



Probabilistic Disclosure: Maximisation vs. Minimisation

Béatrice Bérard, Serge Haddad, Engel Lefauchaux

► **To cite this version:**

Béatrice Bérard, Serge Haddad, Engel Lefauchaux. Probabilistic Disclosure: Maximisation vs. Minimisation . FSTTCS 2017, Dec 2017, Kanpur, India. <10.4230/LIPIcs.FSTTCS.2017>. <hal-01618955>

HAL Id: hal-01618955

<https://hal.inria.fr/hal-01618955>

Submitted on 18 Oct 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Probabilistic Disclosure: Maximisation vs. Minimisation*

Béatrice Bérard^{1,2}, Serge Haddad², and Engel Lefaucheu^{2,3}

1 Sorbonne Universités, UPMC Univ. Paris 06, CNRS UMR 7606, LIP6, Paris, France, beatrice.berard@lip6.fr

2 LSV, ENS Paris-Saclay & CNRS & Inria, Université Paris-Saclay, France
serge.haddad@lsv.fr

3 Inria, Campus Universitaire de Beaulieu, Rennes, France
engel.lefaucheu@inria.fr

Abstract

We consider opacity questions where an observation function provides to an external attacker a view of the states along executions and secret executions are those visiting some state from a fixed subset. Disclosure occurs when the observer can deduce from a finite observation that the execution is secret, the ε -disclosure variant corresponding to the execution being secret with probability greater than $1 - \varepsilon$. In a probabilistic and non deterministic setting, where an internal agent can choose between actions, there are two points of view, depending on the status of this agent: the successive choices can either help the attacker trying to disclose the secret, if the system has been corrupted, or they can prevent disclosure as much as possible if these choices are part of the system design. In the former situation, corresponding to a worst case, the disclosure value is the supremum over the strategies of the probability to disclose the secret (maximisation), whereas in the latter case, the disclosure is the infimum (minimisation). We address quantitative problems (comparing the optimal value with a threshold) and qualitative ones (when the threshold is zero or one) related to both forms of disclosure for a fixed or finite horizon. For all problems, we characterise their decidability status and their complexity. We discover a surprising asymmetry: on the one hand optimal strategies may be chosen among deterministic ones in maximisation problems, while it is not the case for minimisation. On the other hand, for the questions addressed here, more minimisation problems than maximisation ones are decidable.

1998 ACM Subject Classification D.2.4 Software/Program Verification

Keywords and phrases Partially observed systems – Opacity – Markov chain – Markov decision process

Digital Object Identifier 10.4230/LIPIcs.FSTTCS.2017.

1 Introduction

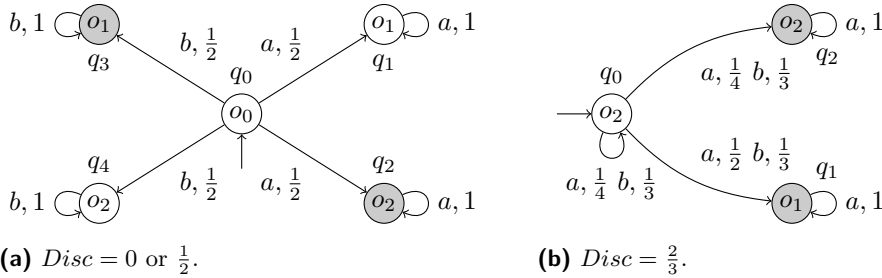
Opacity. Opacity of an information system is a key security property: an external user should not, by observing an execution of a system, acquire the guarantee that it is a secret one. This property was first formalised for labelled transition systems [9], by specifying a subset of secret paths and requiring that, for any secret path, there is a non-secret one with the same observation. The *disclosure set* of a system is then the set of (secret) paths violating opacity.

* The work of S. Haddad has been supported by ERC project EQualIS (FP7-308087).



Opacity raises challenging research issues like (1) formal specification in various frameworks [15, 9], (2) the design of mechanisms to ensure opacity while preserving functionality and performance [2], and (3) the verification of opacity properties [16, 9].

Attacks against opacity. In order to quantify the size of the leak, various measures for the disclosure set, called probabilistic disclosure, were introduced in [5, 19, 3, 6, 4]. For probabilistic and non deterministic systems like Markov Decision Processes (MDP), where an internal agent can choose between several actions, the disclosure can be either maximised or minimised depending on the status of the agent. In previous works, the situation considered corresponded to maximisation, where the system has been corrupted (*e.g.*, by a virus), and the internal agent cooperates with the attacker to disclose the secret. In [3], the set of secret paths is specified by some deterministic automaton and the attacker does not know which strategy is applied, hence the set of executions leaking the secret is fixed by the structure of the system and does not vary according to the strategies. This is illustrated in Figure 1a with an MDP where actions a and b are possible from state q_0 . Action a (resp. b) has a uniform distribution over states q_1 and q_2 (resp. q_3 and q_4). States q_1 and q_3 produce observation o_1 , while q_2 and q_4 produce o_2 . The secret paths are those in $q_0q_2^\omega \cup q_0q_3^\omega$ (reaching either q_2 or q_3 , in grey). If the observer is not aware of the strategy and thus has to consider all possible paths whatever the strategy, the non secret paths are $q_0q_1^\omega$ and $q_0q_4^\omega$, hence there is no disclosing path, leading to a null disclosure (as in [3]). On the other hand, if the observer is aware of the strategy, assuming that initially a is chosen produces a maximal disclosure of $\frac{1}{2}$. This example also shows that deterministic strategies are not sufficient to achieve minimisation: value 0 for the disclosure can only be obtained by randomised strategies, choosing a with probability p and b with probability $1 - p$ (for $0 < p < 1$) in q_0 .



■ **Figure 1** When strategies help (or not) to disclose a secret.

In contrast, Figure 1b represents an MDP where both actions a and b have the same support in state q_0 , hence choosing a or b does not change the states that can be reached. Under a strategy which always plays a , the disclosure is equal to $\frac{2}{3}$ which is the probability of reaching q_1 .

Given some $\varepsilon > 0$, a path could alternatively be considered as disclosing the secret, if the measure of the set of paths with same observation that are not secret is less than ε . This notion is called ε -disclosure. For instance in Figure 1a with $\varepsilon < \frac{1}{2}$ in order to achieve a minimal ε -disclosure of 0 the strategy must select p between ε and $1 - \varepsilon$. However in Figure 1b for every $\varepsilon > 0$ the ε -disclosure is equal to 1 as the probability to be in a secret state converges to 1 on every path.

This figure also illustrates the drastic restriction used in [4] where no edge in an MDP can be blocked by a strategy. With this restricted power of the internal component, the authors can assume that the observer knows the strategy, which is an important requirement since the security of a system should not be based on hiding its design. The model used

in [4] is in fact a restricted case of Interval Markov Decision Processes (IMDPs) so that after some transformation, the problem boils down to IMDP model checking [12]. The general decidability status of the disclosure problem is left open.

Contributions. Here we focus on several problems in MDPs under partial observation that cannot be formalised as problems for classical POMDPs (Partially Observable MDPs). The notion of disclosure is defined with respect to a fixed subset Sec of states: A (finite or infinite) path is secret if it has visited some state of Sec . Other variants of secret path specifications have been proposed with deterministic finite automata accepting finite or infinite paths. The former case can be easily translated in our setting while we believe that the latter one is debatable: a system is not really vulnerable if the attacker can only know the secret at infinite horizon!

Once a strategy is fixed, the behaviour of the system is described by a possibly infinite partially observable Markov chain, so we start in Section 2 by establishing several results on the semantical aspects of disclosure in Markov chains. In addition, we prove undecidability for the positive ε -disclosure problem (deciding if ε -disclosure is positive) within finite horizon. We then consider two different settings depending on the status of the strategies. Like in previous work, maximisation of disclosure corresponds to the internal agent cooperating with the attacker to disclose a secret. Dually, minimisation is interesting to study during the system design process, in order to optimise the choices of the internal agent to defend the system. We address various problems in these settings, for a finite horizon but also for a fixed horizon (given in unary representation), corresponding to real-time constraints requiring the number of steps to be fixed in advance. The quantitative decision problem asks whether the disclosure is above or below some threshold, while qualitative problems consider extremal values (0 or 1) of the disclosure. We prove that observation-based strategies (*i.e.*, which only depend on the sequence of observations and the current state) are dominant in both cases.

The main complexity results for decision problems are gathered in Table 1. For the maximisation objective (Section 3), we show that deterministic strategies are dominant. We answer negatively to the decidability issues left open in [4], proving that both the quantitative problem and the limit-sure problem (asking whether the supremum over all strategies is 1) are undecidable for a finite horizon. Then, we show that the almost-sure problem (asking whether there is a strategy producing a value 1 for disclosure) is EXPTIME-complete. For minimisation (Section 4), we introduce families of randomised strategies, necessary to asymptotically reach minimal disclosure, even within fixed horizon. For finite horizon, we show that the computation and decision problems belong to EXPTIME. Hence surprisingly, although the problem seems more difficult due to the necessity of randomised strategies, the disclosure problem for minimisation is decidable whereas it is not for maximisation. Section 5 is devoted to the fixed horizon problems. For maximisation, we prove that the disclosure value can be computed in PSPACE (while its associated strategy can be computed in EXPTIME) and also establish that the corresponding decision problem is PSPACE-complete. The almost-sure and limit-sure decision problem however are easier and can be solved in PTIME. Refining the techniques to take randomised strategies into account, we obtain PSPACE-completeness of the various decision problems for minimisation. Most of the proofs are in appendix.

2 Specification

We denote by \mathbb{N} the set of natural numbers. For a finite alphabet Σ , we denote by Σ^* (resp. Σ^ω) the set of finite (resp. infinite) words over Σ , with $\Sigma^\infty = \Sigma^* \cup \Sigma^\omega$ and ε the empty word. The length of a word w is denoted by $|w| \in \mathbb{N} \cup \{\infty\}$ and for $n \in \mathbb{N}$, Σ^n is the set of words of

■ **Table 1** Complexity results for maximisation and minimisation of disclosure.

Disclosure	General	Limit-sure	Almost-sure
Maximisation finite horizon	undecidable	undecidable	EXPTIME-c
Minimisation finite horizon	PSPACE-hard \leq Min \leq EXPTIME		
Maximisation fixed horizon	PSPACE-c.	PTIME	PTIME
Minimisation fixed horizon	PSPACE-c	PSPACE-c	PSPACE-c

length n . A word $u \in \Sigma^*$ is a prefix of $v \in \Sigma^\infty$, written $u \leq v$, if $v = uw$ for some $w \in \Sigma^\infty$. The prefix is strict if $w \neq \varepsilon$. Given a countable set Z , a distribution on Z is a mapping $\mu : Z \rightarrow [0, 1]$ such that $\sum_{z \in Z} \mu(z) = 1$. The support of μ is $Supp(\mu) = \{z \in Z \mid \mu(z) > 0\}$. If $Supp(\mu) = \{z\}$ is a single element, μ is a Dirac distribution on z written $\mathbf{1}_z$. We denote by $\text{Dist}(S)$ the set of distributions on S .

2.1 Opacity for Markov chains

For the purpose of opacity questions, the models are equipped with a labelling function on states, called *observation function*, describing what an external observer can see. We first define observable Markov chains (MCs for short).

► **Definition 1** (Markov chains). An observable Markov chain (MC) over alphabet Σ is a tuple $\mathcal{M} = (S, p, \mathbf{O})$ where S is a countable set of states, $p : S \rightarrow \text{Dist}(S)$ is the transition function, and $\mathbf{O} : S \rightarrow \Sigma \cup \{\varepsilon\}$ is the observation function.

We write $p(s'|s)$ instead of $p(s)(s')$ to emphasise the probability of going to state s' conditioned by being in state s . Given a distribution μ_0 on S , we denote by $\mathcal{M}(\mu_0)$ the chain with initial distribution μ_0 . An infinite path of $\mathcal{M}(\mu_0)$ is a sequence of states $\rho = s_0 s_1 \dots \in S^\omega$ such that $\mu_0(s_0) > 0$ and for each $i \geq 0$, $p(s_{i+1}|s_i) > 0$. A finite path of length n is a prefix $\rho = s_0 s_1 \dots s_n$ of an infinite path, ending in state $\text{last}(\rho) = s_n$. We denote by $\text{Path}(\mathcal{M}(\mu_0))$ (resp. $\text{FPath}(\mathcal{M}(\mu_0))$) the set of infinite (finite) paths of $\mathcal{M}(\mu_0)$. The observation of path $\rho = s_0 s_1 \dots$ is the word $\mathbf{O}(\rho) = \mathbf{O}(s_0)\mathbf{O}(s_1)\dots \in \Sigma^\infty$. For a set R of paths, $\mathbf{O}(R) = \{\mathbf{O}(\rho) \mid \rho \in R\}$ and for a set W of observations, $\mathbf{O}^{-1}(W) = \{\rho \mid \mathbf{O}(\rho) \in W\}$. The observation function is called *non erasing* if $\mathbf{O}(S) \subseteq \Sigma$ (all states are visible).

A probability measure $\mathbf{P}_{\mathcal{M}(\mu_0)}$ is defined on $\text{Path}(\mathcal{M}(\mu_0))$, where the measurable sets are generated by the cylinders $\text{Cyl}(\rho)$, for $\rho \in \text{FPath}(\mathcal{M}(\mu_0))$, containing the infinite paths having ρ as prefix. Then $\mathbf{P}_{\mathcal{M}(\mu_0)}$ is inductively defined by: $\mathbf{P}_{\mathcal{M}(\mu_0)}(s) = \mu_0(s)$ for $s \in S$ and for $\rho' = \rho s'$, with $\text{last}(\rho) = s$, $\mathbf{P}_{\mathcal{M}(\mu_0)}(\text{Cyl}(\rho')) = \mathbf{P}_{\mathcal{M}(\mu_0)}(\text{Cyl}(\rho))p(s'|s)$. We sometimes write $\mathbf{P}_{\mathcal{M}(\mu_0)}(\rho)$ instead of $\mathbf{P}_{\mathcal{M}(\mu_0)}(\text{Cyl}(\rho))$ for $\rho \in \text{FPath}(\mathcal{M})$ and for $w \in \Sigma^*$, $\mathbf{P}_{\mathcal{M}(\mu_0)}(w)$ instead of $\mathbf{P}_{\mathcal{M}(\mu_0)}(\cup_{\rho \in \mathbf{O}^{-1}(w)} \text{Cyl}(\rho))$.

We consider here the particular case where the secret is given by a subset of states $\text{Sec} \subseteq S$ of the model: a (finite or infinite) path $s_0 s_1 \dots$ is secret if $s_i \in \text{Sec}$ for some i . We first define a probabilistic version of disclosure w.r.t. some $\varepsilon > 0$ to answer the question: Is there non-zero probability of observing some w that has probability more than $1 - \varepsilon$ of coming from a secret path?

► **Definition 2** (ε -Disclosure). Given an MC $\mathcal{M} = (S, p, \mathbf{O})$, an initial distribution μ_0 , $\text{Sec} \subseteq S$ and an observation $w \in \Sigma^*$, the proportion of secret paths with observation w is:

$$\text{P}_{\text{sec}\mathcal{M}(\mu_0)}(w) = \frac{\mathbf{P}_{\mathcal{M}(\mu_0)}(\{\rho \in \mathbf{O}^{-1}(w) \mid \rho \text{ is secret}\})}{\mathbf{P}_{\mathcal{M}(\mu_0)}(w)}.$$

For $\varepsilon > 0$, w is ε -min-disclosing if $\mathbf{P}_{\text{sec}\mathcal{M}(\mu_0)}(w) > 1 - \varepsilon$ and no prefix of w satisfies this inequality. Writing D_{\min}^ε for the set of ε -min-disclosing observations, the ε -disclosure is defined by $\text{Disc}^\varepsilon(\mathcal{M}(\mu_0)) = \sum_{w \in D_{\min}^\varepsilon} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$. The positive ε -disclosure problem consists in deciding if $\text{Disc}^\varepsilon(\mathcal{M}(\mu_0)) > 0$.

While being the most realistic notion of probabilistic disclosure, unfortunately the problem is already undecidable for Markov chains:

► **Theorem 3** (Undecidability of ε -disclosure). *The positive ε -disclosure problem is undecidable for MCs.*

Like in further proofs, we use a reduction from a problem on Probabilistic Automata (PA). Recall that a PA is a tuple $\mathcal{A} = (Q, q_0, \text{Act}, T, F)$ where Q is a finite set of states with $q_0 \in Q$ the initial state, Act is a finite set of actions, $T : Q \times \text{Act} \rightarrow \text{Dist}(Q)$ is the transition function and $F \subseteq Q$ is the set of final states.

