

Upper Expected Meeting Times for Interdependent Stochastic Agents

Marco Sangalli^[0009-0005-9047-5839], Erik Quaeghebeur^[0000-0003-1462-401X],
and Thomas Krak^[0000-0002-6182-7285]

Uncertainty in AI, Eindhoven University of Technology, The Netherlands

Abstract. We analyse the problem of meeting times for interdependent stochastic agents: random walkers whose behaviour is stochastic but controlled by their selections from some set of allowed actions, and the inference problem of when these agents are all in the same state for the first time. We consider the case where we are epistemically uncertain about the selected actions of these agents, and show how their behaviour can be modelled using imprecise Markov chains. This allows us to use results and algorithms from the literature, to exactly compute bounds on their meeting time, which are tight with respect to our epistemic uncertainty models. After focussing on the two-agent case, we analyse and discuss how it can be extended to an arbitrary number of agents, and how the corresponding combinatorial explosion can be partly mitigated by exploiting symmetries inherent in the problem.

Keywords: Markov Chain · Upper Expectation · Meeting Time · Imprecise Probability

1 Introduction

Markov chains are a cornerstone of stochastic modelling, finding applications in fields as diverse as statistical physics [15], queuing theory [11], and network science [21]. A fundamental question in this context is the *hitting time*: the first moment at which a stochastic process enters a given set of states [16]. When considering multiple independent Markov chains evolving on the same finite state space, a natural extension is the *meeting time*, defined as the first moment at which all chains occupy the same state. Equivalently, this can be viewed as a hitting-time problem on the Cartesian product of their state spaces, with the target set being the diagonal [8, 16, 14].

In this work, we generalize this problem to that of computing the expected meeting time of what we call *interdependent stochastic agents*. Broadly, these are random walkers on some shared state space, whose behaviour is stochastic, but controlled by their selection of transition probabilities from some set representing their allowed actions. Crucially, we allow the selections that these agents make to depend on the state(s) of the other agent(s) in the system; this induces the interdependency.

On top of this construction, we consider the case where we are *epistemically uncertain* about the selections that the agents use, and show that meaningful and conservative estimates for the expected meeting time can still be computed in this setting. We focus in particular on three such uncertainty models: the degenerate belief, which simply describes the case where we *do* know the selections exactly; the vacuous belief, in which we are fully ignorant; and the degenerate-vacuous mixture, which combines these other models and may be particularly useful in some practical settings.

We show how the vacuous belief model allows us to describe the joint behaviour of the agents as an *imprecise Markov chain* [10, 5, 4] on their product space. This allows us to leverage known results from the literature, and in particular iterative algorithms for computing upper- and lower expected hitting times [13], to exactly solve the expected meeting time problem in this setting.

We largely focus on the two-agent case, but show how the formalization can be naturally extended to an arbitrary number of agents using a k -fold product construction. Moreover, we discuss how we can partly mitigate the combinatorial explosion of this problem, by quotienting the (exponentially large) product space over the symmetries obtained by permuting the order of the agents.

The remainder of this paper is structured as follows. Section 2 reviews preliminary notions of precise and imprecise Markov chains, known theory about hitting- and meeting times, and presents an extension of Krak’s characterizations and algorithms [12] to hold under weaker conditions than their original assumptions. Section 3 formalizes the setting of two interdependent stochastic agents with epistemic uncertainty models for their selections, and the characterization of meeting times in terms of the hitting times of an imprecise Markov chain on the product space. In Section 4 we extend the theory to an arbitrary number of interdependent stochastic agents, and Section 5 concludes the paper and gives directions for future work.

2 Preliminaries

This section introduces the basic concepts of stochastic processes and Markov chains [16], hitting and meeting times for Markov chains [16, 8, 7] and imprecise Markov chains [10, 5, 4, 13].

2.1 Stochastic Processes and Markov Chains

Let us denote by \mathbb{N} the set of positive integers and define $\mathbb{N}_0 := \mathbb{N} \cup \{0\}$. Let \mathcal{Z} be a finite state space of cardinality $N \geq 2$. A discrete-time stochastic process on \mathcal{Z} is a sequence of \mathcal{Z} -valued random variables $(Z_n)_{n \in \mathbb{N}_0}$, and we write \mathbb{P}_Z for its associated probability measure. Throughout, we use the notations “ (Z_n) ” and “ \mathbb{P}_Z ” interchangeably to refer to this process.

The process (Z_n) is called a *Markov chain* if it satisfies the Markov property, i.e. if for every $n \in \mathbb{N}_0$ and every $z_0, \dots, z_{n+1} \in \mathcal{Z}$, it holds that

$$\mathbb{P}_Z(Z_{n+1} = z_{n+1} \mid Z_0 = z_0, \dots, Z_n = z_n) = \mathbb{P}_Z(Z_{n+1} = z_{n+1} \mid Z_n = z_n).$$

The chain (Z_n) is called (time-)homogeneous if the one-step transition probabilities do not depend on n . Equivalently, there exists a single matrix $T = (T(z, z'))_{z, z' \in \mathcal{Z}}$, with

$$T(z, z') = \mathbb{P}_Z(Z_{n+1} = z' \mid Z_n = z) \quad \text{for all } n \in \mathbb{N}_0,$$

and each row of T summing to unity. Similarly, a non-homogeneous Markov chain (Z_n) is identified by a family of transition matrices (T_n) .

2.2 Hitting and Meeting Times in the Precise Setting

Hitting Times. Let (Z_n) be a homogeneous Markov chain on \mathcal{Z} with transition matrix $T = (T(z, z'))_{z, z' \in \mathcal{Z}}$. For any target set $A \subset \mathcal{Z}$, define the *hitting time*

$$\tau_A := \inf\{n \geq 0 : Z_n \in A\} \in \mathbb{N}_0 \cup \{+\infty\}.$$

This random variable represents the number of steps the process needs to do to first reach A . Conditioned on the chain starting at $z \in \mathcal{Z}$, the *expected hitting time* is

$$h_A^T(z) := \mathbb{E}_{\mathbb{P}_Z}[\tau_A \mid Z_0 = z].$$

The quantity $h_A^T(z)$ ¹ represents the expected number of steps that the process (Z_n) , starting in $z \in \mathcal{Z}$, takes before reaching the set A .

