

Selective Monitoring[☆]

Radu Grigore¹

University of Kent

Stefan Kiefer²

University of Oxford

Abstract

We study selective monitors for labelled Markov chains. Monitors observe the outputs that are generated by a Markov chain during its run, with the goal of identifying runs as correct or faulty. A monitor is selective if it skips observations in order to reduce monitoring overhead. We are interested in monitors that minimize the expected number of observations. We establish an undecidability result for selectively monitoring general Markov chains. On the other hand, we show for non-hidden Markov chains (where any output identifies the state the Markov chain is in) that simple optimal monitors exist and can be computed efficiently, based on DFA language equivalence. These monitors do not depend on the precise transition probabilities in the Markov chain. We report on experiments where we compute these monitors for several open-source Java projects.

Keywords: runtime monitoring, probabilistic systems, Markov chains, automata, language equivalence

1. Introduction

Consider an MC (Markov chain) whose transitions are labelled with letters, and a finite automaton that accepts languages of infinite words. Computing the probability that the random word emitted by the MC is accepted by the automaton is a classical problem at the heart of probabilistic verification. A finite prefix may already determine whether the random infinite word is accepted, and computing the probability that such a *deciding* finite prefix is produced is a nontrivial *diagnosability* problem. The theoretical problem we study in this paper is how to catch deciding prefixes without observing the whole prefix; i.e.,

[☆]A preliminary version of this paper appeared under the same title at Concur 2018.

¹supported by EPSRC grant EP/R012261/1

²supported by a Royal Society University Research Fellowship

10 we want to minimize the expected number of observations and still catch all
deciding prefixes.

Motivation. In runtime verification a program sends messages to a monitor, which decides if the program run is faulty. Usually, runtime verification is turned off in production code because monitoring overhead is prohibitive. QVM (quality
15 virtual machine) and ARV (adaptive runtime verification) are existing pragmatic solutions to the overhead problem, which perform best-effort monitoring within a specified overhead budget [1, 2]. ARV relies on RVSE (runtime verification with state estimation) to also compute a probability that the program run is faulty [3, 4]. We take the opposite approach: we ask for the smallest overhead
20 achievable without compromising precision at all.

Previous Work. Before worrying about the performance of a monitor, one might want to check if faults in a given system can be diagnosed at all. This problem has been studied under the term *diagnosability*, first for non-stochastic finite discrete event systems [5], which are labelled transition systems. It was
25 shown in [6] that diagnosability can be checked in polynomial time, although the associated monitors may have exponential size. Later the notion of diagnosability was extended to stochastic discrete-event systems, which are labelled Markov chains [7]. Several notions of diagnosability in stochastic systems exist, and some of them have several names, see, e.g., [8, 9] and the references therein.
30 Bertrand et al. [9] also compare the notions. For instance, they show that for one variant of the problem (referred to as A-diagnosability or SS-diagnosability or IF-diagnosability) a previously proposed polynomial-time algorithm is incorrect, and prove that this notion of diagnosability is PSPACE-complete. Indeed, most variants of diagnosability for stochastic systems are PSPACE-complete [9], with
35 the notable exception of AA-diagnosability (where the monitor is allowed to diagnose wrongly with arbitrarily small probability), which can be solved in polynomial time [10].

Selective Monitoring. In this paper, we seem to make the problem harder: since observations by a monitor come with a performance overhead, we allow the
40 monitor to skip observations. In order to decide how many observations to skip, the monitor employs an *observation policy*. Skipping observations might decrease the probability of deciding (whether the current run of the system is faulty or correct). We do not study this tradeoff: we require policies to be *feasible*, i.e., the probability of deciding must be as high as under the policy that observes
45 everything. We do not require the system to be diagnosable; i.e., the probability of deciding may be less than 1. Checking whether the system is diagnosable is PSPACE-complete ([9], [Theorem 8](#)).

The Cost of Decision in General Markov Chains. The *cost* (of decision) is the number of observations that the policy makes during a run of the
50 system. We are interested in minimizing the expected cost among all feasible

policies. We show that if the system is diagnosable then there exists a policy with finite expected cost, i.e., the policy may stop observing after finite expected time. (The converse is not true.) Whether the infimum cost (among feasible policies) is finite is also PSPACE-complete (Theorem 14). Whether there is a feasible
 55 policy whose expected cost is smaller than a given threshold is undecidable (Theorem 15), even for diagnosable systems.

Non-Hidden Markov Chains. We identify a class of MCs, namely non-hidden MCs, where the picture is much brighter. An MC is called *non-hidden* when each label identifies the state. Non-hidden MCs are always diagnosable.
 60 Moreover, we show that *maximally procrastinating* policies are (almost) optimal (Theorem 27). A policy is called maximally procrastinating when it skips observations up to the point where one further skip would put a decision on the current run in question. We also show that one can construct an (almost) optimal maximally procrastinating policy in polynomial time. This policy *does*
 65 *not depend* on the exact probabilities in the MC, although the expected cost under that policy does. That is, we efficiently construct a policy that is (almost) optimal regardless of the transition probabilities on the MC transitions. We also show that the infimum cost (among all feasible policies) can be computed in polynomial time (Theorem 28). Underlying these results is a theory based on
 70 automata, in particular, checking language equivalence of DFAs.

Experiments. We evaluated the algorithms presented in this paper by implementing them in Facebook Infer, and trying them on 11 of the most forked Java projects on GitHub. We found that, on average, selective monitoring can reduce the number of observations to a half.

75 2. Preliminaries

Let S be a finite set. We view elements of \mathbb{R}^S as *vectors*, more specifically as row vectors. We write $\mathbf{1}$ for the all-1 vector, i.e., the element of $\{1\}^S$. For a vector $\mu \in \mathbb{R}^S$, we denote by μ^\top its transpose, a column vector. A vector $\mu \in [0, 1]^S$ is a *distribution over S* if $\mu\mathbf{1}^\top = 1$. For $s \in S$ we write e_s for the
 80 (*Dirac*) distribution over S with $e_s(s) = 1$ and $e_s(t) = 0$ for $t \in S \setminus \{s\}$. We view elements of $\mathbb{R}^{S \times S}$ as *matrices*. A matrix $M \in [0, 1]^{S \times S}$ is called *stochastic* if each row sums up to one, i.e., $M\mathbf{1}^\top = \mathbf{1}^\top$.

For a finite alphabet Σ , we write Σ^* and Σ^ω for the finite and infinite words over Σ , respectively. We write ε for the empty word. We represent languages
 85 $L \subseteq \Sigma^\omega$ using deterministic finite automata, and we represent probability measures Pr over Σ^ω using Markov chains.

A (discrete-time, finite-state, labelled) *Markov chain (MC)* is a quadruple (S, Σ, M, s_0) where S is a finite set of states, Σ a finite alphabet, s_0 an initial state, and $M : \Sigma \rightarrow [0, 1]^{S \times S}$ specifies the transitions, such that $\sum_{a \in \Sigma} M(a)$ is a stochastic matrix. Intuitively, if the MC is in state s , then with probability $M(a)(s, s')$ it emits a and moves to state s' . For the complexity results in this

paper, we assume that all numbers in the matrices $M(a)$ for $a \in \Sigma$ are rationals given as fractions of integers represented in binary. We extend M to the mapping $M : \Sigma^* \rightarrow [0, 1]^{S \times S}$ with $M(a_1 \cdots a_k) = M(a_1) \cdots M(a_k)$ for $a_1, \dots, a_k \in \Sigma$. Intuitively, if the MC is in state s then with probability $M(u)(s, s')$ it emits the word $u \in \Sigma^*$ and moves (in $|u|$ steps) to state s' . An MC is called *non-hidden* if for each $a \in \Sigma$ all non-zero entries of $M(a)$ are in the same column. Intuitively, in a non-hidden MC, the emitted letter identifies the next state. An MC (S, Σ, M, s_0) defines the standard probability measure \Pr over Σ^ω , uniquely defined by assigning probabilities to cylinder sets $\{u\}\Sigma^\omega$, with $u \in \Sigma^*$, as follows:

$$\Pr(\{u\}\Sigma^\omega) := e_{s_0} M(u) \mathbf{1}^\top$$

A *deterministic finite automaton (DFA)* is a quintuple $(Q, \Sigma, \delta, q_0, F)$ where Q is a finite set of states, Σ a finite alphabet, $\delta : Q \times \Sigma \rightarrow Q$ a transition function, q_0 an initial state, and $F \subseteq Q$ a set of accepting states. We extend δ to $\delta : Q \times \Sigma^* \rightarrow Q$ as usual. A DFA defines a language $L \subseteq \Sigma^\omega$ as follows:

$$L := \{w \in \Sigma^\omega \mid \delta(q_0, u) \in F \text{ for some prefix } u \text{ of } w\}$$

Note that we do not require accepting states to be visited infinitely often: just once suffices. Therefore we can and will assume without loss of generality that there is f with $F = \{f\}$ and $\delta(f, a) = f$ for all $a \in \Sigma$.

90 For the rest of the paper we fix an MC $\mathcal{M} = (S, \Sigma, M, s_0)$ and a DFA $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$. We define their composition as the MC $\mathcal{M} \times \mathcal{A} := (S \times Q, \Sigma, M', (s_0, q_0))$ where $M'(a)((s, q), (s', q'))$ equals $M(a)(s, s')$ if $q' = \delta(q, a)$ and 0 otherwise. Thus, \mathcal{M} and $\mathcal{M} \times \mathcal{A}$ induce the same probability measure \Pr .

An *observation* $o \in \Sigma_\perp$ is either a letter or the special symbol $\perp \notin \Sigma$, which stands for ‘not seen’. An *observation policy* $\rho : \Sigma_\perp^* \rightarrow \{0, 1\}$ is a (not necessarily computable) function that, given the observations made so far, says whether we should observe the next letter. An observation policy ρ determines a projection $\pi_\rho : \Sigma^\omega \rightarrow \Sigma_\perp^\omega$: we have $\pi_\rho(a_1 a_2 \dots) = o_1 o_2 \dots$ when

$$o_{n+1} = \begin{cases} a_{n+1} & \text{if } \rho(o_1 \dots o_n) = 1 \\ \perp & \text{if } \rho(o_1 \dots o_n) = 0 \end{cases} \quad \text{for all } n \geq 0$$

We denote the *see-all policy* by \bullet ; thus, $\pi_\bullet(w) = w$.

95 In the rest of the paper we reserve a for letters, o for observations, u for finite words, w for infinite words, v for finite observation prefixes, s for states from an MC, and q for states from a DFA. We write $o_1 \sim o_2$ when o_1 and o_2 are the same or at least one of them is \perp . We lift this relation to (finite and infinite) sequences of observations (of the same length). We write $w \gtrsim v$ when $u \sim v$ holds for the length- $|v|$ prefix u of w .

100 In the following, L is the language of a DFA (the specification) and \Pr the probability measure from the Markov chain (system) to be monitored. We say that v is *negatively deciding* when $\Pr(\{w \gtrsim v \mid w \in L\}) = 0$. Intuitively, v is negatively deciding when v is incompatible with L (up to a null set). Similarly, 105 we say that v is *positively deciding* when $\Pr(\{w \gtrsim v \mid w \notin L\}) = 0$. An

observation prefix v is *deciding* when it is positively or negatively deciding. An observation policy ρ *decides* w when $\pi_\rho(w)$ has a deciding prefix. A *monitor* is an interactive (online) algorithm that implements an observation policy: it processes a stream of letters and, after each letter, it replies with one of ‘yes’, ‘no’, or ‘skip n letters’, where $n \in \mathbb{N} \cup \{\infty\}$. (It may even skip letters at the very beginning.)

Lemma 1. *For any w , if some policy decides w then \bullet decides w .*

PROOF. Let ρ decide w . Then there is a deciding prefix v of $\pi_\rho(w)$. Suppose v is positively deciding, i.e., $\Pr(\{w' \succeq v \mid w' \notin L\}) = 0$. Let u be the length- $|v|$ prefix of w . Then $\Pr(\{w' \succeq u \mid w' \notin L\}) = 0$, since v can be obtained from u by possibly replacing some letters with \perp . Hence u is also positively deciding. Since u is a prefix of $w = \pi_\bullet(w)$, we have that \bullet decides w . The case where v is negatively deciding is similar. \square

It follows that $\max_\rho \Pr(\{w \mid \rho \text{ decides } w\}) = \Pr(\{w \mid \bullet \text{ decides } w\})$. We say that a policy ρ is *feasible* when it also attains the maximum, i.e., when

$$\Pr(\{w \mid \rho \text{ decides } w\}) = \Pr(\{w \mid \bullet \text{ decides } w\}).$$

Equivalently, ρ is feasible when $\Pr(\{w \mid \bullet \text{ decides } w \text{ implies } \rho \text{ decides } w\}) = 1$, i.e., almost all words that are decided by the see-all policy are also decided by ρ . If $v = o_1 o_2 \dots$ is the shortest prefix of $\pi_\rho(w)$ that is deciding, then the *cost of decision* $C_\rho(w)$ is $\sum_{k=0}^{|v|-1} \rho(o_1 \dots o_k)$. This paper is about finding feasible observation policies ρ that minimize $\text{Ex}(C_\rho)$, the expectation of the cost of decision with respect to \Pr .

3. Qualitative Analysis of Observation Policies

In this section we study properties of observation policies that are qualitative, i.e., not directly related to the cost of decision. We focus on properties of observation prefixes that a policy may produce.

Observation Prefixes. We have already defined deciding observation prefixes. We now define several other types of prefixes: enabled, confused, very confused, and finitary. A prefix v is *enabled* if it occurs with positive probability, $\Pr(\{w \succeq v\}) > 0$. Intuitively, the other types of prefixes v are defined in terms of what would happen if we were to observe all from now on: if it is not almost sure that eventually a deciding prefix is reached, then we say v is confused; if it is almost sure that a deciding prefix will not be reached, then we say v is very confused; if it is almost sure that eventually a deciding or very confused prefix is reached, then we say v is finitary. To say this formally, let us make a few notational conventions: for an observation prefix v , we write $\Pr(v)$ as a shorthand for $\Pr(\{uw \mid u \sim v\})$; for a set Υ of observation prefixes, we write $\Pr(\Upsilon)$ as a shorthand for $\Pr(\bigcup_{v \in \Upsilon} \{uw \mid u \sim v\})$. With these conventions, we define:


1. v is *confused* when $\Pr(\{vu \mid vu \text{ deciding}\}) < \Pr(v)$
2. v is *very confused* when $\Pr(\{vu \mid vu \text{ deciding}\}) = 0$
3. v is *finitary* when $\Pr(\{vu \mid vu \text{ deciding or very confused}\}) = \Pr(v)$

Observe that (a) confused implies enabled, (b) deciding implies not confused, and
 145 (c) enabled and very confused implies confused. The following are alternative equivalent definitions (see [Appendix A.1](#)):

1. v is *confused* when $\Pr(\{uw \mid u \sim v, \text{ no prefix of } vw \text{ is deciding}\}) > 0$
2. v is *very confused* when vu' is non-deciding for all enabled vu'
3. v is *finitary* when

$$\Pr(\{uw \mid u \sim v, \text{ no prefix of } vw \text{ is deciding or very confused}\}) = 0.$$

Example 2. Consider the MC and the DFA depicted here:



150

All observation prefixes that do not start with b are enabled. The observation prefixes ab and $\perp b$ and, in fact, all observation prefixes that contain b , are positively deciding. For all $n \in \mathbb{N}$ we have $\Pr(\{w \succeq a^n \mid w \in L\}) > 0$ and $\Pr(\{w \succeq a^n \mid w \notin L\}) > 0$, so a^n is not deciding. If the MC takes the right
 155 transition first then almost surely it emits b at some point. Thus $\Pr(\{aaa \dots\}) = \frac{1}{2}$. Hence ε is confused. In this example only non-enabled observation prefixes are very confused. It follows that ε is not finitary.