For a finite path $\rho = q_0 \xrightarrow{a_1} q_1 \dots \xrightarrow{a_n} q_n$ of \mathcal{A} , the word $a_1 \dots a_n \in \text{Act}^*$ is called the *trace* of ρ and denoted by $tr(\rho)$. Writing $\text{FPath}(w, q) = \{\rho \in \text{FPath} \mid tr(\rho) = w \text{ and } \text{last}(\rho) = q\}$ for $w \in \text{Act}^*$ and $q \in Q$, we define $\mathbf{P}_{\mathcal{A}}^q(w) = \mathbf{P}_{\mathcal{A}}(\cup_{\rho \in \text{FPath}(w, q)} \text{Cyl}_\rho)$, $\mathbf{P}_{\mathcal{A}}^F(w) = \sum_{q \in F} \mathbf{P}_{\mathcal{A}}^q(w)$ and $val(\mathcal{A}) = \sup_{w \in \text{Act}^*} \mathbf{P}_{\mathcal{A}}^F(w)$.

Given a threshold $\theta \in [0, 1[$, we set $\mathcal{L}_{>\theta}(\mathcal{A}) = \{w \in \text{Act}^* \mid \mathbf{P}_{\mathcal{A}}^F(w) > \theta\}$. The strict emptiness problem for \mathcal{A} , asking whether this set is empty or not, is known to be undecidable for $\theta > 0$ [18]. The value 1 problem, asking whether $val(\mathcal{A}) = 1$ is undecidable as well [13].

Sketch of Proof. Given a PA \mathcal{A} , we build a Markov chain $\mathcal{M}_{\mathcal{A}}$ with initial distribution μ_0 and secret Sec such that for any ε , $0 < \varepsilon < 1$, $\mathcal{L}_{>1-\varepsilon}(\mathcal{A})$ is not empty iff $\text{Disc}^\varepsilon(\mathcal{M}_{\mathcal{A}}(\mu_0)) > 0$. ◀

This leads us to return to the simpler case where the disclosure is the probability of the set of paths leaking the secret, *i.e.*, such that *all* paths with the same observation are secret. The ω -disclosure (corresponding to measures in [3, 6, 4]) was defined for a Markov chain $\mathcal{M} = (S, p, \text{O})$ with initial distribution μ_0 by considering a measurable set of secret paths $\text{SPath} \subseteq \text{Path}(\mathcal{M}(\mu_0))$. Here, as mentioned above, SPath is $\text{Reach}(Sec)$, the set of infinite paths visiting a state from Sec , and an infinite observation $w \in \Sigma^\omega$ discloses the secret if all paths $\rho \in \text{O}^{-1}(w)$ are secret. Setting $\overline{\text{SPath}} = \text{Path}(\mathcal{M}(\mu_0)) \setminus \text{SPath}$, we define:

► **Definition 4** (ω -Disclosure). For an MC $\mathcal{M} = (S, p, \text{O})$, an initial distribution μ_0 and a subset $Sec \subseteq S$, with $\text{SPath} = \text{Reach}(Sec)$, the ω -disclosure is defined by:

$$\text{Disc}_\omega(\mathcal{M}(\mu_0)) = \mathbf{P}_{\mathcal{M}(\mu_0)}(\text{SPath} \setminus \text{O}^{-1}(\text{O}(\overline{\text{SPath}}))).$$

To obtain measures directly related to the finite observation of a possible attacker, we assume that $\mathcal{M} = (S, p, \text{O})$ is *convergent*: each infinite path ρ has an infinite observation $\text{O}(\rho) \in \Sigma^\omega$. Two measures can then be defined, when considering a fixed or finite horizon. In the former case, we consider a non-erasing function O to obtain real-time observations.

► **Definition 5** (Disclosure of MCs). Let $\mathcal{M} = (S, p, \text{O})$ be an MC, μ_0 an initial distribution and $Sec \subseteq S$. A finite observation $w \in \Sigma^*$ discloses the secret if all paths $\rho \in \text{O}^{-1}(w)$ are secret. It is min-disclosing if it discloses the secret and no strict prefix of w does.


n -disclosure : When O is non-erasing, we denote by D_n , for $n \in \mathbb{N}$, the set of disclosing observations of length n . The n -disclosure is $\text{Disc}_n(\mathcal{M}(\mu_0)) = \sum_{w \in D_n} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$;

Disclosure : Writing D_{\min} for the set of min-disclosing observations, the disclosure (w.r.t. finite horizon) is defined by $\text{Disc}(\mathcal{M}(\mu_0)) = \sum_{w \in D_{\min}} \mathbf{P}_{\mathcal{M}(\mu_0)}(w)$.

Note that if D is the set of disclosing observations, and $\mathcal{V}(\mu_0) = \cup_{w \in D} \cup_{\rho \in \text{O}^{-1}(w)} \text{Cyl}(\rho)$ the set of paths disclosing the secret, then $\text{Disc}(\mathcal{M}(\mu_0))$ is also equal to $\mathbf{P}_{\mathcal{M}(\mu_0)}(\mathcal{V}(\mu_0))$.

► Remark. Without loss of generality, we can assume that once a secret state has been reached by an execution, all subsequent states remain secret. For this, a new Markov chain $\mathcal{M}' = (S', p', O')$ is defined from \mathcal{M} by: $S' = Sec \uplus ((S \setminus Sec) \times \{0, 1\})$, where $(s, 0)$ represents state s where the secret has not been visited while $(s, 1)$ represents the opposite situation. The transitions are then duplicated accordingly: (1) $p'((s', i)|(s, i)) = p(s'|s)$ for all $s, s' \in S \setminus Sec$, and $i = 0, 1$, (2) $p'((s', 1)|s) = p(s'|s)$ for all $s \in Sec$, and $s' \in S \setminus Sec$, (3) $p'(s'|((s, i))) = p(s'|s)$ for all $s \in S \setminus Sec$, $i = 0, 1$, and $s' \in Sec$, and (4) $p'(s'|s) = p(s'|s)$ for all $s, s' \in Sec$. The observation function is extended by $O'((s, i)) = O(s)$ for all $s \in S \setminus Sec$ and $i = 0, 1$ and the new set of secrets is $Sec \uplus ((S \setminus Sec) \times \{1\})$. There is a one-to-one probability-preserving correspondence between the paths in \mathcal{M} and those in \mathcal{M}' .

We show that disclosure and ω -disclosure may be different by consider the infinitely branching MC of Figure 2, with initial distribution $\mathbf{1}_{q_0}$, $Sec = \{q_s\}$ hence $SPath = \{q_0 q_s^\omega\}$, $O(SPath) = o_1^+ o_2^\omega$. Then $Disc_\omega = \frac{1}{2}$ but since no finite observation is disclosing, $Disc = 0$.



■ Figure 2 An infinitely branching MC with $Sec = \{q_s\}$, $Disc_\omega = \frac{1}{2}$ and $Disc = 0$.

However, both notions coincide for convergent finitely branching MCs.

► Lemma 6 (Comparison of Disclosure Notions). Let $\mathcal{M} = (S, p, O)$ be a Markov chain, μ_0 an initial distribution and $Sec \subseteq S$. For $SPath = Reach(Sec)$, $Disc(\mathcal{M}(\mu_0)) \leq Disc_\omega(\mathcal{M}(\mu_0))$ and equality holds if \mathcal{M} is convergent and finitely branching.

2.2 Opacity for Markov Decision Processes

We now turn to MDPs that combine non determinism with probabilistic transitions.

► Definition 7 (MDP). A Markov Decision Process (MDP) over alphabet Σ is a tuple $M = (S, Act, p, O)$ where S is a finite set of states, $Act = \cup_{s \in S} A(s)$ where $A(s)$ is a finite non-empty set of actions for each state $s \in S$, $p : S \times Act \rightarrow Dist(S)$ is the (partial) transition function defined for (s, a) when $a \in A(s)$ and $O : S \rightarrow \Sigma \cup \{\varepsilon\}$ is the observation function.

As before, we write $p(s'|s, a)$ instead of $p(s, a)(s')$. Given an initial distribution μ_0 , an infinite path of M is a sequence $\rho = s_0 a_0 s_1 a_1 \dots$ where $\mu_0(s_0) > 0$ and $p(s_{i+1}|s_i, a_i) > 0$, for $s_i \in S$, $a_i \in A(s_i)$, for all $i \geq 0$. Finite paths (ending in a state) and observation of a path are defined like for Markov chains, and we use similar notations for the various sets of paths. For decidability and complexity results, we assume that all probabilities occurring in the model (transition probabilities and initial distribution) are rationals.

Nondeterminism is resolved by strategies. Given a finite path ρ with $last(\rho) = s$, a decision rule for ρ is a distribution on the possible actions in $A(s)$ chosen at this point. For such a decision rule δ , we write $p(s'|s, \delta) = \sum_{a \in A(s)} \delta(a) p(s'|s, a)$.

► **Definition 8** (Strategy). A strategy of MDP $M = (S, \text{Act}, p, O)$ with initial distribution μ_0 is a mapping $\sigma : \text{FPath}(M(\mu_0)) \rightarrow \text{Dist}(\text{Act})$ associating with ρ a decision rule $\sigma(\rho)$.

Given a strategy σ , a path $\rho = s_0a_0s_1a_1\dots$ of M is σ -compatible if for all i , $a_i \in \text{Supp}(\sigma(s_0a_0s_1a_1\dots s_i))$. A strategy σ is *deterministic* if $\sigma(\rho)$ is a Dirac distribution for each finite path ρ . In this case, we denote by $\sigma(\rho)$ the single action $a \in A(\text{last}(\rho))$ such that $\sigma(\rho) = \mathbf{1}_a$. A strategy σ is *observation-based* if for any finite path ρ , $\sigma(\rho)$ only depends on the observation sequence $O(\rho)$ and the current state $\text{last}(\rho)$, writing $\sigma(O(\rho), \text{last}(\rho))$ for $\sigma(\rho)$.

Let σ be a strategy and ρ be a σ -compatible path. We define B_ρ^σ the *belief* of ρ w.r.t. σ about states corresponding to the last observation as follows:

$$B_\rho^\sigma = \{s \mid \exists \rho' \text{ } \sigma\text{-compatible, } O(\rho') = O(\rho) \wedge s = \text{last}(\rho') \wedge O(s) \neq \varepsilon\}$$

A strategy σ is *belief-based* if for all ρ , $\sigma(\rho)$ only depends on its belief B_ρ^σ and its current state $\text{last}(\rho)$. Observe that a belief-based strategy is observation-based since B_ρ^σ only depends on $w = O(\rho)$. So we also write B_w^σ for B_ρ^σ . A strategy σ is *memoryless* if $\sigma(\rho)$ only depends on $\text{last}(\rho)$ for all ρ .

A strategy σ on $M(\mu_0)$ defines a (possibly infinite) Markov chain $M_\sigma(\mu_0)$ with set of states $\text{FPath}(M_\sigma(\mu_0))$ (the finite σ -compatible paths), that can be equipped with the observation function associating $O(\text{last}(\rho))$ with the finite path ρ . The transition function p_σ is defined for $\rho \in \text{FPath}(M_\sigma(\mu_0))$ and $\rho' = \rho a s'$ by $p_\sigma(\rho'|\rho) = \sigma(\rho)(a)p(s'|s, a)$ and we denote by $\mathbf{P}_{M_\sigma(\mu_0)}$ (or \mathbf{P}_σ for short when there is no ambiguity) the associated probability measure. Writing $\mathcal{V}_\sigma(\mu_0)$ for the set of paths disclosing the secret in $M_\sigma(\mu_0)$, we have $\text{Disc}(M_\sigma(\mu_0)) = \mathbf{P}_{M_\sigma(\mu_0)}(\mathcal{V}_\sigma(\mu_0))$.

Disclosure values for MDPs are defined according to the status of the strategies, by considering them as adversarial or cooperative with respect to the system (we only consider ε -disclosure for fixed horizon in view of the undecidability result of Theorem 3).

► **Definition 9** (Disclosure of an MDP). Given an MDP $M = (S, \text{Act}, p, O)$, an initial distribution μ_0 and a secret $\text{Sec} \subseteq S$, the maximal disclosure of Sec in M is $\text{disc}_{\max}(M(\mu_0)) = \sup_\sigma \text{disc}(M_\sigma(\mu_0))$ and the minimal disclosure is $\text{disc}_{\min}(M(\mu_0)) = \inf_\sigma \text{disc}(M_\sigma(\mu_0))$ for $\text{disc} \in \{\text{Disc}, \text{Disc}_n, \text{Disc}_n^\varepsilon\}$, $n \in \mathbb{N}$ and $0 < \varepsilon < 1$.

Note that the construction ensuring that once a secret state is visited, the path remains secret forever, extends naturally from Markov chains to MDPs. We consider only MDPs of this form in the sequel. We now show that for disclosure problems we can restrict strategies to observation-based ones.

► **Proposition 10** (Observation-based strategies). *Given an MDP, a secret and a strategy σ , there exists an observation-based strategy σ' with the same disclosure values.*

Erasing observations leads to technical and cumbersome developments. In order to avoid them in the design of procedures for the finite horizon case, we apply the preliminary transformation described in the next proposition. We precisely state the size of the transformed MDP in view of complexity results.

► **Proposition 11** (Avoiding erasing observations). *Given an MDP $M = (S, \text{Act}, p, O)$, an initial distribution μ_0 and a secret Sec , one can build in exponential time an MDP $M' = (S', \text{Act}', p', O')$, an initial distribution μ'_0 and a secret Sec' where O' is non-erasing and for $\text{disc} \in \{\text{Disc}_{\min}, \text{Disc}_{\max}\}$ $\text{disc}(M(\mu_0)) = \text{disc}(M'(\mu'_0))$. In addition, the size of S' , p' and μ'_0 is polynomial w.r.t. those of S , p and μ_0 . The size of Act' is polynomial w.r.t. the size of Act and exponential w.r.t. the size of S .*

We study the following problems over MDPs:

- **Computation problems.** The *value problem*: compute the disclosure and the *strategy problem*: compute an optimal strategy whenever it exists;
- **Quantitative decision problems.** The *disclosure problem*: Given M and a threshold $\theta \in [0, 1]$, is $\text{disc}(M) \bowtie \theta$? with $\bowtie = \geq$ for maximisation and $\bowtie = \leq$ for minimisation, and the more demanding *strategy decision problem*: does there exist a strategy σ such that $\text{disc}(M_\sigma) \bowtie \theta$?
- **Qualitative decision problems.** The *limit-sure disclosure problem*: the disclosure problem when $\theta = 1$ for maximisation and $\theta = 0$ for minimisation and the *almost-sure disclosure problem*: the strategy decision problem with the same restrictions.

For the complexity results regarding a fixed horizon n , we will assume that n is written in unary representation or bounded by a polynomial in the size of the model where the polynomial is independent of the model as done in classical studies (see for instance [17]).

3 Maximisation with finite horizon

While strategies may be randomised, this additional power is not necessary for maximisation:

► **Proposition 12** (Dominance of deterministic strategies). *Given an MDP, a secret and an observation-based strategy σ there exists a deterministic observation-based strategy σ' with greater or equal disclosure of the secret.*

Sketch of Proof. The Lemma 1 of [10] (or alternatively [14]) does not directly give the result as, contrary to the objectives used in their paper, disclosure depends on the strategy. However, as a disclosing path for a randomised strategy is also a disclosing path for a deterministic strategy that does not introduce new paths, we can use parts of their proof to show our result. ◀

An edge can be completely blocked by some strategy, modifying the set of paths that disclose the secret. This was illustrated in Figure 1a, where choosing action a in state q_0 removes the edges to q_3 and q_4 . This situation was excluded in the computation of the disclosure presented in [4], where the general problem was left open for Interval Markov Chains (IMCs). We answer negatively by proving undecidability of the disclosure problem, hence the disclosure cannot be computed in general. Undecidability also holds for limit-sure disclosure.

Writing \mathbb{I} for the set of intervals in $[0, 1]$, an IMC (with observation) is a tuple $M = (S, s_{init}, I, O)$ where S is the set of states, s_{init} is the initial state, $I : S \rightarrow \mathbb{I}^S$ associates with any state $s \in S$ a mapping from S into \mathbb{I} , and $O : S \rightarrow \Sigma \cup \{\varepsilon\}$ is the observation function. An IMC can be transformed into an (exponentially larger) MDP where actions are the basic feasible solutions of the linear program specified by the constraints associated with intervals [20]. Thus undecidability results for IMCs also hold for MDPs.

