{Q} + \varepsilon\mathbf{Q}'$ is a power series in ε with coefficients that can be calculated from $\mathbf{Q}'$ and both the stationary vector and the deviation matrix of the original chain $\mathbf{Q}$ [20]. Such methods provide a sensitivity analysis [5,19,31,36]. Although very useful in many cases, they cover only small perturbations in one direction (that of $\mathbf{Q}'$) and are generally not able to relate credal sets of input variable distributions to performance bounds in case of epistemic irrelevance, as we will discuss further on. To the best of our knowledge, the imprecise probability approach has not been applied to queueing models before.

2 Concepts of independence for a queue with imprecise local models

We consider a discrete-time queue that is in exactly one of several states at each time step n. For a particular time step, we refer to the imprecise probabilities (i.e. the credal sets) of the transitions out of each state as the *local model*. For a precise Markov chain, the local models are independent of each other and the question stands how this independence transfers to imprecise local models. As mentioned previously, we consider two different approaches in constructing a global imprecise probability model (about the states at all time points) from the local models (describing the initial state and transitions between states). We apply both approaches to a discrete-time single-arrival single-departure queueing system, *Geo/Geo/1/L*, and study the effect on the system's performance of the local models, i.e. the upper and lower probabilities for both an arrival and a departure during a slot. In the remainder, we aptly, though somewhat inaccurately, refer to a particular distribution in a credal set of a local model as the 'parameters of the model'. Since a local model involves Bernoulli variables, there is no ambiguity between 'parameter', 'distribution' and 'probability'.

The first approach is the one most closely related to conventional sensitivity analysis. It is based on the notion of *repetition independence* [8] (RI) in the theory of imprecise probabilities. Basically, we assume the parameters in the local models at different times lie within a certain interval, but nevertheless be identical copies of each other. We are interested in the robustness of the model against this imprecision. This is quantified by calculating lower and upper expectations of certain functions of interest of the output stochastic variables, i.e. the minimum and maximum expectation that can be obtained by varying the parameters within the given ranges. The advantage of this approach is the direct relation to sensitivity analysis, since all possible underlying global models are precise Markov chains. A major disadvantage, however, is the complexity of the model for a huge number of uncertain parameters and/or a function that depends on many variables. Finding the upper expectation, for instance, is, in essence, an optimisation problem: calculate the maximum expected value of a given function for parameters in their respective ranges.

The second approach is more general. In the theory of imprecise probabilities, it is referred to as *epistemic irrelevance* [10,27] (EI). We still assume that the parameters lie within some intervals, but in contrast with what is required in the RI approach, the parameters in the local models at different time points no longer have to be identical. Perhaps it is best to illustrate this difference with an example. Assume a *Geo/Geo/1/L* queue with an imprecise arrival probability, i.e. a range of possible values of the arrival probability is given. With RI, we assume that we do not know the arrival rate precisely, but we know (or assume) it is the same at each instant. With the second approach, we *only* assume that the arrival rate lies in the specified range at each instant. It does not have to be the same at each instant. In fact, it can depend on the complete history of the system. Now the underlying global models no longer have to be Markov chains, in fact they could be *any* process on the state space with the only requirement that at each instant n, and for any given history, the probabilities of arrival and departure lie within the specified ranges. The advantage of the EI approach is that efficient algorithms exist to calculate the corresponding lower and upper expectations [10]. These algorithms are basically linear in time. No such algorithms exist for the first approach.

Clearly, the models that are included in the RI approach are included in the EI approach as well (if the imprecise parameters lie in the same intervals). Therefore, the latter approach typically leads to more conservative results, with larger upper expectations and smaller lower expectations. However, in some cases, as we will numerically and formally show, they lead to the same upper and lower expectations. A conclusion that can then be drawn for EI is that, although extra imprecision is included in the model, the worst-case scenario was already included in the RI approach. However, as mentioned above, it is easier to do calculations assuming EI.

On the EI approach, however, some well-known results start to break down, as we shall see further on. For instance, in the precise case, the stationary distribution of a given output variable can be found from the transient distribution in two ways, namely, via (i) taking the limit for time going to infinity or (ii) taking the average over time. As already mentioned, we will show that in the imprecise case, for the EI approach, this is no longer necessarily true. We provide some discussion on this and argue that both approaches can be interesting from a practical point of view.

In the next sections, we discuss calculating expectations and probabilities in the *Geo/Geo/1/L* queue. First, in Sect. 3, we consider the probabilities of arrival and departure to be precisely known. We add imprecision in Sects. 4 and 5, meaning that arrival and departure probabilities lie within an interval, and we are interested in finding minima and maxima of expectations and probabilities under both RI and EI. Finally, in Sect. 6, we run various experiments using both approaches, compare the results and prove some properties.

3 Model description

We focus on a *Geo/Geo/1/L* queue, a simple but quite common example in queueing theory. Following the notation, our system is a single-server queue of maximum capacity (length) L, where $L \in \mathbb{N}$ and $\mathbb{N}$ is the set of all natural numbers, excluding zero. We denote by $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$ the set of all non-negative integers. Arrivals and departures occur according to geometric distributions, which we assume to be (stochastically) independent. The queue content is observed during consecutive slots and we assume that on slot boundaries a departure occurs prior to an arrival, a convention called either Departures-First (DF) or Early Arrival System (EAS) [7]. We choose this priority in order to avoid zero time servicing when the queue is empty. Furthermore, we assume that an item stays in the queue until served and that the service discipline is work-conserving.

In the next section, we discuss ways to calculate expectations and probabilities in a *Geo/Geo/1/L* queue, when the probabilities of arrival and departure are precisely known. In such systems, it is typical to use the so-called *balance equations* [13], in order to find the distribution of the queue lengths when the system is in a steady state. For our purposes, we also require the transient solution [1], which is why we also show another method for calculating probabilities or expectations in such queues. It has the added advantage that it paves the way for our treatment of the EI approach.

3.1 The precise case

We start with precise parameters in our model, as the approaches that will be described in the next sections are extensions of the precise case. Assume that, at any time point and given any possible queue length, we have a probability of arrival and a probability of departure, which are denoted by a and d, respectively. At each time point $n \in \mathbb{N}$, the queue length can then be represented by a random variable X_n taking values in a set $\mathscr{X} = \{0, \ldots, L\}$. The probability of any length at time point $n + 1$ given that $X_n = x_n$, with $x_n \in \mathscr{X}$, is time-homogeneous and depends only on the value x_n.

Hence, the conditional probability distribution for X_{n+1} can be fully described by x_n and by the probabilities of arrival and departure. The conditional probability mass function for this local model is therefore denoted by $q(\cdot|x_n, a, d)$. If $x_n = 0$, then

$$q(x_{n+1}|0, a, d) := \begin{cases} 1 - a & \text{if } x_{n+1} = 0 \\ a & \text{if } x_{n+1} = 1 \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } x_{n+1} \in \mathscr{X}. \tag{1}$$

If $0 < x_n < L$, then for all $x_{n+1} \in \mathscr{X}$:

$$q(x_{n+1}|x_n, a, d) := \begin{cases} d(1 - a) & \text{if } x_{n+1} = x_n - 1 \\ da + (1 - d)(1 - a) & \text{if } x_{n+1} = x_n \\ (1 - d)a & \text{if } x_{n+1} = x_n + 1 \\ 0 & \text{otherwise.} \end{cases} \tag{2}$$

Finally, if $x_n = L$, then

$$q(x_{n+1}|L, a, d) := \begin{cases} d(1 - a) & \text{if } x_{n+1} = L - 1 \\ 1 - d(1 - a) & \text{if } x_{n+1} = L \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } x_{n+1} \in \mathscr{X}. \tag{3}$$

Figure 1 depicts these local probabilities at any time point for various queue lengths. As will be readily inferred from this figure, we can model the *Geo/Geo/1/L* queue as a tree, where at each node, we have conditional probabilities for the transition to a child node. This yields a so-called probability tree [23].

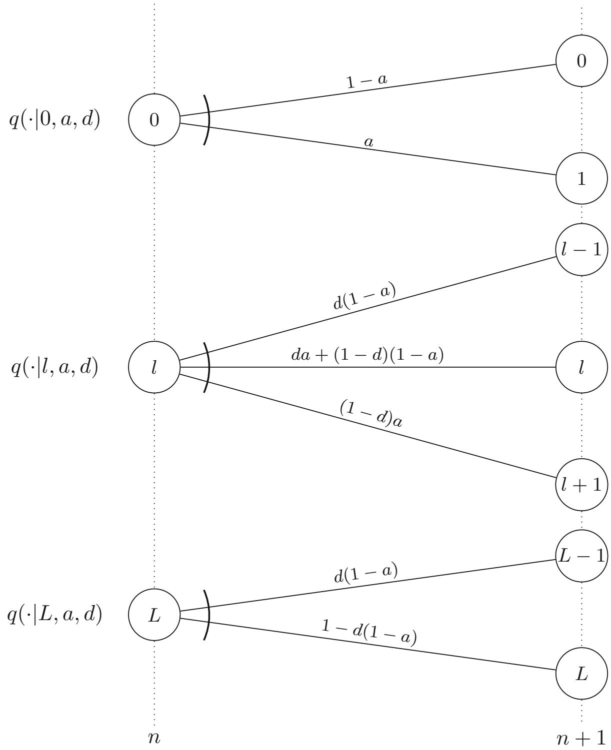

Fig. 1 At any time point n and for each possible queue length $x_n \in \{0, \ldots, L\}$, we have a conditional probability mass function $q(\cdot|x_n, a, d)$, where a and d are the time-homogeneous probabilities of arrival and departure, respectively