Beliefs. For any s we write \Pr_s for the probability measure of the MC \mathcal{M}_s obtained from \mathcal{M} by making s the initial state. For any q we write $L_q \subseteq \Sigma^\omega$ for the language of the DFA \mathcal{A}_q obtained from \mathcal{A} by making q the initial state. We call a pair (s, q) *negatively deciding* when $\Pr_s(L_q) = 0$; similarly, we call (s, q) *positively deciding* when $\Pr_s(L_q) = 1$. A subset of $S \times Q$ is called *belief*. We call a belief *negatively (positively, respectively) deciding* when all its elements are. We fix the notation $B_0 := \{(s_0, q_0)\}$ (for the *initial belief*) for the remainder of the paper. Define the *belief NFA* as the NFA $\mathcal{B} = (S \times Q, \Sigma_\perp, \Delta, B_0, \emptyset)$ with:

$$\begin{aligned} \Delta((s, q), a) &= \{(s', q') \mid M(a)(s, s') > 0, \delta(q, a) = q'\} \quad \text{for } a \in \Sigma \\ \Delta((s, q), \perp) &= \bigcup_{a \in \Sigma} \Delta((s, q), a) \end{aligned}$$

We extend the transition function $\Delta : (S \times Q) \times \Sigma_\perp \rightarrow 2^{S \times Q}$ to $\Delta : 2^{S \times Q} \times \Sigma_\perp^* \rightarrow 2^{S \times Q}$ in the way that is usual for NFAs. Intuitively, if belief B is the set of states
 160 where the product $\mathcal{M} \times \mathcal{A}$ could be now, then $\Delta(B, v)$ is the belief adjusted by additionally observing v . To reason about observation prefixes v algorithmically, it will be convenient to reason about the belief $\Delta(B_0, v)$.

We define confused, very confused, and finitary beliefs as follows:

1. B is *confused* when $\Pr_s(\{uw \mid \Delta(B, u) \text{ deciding}\}) < 1$ for some $(s, q) \in B$
- 165 2. B is *very confused* when, for all u , we have $\Delta(B, u)$ is empty or not deciding
3. B is *finitary* when $\Pr_s(\{uw \mid \Delta(B, u) \text{ deciding or very confused}\}) = 1$ for all $(s, q) \in B$

Example 3. In [Example 2](#) we have $B_0 = \{(s_0, q_0)\}$, and $\Delta(B_0, a^n) = \{(s_1, q_0), (s_2, q_0)\}$ for all $n \geq 1$, and $\Delta(B_0, b) = \emptyset$, and $\Delta(B_0, a\perp) = \{(s_1, q_0), (s_2, q_0), (s_2, f)\}$, and $\Delta(B_0, \perp v) = \{(s_2, f)\}$ for all v that contain b . The latter belief $\{(s_2, f)\}$ is *positively deciding*. We have $\Pr_{s_1}(\{uw \mid \Delta(\{(s_1, q_0), (s_2, q_0)\}, u) \text{ is deciding}\}) = 0$, so any belief that contains both (s_1, q_0) and (s_2, q_0) is *confused*. Also, B_0 is *confused* as $\Pr_{s_0}(\{uw \mid \Delta(\{(s_0, q_0)\}, u) \text{ is deciding}\}) = \frac{1}{2}$.

175 *Relation Between Observation Prefixes and Beliefs.* By the following lemma, the corresponding properties of observation prefixes and beliefs are closely related.

Lemma 4. *Let v be an observation prefix.*

1. v is *enabled* if and only if $\Delta(B_0, v) \neq \emptyset$.
2. v is *negatively deciding* if and only if $\Delta(B_0, v)$ is *negatively deciding*.
- 180 3. v is *positively deciding* if and only if $\Delta(B_0, v)$ is *positively deciding*.
4. v is *confused* if and only if $\Delta(B_0, v)$ is *confused*.
5. v is *very confused* if and only if $\Delta(B_0, v)$ is *very confused*.
6. v is *finitary* if and only if $\Delta(B_0, v)$ is *finitary*.

PROOF. We have:

$$\begin{aligned} \Pr(\{w \succeq v\}) &= \sum_{u \sim v} \Pr(\{u\} \Sigma^\omega) \\ &= \sum_{u \sim v} \sum_s M(u)(s_0, s) \end{aligned} \quad (1)$$

Similarly we have:

$$\begin{aligned} \Pr(\{w \succeq v \mid w \in L\}) &= \sum_{u \sim v} \Pr(\{u\} \Sigma^\omega \cap L) \\ &= \sum_{u \sim v} \sum_s M(u)(s_0, s) \cdot \Pr_s(L_{\delta(q_0, u)}) \end{aligned} \quad (2)$$

Further, we have for all $B \subseteq S \times Q$:

$$\begin{aligned} \Delta(B, v) &= \bigcup_{u \sim v} \Delta(B, u) \\ &= \bigcup_{u \sim v} \{(s', q') \mid \exists (s, q) \in B : M(u)(s, s') > 0, \delta(q, u) = q'\} \end{aligned} \quad (3)$$

We prove item 1:

$$\begin{aligned}
& v \text{ is enabled} \\
\iff & \Pr(\{w \succeq v\}) > 0 && \text{definition} \\
\iff & \exists u \sim v. \exists s \text{ with } M(u)(s_0, s) > 0 && \text{by (1)} \\
\iff & \exists (s, q) \in \Delta(\{(s_0, q_0)\}, v) && \text{by (3)} \\
\iff & \Delta(B_0, v) \neq \emptyset && \text{definition}
\end{aligned}$$

We prove item 2:

$$\begin{aligned}
& v \text{ is negatively deciding} \\
\iff & \Pr(\{w \succeq v \mid w \in L\}) = 0 && \text{definition} \\
\iff & \forall u \sim v. \forall s \text{ with } M(u)(s_0, s) > 0 : \Pr_s(L_{\delta(q_0, u)}) = 0 && \text{by (2)} \\
\iff & \forall (s, q) \in \Delta(B_0, v) : \Pr_s(L_q) = 0 && \text{by (3)} \\
\iff & \text{all } (s, q) \in \Delta(B_0, v) \text{ are negatively deciding} && \text{definition} \\
\iff & \Delta(B_0, v) \text{ is negatively deciding} && \text{definition}
\end{aligned}$$

The proof of item 3 is similar. Towards item 4 note that if v is not deciding then none of its prefixes is. This property explains the second equivalence in the following proof of item 4. There we write $u' < w$ to denote that u' is a finite prefix of w .

$$\begin{aligned}
& v \text{ is confused} \\
\iff & \Pr(\{uw \mid u \sim v, \text{ no prefix of } vw \text{ is deciding}\}) > 0 && \text{definition} \\
\iff & \Pr(\{uw \mid u \sim v, \forall u' < w : vu' \text{ is not deciding}\}) > 0 && \text{see above} \\
\iff & \sum_{u \sim v} \sum_s M(u)(s_0, s) \cdot && \text{as before} \\
& \Pr_s(\{w \mid \forall u' < w : vu' \text{ is not deciding}\}) > 0 \\
\iff & \exists u \sim v. \exists s : M(u)(s_0, s) > 0 \text{ and} && \text{as before} \\
& \Pr_s(\{w \mid \forall u' < w : vu' \text{ is not deciding}\}) > 0 \\
\iff & \exists (s, q) \in \Delta(B_0, v) : && \text{by (3)} \\
& \Pr_s(\{w \mid \forall u' < w : vu' \text{ is not deciding}\}) > 0 \\
\iff & \exists (s, q) \in \Delta(B_0, v) : && \text{items 2, 3} \\
& \Pr_s(\{w \mid \forall u' < w : \Delta(\Delta(B_0, v), u') \text{ is not deciding}\}) > 0 \\
\iff & \Delta(B_0, v) \text{ is confused} && \text{definition}
\end{aligned}$$

We prove item 5:

$$\begin{aligned}
& v \text{ is very confused} \\
\iff & \forall u : (vu \text{ is enabled} \rightarrow vu \text{ is not deciding}) && \text{definition} \\
\iff & \forall u : (\Delta(B_0, vu) \neq \emptyset \rightarrow \Delta(B_0, vu) \text{ is not deciding}) && \text{items 2, 3} \\
\iff & \Delta(B_0, v) \text{ is very confused} && \text{definition}
\end{aligned}$$

The proof of item 6 is similar to the proof of item 4. We abbreviate “not deciding and not very confused” to “continuing”. Note that if v is continuing then all its prefixes are. This property explains the second equivalence in the following proof of item 6. There we write $u' < w$ to denote that u' is a finite prefix of w .

$$\begin{aligned}
& v \text{ is finitary} \\
\iff & \Pr(\{uw \mid u \sim v, \text{ all } u' < vw \text{ are continuing}\}) = 0 && \text{definition} \\
\iff & \Pr(\{uw \mid u \sim v, \forall u' < w : vu' \text{ is continuing}\}) = 0 && \text{see above} \\
\iff & \sum_{u \sim v} \sum_s M(u)(s_0, s) \cdot && \text{as before} \\
& \quad \cdot \Pr_s(\{w \mid \forall u' < w : vu' \text{ is continuing}\}) = 0 \\
\iff & \forall u \sim v. \forall s : M(u)(s_0, s) = 0 \text{ or} && \text{as before} \\
& \quad \Pr_s(\{w \mid \forall u' < w : vu' \text{ is continuing}\}) = 0 \\
\iff & \forall (s, q) \in \Delta(B_0, v) : && \text{by (3)} \\
& \quad \Pr_s(\{w \mid \forall u' < w : vu' \text{ is continuing}\}) = 0 \\
\iff & \forall (s, q) \in \Delta(B_0, v) : && \text{items 2, 3, 5} \\
& \quad \Pr_s(\{w \mid \forall u' < w : \Delta(B_0, vu') \text{ is continuing}\}) = 0 \\
\iff & \forall (s, q) \in \Delta(B_0, v) : && \text{definitions} \\
& \quad \Pr_s(\{u'w \mid \Delta(\Delta(B_0, v), u') \text{ is not continuing}\}) = 1 \\
\iff & \Delta(B_0, v) \text{ is finitary} && \text{definition}
\end{aligned}$$

This proves item 6 and completes the proof of the lemma. \square

185 The following lemma gives complexity bounds for computing these properties.

Lemma 5. *Let v be an observation prefix, and B a belief.*

1. *Whether v is enabled can be decided in P .*
2. *Whether v (or B) is negatively deciding can be decided in P .*
3. *Whether v (or B) is positively deciding can be decided in P .*
- 190 4. *Whether v (or B) is confused can be decided in $PSPACE$.*
5. *Whether v (or B) is very confused can be decided in $PSPACE$.*
6. *Whether v (or B) is finitary can be decided in $PSPACE$.*

PROOF (SKETCH, FULL PROOF IN [APPENDIX A.2](#)). The belief NFA \mathcal{B} and the MC $\mathcal{M} \times \mathcal{A}$ can be computed in polynomial time (even in deterministic logspace).
195 For items 1–3, there are efficient graph algorithms that search these product structures. For instance, to show that a given pair (s_1, q_1) is not negatively deciding, it suffices to show that \mathcal{B} has a path from (s_1, q_1) to a state (s_2, f) for some s_2 . This can be checked in polynomial time (even in NL).

For items 4–6, one searches the (exponential-sized) product of \mathcal{M} and the *determinization* of \mathcal{B} . This can be done in $PSPACE$. For instance, to show that
200 a given belief B is confused, it suffices to show that there are $(s_1, q_1) \in B$ and u_1 and s_2 such that \mathcal{M} has a u_1 -labelled path from s_1 to s_2 such that there

do not exist u_2 and s_3 such that \mathcal{M} has a u_2 -labelled path from s_2 to s_3 such that $\Delta(B, u_1 u_2)$ is deciding. This can be checked in $\text{NPSpace} = \text{PSPACE}$ by
 205 nondeterministically guessing paths in the product of \mathcal{M} and the determinization of \mathcal{B} . \square

Diagnosability. We call a policy a *diagnoser* when it decides almost surely.

Example 6. In [Example 2](#) a diagnoser does not exist. Indeed, the policy \bullet does not decide when the MC takes the left transition, and decides (positively)
 210 almost surely when the MC takes the right transition in the first step. Hence $\Pr(\{w \mid \bullet \text{ decides } w\}) = \Pr(\Sigma^* \{b\} \Sigma^\omega) = \frac{1}{2}$. So \bullet is not a diagnoser. By [Lemma 1](#), it follows that there is no diagnoser.

Diagnosability can be characterized by the notion of confusion:

Proposition 7. *There exists a diagnoser if and only if ε is not confused.*

PROOF. We have:

there exists a diagnoser	
\iff there exists a policy that decides almost surely	definition of diagnoser
\iff \bullet decides almost surely	Lemma 1
\iff $\Pr(\{w \mid w \text{ has a deciding prefix}\}) = 1$	definitions
\iff ε is not confused	definition of confused

215 \square


The following proposition shows that diagnosability is hard to check.

Theorem 8 (cf. [9, Theorem 6]). *Given an MC \mathcal{M} and a DFA \mathcal{A} , it is PSPACE-complete to check if there exists a diagnoser.*

[Theorem 8](#) essentially follows from a result by Bertrand et al. [9]. They study
 220 several different notions of diagnosability; one of them (*FA-diagnosability*) is very similar to our notion of diagnosability. There are several small differences; e.g., their systems are not necessarily products of an MC and a DFA. Therefore we give a self-contained proof of [Theorem 8](#).

PROOF (OF [THEOREM 8](#)). By [Proposition 7](#) it suffices to show PSPACE-completeness of checking whether ε is confused. Membership in PSPACE follows
 225 from [Lemma 5.4](#).

For hardness we reduce from the following problem: given an NFA \mathcal{U} over $\Sigma = \{a, b\}$ where all states are initial and accepting, does \mathcal{U} accept all (finite) words? This problem is PSPACE-complete [11, Lemma 6]. Let $\mathcal{U} = (Q, \Sigma, \delta, Q, Q)$ be
 230 the given NFA. We construct an MC $\mathcal{M} = (Q \cup \{s_0, s_1, s_2\}, \Sigma \cup \{\#\}, M, s_0)$ where $s_0, s_1, s_2 \notin Q$, and $\# \notin \Sigma = \{a, b\}$, and the transitions are as follows:



In this picture, the rectangle labelled with Q indicates states and transitions that involve the states in Q , i.e., those coming from the NFA \mathcal{U} . In more detail we define:

- $M(a)(s_0, q) := \frac{1}{2|Q|}$ for all $q \in Q$;
- $x(q) := |\delta(q, a)| + |\delta(q, b)| + 1$ for all $q \in Q$;
- $M(\sigma)(q, q') := \frac{1}{x(q)}$ for all $q \in Q$ and both $\sigma \in \{a, b\}$ and all $q' \in \delta(q, \sigma)$;
- $M(\#)(q, s_2) := \frac{1}{x(q)}$ for all $q \in Q$.

So for each $q \in Q$ there is a uniform distribution over the outgoing transitions. Define \mathcal{A} to be a DFA that accepts $L = \Sigma^* \{\#\} (\Sigma \cup \{\#\})^*$. We show that ε is confused if and only if \mathcal{U} accepts all words.

Suppose \mathcal{U} does not accept the word $u \in \Sigma^*$. Then one of the following two events happens almost surely, depending on whether \mathcal{M} takes the left or the right transition in the first step:

- \mathcal{M} emits $au_0\#$ for some $u_0 \in \Sigma^*$;
- \mathcal{M} emits au_0u for some $u_0 \in \Sigma^*$.