Theorem 1 ([16]). *Let $A \subset \mathcal{Z}$ and let (Z_n) be a homogeneous Markov chain with transition matrix T . The vector of expected hitting times $h_A^T = (h_A^T(z))_{z \in \mathcal{Z}}$ is the minimal non-negative solution to the following system of equations:*

$$\begin{cases} h_A^T(z) = 0 & \text{for all } z \in A, \\ h_A^T(z) = 1 + \sum_{z' \in \mathcal{Z}} T(z, z') h_A^T(z') & \text{for all } z \notin A, \end{cases} \quad (1)$$

which can be rewritten as

$$h_A^T = \mathbb{1}_{A^c} + \mathbb{1}_{A^c} \cdot T h_A^T, \quad (2)$$

where $\mathbb{1}$ is the indicator function and \cdot represents the element-wise multiplication.

Meeting Times. Consider a joint stochastic process $(X_n, Y_n)_{n \in \mathbb{N}_0}$ on $\mathcal{Z}^2 := \mathcal{Z} \times \mathcal{Z}$. Define the *meeting time* for X_n and Y_n as

$$\mu := \inf\{n \geq 0 : X_n = Y_n\} \in \mathbb{N}_0 \cup \{+\infty\}.$$

If the joint process starts at $(X_0, Y_0) = (x, y)$ then the *expected meeting time* is

$$m(x, y) := \mathbb{E}_{\mathbb{P}_{(X, Y)}}[\mu \mid X_0 = x, Y_0 = y]. \quad (3)$$

¹ The superscript T does *not* denote matrix transpose; rather, it indicates the dependency of the hitting time on the transition matrix T .

The quantity $m(x, y)$ represents the expected number of steps that the two processes, when starting in x and y respectively, take before meeting. In the case where the joint process (X_n, Y_n) is obtained as the independent product of two homogenous Markov chains (X_n) and (Y_n) with transition matrices T and S , respectively, the expected meeting time has a particularly nice characterisation.

Theorem 2 ([8, 7]). *Let (X_n) and (Y_n) be two homogenous Markov chains with transition matrices T and S respectively. The matrix of expected meeting times $m = (m(x, y))_{x, y \in \mathcal{Z}}$ is the minimal non-negative solution to the following system of equations:*

$$\begin{cases} m(z, z) = 0 & \text{for all } z \in \mathcal{Z}, \\ m(x, y) = 1 + \sum_{x', y' \in \mathcal{Z}} T(x, x')S(y, y')m(x', y') & \text{for all } x \neq y, \end{cases} \quad (4)$$

which can be rewritten as

$$m = J + J \cdot (TmS^\top), \quad (5)$$

where J is a $N \times N$ matrix that has zeros on the main diagonal and ones everywhere else, and S^\top is the transpose of S .

The meeting time problem can be reformulated as a hitting time problem on the product space \mathcal{Z}^2 , with target set being $D := \{(z, z) \mid z \in \mathcal{Z}\}$ and transition probability from state (x, y) to state (x', y') being $T(x, x')S(y, y')$.

Using the operation of vectorization and the Kronecker product [22, 8], denoted by the symbol \otimes , we may rewrite (5) as

$$\text{vec}(m) = \mathbb{1}_{D^c} + \mathbb{1}_{D^c} \cdot (T \otimes S) \text{vec}(m),$$

where the operator vec stacks the columns of m to make an N^2 long vector.

2.3 Imprecise Markov Chains

Let us consider a set of transition matrices \mathcal{T} on the finite space \mathcal{Z} , with $|\mathcal{Z}| = N$. Throughout this paper, we assume that \mathcal{T} is a non-empty, compact, and convex set with rows that are independently specified—a property known in the literature as having separately specified rows (SSR) [10, 13]. The assumption of having separately specified rows is equivalent to defining N separate convex and compact sets of probability distributions, one for each state $z \in \mathcal{Z}$.

An imprecise Markov chain (IMC) can be viewed as a collection of Markov chains. A natural example is $\mathcal{P}_{\mathcal{T}}^H$, the collection of all homogeneous Markov chains whose transition matrix belongs to \mathcal{T} . We are interested in computing hitting times for imprecise Markov chains, so we introduce the following upper and lower expectation operators:

$$\overline{\mathbb{E}}_{\mathcal{T}}[\cdot] := \sup_{\mathbb{P} \in \mathcal{P}_{\mathcal{T}}^H} \mathbb{E}_{\mathbb{P}}[\cdot], \quad \underline{\mathbb{E}}_{\mathcal{T}}[\cdot] := \inf_{\mathbb{P} \in \mathcal{P}_{\mathcal{T}}^H} \mathbb{E}_{\mathbb{P}}[\cdot].$$

These definitions allow us to define the upper and lower expected hitting time of a target set $A \subset \mathcal{Z}$ as

$$\bar{h}_A^\mathcal{T}(z) := \bar{\mathbb{E}}_\mathcal{T}[\tau_A \mid Z_0 = z] \quad \text{and} \quad \underline{h}_A^\mathcal{T}(z) := \underline{\mathbb{E}}_\mathcal{T}[\tau_A \mid Z_0 = z],$$

respectively.

There exists a characterization for upper expected hitting times similar to Equation (2). Let $\bar{T} : \mathbb{R}^N \rightarrow \mathbb{R}^N$ be the (possibly) nonlinear map defined for all $f \in \mathbb{R}^N$ and all $z \in \mathcal{Z}$ as

$$[\bar{T}f](z) := \sup_{T \in \mathcal{T}} [Tf](z).$$

Then, we have that $\bar{h}_A^\mathcal{T} = \mathbb{1}_{A^c} + \mathbb{1}_{A^c} \cdot \bar{T} \bar{h}_A^\mathcal{T}$, with a similar characterization for the lower expected hitting times [13].