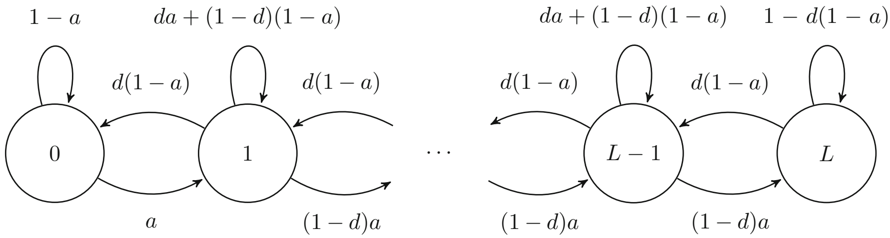

Fig. 2 All the possible transitions from one state to another when the probability of arrival is a and the one for departure is d

In a *Geo/Geo/1/L* queue, we have analytical formulas for calculating the probability $\lim_{n\to\infty} \Pr[X_n = k]$, $\forall k \in \mathscr{X}$. Since we have time-homogeneous probabilities a and d for arrival and departure, respectively, our queue has the birth–death chain representation depicted in Fig. 2. Let X be the length of the queue in steady state, so $X = \lim_{n\to\infty} X_n$; then the balance equations are

$$a\Pr[X = 0] = d(1-a)\Pr[X = 1]$$

and

$$(1-d)a\Pr[X=k-1] = d(1-a)\Pr[X=k]\ ,\ \text{ for } 2 \leq k \leq L.$$

By combining these equations with the unitary constraint and solving the resulting system of equations, we find—if $0 < a < 1$ and $0 < d < 1$—that

$$\Pr[X=0] = \frac{d-a}{d - \frac{(1-d)^L a^{L+1}}{(d(1-a))^L}}. \tag{4}$$

and

$$\Pr[X=k] = \frac{(1-d)^{k-1}a^k}{(d(1-a))^k}\Pr[X=0]\ ,\ \text{ for } 1 \leq k \leq L. \tag{5}$$

Alternatively, we can use the law of iterated expectations to compute expectations of functions, and regard probabilities as special cases. Although this approach is more involved, it has the advantage that it does not require us to restrict attention to the steady state. Also, as we will see further on, it is very similar to the imprecise approach with EI. For the sake of notational convenience, for any finite set $\mathscr{Y}$, we denote the linear space of all real-valued functions f on $\mathscr{Y}$ by $\mathscr{L}(\mathscr{Y})$.

For any function $f \in \mathscr{L}(\mathscr{X}^n)$, with $n \in \mathbb{N}$ and $\mathscr{X}^n := \overbrace{\mathscr{X} \times \cdots \times \mathscr{X}}^{n}$, we define its expectation as

$$E_{1:n}(f) := E(f(X_{1:n})) = \sum_{x_{1:n} \in \mathscr{X}^n} f(x_{1:n})p(x_{1:n}), \tag{6}$$

where $X_{1:n} := X_1, \ldots, X_n$ is a sequence of queue lengths taking values in $\mathscr{X}^n$, $x_{1:n} := x_1, \ldots, x_n$ is an realization of $X_{1:n}$ and

$$p(x_{1:n}) := \Pr[X_1 = x_1, \ldots, X_n = x_n].$$

In many cases, $p(x_{1:n})$ is not directly available and all we have to start from are (i) the conditional probabilities

$$p(x_{i+1}|x_{1:i}) := \Pr[X_{i+1} = x_{i+1}|X_1 = x_1, \ldots, X_i = x_i]$$

and the corresponding conditional expectation operators $E_{i+1}(\cdot|x_{1:i})$, defined for all $g \in \mathscr{L}(\mathscr{X}^{i+1})$ by

$$E_{i+1}(g|x_{1:i}) := \sum_{x_{i+1} \in \mathscr{X}} g(x_{1:i+1})p(x_{i+1}|x_{1:i});$$

and (ii) the initial probabilities

$$p_{X_1}(x_1) := \Pr[X_1 = x_1]$$

with corresponding expectation operators E_1, as defined for all $g \in \mathscr{L}(\mathscr{X})$ by

$$E_1(g) := \sum_{x_1 \in \mathscr{X}} g(x_1)p(x_1).$$

In those cases, since $p(x_{1:n}) = p(x_1) \prod_{i=1}^{n-1} p(x_{i+1}|x_{1:i})$, it follows that

$$\begin{aligned} E_{1:n}(f) &= \sum_{x_{1:n} \in \mathscr{X}^n} f(x_{1:n})p(x_1) \prod_{i=1}^{n-1} p(x_{i+1}|x_{1:i}) \\ &= \sum_{x_1 \in \mathscr{X}} p(x_1) \sum_{x_2 \in \mathscr{X}} p(x_2|x_1) \cdots \sum_{x_n \in \mathscr{X}} f(x_{1:n})p(x_n|x_{n-1}) \\ &= E_1(E_2(E_3(\ldots E_n(f|X_{1:n-1}) \ldots |X_{1:2})|X_1)), \end{aligned} \tag{7}$$

a version of the law of total probability that is also known as the *law of iterated expectations*.

Similarly, for any $h \in \mathscr{L}(\mathscr{X})$, we let $E_n(h) := E(h(X_n))$. This expected value can again be computed by means of the law of iterated expectations; it suffices to particularise $f(X_{1:n})$ to $h(X_n)$:

$$E_n(h) = E_1(E_2(E_3(\ \ldots\ E_n(g(X_n)|X_{1:n-1})\ \ldots\ |X_{1:2})|X_1)) \tag{8}$$

Although these equations apply generally, we are mainly interested in our specific case, where, for all $i \in \mathbb{N}$ and $x_{1:i} \in \mathscr{X}^i$, the conditional probability mass function

$$p(\cdot|x_{1:i}) = q(\cdot|x_i, a, d) \tag{9}$$

only depends on x_i, since a and d are considered to be fixed. In that case, $E_{i+1}(\cdot|X_{1:i})$ can be identified with a map $E_{i+1}(\cdot|X_i)$ from $\mathscr{L}(\mathscr{X})$ to $\mathscr{L}(\mathscr{X})$ and evaluating Eq. (8) therefore has a computational complexity that is linear in n. For any fixed a and d, we will denote the models in probability trees of the form (9) by $p_{a,d}$ and use the notation $E_{1:n}^{p_{a,d}}$ and $E_n^{p_{a,d}}$ to refer to the corresponding expectation operators.

3.2 Introduction to the imprecise case

In the next sections, we will add imprecision to our assessments, meaning that we assume that arrival and departure probabilities lie within an interval. Instead of a single pair a and d, we have intervals of arrival and departure probabilities, denoted by $[\underline{a}, \overline{a}]$ and $[\underline{d}, \overline{d}]$, respectively. The local models are then described as follows. For any $i \in \mathbb{N}$ and any $x_{1:i} \in \mathscr{X}^i$, we assume that $p(\cdot|x_{1:i})$ belongs to the set

$$\mathscr{Q}_{x_i} := \left\{q(\cdot|x_i, a, d)\colon a \in [\underline{a}, \overline{a}], d \in [\underline{d}, \overline{d}]\right\}. \tag{10}$$

However, we do not specify which value $p(\cdot|x_{1:i})$ takes within this set. Similarly, we assume that p_{X_1} is known to belong to some closed set $\mathscr{Q}_1$ of probability mass functions, but we do not specify which value it takes within this set. Since we do not specify exact values for these local probabilities, we cannot provide exact values for derived global probabilities or expectations either. Instead, we will assess the robustness of a probability or an expected value with respect to variations of $p(\cdot|x_{1:i})$ within $\mathscr{Q}_{x_i}$—and $p_{X_1} \in \mathscr{Q}_1$—by computing tight lower and upper bounds for them. Since probabilities are just special cases of expectations, we focus on expectations, leading us to introduce the notion of a *lower* and *upper expectation*, denoted by $\underline{E}$ and $\overline{E}$, respectively.

Depending on how the probability mass functions $p(\cdot|x_{1:i})$, $x_{1:i} \in \mathscr{X}^i$, are allowed to be chosen from the sets $\mathscr{Q}_{x_i}$, $x_i \in \mathscr{X}$, different lower and upper expectations are obtained. We distinguish between two extreme cases: repetition independence and epistemic irrelevance.

4 The RI approach: repetition independence

We first introduce the case of repetition independence (RI). Under the RI approach, we assume that there are time-homogeneous $a \in [\underline{a}, \overline{a}]$ and $d \in [\underline{d}, \overline{d}]$ such that $p(\cdot|x_{1:i}) = q(\cdot|x_i, a, d) \in \mathscr{Q}_{x_i}$ for all $i \in \mathbb{N}$ and $x_{1:i} \in \mathscr{X}^i$. p_{X_1} is taken to be an element of $\mathscr{Q}_1$. As before, generic probability trees of this form are denoted by $p_{a,d}$ and the corresponding expectation operators are denoted by $E_{1:n}^{p_{a,d}}$ and $E_n^{p_{a,d}}$. We use $\mathscr{T}^{\mathrm{RI}}$ to refer to the set of all probability trees $p_{a,d}$. The lower and upper expectations that are obtained by optimising over the elements of this set are denoted by $\underline{E}^{\mathrm{RI}}$ and $\overline{E}^{\mathrm{RI}}$, respectively. This type of optimisation is in line with the classical approach to sensitivity analysis in such systems.