► **Theorem 13** (Undecidability of maximal finite horizon disclosure). *The maximal finite horizon disclosure problem is undecidable for MDPs, even when the secret is reached with probability 1 and for a non-erasing observation function. The maximal finite horizon disclosure problem when restricted to finite-memory strategies is also undecidable (with the same additional assumptions).*

Sketch of Proof. Starting from a PA \mathcal{A} , we build an IMC $M_{\mathcal{A}} = (S, s_0, I, O)$ such that there exists a word $w \in \{a, b\}^*$ with $\mathbf{P}_{\mathcal{A}}^F(w) > \frac{1}{2}$ if and only if $\text{Disc}_{\max}(M_{\mathcal{A}}) > \frac{1}{4}$. While similar

to the one of Theorem 3, the proof is more involved because the strategies must be taken into account. ◀

As a consequence, we obtain:

► **Corollary 14.** *The maximal finite horizon disclosure of an MDP cannot be computed.*

Using a reduction of the value 1 problem in PA, we also have:

► **Theorem 15** (Undecidability of maximal finite horizon limit-sure disclosure). *The maximal finite horizon limit-sure disclosure problem is undecidable for MDPs.*

Fortunately the maximal finite-horizon almost-sure disclosure problem is decidable. The proof relies on results for partially observable MDPs (POMDPs): a POMDP is an MDP where the strategies resolving the non determinism only depend on the observation sequence and do not take the current state into account.

► **Theorem 16** (Decidability of maximal finite-horizon almost-sure disclosure). *The maximal finite-horizon almost-sure disclosure problem in MDPs is EXPTIME-complete. Moreover, if the system is almost-surely disclosing, one can build a belief-based strategy with disclosure 1.*

Sketch of Proof. We reduce the almost-sure disclosure problem for maximisation in MDPs to almost-sure reachability in a POMDP. The POMDP we build is exponential in the size of the original MDP and the algorithm to solve almost-sure reachability is exponential in the size of the POMDP [11]. This gives an EXPTIME algorithm as those two exponentials do not stack. The hardness is obtained by a reduction from the safety problem in games with imperfect information that was shown to be EXPTIME-complete in [8]. ◀

4 Minimisation with finite horizon

Recall from the example illustrated in Figure 1a of introduction, that randomised strategies are necessary for minimisation. To address this issue we introduce *families of almost deterministic strategies* based on ε -decision rules, that will be used in the decision procedures.

► **Definition 17.** Let δ be the deterministic decision rule for state s selecting action $a \in A(s)$. Then $\delta_\varepsilon \in \text{Dist}(A(s))$ is a (randomised) ε -decision rule, said to *favour* a , and defined by:

1. If $|A(s)| > 1$ then $\delta_\varepsilon(a) = 1 - \varepsilon$ and for all $b \in A(s) \setminus \{a\}$, $\delta_\varepsilon(b) = \frac{\varepsilon}{|A(s)|-1}$;
2. Else $\delta_\varepsilon(a) = 1$.

► **Definition 18.** Let σ be an observation-based deterministic strategy. Then $\{\sigma_\varepsilon\}_{\varepsilon>0}$ is a family of *observation-based almost deterministic strategies* defined for any state s and $w \in \Sigma^n$, an observation of length $n \in \mathbb{N}$, by: $\sigma_\varepsilon(w, s) = \sigma(w, s)_{2^{-n\varepsilon}}$.

Using Proposition 11, we assume that the observation function O is non-erasing. The complexity of the transformation does not affect the results since the complexities are all polynomial in the number of actions. To compute the minimal disclosure value, we build from an MDP M , another MDP M_{\min} which is a “correct abstraction” (as stated by Proposition 19) for reducing minimal disclosure problems to minimal reachability problems, by enlarging states with the maximal belief that can occur independently of the action that has been selected.

Given a set of potential current states B and a new observation o , we define the maximal set of potential next states $\text{NextMax}(B, o)$ over decision rules applied to B by:

$$\text{NextMax}(B, o) = \{s' \in O^{-1}(o) \mid \exists s \in B \exists a \in A(s) p(s'|s, a) > 0\}$$

XX:10 Probabilistic Disclosure: Maximisation vs. Minimisation

Observe that given a family of almost deterministic strategies $\{\sigma_\varepsilon\}$ and a path ρ_{as} of M with $O(s) = o$, one has $B_{\rho_{as}}^{\sigma_\varepsilon} = \text{NextMax}(B_\rho^{\sigma_\varepsilon}, o)$. Then M_{\min} is formally defined as follows:

- S_{\min} , the set of states, is defined by: $S_{\min} = \{(s, B) \mid s \in B \subseteq O^{-1}(O(s))\}$;
- Let $(s, B) \in S_{\min}$. Then $A(s, B) = A(s)$;
- Let $(s, B), (s', B') \in S_{\min}$. If $B' = \text{NextMax}(B, O(s'))$ then $p((s', B') \mid (s, B), a) = p(s' \mid s, a)$ else $p((s', B') \mid (s, B), a) = 0$.

Given μ_0 an initial distribution over S , the associated initial distribution μ_{\min} over S_{\min} is defined by $\mu_{\min}(s, \text{Supp}(\mu_0) \cap O^{-1}(O(s))) = \mu_0(s)$ and $\mu_{\min}(s, B) = 0$ for all other B . We define the subset $\text{Avoid}(\text{Sec}) \subseteq S_{\min}$ by $\text{Avoid}(\text{Sec}) = \{(s, B) \mid B \subseteq \text{Sec}\}$.

► **Proposition 19.** *The minimal disclosure value for Sec in $M(\mu_0)$ is equal to the minimal probability to reach $\text{Avoid}(\text{Sec})$ in $M_{\min}(\mu_{\min})$. Furthermore it is asymptotically reached by a family of belief-based almost deterministic strategies.*

Since minimal reachability probability in MDPs can be computed in polynomial time we immediately obtain the first part of the next theorem. We establish the second part (PSPACE-hardness) in the proof of Theorem 23.

► **Theorem 20.** *The minimal disclosure value of $M(\mu_0)$ can be computed in EXPTIME. The associated decision problem is PSPACE-hard.*

We now turn to the existence of a strategy that achieves the minimal value and establish that it can be analysed without additional complexity. The main ingredient of the proof is an equation system over states of the MDP whose unique solution is the minimal reachability probability vector.

Notations. Given μ a distribution over states and $\vec{\delta}$ a vector of decision rules over states in the support of μ , we define $\text{NextDist}(\mu, \vec{\delta})$ the next distribution over S when applying $\vec{\delta}$ by:

$$\text{NextDist}(\mu, \vec{\delta})(s') = \sum_{s \in \text{Supp}(\mu)} \mu(s) p(s' \mid s, \vec{\delta}[s]) \quad \text{for any } s' \in S.$$

For a distribution μ over S and $o \in \Sigma$, we write $\mu(o)$ for $\mu(O^{-1}(o))$. If $\text{Supp}(\mu) \cap O^{-1}(o) \neq \emptyset$, the relative distribution μ_o over $O^{-1}(o)$ is defined by: $\mu_o(s) = \frac{\mu(s)}{\mu(o)}$ for $s \in O^{-1}(o)$ and $\mu_o(s) = 0$ otherwise.

We have $\text{Disc}_{\min}(M(\mu)) = \sum_{o \in \Sigma} \mu(o) \text{Disc}_{\min}(M(\mu_o))$. For use in the next proof, we define $\text{disc}^*(M(s, B))$ as the minimal disclosure value when starting in M in state s with belief B . Given some belief B and some decision rule vector $\vec{\delta}$ over B we introduce the possible successors of B when applying $\vec{\delta}$: $\text{Next}(B, \vec{\delta}) = \{s' \mid \exists s \in B \exists a \in \text{Supp}(\vec{\delta}[s]) p(s' \mid s, a) > 0\}$ and $\text{Next}(B, \vec{\delta}, o) = \text{Next}(B, \vec{\delta}) \cap O^{-1}(o)$.

► **Theorem 21.** *The existence of a strategy that achieves the minimal disclosure value can be decided in EXPTIME. In the positive case, this strategy can be computed in EXPTIME.*

Proof. The algorithm simultaneously solves the existence and the synthesis problem.

Using proposition 19, the algorithm computes for all $(s, B) \in S_{\min}$, $\text{disc}^*(M(s, B))$.

Then it maintains a set Win of beliefs initially set to all beliefs from which it iteratively eliminates items and stops when no more elimination is possible.

Given $B \in \text{Win}$, it looks for a decision rule vector $\vec{\delta}$ over B such that:

- for all $o \in O(\text{Next}(B, \vec{\delta}))$, $\text{Next}(B, \vec{\delta}, o) \in \text{Win}$;
- for all $s \in B$, $\text{disc}^*(M(s, B)) = \sum_{o \in \Sigma} \sum_{s' \in O^{-1}(o)} p(s' \mid s, \vec{\delta}[s]) \text{disc}^*(M(s', \text{Next}(B, \vec{\delta}, o)))$.

If such a $\vec{\delta}$ does not exist then B is eliminated from Win . Each iteration can be performed in polynomial time w.r.t. $|S_{\min}|$ and the number of iterations is at most $|S_{\min}|$. Observe that when a belief is eliminated, it should not be “reached” by a strategy that obtains the minimal disclosure value. So the elimination is sound.

When the elimination stops, the algorithm answers positively iff for all $o \in \mathcal{O}(Supp(\mu_0))$, $Supp(\mu_0) \cap \mathcal{O}^{-1}(o) \in Win$. Thus by the soundness of the elimination step, if the answer is negative there is no optimal strategy for minimal disclosure value.

If the answer is positive, let us consider the belief-based strategy σ defined by applying the decision rules obtained during the last iteration of the algorithm. On the one hand, under σ when visiting a state s with belief B such that $disc^*(M(s, B)) = 0$, one never leaves such kind of pairs of states and beliefs. So the secret is never disclosed, showing that the disclosure value obtained by σ for such (s, B) is null. Under σ the disclosure value of all the other pairs of state and belief fulfill the equations of the elimination step. It is known that the single solution of this system is the vector of minimal reachability probabilities of Avoid in $M_{\min}(\mu_{\min})$ (see [1] for instance) which yields the result. ◀

5 Fixed horizon problems

5.1 Maximal disclosure

In order to compute the value of the maximal disclosure within a fixed horizon, one could build the POMDP described in the proof of Theorem 16 then use pre-existing results on POMDPs. This would result in an EXPTIME algorithm, whereas we obtain in the result below an algorithm with a better complexity in PSPACE.

► **Theorem 22** (Computation of the maximal disclosure value within fixed-horizon). *The fixed-horizon maximal value (when the horizon n is described in unary representation) is computable in PSPACE and the fixed-horizon maximal disclosure problem is PSPACE-complete.*

Sketch of proof - value and membership. We first order the observation alphabet Σ . Then a non deterministic decision procedure operating in PSPACE orderly reads every observation sequence of length n while maintaining the sets of states that were possible after every prefix of this observation, the actions that were chosen nondeterministically in those states and values used in the computation of the disclosure. The information kept is of polynomial size and when every observation has been read, one of the values computed will be exactly the disclosure of the system at time n . We then remove the non determinism using Savitch’s Theorem. In order to get the value we observe that we can compute the polynomially sized denominator of this value and then we proceed by iterations of the decision algorithm. ◀

As can be seen in the proof, the optimal strategy could be computed when solving the value problem. However the size of this strategy may be exponential due to the beliefs and thus this strategy is computable in EXPTIME.

For the hardness result, we reduce the truth of a Quantified Boolean Formula (QBF). Recall that QBFs are extension of propositional formulas where boolean variables can be quantified. Syntactically, the formulas are described by the following grammar:

$$\begin{aligned} \phi &::= \psi \mid \exists x.\phi \mid \forall x.\phi \\ \psi &::= x \mid \psi \wedge \psi \mid \psi \vee \psi \mid \neg\psi \mid \mathbf{true} \end{aligned}$$

A QBF is *closed* if every boolean variable is bound by a quantifier. Deciding if a closed QBF is equivalent to \mathbf{true} is PSPACE-hard [21].

XX:12 Probabilistic Disclosure: Maximisation vs. Minimisation

Sketch of proof - hardness. Given ϕ a closed QBF (w.l.o.g. in 3CNF with n variables and m clauses), we build an MDP M such that ϕ is true iff the disclosure of M is greater or equal to $\frac{1}{2^{2n}}$ in $2(n+m)+3$ steps. In fact, $\frac{1}{2^{2n}}$ is exactly the measure of paths reaching the secret in $2(n+m)+3$ steps, thus every path reaching the secret must be disclosing. Such a path discloses the secret iff a boolean variable of ϕ and its negation (x and $\neg x$ for example) do not occur in its observation.

In M , during the first $2n$ steps, an assignment will be ‘given’ to each boolean variable: (i) for each existentially quantified boolean variable x , the strategy chooses whether x or $\neg x$ occurs in the observation and (ii) for each universally quantified boolean variable y , by a random choice with probability $\frac{1}{2}$. During the last $2m$ steps, the strategy must trigger a boolean variable in every clause of ϕ so that if a clause is not satisfied by the current assignment, then a boolean variable will be observed as both true and false during the path. Thus the observation would not disclose the secret. ◀

The existence of an optimal strategy here implies that the limit-sure and the almost-sure problem are equivalent. Moreover, the secret being revealed with probability 1 in a given number of steps implies that every path reaches the secret in this number of steps. Therefore the almost-sure problem can be seen as a reachability problem in an MDP which can be solved in polynomial time.

The proof of hardness can be adapted for ε -disclosure, but the algorithm for membership can not be directly applied. The ε -disclosure could however be computed by minimising an exponential system of equations, resulting in an exponential time algorithm.

5.2 Minimal disclosure

The proofs of the two first assertions of the next theorem are similar to the proof of Theorem 22. However in order to get the same complexity for the last assertion, we establish that when a randomised decision rule must be selected in the optimal strategy, it can always be uniformly distributed over its support.

► **Theorem 23** (Minimal disclosure within fixed horizon). *The fixed horizon minimal value is computable in PSPACE. The fixed horizon minimal disclosure problem is PSPACE-complete. In addition, the strategy decision problem is also decidable in PSPACE.*

Contrary to the case of maximisation, the above proof implies PSPACE-completeness for the limit-sure and almost-sure problem for minimisation.

The remark on ε -disclosure of the previous subsection holds again here.

6 Conclusion

We revisit the problems of disclosure for MDPs by (1) taking into account general actions contrary to previous work and (2) considering both maximisation and minimisation problems. We almost fully characterise the decidability and complexity of those problems establishing an asymmetry between minimisation and maximisation problems: the former ones being easier although they require families of randomised strategies for reaching the optimal value.

There remains a complexity gap (PSPACE versus EXPTIME) for the finite-horizon minimisation problem that we want to fill. From a qualitative point of view, observe that disclosure is a hyperproperty as its truth value is defined relatively to a set of paths. Thus we plan to address such kinds of properties in a restricted setting in order to get other decidability results. Another direction would be to strengthen the requirement for approximate ε -disclosure to regain decidability within finite horizon.