Reachability Condition. Given a set of transition matrices \mathcal{T} we say that a set of states $\mathcal{C} \subseteq \mathcal{Z}$ is *upper reachable* [2, 4] from $z \in \mathcal{Z}$ if there exists $n \in \mathbb{N}$ such that $[\bar{T}^n \mathbb{1}_\mathcal{C}](z) > 0$, i.e. there exists $n \in \mathbb{N}$ and $T \in \mathcal{T}$ such that $[T^n \mathbb{1}_\mathcal{C}](z) > 0$. We denote by $z \rightarrow \mathcal{C}$ if \mathcal{C} is upper reachable from z . Analogously, the set $\mathcal{C} \subseteq \mathcal{Z}$ is *lower reachable* [3, 20] from $z \in \mathcal{Z}$ if there exists $n \in \mathbb{N}$ such that $[\underline{T}^n \mathbb{1}_\mathcal{C}](z) > 0$. We denote by $z \rightarrow \mathcal{C}$ if \mathcal{C} is lower reachable from z .

Given a set of transition matrices \mathcal{T} , let $A \subset \mathcal{Z}$ be a closed target set, i.e. for all $x \in A, y \notin A$ and all $T \in \mathcal{T}$ it holds that $T(x, y) = 0$. We say that the reachability condition (R1) holds if

$$(R1) : \quad \forall z \in A^c : z \rightarrow A$$

Under (R1), the upper expected hitting time $\bar{h}_A^\mathcal{T}(z)$ is finite for all $z \in A^c$ [12]. If (R1) does not hold, then there exists a non-empty set of states $\mathcal{A}_\mathcal{T} \subset A^c$ such that

$$\mathcal{A}_\mathcal{T} := \{z \in A^c \mid z \not\rightarrow A\}.$$

For all states $z \in \mathcal{A}_\mathcal{T}$ it naturally holds that $\bar{h}_A^\mathcal{T}(z) = +\infty$. This follows from the fact that, if A is closed, then for all $z \in \mathcal{A}_\mathcal{T}$ there exists $T \in \mathcal{T}$ such that, for all $n \in \mathbb{N}$, $[T^n \mathbb{1}_A](z) = 0$ ².

Let us denote with $\mathcal{U}_\mathcal{T}$ the set of all states of $A^c \setminus \mathcal{A}_\mathcal{T}$ that upper reach a state in $\mathcal{A}_\mathcal{T}$, i.e.

$$\mathcal{U}_\mathcal{T} := \{z \in A^c \setminus \mathcal{A}_\mathcal{T} \mid z \rightarrow \mathcal{A}_\mathcal{T}\}.$$

The following theorem characterizes upper hitting times for all states of \mathcal{Z} .

Theorem 3. *Let \mathcal{T} be a set of transition matrices and let $\mathcal{A}_\mathcal{T}$ and $\mathcal{U}_\mathcal{T}$ be as defined above. Then, for all $z \in \mathcal{B}_\mathcal{T} := \mathcal{A}_\mathcal{T} \cup \mathcal{U}_\mathcal{T}$ it holds that $\bar{h}_A^\mathcal{T}(z) = +\infty$. Moreover, there exists a matrix $\tilde{T} \in \mathcal{T}$ such that*

$$\bar{h}_A^\mathcal{T}|_{A^c \setminus \mathcal{B}_\mathcal{T}} = (I - \tilde{T}|_{A^c \setminus \mathcal{B}_\mathcal{T}})^{-1} \mathbf{1},$$

² This implication is not entirely trivial, but we omit the full derivation due to page limit constraints.

where I denotes the identity matrix and $\mathbf{1}$ the vector of all ones.

Proof. It follows from Krak et al. [13, Theorem 12] that there always exists a matrix $\tilde{T} \in \mathcal{T}$ such that $\bar{h}_A^{\tilde{T}} = h^{\tilde{T}}$, where $h^{\tilde{T}}$ is the minimal non-negative solution of the linear system $h^{\tilde{T}} = \mathbb{1}_{A^c} + \mathbb{1}_{A^c} \cdot \tilde{T} h^{\tilde{T}}$. If $z \in \mathcal{U}_{\tilde{T}}$, the Markov chain $\mathbb{P}_{\tilde{T}}$ that starts in z has positive probability of hitting the set $\mathcal{A}_{\tilde{T}}$ thus positive probability of never reaching A . This implies that the upper expected hitting time for all $z \in \mathcal{U}_{\tilde{T}}$ is infinite.

Observe that, for all $z \in A^c \setminus \mathcal{B}_{\tilde{T}}$, for all $T \in \mathcal{T}$, and for all $n \in \mathbb{N}$, it holds that $[T^n \mathbb{1}_{\mathcal{B}_{\tilde{T}}}] (z) = 0$. Therefore, we can restrict \tilde{T} on $\tilde{\mathcal{Z}} := \mathcal{Z} \setminus \mathcal{B}_{\tilde{T}}$ as the Markov chain $\mathbb{P}_{\tilde{T}}$ starting from any state in $\tilde{\mathcal{Z}}$ never leaves $\tilde{\mathcal{Z}}$. Thus, starting from any state in $\tilde{\mathcal{Z}} \setminus A$, the process eventually reaches A , and this yields the invertibility of $I - \tilde{T}|_{\tilde{\mathcal{Z}} \setminus A}$ [16]. We are then able to conclude that $h^{\tilde{T}}|_{\tilde{\mathcal{Z}} \setminus A} = (I - \tilde{T}|_{\tilde{\mathcal{Z}} \setminus A})^{-1} \mathbf{1}$ which is what we wanted. \square

An analogous result holds for lower expected hitting times: in this case, the definitions of $\mathcal{A}_{\tilde{T}}$ and $\mathcal{U}_{\tilde{T}}$ are obtained by swapping upper and lower reachability in their respective constructions.

A Computational Method. Upper- and lower expected hitting times can be computed efficiently using iterative algorithms [12, 13]. One such algorithm is the following: starting from any extreme point T_1 of (the convex set) \mathcal{T} , let h_n be the unique solution of the linear system $h_n = \mathbb{1}_{A^c} + \mathbb{1}_{A^c} \cdot T_n h_n$ and let T_{n+1} be an extreme point of \mathcal{T} such that $\bar{T} h_n = T_{n+1} h_n$. Krak [12] showed that, under (R1), the sequence $(h_n)_{n \in \mathbb{N}}$ converges to $\bar{h}_A^{\tilde{T}}$. We refer to this procedure as Krak's algorithm.