For example, for any time point n and any $f \in \mathscr{L}(\mathscr{X}^n)$, we define

$$\begin{aligned} \underline{E}_{1:n}^{\mathrm{RI}}(f) &:= \min\left\{E_{1:n}^{p_{a,d}}(f) \colon p_{a,d} \in \mathscr{T}^{\mathrm{RI}}\right\}, \\ \overline{E}_{1:n}^{\mathrm{RI}}(f) &:= \max\left\{E_{1:n}^{p_{a,d}}(f) \colon p_{a,d} \in \mathscr{T}^{\mathrm{RI}}\right\}, \end{aligned} \tag{11}$$

with $p_{a,d}$ as in Sect. 3.1. Similarly, for any time point n and any $h \in \mathscr{L}(\mathscr{X})$, we let

$$\begin{aligned} \underline{E}_{n}^{\mathrm{RI}}(h) &:= \min\left\{E_{n}^{p_{a,d}}(h) \colon p_{a,d} \in \mathscr{T}^{\mathrm{RI}}\right\}, \\ \overline{E}_{n}^{\mathrm{RI}}(h) &:= \max\left\{E_{n}^{p_{a,d}}(h) \colon p_{a,d} \in \mathscr{T}^{\mathrm{RI}}\right\}. \end{aligned} \tag{12}$$

Other types of lower and upper expectations—for example, conditional ones—are defined analogously.

If $0 < a < 1$ and $0 < d < 1$, it follows from the discussion in the previous section that

$$\lim_{n\to\infty} \underline{E}_n^{\mathrm{RI}}(h) = \min\left\{\sum_{x\in\mathscr{X}} h(x)\Pr[X = x] \colon a \in [\underline{a}, \overline{a}], d \in [\underline{d}, \overline{d}]\right\}, \tag{13}$$

where the stationary probabilities $\Pr[X = x]$ are given by Eqs. (4) and (5). A similar expression holds for $\lim_{n\to\infty} \overline{E}_n^{\mathrm{RI}}(h)$.

Computationally, we are able to evaluate Eq. (13) by means of brute force solvers. However, for Eqs. (11) and (12), for large n, this is no longer tractable. In those cases, we can approximate the lower (or upper) expectation in (11) and (12) by selecting a finite number of probabilities a in $[\underline{a}, \overline{a}]$ and d in $[\underline{d}, \overline{d}]$ and calculating the expectations for all possible combinations. For a single such combination of a and d, we compute the corresponding expectation by applying the law of iterated expectations—see Eqs. (7) and (8)—with $p(x_{i+1}|x_{1:i})$ as in Eq. (9).

Unfortunately, this approximation method is not guaranteed to return the global minimum (or maximum) expectation of the function under study.

5 The EI approach: epistemic irrelevance

The case of epistemic irrelevance (EI) is more general than that of RI. In the EI approach, we do not assume time-homogeneous probabilities of arrival and departure. Instead, for all $i \in \mathbb{N}$ and all $x_{1:i} \in \mathscr{X}^i$, we consider (possibly different) probabilities of arrival and departure $a_{x_{1:i}} \in [\underline{a}, \overline{a}]$ and $d_{x_{1:i}} \in [\underline{d}, \overline{d}]$ and let $p(\cdot|x_{1:i}) = q\left(\cdot|x_i, a_{x_{1:i}}, d_{x_{1:i}}\right) \in \mathscr{Q}_{x_i}$. p_{X_1} is taken to be an element of $\mathscr{Q}_1$. Generic probability trees of this form are denoted by $p_{A,D}$, and the corresponding expectation operators are denoted by $E_{1:n}^{p_{A,D}}$ and $E_n^{p_{A,D}}$. We use $\mathscr{T}^{\mathrm{EI}}$ to refer to the set of all probability trees $p_{A,D}$. The lower and upper expectations that are obtained by optimising over the elements of this set are denoted by $\underline{E}^{\mathrm{EI}}$ and $\overline{E}^{\mathrm{EI}}$, respectively.

For example, for any time point n and any $f \in \mathscr{L}(\mathscr{X}^n)$, we define

$$\begin{aligned}\underline{E}_{1:n}^{\mathrm{EI}}(f) &:= \min\left\{E_{1:n}^{p_{A,D}}(f)\colon p_{A,D} \in \mathscr{T}^{\mathrm{EI}}\right\},\\ \overline{E}_{1:n}^{\mathrm{EI}}(f) &:= \max\left\{E_{1:n}^{p_{A,D}}(f)\colon p_{A,D} \in \mathscr{T}^{\mathrm{EI}}\right\}.\end{aligned} \tag{14}$$

Similarly, for any time point n and any $h \in \mathscr{L}(\mathscr{X})$, we let

$$\begin{aligned}\underline{E}_{n}^{\mathrm{EI}}(h) &:= \min\left\{E_{n}^{p_{A,D}}(h)\colon p_{A,D} \in \mathscr{T}^{\mathrm{EI}}\right\},\\ \overline{E}_{n}^{\mathrm{EI}}(h) &:= \max\left\{E_{n}^{p_{A,D}}(h)\colon p_{A,D} \in \mathscr{T}^{\mathrm{EI}}\right\}.\end{aligned} \tag{15}$$

Other types of lower and upper expectations—for example, conditional ones—are defined analogously.

Interestingly, these lower and upper expectations can be computed by means of a generalised version of the law of iterated expectations. It makes use of the (non-linear) operator $\underline{Q}_1$, defined for all $f \in \mathscr{L}(\mathscr{X})$ by

$$\underline{Q}_1(f) = \min\left\{\sum_{x_1\in\mathscr{X}} p(x_1)f(x_1)\colon p_{X_1} \in \mathscr{Q}_1\right\},$$

and, for all $n \in \mathbb{N}$, the (non-linear) projection operator $\underline{Q}_{n+1}(\cdot|X_{1:n})\colon \mathscr{L}(\mathscr{X}^{n+1}) \to \mathscr{L}(\mathscr{X}^n)$, defined for all $f \in \mathscr{L}(\mathscr{X}^{n+1})$ and $x_{1:n} \in \mathscr{X}^n$ by

$$\begin{aligned}\underline{Q}_{n+1}(f|x_{1:n}) &= \min\left\{\sum_{x_{n+1}\in\mathscr{X}} q(x_{n+1}|x_n, a, d) f(x_{1:n+1})\colon a \in [\underline{a}, \overline{a}], d \in [\underline{d}, \overline{d}]\right\}\\ &= \min\left\{\sum_{x_{n+1}\in\mathscr{X}} q(x_{n+1}|x_n, a, d) f(x_{1:n+1})\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\},\end{aligned} \tag{16}$$

where the last equality follows because, if we fix either of the parameters a or d, then $\sum_{x_{n+1}\in\mathscr{X}} q(x_{n+1}|x_n, a, d) f(x_{1:n+1})$ becomes a linear function of the other parameter.

Theorem 1 ([11, Theorem 3.2] Concatenation Formula) *For any $n \in \mathbb{N}$ and $m \in \mathbb{N}_0$, with $m < n$, and any $f \in \mathscr{L}(\mathscr{X}^n)$:*

$$\underline{E}^{EI}_{1:n}(f|X_{1:m}) = \underline{Q}_{m+1}(\underline{Q}_{m+2}(\dots \underline{Q}_n(f|X_{1:n-1})\dots|X_{1:m+1})|X_{1:m}).$$

This result is essentially well known [11]; for ease of comprehension, we reconstruct the proof in the Appendix using our notation. In summary, given any sequence of queue lengths, we can use different probabilities of arrival and departure, as long as they lie within $[\underline{a}, \overline{a}]$ and $[\underline{d}, \overline{d}]$, respectively. Therefore, we can optimise with respect to every $p(\cdot|x_{1:i})$ separately.

The advantage of Theorem 1 is that it allows for extremely efficient computations. For example, since $\underline{Q}_{i+1}(h(X_{i+1})|X_{1:i})$ only depends on X_i, it follows from Theorem 1 that evaluating $\underline{E}^{\mathrm{EI}}_n(f)$ has a computational complexity that is linear in n. On the other hand, in general, evaluating $\underline{E}^{\mathrm{RI}}_n(f)$ can only be done approximately, and becomes increasingly inefficient as the approximation becomes better. Other examples of optimisations problems that are easy for the case of epistemic irrelevance, but hard for the case of repetition independence, are discussed further on.

6 Properties and results

In this section, we present some bounds on various expectations and probabilities that are essential for queues and we discuss similarities and differences among the two independence approaches.

6.1 RI included in EI

Figure 3 illustrates how the two approaches differ in terms of arrival and departure probabilities in the probability tree. Because of the difference in the selection of the local probabilities, we easily get the following property.