We claim that both $au_0\#$ and au_0u are deciding, implying that ε is not confused. It is clear that $au_0\#$ is positively deciding. We argue that au_0u is negatively deciding. Indeed, since \mathcal{U} does not accept u , we have $\delta(q, u) = \emptyset$ for all $q \in Q$. Hence, starting from any $q \in Q$, the MC \mathcal{M} cannot emit u . So if \mathcal{M} emits au_0u , it must have taken the right transition in the first step and thus will never emit $\#$, which means that au_0u is negatively deciding.

Conversely, suppose that ε is not confused. Then, almost surely, \mathcal{M} emits a deciding prefix u' . By our construction, the word u' is positively deciding with probability $\frac{1}{2}$, and negatively deciding with probability $\frac{1}{2}$ (depending on which transition \mathcal{M} takes in the first step). Hence there exists $u \in \Sigma^*$ such that au is negatively deciding. So \mathcal{M} cannot take the left transition in the first step and then emit u . It follows that there is no $q \in Q$ with $\delta(q, u) \neq \emptyset$, i.e., \mathcal{U} does not accept u . \square

Allowing Confusion. We say an observation policy *allows confusion* when, with positive probability, it produces an observation prefix $v\perp$ such that $v\perp$ is confused but v is not.

Proposition 9. *A feasible observation policy does not allow confusion.*

PROOF. Suppose that ρ allows confusion, i.e., with positive probability, it produces an observation prefix $v\perp$ such that $v\perp$ is confused but v is not. We have:

$$\begin{aligned}
& \Pr(\{w \succsim v \mid \bullet \text{ decides } w\}) \\
&= \Pr(\{uw \mid u \sim v, uw \text{ has a deciding prefix}\}) && \text{definitions} \\
&\geq \Pr(\{uw \mid u \sim v, vw \text{ has a deciding prefix}\}) && (4) \\
&= \Pr(\{uw \mid u \sim v\}) && v \text{ is not confused} \\
&= \Pr(\{w \succsim v\}) && \text{definition of } \succsim
\end{aligned}$$

Thus we have:

$$\begin{aligned}
& \Pr(\{w \mid \bullet \text{ decides } w, \rho \text{ does not decide } w\}) \\
&\geq \Pr(\{w \succsim v \mid \bullet \text{ decides } w, \rho \text{ does not decide } w\}) \\
&\geq \Pr(\{w \succsim v \mid \rho \text{ does not decide } w\}) && \text{by (4)} \\
&= \Pr(\{w \succsim v\perp \mid \rho \text{ does not decide } w\}) \\
&= \Pr(\{uw \mid u \sim v\perp, \rho \text{ does not decide } w\}) && \text{definition of } \succsim \\
&= \Pr(\{uw \mid u \sim v\perp, \pi_\rho(uw) \text{ has no deciding prefix}\}) && \text{definition of decide} \\
&\geq \Pr(\{uw \mid u \sim v\perp, v\perp w \text{ has no deciding prefix}\}) && \text{definition of } \rho \\
&> 0 && v\perp \text{ is confused}
\end{aligned}$$

265 Hence ρ is not feasible. □

Hence, in order to be feasible, a policy must observe when it would get confused otherwise. In [Section 5](#) we show that in the non-hidden case there is almost a converse of [Proposition 9](#); i.e., in order to be feasible, a policy need not do much more than not allow confusion.

270 4. Analyzing the Cost of Decision

In this section we study the computational complexity of finding feasible policies that minimize the expected cost of decision. We focus on the decision version of the problem: *Is there a feasible policy whose expected cost is smaller than a given threshold?* Define:

$$c_{inf} := \inf_{\text{feasible } \rho} \text{Ex}(C_\rho)$$

Since the see-all policy \bullet never stops observing, we have $\Pr(C_\bullet = \infty) = 1$, so $\text{Ex}(C_\bullet) = \infty$. However, once an observation prefix v is deciding or very confused, there is no point in continuing observation. Hence, we define a *light see-all* policy \circ , which observes until the observation prefix u is deciding or very confused; formally, $\circ(v) = 0$ if and only if v is deciding or very confused.

It follows from the definition of very confused that the policy \circ is feasible. Concerning the cost C_\circ we have for all w

$$C_\circ(w) = \sum_{n=0}^{\infty} (1 - D_n(w)), \quad (5)$$

where $D_n(w) = 1$ if the length- n prefix of w is deciding or very confused, and $D_n(w) = 0$ otherwise.

The following lemma holds because, in $\mathcal{M} \times \mathcal{A}$, a bottom strongly connected component is reached in expected finite time.

275 **Lemma 10.** *If ε is finitary then $\text{Ex}(C_\circ)$ is finite.*

PROOF. We abbreviate “deciding or very confused” to “dv”. Note that if a belief B is dv then all $\Delta(B, a)$ are dv. We use this to show the following similar property:

Claim. For any finitary belief B , all $\Delta(B, a)$ are finitary.

We prove the claim. Let B be a finitary belief. Then we have:

$$\begin{aligned} 1 &= \Pr_s(\{uw \mid \Delta(B, u) \text{ is dv}\}) \\ &= \sum_{a \in \Sigma} \sum_{s'} M(a)(s, s') \cdot \Pr_{s'}(\{uw \mid \Delta(B, au) \text{ is dv}\}) \end{aligned}$$

280 Here, the second equality follows from the property of dv mentioned above. Hence, for all $a \in \Sigma$ and all $(s', q') \in \Delta(B, a)$ we have $\Pr_{s'}(\{uw \mid \Delta(B, au) \text{ is dv}\}) = 1$, so $\Delta(B, a)$ is finitary for all a . This proves the claim.

Define:

$$\mathcal{F} := \{(s, B) \in S \times 2^{S \times Q} \mid B \text{ finitary}, \exists q : (s, q) \in B\}$$

For any belief B define:

$$e_{dv}(B) := \begin{cases} 1 & \text{if } B \text{ is dv} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

For any $(s, B) \in \mathcal{F}$ and any $n \in \mathbb{N}$ define

$$c_n(s, B) := \text{Ex}_s(1 - D_{B,n}), \quad (7)$$

where by Ex_s we denote the expectation with respect to \Pr_s and we define $D_{B,n}(w) := e_{dv}(\Delta(B, u_n))$ where u_n is the length- n prefix of w . We have $c_0(s, B) = 1 - e_{dv}(B)$. Further, due to the claim above, we have for all $n \in \mathbb{N}$:

$$c_{n+1}(s, B) = \sum_{a \in \Sigma} \sum_{\substack{s' \\ M(a)(s, s') > 0}} M(a)(s, s') \cdot c_n(s', \Delta(B, a))$$

We may write those equations in vector form:

$$\begin{aligned} c_0^\top &= \mathbf{1}^\top - e_{dv}^\top \\ c_{n+1}^\top &= U \cdot c_n^\top \end{aligned} \tag{8}$$

where

- $e_{dv} \in \{0, 1\}^\mathcal{F}$ where $e_{dv}(s, B) = e_{dv}(B)$ as in (6);
- $c_n \in \mathbb{R}^\mathcal{F}$ where $c_n(s, B)$ is as in (7);
- $U \in [0, 1]^{\mathcal{F} \times \mathcal{F}}$ where

$$U((s, B), (s', B')) = \begin{cases} \sum_{\substack{a \in \Sigma \\ \Delta(B, a) = B'}} M(a)(s, s') & \text{if } B \text{ is not } dv \\ 0 & \text{otherwise.} \end{cases}$$

The matrix U satisfies $U\mathbf{1}^\top = \mathbf{1}^\top - e_{dv}^\top$. By a straightforward induction it follows:

$$U^n \mathbf{1}^\top = \mathbf{1}^\top - \sum_{i=0}^{n-1} U^i e_{dv}^\top \quad \text{for all } n \in \mathbb{N} \tag{9}$$

From the definition of finitary, for all $(s, B) \in \mathcal{F}$ there is u such that $\Pr_s(\{u\}^\Sigma) > 0$ and $\Delta(B, u)$ is dv . It follows that there is $n \in \mathbb{N}$ such that $\sum_{i=0}^{n-1} U^i e_{dv}^\top > \mathbf{0}^\top$ where the inequality is strict in all entries. By (9), we have $U^n \mathbf{1}^\top < \mathbf{1}^\top$ where the inequality is strict in all entries. Hence the spectral radius of U is less than one, and so the matrix series $\sum_{n=0}^{\infty} U^n$ converges to a finite matrix, say $U^* \in \mathbb{R}^{\mathcal{F} \times \mathcal{F}}$.

Suppose ε is finitary. By Lemma 4.6, the belief $B_0 = \Delta(B_0, \varepsilon)$ is finitary. Hence $(s_0, B_0) \in \mathcal{F}$. Further we have:

$$\begin{aligned} \text{Ex}(C_o) &= \sum_{n=0}^{\infty} \text{Ex}(1 - D_n) && \text{by (5)} \\ &= \sum_{n=0}^{\infty} \text{Ex}_{s_0}(1 - D_{B_0, n}) && \text{def. of } D_n, D_{B_0, n}, \text{ Lemma 4} \\ &= \sum_{n=0}^{\infty} c_n(s_0, B_0) && \text{by (7)} \\ &= \sum_{n=0}^{\infty} \left(U^n (\mathbf{1}^\top - e_{dv}^\top) \right) (s_0, B_0) && \text{by (8)} \\ &= \left(U^* (\mathbf{1}^\top - e_{dv}^\top) \right) (s_0, B_0) && \text{definition of } U^* \\ &< \infty && U^* \text{ is finite} \end{aligned}$$

This completes the proof. \square

The following lemma says that a kind of converse holds for feasible policies.

Lemma 11. *Let ρ be a feasible observation policy. If $\Pr(C_\rho < \infty) = 1$ then ε is finitary.*

PROOF. Let ρ be an observation policy. Suppose ε is not finitary and $\Pr(C_\rho < \infty) = 1$. We show that ρ is not feasible.

Observe that for any w , if $C_\rho(w) < \infty$ then there is a (unique) shortest finite prefix of w , say $\tilde{u}(w)$, such that ρ never observes a letter after $\tilde{u}(w)$. Abbreviating “deciding or very confused” to “dv”, we have:

$$\begin{aligned}
0 &< \Pr(\{w \mid \text{no prefix of } w \text{ is dv}\}) && \varepsilon \text{ is not finitary} \\
&= \Pr(\{w \mid \text{no prefix of } w \text{ is dv, } C_\rho(w) < \infty\}) && \Pr(C_\rho < \infty) = 1 \\
&= \Pr(\{w \mid \text{no prefix of } w \text{ is dv, } \exists u \text{ s.t. } u = \tilde{u}(w)\}) && \text{observation above} \\
&= \Pr(\{uw \mid \text{no prefix of } uw \text{ is dv, } u = \tilde{u}(uw)\}) \\
&\leq \Pr(\{uw \mid u \text{ is not dv, } u = \tilde{u}(uw)\}) && u \text{ is a prefix of } uw \\
&\leq \sum_{u \text{ is not dv}} \Pr(\{uw \mid u = \tilde{u}(uw)\}) && \Sigma^* \text{ is countable}
\end{aligned}$$

It follows that there is u that is (i) not deciding, (ii) not very confused, and (iii) such that ρ never observes a letter after u . Since u is not very confused, there is u' such that uu' is enabled and deciding. Hence \bullet decides all words of the form $uu'w$. Since ρ never observes a letter after u , we have that ρ does not decide any words of the form $uu'w$. It follows:

$$\begin{aligned}
&\Pr(\{w \mid \bullet \text{ decides } w, \rho \text{ does not decide } w\}) \\
&\geq \Pr(\{uu'w \mid \bullet \text{ decides } uu'w, \rho \text{ does not decide } uu'w\}) \\
&= \Pr(\{uu'w\}) && \text{as explained above} \\
&> 0 && uu' \text{ is enabled}
\end{aligned}$$

Hence ρ is not feasible. □

Now we can prove the following, justifying the term “finitary”:

Proposition 12. *c_{inf} is finite if and only if ε is finitary.*

PROOF. Suppose ε is finitary. By Lemma 10 we then have that $\text{Ex}(C_o)$ is finite. Hence $c_{inf} \leq \text{Ex}(C_o)$ is finite.

Conversely, suppose ε is not finitary. By Lemma 11 we then have $\Pr(C_\rho = \infty) > 0$ for all feasible observation policies. Thus $\text{Ex}(C_\rho) = \infty$ holds for all feasible observation policies. Hence $c_{inf} = \infty$. □

Proposition 13. *If a diagnoser exists then c_{inf} is finite.*

PROOF. We have:

$$\begin{aligned}
& \text{a diagnoser exists} \\
\iff & \varepsilon \text{ is not confused} && \text{Prop. 7} \\
\iff & \Pr(\{w \mid \text{some prefix of } w \text{ is deciding}\}) = 1 && \text{definition} \\
\implies & \Pr(\{w \mid \text{some prefix of } w \text{ is deciding or very confused}\}) = 1 \\
\iff & \varepsilon \text{ is finitary} && \text{definition} \\
\iff & c_{inf} \text{ is finite} && \text{Prop. 12}
\end{aligned}$$

□

Theorem 14. *It is PSPACE-complete to check if $c_{inf} < \infty$.*

PROOF. By [Proposition 12](#) it suffices to show PSPACE-completeness of checking whether ε is finitary. Membership in PSPACE follows from [Lemma 5.6](#).

For hardness we use the same reduction as in the proof of [Theorem 8](#). There we reduce from the following PSPACE-complete problem: given an NFA \mathcal{U} over $\Sigma = \{a, b\}$ where all states are initial and accepting, does \mathcal{U} accept all (finite) words? The reduction produces an MC \mathcal{M} and a DFA \mathcal{A} .

First we show that in the result of that reduction, no enabled word is very confused. If a word u contains the letter $\#$ then u is negatively deciding, hence not very confused. Let $u_0 \in \Sigma^*$.

- If the NFA \mathcal{U} accepts all words then $au_0\#$ is enabled and positively deciding;
- if \mathcal{A} does not accept some word u then au_0u is enabled and negatively deciding.

In either case it follows that au_0 is not very confused and also that ε is not very confused. We conclude that no enabled word is very confused.

Thus we have:

$$\begin{aligned}
& \varepsilon \text{ is finitary} \\
\iff & \Pr(\{w \mid \text{no prefix of } w \text{ is deciding or very confused}\}) = 0 && \text{definition} \\
\iff & \Pr(\{w \mid \text{no prefix of } w \text{ is deciding}\}) = 0 && \text{as argued} \\
\iff & \varepsilon \text{ is not confused} && \text{definition} \\
\iff & \text{NFA } \mathcal{U} \text{ does not accept all words,}
\end{aligned}$$

where the last equivalence was shown in the proof of [Theorem 8](#). It follows that checking if c_{inf} is finite is PSPACE-complete. □

The main negative result of the paper is that one cannot compute c_{inf} :

Theorem 15. *It is undecidable to check if $c_{inf} < 3$, even when a diagnoser exists.*

The proof is by a reduction from the undecidable problem whether a given
 330 probabilistic automaton accepts some word with probability $> \frac{1}{2}$. The proof is
 somewhat complicated. First we give a wrong proof, which contains useful ideas
 and is similar to and simpler than the correct one, but is flawed. We point out
 the flaw. In [Appendix A.3](#) we amend the proof, making it more complicated but
 correct. The incorrect proof contains only ideas that feature also in the correct
 335 one; but the correct proof can be read without having read the incorrect one.