References


- 1 C. Baier and J.-P. Katoen. *Principles of model checking*. MIT Press, 2008.
- 2 B. Barak, O. Goldreich, R. Impagliazzo, S. Rudich, A. Sahai, S. P. Vadhan, and K. Yang. On the (Im)Possibility of Obfuscating Programs. In *Proc. of CRYPTO'01*, pages 1–18. Springer, 2001.
- 3 B. Bérard, K. Chatterjee, and N. Sznajder. Probabilistic Opacity for Markov decision processes. *Information Processing Letters*, 115(1):52–59, 2015.
- 4 B. Bérard, O. Kouchnarenko, J. Mullins, and M. Sassolas. Preserving opacity on interval Markov chains under simulation. In *Proc. of WODES'16*, pages 319–324. IEEE, 2016.
- 5 B. Bérard, J. Mullins, and M. Sassolas. Quantifying opacity. In *Proc. of QEST'10*, pages 263–272. IEEE, 2010.
- 6 B. Bérard, J. Mullins, and M. Sassolas. Quantifying opacity. *Mathematical Structures in Computer Science*, 25(2):361–403, 2015.
- 7 N. Bertrand, S. Haddad, and E. Lefaucheux. Foundation of Diagnosis and Predictability in Probabilistic Systems. In *Proc. of FSTTCS'14*, volume 29 of *LIPICs*, pages 417–429. Leibniz-Zentrum für Informatik, 2014.
- 8 D. Berwanger and L. Doyen. On the power of imperfect information. In *Proc. of FSTTCS'08*, volume 2 of *LIPICs*, Bangalore, India, 2008. Leibniz-Zentrum fuer Informatik.
- 9 J. W. Bryans, M. Koutny, L. Mazaré, and P. Y. A. Ryan. Opacity Generalised to Transition Systems. *International Journal of Information Security*, 7(6):421–435, 2008.
- 10 K. Chatterjee, L. Doyen, H. Gimbert, and T. A. Henzinger. Randomness for free. In *Proc. of MFCS'10*, volume 6281 of *LNCS*, pages 246–257. Springer, 2010.
- 11 K. Chatterjee, L. Doyen, and T. A. Henzinger. Qualitative Analysis of Partially-Observable Markov Decision Processes. In *Proc. of MFCS'10*, pages 258–269. Springer, 2010.
- 12 K. Chatterjee, K. Sen, and T. A. Henzinger. Model-Checking omega-Regular Properties of Interval Markov Chains. In *Proc. of FOSSACS'08*, volume 4962 of *LNCS*, pages 302–317. Springer, 2008.
- 13 H. Gimbert and Y. Oualhadj. Probabilistic Automata on Finite Words: Decidable and Undecidable Problems. In *Proc. of ICALP'10, Part II*, volume 6199 of *LNCS*, pages 527–538. Springer, 2010.
- 14 J. Goubault-Larrecq and R. Segala. Random measurable selections. In *Horizons of the Mind. A Tribute to Prakash Panangaden*, volume 8464 of *LNCS*, pages 343–362. Springer, 2014.
- 15 D. J. D. Hughes and V. Shmatikov. Information Hiding, Anonymity and Privacy: a Modular Approach. *Journal of Computer Security*, 12(1):3–36, 2004.
- 16 L. Mazaré. Decidability of Opacity with Non-Atomic Keys. In *Proc. of FAST'04*, volume 173 of *IFIP*, pages 71–84. Springer, 2005.
- 17 C. H. Papadimitriou and J. N. Tsitsiklis. The complexity of Markov decision processes. *Math. Oper. Res.*, 12(3):441–450, 1987.
- 18 A. Paz. *Introduction to Probabilistic Automata*. Academic Press, 1971.
- 19 A. Saboori and Ch. N. Hadjicostis. Current-State Opacity Formulations in Probabilistic Finite Automata. *IEEE Transactions on Automatic Control*, 59(1):120–133, 2014.
- 20 K. Sen, M. Viswanathan, and G. Agha. Model-Checking Markov Chains in the Presence of Uncertainties. In *Proc. of TACAS'06*, volume 3920 of *LNCS*, pages 394–410. Springer, 2006.
- 21 M. Sipser. *Introduction to the theory of computation*. Thomson Course Technology, 2006.

A Semantics


Recall that the strict emptiness and the value 1 problems are undecidable for PA already for an alphabet with two letters. Hence in the various reductions we use the alphabet $\{a, b\}$.

► **Theorem 3** (Undecidability of ε -disclosure). *The positive ε -disclosure problem is undecidable for MCs.*

Proof. Given a PA $\mathcal{A} = (Q, q_0, \{a, b\}, T, F)$ that we suppose complete without loss of generality, we first transform \mathcal{A} into an incomplete MC $\hat{\mathcal{A}}$ where $\{a, b, \#$ is the observation alphabet (an illustration is given in Figure 3). The set of states is $\hat{Q} = Q \cup \{q_c \mid q \in Q \wedge c \in \{a, b\}\}$, with initial distribution $\mathbf{1}_{q_0}$. The observation function \hat{O} is defined by $\hat{O}(q) = b$ and $\hat{O}(q_c) = c$ for $q \in Q$ and $c \in \{a, b\}$. The transition function \hat{p} is defined for $q, q' \in Q$ and $c \in \{a, b\}$ by $\hat{p}(q' \mid q, c) = T(q' \mid q, c)$ and $\hat{p}(q_c \mid q) = \frac{1}{4}$.



■ **Figure 3** From PA \mathcal{A} to incomplete MC $\hat{\mathcal{A}}$.



■ **Figure 4** Reduction to the positive ε -disclosure problem.

We now build the MC $\mathcal{M}_{\mathcal{A}} = (S, p, O)$ over alphabet $\{a, b, \#, \#\}$ by adding two states to complete $\hat{\mathcal{A}}$ (see Figure 4 where doubly circled state q_f is a final state of \mathcal{A}):

- $S = \{s_{pub}, s_{sec}\} \cup \hat{Q}$, with $Sec = \{s_{sec}\}$;
- The function p is obtained from \hat{p} by adding the transitions: For $q \in F$, $p(s_{sec} \mid q) = \frac{1}{2}$, for $q \in Q \setminus F$, $p(s_{pub} \mid q) = \frac{1}{2}$, and $p(s_{pub} \mid s_{pub}) = p(s_{sec} \mid s_{sec}) = 1$;
- O extends \hat{O} by $O(s_{sec}) = O(s_{pub}) = \#\}$.

We now prove that, given $\varepsilon \in]0, 1[$, \mathcal{A} admits a word with probability strictly greater than $1 - \varepsilon$ iff $Disc^\varepsilon(\mathcal{M}(\mu_0)) > 0$. First assume that there exists a word $w = a_1 \dots a_n \in \{a, b\}^*$ with $\mathbf{P}_{\mathcal{A}}^F(w) > 1 - \varepsilon$. Then w corresponds to a non secret path with observation $\hat{w} = ba_1b \dots a_nb$ in $\mathcal{M}_{\mathcal{A}}$ and $\mathbf{P}_{sec, \mathcal{M}(\mu_0)}(\hat{w}\#) = \mathbf{P}_{\mathcal{A}}^F(w) > 1 - \varepsilon$, which implies $Disc^\varepsilon(\mathcal{M}(\mu_0)) > 0$. Conversely, if $Disc^\varepsilon(\mathcal{M}(\mu_0)) > 0$, then there exists an observation w' in $(\{a, b, \#, \#\})^*$ such

that $\text{Psec}_{\mathcal{M}(\mu_0)}(w') > 1 - \varepsilon$. In this case, w' belongs to $\text{ba}_1\text{b} \dots \text{a}_n\text{b}\#\#^*$ for some $w = a_1 \dots a_n$ such that $\text{Psec}_{\mathcal{M}(\mu_0)}(w') = \text{Psec}_{\mathcal{M}(\mu_0)}(\text{ba}_1\text{b} \dots \text{a}_n\text{b}\#) = \mathbf{P}_{\mathcal{A}}^F(w)$ and $\mathcal{L}_{>1-\varepsilon}(\mathcal{A})$ is not empty. \blacktriangleleft

► **Lemma 6** (Comparison of Disclosure Notions). *Let $\mathcal{M} = (S, p, \mathbf{O})$ be a Markov chain, μ_0 an initial distribution and $\text{Sec} \subseteq S$. For $\text{SPath} = \text{Reach}(\text{Sec})$, $\text{Disc}(\mathcal{M}(\mu_0)) \leq \text{Disc}_\omega(\mathcal{M}(\mu_0))$ and equality holds if \mathcal{M} is convergent and finitely branching.*

Proof. We first establish:

Claim. If \mathcal{M} is convergent and finitely branching, then the set of paths ρ such that $\mathbf{O}(\rho)$ is of length n is finite for any $n > 0$.

In this context, similarly to [7], we define a *signaling* path as a finite path ρ such that $\mathbf{O}(\text{last}(\rho)) \neq \varepsilon$ and we denote by SP the set of signaling paths. We first prove the claim for such paths. By induction, we start with $n = 1$ and consider the tree formed by the set $O_1 = \{\rho \in \text{SP} \mid |\mathbf{O}(\rho)| = 1\}$ by sharing common prefixes. Internal nodes of this tree correspond to unobservable states while all leaves are observable. Since the chain is finitely branching, the tree is of bounded degree. By contradiction, assume that the tree is infinite. König's lemma yields an infinite branch containing only unobservable states, which contradicts the hypothesis of convergence. The induction step is obtained by a similar argument. The general case also follows the same lines since any path ρ with an observation of length n has a signaling prefix ρ' with the same observation. Hence the tree of suffixes of ρ' must also be finite.

We now prove that $\mathcal{V} = \bigcup_{w \in D} \bigcup_{\rho \in \mathbf{O}^{-1}(w)} \text{Cyl}(\rho)$ is contained in $\text{SPath} \setminus \mathbf{O}^{-1}(\mathbf{O}(\overline{\text{SPath}}))$. Let ρ_1 be an infinite path in \mathcal{V} . Then there is a disclosing observation $w_1 \in \Sigma^*$ and a signaling prefix ρ'_1 of ρ_1 such that $\mathbf{O}(\rho'_1) = w_1$ and ρ'_1 is secret. For any infinite path ρ_2 such that $\mathbf{O}(\rho_1) = \mathbf{O}(\rho_2)$, the observation w_1 is also a prefix of $\mathbf{O}(\rho_2)$, hence there is a finite signaling prefix ρ'_2 of ρ_2 such that $\mathbf{O}(\rho'_2) = w_1$. Since w_1 is disclosing, ρ'_2 is also secret, hence ρ_1 belongs to $\text{SPath} \setminus \mathbf{O}^{-1}(\mathbf{O}(\overline{\text{SPath}}))$ and $\text{Disc}(\mathcal{M}(\mu_0)) \leq \text{Disc}_\omega(\mathcal{M}(\mu_0))$.

For the converse inclusion, let ρ be an infinite path in $\text{SPath} \setminus \mathbf{O}^{-1}(\mathbf{O}(\overline{\text{SPath}}))$ with observation $\mathbf{O}(\rho) = w = o_1 o_2 \dots \in \Sigma^\omega$. We prove by contradiction that there is a finite disclosing prefix \hat{w} of w and a signaling prefix $\hat{\rho}$ of ρ such that $\rho \in \text{Cyl}(\hat{\rho})$ and $\mathbf{O}(\hat{\rho}) = \hat{w}$. Otherwise, for any $n \geq 1$, $w_n = o_1 \dots o_n$ is not disclosing and there exists a signaling path ρ_n such that $\mathbf{O}(\rho_n) = w_n$ but ρ_n is not secret. The set $\mathcal{T} = \{\rho' \in \text{SP} \mid \exists n \rho' \leq \rho_n\}$ of all signaling prefixes of the ρ_n 's form a tree: the root of the tree is ε and the nodes at level k are the prefixes with observation w_k : $\{\rho' \in \mathcal{T} \mid |\mathbf{O}(\rho')| = k\}$. A node ρ'' is a child of ρ' if $|\mathbf{O}(\rho')| = w_k$, $|\mathbf{O}(\rho'')| = w_{k+1}$ for some k and $\rho' \leq \rho''$. From the claim, we know that \mathcal{T} is of bounded degree. Assuming that it is infinite, König's lemma again yields an infinite branch ρ_∞ such that each prefix of length k is not secret and has observation w_k . Hence ρ_∞ is not secret and has observation $\mathbf{O}(\rho_\infty) = w$, which is a contradiction. \blacktriangleleft

► **Proposition 10** (Observation-based strategies). *Given an MDP, a secret and a strategy σ , there exists an observation-based strategy σ' with the same disclosure values.*

Proof. Let $\mathcal{M} = (S, \text{Act}, p, \mathbf{O})$ be an MDP with initial distribution μ_0 . Given a strategy σ , we note \mathbf{P}_σ instead of $\mathbf{P}_{\mathcal{M}_\sigma(\mu_0)}$ for the associated probability measure. For an observation $w \in \Sigma^*$ and a state $s \in S$, we define the sets (which are finite from the claim in the lemma above) $R(w, s) = \{\rho \in \text{FPATH}(\mathcal{M}_\sigma(\mu_0)) \mid \mathbf{O}(\rho) = w \wedge \text{last}(\rho) = s\}$.

We now define a mapping $\hat{\sigma}$ from $\Sigma^* \times S$ into $\text{Dist}(\text{Act})$ by

$$\hat{\sigma}(w, s) = \frac{1}{\sum_{\rho \in R(w, s)} \mathbf{P}_\sigma(\rho)} \sum_{\rho \in R(w, s)} \mathbf{P}_\sigma(\rho) \sigma(\rho)$$

XX:16 Probabilistic Disclosure: Maximisation vs. Minimisation

and a new strategy σ' for a finite path ρ by $\sigma'(\rho) = \hat{\sigma}(\mathbf{O}(\rho), \text{last}(\rho))$. We prove that $P_{\sigma'}(R(w, s)) = P_{\sigma}(R(w, s))$ for any observation w and any state s , which entails equality of disclosure.

Partitioning the set of states into $S = S_{ob} \uplus S_u$ where $S_u = \mathbf{O}^{-1}(\varepsilon)$, we can assume a topological sort on the subgraph obtained by removing all edges in $S \times S_{ob}$ (this subgraph is acyclic due to the hypothesis of convergence). We denote by η the mapping numbering all states according to the total order and we proceed to prove the claim above by a joint induction on the pairs (w, s) using $|w|$ and $\eta(s)$.

For the base cases, we need to establish the property for $w = \varepsilon$ with $s \in S_u$, and for $w \in \Sigma$ with $s \in S_{ob}$, where $\mu_0(s) > 0$ in both cases.

Case 1. By induction on $\eta(s)$, we consider a state $s \in S_u$ such that $\eta(s) = \min_{s' \in S_u}(\eta(s'))$. Then $P_{\sigma'}(R(\varepsilon, s)) = \mu_0(s) = P_{\sigma}(R(\varepsilon, s))$. Assuming the property holds for (ε, s) with $\eta(s) \leq n$, we prove it for s' with $\eta(s') = n + 1$. We have:

$$\mathbf{P}_{\sigma'}(R(\varepsilon, s')) = \mu_0(s') + \sum_{s \in S_u, \eta(s) < \eta(s')} \sum_{a \in A(s)} p(s'|s, a) \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma'}(\rho) \sigma'(\rho)(a)$$

and using the definition of σ' yields:

$$\begin{aligned} \mathbf{P}_{\sigma'}(R(\varepsilon, s')) &= \mu_0(s') + \sum_{s \in S_u, \eta(s) < \eta(s')} \sum_{a \in A(s)} p(s'|s, a) \hat{\sigma}(\varepsilon, s)(a) \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma'}(\rho) \\ &= \mu_0(s') + \sum_{s \in S_u, \eta(s) < \eta(s')} \sum_{a \in A(s)} p(s'|s, a) \frac{\sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma}(\rho) \sigma(\rho)(a)}{\sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma}(\rho)} \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma'}(\rho). \end{aligned}$$

Applying the induction hypothesis on (ε, s) yields $\sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma'}(\rho) = \mathbf{P}_{\sigma'}(R(\varepsilon, s)) = \mathbf{P}_{\sigma}(R(\varepsilon, s)) = \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma}(\rho)$ thus:

$$\mathbf{P}_{\sigma'}(R(\varepsilon, s')) = \mu_0(s') + \sum_{s \in S_u, \eta(s) < \eta(s')} \sum_{a \in A(s)} p(s'|s, a) \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma}(\rho) \sigma(\rho)(a) = \mathbf{P}_{\sigma}(R(\varepsilon, s')).$$

Case 2. We now consider $w = o \in \Sigma$ and $s' \in S_{ob}$, hence $\mathbf{O}(s') = o$. Then:

$$\mathbf{P}_{\sigma'}(R(o, s')) = \mu_0(s') + \sum_{s \in S_u} \sum_{a \in A(s)} p(s'|s, a) \sum_{\rho \in R(\varepsilon, s)} \mathbf{P}_{\sigma'}(\rho) \sigma'(\rho)(a)$$

and a reasoning similar as above yields the result.

For the induction step, we first need to prove the property for (w, s') with $s' \in S_u$, assuming it holds for all (w, s) with $s \in S_{ob}$ and for all (w, s) with $s \in S_u$ and $\eta(s) < \eta(s')$. Then we have:

$$\mathbf{P}_{\sigma'}(R(w, s')) = \sum_{\substack{s \in S_{ob} \\ s \in S_u, \eta(s) < \eta(s')}} \sum_{a \in A(s)} p(s'|s, a) \sum_{\rho \in R(w, s)} \mathbf{P}_{\sigma'}(\rho) \sigma'(\rho)(a)$$

and we can conclude along the same lines as above.

Finally, we consider (w', s') with $w' = wo$ and $s \in S_{ob}$, with:

$$\mathbf{P}_{\sigma'}(R(w', s')) = \sum_{s \in S} \sum_{a \in A(s)} p(s'|s, a) \sum_{\rho \in R(w, s)} \mathbf{P}_{\sigma'}(\rho) \sigma'(\rho)(a)$$

which again implies the desired result. ◀

► **Proposition 11** (Avoiding erasing observations). *Given an MDP $M = (S, \text{Act}, p, O)$, an initial distribution μ_0 and a secret Sec , one can build in exponential time an MDP $M' = (S', \text{Act}', p', O')$, an initial distribution μ'_0 and a secret Sec' where O' is non-erasing and for $\text{disc} \in \{\text{Disc}_{\min}, \text{Disc}_{\max}\}$ $\text{disc}(M(\mu_0)) = \text{disc}(M'(\mu'_0))$. In addition, the size of S' , p' and μ'_0 is polynomial w.r.t. those of S , p and μ_0 . The size of Act' is polynomial w.r.t. the size of Act and exponential w.r.t. the size of S .*

Proof. We first build the new MDP and then explain the correspondence between strategies in both models.