If there is a non-empty set $\mathcal{A}_{\tilde{T}}$ then, as Theorem 3 shows, we may restrict our analysis to $A^c \setminus \mathcal{B}_{\tilde{T}}$ as elsewhere the upper expected hitting time is either zero or infinite. For each $T \in \mathcal{T}$ consider the restriction of T on $\mathcal{B}_{\tilde{T}}^c$

$$T|_{\mathcal{B}_{\tilde{T}}^c} = (T(i, j))_{i, j \in \mathcal{B}_{\tilde{T}}^c} \in \mathbb{R}^{|\mathcal{B}_{\tilde{T}}^c| \times |\mathcal{B}_{\tilde{T}}^c|},$$

the submatrix of T indexed by $\mathcal{B}_{\tilde{T}}^c$. Then, the set of matrices $T|_{\mathcal{B}_{\tilde{T}}^c}$ containing the restriction on $\mathcal{B}_{\tilde{T}}^c$ of all matrices in \mathcal{T} , satisfies (R1)³. Therefore, we may apply Krak's algorithm to compute the upper expected hitting time for all nontrivial states of \mathcal{Z} .

Before applying Krak's algorithm, we need to first identify sets $\mathcal{A}_{\tilde{T}}$ and $\mathcal{U}_{\tilde{T}}$ (in this order). To obtain the former we make use of the following lemma.

Lemma 1 ([3, 2, 9]). *Let \mathcal{C} be a closed subset of \mathcal{Z} . Let $(C_n)_{n \in \mathbb{N}_0}$ be the non-decreasing sequence given by $C_0 = \mathcal{C}$ and*

$$C_{n+1} = C_n \cup \{z \in \mathcal{Z} \setminus C_n \mid [T \mathbb{1}_{C_n}](z) > 0\},$$

³ To compute the lower expected meeting time, we need to further restrict \mathcal{T} excluding every matrix that has a positive probability of reaching $\mathcal{B}_{\tilde{T}}$.

for all $n \in \mathbb{N}_0$. Let n^* be the first index for which $C_{n^*} = C_{n^*+1}$. Then, the set C_{n^*} contains all states of \mathcal{Z} that lower reach \mathcal{C} .

This lemma enables us to implement a straightforward algorithm to identify $\mathcal{A}_{\mathcal{T}}^c$ (and therefore $\mathcal{A}_{\mathcal{T}}$). This algorithm is a modification of Algorithm 2 in [2].

Algorithm 1: Finding the set $\mathcal{A}_{\mathcal{T}}$

```

1  $\mathcal{C} \leftarrow A$ ;
2 repeat
3    $\mathcal{M} \leftarrow \emptyset$ ;
4   foreach  $z \in \mathcal{Z} \setminus \mathcal{C}$  do
5     if  $[\underline{T}\mathbf{1}_{\mathcal{C}}](z) > 0$  then
6        $\mathcal{M} \leftarrow \mathcal{M} \cup \{z\}$ ;
7    $\mathcal{C} \leftarrow \mathcal{C} \cup \mathcal{M}$ ;
8 until  $\mathcal{M} = \emptyset$ ;
9 return  $\mathcal{A}_{\mathcal{T}} = \mathcal{C}^c$ ;

```

Finding the set $\mathcal{U}_{\mathcal{T}}$ of states that upper reach $\mathcal{A}_{\mathcal{T}}$ is even more straightforward. First, we build the directed graph $G_{\overline{\mathcal{T}}} = (\mathcal{Z}, E)$, where $(x, y) \in E$ if $[\overline{T}\mathbf{1}_{\{y\}}](x) > 0$. Then, for all $z \in A^c \setminus \mathcal{A}_{\mathcal{T}}$, we check (e.g. by performing a breadth-first search on $G_{\overline{\mathcal{T}}}$) whether there exists a path in $G_{\overline{\mathcal{T}}}$ from z to $\mathcal{A}_{\mathcal{T}}$. If such path exists then $z \in \mathcal{U}_{\mathcal{T}}$.

3 Upper and Lower Meeting Times for Interdependent Stochastic Agents

We now move on to the formalization of what we call *stochastic agents*. The basic idea is that we have agents—essentially random walkers on the state space \mathcal{Z} —whose behaviour can be influenced by making selections from some allowed set of possible actions.

Formally, for each $z \in \mathcal{Z}$, we consider a set \mathcal{T}_z of probability mass functions on \mathcal{Z} , which we interpret as the set of allowed actions that an agent can select from, whenever they are in state z . We then say that a stochastic agent is a random walker on \mathcal{Z} whose behaviour is determined by her specifying some selection that is compatible with these sets \mathcal{T}_z . Such a selection, say $T(z, \cdot) \in \mathcal{T}_z$, thus corresponds to a probability distribution that governs the stochastic behaviour of the agent: the probability that she will be in state z' at the next time point is given by $T(z, z')$.

To clean up the notation, we gather all these allowed actions in a set \mathcal{T} of transition matrices, such that

$$\mathcal{T} = \{T \text{ a trans. mat.} \mid T(z, \cdot) \in \mathcal{T}_z \text{ for all } z \in \mathcal{Z}\}, \quad (6)$$

where $T(z, \cdot)$ denotes the z row of the transition matrix T . In this way, any state-dependent selection can be expressed as a single element of \mathcal{T} .

Note that so far we have not specified how these selections are made and what they can depend on. Indeed, a simple example would be an agent who simply selects some $T \in \mathcal{T}$ and maintains this forever; the associated walker can then be represented by a homogeneous Markov chain with transition matrix T . Another example might be an agent whose selections depend on time: these selections $(T_n)_{n \in \mathbb{N}_0}$ then determine a non-homogeneous Markov chain that describes the stochastic evolution of the agent. However, the selections might also be functions of additional things.