PROOF (WRONG PROOF). We reduce from the emptiness problem for proba-
 bilistic automata. A *probabilistic automaton (PA)* is a tuple $\mathcal{P} = (S, \Sigma, M, s_0, \eta)$
 where S is a finite set of states, Σ is a finite alphabet, the mapping $M : \Sigma \rightarrow$
 $[0, 1]^{S \times S}$ specifies the transitions such that $M(a)$ is stochastic for each $a \in \Sigma$, the
 340 state s_0 is an initial state, and $\eta \in [0, 1]^S$ is a vector of acceptance probabilities.
 Extend M to $M : \Sigma^* \rightarrow [0, 1]^{S \times S}$ as in the case of MCs. In the case of PAs, $M(u)$
 is stochastic for each $u \in \Sigma^*$. For each u define $\text{Pr}_{\mathcal{P}}(u) := e_{s_0} M(u) \eta^\top$. The
 probability $\text{Pr}_{\mathcal{P}}(u)$ can be interpreted as the probability that \mathcal{P} accepts u . The
emptiness problem asks, given a PA \mathcal{P} , whether there is u such that $\text{Pr}_{\mathcal{P}}(u) > \frac{1}{2}$.
 345 This problem is undecidable [[12](#), p. 190, Theorem 6.17].

Let $\mathcal{P} = (S_{\mathcal{P}}, \Sigma_{\mathcal{P}}, M_{\mathcal{P}}, s_{0\mathcal{P}}, \eta)$ be the given PA. We will construct a DFA \mathcal{A}
 and an MC \mathcal{M} over the alphabet $\Sigma := \Sigma_{\mathcal{P}} \cup \{0, 1, ?\}$, where $0, 1, ?$ are fresh letters.
 Define \mathcal{A} to be a DFA that accepts $L = 0(\Sigma\Sigma)^*\Sigma 1\Sigma^\omega$. We might characterize L
 by saying that 1 appears on an odd position.

We construct an MC $\mathcal{M} = (S, \Sigma, M, s_0)$ with

$$S := \{s_0\} \cup (S_{\mathcal{P}} \times \{0, 1, 2, 3\}),$$

350 such that s_0 is a fresh state. The MC \mathcal{M} initially splits randomly into a “0-copy”
 (with states in $S \times \{0, 2\}$) and a “1-copy” (with states in $S \times \{1, 3\}$), in either
 case emitting the letter 0. Formally, $M(0)(s_0, (s_{0\mathcal{P}}, 0)) = M(0)(s_0, (s_{0\mathcal{P}}, 1)) = \frac{1}{2}$.
 The MC \mathcal{M} is constructed such that if it goes into the 0-copy then surely it
 emits an infinite word that is not in L ; and if it goes into the 1-copy then almost
 355 surely it emits an infinite word in L . Therefore, for an observation policy it
 suffices to identify which copy \mathcal{M} has entered.

The transitions in \mathcal{M} depend on $M_{\mathcal{P}}$ and η . In [Figure 1](#) we illustrate this
 dependence with an example, where $\Sigma_{\mathcal{P}} = \{a, b\}$. For all $s, t \in S_{\mathcal{P}}$ and both
 $i \in \{0, 1\}$ we have the following transitions:

$$\begin{aligned} M(i)((s, i), (s, i+2)) &:= \frac{2}{3}\eta(s) \\ M(?)((s, i), (s, i+2)) &:= 1 - \frac{2}{3}\eta(s) \\ M(a)((s, i+2), (t, i)) &:= \frac{2}{3} \cdot \frac{1}{|\Sigma_{\mathcal{P}}|} \cdot M_{\mathcal{P}}(a)(s, t) && \forall a \in \Sigma_{\mathcal{P}} \\ M(i)((s, i+2), (t, i)) &:= \frac{1}{3} \cdot \frac{1}{|\Sigma_{\mathcal{P}}|} \cdot \sum_{a \in \Sigma_{\mathcal{P}}} M_{\mathcal{P}}(a)(s, t) \end{aligned}$$

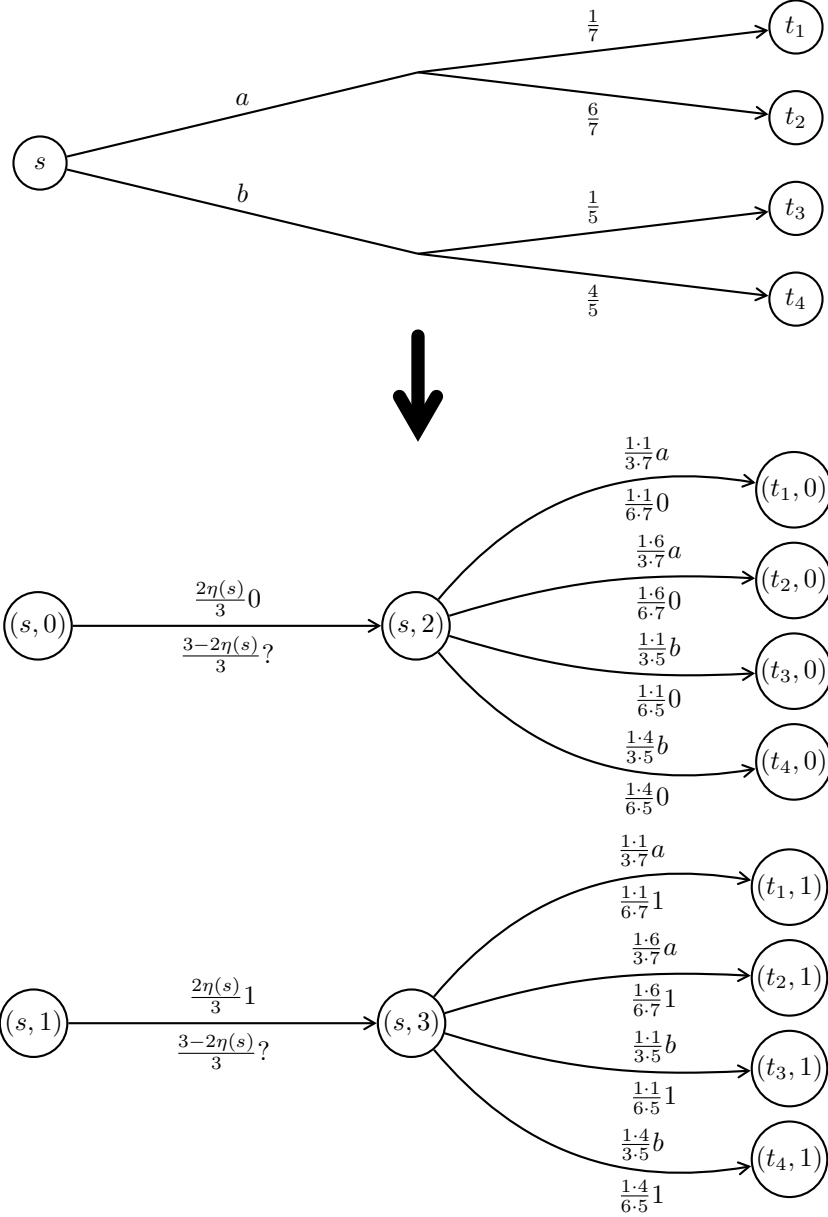


Figure 1: Illustration of the wrong reduction, based on an example PA, at the top, with a state s and its outgoing transitions. The outgoing transitions are such that $M(a)(s, t_1) = \frac{1}{7}$, $M(a)(s, t_2) = \frac{6}{7}$, $M(b)(s, t_3) = \frac{1}{5}$, $M(b)(s, t_4) = \frac{4}{5}$. The resulting MC has two corresponding states, $(s, 0)$ and $(s, 1)$. From such a state (s, i) it first emits, with a probability depending on $\eta(s)$, either the letter i (thus giving away in which copy the MC is) or the letter $?$, and then emits, with probability $\frac{1}{3}$, the letter i (thus giving away in which copy the MC is), or, with probability $\frac{2}{3} \cdot \frac{1}{|\Sigma_{\mathcal{P}}|} = \frac{1}{3}$ each, the letters a or b .

Note that any state (s, i) emits, with a probability depending on $\eta(s)$, the letter i (thus giving away in which copy \mathcal{M} is), or, with the remaining probability, the letter $?$. Any state $(s, i + 2)$ emits, with probability $\frac{1}{3}$, the letter i (thus
360 giving away in which copy \mathcal{M} is), or, with probability $\frac{2}{3}$, a letter in $\Sigma_{\mathcal{P}}$ that is chosen uniformly at random. (Thus, if $|\Sigma_{\mathcal{P}}| = 2$, all letters in $\Sigma_{\mathcal{P}} \cup \{i\}$ have probability $\frac{1}{3}$.) In more detail, one can view the behaviour of \mathcal{M} in a state $(s, i + 2)$ as follows: \mathcal{M} first samples a letter a from $\Sigma_{\mathcal{P}}$ uniformly at random, then samples a successor state from $S \times \{i\}$ (each (t, i) with probability $M_{\mathcal{P}}(a)(s, t)$),
365 and then, upon transitioning to (t, i) , emits, with probability $\frac{2}{3}$, the sampled letter a , or, with probability $\frac{1}{3}$, the letter i . This completes the description of the construction.

There exists a diagnoser: Consider the policy $\bar{\rho}$, which, after the initial 0, observes every 2nd letter; each time this letter will be either from $\Sigma_{\mathcal{P}}$ (with
370 probability $\frac{2}{3}$) or from $\{0, 1\}$ (with probability $\frac{1}{3}$). With probability 1 the latter case occurs eventually. Once it has occurred, the copy has been revealed; thus the observation prefix becomes deciding and $\bar{\rho}$ stops making any observations. The required number of observations under $\bar{\rho}$ is geometrically distributed, with expectation $1/\frac{1}{3} = 3$. Hence $\text{Ex}(C_{\bar{\rho}}) = 3$ and thus $c_{inf} \leq 3$.

It remains to show that $c_{inf} < 3$ if and only if there is u such that $\Pr_{\mathcal{P}}(u) > \frac{1}{2}$.
375 Suppose u_0 is such that $\Pr_{\mathcal{P}}(u_0) > \frac{1}{2}$. Consider the following policy ρ_0 :

- do not observe the initial 0 and then observe every 2nd letter until either the prefix is deciding (because 0 or 1 has been observed) or the observation prefix becomes equal to u_0 ;
- 380 • in the former case: stop observing forever;
- in the latter case: observe the immediately following letter;
 - if this letter is 0 or 1 then the observation prefix has become deciding, so stop observing;
 - if this letter is ? then observe the next letter and then every 2nd letter
385 as before.

This policy coincides mostly with the diagnoser $\bar{\rho}$ described above. It only deviates if and when u_0 has been observed. One can show that this deviation improves the expected cost: the conditional probability that 0 or 1 (hence not the letter ?) will be observed conditioned under having observed u_0 is equal to
390 $\frac{2}{3} \cdot \Pr_{\mathcal{P}}(u_0) > \frac{2}{3} \cdot \frac{1}{2} > \frac{1}{3}$. It follows that $c_{inf} < 3$.

It remains to show the converse. To this end, suppose all u satisfy $\Pr_{\mathcal{P}}(u) \leq \frac{1}{2}$. Then one is tempted to think that it is not beneficial to deviate from the diagnoser $\bar{\rho}$ described above: conditioned under having observed some observation prefix u using $\bar{\rho}$, the conditional probability that the immediately following letter
395 is 0 or 1 (hence not the letter ?) is equal to $\frac{2}{3} \cdot \Pr_{\mathcal{P}}(u) \leq \frac{2}{3} \cdot \frac{1}{2} \leq \frac{1}{3}$, so not better than what one would get by proceeding with $\bar{\rho}$.

The last argument is flawed though; it is valid only for the very first deviation from $\bar{\rho}$. Indeed, suppose there is some u_0 with $\Pr_{\mathcal{P}}(u_0) = \frac{1}{2}$. Suppose further we


follow $\bar{\rho}$ until the observation prefix is u_0 . At this point it does not come with a
 400 risk to observe the immediately following letter, as the probability to observe
 0 or 1 is $\frac{1}{3}$, which is equal to the success probability of each single observation
 of $\bar{\rho}$. Suppose we do that and observe ?. Now we have not learned anything
 about which copy the MC is in; however, after this observation, the conditional
 probability (conditioned on all observations so far) that we are in a state $(s, 2)$ or
 405 $(s, 3)$ with *low* $\eta(s)$ has increased, as a state (s, i) with high $\eta(s)$ would probably
 have produced the letter i instead of ?. This information might be exploitable
 later: For instance, suppose that the PA is such that states s with low and high
 values of $\eta(s)$ alternate. Then in the next round we believe to be in a state (s, i)
 with high $\eta(s)$. Now it might be beneficial to make another non- $\bar{\rho}$ observation,
 410 even if $\Pr_{\mathcal{P}}(u_0a) \leq \frac{1}{2}$ holds for all $a \in \Sigma_{\mathcal{P}}$.

To fix this problem we need to change the construction in a way that a single
 deviation from $\bar{\rho}$ leaks less information. We do this in [Appendix A.3](#). \square

5. The Non-Hidden Case

Now we turn to positive results. In the rest of the paper we assume that
 415 the MC \mathcal{M} is non-hidden, i.e., there exists a function $\vec{\cdot} : \Sigma \rightarrow S$ such that
 $M(a)(s, s') > 0$ implies $s' = \vec{a}$. We extend $\vec{\cdot}$ to finite nonempty words so that
 $\vec{u\vec{a}} = \vec{a}$. We write $s \xrightarrow{u}$ to indicate that there is s' with $M(u)(s, s') > 0$.

Example 16. Consider the following non-hidden MC and DFA:



420

$$\begin{aligned}
 B_0 &:= \{(\vec{a}, q_0)\} \\
 B_1 &:= \Delta(B_0, \perp) = \{(\vec{b}, q_0), (\vec{c}, f)\} \\
 B_2 &:= \Delta(B_0, \perp^2) = \{(\vec{b}, q_0), (\vec{a}, f)\} \\
 B_3 &:= \Delta(B_0, \perp^2 b) = \{(\vec{b}, q_0), (\vec{b}, f)\}
 \end{aligned}$$


B_0 is the initial belief. The beliefs B_0 and B_1 are not confused: indeed, $\Delta(B_1, b) = \{(\vec{b}, q_0)\}$ is negatively deciding, and $\Delta(B_1, a) = \{(\vec{a}, f)\}$ is positively deciding. The belief B_2 is confused, as there is no $i \in \mathbb{N}$ for which $\Delta(B_2, b^i)$ is deciding. Finally, B_3 is very confused.

425 We will show that in the non-hidden case there always exists a diagnoser
 ([Lemma 23](#)). It follows that feasible policies need to decide almost surely
 and, by [Proposition 13](#), that c_{inf} is finite. We have seen in [Proposition 9](#) that

feasible policies do not allow confusion. In this section we construct policies that procrastinate so much that they avoid confusion just barely. We will see that such policies have an expected cost that comes arbitrarily close to c_{inf} .

Language Equivalence. We characterize confusion by language equivalence in a certain DFA. Consider the belief NFA \mathcal{B} . In the non-hidden case, if we disallow \perp -transitions then \mathcal{B} becomes a DFA \mathcal{B}' . For \mathcal{B}' we define a set of accepting states by $F_{\mathcal{B}'} := \{(s, q) \mid \Pr_s(L_q) = 1\}$.

Example 17. For the previous example, a part of the DFA \mathcal{B}' looks as follows:



States that are unreachable from (\vec{a}, q_0) are not drawn here.

We associate with each (s, q) the language $L_{s,q} \subseteq \Sigma^*$ that \mathcal{B}' accepts starting from initial state (s, q) . We call $(s, q), (s', q')$ *language equivalent*, denoted by $(s, q) \approx (s', q')$, when $L_{s,q} = L_{s',q'}$.