Construction of the MDP. We start from MDP $M = (S, \text{Act}, p, O)$ with $\text{Act} = \cup_{s \in S} A(s)$, observation alphabet Σ , and a set of secret states $\text{Sec} \subseteq S$. Choosing a fresh observation symbol \sharp and a fresh state s_\sharp , we build a MDP $M' = (S', \text{Act}', p', O')$ with set of states $S' = \{s_\sharp\} \cup (S \setminus O^{-1}(\varepsilon))$, and observation alphabet $\Sigma \cup \{\sharp\}$, where the initial distribution is $\mathbf{1}_{s_\sharp}$.

The observation function O' is defined by $O'(s_\sharp) = \sharp$ and $O'(s) = O(s)$ otherwise. The set of actions of M' is $\text{Act}' = \text{DR}$ where DR is the set of vectors of deterministic decision rules $\vec{\delta}$ over S , i.e. such that $\vec{\delta}(s) \in A(s)$.

For a path $\rho = s_0 a_1 \dots a_n s_n$, we write $\pi(\rho) = \prod_{i=1}^n p(s_i | s_{i-1}, a_i)$ and $\text{first}(\rho) = s_0$. Given a state $s \neq s_\sharp$, an action $\vec{\delta} \in \text{DR}$ and an observable state s' , we consider the finite set $\hat{P}(s, \vec{\delta}, s')$ of paths $\rho = s_0 a_1 \dots a_n s_n$ starting in $s_0 = s$ and ending in $s_n = s'$ such that all intermediate states are unobservable (and so distinct by convergence of the MDP) and for each i , $1 < i \leq n$, $a_i = \vec{\delta}(s_{i-1})$. We set $\hat{p}(s' | s, \vec{\delta}) = \sum_{\rho \in \hat{P}(s, \vec{\delta}, s')} \pi(\rho)$. Note that $\hat{P}(s, \vec{\delta}, s')$ may include paths like $\rho = s \vec{\delta}(s) s'$. The transition function p' is defined by: $p'(s' | s, \vec{\delta}) = \hat{p}(s' | s, \vec{\delta})$.

Since the initial distribution μ'_0 of M' is Dirac on s_\sharp , transitions from this state are defined similarly as above but they must take into account the initial distribution μ_0 of M . Given an observable state $s \in S$ and $\vec{\delta} \in \text{DR}$, the set $\hat{P}(s_\sharp, \vec{\delta}, s)$ contains the finite paths $\rho = s_0 a_1 s_1 \dots a_n s_n$ starting from some $s_0 \in \text{Supp}(\mu_0)$ and ending in $s_n = s$ such that all states s_0, \dots, s_{n-1} are unobservable, and $a_i = \vec{\delta}(s_{i-1})$ for all i , $1 \leq i \leq n$. In this case, we set $\hat{p}(s | s_\sharp, \vec{\delta}) = \sum_{\rho \in \hat{P}(s_\sharp, \vec{\delta}, s)} \mu_0(\text{first}(\rho)) \pi(\rho)$. If $s \in \text{Supp}(\mu_0)$, the set $\hat{P}(s_\sharp, \vec{\delta}, s)$ contains the path reduced to $\rho = s$. Then $p'(s | s_\sharp, \vec{\delta}) = \hat{p}(s | s_\sharp, \vec{\delta})$.

In order to efficiently compute the transition function of some $\vec{\delta}$, one uses a topologic sort of the unobservable states thanks to the convergence hypothesis, and then compute the probability from observable states to reach first the unobservable states topologically sorted and then the observable states. This allows for a polynomial time computation of the transition function of $\vec{\delta}$. The building of M' is thus polynomial in number of states and of the probabilities except for Act' which is polynomial in $|\text{Act}|$ and exponential in $|S|$.

Correspondence between strategies. The construction above ensures that any path $\rho' = s_\sharp \vec{\delta}_1 s_1 \dots \vec{\delta}_k s_k$ of M' corresponds to the set of paths ρ of M containing the sequence $s_1 \dots s_k$ of observable states, with subpaths $\rho_1 \in \hat{P}(s_\sharp, \vec{\delta}_1, s_1)$ and $\rho_i \in \hat{P}(s_{i-1}, \vec{\delta}_i, s_i)$ for $1 < i \leq k$. All paths in the set have the same observation $w = O(s_1) \dots O(s_k)$ with $O'(\rho') = \sharp w$.

To show that disclosure over finite horizon is the same in both MDP, we establish correspondences between the strategies of M and M' and the associated disclosure value. From Proposition 10, we can restrict to observation-based strategies.

Let σ' be an observation based strategy of M' , defined on $\sharp \Sigma^* \times S'$. Given an observation $w \in \Sigma^*$ there exists $\vec{\delta}$ such that for all state $s \in S'$, we have $\sigma'(\sharp w, s) = \vec{\delta}$. We define for $s \in S$ $\sigma(w, s) = \vec{\delta}(s)$. Conversely, given an observation-based strategy σ of M , we build an observation based strategy σ' of M' as follows: Given $w \in \Sigma^*$, we define the mapping

$\sigma'(\#w) : S \rightarrow \text{Act}$ by $\sigma'(\#w)(s) = \sigma(w, s)$ for any $s \in S$.

Then, writing \mathbf{P}_σ (resp. $\mathbf{P}_{\sigma'}$) instead of $\mathbf{P}_{M_\sigma(\mu_0)}$ (resp. $\mathbf{P}_{M_{\sigma'}(\mu'_0)}$), and defining for $w \in \Sigma^*$ and $s \in S \setminus \mathcal{O}^{-1}(\varepsilon)$, $R(w, s) = \{\rho \in \text{FPath}(M_\sigma(\mu_0)) \mid \mathcal{O}(\rho) = w \wedge \text{last}(\rho) = s\}$ and $R'(w, s) = \{\rho' \in \text{FPath}(M_{\sigma'}(\mu'_0)) \mid \mathcal{O}(\rho') = \#w \wedge \text{last}(\rho') = s\}$, we have $\mathbf{P}_\sigma(R(w, s)) = \mathbf{P}_{\sigma'}(R'(w, s))$. Therefore, by choosing as set of secret state of S' , $\text{Sec}' = \text{Sec} \cap S'$ and as the set of secret state is absorbing, the disclosures over finite horizon are equal for σ and σ' . \blacktriangleleft

B Maximisation

► **Proposition 12** (Dominance of deterministic strategies). *Given an MDP, a secret and an observation-based strategy σ there exists a deterministic observation-based strategy σ' with greater or equal disclosure of the secret.*

Proof. In the proof of Lemma 1 of [10], the authors show that a randomised observation based strategy can be seen as an average over a family of deterministic observation based strategy. As a consequence of their equation (2), in our framework, given an observation based strategy σ and a disclosure notion disc , there exists an observation based deterministic strategy σ_{det} such that for all $\rho \in \text{FPath}$, $\text{Supp}(\sigma_{det}(\rho)) \subseteq \text{Supp}(\sigma(\rho))$ and $\mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_\sigma(\mu_0)) \geq \mathbf{P}_{M_\sigma(\mu_0)}(\mathcal{V}_\sigma(\mu_0))$.

The first property implies that $\mathcal{V}_\sigma(\mu_0) \cap \text{Path}(M_{\sigma_{det}}(\mu_0)) \subseteq \mathcal{V}_{\sigma_{det}}(\mu_0)$. Indeed, as σ allows more possibilities than σ_{det} , $\text{Path}(M_{\sigma_{det}}(\mu_0)) \subseteq \text{Path}(M_\sigma(\mu_0))$. This implies that given a path ρ , if $\mathcal{O}(\rho)$ discloses the secret with the strategy σ then either $\mathcal{O}(\rho)$ discloses the secret with the strategy σ_{det} or $\mathcal{O}(\rho)$ cannot be observed with σ_{det} . Hence the result.

This implies:

$$\mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_\sigma(\mu_0)) = \mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_\sigma(\mu_0) \cap \text{Path}(M_{\sigma_{det}}(\mu_0))) \leq \mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_{\sigma_{det}}(\mu_0))$$

Therefore,

$\text{disc}(M_{\sigma_{det}}(\mu_0)) = \mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_{\sigma_{det}}(\mu_0)) \geq \mathbf{P}_{M_{\sigma_{det}}(\mu_0)}(\mathcal{V}_\sigma(\mu_0)) \geq \mathbf{P}_{M_\sigma(\mu_0)}(\mathcal{V}_\sigma(\mu_0))$ and the result holds since $\mathbf{P}_{M_\sigma(\mu_0)}(\mathcal{V}_\sigma(\mu_0)) = \text{disc}(M_\sigma(\mu_0))$. \blacktriangleleft

► **Theorem 13** (Undecidability of maximal finite horizon disclosure). *The maximal finite horizon disclosure problem is undecidable for MDPs, even when the secret is reached with probability 1 and for a non-erasing observation function.*

The maximal finite horizon disclosure problem when restricted to finite-memory strategies is also undecidable (with the same additional assumptions).


Proof. Recall that we prove the result for (more compact) IMCs since it carries over to MDPs.

For an IMC $M = (S, s_{init}, I, \mathcal{O})$, we abuse notations by writing $\mu \in I(s)$ to denote any distribution $\mu : S \rightarrow [0, 1]$ such that for all $s' \in S$, $\mu(s') \in I(s' \mid s)$. The notion of path ρ is the same as for a Markov chain but a transition from $s = \text{last}(\rho)$ to some successor requires the choice of a distribution $\mu \in I(s)$. A strategy of IMC M is thus a mapping $\sigma : \text{FPath}(M) \rightarrow \text{Dist}(S)$ such that for each path ρ with $s = \text{last}(\rho)$, $\sigma(\rho) \in I(s)$.

We first give the construction. The preliminary step is very similar to what was done in the proof of Theorem 3. Starting from a PA $\mathcal{A} = (Q, q_0, \{a, b\}, T, F)$ that is supposed complete, we build an IMC $\hat{\mathcal{A}}$ where $\{a, b\}$ is the observation alphabet: The set of states is $\hat{Q} = Q \cup \{q_c \mid q \in Q \wedge c \in \{a, b\}\}$, with initial state q_0 . The observation function $\hat{\mathcal{O}}$ is defined by $\hat{\mathcal{O}}(q) = \varepsilon$ and $\hat{\mathcal{O}}(q_c) = c$ for $q \in Q$ and $c \in \{a, b\}$. The interval mapping $\hat{I} : \hat{Q} \rightarrow \mathbb{I}^{\hat{Q}}$ is defined for $q, q' \in Q$ and $c \in \{a, b\}$ by:

- $\hat{I}(q' | q_c) = T(q' | q, c)$ is a point interval;
- $\hat{I}(q_c | q) = [0, 1]$.

Compared to the illustration given in Figure 3, this construction amounts to replacing all b by ε (making the states non observable) and the probabilities $\frac{1}{4}$ from original states to new ones by the interval $[0, 1]$.



■ **Figure 5** Reduction from the strict emptiness problem to the disclosure problem. An edge outgoing from a dotted box should be duplicated to originate from all states in the box and an edge entering a dotted box from a state s should be duplicated from s to any state in the box. Hence a loop on a dotted box means a complete graph inside the box (including self-loops).

However, the construction of the complete IMC $M_{\mathcal{A}} = (S, s_0, I, O)$ from \mathcal{A} is more involved and requires to add an upper gadget limiting the power of the strategy. This is why we first use an observation function which can erase states and explain at the end how to lift this hypothesis. The construction is illustrated in Figure 5 with some conventions to avoid too many edges, a final state from \mathcal{A} (like q_f) is doubly circled.

- $S = \{s_0, s_1, q_{\#}^1, q_{\#}^2, q_b, q_s\} \cup \hat{Q} \cup \{s_c \mid c \in \{a, b, \#\}\} \cup \{r_c \mid c \in \{a, b, \#, b\}\}$;
- $I(s_1 | s_0) = I(q_0 | s_0) = \frac{1}{2}$ and the restriction of I to \hat{Q} is \hat{I} . For all $c \in \{1, a, b, \#\}$, $c' \in \{a, b, \#\}$, $I(s_{c'} | s_c) = \frac{1}{6}$ and $I(r_{c'} | r_c) = [0, \frac{1}{4}]$, for all $c, c' \in \{a, b, \#, b\}$, $I(r_{c'} | r_c) = \frac{1}{5}$ and $I(q_s | r_c) = \frac{1}{5}$. For all $q \in Q \setminus F$, $I(q_{\#}^1 | q) = [0, 1]$, for all $q \in F$, $I(q_{\#}^2 | q) = [0, 1]$, and $I(q_s | q_{\#}^1) = I(q_b | q_{\#}^2) = I(q_s | q_b) = I(q_s | q_s) = 1$.
- O extends \hat{O} by: $O(s_0) = O(s_1) = \varepsilon$, $O(q_{\#}^1) = O(q_{\#}^2) = O(q_s) = \#$, $O(q_b) = b$, for all $c \in \{a, b, \#, b\}$, $O(r_c) = c$, and for all $c \in \{a, b, \#\}$, $O(s_c) = c$.

Informally, for $Sec = \{q_s\}$, the upper gadget ensures that for any strategy σ there is at most one word $w \in \{a, b\}^*$ such that the observation $w\#\#$ discloses the secret. The lower

gadget allows to generate secret paths of observation $w\sharp b\sharp$ with half the probability as the one assigned by the PA to w .

We now formally prove that there exists a word $w \in \{a, b\}^*$ with $\mathbf{P}_{\mathcal{A}}^F(w) > \frac{1}{2}$ if and only if $\text{Disc}_{\max}(\mathbf{M}_{\mathcal{A}}) > \frac{1}{4}$.

First suppose there exists a word $w = a_1 \dots a_n \in \{a, b\}^*$ accepted with probability greater than $\frac{1}{2}$ in \mathcal{A} . We define the strategy σ for a finite path ρ in both parts of $\mathbf{M}_{\mathcal{A}}$ (when relevant) as follows:

- In the upper part, assume that ρ ends in a state s_c with $c \in \{1, a, b, \sharp\}$. If there exists $i < n$ such that $\mathbf{O}(\rho) = a_1 \dots a_i$ then $\sigma(\rho)(r_{a_{i+1}}) = 0$, leaving no choice for the rest of the distribution: In order for the sum of probabilities to be equal to 1 we have for $b \neq a_{i+1}$, $\sigma(\rho)(r_b) = \frac{1}{4}$. If $\mathbf{O}(\rho) = w$, then $\sigma(\rho)(r_{\sharp}) = 0$, which also leaves no choice for the rest of the distribution.
- In the bottom part, we can assume that ρ ends in a state $q \in Q$. If there exists $i < n$ such that $\mathbf{O}(\rho) = a_1 \dots a_i$ then $\sigma(\rho)(q_{a_{i+1}}) = 1$. Finally, if $\mathbf{O}(\rho) = w$ then $\sigma(\rho)(q_{\sharp}^2) = 1$ if $q \in F$, and $\sigma(\rho)(q_{\sharp}^1) = 1$ otherwise.

At the beginning the system will go with probability $\frac{1}{2}$ to $\hat{\mathcal{A}}$, where the strategy ensures that the word $w\sharp$ is observed. This leads to the state q_{\sharp}^1 with probability $\frac{1}{2}\mathbf{P}_{\mathcal{A}}^F(w)$ and thus the next observations belong to $b\sharp^*$ and the paths with observations in $w\sharp b\sharp^+$ belong to the secret. On the other hand, the system can also go to s_1 with probability $\frac{1}{2}$ from where $w\sharp$ cannot be observed because of the choices of the strategy. This implies that $w\sharp b\sharp$ is a min-disclosing observation in $\mathbf{M}_{\mathcal{A}, \sigma}$, hence $\text{Disc}(\mathbf{M}_{\mathcal{A}, \sigma}) \geq \frac{1}{2}\mathbf{P}_{\mathcal{A}}^F(w) > \frac{1}{4}$. Since $\text{Disc}_{\max}(\mathbf{M}_{\mathcal{A}}) = \sup_{\sigma} \text{Disc}(\mathbf{M}_{\mathcal{A}, \sigma})$, we can conclude that $\text{Disc}_{\max}(\mathbf{M}_{\mathcal{A}}) > \frac{1}{4}$.