In particular, we are interested in studying the multi-agent setting, in which the behaviour of agents is interdependent. We start the exposition here with the simplest case, in which we are only dealing with two agents. We say that they are *interdependent*, when the selection that an agent makes at some point in time, depends on the state of the other agent at that same point in time (and vice versa). For each state y that one agent can be in, the other agent needs to specify a selection $T^y(x, \cdot) \in \mathcal{T}_x$ for each state x that *they* can be in. Hence, the selection for this agent can be summarised as a collection $(T^y)_{y \in \mathcal{Z}}$ of transition matrices in \mathcal{T} . Similarly, we denote the selection of the other agent as $(S^x)_{x \in \mathcal{Z}}$.

Given the selections of the two agents, their joint behaviour can be modelled by a stochastic process (X_n, Y_n) on \mathcal{Z}^2 , in which X_n and Y_n represent the states of the two agents at time n . This process is Markovian and homogeneous, and satisfies

$$\mathbb{P}_{(X,Y)}((X_{n+1}, Y_{n+1}) = (x', y') \mid (X_n, Y_n) = (x, y)) = T^y(x, x')S^x(y, y'), \quad (7)$$

where we write (T^y) and (S^x) for the selections driving X_n and Y_n , respectively.

To summarize the above construction, we have described a multi-agent setting, in which the behaviour of each agent is stochastic, but influenced by their selections from some set of allowed actions. There are potentially many interesting inference problems that we could consider, but here we focus on the *expected meeting time* for the two agents.

To conclude the formalization of the problem that we want to study, we now consider the case where we are *epistemically uncertain* about the selections of these agents. Clearly, this uncertainty affects our ability to compute these expected meeting times. However, depending on the exact uncertainty model that we use to represent these beliefs, we may still be able to make useful conclusions about this quantity of interest. In the remainder of this section, we consider this problem under various choices of such uncertainty models.

3.1 Degenerate Belief

Let us begin by considering what is arguably the simplest uncertainty model: the degenerate belief, in which we are certain that the two agents use selections $(T^y)_{y \in \mathcal{Z}}$ and $(S^x)_{x \in \mathcal{Z}}$, respectively.⁴

It follows immediately that this case is described exactly by the single joint process (X_n, Y_n) characterised in Equation (7). Since this joint process is a homogenous Markov chain, it is determined by a transition matrix Δ on \mathcal{Z}^2 . We see from the above that the $((x, y), (x', y'))$ entry, which represents the probability of transitioning from (x, y) to (x', y') , is given by

$$\Delta((x, y), (x', y')) = T^y(x, x')S^x(y, y'). \quad (8)$$

Moreover, the expected meeting time for the agents coincides with the expected hitting time of the Markov chain (X_n, Y_n) on target set $D = \{(z, z) \mid z \in \mathcal{Z}\}$. The corresponding matrix of expected meeting times m satisfies a system of equations analogous to (4), where each occurrence of T and S is replaced by T^y and S^x , respectively; using the definition of Δ yields

$$\begin{cases} m(x, x) = 0 & \text{for all } x \in \mathcal{Z}, \\ m(x, y) = 1 + \sum_{x', y' \in \mathcal{Z}} \Delta((x, y), (x', y'))m(x', y') & \text{for all } x \neq y. \end{cases} \quad (9)$$

So, in this case, we can compute the expected meeting times exactly, by solving this linear system of equations.

3.2 Vacuous Belief

The next uncertainty model that we consider is the *vacuous belief* over the space of all possible selections; this corresponds to the case in which we are fully ignorant about which selections the agents use.

We know from our discussion in Section 3.1 that any specific pair of selections, say $(T^y)_{y \in \mathcal{Z}}$ and $(S^x)_{x \in \mathcal{Z}}$, determines a homogenous Markov chain on \mathcal{Z}^2 that is characterized by a transition matrix Δ , as in Equation (8). By following this construction for every possible pair of selections that the agents might choose, we induce a set of transition matrices on \mathcal{Z}^2 :

$$\mathcal{T}^2 := \left\{ \Delta \in \mathbb{R}^{N^2 \times N^2} \mid \Delta((x, y), (x', y')) := T^y(x, x')S^x(y, y') \right\},$$

where $N = |\mathcal{Z}|$ and we are varying all possible selections (T^y) and (S^x) in \mathcal{T} .

Each element Δ of \mathcal{T}^2 exactly corresponds to a pair of selections (T^y) and (S^x) , which makes this set a convenient representation of the space over which we want to model our beliefs. In particular, each $\Delta \in \mathcal{T}^2$ induces a homogeneous Markov chain \mathbb{P}_Δ , which we collect in the set $\mathcal{P}_{\mathcal{T}^2} = \{\mathbb{P}_\Delta \mid \Delta \in \mathcal{T}^2\}$. Our vacuous

⁴ Clearly, in this case we are not actually ‘uncertain’ at all, but this framing allows us to represent all cases that we consider in a unified framework.

beliefs are now represented by the upper (and lower) expectation operators given by

$$\bar{\mathbb{E}}_{\mathcal{T}^2}[\cdot] := \sup_{\mathbb{P} \in \mathcal{P}_{\mathcal{T}^2}} \mathbb{E}_{\mathbb{P}}[\cdot], \quad \text{and} \quad \underline{\mathbb{E}}_{\mathcal{T}^2}[\cdot] := \inf_{\mathbb{P} \in \mathcal{P}_{\mathcal{T}^2}} \mathbb{E}_{\mathbb{P}}[\cdot]. \quad (10)$$

We immediately recognize the set $\mathcal{P}_{\mathcal{T}^2}$ as an *imprecise Markov chain* on the product space \mathcal{Z}^2 , that is characterised by the set \mathcal{T}^2 of transition matrices. Moreover, our vacuous beliefs in Equation (10) correspond simply to the upper- and lower expectation operators for this imprecise Markov chain.