Lemma 18. One can compute the relation \approx in polynomial time.

PROOF. For any (s, q) one can use standard MC algorithms to check in polynomial time if $\Pr_s(L_q) = 1$ (using a graph search in the composition $\mathcal{M} \times \mathcal{A}$, as in the proof of Lemma 5.3). Language equivalence in the DFA \mathcal{B}' can be computed in polynomial time by minimization. \square

We call a belief $B \subseteq S \times Q$ *settled* when all $(s, q) \in B$ are language equivalent.

Lemma 19. A belief $B \subseteq S \times Q$ is confused if and only if there is a $a \in \Sigma$ such that $\Delta(B, a)$ is not settled.

PROOF. Since singleton beliefs cannot be confused in a non-hidden MC, we have:

$$\Pr_{\vec{a}}(\{uw \mid \Pr_{\vec{a}u}(L_{\delta(q,u)}) \in \{0, 1\}\}) = 1 \quad \forall a \forall q \quad (10)$$

Suppose $\Delta(B, a)$ is settled for all a . Then for all a and all $(\vec{a}, q_1), (\vec{a}, q_2) \in \Delta(B, a)$ and all u with $\vec{a} \xrightarrow{u}$ we have:

$$\Pr_{\vec{a}u}(L_{\delta(q_1,u)}) = 1 \iff \Pr_{\vec{a}u}(L_{\delta(q_2,u)}) = 1$$

By combining this with (10) we get:

$$\Pr_{\vec{a}}(\{uw \mid \Delta(B, au) \text{ is deciding}\}) = 1 \quad \forall a$$

Hence:

$$\Pr_s(\{u'w \mid \Delta(B, u') \text{ is deciding}\}) = 1 \quad \forall (s, q) \in B$$

450 Hence B is not confused.

Conversely, suppose there is a such that $\Delta(B, a)$ is not settled. Then there are $(\vec{a}, q_1), (\vec{a}, q_2) \in \Delta(B, a)$ and u_0 with $\vec{a} \xrightarrow{u_0}$ such that:

$$\Pr_{\vec{a}u_0} (L_{\delta(q_1, u_0)}) = 1 > \Pr_{\vec{a}u_0} (L_{\delta(q_2, u_0)})$$

It follows from (10) that there is u_1 with $\vec{a}u_0 \xrightarrow{u_1}$ such that:

$$\Pr_{\vec{a}u_0u_1} (L_{\delta(q_1, u_0u_1)}) = 1 > 0 = \Pr_{\vec{a}u_0u_1} (L_{\delta(q_2, u_0u_1)})$$

Hence for all u_2 with $\vec{a}u_0u_1 \xrightarrow{u_2}$ we have that $(\vec{a}u_0u_1u_2, \delta(q_1, u_0u_1u_2))$ is positively deciding and $(\vec{a}u_0u_1u_2, \delta(q_2, u_0u_1u_2))$ is negatively deciding, and therefore also that

$$\Delta(B, \vec{a}u_0u_1u_2) \supseteq \{(\vec{a}u_0u_1u_2, \delta(q_1, u_0u_1u_2)), (\vec{a}u_0u_1u_2, \delta(q_2, u_0u_1u_2))\}$$

is not deciding. So we have:

$$\Pr_{\vec{a}u_0u_1} (\{u_2w \mid \Delta(B, \vec{a}u_0u_1u_2) \text{ is deciding}\}) = 0$$

It follows that there is $(s, q) \in B$ with

$$\Pr_s (\{uw \mid \Delta(B, u) \text{ is deciding}\}) < 1$$

Hence B is confused. □

It follows that one can check in polynomial time whether a given belief is confused. We generalize this fact in [Lemma 22](#) below.

Example 20. In [Example 16](#) the belief B_3 is not settled. Indeed, from the DFA in [Example 17](#) we see that $L_{\vec{b}, q_0} = \emptyset \neq \{b\}^* = L_{\vec{b}, f}$. Since $B_3 = \Delta(B_2, b)$, by [Lemma 19](#), the belief B_2 is confused.

Procrastination. For a belief $B \subseteq S \times Q$ and $k \in \mathbb{N}$, if $\Delta(B, \perp^k)$ is confused then so is $\Delta(B, \perp^{k+1})$. We define:

$$\text{cras}(B) := \sup\{k \in \mathbb{N} \mid \Delta(B, \perp^k) \text{ is not confused}\} \in \mathbb{N} \cup \{-1, \infty\}$$

We set $\text{cras}(B) := -1$ if B is confused. We may write $\text{cras}(s, q)$ for $\text{cras}(\{(s, q)\})$.

Example 21. In [Example 16](#) we have $\text{cras}(B_0) = \text{cras}(\vec{a}, q_0) = 1$ and $\text{cras}(B_1) = 0$ and $\text{cras}(B_2) = \text{cras}(B_3) = -1$ and $\text{cras}(\vec{b}, q_0) = \text{cras}(\vec{a}, f) =$
460 ∞ .

Lemma 22. Given a belief B , one can compute $\text{cras}(B)$ in polynomial time. Further, if $\text{cras}(B)$ is finite then $\text{cras}(B) < |S|^2 \cdot |Q|^2$.

PROOF. Let $k \in \mathbb{N}$. By [Lemma 19](#), $\Delta(B, \perp^k)$ is confused if and only if:

$$\exists a. \exists (s, q), (t, r) \in \Delta(B, \perp^k) : s \xrightarrow{a}, t \xrightarrow{a}, (\vec{a}, \delta(q, a)) \not\approx (\vec{a}, \delta(r, a))$$

This holds if and only if there is a two-element subset of B , say $B_2 \subseteq B$ with $|B_2| \leq 2$, such that:

$$\exists a. \exists (s, q), (t, r) \in \Delta(B_2, \perp^k) : s \xrightarrow{a}, t \xrightarrow{a}, (\vec{a}, \delta(q, a)) \not\approx (\vec{a}, \delta(r, a))$$

Let G be the directed graph with nodes in $S \times Q \times S \times Q$ and edges

$$((s, q, t, r), (s', q', t', r')) \iff \Delta(\{(s, q), (t, r)\}, \perp) \supseteq \{(s', q'), (t', r')\}.$$

Also define the following set of nodes:

$$U := \{(s, q, t, r) \mid \exists a : s \xrightarrow{a}, t \xrightarrow{a}, (\vec{a}, \delta(q, a)) \not\approx (\vec{a}, \delta(r, a))\}$$

By [Lemma 18](#) one can compute U in polynomial time. By the argument above, it is enough to track two-element beliefs. Hence $\Delta(B, \perp^k)$ is confused if and only if there are $(s, q), (t, r) \in B$ such that there is a length- k path in G from (s, q, t, r) to a node in U . Let $k \leq |S \times Q \times S \times Q|$ be the length of the shortest such path, and set $k := \infty$ if no such path exists. Then k can be computed in polynomial time by a search of the graph G , and we have $cras(B) = k - 1$. \square

The Procrastination Policy. For any belief B and any observation prefix v , the language equivalence classes represented in $\Delta(B, v)$ depend only on v and the language equivalence classes in B . Therefore, when tracking beliefs along observations, we may restrict B to a single representative of each equivalence class. We denote this operation by $B \downarrow$. A belief B is settled if and only if $|B \downarrow| \leq 1$.

In the following we define a *procrastination policy* $\rho_{pro}(K)$, which is parameterized by (a large) $K \in \mathbb{N}$. This parameter indicates how many observations the monitor skips in the situation where $cras(s, q)$ is infinite but at least one further observation is needed. We will show in [Lemma 25](#) that in this case for each $\varepsilon > 0$ one can choose K large enough so that the expected cost is at most $1 + \varepsilon$. The proof of [Lemma 25](#) indicates how K depends on ε , on the DFA \mathcal{A} , and on the MC \mathcal{M} , especially its transition probabilities.

Define (and precompute) $k(s, q) := \min\{K, cras(s, q)\}$ for all (s, q) . We define $\rho_{pro}(K)$ by the following monitor that implements it:

1. $i := 0$
2. while (s_i, q_i) is not deciding:
 - (a) skip $k(s_i, q_i)$ observations, then observe a letter a_i
 - (b) $\{(s_{i+1}, q_{i+1})\} := \Delta((s_i, q_i), \perp^{k(s_i, q_i)} a_i) \downarrow$;
 - (c) $i := i + 1$;
3. output yes/no decision

It follows from the definition of $cras$ and [Lemma 19](#) that $\Delta((s_i, q_i), v_i) \downarrow$ is indeed a singleton for all i . We have:

Lemma 23. For all $K \in \mathbb{N}$ the procrastination policy $\rho_{pro}(K)$ is a diagnoser.


PROOF. For a non-hidden MC \mathcal{M} and a DFA \mathcal{A} , there is at most one successor for (s, q) on letter a in the belief NFA \mathcal{B} , for all s, q, a . Then, by Lemma 19, singleton beliefs are not confused, and in particular the initial belief B_0 is not confused. By Lemma 4.4, ε is not confused, which means that $\Pr(\{u \mid u \text{ deciding}\}) = \Pr(\varepsilon) = 1$. Since almost surely a deciding word u is produced and since $\Delta(B_0, u) \subseteq \Delta(B_0, v)$ whenever $u \sim v$, it follows that eventually an observation prefix v is produced such that $\Delta(B_0, v)$ contains a deciding pair (s, q) . But, as remarked above, $\Delta(B_0, v)$ is settled, so it is deciding. \square

The Procrastination MC $\mathcal{M}_{pro}(K)$. The policy $\rho_{pro}(K)$ produces a (random, almost surely finite) word $a_1 a_2 \cdots a_n$ with $n = C_{\rho_{pro}(K)}$. Indeed, the observations that $\rho_{pro}(K)$ makes can be described by an MC. Recall that we have previously defined a composition MC $\mathcal{M} \times \mathcal{A} = (S \times Q, \Sigma, M', (s_0, q_0))$. Now define an MC $\mathcal{M}_{pro}(K) := (S \times Q, \Sigma \cup \{\$, \}, M_{pro}(K), (s_0, q_0))$ where $\$ \notin \Sigma$ is a fresh letter and the transitions are as follows: when (s, q) is deciding then $M_{pro}(K)(\$)((s, q), (s, q)) := 1$, and when (s, q) is not deciding then

$$M_{pro}(K)(a)((s, q), (\vec{a}, q')) := \left(M'(\perp)^{k(s, q)} M'(a) \right) ((s, q), (\vec{a}, q')),$$

where the matrix $M'(\perp) := \sum_a M'(a)$ is powered by $k(s, q)$. The MC $\mathcal{M}_{pro}(K)$ may not be non-hidden, but could be made non-hidden by (i) collapsing all language equivalent $(s, q_1), (s, q_2)$ in the natural way, and (ii) redirecting all $\$$ -labelled transition to a new state $\vec{\$}$ that has a self-loop. In the understanding that $\$\$\$\dots$ indicates ‘decision made’, the probability distribution defined by the MC $\mathcal{M}_{pro}(K)$ coincides with the probability distribution on sequences of non- \perp observations made by $\rho_{pro}(K)$.

Example 24. For Example 16 the MC $\mathcal{M}_{pro}(K)$ for $K \geq 1$ is as follows:



Here the lower number in a state indicate the cras number. The left state is negatively deciding, and the right state is positively deciding. The policy $\rho_{pro}(K)$ skips the first observation and then observes either b or a , each with probability $\frac{1}{2}$, each leading to a deciding belief.

Maximal Procrastination is Optimal. The following lemma states, loosely speaking, that when a belief $\{(s, q)\}$ with $cras(s, q) = \infty$ is reached and K is large, then a single further observation is expected to suffice for a decision.

Lemma 25. Let $c(K, s, q)$ denote the expected cost of decision under $\rho_{pro}(K)$ starting in (s, q) . For each $\varepsilon > 0$ there exists $K \in \mathbb{N}$, depending on ε , \mathcal{M} and \mathcal{A} , such that for all (s, q) with $cras(s, q) = \infty$ we have $c(K, s, q) \leq 1 + \varepsilon$.

520 **PROOF (SKETCH).** The proof is a quantitative version of the proof of [Lemma 23](#). The singleton belief $\{(s, q)\}$ is not confused. Thus, if K is large then with high probability the belief $B := \Delta(\{(s, q)\}, \perp^K a)$ (for the observed next letter a) contains a deciding pair (s', q') . But if $cras(s, q) = \infty$ then, by [Lemma 19](#), B is settled, so if B contains a deciding pair then B is deciding. \square

PROOF (FULL). Let $(s, q) \in S \times Q$. The singleton belief $\{(s, q)\}$ cannot be confused in a non-hidden MC, i.e., we have:

$$1 = \Pr_s(\{uaw \mid (\vec{a}, \delta(q, ua)) \text{ is deciding}\})$$

It follows that for each $p < 1$ there is $K \in \mathbb{N}$ such that

$$\begin{aligned} p &\leq \Pr_s(\{uaw \mid (\vec{a}, \delta(q, ua)) \text{ is deciding, } |u| \leq K\}) \\ &= \Pr_s(\{uaw \mid (\vec{a}, \delta(q, ua)) \text{ is deciding, } |u| = K\}) \end{aligned}$$

Suppose $cras(s, q) = \infty$. Then for all $k \in \mathbb{N}$ the belief $\Delta(\{(s, q)\}, \perp^k)$ is not confused, and so, by [Lemma 19](#), for all a the belief $\Delta(\{(s, q)\}, \perp^k a)$ is settled. It follows that for each $p < 1$ there is $K \in \mathbb{N}$ such that

$$p \leq \Pr_s(\Sigma^K \{aw \mid \Delta(\{(s, q)\}, \perp^K a) \text{ is deciding}\}) \quad (11)$$


525 Since there are finitely many (s, q) , it is also the case that for each $p < 1$ there is $K \in \mathbb{N}$ such that for all (s, q) Equation (11) holds.

Let $\varepsilon > 0$. Choose $K \in \mathbb{N}$ such that (11) holds for $p := 1/(1 + \varepsilon)$. Let (s, q) be such that $cras(s, q) = \infty$ and let X denote the random number of observations that $\rho_{pro}(K)$ needs to make starting in (s, q) . From the argument above, we have $\Pr_s(X > 1) \leq 1 - p$. Since for the (random) pair (s', q') after the next observation (i.e., $\{(s', q')\} = \Delta(\{(s, q)\}, \perp^K a)$ for the observed a) we have $cras(s', q') = \infty$ again, it follows that $\Pr_s(X > i) \leq (1 - p)^i$ holds for all $i \in \mathbb{N}$. Hence we have:


$$c(K, s, q) := \text{Ex}_s(X) = \sum_{i=0}^{\infty} \Pr_s(X > i) \leq \sum_{i=0}^{\infty} (1 - p)^i = \frac{1}{p} = 1 + \varepsilon$$

\square

Example 26. Consider the following variant of the previous example:



530 The MC $\mathcal{M}_{pro}(K)$ for $K \geq 0$ is as follows:



The left state is negatively deciding, and the right state is positively deciding. We have $c(K, \vec{b}, q_0) = c(K, \vec{c}, f) = 0$ and $c(K, \vec{a}, q_0) = 1/(1 - (\frac{1}{3})^{K+1})$.