Conversely suppose that the disclosure is strictly greater than $\frac{1}{4}$ and let σ be a strategy such that $\text{Disc}(\mathbf{M}_{\mathcal{A}, \sigma}) > \frac{1}{4}$. Then σ must forbid states in $\{r_c \mid c \in \{a, b, \sharp\}\}$, otherwise there would be no disclosing observation since every observation can be simulated once a state r_c is reached and followed by a b . Writing $\bar{\Sigma} = \{a, b, \sharp\}$, we inductively define the word $\bar{w} \in \bar{\Sigma}^{\infty}$ by a sequence $(\bar{w}_i)_{i \geq 0}$, such that $\bar{w}_i \leq \bar{w}_{i+1} \leq \bar{w}$ for all $i \geq 0$:

- We start with $\bar{w}_0 = \varepsilon$;
- Assume \bar{w}_i is built and let ρ_i be a path ending in state s_x for some $x \in \{1, a, b, \sharp\}$, with $\mathbf{O}(\rho_i) = \bar{w}_i$. If $\sigma(\rho_i)(r_c) = 0$ for some $c \in \{a, b, \sharp\}$, then $\bar{w}_{i+1} = \bar{w}_i c$, otherwise $\bar{w}_{i+1} = \bar{w}_i$. The set of ambiguous observations (i.e. corresponding to both secret and non-secret paths) are those reaching the set of states $\{r_c \mid c \in \{a, b, \sharp, b\}\}$:

$$\bigcup_{\substack{\bar{w}_i x \neq \bar{w}_{i+1} \\ x \neq b}} \bar{w}_i x (\bar{\Sigma} \cup \{b\})^* \sharp^{\omega}.$$


Hence, the set of disclosing observations is reduced to either $w\sharp b\sharp^{\omega}$, where w is the largest prefix of \bar{w} in $\{a, b\}^*$ if \sharp occurs in \bar{w} , and empty otherwise. Since the disclosure is greater than 0, we obtain $w\sharp b\sharp$ as the single min-disclosing observation with $\text{Disc}(\mathbf{M}_{\mathcal{A}, \sigma}) = \mathbf{P}_{\mathbf{M}_{\mathcal{A}, \sigma}}(w\sharp b\sharp)$. Since $\mathbf{P}_{\mathcal{A}}^F(w) \geq \mathbf{P}_{\mathbf{M}_{\mathcal{A}, \sigma}}(w\sharp b\sharp)$, we can conclude that $\mathbf{P}_{\mathcal{A}}^F(w) > \frac{1}{2}$.

The proof can be extended with a non-erasing observation function by replacing ε with a fixed additional symbol (like in the proof of Theorem 3). This requires to slightly modify the boxes with states $\{s_c \mid c \in \{a, b, \sharp\}\}$ and $\{r_c \mid c \in \{a, b, \sharp\}\}$ to ensure alternation of letters from $\{a, b\}$ and this new symbol.

The result on finite memory strategies holds as the strategy we build from the word of the automaton in the first direction of the proof only uses finite memory. ◀

► **Theorem 15** (Undecidability of maximal finite horizon limit-sure disclosure). *The maximal finite horizon limit-sure disclosure problem is undecidable for MDPs.*

Proof. The proof is a reduction from the value 1 problem for PA. The construction of $M_{\mathcal{A}}$ depicted in Figure 5 for the proof of Theorem 13 is slightly modified as follows (see Figure 6): a new state q_{\natural} with $O(q_{\natural}) = \natural$ is added in the upper part just before reaching the secret state q_s . In this case, the paths reaching the secret in the upper part disclose the secret as they end with \natural^{ω} . In the bottom part, the disclosure depends on $val(\mathcal{A})$ and, similarly as above, we can prove that $Disc_{\max}(M_{\mathcal{A}}) = 1$ if and only $val(\mathcal{A}) = 1$, which yields the result. ◀



■ **Figure 6** Modification of Figure 5 for limit disclosure.

► **Theorem 16** (Decidability of maximal finite-horizon almost-sure disclosure). *The maximal finite-horizon almost-sure disclosure problem in MDPs is EXPTIME-complete. Moreover, if the system is almost-surely disclosing, one can build a belief-based strategy with disclosure 1.*

Proof. We reduce the almost-sure disclosure problem for maximisation in MDPs to the almost-sure reachability of a partially observable MDP (POMDP). A POMDP is an MDP where the strategies resolving the non determinism depend only on the observation sequence and do not take the current state into account. The POMDP we build is exponential in the size of the original MDP and the algorithm to solve almost-sure reachability is exponential in the size of the POMDP [11]. This gives an EXPTIME algorithm as those two exponentials do not stack.

Construction of the POMDP. We start from MDP $M = (S, Act, p, O)$ with $Act = \cup_{s \in S} A(s)$, observation alphabet Σ , and a set of secret states $Sec \subseteq S$. Thanks to Proposition 11 we can suppose O to be non-erasing. We build a POMDP $M' = (S', Act', p', O')$ with set of states $S' \subseteq S \times 2^S$, and observation alphabet Σ . The observation function O' is defined by $O'((s, B)) = O(s)$. The set of actions of M' is $Act' = DR$ where DR is the set of vectors of deterministic decision rules $\vec{\delta}$ over S . The initial distribution μ'_0 satisfy $\mu'_0(s, B) = \mu_0(s)$ iff $B = O^{-1}(s) \cap Supp(\mu_0)$, and 0 otherwise. Given a state $(s, B) \in S'$, an action $\vec{\delta} \in DR$ and an observable state s' ,
 $p'((s', B') | (s, B), \vec{\delta}) = p(s' | s, \vec{\delta}(s))$ for $B' = \{s'' | O(s'') = O(s') \wedge \exists \hat{s} \in B, p(s'' | \hat{s}, \vec{\delta}(s)) > 0\}$, and 0 otherwise.

Correspondence between strategies. To show that M is almost-surely disclosing for Sec iff $Sec \times 2^{Sec}$ can almost-surely be reached in M' , we have to establish correspondences between the strategies of M and M' . From Proposition 12, we can restrict to deterministic observation-based strategies for M , and from [11], we also restrict to deterministic strategies for M' .

Let σ' be a strategy of M' , defined on Σ^* . Then σ is defined for any observation $w \in \Sigma^*$ and state $s \in S$ by $\sigma(w, s) = \sigma'(w)(s)$. Remark that a σ' -compatible path ρ' of M' ends in a state (s, B) where $B = B_w^\sigma$ (the belief w.r.t. σ) if $O(\rho') = w$. Conversely, given an

observation-based strategy σ of M , we build a strategy σ' of M' as follows: Given $w \in \Sigma^*$, we define the mapping $\sigma'(w) : S \rightarrow \text{Act}$ by $\sigma'(w)(s) = \sigma(w, s)$ for any $s \in S$.

Then, writing \mathbf{P}_σ (resp. $\mathbf{P}_{\sigma'}$) instead of $\mathbf{P}_{M_\sigma(\mu_0)}$ (resp. $\mathbf{P}_{M_{\sigma'}(\mu'_0)}$), and defining $R(w, s) = \{\rho \in \text{FPath}(M_\sigma(\mu_0)) \mid \text{O}(\rho) = w \wedge \text{last}(\rho) = s\}$ and $R'(w, s) = \{\rho' \in \text{FPath}(M_{\sigma'}(\mu'_0)) \mid \text{O}(\rho') = w \wedge \text{last}(\rho') = (s, B_w^\sigma)\}$, we have $\mathbf{P}_\sigma(R(w, s)) = \mathbf{P}_{\sigma'}(R'(w, s))$.

Now let $\text{Reach}(Sec \times 2^{Sec})$ be the set of paths reaching $Sec \times 2^{Sec}$ in $M_{\sigma'}(\mu'_0)$. Then we claim that $\text{Disc}_{\max}(M_\sigma(\mu_0)) = \mathbf{P}_{\sigma'}(\text{Reach}(Sec \times 2^{Sec}))$. Indeed, an observation w discloses the secret under strategy σ iff all observable states reachable with observation w belong to the secret, i.e. iff $B_w^\sigma \subseteq Sec$. Thus the paths ρ' in M' with a disclosing observation are those for which $\text{last}(\rho') \in Sec \times 2^{Sec}$. Therefore the probability of reaching $Sec \times 2^{Sec}$ in M' under strategy σ' is also the probability of disclosing Sec in M under strategy σ .

We can conclude that M is almost-surely disclosing if and only if the paths of M' reach almost-surely the set of states $Sec \times 2^{Sec}$. Moreover, if M' almost-surely reaches the set of states $Sec \times 2^{Sec}$, we can build a strategy σ' doing so. Using the transformation described above, we extract from σ' a belief-based strategy σ of M that almost-surely discloses the secret.

Hardness is shown with a reduction from *safety games with imperfect information*.

► **Definition 24.** A safety game with imperfect information is a tuple $\mathcal{G} = (L, \ell_0, \Sigma, \Delta, O, F, obs)$ where

- L is a finite set of locations with initial location $\ell_0 \in L$;
- Σ is a finite alphabet;
- $\Delta \subseteq L \times \Sigma \times L$ is the transition relation such that for all $\ell \in L$ and $a \in \Sigma$ there exists at least one ℓ' with $(\ell, a, \ell') \in \Delta$;
- O is a finite set of observations, and $F \subseteq O$ are the final observations;
- $obs : L \rightarrow O$ is the observation mapping.

A safety game with imperfect information \mathcal{G} is a turn-based game played by two players Control and Environment. It starts in location ℓ_0 with Control to play. In the first round, Control chooses a letter $a_0 \in \Sigma$, then Environment chooses a location ℓ_1 such that $(\ell_0, a_0, \ell_1) \in \Delta$ and Control only observes $o_1 = obs(\ell_1)$. The next rounds are played similarly and Control wins if for all i , $o_i \notin F$.

The problem of existence of a winning strategy for Control is EXPTIME-complete [8]. We now describe a reduction from this problem to the almost-sure disclosure problem of MDPs.

The reduction is similar to the proof of Theorem 13 except that we replace the probabilistic automaton by a safety game $\mathcal{G} = (L, \ell_0, \Sigma, \Delta, O, F, obs)$ with imperfect information and directly build an MDP $M = (S, \text{Act}, p, \text{O})$ over alphabet $(O \cup \{\#\, b, \dagger\}) \cup \Sigma \times (O \cup \{\#\, b\})$, with:

- $S = \{s_0, \ell_0, q_\#^1, q_\#^2, q_b, q_s^1, q_s^2\} \cup \{\ell_c \mid \ell \in L, c \in \Sigma \times O\} \cup \{s_c \mid c \in \Sigma \times (O \cup \{\#\, b\})\} \cup \{r_c \mid c \in \Sigma \times (O \cup \{\#\, b\})\}$;
- $\text{Act} = \Sigma$;
- For all $a \in \Sigma, o \in O, \ell_c \in S, \ell' \in obs^{-1}(o)$, $p(\ell'_{a,o} \mid \ell_c, a) > 0$ iff $(\ell, a, \ell') \in \Delta$.
If $\ell \in obs^{-1}(F)$ then $p(q_\#^1 \mid \ell_c, a) > 0$ and if $\ell \notin obs^{-1}(F)$ then $p(q_\#^2 \mid \ell_c, a) > 0$.
For all $a \in \Sigma, c \in \{0\} \cup (\Sigma \times (O \cup \{\#\, b\}))$, $(b', o') \in \Sigma \times (O \cup \{\#\, b\})$, $p(s_{(b', o')} \mid s_c, a) > 0$ and if $b' \neq a$, $p(r_{(b', o')} \mid s_c, a) > 0$.
For all $c, c' \in \Sigma \times (O \cup \{\#\, b\})$, $p(r_{c'} \mid r_c, a) > 0$ and $p(q_s^2 \mid r_c, a) > 0$. For all $a, a' \in \Sigma$, $p(q_s^1 \mid q_\#^1, a) = p(q_b \mid q_\#^2, a) = p(q_s^1 \mid q_b, a) = p(q_s^1 \mid q_s^1, a) = p(q_s^2 \mid q_s^2, a) = 1$.
- $\text{O}(s_0) = \text{O}(\ell_0) = obs(\ell_0)$; For $z \in L, s, r, a \in \Sigma, o \in O$, $\text{O}(z_{a,o}) = (a, o)$; For $o \in \{\#\, b\}$, $\text{O}(z_{a,o}) = o$, and $\text{O}(q_\#^1) = \text{O}(q_\#^2) = \# = \text{O}(q_s^1)$, $\text{O}(q_b) = b$, $\text{O}(q_s^2) = \dagger$.

The initial distribution is $\mu_0(s_0) = 1/2 = \mu_0(\ell_0)$ and the secret is $Sec = \{q_s^1, q_s^2\}$.

This proof being similar to the one of Theorem 13, we only detail here the differences. A path starting in s_0 will almost surely trigger a \dagger and disclose the secret. A path starting in ℓ_0 will almost surely reach q_s^1 as after any action in the copy of \mathcal{G} there is a positive probability to reach $q_\#^1$ or $q_\#^2$. In order for a finite path starting in ℓ_0 to disclose the secret, it can not go through $q_\#^1$ and should not have the same observation as the one of a path ending in a state r_c . Given a strategy σ of M , if there exists a σ -compatible path ρ visiting a state ℓ_c with an observation $O(\ell_c) \in \Sigma \times F$, then there is a σ -compatible path ρ' visiting $q_\#^1$, therefore a set of paths with positive probability do not visit the secret. Thus a deterministic strategy almost surely disclosing the secret in M never visits a state triggering an observation of the form $\Sigma \times F$. Moreover such a strategy does not take the current state into account. Indeed, let ρ and ρ' be two paths such that $O(\rho) = O(\rho')$ and ending in two states ℓ_c and s_c . If $\sigma(\rho) = a$ and $\sigma(\rho') = a'$ are two actions in Σ with $a \neq a'$ then there exists $o \in O$ such that $\rho a \ell_{a,o}$ is a σ -compatible path. Since $a \neq a'$, $\rho' a r_{a,o}$ is also a σ -compatible path with same observation than $\rho a \ell_{a,o}$. Hence no observation prefixed by $O(\rho a \ell_{a,o})$ would disclose the secret.

Therefore, similarly as in Theorem 13, Control has a winning strategy iff there exists a deterministic strategy considering only the sequence of observation that almost-surely discloses the secret. This implies EXPTIME-hardness. \blacktriangleleft

C Minimisation

► **Proposition 19.** *The minimal disclosure value for Sec in $M(\mu_0)$ is equal to the minimal probability to reach $Avoid(Sec)$ in $M_{\min}(\mu_{\min})$. Furthermore it is asymptotically reached by a family of belief-based almost deterministic strategies.*

Proof. We know that the minimal reachability probability for $Avoid(Sec)$ in $M_{\min}(\mu_{\min})$ is obtained by a memoryless deterministic strategy σ_{\min} that selects some $a_{s,B}$ in state (s, B) . Consider $\{\sigma_\varepsilon\}$ the family of belief-based almost-deterministic strategies defined by favouring $a_{s,B}$ in state s after a path ρ such that $B_\rho^{\sigma_\varepsilon} = B$. Given a path $\rho = s_0 a_0 \dots a_{n-1} s_n$ in $M(\mu_0)$ we inductively define the path $b(\rho) = (s_0, S_0) a_0 \dots a_{n-1} (s_n, S_n)$ in $M_{\min}(\mu_{\min})$ by: $S_0 = Supp(\mu_0) \cap O^{-1}(O(s_0))$ and $S_{i+1} = NextMax(S_i, O(s_{i+1}))$. Due to the observation given when introducing $NextMax$, with strategy σ_ε , the observation of path ρ discloses the secret iff $b(\rho)$ reaches $Avoid(Sec)$. Consider under strategy σ_ε the probability to disclose the secret with paths ρ such that $b(\rho)$ includes at least once an action not selected by σ_{\min} . By construction, at each step i , the probability of not choosing the action favoured by σ_ε is $\frac{\varepsilon}{2^i}$, hence the probability of those paths is $\sum_{i \geq 0} (1 - \varepsilon)^i \frac{\varepsilon}{2^i} \leq 2\varepsilon$. Consider now a path $s_0 a_0 \dots a_{n-1} s_n$ such that $b(\rho)$ is σ_{\min} -compatible. Then the probability of the original path is less than or equal to the probability of its corresponding path. So we deduce that the minimal disclosure value of $M(\mu_0)$ is bounded above by $\nu + 2\varepsilon$ where ν is the minimal reachability probability for $Avoid(Sec)$ in $M_{\min}(\mu_{\min})$. Since this holds for all $\varepsilon > 0$, we obtain that the minimal disclosure value of $M(\mu_0)$ is bounded above by the minimal reachability probability for $Avoid(Sec)$ in $M_{\min}(\mu_{\min})$.