This also means that we can leverage the known results from Section 2.3, provided that \mathcal{T}^2 satisfies the required regularity conditions; let us analyze how these carry over from the set \mathcal{T} of allowed actions. If \mathcal{T} is compact and with SSR⁵, the same holds for \mathcal{T}^2 . This follows from the fact that for any fixed x and y , rows $T^y(x, \cdot)$ and $S^x(y, \cdot)$ appear only in one row of Δ . Convexity does not carry over from \mathcal{T} to \mathcal{T}^2 . This is due to the fact that the map

$$(T^1, \dots, T^N, S^1, \dots, S^N) \mapsto \Delta$$

is bilinear (i.e. linear in each family (T^y) or (S^x) separately) rather than affine, and only affine maps are guaranteed to preserve convexity [17]. Restoring convexity by replacing \mathcal{T}^2 with its convex hull will prove useful later.

Moving on, we can now characterise the upper expected meeting time as

$$\bar{m}^{\mathcal{T}^2}(x, y) := \bar{\mathbb{E}}_{\mathcal{T}^2}[\mu \mid (X_0, Y_0) = (x, y)] = \bar{h}_D^{\mathcal{T}^2}((x, y)). \quad (11)$$

Since the supremum in Equation (11) is always achieved on a extreme point [1] of \mathcal{T}^2 , and since \mathcal{T}^2 and its convex hull $\text{co}(\mathcal{T}^2)$ share the same extreme points, the upper expected meeting time can be computed by optimizing over $\text{co}(\mathcal{T}^2)$, obtaining the same maximal value.

Similarly as before, let us define the upper transition operator $\bar{\Delta} : \mathbb{R}^{N^2} \rightarrow \mathbb{R}^{N^2}$ for $\text{co}(\mathcal{T}^2)$ as $[\bar{\Delta}f](z) := \sup_{\Delta \in \text{co}(\mathcal{T}^2)} [\Delta f](z)$, for all $f \in \mathbb{R}^{N^2}$ and all $z \in \mathcal{Z}^2$. Then, $\bar{m}^{\mathcal{T}^2}$ is the minimal non-negative solution to

$$\text{vec}(\bar{m}^{\mathcal{T}^2}) = \mathbb{1}_{D^c} + \mathbb{1}_{D^c} \cdot \bar{\Delta} \text{vec}(\bar{m}^{\mathcal{T}^2}). \quad (12)$$

An analogous construction characterises lower expected meeting times.


Computing Upper Meeting Times To compute upper (and lower) expected meeting times under our vacuous belief model, we may use Krak's algorithm [12] for upper (and lower) expected hitting times, restricted to $\mathcal{B}_{\mathcal{T}^2}^c$, as introduced at the end of Section 2.3. First of all, we need to identify the sets $\mathcal{A}_{\mathcal{T}^2}$ and $\mathcal{U}_{\mathcal{T}^2}$ that have infinite upper expected meeting time. To do so, we apply the algorithms from the end of Section 2.3 to the product space \mathcal{Z}^2 , replacing \mathcal{T} with $\text{co}(\mathcal{T}^2)$.

⁵ Note that \mathcal{T} always has SSR, when it is constructed as in Equation (6).

The Corresponding Optimal Control Problem We feel that it is relevant to note that while we phrased the problem here in terms of our epistemic uncertainty about the selections of the agents, our formalization has a natural other interpretation as an optimal control problem. That is, the agents might be strategically choosing their selections—or even have these selections imposed by some external controller—in such a way as to minimize or maximize their meeting time, subject to the constraints on the allowed actions \mathcal{T} . Finding the optimal pair of selections then amounts to finding the elements of \mathcal{T}^2 that attain the extremal values in Equation (10). Indeed, Krak’s algorithm [12] may be used precisely to find such an optimal control policy.

Exploiting the Symmetry. We shall note that we can exploit the symmetry of the problem and restrict our analysis on the *symmetric product space* \mathcal{Z}^2/S_2 , where S_2 is the symmetric group [18]. Explicitly, we identify each pair (x, y) with (y, x) , as the two agents are indistinguishable, so that the space \mathcal{Z}^2/S_2 consists of all unordered pairs of states of \mathcal{Z} . As a consequence the optimal selections $(T^y)_{y \in \mathcal{Z}}$ and $(S^x)_{x \in \mathcal{Z}}$ coincide, i.e. $T^z = S^z$ for all $z \in \mathcal{Z}$. As we pass to the symmetric product space \mathcal{Z}^2/S_2 , \mathcal{T}^2 naturally becomes the set of transition matrices over \mathcal{Z}^2/S_2 . For brevity, we continue to denote it by \mathcal{T}^2 ; the intended domain—whether the full product space or its symmetric quotient—is always clear from context.

Example 1. The figure below shows the graph $G_{\bar{\mathcal{T}}}$ that encodes all possible transitions on the state space $\mathcal{Z} = \{1, 2, 3, 4, 5\}$ (left) and the corresponding transitions on the symmetric product space \mathcal{Z}^2/S_2 (right). An arrow that connects node i to node j in the left graph indicates that $[\bar{\mathcal{T}}\mathbb{1}_{\{j\}}](i) > 0$; an analogous interpretation applies for the right graph with $\bar{\Delta}$ in place of $\bar{\mathcal{T}}$.



Depending on the choice of the set of allowed actions \mathcal{T} , some pair of states in \mathcal{Z} may or may not have infinite upper expected meeting time. For example, if the set \mathcal{T} is such that $[\underline{\mathcal{T}}\mathbb{1}_{\{4\}}](2) = 0$, i.e. an agent in state 2 can choose not to go to state 4, then the pair $(1, 2)$ would become an absorbing state for \mathcal{Z}^2/S_2 , and naturally $(1, 2) \in \mathcal{A}_{\mathcal{T}^2}$. In this case, every pair of states, e.g. $(2, 3)$, that has positive lower probability of eventually hitting $(1, 2)$ would be part of the set $\mathcal{U}_{\mathcal{T}^2}$ whose states also have infinite upper expected meeting time.