Now we can prove the main positive result of the paper:

Theorem 27. For any feasible policy ρ there is $K \in \mathbb{N}$ such that:

$$\text{Ex}(C_{\rho_{pro}(K)}) \leq \text{Ex}(C_{\rho})$$

PROOF (SKETCH). Let ρ be a feasible policy. We choose $K > |S|^2 \cdot |Q|^2$, so, by Lemma 22, $\rho_{pro}(K)$ coincides with $\rho_{pro}(\infty)$ (the latter is defined like $\rho_{pro}(K')$ but with K' replaced by ∞) until time, say, n_{∞} when $\rho_{pro}(K)$ encounters a pair (s, q) with $cras(s, q) = \infty$. (The time n_{∞} may, with positive probability, never come.) Let us compare $\rho_{pro}(K)$ with ρ up to time n_{∞} . For $n \in \{0, \dots, n_{\infty}\}$, define $v_{pro}(n)$ and $v_{\rho}(n)$ as the observation prefixes obtained by ρ_{pro} and ρ , respectively, after n steps. Write $\ell_{pro}(n)$ and $\ell_{\rho}(n)$ for the number of non- \perp observations in $v_{pro}(n)$ and $v_{\rho}(n)$, respectively. Note that $\ell_{pro}(n)$ and $\ell_{\rho}(n)$ are random variables. For beliefs B, B' we write $B \preceq B'$ when for all $(s, q) \in B$ there is $(s', q') \in B'$ with $(s, q) \approx (s', q')$. One can show by induction that we have for all $n \in \{0, \dots, n_{\infty}\}$:

$$\ell_{pro}(n) \leq \ell_{\rho}(n) \quad \text{and} \quad (\Delta(B_0, v_{pro}(n)) \preceq \Delta(B_0, v_{\rho}(n)) \quad \text{or} \quad \ell_{pro}(n) < \ell_{\rho}(n))$$

535 If time n_{∞} does not come then the inequality $\ell_{pro}(n) \leq \ell_{\rho}(n)$ from above suffices. Similarly, if at time n_{∞} the pair (s, q) is deciding, we are also done. If after time n_{∞} the procrastination policy $\rho_{pro}(K)$ observes at least one more letter then ρ also observes at least one more letter. By Lemma 25, one can choose K large so that for $\rho_{pro}(K)$ one additional observation probably suffices. If it is
540 the case that ρ almost surely observes only one letter after n_{∞} , then $\rho_{pro}(K)$ also needs only one more observation, since it has observed at time n_{∞} . \square

PROOF (FULL). Let ρ be a feasible policy. First we compare $\rho_{pro} := \rho_{pro}(\infty)$ with ρ . Let w be such that all its prefixes are enabled. For $n \in \mathbb{N}$, define $v_{pro}(n)(w)$ and $v_{\rho}(n)(w)$ as the observation prefixes obtained by ρ_{pro} and ρ , respectively, after n steps. (Thus, $v_{pro}(n)(w)$ and $v_{\rho}(n)(w)$ are, respectively, the length n prefixes of $\pi_{\rho_{pro}}(w)$ and $\pi_{\rho}(w)$.) We write $\ell_{pro}(n)(w)$ and $\ell_{\rho}(n)(w)$ for the number of non- \perp observations in $v_{pro}(n)(w)$ and $v_{\rho}(n)(w)$, respectively.

Define $B_{pro}(n)(w) := \Delta(B_0, v_{pro}(n)(w))$ and $B_\rho(n)(w) := \Delta(B_0, v_\rho(n)(w))$. For beliefs B, B' we write $B \preceq B'$ when for all $(s, q) \in B$ there is $(s', q') \in B'$ with $(s, q) \approx (s', q')$. In the following we suppress w in the notation to avoid clutter. We show for all w and all $n \in \mathbb{N}$:

$$\ell_{pro}(n) \leq \ell_\rho(n) \quad \text{and} \quad (B_{pro}(n) \preceq B_\rho(n) \quad \text{or} \quad \ell_{pro}(n) < \ell_\rho(n)) \quad (12)$$

We proceed by induction on n . In the base case, $n = 0$, we have $\ell_{pro}(0) = 0 = \ell_\rho(0)$ and $B_{pro}(0) = B_0 = B_\rho(0)$. For the inductive step, suppose (12) holds for n . (Note that, if $v_1 \sim u$ and $v_2 \sim u$ and $B_1 \preceq B_2$ and $|\Delta(B_1, v_1) \downarrow| \leq 1$, then $\Delta(B_1, v_1) \preceq \Delta(B_2, v_2)$.)

1. Suppose ρ_{pro} does not observe the $(n + 1)$ st letter. Then $\ell_{pro}(n + 1) = \ell_{pro}(n)$.
 - (a) Suppose $\ell_{pro}(n + 1) < \ell_\rho(n + 1)$. Then (12) holds for $n + 1$.
 - (b) Otherwise, $\ell_\rho(n + 1) \leq \ell_{pro}(n + 1) = \ell_{pro}(n) \leq \ell_\rho(n) \leq \ell_\rho(n + 1)$ by the induction hypothesis. So all these numbers are equal and we conclude that ρ does not observe the $(n + 1)$ st letter either and that $\ell_{pro}(n) = \ell_\rho(n)$. It follows by the induction hypothesis that $B_{pro}(n) \preceq B_\rho(n)$. Thus $B_{pro}(n + 1) = \Delta(B_{pro}(n), \perp) \preceq \Delta(B_\rho(n), \perp) = B_\rho(n + 1)$. Hence (12) holds for $n + 1$.
2. Otherwise ρ_{pro} observes the $(n + 1)$ st letter. From the definition of ρ_{pro} it follows that $B_{pro}(n + 1)$ is settled and thus $B_{pro}(n + 1) \preceq B_\rho(n + 1)$.
 - (a) Suppose $\ell_{pro}(n) < \ell_\rho(n)$. Then $\ell_{pro}(n + 1) \leq \ell_\rho(n) \leq \ell_\rho(n + 1)$, i.e., (12) holds for $n + 1$.
 - (b) Otherwise, by the induction hypothesis, we have $B_{pro}(n) \preceq B_\rho(n)$. Since ρ_{pro} observes the $(n + 1)$ st letter, the belief $\Delta(B_{pro}(n), \perp)$ is confused. Since $B_{pro}(n) \preceq B_\rho(n)$, the belief $\Delta(B_\rho(n), \perp)$ is also confused. Hence, by Lemma 4.4, $v_\rho(n) \perp$ is confused. By Proposition 9, ρ does not allow confusion. So ρ observes the $(n + 1)$ st letter. Thus we have $\ell_{pro}(n + 1) = \ell_{pro}(n) + 1 \leq \ell_\rho(n) + 1 = \ell_\rho(n + 1)$, where the inequality is by the induction hypothesis. Hence (12) holds for $n + 1$.

Thus we have shown (12).

Define $I := \{(s, q) \mid (s, q) \text{ is not deciding, } cras(s, q) = \infty\}$. We will choose $K > |S|^2 \cdot |Q|^2$, so, by Lemma 22, $\rho_{pro}(K)$ coincides with $\rho_{pro}(\infty)$ until possibly $\rho_{pro}(K)$ encounters a pair $(s, q) \in I$. Define:

$$D := \{w \mid \rho_{pro}(\infty) \text{ does not encounter any element of } I \text{ on } w\}$$

For each $(s, q) \in I$ define:

$$E(s, q) := \{w \mid (s, q) \text{ is the first element of } I \text{ that } \rho_{pro}(\infty) \text{ encounters on } w\}$$

So $D \cup \bigcup_{(s, q) \in I} E(s, q) = \Sigma^\omega$. This partition is convenient, as the policy $\rho_{pro}(K)$, and hence its expected cost, depends on whether I is encountered. We show $\text{Ex}(C_{\rho_{pro}(K)}) \leq \text{Ex}(C_\rho)$ by conditioning separately on D and on each $E(s, q)$ that has positive probability.

There is, by [Lemma 23](#), almost surely a time, n_{dec} , when $\rho_{pro}(K)$ encounters a deciding (s, q) . In the event D , it follows $C_{\rho_{pro}(K)} = \ell_{pro}(n_{dec})$; i.e., we have almost surely

$$C_{\rho_{pro}(K)} = \ell_{pro}(n_{dec}) \stackrel{\text{by (12)}}{\leq} \ell_{\rho}(n_{dec}) \leq C_{\rho} \quad \text{in the event } D. \quad (13)$$

Let $(s, q) \in I$ and consider the event $E(s, q)$. Let us write n for the time when $\rho_{pro}(\infty)$, and thus $\rho_{pro}(K)$, first encounters (s, q) . Since (s, q) is not deciding, the belief $B_{pro}(n)$ is not deciding. Since $\rho_{pro}(K)$ observes a letter at time n , the belief $B_{\rho}(n) \succeq B_{pro}(n)$ is not deciding either, and ρ needs to make at least one more observation. So we have almost surely

$$C_{\rho} - \ell_{pro}(n) \geq \ell_{\rho}(n) + 1 - \ell_{pro}(n) \stackrel{\text{by (12)}}{\geq} 1 \quad \text{in the event } E(s, q). \quad (14)$$

1. Suppose $\text{Ex}(C_{\rho} - \ell_{pro}(n) \mid E(s, q)) \leq 1$. Then, by (14), we have $C_{\rho} = \ell_{\rho}(n) + 1$ almost surely in the event $E(s, q)$. Thus, almost surely in $E(s, q)$, there exists $k \in \mathbb{N}$ such that $\Delta(B_{\rho}, \perp^k a) \succeq \Delta((s, q), \perp^k a)$ is deciding for all a . It follows that there exists $k \in \mathbb{N}$ such that $\Delta((s, q), \perp^k a)$ is deciding for all a . We will choose $K \geq k$, so $\Delta((s, q), \perp^K a)$ is deciding for all a . Hence, by (12), we have $C_{\rho_{pro}(K)} = \ell_{pro}(n) + 1 \leq \ell_{\rho}(n) + 1 = C_{\rho}$ almost surely in the event $E(s, q)$.
2. Otherwise $\text{Ex}(C_{\rho} - \ell_{pro}(n) \mid E(s, q)) > 1$. It follows from [Lemma 25](#) that one can choose K large enough so that $\text{Ex}(C_{\rho_{pro}(K)} - \ell_{pro}(n) \mid E(s, q)) < \text{Ex}(C_{\rho} - \ell_{pro}(n) \mid E(s, q))$ (choose $\varepsilon := (\text{Ex}(C_{\rho} - \ell_{pro}(n) \mid E(s, q)) - 1)/2$). We will choose K in this way, so that we have $\text{Ex}(C_{\rho_{pro}(K)} \mid E(s, q)) < \text{Ex}(C_{\rho} \mid E(s, q))$.

By choosing K large enough so that for each $(s, q) \in I$ the respective constraint in item 1 or 2 is satisfied, we obtain $\text{Ex}(C_{\rho_{pro}(K)} \mid E(s, q)) \leq \text{Ex}(C_{\rho} \mid E(s, q))$ for each $(s, q) \in I$. Combining this with (13) yields the result. \square

It follows that, in order to compute c_{inf} , it suffices to analyze $\text{Ex}(C_{\rho_{pro}(K)})$ for large K . This leads to the following theorem:

Theorem 28. *Given a non-hidden MC \mathcal{M} and a DFA \mathcal{A} , one can compute c_{inf} in polynomial time.*

PROOF. For each (s, q) define $c(K, s, q)$ as in [Lemma 25](#), and define $c(s, q) := \lim_{K \rightarrow \infty} c(K, s, q)$. By [Lemma 25](#), for each non-deciding (s, q) with $cras(s, q) = \infty$ we have $c(s, q) = 1$. Hence the $c(s, q)$ satisfy the following system of linear equations where some coefficients come from the procrastination MC $\mathcal{M}_{pro}(\infty)$ (the latter being defined like $\mathcal{M}_{pro}(K)$ but with $K = \infty$):

$$c(s, q) = \begin{cases} 0 & \text{if } (s, q) \text{ is deciding} \\ 1 & \text{if } (s, q) \text{ is not deciding and } cras(s, q) = \infty \\ 1 + c'(s, q) & \text{otherwise} \end{cases}$$

$$c'(s, q) = \sum_a \sum_{q'} M_{pro}(\infty)((s, q), (\vec{a}, q')) \cdot c(\vec{a}, q') \quad \text{if } cras(s, q) < \infty$$

By solving the system one can compute $c(s_0, q_0)$ in polynomial time. We have:

$$c_{inf} = \inf_{\text{feasible } \rho} \text{Ex}(C_\rho) \stackrel{\text{Theorem 27}}{=} \lim_{K \rightarrow \infty} \text{Ex}(C_{\rho_{pro}(K)}) = c(s_0, q_0)$$

590 Hence one can compute c_{inf} in polynomial time. □

6. Empirical Evaluation of the Expected Optimal Cost

We have shown that maximal procrastination is optimal in the non-hidden case (Theorem 27). However, we have not shown *how much* better the optimal policy is than the see-all baseline. It appears difficult to answer this question analytically, so we address it empirically. We implemented our algorithms in a fork of the Facebook Infer static analyzer [13], and applied them to 10 open-source projects, totaling 120 thousand Java methods. We found that in > 95% of cases the maximally procrastinating monitor is trivial and thus the optimal cost is 0, because Infer decides statically whether the property is violated. In the remaining cases, we found that the optimal cost is roughly half of the see-all cost, but the variance is high.

Design. Our setting requires a DFA and an MC representing, respectively, a program property and a program. For this empirical estimation of the expected optimal cost, the DFA is fixed, the MC shape is the symbolic flowgraph of a real program, and the MC probabilities are sampled from Dirichlet distributions.

The DFA represents the following property: ‘there are no two calls to *next* without an intervening call to *hasNext*’. To understand how the MC shape is extracted from programs, some background is needed. Infer [13, 14] is a static analyzer that, for each method, infers several preconditions and, attached to each precondition, a symbolic path. For a simple example, consider a method whose body is ‘if (*b*) *x.next*(); if (!*b*) *x.next*()’. Infer would generate two preconditions for it, *b* and $\neg b$. In each of the two attached symbolic paths, we can see that *next* is not called twice, which we would not notice with a control flowgraph. The symbolic paths are inter-procedural. If a method *f* calls a method *g*, then the path of *f* will link to a path of *g* and, moreover, it will pick one of the paths of *g* that corresponds to what is currently known at the call site. For example, if *g(b)* is called from a state in which $\neg b$ holds, then Infer will select a path of *g* compatible with the condition $\neg b$.

The symbolic paths are finite because abstraction is applied, including across mutually recursive calls. But, still, multiple vertices of the symbolic path correspond to the same vertex of the control flowgraph. For example, Infer may go around a for-loop five times before noticing the invariant. By coalescing those vertices of the symbolic path that correspond to the same vertex of the control flowgraph we obtain an *SFG* (*symbolic flowgraph*). We use such SFGs as the skeleton of MCs. Intuitively, one can think of SFGs as inter-procedural control flowgraphs restricted based on semantic information. Vertices correspond to locations in the program text, and transitions correspond to method calls

or returns; that is labels are of the form `call m` or `return m`, where m is a method name. We do not track other events, such as assignments. Transition probabilities should be interpreted as a form of static branch prediction. One could learn these probabilities by observing many runs of the program on typical input data, for example by using the Baum–Welch algorithm [15]. Instead, we opt to show that the improvement in expected observation cost is robust over a wide range of possible transition probabilities, which we do by drawing several samples from Dirichlet distributions. Besides, recall that the (optimal) procrastination policy does not depend on transition probabilities.

Once we have a DFA and an MC we compute their product. In some cases, it is clear that the product is empty or universal. These are the cases in which we can give the verdict right away, because no observation is necessary. We then focus on the non-trivial cases.