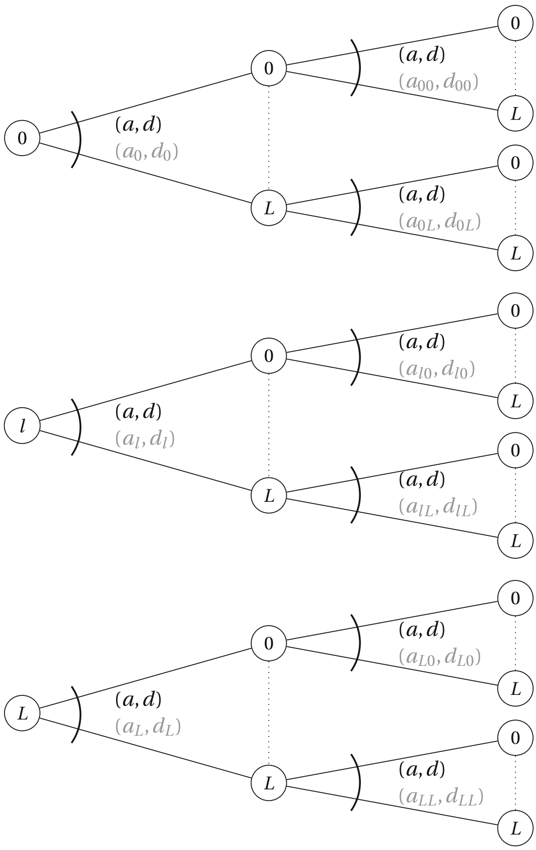

Fig. 3 The difference between the RI and the EI approaches in the local conditional models: in the first approach, we require a time-homogeneous pair of arrival and departure probabilities, whereas in the second, we can have different arrival and departure probabilities at any sequence of queue lengths

Lemma 1 *For any $n \in \mathbb{N}$ and $m \in \mathbb{N}_0$, with $m < n$, and any $f \in \mathscr{L}(\mathscr{X}^n)$:*

$$\underline{E}^{EI}_{1:n}(f|X_{1:m}) \leq \underline{E}^{RI}_{1:n}(f|X_{1:m}) \leq \overline{E}^{RI}_{1:n}(f|X_{1:m}) \leq \overline{E}^{EI}_{1:n}(f|X_{1:m})$$

Proof This is trivial because $\mathscr{T}^{\mathrm{RI}}$ is clearly a subset of $\mathscr{T}^{\mathrm{EI}}$. □

In the remainder, we calculate various interesting performance measures of queues using both approaches and we compare the results. More specifically, we examine the expected queue length and the expected average queue length, for which we find that the two approaches coincide. Next, we calculate the lower and upper (average) probability for each possible queue length. We finish by calculating the lower and upper (average) probability of ‘turning on the server’—a transition from queue length 0 to 1. For some of these performance measures, we also prove a number of useful theoretical properties.

All our experiments concern a queue with $L = 7$, where the probability interval for an arrival is $[0.5, 0.6]$ and the one for a departure $[0.7, 0.8]$. For the initial set of probabilities $\mathscr{Q}_1$, we use a *vacuous* model—the set of all probability mass functions. For the RI approach, we use eleven different values for the arrival and departure

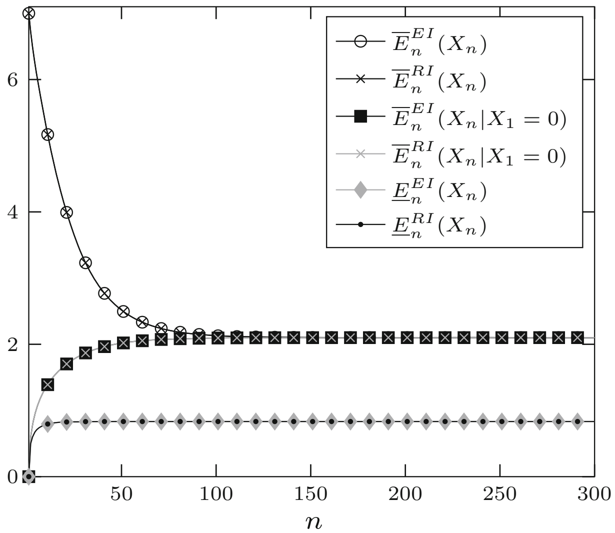

Fig. 4 Lower and upper expected queue length

probabilities and we calculated the corresponding expectations for all possible combinations. In particular, arrival probabilities take values in the set $\{0.5, 0.51, \ldots, 0.6\}$ and departures in $\{0.7, 0.71, \ldots, 0.8\}$. This implies that the lower (and upper) expectations we obtain for these values might not correspond exactly to the actual minimum (or maximum) expectation of the function under study.

6.2 Expected queue length

Figure 4 depicts lower and upper expected queue lengths at time n, for increasing values of n, and for both approaches. We observe convergence for both the lower and the upper expected value, regardless of whether we use the vacuous initial model or start from an empty queue. This is not surprising because, under very weak assumptions, such convergence will always happen, and is furthermore independent of the choice of the initial model; see [15] for the RI case and [11, Theorem 6.1] for the EI case.

6.2.1 Monotonicity

In Fig. 4, we also observe that the results under the EI approach coincide with those of the RI. Under the RI approach, we obtain, reasonably, the lower expected queue length for the smallest probability of arrival (0.5) and the largest probability of departure (0.8). Under the EI approach, although we do not require the use of a single time-homogeneous probability of arrival and probability of departure, one can show that in this case, the optimal choice for the arrival probability is to always consider the minimum value, and the optimal choice for the departure probability is to always consider the maximum value. This is due to the 'monotonicity' of the argument func-

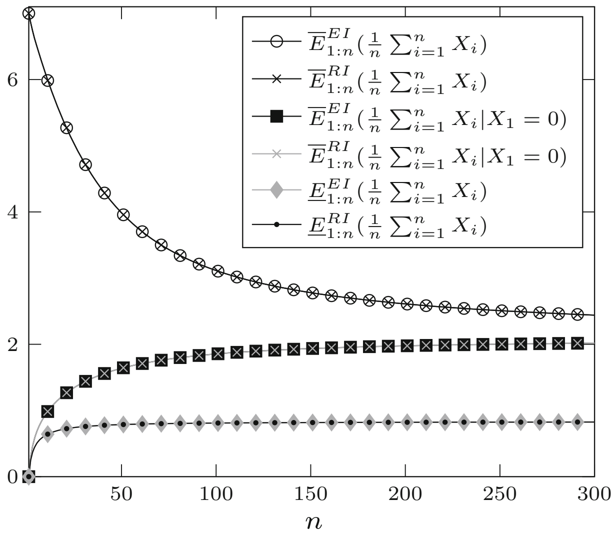

Fig. 5 Lower and upper expected average queue length

tion, as made clear in the following theorem whose proof we have delegated to the Appendix. Similar (suitably adapted) results hold for the upper expectation, and for functions h that are non-increasing in k rather than non-decreasing.

Theorem 2 *Consider any $n \in \mathbb{N}$ and any $h \in \mathscr{L}(\mathscr{X})$ such that*

$$h(k) \leq h(k+1) \text{ for all } k \in \{0, \ldots, L-1\}. \tag{17}$$

Then, in a Geo/Geo/1/L *queue with parameters in intervals $[\underline{a}, \overline{a}]$ and $[\underline{d}, \overline{d}]$, the lower expected value $\underline{E}^{EI}_n(h)$ is obtained for time-homogeneous parameters $\underline{a}$ and $\overline{d}$, and the upper expected value $\overline{E}^{EI}_n(h)$ for $\overline{a}$ and $\underline{d}$.*

6.2.2 Expected average queue length

In Fig. 5, we depict the lower and upper expected average queue length, where the average is being taken over the time points 1 through n. In the precise case, it is well known that—if the expected queue length converges—this expected average queue length converges, at a slower rate, to the same value as the expected queue length. Although the convergence is very slow, our experiments seem to suggest that similar behaviour occurs for our lower and upper bounds on the expected (average) queue length, regardless of whether we use RI or EI.

For the EI approach, computing lower expected average queue lengths is easy. For example, for any $h \in \mathscr{L}(\mathscr{X})$, it follows from Theorem 1 that

$$\begin{aligned}&\underline{E}^{\mathrm{EI}}_{1:n}\left(\frac{1}{n}\sum_{i=1}^{n} h(X_i)\right)\\&\quad= \frac{1}{n}\underline{E}^{\mathrm{EI}}_{1:n}\left(\sum_{i=1}^{n} h(X_i)\right)\\&\quad= \frac{1}{n}\underline{Q}_1\left(\dots\underline{Q}_{n-1}\left(\underline{Q}_n\left(\sum_{i=1}^{n} h(X_i)|X_{1:n-1}\right)|X_{1:n-2}\right)\dots\right)\\&\quad= \frac{1}{n}\underline{Q}_1\Big(h(X_1)+\dots\underline{Q}_{n-1}\Big(h(X_{n-1})+\underline{Q}_n(h(X_n)|X_{1:n-1})|X_{1:n-2}\Big)\dots\Big).\end{aligned}$$

Therefore, if we let $\underline{Q}$ be an operator from $\mathscr{L}(\mathscr{X})$ to $\mathscr{L}(\mathscr{X})$, defined for all $f \in \mathscr{L}(\mathscr{X})$ by

$$\begin{aligned}\underline{Q}f(y) &= \min\left\{\sum_{x\in\mathscr{X}} q(x|y,a,d)f(x)\colon a\in[\underline{a},\overline{a}], d\in[\underline{d},\overline{d}]\right\}\\&= \min\left\{\sum_{x\in\mathscr{X}} q(x|y,a,d)f(x)\colon a\in\{\underline{a},\overline{a}\}, d\in\{\underline{d},\overline{d}\}\right\} \text{ for all } y\in\mathscr{X},\end{aligned} \tag{18}$$

then

$$\underline{E}^{\mathrm{EI}}_{1:n}\left(\frac{1}{n}\sum_{i=1}^{n} h(X_i)\right) = \frac{1}{n}\underline{Q}_1\big(h+\underline{Q}\big(h+\dots\underline{Q}\big(h+\underline{Q}(h)\big)\dots\big)\big),$$

where the right-hand side consists of n nested simple optimisation problems. The lower expected average queue length is obtained by choosing $h(X_i) = X_i$. The upper expected average queue length can be computed analogously; similar expressions hold for the precise case as well and, by evaluating this precise version for different values of a and d, the RI approach can be computed.