Conversely consider an arbitrary strategy σ in $M(\mu_0)$. This strategy may be also applied in $M_{\min}(\mu_{\min})$ by forgetting the second component of the state, defining a strategy σ' . For any path $s_0 \dots s_n$ in $M_\sigma(\mu_0)$, there is a single path $(s_0, S_0) a_0 \dots (s_n, S_n)$ in $M_{\min}(\mu_{\min})$ under σ' with the same probability. Given the path $s_0 a_0 \dots s_n$, consider the successive associated subsets of beliefs according to σ , B_0, \dots, B_n . By induction (and definition of M_{\min}) it is straightforward to show that $B_i \subseteq S_i$. So $s_0 a_0 \dots s_n$ does not disclose the secret in M under σ

implies that $(s_0, S_0) \dots (s_n, S_n)$ does not reach $\text{Avoid}(\text{Sec})$. This entails that the reachability probability of $\text{Avoid}(\text{Sec})$ in $M_{\min}(\mu_{\min})$ under σ' is less than or equal to the disclosure probability in $M(\mu_0)$ under σ . ◀

D Fixed horizon

► **Theorem 22** (Computation of the maximal disclosure value within fixed-horizon). *The fixed-horizon maximal value (when the horizon n is described in unary representation) is computable in PSPACE and the fixed-horizon maximal disclosure problem is PSPACE-complete.*

Value and membership. We first present a non deterministic procedure that decides in PSPACE the disclosure problem. It can then be determined using Savitch's Theorem.

From an arbitrarily ordered observation alphabet Σ , the procedure operates as follows for horizon n :

- It maintains a disclosure value v , a sequence of observations $o_0 \dots o_i$ with $i \leq n$, a sequence of sets of states $B_1 \dots B_i$ with $B_j \subseteq \mathcal{O}^{-1}(o_j)$ for all $j \leq i$, an action $a_{j,s} \in A(s)$ for all (j, s) with $j < i$ and $s \in S_j$, and for all (j, s) with $j \leq i$ and $s \in B_j$ the probability $p_{j,s}$ to reach s after the sequence of observations $o_0 \dots o_i$;
- Initially $v = 0$, o_0 is the smallest observation in $\mathcal{O}(\text{Supp}(\mu_0))$, where μ_0 is the initial distribution, $B_0 = \text{Supp}(\mu_0) \cap \mathcal{O}^{-1}(o_0)$ and $p_{0,s} = \mu_0(s)$ for $s \in B_0$;
- If $i < n$ then for all $s \in B_i$, the procedure *guesses* an action $a_{i,s} \in A(s)$. Let o_{i+1} be the smallest observation such that there exists a state $s \in B_i$ and a state $s' \in \mathcal{O}^{-1}(o_{i+1})$ with $p(s'|s, a_{i,s}) > 0$. Then B_{i+1} is set to $\{s' \in \mathcal{O}^{-1}(o_{i+1}) \mid \exists s \in B_i p(s'|s, a_{i,s}) > 0\}$ and for all $s' \in B_{i+1}$, $p_{i+1,s'} = \sum_{s \in B_i} p_{i,s} p(s'|s, a_{i,s})$;
- If $i = n$, the procedure examines B_n . If $B_n \subseteq \text{Sec}$ then $v = v + \sum_{s \in B_n} p_{n,s}$ otherwise v is unchanged. Afterwards it “backtracks” to the greatest $0 < i \leq n$ such that there exists $o'_i > o_i$ with some $s \in B_{i-1}$ and a state $s' \in \mathcal{O}^{-1}(o'_i)$ with $p(s'|s, a_{i-1,s}) > 0$. Then B_i and the $p_{i,s'}$'s are updated accordingly and the procedure carries on. If there is no such i , the procedure returns to $i = 0$ and similarly looks for some $o'_0 > o_0$, where the initialization step is again performed except for the value of v which is unchanged. When the maximal observation in $\Sigma \cap \mathcal{O}(\text{Supp}(\mu_0))$ is handled, the procedure terminates by comparing v to the threshold.


The correctness of the procedure follows from the fact that there exists an optimal deterministic strategy where the selection of the action for the current state depends only on the sequence of observations (and not on the sequence of visited states).

The space complexity of the procedure is in $O(n|S|(\log(|A|) + nb))$ where b is the maximal number of bits used to represent a transition probability of the MDP.


Observe now that since the maximal value is obtained by a deterministic strategy, one knows a denominator of this value: d^n where d is the lcm of the denominators for probabilities occurring in the model. Its bit size is polynomial w.r.t. the size of the model. So by iteratively solving the disclosure problem for $\frac{i}{d^n}$ for increasing values of i , one computes the maximal value in PSPACE. ◀

Hardness. We reduce the truth of a quantified boolean formula: Given a closed QBF in 3CNF $\phi = \exists x_1 \forall y_1 \exists x_2 \dots \forall y_n \psi$ with $\psi = \bigwedge_{i=1 \dots m} (z_{i_1} \vee z_{i_2} \vee z_{i_3})$, we build an MDP M such that ϕ is true if and only if $\text{Disc}_{2(n+m)+3, \max}(M) \geq \frac{1}{2^{2^n}}$.

The MDP $M = (S, \text{Act}, p, \mathcal{O})$ (depicted in Figure 7 with some conventions to avoid having too many edges) is defined by:



■ **Figure 7** Reduction of the satisfiability problem to the disclosure on a fixed horizon. The box S_{x_1} is represented in Figure 8.



■ **Figure 8** Representation of the box S_{x_1} . We use the convention of Figure 5.

- $S = \{s_{init}, s_{end}, s_{\#}\} \cup \{s_{z_i} \mid z \in \{x, y\}, i = 1 \dots n\} \cup \{s_{\neg z_i} \mid z \in \{x, y\}, i = 1 \dots n\} \cup \{s_t^i \mid i \in \{1, \dots, m\}\} \cup \{s^{z_i j} \mid i \in \{1, \dots, m\}, j \in \{1, 2, 3\}\} \cup_{i \in \{1, \dots, n\}} (S_{x_i} \cup S_{y_i} \cup S_{\neg x_i} \cup S_{\neg y_i})$ where for z_i one of the boolean variable, $S_{z_i} = \{s_a^{z_i, 1} \mid a \in \{\#, z_i, \neg z_i \mid z \in \{x, y\}, i \in \{1, \dots, n\}\}\} \cup \{s_a^{z_i, 2} \mid a \in \{\#, b, z_i, \neg z_i \mid z \in \{x, y\}, i \in \{1, \dots, n\}\}\}$. Similar for the box $S_{\neg z_i}$ of the negation of a variable.
- $A(s_{init}) = \{x_1, \neg x_1\}$, and for all $i < n$, $A(s_{y_i}) = A(s_{\neg y_i}) = \{x_{i+1}, \neg x_{i+1}\}$. For $i \in \{1, \dots, m\}$, $A(s_t^i) = \{1, 2, 3\}$ and $A(s) = \{next\}$ for all other states (even in boxes).
- $p(s_a \mid s_{init}, a) = p(s_a^{a, 1} \mid s_{init}, a) = 1/2$. For all $i < n$, $p(s_a \mid s_{y_i}, a) = p(s_a^{a, 1} \mid s_{y_i}, a) = p(s_a \mid s_{\neg y_i}, a) = p(s_a^{a, 1} \mid s_{\neg y_i}, a) = 1/2$. For all $i \leq n$, $p(s_{y_i} \mid s_{x_i}, next) = p(s_{y_i}^{y_i, 1} \mid s_{x_i}, next) = p(s_{\neg y_i} \mid s_{x_i}, next) = p(s_{\neg y_i}^{y_i, 1} \mid s_{x_i}, next) = p(s_{y_i} \mid \neg s_{x_i}, next) = p(s_{y_i}^{y_i, 1} \mid s_{\neg x_i}, next) = p(s_{\neg y_i} \mid s_{\neg x_i}, next) = p(s_{\neg y_i}^{y_i, 1} \mid s_{\neg x_i}, next) = 1/4$, and $p(s_t^1 \mid s_{\#}, next) = 1$. For all $i = 1 \dots m, j \in \{1, 2, 3\}, p(s^{z_i j} \mid s_t^i, j) = 1$, and if

$i < m$, $p(s_t^{i+1} | s^{z_{i_j}}, next) = 1$. Finally, $p(s_{end} | s^{z_{m_j}}, next) = p(s_{end} | s_{end}, next) = 1$.

We now describe p for the box S_{x_1} other boxes being similar. For all $a, b \in \{\sharp, x_1, z_i, \neg z_i | z \in \{x, y\}, i \in \{2, \dots, n\}\}$, $p(s_b^{x_1,1} | s_a^{x_1,1}, next) = p(s_{\neg x_1}^{x_1,2} | s_a^{x_1,1}, next) = 1/(4n + 1)$ and for all $c, d \in \{\sharp, b, z_i, \neg z_i | z \in \{x, y\}, i \in \{1, \dots, n\}\}$, $p(s_d^{x_1,2} | s_c^{x_1,2}, next) = 1/(4n + 2)$.

- $O(s_{end}) = b$, $O(s_a^b) = a$ and $O(s^a) = a$ when a is a boolean variable or its negation and for all other state s , $O(s) = \sharp$.

The initial distribution μ_0 is $\mathbf{1}_{s_{init}}$ and the set of secret paths SPath contains those reaching s_{end} .

We show that ϕ is true iff the disclosure of \mathbf{M} for observations of length $2(n + m) + 3$ is greater than or equal to $\frac{1}{2^{2n}}$. First remark that for any strategy, the measure of paths reaching state s_{end} with observation of length $2(n + m) + 3$ is exactly $\frac{1}{2^{2n}}$. Indeed, during each of the first $2n$ actions, whatever the choices of the strategy, there is a probability $\frac{1}{2}$ to go in one of the boxes and $\frac{1}{2}$ to continue advancing, thus a probability $\frac{1}{2^{2n}}$ to reach the state s_{\sharp} . From there every path reaches s_{end} in $2(m + 1)$ steps. If the strategy is such that some variable and its negation are read on the way to s_{end} , then there exists a path with same observation reaching the second part of a box where every observation can be triggered, and thus the path reaching s_{end} will not disclose the secret.

Intuitively, during the first $2n$ steps, every boolean variable will be assigned a value: either chosen by the strategy as it chooses whether x_i or $\neg x_i$ occurs in the observation for all $1 \leq i \leq n$, or randomly as y_i and $\neg y_i$ both have half chance of being triggered. During the last $2m$ steps, the strategy must trigger a boolean formula in every clause of the disjunction so that if a clause is not satisfied by the current assignment, then a boolean variable will be observed as both true and false during the path. Thus the observation would not disclose the secret. In order for a measure of $\frac{1}{2^{2n}}$ of paths to disclose the secret, for every assignment of the y_i the controller must force the path reaching s_{end} to disclose the secret.

Suppose that ϕ is equivalent to true. Thus there exist functions $(f_i)_{i=1\dots n}$ (expressing the choices for x_1, \dots, x_n) such that for every set of assignments (a_1, \dots, a_n) of the variables y_1, \dots, y_n the boolean formula $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is true. We choose a strategy σ such that for every possible set of assignments (a_1, \dots, a_n) for the variables y_1, \dots, y_n , for all $i \in \{0, \dots, n - 1\}$, $\sigma(\sharp f_1() a_1 f_2(a_1) \dots a_i) = f_{i+1}(a_1, \dots, a_i)$. Moreover for $i_1, \dots, i_k \in \{1, 2, 3\}$, there exists $z_{k+1}^{i_{k+1}} \in \{f_1(), a_1, f_2(a_1), \dots, a_n\}$ such that $\sigma(\sharp f_1() a_1 f_2(a_1) \dots a_n \sharp z_{1_{i_1}} z_{2_{i_2}} \dots z_{k_{i_k}}) = z_{k+1}^{i_{k+1}}$. The choice of the strategy is arbitrary in the other cases. Such a strategy can be defined since $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is true and thus every clause is satisfied by this choice of assignments.

With this strategy, the fixed horizon disclosure in $2(n + m + 1)$ steps is $\frac{1}{2^{2n}}$. In other words, all the paths reaching the secret disclose it. Indeed let ρ be a secret path of length $2(n + m + 1)$. There exists an assignment $a_1, \dots, a_n \in \{y_1, \neg y_1, \dots, y_n, \neg y_n\}$ such that $O(\rho) = \sharp f_1() a_1 f_2(a_1) \dots a_n \sharp z_{1_{i_1}} z_{2_{i_2}} \dots z_{m_{i_m}} b$. By choice of σ , if, for $z \in \{x, y\}$ and $i \in \{1, \dots, n\}$, z_i appears in the observation of ρ , $\neg z_i$ does not appear, and vice versa. Therefore as b can be read either in s_{end} or in a state reachable only by paths observing a boolean variable and its negation, ρ discloses the secret.

Conversely, suppose that ϕ is not equivalent to true and let σ be a strategy, which can be assumed to be deterministic thanks to Proposition 12. We build partial functions $f_i : \Sigma^{2i} \mapsto \text{Act}$ consistent with σ : for every observation $\sharp w \in \sharp \Sigma^{2i}$ of some path ρ , if σ chooses action $a \in A(\text{last}(\rho))$ for ρ (i.e., $\sigma(\rho)(a) = 1$) then $f_i(w) = a$. As ϕ is not equivalent to true, there exists an assignment (a_1, \dots, a_n) for the variables y_1, \dots, y_n such that the boolean formula $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is false. We now build a path with non null probability, reaching the secret but not disclosing it.

By construction, there exists ρ such that $\mathbf{O}(\rho) = \#f_1()a_1f_2(a_1)\dots f_n(a_1\dots a_{n-1})a_n\#$, with $\text{last}(\rho) = s_\#$ and $\mathbf{P}_\sigma(\rho) > 0$ (where again \mathbf{P}_σ stands for $\mathbf{P}_{\mathbf{M}_\sigma(\mu_0)}$). Let $i \in \mathbb{N}$ be an integer such that $z_{i_1} \vee z_{i_2} \vee z_{i_3}$ is not true under the assignment $[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$, in other words such that the negations of z_{i_1}, z_{i_2} and z_{i_3} were chosen as assignment. Let ρ' be the path of length $2(n+m)+2$ extending ρ and ending in s_{end} . Then ρ' does not disclose the secret: indeed, there exists $j \in \{1, 2, 3\}$ such that z_{i_j} appears in its last $2m$ observations while its negation (written $\neg z_{i_j}$ here) appears in the first $2n+1$ observations. Thus there exists a path with same observation leading to the second part of the box $S_{\neg z_{i_j}}$ which is outside the secret and where every observation is possible. As the total measure of paths reaching the secret is $\frac{1}{2^{2n}}$ and at least a subset of measure $\mathbf{P}_\sigma(\rho)$ of the paths reaching the secret do not disclose it, the disclosure of \mathbf{M} is strictly smaller than $\frac{1}{2^{2n}}$ in $2(n+m)+3$ observation steps. \blacktriangleleft

To prove Theorem 23, we need the following lemma where disc_h^* stands for $(\text{Disc}_h)_{\min}$:

► **Lemma 25.** *Let h be a fixed horizon. Let μ be a distribution with support $B \subseteq \mathbf{O}^{-1}(o)$ for some observation o . Then $\text{disc}_h^*(\mathbf{M}(\mu)) = \sum_{s \in \text{Supp}(\mu)} \mu(s) \text{disc}_h^*(\mathbf{M}(s, \text{Supp}(\mu)))$ and there exists a family of almost deterministic strategies that asymptotically reaches this value.*

Proof. The proof is done by induction on the horizon length h with some arbitrary initial distribution μ with support $B \subseteq \mathbf{O}^{-1}(o)$ for some o :

$$\text{disc}_{h+1}^*(\mathbf{M}(\mu))$$

$$\begin{aligned} &= \inf_{\vec{\delta}} \text{disc}_h^*(\mathbf{M}(\text{NextDist}(\mu, \vec{\delta}))) \\ &= \inf_{\vec{\delta}} \sum_{o \in \Sigma} \text{NextDist}(\mu, \vec{\delta})(o) \text{disc}_h^*(\mathbf{M}(\text{NextDist}(\mu, \vec{\delta}), o)) \\ &= \inf_{\vec{\delta}} \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} \text{NextDist}(\mu, \vec{\delta})(s) \text{disc}_h^*(\mathbf{M}(s, \text{Supp}(\text{NextDist}(\mu, \vec{\delta}), o))) \end{aligned}$$

(applying the induction hypothesis)

$$\geq \inf_{\vec{\delta}} \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} \text{NextDist}(\mu, \vec{\delta})(s) \text{disc}_h^*(\mathbf{M}(s, \text{NextMax}(B, o)))$$

(since enlarging the belief decreases disclosure)

$$\begin{aligned} &= \inf_{\vec{\delta}} \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} \sum_{\hat{s} \in B} \mu(\hat{s}) p(s|\hat{s}, \vec{\delta}[\hat{s}]) \text{disc}_h^*(\mathbf{M}(s, \text{NextMax}(B, o))) \\ &= \inf_{\vec{\delta}} \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} \sum_{\hat{s} \in B} \sum_{a \in A(\hat{s})} \mu(\hat{s}) \vec{\delta}[\hat{s}](a) p(s|\hat{s}, a) \text{disc}_h^*(\mathbf{M}(s, \text{NextMax}(B, o))) \end{aligned}$$

(applying the definition of NextDist and summing over the possible actions).