3.3 Degenerate-Vacuous Mixtures

The final uncertainty model that we consider is the *degenerate-vacuous mixture*. Note that this is a special case of the well-known *linear-vacuous mixture* [6], in which the linear component is represented by a degenerate (point mass) belief. Consequently, to characterize this model we need to specify both a pair of selections—which we conveniently represent through a single transition matrix $\Delta \in \mathcal{T}^2$ —and some numerical parameter $\epsilon \in [0, 1]$. The associated upper expectation operator is then

$$\bar{\mathbb{E}}_{\epsilon, \Delta, \mathcal{T}^2}[\cdot] := (1 - \epsilon)\mathbb{E}_{\mathbb{P}_\Delta}[\cdot] + \epsilon\bar{\mathbb{E}}_{\mathcal{T}^2}[\cdot], \quad (13)$$

and similarly for the lower expectation. Since for the particular case of meeting times, both terms on the right-hand side of this expression can be computed using the results in Sections 3.1 and 3.2 respectively, it follows that we can also evaluate the upper (and lower) expected meeting times under this model.

We note that this uncertainty model may be particularly useful for conservatively evaluating performance in an “unreliable control” setting. That is, we might consider the degenerate component Δ to represent some optimal control, e.g. a minimizing pair of selections for the meeting time of the two agents, but in which these agents sometimes deviate from the intended control in a manner that is consistent with the allowed actions \mathcal{T} , yet otherwise completely unknown to us. The parameter ϵ measures this degree of unreliability, and the associated upper expected hitting time then provides a conservative estimate for the attained performance of this system.

4 Meeting Times for Multiple Interdependent Agents

In this section, we extend the meeting-time problem to multiple interdependent stochastic agents seeking to all gather at a single common location.

4.1 Imprecise Markov Chain Formulation for k Stochastic Agents

Consider $k \in \mathbb{N}, k > 2$ interdependent stochastic agents on the same state space \mathcal{Z} , with $|\mathcal{Z}| = N$, and with set of transition matrices \mathcal{T} . Each agent looks at the position of all other agents to determine which selection to make. For every agent $j \in [k] := \{1, \dots, k\}$, her selections are identified with a collection of N^{k-1} transition matrices, $(T_j^{(z_i)_{i \in [k], i \neq j}})$ where $(z_i)_{i \in [k], i \neq j}$ represents the states of all other agents. To find the upper (and lower) expected meeting time of these stochastic agents, we consider the joint homogenous Markov chain (Z^1, \dots, Z^k) on $\mathcal{Z}^k = \mathcal{Z} \times \dots \times \mathcal{Z}$, k times. Similarly as before, we define a set of transition matrices \mathcal{T}^k on the space \mathcal{Z}^k as follows

$$\mathcal{T}^k := \left\{ \Delta_k \in \mathbb{R}^{N^k \times N^k} \mid \Delta_k((z_{1:k}), (z'_{1:k})) = \prod_{j=1}^k T_j^{(z_i)_{i \neq j}}(z_j, z'_j) \right\},$$

where, for each $j = 1, \dots, k$, we vary over all possible selections of the k agents. Thus, each matrix $\Delta_k \in \mathcal{T}^k$ is defined using kN^{k-1} transition matrices on \mathcal{Z} .

As before, under the assumption that \mathcal{T} is compact, convex and has SSR, the set \mathcal{T}^k is compact and has separately specified rows. The upper (and lower) expected meeting time of k interdependent stochastic agents with set of transition matrices \mathcal{T} can be seen as the upper (and lower) hitting time of an imprecise Markov chain on the space \mathcal{Z}^k with set of transition matrices \mathcal{T}^k and target set $D^k := \{(z, \dots, z) \mid z \in \mathcal{Z}\}$. As in the case with two agents, by taking the convex hull of \mathcal{T}^k and defining the corresponding upper (and lower) transition operator, we obtain a characterization of the upper (and lower) meeting time analogous to that in Equation (12). Moreover, this also allows us to apply Krak's algorithm [12] to compute the upper (and lower) expected meeting time, as taking the convex hull leaves the upper (and lower) expected meeting time unchanged.

4.2 Symmetry Reduction

The state space \mathcal{Z}^k grows exponentially with k . Exploiting the symmetry drastically reduces the number of states to consider as

$$|\mathcal{Z}^k/S_k| = \binom{N+k-1}{k} \ll N^k = |\mathcal{Z}^k|,$$

where S_k is the symmetric group of degree k . In other words, in the space \mathcal{Z}^k/S_k all permutations of a k -tuple are identified. As in the case $k = 2$, every matrix $\Delta_k : \mathbb{R}^{|\mathcal{Z}^k/S_k|} \rightarrow \mathbb{R}^{|\mathcal{Z}^k/S_k|}$ is now defined using only N^{k-1} matrices from \mathcal{T} , as the k families of transition matrices coincide.

Moreover, if $2 \leq i < k$ stochastic agents are in the same state $z \in \mathcal{Z}$ their selected action that minimizes the meeting time is the same for all i agents and depends only on the position of the other $k-i$ agents. Similarly, these other $k-i$ agents have the same optimal selection whether in z there are 1 or i distinct agents.

Hence, if we want to compute the lower expected meeting time of $k > N$ interdependent stochastic agents, we can reduce to the same problem with N such agents. From the optimal matrix Δ_N we can obtain the optimal selections for each agent, which we can then use to build the optimal matrix Δ_k . Therefore, the cost to compute lower expected meeting times for $k > N$ interdependent stochastic agents is equivalent to that of N such agents, thus drastically reducing the computational burden for large values of k .

5 Conclusion and Future Work

In this work we considered the meeting-time problem for multiple interdependent stochastic agents, about whose selected actions we are epistemically uncertain. By establishing that these meeting times correspond to hitting the diagonal in the product space, and by focusing on specific uncertainty models for our beliefs

about these agents’ selected actions, we were able to recast the problem using the existing theory of imprecise Markov chains. We used their associated upper expectation operators to derive explicit characterizations for these quantities of interest, and leveraged Krak’s iterative algorithm [12] to compute them. Finally, we showed that our approach extends naturally to an arbitrary number of agents via a k -fold product construction, and we exploited the problem’s inherent symmetries to mitigate the exponential growth of the state space.