For non-trivial $\text{MC} \times \text{DFA}$ products, we compute the expected cost of the light see-all policy $\text{Ex}(C_\circ)$, which observes all letters until a decision is made and then stops. We can do so by using standard algorithms [16, Chapter 10.5]. Then, we compute \mathcal{M}_{pro} , which we use to compute the expected observation cost c_{inf} of the procrastination policy (Theorem 28). Recall that in order to compute \mathcal{M}_{pro} , one needs to compute the *cras* function, and also to find language equivalence classes. Thus, computing \mathcal{M}_{pro} entails computing all the information necessary for implementing a procrastinating monitor.

We note that Infer produces one symbolic path for each method precondition. This means that in our experiments we have one, fixed DFA and many MCs. When we look at the MC of some precondition P for some method m , we are checking whether the property described by the DFA is violated under the assumption that we start the DFA in its initial state and we simultaneously start method m in a program state satisfying P .

Figure 2 gives a small example of how MCs are extracted from Java code. In Figure 2a, method `g` creates an iterator `i`, and keeps extracting data from it as long as `f(i)` returns true. The fact that `f(i)` sometimes skips the call to `i.hasNext()` is reflected in the MC structure of Figure 2b. In this example, we have a single precondition for `g`, requiring that `xs` is non-null.

Methodology. We selected 10 Java projects among those that are most forked on GitHub (Table 1). We ran Infer on each of these projects. From the inferred specifications, we built SFGs and monitors that employ light see-all policies and maximal procrastination policies. From these monitors, we computed the respective expected costs, solving the linear systems using Gurobi [17]. Our implementation is in a fork of Infer, on GitHub³.

Results. The results are given in Table 2. We first note that the number of monitors is much smaller than the number of methods, by a factor of 10 or 100.

³<https://github.com/rgrig/infer/tree/selmon-journal/>

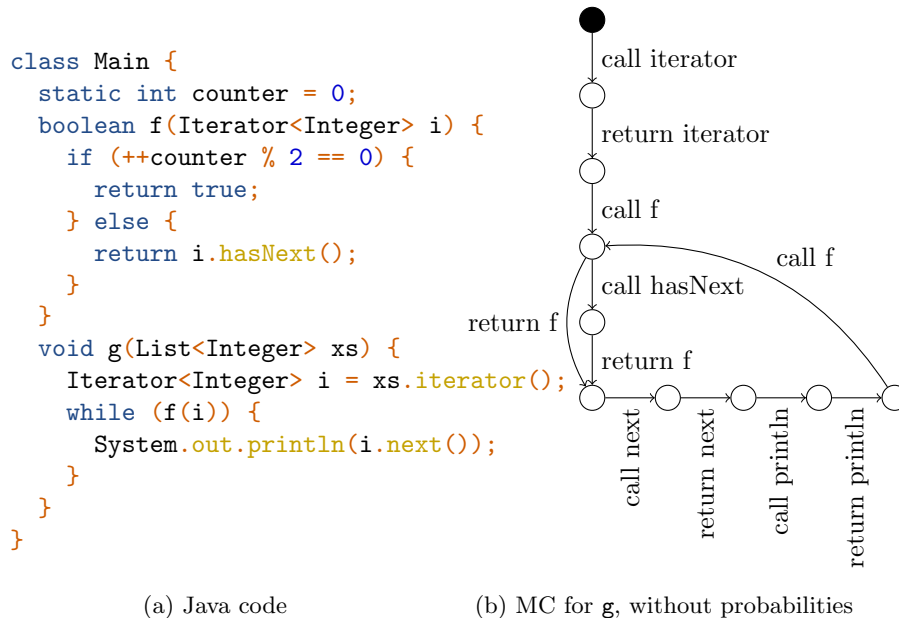


Figure 2: Example of extracting an MC from Java code.

This is because in most cases we are able to determine the answer statically, by analyzing the symbolic paths produced by Infer. The large factor should not be too surprising: we are considering a fixed property about iterators, not all Java methods use iterators, and, when they do, it is usually easy to tell that they do so correctly. Still, each project has some monitors, which handle the cases that are not so obvious.

We note that $\frac{c_{inf}}{\text{Ex}(C_c)} \approx 0.5$. The table supports this by presenting the median and the geometric average, which are close to each-other; the arithmetic average is also close. There is, however, quite a bit of variation from monitor to monitor, as shown in Figure 3. We conclude that selective monitoring has the potential to substantially reduce the overhead of runtime monitoring.

7. Future Work

In this paper we required policies to be feasible, which means that our selective monitors are as precise as non-selective monitors. One may relax this and study the tradeoff between efficiency (skipping even more observations) and precision (probability of making a decision). Further, one could replace the diagnosability notion of this paper by other notions from the literature; one could investigate how to compute c_{inf} for other classes of MCs, such as acyclic MCs; one could study the sensitivity of c_{inf} to changes in transition probabilities; and one could identify classes of MCs for which selective monitoring helps and classes of MCs for which selective monitoring does not help.

Table 1: Open Source projects, from GitHub.

Owner	Project	Version	Description
Alibaba	dubbo	048fadf	a high-performance RPC framework
Alibaba	fastjson	1196251	a fast JSON parser/generator
Apache	tomcat	22ad695	Java Servlet, Server Pages, and WebSocket
Clojure	clojure	653b846	the Clojure programming language
Google	guava	03f4322	Google core libraries for Java
ReactiveX	RxJava	042dee3	Reactive Extensions for the JVM
redisson	redisson	824bed3	Java Redis client
seata	seata	ac6521f	distributed transactions
Skylot	jadx	fa0a38d	Dex to Java decompiler
TeamNewPipe	NewPipe	07cead7	a libre streaming front-end for Android

Table 2: Reduction in expected observation cost, on real-world data. Each SFG (symbolic flowgraph) corresponds to one inferred precondition of a method. The size of monitors is measured in number of language equivalence classes. (LOC = lines of code; GAvg = geometric average.)

Name	Project Size			Monitors			$c_{inf}/Ex(C_o)$	
	Methods	SFGs	LOC	Count	Avg-Size	Max-Size	Med	GAvg
dubbo	12K	23K	153K	221	217	1149	0.42	0.42
fastjson	2K	6K	54K	136	96	899	0.59	0.54
tomcat	27K	61K	471K	555	106	711	0.44	0.45
clojure	5K	15K	60K	195	84	583	0.55	0.53
guava	21K	43K	1115K	1005	143	979	0.36	0.36
RxJava	12K	45K	186K	80	33	250	0.59	0.56
redisson	13K	25K	191K	515	231	1326	0.45	0.42
seata	12K	21K	124K	214	129	450	0.43	0.42
jadx	8K	15K	61K	349	105	1015	0.54	0.48
NewPipe	6K	13K	76K	13	110	256	0.44	0.43

A nontrivial extension to the formal model would be to include some notion
 of data, which is pervasive in practical specification languages used in runtime
 verification [18]. This would entail replacing the DFA with a more expressive
 device, such as a nominal automaton [19], a symbolic automaton [20], or a logic
 with data (e.g., [21]). Alternatively, one could side-step the problem by using
 the slicing idea [22], which separates the concern of handling data at the expense
 of a mild loss of expressive power. Finally, the monitors we computed could be
 used in a runtime verifier, or even in session type monitoring where the setting
 is similar [23].

[1] M. Arnold, M. T. Vechev, E. Yahav, QVM: an efficient runtime for detecting
 defects in deployed systems, ACM Trans. Softw. Eng. Methodol. 21 (1)
 (2011) 2:1–2:35.

[2] E. Bartocci, R. Grosu, A. Karmarkar, S. A. Smolka, S. D. Stoller, E. Zadok,




Figure 3: Empirical distribution of $c_{inf}/\text{Ex}(C_o)$, across all projects.

J. Seyster, Adaptive runtime verification, in: Runtime Verification, Third International Conference, RV 2012, Vol. 7687 of Lecture Notes in Computer Science, Springer, 2013, pp. 168–182.

- 705 [3] S. Stoller, E. Bartocci, J. Seyster, R. Grosu, K. Havelund, S. Smolka, E. Zadok, Runtime verification with state estimation, in: Runtime Verification - Second International Conference, RV 2011, Vol. 7186 of Lecture Notes in Computer Science, Springer, 2012, pp. 193–207.
- 710 [4] K. Kalajdzic, E. Bartocci, S. A. Smolka, S. D. Stoller, R. Grosu, Runtime verification with particle filtering, in: Runtime Verification - 4th International Conference, RV 2013. Proceedings, Vol. 8174 of Lecture Notes in Computer Science, Springer, 2013, pp. 149–166.
- 715 [5] M. Sampath, R. Sengupta, S. Lafortune, K. Sinnamohideen, D. Teneketzis, Diagnosability of discrete-event systems, IEEE Transactions on Automatic Control 40 (9) (1995) 1555–1575.
- [6] S. Jiang, Z. Huang, V. Chandra, R. Kumar, A polynomial algorithm for testing diagnosability of discrete-event systems, IEEE Transactions on Automatic Control 46 (8) (2001) 1318–1321.
- 720 [7] D. Thorsley, D. Teneketzis, Diagnosability of stochastic discrete-event systems, IEEE Transactions on Automatic Control 50 (4) (2005) 476–492.
- [8] A. P. Sistla, M. Zefran, Y. Feng, Monitorability of stochastic dynamical systems, in: Computer Aided Verification - 23rd International Conference, CAV 2011. Proceedings, Vol. 6806 of Lecture Notes in Computer Science, Springer, 2011, pp. 720–736.
- 725 [9] N. Bertrand, S. Haddad, E. Lefauchaux, Foundation of diagnosis and predictability in probabilistic systems, in: 34th International Conference on

Foundation of Software Technology and Theoretical Computer Science, FSTTCS 2014, Vol. 29 of LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014, pp. 417–429.

- 730 [10] N. Bertrand, S. Haddad, E. Lefaucheu, Accurate approximate diagnosability of stochastic systems, in: Language and Automata Theory and Applications - 10th International Conference, LATA 2016, Proceedings, Vol. 9618 of Lecture Notes in Computer Science, Springer, 2016, pp. 549–561.
- [11] J.-Y. Kao, N. Rampersad, J. Shallit, On NFAs where all states are final, initial, or both, Theoretical Computer Science 410 (47) (2009) 5010–5021.
- 735 [12] A. Paz, Introduction to Probabilistic Automata, Academic Press, 1971.
- [13] C. Calcagno, D. Distefano, J. Dubreil, D. Gabi, P. Hooimeijer, M. Luca, P. W. O’Hearn, I. Papakonstantinou, J. Purbrick, D. Rodriguez, Moving fast with software verification, in: NASA Formal Methods - 7th International Symposium, NFM 2015, Proceedings, Vol. 9058 of Lecture Notes in Computer Science, Springer, 2015, pp. 3–11.
- 740 [14] C. Calcagno, D. Distefano, P. W. O’Hearn, H. Yang, Compositional shape analysis by means of bi-abduction, J. ACM 58 (6) (2011) 26:1–26:66.
- [15] B. G. Leroux, Maximum-likelihood estimation for hidden Markov models, Stochastic Processes and their Applications 40 (1) (1992) 127–143.
- 745 [16] C. Baier, J.-P. Katoen, Principles of model checking, MIT Press, 2008.
- [17] Gurobi Optimization, Inc., Gurobi optimizer reference manual, <http://www.gurobi.com> (2017).
- [18] K. Havelund, M. Leucker, G. Reger, V. Stolz, A Shared Challenge in Behavioural Specification (Dagstuhl Seminar 17462), Dagstuhl Reports [doi: 10.4230/DagRep.7.11.59](https://doi.org/10.4230/DagRep.7.11.59).
- 750 [19] M. Bojańczyk, B. Klin, S. Lasota, Automata theory in nominal sets, Logical Methods in Computer Science 10 (3).
- [20] L. D’Antoni, M. Veanes, The power of symbolic automata and transducers, in: Computer Aided Verification - 29th International Conference, CAV 2017, Proceedings, Part I, Vol. 10426 of Lecture Notes in Computer Science, Springer, 2017, pp. 47–67.
- 755 [21] S. Demri, R. Lazic, LTL with the freeze quantifier and register automata, in: 21th IEEE Symposium on Logic in Computer Science (LICS) 2006), Proceedings, IEEE Computer Society, 2006, pp. 17–26.
- 760 [22] G. Rosu, F. Chen, Semantics and algorithms for parametric monitoring, Logical Methods in Computer Science 8 (1) (2012) .
- [23] L. Bocchi, T.-C. Chen, R. Demangeon, K. Honda, N. Yoshida, Monitoring networks through multiparty session types, Theor. Comput. Sci. 669 (2017) 33–58.
- 765

Appendix A. Missing Proofs

Appendix A.1. Equivalence of Definitions for Observation Prefixes

Confused. Recall the two definitions for v confused:

- (a) $\Pr(\{vu \mid vu \text{ deciding}\}) < \Pr(v)$
 770 (b) $\Pr(\{uw \mid u \sim v, \text{ no prefix of } vw \text{ is deciding}\}) > 0$

If $A \subseteq B$, then $\Pr(A) < \Pr(B)$ is equivalent to $\Pr(B \setminus A) > 0$. The plan is to pick A and B such that definition (a) has the form $\Pr(A) < \Pr(B)$, and then show that $A \subseteq B$ and that definition (b) has the form $\Pr(B \setminus A) > 0$.

$$\begin{aligned} & \Pr(\{vu \mid vu \text{ deciding}\}) \\ = & \Pr\left(\bigcup_{\substack{vu \\ \text{deciding}}} \{u'w \mid u' \sim vu\}\right) && \text{the } \Pr(\Upsilon) \text{ notation} \\ = & \Pr(\{u'w \mid u' \sim vu \text{ and } vu \text{ deciding}\}) \\ = & \Pr(\{uw \mid u \sim v \text{ and } vw \text{ has a deciding prefix}\}) \end{aligned}$$

Thus, we set

$$\begin{aligned} A & := \{uw \mid u \sim v \text{ and } vw \text{ has a deciding prefix}\} \\ B & := \{uw \mid u \sim v\} \end{aligned}$$

It is clear that $A \subseteq B$, and that $B \setminus A$ is the event from definition (b).

Very Confused. Recall the two definitions for v very confused:

- (a) $\Pr(\{vu \mid vu \text{ deciding}\}) = 0$
 (b) vu' is non-deciding for all enabled vu'

To see that (a) implies (b), assume that vu' is deciding. Because

$$0 \stackrel{(a)}{=} \Pr(\{vu \mid vu \text{ deciding}\}) \geq \Pr(vu')$$

775 it follows that vu' is not enabled.

To see that (b) implies (a), recall that observation prefixes are finite, and calculate:

$$\begin{aligned} & \Pr(\{vu \mid vu \text{ deciding}\}) \\ = & \Pr\left(\bigcup_{\substack{vu \\ \text{deciding}}} \{u'w \mid u' \sim vu\}\right) && \text{unfold the } \Pr(\Upsilon) \text{ notation} \\ \leq & \sum_{\substack{vu \\ \text{deciding}}} \Pr(\{u'w \mid u' \sim vu\}) && \text{union bound} \\ = & \sum_{\substack{vu \\ \text{deciding}}} \Pr(vu) && \text{fold the } \Pr(vu) \text{ notation} \\ = & 0 && \text{as } \Sigma_{\perp}^* \text{ is countable, and we assume (b)} \end{aligned}$$

Finitary. Recall the two definitions for v finitary:

$$(a) \Pr(\{vu \mid vu \text{ very confused or deciding}\}) = \Pr(v)$$

$$(b) \Pr(\{uw \mid u \sim v, \text{ no prefix of } vw \text{ is deciding or very confused}\}) = 0$$

The proof is similar to the confused case. Definition (a) has the form $\Pr(A) = \Pr(B)$ with $A \subseteq B$, and definition (b) has the form $\Pr(B \setminus A) = 0$. We just need to make the sets A and B explicit. Below, ‘dv’ stands for ‘deciding or very confused’.

$$\begin{aligned} & \Pr(\{vu \mid vu \text{ deciding or very confused}\}) \\ = & \Pr\left(\bigcup_{\substack{vu \\ vu \text{ dv}}} \{u'w \mid u' \sim vu\}\right) && \text{unfold the } \Pr(\Upsilon) \text{ notation} \\ = & \Pr(\{u'w \mid u' \sim vu \text{ and } vu \text{ is dv}\}) \\ = & \Pr(\{uw \mid u \sim v \text{ and } vw \text{ has a dv prefix}\}) \end{aligned}$$

So, we pick

$$\begin{aligned} A & := \{uw \mid u \sim v \text{ and } vw \text{ has a dv prefix}\} \\ B & := \{uw \mid u \sim v\} \end{aligned}$$

from which $A \subseteq B$ and $B \setminus A = \{uw \mid u \sim v \text{ and } uw \text{ has no dv prefix}\}$, as
780 desired.