Observe now that we have a linear expression over the unknowns $\{\vec{\delta}[\hat{s}](a)\}$. So the last infimum is obtained for some $\vec{\delta}$ such that there exists a family of $\{a_{\hat{s}}\}$ with $\vec{\delta}[\hat{s}] = \mathbf{1}_{a_{\hat{s}}}$ for each $\hat{s} \in B$. We denote this value by $I(\mathbf{M}, \mu, h)(\vec{\delta})$ and consider $\vec{\delta}_\varepsilon$ defined by $\vec{\delta}_\varepsilon[s] = \vec{\delta}[s]_\varepsilon$. Then:

$$I(\mathbf{M}, \mu, h)(\vec{\delta}) \leq \text{disc}_{h+1}^*(\mathbf{M}(\mu)) \leq I(\mathbf{M}, \mu, h)(\vec{\delta}_\varepsilon) \leq I(\mathbf{M}, \mu, h)(\vec{\delta}) + \varepsilon$$

Thus we establish the induction by selecting initially the family of decision rules $\{\vec{\delta}_\varepsilon\}_\varepsilon$ which asymptotically leads to the optimal choice. In addition the optimal value can be rewritten as:

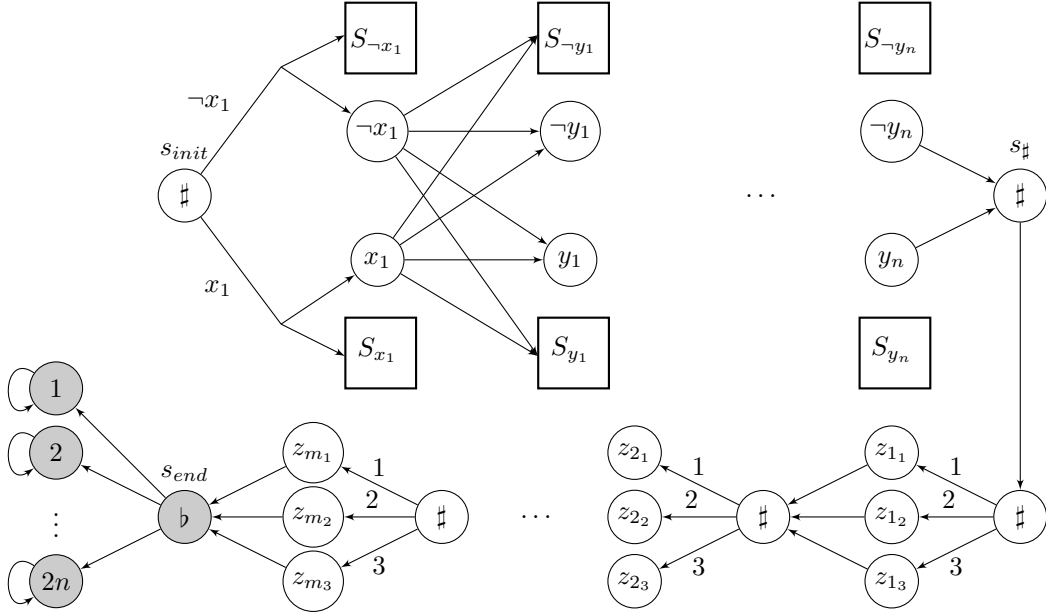
$$\begin{aligned} I(\mathbf{M}, \mu, h)(\vec{\delta}) &= \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} \sum_{\hat{s} \in B} \mu(\hat{s}) p(s|\hat{s}, a_{\hat{s}}) \text{disc}_h^*(\mathbf{M}(s, \text{NextMax}(B, o))) \\ &= \sum_{\hat{s} \in B} \mu(\hat{s}) \sum_{o \in \Sigma} \sum_{s \in \mathbf{O}^{-1}(o)} p(s|\hat{s}, a_{\hat{s}}) \text{disc}_h^*(\mathbf{M}(s, \text{NextMax}(B, o))) \end{aligned}$$

where $\sum_{o \in \Sigma} \sum_{s \in O^{-1}(o)} p(s|\hat{s}, a_s) disc_h^*(M(s, \text{NextMax}(B, o)))$ can be shown by a similar reasoning to be equal to $disc_{h+1}^*(M(\hat{s}, B))$. ◀

► **Theorem 23** (Minimal disclosure within fixed horizon). *The fixed horizon minimal value is computable in PSPACE. The fixed horizon minimal disclosure problem is PSPACE-complete. In addition, the strategy decision problem is also decidable in PSPACE.*


Proof. The procedures for the first two problems are very similar to the ones used in Theorem 22. There are only two differences. First given B_i the current belief and o_{i+1} one computes $B_{i+1} = \text{NextMax}(B_i, o_{i+1})$ (independently of the guessed actions $a_{i,s}$). Second the computation procedure operates by decreasing values of i when the value is less or equal than $\frac{i}{d^n}$.

In order to decide whether a strategy exists that provides the minimal value, one guesses this strategy in PSPACE as before. However there is an additional difficulty since the (possible) optimal strategy may be randomised. Thus during the procedure, given some belief B and some state s , one guesses the support $A' \subseteq A_s$ of the decision rule and one defines the decision rule say δ as a uniform choice over A' . We claim that this restriction is sound. Assume another decision rule δ' with same support would provide a smaller value then since the support are unchanged the decision rule uniformly described as $(1 + \varepsilon)\delta' - \varepsilon\delta$ for small enough ε would still provide a better value meaning that the support A' cannot be used to find an optimal strategy.



■ **Figure 9** Reduction of the satisfiability problem to the disclosure on a fixed horizon. The box S_{x_1} is represented in Figure 10

Like for the case of maximisation, the hardness of the fixed-horizon minimal disclosure problem is obtained by a reduction from the truth of a quantified boolean formula. Let $\phi = \exists x_1 \forall y_1 \exists x_2 \dots \forall y_n \psi$ with $\psi = \bigwedge_{i=1 \dots m} (z_{i1} \vee z_{i2} \vee z_{i3})$ a closed QBF where assume w.l.o.g. that the literals of any clause are distinct. We build the MDP $M = (S, \text{Act}, p, O)$ (represented in Figure 9) where:



■ **Figure 10** Representation of the box S_{x_1} (with the conventions of Figure 5).

- $S = \{s_{init}, s_{end}, s_{\#}\} \cup \{s_{z_i} \mid z \in \{x, y\}, i = 1 \dots n\} \cup \{s_{\neg z_i} \mid z \in \{x, y\}, i = 1 \dots n\} \cup \{s_t^i \mid i \in \{1, \dots, m\}\} \cup \{s^{z_{ij}} \mid i \in \{1, \dots, m\}, j \in \{1, 2, 3\}\} \cup \{s_{end}^{z_i} \mid z \in \{x, y\}, i = 1 \dots n\} \cup_{i \in \{1, \dots, n\}} (S_{x_i} \cup S_{y_i} \cup S_{\neg x_i} \cup S_{\neg y_i})$ where for z_i one of the boolean variable, $S_{z_i} = \{s_a^{z_i} \mid a \in \{\#, z_i, \neg z_i \mid z \in \{x, y\}, i \in \{1, \dots, n\}\}\} \cup \{s_b^{z_i}, s_f^{z_i}\}$. Similar for the box $S_{\neg z_i}$ of the negation of a variable.
- $A(s_{init}) = \{x_1, \neg x_1\}$, and for all $i < n$, $A(s_{y_i}) = A(s_{\neg y_i}) = \{x_{i+1}, \neg x_{i+1}\}$. For $i \in \{1, \dots, m\}$, $A(s_t^i) = \{1, 2, 3\}$ and for every other state (even in boxes), $A(s) = \{next\}$.
- $p(s_a \mid s_{init}, a) = p(s_a^{a,1} \mid s_{init}, a) = 1/2$. For all $i < n$, $p(s_a \mid s_{y_i}, a) = p(s_a^{a,1} \mid s_{y_i}, a) = p(s_a \mid s_{\neg y_i}, a) = p(s_a^{a,1} \mid s_{\neg y_i}, a) = 1/2$. For all $i \leq n$,
 $p(s_{y_i} \mid s_{x_i}, next) = p(s_{y_i}^{y_i,1} \mid s_{x_i}, next) = p(s_{\neg y_i} \mid s_{x_i}, next) = p(s_{\neg y_i}^{\neg y_i,1} \mid s_{x_i}, next) =$
 $p(s_{s_{y_i} \mid \neg x_i}, y_i) = p(s_{y_i}^{y_i,1} \mid s_{\neg x_i}, y_i) = p(s_{\neg y_i} \mid s_{\neg x_i}, next) = p(s_{\neg y_i}^{\neg y_i,1} \mid s_{\neg x_i}, next) = 1/4$.
 $p(s_t^1 \mid s_{\#}, next) = 1$. For all $i = 1 \dots m$, $j \in \{1, 2, 3\}$, $p(s^{z_{ij}} \mid s_t^i, j) = 1$, and if $i < m$,
 $p(s_t^{i+1} \mid s^{z_{ij}}, next) = 1$. Finally $p(s_{end} \mid s^{z_{mj}}, next) = 1$, and for all $z \in \{x, y\}$, $i = 1 \dots n$,
 $p(s_{end}^{z_i} \mid s_{end}, next) = 1/(2n)$.
 We now describe p for the box S_{x_1} other boxes being similar. For all $a \in \{\#, x_1, z_i, \neg z_i \mid z \in \{x, y\}, i \in \{2, \dots, n\}\}$, $b \in \{\#, b, z_i, \neg z_i \mid z \in \{x, y\}, i \in \{1, \dots, n\}\}$ $p(s_b^{x_1} \mid s_a^{x_1}, next) = 1/(4n+2)$, and $p(s_{x_1}^{x_1} \mid s_{\neg x_1}^{x_1}, next) = p(s_f^{x_1} \mid s_b^{x_1}, next) = p(s_f^{x_1} \mid s_f^{x_1}, next) = 1$.
- $O(s_{end}) = b$, $O(s_a^b) = a$ and $O(s^a) = a$ when a is a boolean variable, its negation or b . For $i = 1 \dots n$, $O(s_{end}^{x_i}) = O(s_f^{x_i}) = O(s_f^{\neg x_i}) = 2i - 1$ and $O(s_{end}^{y_i}) = O(s_f^{y_i}) = O(s_f^{\neg y_i}) = 2i$, and for any other state s , $O(s) = \#$.

The initial distribution μ_0 is $\mathbf{1}_{s_{init}}$ and the secret paths are those visiting s_{end} ($Sec = \{s_{end}\} \cup \{s_{end}^{z_i} \mid z \in \{x, y\}, i = 1 \dots n\}$).

In a similar fashion as what was done in the hardness part of the proof of Proposition 22, we show that ϕ is true iff the disclosure of M is equal to 0 in $2(n+m+2)$ steps. First remark that a path ρ reaching s_{end} can be extended for all $j \in \{1, \dots, 2n\}$ in a path ρ_j such that $O(\rho_j) = O(\rho)j$. Moreover, ρ_1 discloses the secret iff x_1 and $\neg x_1$ both occur in $O(\rho_1)$ (and similarly for the other ρ_j s). Indeed a path reaching S_{x_1} or $S_{\neg x_1}$ cannot have triggered both observations x_1 and $\neg x_1$ and also end with observation 1.

Intuitively, during the first $2n$ steps, all boolean variables will be assigned a value: either chosen by the strategy as it chooses whether x_i or $\neg x_i$ occurs in the observation for each $1 \leq i \leq n$, or randomly as y_i and $\neg y_i$ both have half a chance of being triggered. During the last $2m+1$ steps, the strategy must choose a boolean formula in every clause so that if a clause is not satisfied by the current assignment, then a boolean variable will be observed as both true and false during the path. The last step triggers then randomly the observation j for $j \in \{1, \dots, 2n\}$.

Suppose that ϕ is equivalent to true. Then there exist functions $(f_i)_{i=1\dots n}$ such that for every set of assignments (a_1, \dots, a_n) for the variables y_1, \dots, y_n the boolean formula $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is true. We choose a strategy σ such that for every possible set of assignments (a_1, \dots, a_n) for the variables y_1, \dots, y_n , and for all i , $0 \leq i \leq n-1$, $\sigma(\#f_1()a_1f_2(a_1)\dots a_i) = f_{i+1}(a_1, \dots, a_i)$. Moreover for $k \in \{1, \dots, m\}$ and $i_1, \dots, i_k \in \{1, 2, 3\}$, there exists $z_{k+1_{i_k+1}} \in \{f_1(), a_1, f_2(a_1), \dots, a_n\}$ such that $\sigma(\#f_1()a_1f_2(a_1)\dots a_n\#z_{i_1}z_{i_2}\dots z_{i_k}) = z_{k+1_{i_k+1}}$. The choice of the strategy is arbitrary in the other cases. Since $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is true, every clause is satisfied by this choice of assignments, hence it is possible to define such a strategy.

With this strategy, the fixed horizon disclosure in $2(n+m+2)$ steps is 0. In other words, none of the paths reaching the secret discloses it. Indeed let ρ be a secret path of length $2(n+m+2)$, then there exist $a_1, \dots, a_n \in \{y_1, \neg y_1, \dots, y_n, \neg y_n\}$ and $j \in \{1, \dots, 2n\}$ such that $O(\rho) = \#f_1()a_1f_2(a_1)\dots a_n\#z_{i_1}z_{i_2}\dots z_{i_m} \flat j$. By choice of σ , if, for $z \in \{x, y\}$ and $i \in \{1, \dots, n\}$, z_i appears in the observation of ρ , $\neg z_i$ does not, and vice versa. Therefore as $\flat j$ can be read either from s_{end} or in a box state outside of the secret reachable only by paths that do not observe a boolean variable and its opposite, ρ does not disclose the secret.

Conversely, suppose that ϕ is not equivalent to true and let σ be an arbitrary strategy. We first build a deterministic strategy σ' with smaller or equal disclosure. The first choice concerns $\{x_1, \neg x_1\}$ and the next observation in a path corresponds to the choice. Consider σ_1 (resp. σ'_1) the strategy that selects x_1 (resp. $\neg x_1$) and then plays like σ . Due to the fact that observations are distinct, the disclosure value w.r.t. σ is a convex combination of those of σ_1 and σ'_1 . So one substitutes σ by the one with smaller or equal disclosure. A similar pattern applies for every choice until reaching the horizon. Thus by iterating this transformation we obtain a deterministic strategy. So we assume now that σ is deterministic. Since there is a finite number of such strategies for fixed horizon, it only remains to prove that the disclosure value under σ is positive. We build partial functions $f_i : \Sigma^{2i} \mapsto \text{Act}$ consistent with σ : for every observation $\#w \in \#\Sigma^{2i}$ of some path ρ , if σ chooses action $a \in A(\text{last}(\rho))$ for ρ , then we set $f_i(w) = a$. As ϕ is not equivalent to true, there exists an assignments (a_1, \dots, a_n) for the variables y_1, \dots, y_n such that the boolean formula $\psi[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$ is false.

We now build a path disclosing the secret. By construction, there exists ρ such that $O(\rho) = \#f_1()a_1f_2(a_1)\dots f_n(a_1\dots a_{n-1})a_n\#$, leading to $\text{last}(\rho) = s_\#$ with $\mathbf{P}_\sigma(\rho) > 0$. Let $i \in \{1, \dots, m\}$ such that the negations of z_{i_1}, z_{i_2} and z_{i_3} were chosen as assignment hence $z_{i_1} \vee z_{i_2} \vee z_{i_3}$ is not true under the assignment $[f_1(), a_1, f_2(a_1), \dots, f_n(a_1, \dots, a_{n-1}), a_n]$. Let ρ' be a path of length $2(n+m)+3$ extending ρ and ending in s_{end} . Then ρ' does not disclose the secret because there exists $j \in \{1, 2, 3\}$ such that z_{i_j} appears in the previous $2m$ observations while its negation (written $\neg z_{i_j}$ here) appears in the first $2n+1$ observations. Let ρ'' of length $2(n+m+2)$ extending ρ' by ending in $s_{end}^{z_{i_j}}$. There is no other path with the same observation and ρ'' is a secret path, thus ρ'' discloses the secret. Therefore the disclosure of M is strictly greater than 0.

Observe that this reduction also works for finite horizon since no further disclosure may occur after the first occurrence of a state in $\{s_{end}^{x_1}, \dots, s_{end}^{x_n}, s_{end}^{y_1}, \dots, s_{end}^{y_n}\}$. ◀