6.3 (Average) probability of different queue lengths

In Tables 1 and 2, we show the lower and upper (average) probability of every possible queue length as n tends to infinity. For the results under the RI approach, we also provide—between parentheses—the probabilities of arrival and departure (a, d), for which the lower or upper expectation was obtained. Notice that probabilities of queue lengths can be treated as expectations of indicators, where $\mathbb{I}_k$ is the indicator of k, as

Table 1 Probabilities of queue lengths $0, \ldots, 3$ for $n \to \infty$

k	0	1	2	3
$\underline{E}_n^{\mathrm{EI}}(\mathbb{I}_k)$	0.1486	0.2909	0.1081	0.0276
$\underline{E}_{1:n}^{\mathrm{EI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.1486	0.3093	0.1100	0.0289
$\overline{E}_n^{\mathrm{EI}}(\mathbb{I}_k)$	0.3750	0.5344	0.2684	0.1683
$\overline{E}_{1:n}^{\mathrm{EI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.3750	0.5173	0.2577	0.1492
$\underline{E}_n^{\mathrm{RI}}(\mathbb{I}_k)$	0.1486	0.3185	0.1172	0.0293
	(0.6, 0.7)	(0.6, 0.7)	(0.5, 0.8)	(0.5, 0.8)
$\underline{E}_{1:n}^{\mathrm{RI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.1486	0.3185	0.1172	0.0293
	(0.6, 0.7)	(0.6, 0.7)	(0.5, 0.8)	(0.5, 0.8)
$\overline{E}_n^{\mathrm{RI}}(\mathbb{I}_k)$	0.3750	0.4775	0.2065	0.1316
	(0.5, 0.8)	(0.55, 0.8)	(0.6, 0.72)	(0.6, 0.7)
$\overline{E}_{1:n}^{\mathrm{RI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.3750	0.4775	0.2065	0.1316
	(0.5, 0.8)	(0.55, 0.8)	(0.6, 0.72)	(0.6, 0.7)

Table 2 Probabilities of queue lengths $4, \ldots, 7$ for $n \to \infty$

k	4	5	6	7
$\underline{E}_n^{\mathrm{EI}}(\mathbb{I}_k)$	0.0069	0.0017	0.00044	0.0001
$\underline{E}_{1:n}^{\mathrm{EI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.0073	0.0018	0.00046	0.0001
$\overline{E}_n^{\mathrm{EI}}(\mathbb{I}_k)$	0.1053	0.0648	0.0388	0.0225
$\overline{E}_{1:n}^{\mathrm{EI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.0902	0.0559	0.0352	0.0225
$\underline{E}_n^{\mathrm{RI}}(\mathbb{I}_k)$	0.0073	0.0018	0.00046	0.0001
	(0.5, 0.8)	(0.5, 0.8)	(0.5, 0.8)	(0.5, 0.8)
$\underline{E}_{1:n}^{\mathrm{RI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.0073	0.0018	0.00046	0.0001
	(0.5, 0.8)	(0.5, 0.8)	(0.5, 0.8)	(0.5, 0.8)
$\overline{E}_n^{\mathrm{RI}}(\mathbb{I}_k)$	0.0846	0.0544	0.0350	0.0225
	(0.6, 0.7)	(0.6, 0.7)	(0.6, 0.7)	(0.6, 0.7)
$\overline{E}_{1:n}^{\mathrm{RI}}(\frac{1}{n}\sum_{i=1}^n \mathbb{I}_k(X_i))$	0.0846	0.0544	0.0350	0.0225
	(0.6, 0.7)	(0.6, 0.7)	(0.6, 0.7)	(0.6, 0.7)

defined by

$$\mathbb{I}_k(x) := \begin{cases} 1 & \text{if } x = k \\ 0 & \text{otherwise} \end{cases} \quad \text{for all } x \in \mathscr{X}.$$

For $k \in \{0, 7\}$, as in the case of the expected queue length, both approaches yield identical results.

Again, this is due to the ‘monotonicity’ of the function used. However, this is not the case for the other queue lengths, i.e. $k \in \{1, 2, \ldots, 6\}$. In the paragraphs below, we discuss these and other differences between the two approaches.

6.3.1 RI versus EI

The two approaches sometimes use different probability mass functions towards the calculation of lower (or upper) bounds. For example, as we can see in Table 1, for the RI approach, the stationary upper probability of having queue length 1 is obtained for $a = 0.55$ and $d = 0.8$. On the other hand, as we know from Theorem 1 and Eq. (16), the optimisations in the EI approach only consider extreme values of a, that is $\underline{a}$ or $\overline{a}$. This implies that, in this case, the two approaches use different probability trees. It is therefore not surprising that the limiting values of $\overline{E}_n^{\mathrm{EI}}(\mathbb{I}_1)$ and $\overline{E}_n^{\mathrm{RI}}(\mathbb{I}_1)$, for $n \to \infty$, are different. As is to be expected from Lemma 1, we have that $\overline{E}_n^{\mathrm{EI}}(\mathbb{I}_1) \geq \overline{E}_n^{\mathrm{RI}}(\mathbb{I}_1)$. In this case, we find that $\overline{E}_n^{\mathrm{EI}}(\mathbb{I}_1) > \overline{E}_n^{\mathrm{RI}}(\mathbb{I}_1)$.

6.3.2 Marginal versus average

Judging from Tables 1 and 2, we see that the bounds on the limiting values of the averages are included in the stationary bounds.

For the RI approach, this is to be expected. In fact, for that approach, if $0 < \underline{a} \leq \overline{a} < 1$ and $0 < \underline{d} \leq \overline{d} < 1$, the stationary bounds will coincide with the bounds for the limiting values of the averages because this is true for every precise model in our optimisation problem.

For the EI approach, equality is not necessarily obtained. An expected lower (or upper) average value and a respective single one might be obtained for different probability mass functions. However, in the limit, as the following result establishes, we do obtain the following inequalities; the proof can be found in the Appendix.

Theorem 3 *For all $k \in \{0, \ldots, L\}$, it holds that*

$$\begin{aligned}\liminf_{n\to\infty} \underline{E}_n^{EI}(\mathbb{I}_k) &\leq \liminf_{n\to\infty} \underline{E}_{1:n}^{EI}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_k(X_i)\right)\\ &\leq \limsup_{n\to\infty} \overline{E}_{1:n}^{EI}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_k(X_i)\right) \leq \limsup_{n\to\infty} \overline{E}_n^{EI}(\mathbb{I}_k).\end{aligned}$$

In other words, for the EI approach, the 'worst-case' scenario in the limit is never better than the 'average worst-case' scenario. Our experiments in Tables 1 and 2—where the limit inferiors and limit superiors in the theorem are actually limits—confirm this result. Also, as we can see, in some cases, strict inequalities can be observed. We stress that both scenarios are practically relevant. The probability of being in state k at time n is important for a single customer who would arrive at time instant n, while the average probability of being in state k is important from the system operator's point of view. In the precise case, if $0 < a < 1$ and $0 < d < 1$, these are equal in the long run—as n tends to infinity—but this is no longer guaranteed to be the case when imprecision is added. In Figs. 6 and 7, we show the lower and upper probabilities, single and average ones respectively, for queue length 1. In Fig. 8, we show the lower and upper probabilities, single and average ones, under EI for queue length 1.

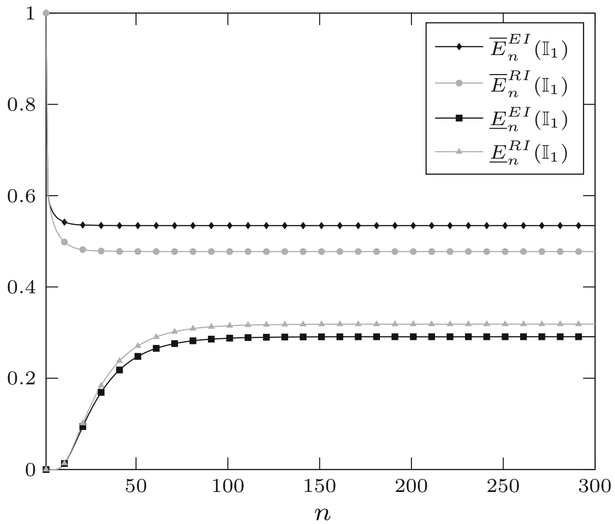

Fig. 6 Lower and upper probability of queue length 1

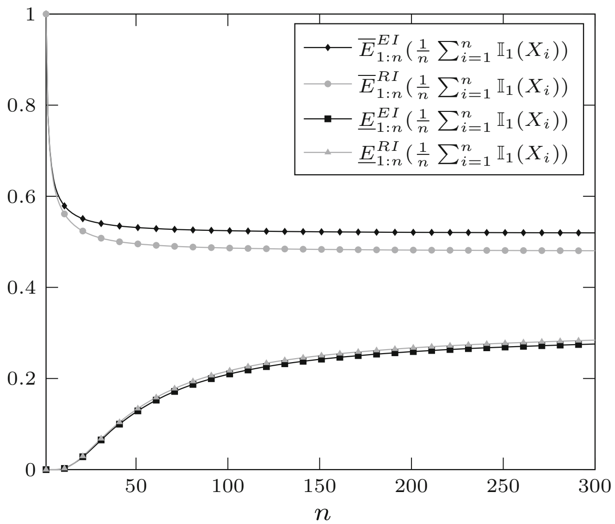

Fig. 7 Lower and upper average probability of queue length 1

6.4 Turning on the server

In Figs. 9 and 10, we depict the lower and upper probability and average probability of ‘turning on the server’. At a single time point $n + 1$, this is taken to be the probability of having queue length 1 at time point $n + 1$ and queue length 0 at time point n.