In future work, we hope to investigate more general interdependency mechanisms between these agents, provide a more in-depth characterisation of the joint imprecise process that describes their behaviour, and consider the case where agents may have different—possibly adversarial—objectives. It would also be interesting to generalize the problem to a continuous-time setting.

Acknowledgments. This work has been partly supported by the PersOn project (P21-03), which has received funding from the Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO). Moreover, we are grateful to the three anonymous reviewers for their careful reading and thoughtful comments. Their input was essential in helping us strengthen the paper and present our ideas more clearly.

Errata

In the published version of this paper [19] the requirement that the target set A must be closed was omitted from Section 2.3, “Imprecise Markov Chains”. In this revised version of the paper we have added a formal definition of a closed set, and updated the surrounding text accordingly. Two missing edges have now been added to the right-hand graph in Example 1. Moreover, some statements in the “Symmetry Reduction” section are true only for lower meeting times, thus also that part has been corrected.

References

- [1] Heinz Bauer. “Minimalstellen von Funktionen und Extrempunkte”. In: *Archiv der Mathematik* 9 (1958), pp. 389–393. DOI: 10.1007/BF01898615.
- [2] Jasper De Bock. “The Limit Behaviour of Imprecise Continuous-Time Markov Chains”. In: *Journal of Nonlinear Science* 27 (2017), pp. 159–196. DOI: 10.1007/s00332-016-9328-3.
- [3] Jasper De Bock, Alexander Erreygers, and Floris Persiau. “A convenient characterisation of convergent upper transition operators”. In: *Proceedings of the 14th International Symposium on Imprecise Probabilities: Theories and Applications (ISIPTA)*. Vol. 290. Proceedings of Machine Learning Research. arXiv:2502.04509. 2025.
- [4] Gert de Cooman, Filip Hermans, and Erik Quaeghebeur. “Imprecise Markov chains and their limit behavior”. In: *Probability in the Engineering and Informational Sciences* 23.4 (2009), pp. 597–635. DOI: 10.1017/S0269964809990039.

- [5] D. J. Hartfiel and E. B. Seneta. “On the theory of Markov set-chains”. In: *Advances in Applied Probability* 26.4 (1994), pp. 947–964. DOI: 10.2307/1427899.
- [6] Sébastien Destercke and Didier Dubois. “Special Cases”. In: *Introduction to Imprecise Probabilities*. Ed. by Thomas Augustin et al. Wiley Series in Probability and Statistics. Wiley, 2014. DOI: 10.1002/9781118763117.ch4.
- [7] Wolfgang Doeblin. “Exposé de la théorie des chaînes simples constantes de Markov à un nombre fini d’états”. In: *Revue de Mathématiques de l’Union Interbalkanique* 2 (1937), pp. 77–105.
- [8] Mishel George, Rushabh Patel, and Francesco Bullo. “The Meeting Time of Multiple Random Walks”. In: *arXiv preprint arXiv:1806.08843* (2018).
- [9] Filip Hermans and Gert de Cooman. “Characterisation of ergodic upper transition operators”. In: *International Journal of Approximate Reasoning* 53.4 (2012), pp. 573–583. DOI: 10.1016/j.ijar.2011.12.008.
- [10] Filip Hermans and Damjan Škulj. “Stochastic Processes”. In: *Introduction to Imprecise Probabilities*. Ed. by Thomas Augustin et al. Wiley, 2014. Chap. 11. DOI: 10.1002/9781118763117.ch11.
- [11] Frank P. Kelly. *Reversibility and Stochastic Networks*. Cambridge University Press, 1979.
- [12] Thomas Krak. “Computing expected hitting times for imprecise Markov chains”. In: *Space Technology Proceedings* 8 (2021), pp. 185–205. DOI: 10.1007/978-3-030-80542-5_12.
- [13] Thomas Krak, Natan T’Joens, and Jasper De Bock. “Hitting Times and Probabilities for Imprecise Markov Chains”. In: *Proceedings of the 14th International Symposium on Imprecise Probabilities: Theories and Applications (ISIPTA)*. Vol. 103. Proceedings of Machine Learning Research. 2019, pp. 265–275. arXiv: 1905.08781.
- [14] Torgny Lindvall. *Lectures on the Coupling Method*. Wiley-Interscience, 1992. ISBN: 9780471540250.
- [15] Nicholas Metropolis et al. “Equation of State Calculations by Fast Computing Machines”. In: *The Journal of Chemical Physics* 21.6 (1953), pp. 1087–1092. DOI: 10.1063/1.1699114.
- [16] James R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge, UK: Cambridge University Press, 1997. ISBN: 9780511810633. DOI: 10.1017/CB09780511810633.
- [17] R. Tyrrell Rockafellar. *Convex Analysis*. Vol. 28. Princeton Mathematical Series. Princeton, NJ: Princeton University Press, 1970. DOI: 10.1515/9781400873173.
- [18] Joseph J. Rotman. *An Introduction to the Theory of Groups*. Vol. 148. Graduate Texts in Mathematics. Springer, 1995. DOI: 10.1007/978-1-4612-4176-8.
- [19] Marco Sangalli, Erik Quaeghebeur, and Thomas Krak. “Upper Expected Meeting Times for Interdependent Stochastic Agents”. In: *Symbolic and Quantitative Approaches to Reasoning with Uncertainty*. Vol. 16099. Lec-

ture Notes in Artificial Intelligence (LNAI). Springer, 2025, pp. 238–252. DOI: 10.1007/978-3-032-05134-9_17.

- [20] Natan T’Joens and Jasper De Bock. “Average behaviour in discrete-time imprecise Markov chains: A study of weak ergodicity”. In: *International Journal of Approximate Reasoning* 132 (2021), pp. 181–205. DOI: 10.1016/j.ijar.2021.03.001.
- [21] Dmitri Volchenkov and Philippe Blanchard. “Markov Chain Methods for Analyzing Complex Transport Networks”. In: *Journal of Statistical Physics* 132.6 (2008), pp. 1051–1069. DOI: 10.1007/s10955-008-9591-2.
- [22] Georg Zehfuss. “Ueber eine gewisse Determinante”. German. In: *Zeitschrift für Mathematik und Physik* 3 (1858), pp. 298–301.