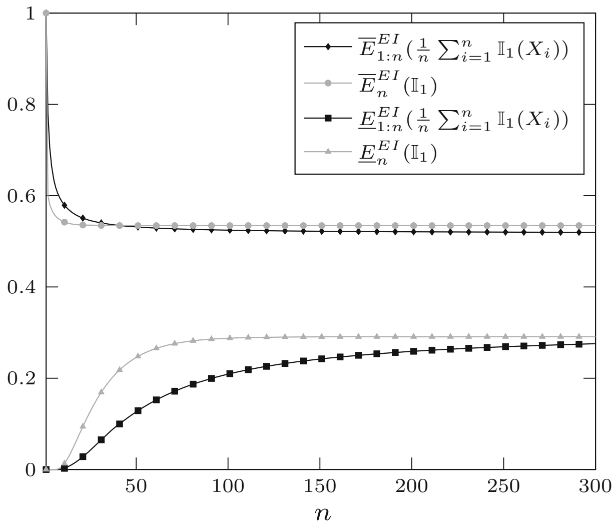

Fig. 8 (Average) probability of queue length 1 in the EI approach

For the EI approach, computing the lower average probability of turning on the server is easy. Indeed, it follows from Theorem 1 that

$$\begin{aligned}
&\underline{E}^{\mathrm{EI}}_{1:n+1}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\right)\\
&=\frac{1}{n}\underline{E}^{\mathrm{EI}}_{1:n+1}\left(\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\right)\\
&=\frac{1}{n}\underline{Q}_1\left(\cdots\underline{Q}_n\left(\underline{Q}_{n+1}\left(\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})|X_{1:n}\right)|X_{1:n-1}\right)\cdots\right)\\
&=\frac{1}{n}\underline{Q}_1\left(\cdots\underline{Q}_n\left(\sum_{i=1}^{n-1}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})+\sum_{x_n\in\mathscr{X}}\underline{Q}(\mathbb{I}_0(x_n)\mathbb{I}_1)|X_{1:n-1}\right)\cdots\right)\\
&=\frac{1}{n}\underline{Q}_1\left(\sum_{x_1\in\mathscr{X}}\underline{Q}\left(\cdots+\sum_{x_{n-1}\in\mathscr{X}}\underline{Q}\left(\mathbb{I}_0(x_{n-1})\mathbb{I}_1+\sum_{x_n\in\mathscr{X}}\underline{Q}(\mathbb{I}_0(x_n)\mathbb{I}_1)\right)\right)\right),
\end{aligned}$$

where the last expression consists of n nested simple optimisation problems. The upper average probability of turning on the server can be computed in an analogous way; similar expressions hold for the precise case as well and, by evaluating this precise version for different values of a and d, the RI approach can be computed.

For $n \to \infty$, we obtain the following results:

(a) $\lim_{n\to\infty}\underline{E}^{\mathrm{EI}}_{1:n+1}(\mathbb{I}_0(X_n)\mathbb{I}_1(X_{n+1})) = 0.0743$
(b) $\lim_{n\to\infty}\underline{E}^{\mathrm{EI}}_{1:n+1}\big(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\big) = 0.0866$

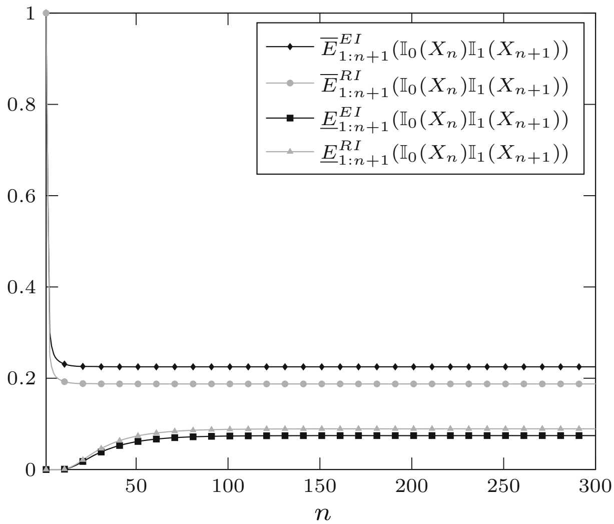

Fig. 9 Lower and upper probability of turning on the server

(c) $\lim_{n\to\infty} \overline{E}^{\mathrm{EI}}_{1:n+1}(\mathbb{I}_0(X_n)\mathbb{I}_1(X_{n+1})) = 0.2250$
(d) $\lim_{n\to\infty} \overline{E}^{\mathrm{EI}}_{1:n+1}\big(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\big) = 0.2000$
(e) $\lim_{n\to\infty} \underline{E}^{\mathrm{RI}}_{1:n+1}(\mathbb{I}_0(X_n)\mathbb{I}_1(X_{n+1})) = 0.0892$
(f) $\lim_{n\to\infty} \underline{E}^{\mathrm{RI}}_{1:n+1}\big(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\big) = 0.0892$
(g) $\lim_{n\to\infty} \overline{E}^{\mathrm{RI}}_{1:n+1}(\mathbb{I}_0(X_n)\mathbb{I}_1(X_{n+1})) = 0.1875$
(h) $\lim_{n\to\infty} \overline{E}^{\mathrm{RI}}_{1:n+1}\big(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_0(X_i)\mathbb{I}_1(X_{i+1})\big) = 0.1875$

As before, we observe that for the EI approach, the ‘worst-case’ scenario in the limit is never better than the ‘average worst-case’ scenario.

7 Conclusions

We have analysed various performance measures of a *Geo/Geo/1/L* queueing system with interval probabilities instead of conventional precise probabilities under two ‘independence’ approaches, namely repetition independence (RI) and epistemic irrelevance (EI). The RI approach, which assumes that the lower (or upper) expectation is achieved by a precise time-homogeneous Markov chain, is not as robust as the EI approach for which the minimisation (or maximisation) is over *all* precise probability trees, as long as the branch probabilities lie in the respective intervals. This implies, in practical terms, that a ‘worst-case’ scenario for the output variable can be worse for EI. However, for monotone functions on the state space $\mathscr{X}$, the EI approach coincides with RI.

Specifically for the EI approach, we witness differences between the bounds on the expected (time-)averages of an output variable and the bounds on the corresponding stationary expectation. The stationary expectations in the limit are more robust than the averages under the EI approach, meaning that a ‘worst-case’ scenario at one specific

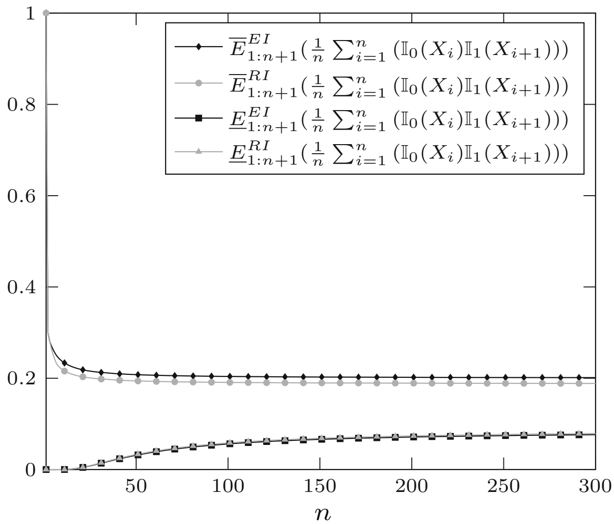

Fig. 10 Lower and upper average probability of turning on the server

time point in the future might be worse than the average one up to that time point. This observation holds practical consequences for the use of queueing models when imprecision is involved. We are used to thinking of queueing systems as ergodic Markov chains with precise parameters, for which we are foremost interested in the unique stationary distribution over its state space. This distribution can then be used to calculate expectations of steady-state performance metrics. In the precise case, these expectations are also predictions for the limiting values of time averages over a long sample path. Under imprecision with epistemic irrelevance, however, these limiting values may have lower and upper expectations that differ from the lower and upper expectations in an arbitrary steady-state instant.

In a recent paper [9], we have tried to shed more theoretical light on this ergodicity issue. It was shown earlier [11,14,26] that there is a large class of so-called Perron–Frobenius-like imprecise Markov chains—to which the queueing models under EI we study here belong, unless $\overline{a} = 0$ or $\overline{d} = 0$—for which the lower expectations $\underline{E}^{\mathrm{EI}}_n(\cdot)$ converge pointwise to a unique steady-state lower expectation $\underline{E}^{\mathrm{EI}}_\infty(\cdot)$. We were able to show that for such imprecise Markov chains—in contradistinction with their precise counterparts—the time average $\frac{1}{n}\sum_{k=1}^{n} h(X_k)$ does not necessarily converge (almost surely), but that nevertheless

$$\underline{E}^{\mathrm{EI}}_\infty(h) \leq \liminf_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} h(X_k) \leq \limsup_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} h(X_k) \leq \overline{E}^{\mathrm{EI}}_\infty(h) \text{ almost surely,}$$

and that (of course) also $\lim_{n\to\infty} \frac{1}{n}\sum_{k=1}^{n} \underline{E}^{\mathrm{EI}}_k(h) = \underline{E}^{\mathrm{EI}}_\infty(h)$ for all $h \in \mathscr{L}(\mathscr{X})$. That we nevertheless observed in Sect. 6.3 that $\lim_{n\to\infty} \underline{E}^{\mathrm{EI}}_{1:n}\left(\frac{1}{n}\sum_{k=1}^{n} h(X_k)\right) \geq \underline{E}^{\mathrm{EI}}_\infty(h)$ is then due to the typical superadditivity of lower expectation operators [25,27].

In future research, we plan to investigate further similarities and divergences of the two approaches from a state-dependent approach [32, Chap. 9, Sect. 4], where again we would assume imprecision in the parameters. Another interesting problem to be tackled is the efficient calculation of the bounds under both approaches for queueing systems that use other types of arrival and/or departure processes.

Acknowledgments Research for this paper by Stavros Lopatatzidis, Gert de Cooman, Stijn De Vuyst and Joris Walraevens was funded through project number 3G012512 of the Research Foundation Flanders (FWO). Jasper De Bock is a PhD Fellow of the FWO and wishes to acknowledge its financial support. The authors would also like to thank three anonymous referees for their many helpful suggestions.

Appendix

Proofs

Proof of Theorem 1 First notice that

$$\begin{aligned}
\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:m}) &= \min\left\{E^{p_{A,D}}_{1:n}(f|X_{1:m})\colon\ p_{A,D}\in\mathscr{T}^{\mathrm{EI}}\right\}\\
&= \min\left\{E^{p_{A,D}}_{1:n-1}\left(E^{p_{A,D}}_{1:n}(f|X_{1:n-1})|X_{1:m}\right)\colon\ p_{A,D}\in\mathscr{T}^{\mathrm{EI}}\right\}\\
&= \min\left\{E^{p_{A,D}}_{1:n-1}\left(E^{p'_{A,D}}_{1:n}(f|X_{1:n-1})|X_{1:m}\right)\colon\ p_{A,D}\in\mathscr{T}^{\mathrm{EI}},\ p'_{A,D}\in\mathscr{T}^{\mathrm{EI}}\right\}\\
&= \min\left\{E^{p_{A,D}}_{1:n-1}\big(\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:n-1})|X_{1:m}\big)\colon\ p_{A,D}\in\mathscr{T}^{\mathrm{EI}}\right\}\\
&= \underline{E}^{\mathrm{EI}}_{1:n-1}\big(\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:n-1})|X_{1:m}\big),
\end{aligned}$$

where the crucial third equality holds because the local probabilities of the probability trees in $\mathscr{T}^{\mathrm{EI}}$ are chosen independently of each other. By continuing in this way, we find that

$$\begin{aligned}
&\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:m})\\
&\quad= \underline{E}^{\mathrm{EI}}_{1:n-1}\left(\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:n-1})|X_{1:m}\right)\\
&\quad= \underline{E}^{\mathrm{EI}}_{1:n-2}\left(\underline{E}^{\mathrm{EI}}_{1:n-1}\left(\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:n-1})|X_{1:n-2}\right)|X_{1:m}\right)\\
&\quad= \dots\\
&\quad= \underline{E}^{\mathrm{EI}}_{1:m+1}\left(\underline{E}^{\mathrm{EI}}_{1:m+2}\left(\dots\underline{E}^{\mathrm{EI}}_{1:n-1}\left(\underline{E}^{\mathrm{EI}}_{1:n}(f|X_{1:n-1})|X_{1:n-2}\right)\dots|X_{1:m+1}\right)|X_{1:m}\right).
\end{aligned}$$

The result now follows because

$$\begin{aligned}
\underline{E}^{\mathrm{EI}}_{1:1}(h) &= \min\left\{E^{p_{A,D}}_{1:1}(h)\colon\ p_{A,D}\in\mathscr{T}^{\mathrm{EI}}\right\}\\
&= \min\left\{\sum_{x_1\in\mathscr{X}} p(x_1)h(x_1)\colon\ p_{X_1}\in\mathscr{Q}_1\right\} = \underline{Q}_1(h) \text{ for all } h\in\mathscr{L}(\mathscr{X}),
\end{aligned}$$

and because, for all $i \in \mathbb{N}$:

$$\begin{aligned}\underline{E}^{\mathrm{EI}}_{1:i+1}(h|X_{1:i}) &= \min\left\{E^{p_{A,D}}_{1:i+1}(h|X_{1:i})\colon p_{A,D} \in \mathscr{T}^{\mathrm{EI}}\right\}\\ &= \min\Bigg\{\sum_{x_{i+1}\in\mathscr{X}} q\left(x_{i+1}|X_i, a_{X_{1:i}}, d_{X_{1:i}}\right)h(X_{1:i}, x_{i+1})\colon\\ &\qquad (\forall x_{1:i} \in \mathscr{X}^i)\, a_{x_{1:i}} \in [\underline{a}, \overline{a}], (\forall x_{1:i} \in \mathscr{X}^i)\, d_{x_{1:i}} \in [\underline{d}, \overline{d}]\Bigg\}\\ &= \underline{Q}_{i+1}(h|X_{1:i}) \text{ for all } h \in \mathscr{L}(\mathscr{X}^{i+1}).\end{aligned}$$

Proof of Theorem 2 We provide the proof for $\underline{E}^{\mathrm{EI}}_n(h)$; the proof for $\overline{E}^{\mathrm{EI}}_n(h)$ is completely analogous. It follows from Eqs. (16) and (18) and Theorem 1 that

$$\underline{E}^{\mathrm{EI}}_n(h) = \underline{Q}_1\underline{Q}^{n-1}h.$$

Therefore, the result follows—by induction—if, for any function $h \in \mathscr{L}(\mathscr{X})$ that satisfies Eq. (17), we can show (a) that $\underline{Q}h$ also satisfies Eq. (17) and (b) that, for all $y \in \mathscr{X}$, the minimum in

$$\underline{Q}h(y) = \min\left\{\sum_{x\in\mathscr{X}} q(x|y, a, d)h(x)\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\} \tag{19}$$

is obtained for $a = \underline{a}$ and $d = \overline{d}$.

For all $y \in \{1, \ldots, L\}$, let $m_y := h(y) - h(y-1) \geq 0$, where the inequality follows from Eq. (17).

We first prove (b). For $y = 0$, Eqs. (1) and (19) imply that

$$\begin{aligned}\underline{Q}h(0) &= \min\left\{(1-a)h(0) + ah(1)\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\}\\ &= \min\left\{h(0) + am_1\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\} = h(0) + \underline{a}m_1,\end{aligned} \tag{20}$$

where the last step holds because $m_1 \geq 0$. Similarly, for $y \in \{1, \ldots, L-1\}$, Eqs. (2) and (19) imply that

$$\begin{aligned}\underline{Q}h(y) &= \min\big\{[d(1-a)]h(y-1) + [(1-d)(1-a) + da]h(y)\\ &\qquad + [(1-d)a]h(y+1)\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\big\}\\ &= \min\left\{h(y) - d(1-a)m_y + (1-d)am_{y+1}\colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\}\\ &= h(y) - \overline{d}(1-\underline{a})m_y + (1-\overline{d})\underline{a}m_{y+1},\end{aligned} \tag{21}$$

where the last step holds because $m_y \geq 0$ and $m_{y+1} \geq 0$.

Finally, for $y = L$, Eqs. (3) and (19) imply that

$$\begin{aligned}\underline{Q}h(L) &= \min\left\{[d(1-a)]h(L-1) + [1-d(1-a)]h(L) \colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\}\\ &= \min\left\{h(L) - d(1-a)m_L \colon a \in \{\underline{a}, \overline{a}\}, d \in \{\underline{d}, \overline{d}\}\right\} = h(L) - \overline{d}(1-\underline{a})m_L,\end{aligned} \tag{22}$$

where the last step holds because $m_L \geq 0$. This concludes the proof of (b).

We now prove (a): $\underline{Q}h(y+1) - \underline{Q}h(y) \geq 0$ for all $y \in \{0, \ldots, L-1\}$. For $y = 0$, this holds because it follows from Eqs. (20) and (21) that

$$\begin{aligned}\underline{Q}h(1) - \underline{Q}h(0) &= \left(h(1) - \overline{d}(1-\underline{a})m_1 + (1-\overline{d})\underline{a}m_2\right) - \left(h(0) + \underline{a}m_1\right)\\ &= m_1 - \overline{d}(1-\underline{a})m_1 + (1-\overline{d})\underline{a}m_2 - \underline{a}m_1\\ &\geq m_1 - \overline{d}(1-\underline{a})m_1 - \underline{a}m_1 = (1-\underline{a})(1-\overline{d})m_1 \geq 0.\end{aligned}$$

For $y \in \{1, \ldots, L-2\}$, this holds because it follows from Eq. (21) that

$$\begin{aligned}&\underline{Q}h(y+1) - \underline{Q}h(y)\\ &\quad= \left(h(y+1) - \overline{d}(1-\underline{a})m_{y+1} + (1-\overline{d})\underline{a}m_{y+2}\right)\\ &\qquad - \left(h(y) - \overline{d}(1-\underline{a})m_y + (1-\overline{d})\underline{a}m_{y+1}\right)\\ &\quad= m_{y+1} - \overline{d}(1-\underline{a})m_{y+1} + (1-\overline{d})\underline{a}m_{y+2} + \overline{d}(1-\underline{a})m_y - (1-\overline{d})\underline{a}m_{y+1}\\ &\quad\geq m_{y+1} - \overline{d}(1-\underline{a})m_{y+1} - (1-\overline{d})\underline{a}m_{y+1} = (1-\underline{a})(1-\overline{d})m_{y+1} \geq 0.\end{aligned}$$

For $y = L-1$, this holds because it follows from Eqs. (21) and (22) that

$$\begin{aligned}\underline{Q}h(L) - \underline{Q}h(L-1) &= \left(h(L) - \overline{d}(1-\underline{a})m_L\right)\\ &\quad - \left(h(L-1) - \overline{d}(1-\underline{a})m_{L-1} + (1-\overline{d})\underline{a}m_L\right)\\ &= m_L - \overline{d}(1-\underline{a})m_L + \overline{d}(1-\underline{a})m_{L-1} - (1-\overline{d})\underline{a}m_L\\ &\geq m_L - \overline{d}(1-\underline{a})m_L - (1-\overline{d})\underline{a}m_L\\ &= \left((1-\overline{d})(1-\underline{a}) + \overline{d}\underline{a}\right) m_L \geq 0.\end{aligned}$$

Proof of Theorem 3 For all $n \in \mathbb{N}$, it follows from subadditivity [27, Chap. 2.6.1(e)] that

$$\underline{E}^{\mathrm{EI}}_{1:n}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_k(X_i)\right) \geq \frac{1}{n}\sum_{i=1}^{n}\underline{E}^{\mathrm{EI}}_{1:n}(\mathbb{I}_k(X_i)) = \frac{1}{n}\sum_{i=1}^{n}\underline{E}^{\mathrm{EI}}_{i}(\mathbb{I}_k),$$

whence

$$\liminf_{n\to\infty}\underline{E}^{\mathrm{EI}}_{1:n}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbb{I}_k(X_i)\right) \geq \liminf_{n\to\infty}\frac{1}{n}\sum_{i=1}^{n}\underline{E}^{\mathrm{EI}}_{i}(\mathbb{I}_k) \geq \liminf_{n\to\infty}\underline{E}^{\mathrm{EI}}_{n}(\mathbb{I}_k),$$

where the last inequality follows from the definition of the limit inferior. The proof for the upper expectations is completely analogous; the inequalities are reversed and subadditivity is replaced by superadditivity.

References

1. Alfa, A.S.: Queueing Theory for Telecommunications: Discrete Time Modelling of a Single Node System, 1st edn. Springer, New York (2010)
2. Altman, E., Avrachenkov, K.E., Núñez-Queija, R.: Perturbation analysis for denumerable Markov chains with application to queueing models. Adv. Appl. Prob. **36**, 839–853 (2004)
3. Augustin, T., Coolen, F.P., De Cooman, G., Troffaes, M.C. (eds.): Introduction to Imprecise Probabilities. Wiley (2014)
4. Boole, G.: Studies in Logic and Probability. Dover Publications, Mineola (2004, reprint of the work originally published by Watts & Co., London, in 1952)
5. Bucklew, J.A., Ney, P., Sadowsky, J.S.: Monte Carlo simulation and large deviations theory for uniformly recurrent Markov chains. J. Appl. Prob. **27**, 44–59 (1990)
6. Cao, X.R., Chen, H.F.: Perturbation realization, potentials, and sensitivity analysis of Markov processes. IEEE Trans. Autom. Control **42**(10), 1382–1393 (1997)
7. Chaudhry, M.L.: On numerical computations of some discrete-time queues. In: Computational Probability, Chap. 10, pp. 365–408. Springer (2000)
8. Couso, I., Moral, S., Walley, P.: A survey of concepts of independence for imprecise probabilities. Risk Decis. Policy **5**(02), 165–181 (2000)
9. De Cooman, G., De Bock, J., Lopatatzidis, S.: A pointwise ergodic theorem for imprecise Markov chains. In: Augustin, T., Doria, S., Miranda, E., Quaeghebeur E. (eds.) ISIPTA '15—Proceedings of the Ninth International Symposium on Imprecise Probability: Theories and Applications, pp. 107–116. Aracne Editrice (2015)
10. De Cooman, G., Hermans, F., Antonucci, A., Zaffalon, M.: Epistemic irrelevance in credal networks: the case of imprecise Markov trees. In: Augustin, T., Coolen, F.P.A., Moral, S., Troffaes, M.C.M. (eds.) ISIPTA '09—Proceedings of the Sixth International Symposium on Imprecise Probability: Theories and Applications, pp. 149–158. SIPTA, SIPTA, Durham,(2009). http://www.sipta.org/isipta09/proceedings/053.html
11. De Cooman, G., Hermans, F., Quaeghebeur, E.: Imprecise Markov chains and their limit behaviour. Prob. Eng. Inf. Sci. **23**(4), 597–635 (2009). doi:10.1017/S0269964809990039. arXiv:0801.0980
12. De Bock, J., De Cooman, G.: An efficient algorithm for estimating state sequences in imprecise hidden Markov models. J. Artif. Intell. Res. **50**, 189–233 (2014)
13. Gross, D., Shortle, J.F., Thompson, J.M., Harris, C.M.: Fundamentals of Queueing Theory. Wiley, New York (2013)
14. Hermans, F., De Cooman, G.: Characterisation of ergodic upper transition operators. Int. J. Approximate Reasoning **53**(4), 573–583 (2012). doi:10.1016/j.ijar.2011.12.008
15. Kemeny, J.G., Snell, J.L.: Finite Markov Chains. Undergraduate Text in Mathematics. Springer, New York (1976). Reprint of the 1960 Edition
16. Keynes, J.M.: A Treatise on Probability. Macmillan, London (1921)
17. Levi, I.: The Enterprise of Knowledge. MIT Press, London (1980)
18. Quaeghebeur, E.: Learning from samples using coherent lower previsions. Ph.D. thesis, Ghent University (2008)
19. Rubinstein, R.Y.: Optimization of computer simulation models with rare events. Eur. J. Oper. Res. **99**(1), 89–112 (1997)
20. Schweitzer, P.J.: Perturbation theory and finite Markov chains. J. Appl. Prob. **5**, 401–413 (1968)
21. Seidenfeld, T., Schervish, M.J., Kadane, J.B.: Rethinking the Foundations of Statistics. Cambridge University Press, Cambridge (1999)
22. Shafer, G.: A Mathematical Theory of Evidence. Princeton University Press, Princeton (1976)
23. Shafer, G.: The Art of Causal Conjecture. The MIT Press, Cambridge (1996)
24. Smith, C.A.B.: Consistency in statistical inference and decision. J. Royal Stat. Soc. A **23**, 1–37 (1961)
25. Troffaes, M.C.M., De Cooman, G.: Lower Previsions. Wiley, New York (2014)

26. Škulj, D., Hable, R.: Coefficients of ergodicity for Markov chains with uncertain parameters. Metrika **76**(1), 107–133 (2013). doi:10.1007/s00184-011-0378-0
27. Walley, P.: Statistical Reasoning with Imprecise Probabilities. Chapman and Hall, London (1991)
28. Walley, P.: Inferences from multinomial data: learning about a bag of marbles. J. Royal Stat. Soc. B **58**, 3–57 (1996). With discussion
29. Walley, P.: Towards a unified theory of imprecise probability. Int. J. Approximate Reasoning **24**, 125–148 (2000)
30. Weichselberger, K.: The theory of interval-probability as a unifying concept for uncertainty. Int. J. Approximate Reasoning **24**, 149–170 (2000)
31. White, J.A.: Analysis of Queueing Systems. Elsevier, Amsterdam (2012)
32. Whittle, P.: Probability Via Expectation, 4th edn. Springer, New York (2000)
33. Williams, P.M.: Notes on conditional previsions. Technical report, School of Mathematical and Physical Science, University of Sussex (1975). Revised journal version: [35]
34. Williams, P.M.: Indeterminate probabilities. In: M. Przelecki, K. Szaniawski, R. Wojcicki (eds.) Proceedings of a 1974 Conference Held in Warsaw. Formal Methods in the Methodology of Empirical Sciences, pp. 229–246. Reidel, Dordrecht (1976)
35. Williams, P.M.: Notes on conditional previsions. Int. J. Approximate Reasoning **44**, 366–382 (2007). Revised journal version of [33]
36. Xia, L., Cao, X.R.: Performance optimization of queueing systems with perturbation realization. Eur. J. Oper. Res. **218**, 293–304 (2012)
37. Zadeh, L.A.: Fuzzy sets. Inf. Control **8**(3), 338–353 (1965)