Appendix A.2. Proof of Lemma 5

Here is Lemma 5 from the main body:

Lemma 5. *Let v be an observation prefix, and B a belief.*

1. *Whether v is enabled can be decided in P.*
- 785 2. *Whether v (or B) is negatively deciding can be decided in P.*
3. *Whether v (or B) is positively deciding can be decided in P.*
4. *Whether v (or B) is confused can be decided in PSPACE.*
5. *Whether v (or B) is very confused can be decided in PSPACE.*
6. *Whether v (or B) is finitary can be decided in PSPACE.*

790 **PROOF.** Let G be the following graph:

- the set of vertices is $S \times Q \times P$ where P is the set of prefixes of v ;
- there is an edge $(s_1, q_1, v_1) \rightarrow (s_2, q_2, v_2)$ if and only if $v_1 = ov_2$ for some o and $\Delta((s_1, q_1), o) \ni (s_2, q_2)$.

This graph G can be computed in deterministic logspace. Also note that the
795 belief NFA can be computed in deterministic logspace.

Proof of item 1. By Lemma 4.1, we have that v is enabled if and only if $\Delta(B_0, v) \neq \emptyset$. We have $\Delta(B_0, v) \neq \emptyset$ if and only if G has a path from (s_0, q_0, v) to a node (s, q, ε) . But graph reachability is in NL. This proves item 1.

800 *Proof of items 2 and 3.* Assume that, for a given pair (s_1, q_1) , one can determine in NL whether (s_1, q_1) is negatively deciding.

- Then we can check in NL whether the belief B is *not* negatively deciding: guess $(s_1, q_1) \in B$ nondeterministically and check whether (s_1, q_1) is *not* negatively deciding. The latter can be done in NL, as NL is closed under complement. It follows that one can check in NL whether B is negatively deciding. 805
- Under the same assumption we can check in NL whether the observation prefix v is negatively deciding: By Lemma 4.2 we have that v is not negatively deciding if and only if $\Delta(B_0, v)$ is not negatively deciding. The latter can be checked in NL by following nondeterministically a path in G from (s_0, q_0, v) to a node (s_1, q_1, ε) and then checking if (s_1, q_1) is not negatively deciding. It follows that one can check in NL whether v is negatively deciding. 810

The same reasoning applies when “negatively deciding” is replaced with “positively deciding”. Therefore, for items 2 and 3, it remains to show that one can determine in NL whether a given pair (s_1, q_1) is negatively (positively, respectively) deciding. We have:

$$\begin{aligned}
& (s_1, q_1) \text{ is not negatively deciding} \\
\iff & \Pr_{s_1}(L_{q_1}) > 0 \\
\iff & \exists u. \exists s_2 : M(u)(s_1, s_2) > 0 \text{ and } \delta(q_1, u) = f \\
\iff & \exists u. \exists s_2 : \Delta(\{(s_1, q_1)\}, u) \ni (s_2, f)
\end{aligned}$$

815 The latter can be checked in NL by nondeterministically guessing a word u letter-by-letter and by nondeterministically following, in the belief NFA, a path that is labelled by u and leads from (s_1, q_1) to a node (s_2, f) . This proves item 2.

For item 3 it remains to show how to determine in NL whether a given pair (s_1, q_1) is positively deciding. We have:

$$\begin{aligned}
& (s_1, q_1) \text{ is not positively deciding} \\
\iff & \Pr_{s_1}(L_{q_1}) < 1 \\
\iff & \exists u. \exists s_2. \exists q_2 : M(u)(s_1, s_2) > 0, \delta(q_1, u) = q_2, \Pr_{s_2}(L_{q_2}) = 0 \\
\iff & \exists u. \exists (s_2, q_2) \in \Delta(\{(s_1, q_1)\}, u) : (s_2, q_2) \text{ is negatively deciding} \\
\iff & \exists u. \exists (s_2, q_2) \in \Delta(\{(s_1, q_1)\}, u) : (s_2, q_2) \text{ is negatively deciding,} \\
& |u| \leq |S \times Q|
\end{aligned}$$

The last equivalence follows from the fact that there are at most $|S \times Q|$ different pairs of the form (s, q) . Therefore, one can check in NL whether (s_1, q_1) is not positively deciding:

1. nondeterministically guess a word u letter-by-letter;
- 820 2. nondeterministically follow, in the belief NFA, a path that is labelled by u and leads from (s_1, q_1) to a node (s_2, q_2) ;

3. check that (s_2, q_2) is negatively deciding; we have shown previously that this can be done in NL.

This proves item 3.

Proof of item 4. By Lemma 4.4 it suffices to prove membership in PSPACE for the case where a belief B is given. For all $B \subseteq S \times Q$, define:

$$V_B := \{uw \mid \Delta(B, u) \text{ is deciding}\}$$

825 We first prove the following claim:

Claim. For given s and $B \subseteq S \times Q$ one can determine in PSPACE whether $\Pr_s(V_B) = 0$.

We prove the claim. We have:

$$\begin{aligned} & \Pr_s(V_B) > 0 \\ \iff & \exists u. \exists s' : M(u)(s, s') > 0, \Delta(B, u) \text{ is deciding} \\ \iff & \exists u. \exists s' : M(u)(s, s') > 0, \Delta(B, u) \text{ is deciding, } |u| \leq 2^{|S \times Q|} \end{aligned}$$

The last equivalence follows from the fact that there are at most $2^{|S \times Q|}$ different beliefs of the form $\Delta(B, u)$. Therefore, one can check in NPSpace = PSPACE
830 whether $\Pr_s(V_B) > 0$:

1. guess a word $u = a_1 \cdots a_n$ with $n \leq 2^{|S \times Q|}$ letter-by-letter and check that there are states s_1, \dots, s_n with $M(a_1)(s, s_1) > 0$ and $M(a_{i+1})(s_i, s_{i+1}) > 0$ for all $i \in \{1, \dots, n-1\}$;
2. compute, on the fly, the belief

$$B' := \Delta(B, u) = \Delta(\cdots \Delta(\Delta(B, a_1), a_2) \cdots a_n);$$

- 835 3. check that B' is deciding; we have shown in items 2 and 3 that this can be done in polynomial time.

Since PSPACE is closed under complement, we have proved the claim.

Towards item 4, let B be any belief. We have:

$$\begin{aligned} & B \text{ is confused} \\ \iff & \exists (s, q) \in B : \Pr_s(V_B) < 1 \\ \iff & \exists (s, q) \in B. \exists u. \exists s' : M(u)(s, s') > 0, \Pr_{s'}(V_{\Delta(B, u)}) = 0 \\ \iff & \exists (s, q) \in B. \exists u. \exists s' : M(u)(s, s') > 0, \Pr_{s'}(V_{\Delta(B, u)}) = 0, \\ & |u| \leq |S| \cdot 2^{|S \times Q|} \end{aligned}$$

The last equivalence follows from the fact that there are at most $|S| \cdot 2^{|S \times Q|}$ different pairs of the form $(s', \Delta(B, u))$. Therefore, one can check in NPSpace = PSPACE whether B is confused:

- 840 1. guess a word $u = a_1 \cdots a_n$ with $n \leq |S| \cdot 2^{|S \times Q|}$ letter-by-letter and check that there are states s_1, \dots, s_n with $M(a_1)(s, s_1) > 0$ and $M(a_{i+1})(s_i, s_{i+1}) > 0$ for all $i \in \{1, \dots, n-1\}$;
2. compute, on the fly, the belief

$$B' := \Delta(B, u) = \Delta(\cdots \Delta(\Delta(B, a_1), a_2) \cdots a_n);$$

3. check that $\Pr_{s_n}(V_{B'}) = 0$; we have shown in the claim above that this can be done in PSPACE.

Proof of item 5. By Lemma 4.5 it suffices to prove membership in PSPACE for the case where a belief B is given. We have:

$$\begin{aligned} & B \text{ is not very confused} \\ \iff & \exists u : \Delta(B, u) \neq \emptyset \text{ is deciding} && \text{definition} \\ \iff & \exists u : \Delta(B, u) \neq \emptyset \text{ is deciding, } |u| \leq 2^{|S \times Q|} \end{aligned}$$

845 The last equivalence follows from the fact that there are at most $2^{|S \times Q|}$ different beliefs of the form $\Delta(B, u)$. Therefore, one can check in NPSpace = PSPACE whether B is not very confused:

1. guess a word u with $|u| \leq 2^{|S \times Q|}$ letter-by-letter and compute, on the fly, the belief $B' := \Delta(B, u)$;
- 850 2. check that $B' \neq \emptyset$ and B' is deciding; we have shown in items 2 and 3 that this can be done in polynomial time.

Since PSPACE is closed under complement, we have proved item 5.

Proof of item 6. By Lemma 4.6 it suffices to prove membership in PSPACE for the case where a belief B is given. The proof is analogous to the one of item 4. We abbreviate “deciding or very confused” to “dv”. For all $B \subseteq S \times Q$, define:

$$W_B := \{uw \mid \Delta(B, u) \text{ is dv}\}$$

We first prove the following claim:

Claim. For given s and $B \subseteq S \times Q$ one can determine in PSPACE whether
855 $\Pr_s(W_B) = 0$.

We prove the claim. We have:

$$\begin{aligned} & \Pr_s(W_B) > 0 \\ \iff & \exists u. \exists s' : M(u)(s, s') > 0, \Delta(B, u) \text{ is dv} \\ \iff & \exists u. \exists s' : M(u)(s, s') > 0, \Delta(B, u) \text{ is dv, } |u| \leq 2^{|S \times Q|} \end{aligned}$$

The last equivalence follows from the fact that there are at most $2^{|S \times Q|}$ different beliefs of the form $\Delta(B, u)$. Therefore, one can check in NPSpace = PSPACE whether $\Pr_s(W_B) > 0$:

- 860
1. guess a word $u = a_1 \cdots a_n$ with $n \leq 2^{|S \times Q|}$ letter-by-letter and check that there are states s_1, \dots, s_n with $M(a_1)(s, s_1) > 0$ and $M(a_{i+1})(s_i, s_{i+1}) > 0$ for all $i \in \{1, \dots, n-1\}$;
 2. compute, on the fly, the belief

$$B' := \Delta(B, u) = \Delta(\cdots \Delta(\Delta(B, a_1), a_2) \cdots a_n);$$

3. check that B' is dv; we have shown in items 2, 3 and 5 that this can be done in PSPACE.

Since PSPACE is closed under complement, we have proved the claim.

Towards item 6, let B be any belief. We have:

$$\begin{aligned} & B \text{ is not finitary} \\ \iff & \exists (s, q) \in B : \Pr_s(W_B) < 1 \\ \iff & \exists (s, q) \in B. \exists u. \exists s' : M(u)(s, s') > 0, \Pr_{s'}(W_{\Delta(B, u)}) = 0 \\ \iff & \exists (s, q) \in B. \exists u. \exists s' : M(u)(s, s') > 0, \Pr_{s'}(W_{\Delta(B, u)}) = 0, \\ & |u| \leq |S| \cdot 2^{|S \times Q|} \end{aligned}$$

865 The last equivalence follows from the fact that there are at most $|S| \cdot 2^{|S \times Q|}$ different pairs of the form $(s', \Delta(B, u))$. Therefore, one can check in NPSpace = PSPACE whether B is not finitary:

1. guess a word $u = a_1 \cdots a_n$ with $n \leq |S| \cdot 2^{|S \times Q|}$ letter-by-letter and check that there are states s_1, \dots, s_n with $M(a_1)(s, s_1) > 0$ and 870 $M(a_{i+1})(s_i, s_{i+1}) > 0$ for all $i \in \{1, \dots, n-1\}$;
2. compute, on the fly, the belief

$$B' := \Delta(B, u) = \Delta(\cdots \Delta(\Delta(B, a_1), a_2) \cdots a_n);$$

3. check that $\Pr_{s_n}(W_{B'}) = 0$; we have shown in the claim above that this can be done in PSPACE.

Since PSPACE is closed under complement, we have proved item 6. □

Appendix A.3. Correct Proof of Theorem 15

875 Here is Theorem 15 from the main body:

Theorem 15. *It is undecidable to check if $c_{inf} < 3$, even when a diagnoser exists.*

PROOF. We reduce from the emptiness problem for probabilistic automata. A probabilistic automaton (PA) is a tuple $\mathcal{P} = (S, \Sigma, M, s_0, \eta)$ where S is a finite set of states, Σ is a finite alphabet, the mapping $M : \Sigma \rightarrow [0, 1]^{S \times S}$ specifies the transitions such that $M(a)$ is stochastic for each $a \in \Sigma$, the state s_0 is an initial state, and $\eta \in [0, 1]^S$ is a vector of acceptance probabilities. Extend M to $M : \Sigma^* \rightarrow [0, 1]^{S \times S}$ as in the case of MCs. In the case of PAs, $M(u)$ is stochastic for each $u \in \Sigma^*$. For each u define $\Pr_{\mathcal{P}}(u) := e_{s_0} M(u) \eta^T$. The 880

885 probability $\Pr_{\mathcal{P}}(u)$ can be interpreted as the probability that \mathcal{P} accepts u . The *emptiness problem* asks, given a PA \mathcal{P} , whether there is u such that $\Pr_{\mathcal{P}}(u) > \frac{1}{2}$. This problem is undecidable [12, p. 190, Theorem 6.17].

We assume that for all u there is s with $M(u)(s_0, s) > 0$ and $0 < \eta(s) < 1$. This is without loss of generality, as we can make the PA branch, in its first
890 transition and with positive probability, to a sub-PA with a single state from which every word is accepted with probability $\frac{1}{2}$. Formally, from the original PA $\mathcal{P} = (S, \Sigma, M, s_0, \eta)$ obtain another PA $\mathcal{P}' = (S' \cup \{t_0, t\}, \Sigma, M', t_0, \eta')$ where:

- t_0, t are fresh states;
- $\eta'(t_0) := \eta(s_0)$ and $\eta'(t) := \frac{1}{2}$ and $\eta'(s) := \eta(s)$ for all s ;
- 895 • $M'(a)(t_0, t) := \frac{1}{2}$ and $M'(a)(t, t) := 1$ for all $a \in \Sigma$;
- $M'(a)(t_0, s) := \frac{1}{2}M(a)(s_0, s)$ for all $a \in \Sigma$ and all s ;
- $M'(a)(s, s') = M(a)(s, s')$ for all s, s' .

We have for all u that $\Pr_{\mathcal{P}'}(u) - \frac{1}{2} = \frac{1}{2}(\Pr_{\mathcal{P}}(u) - \frac{1}{2})$. It follows that $\Pr_{\mathcal{P}}(u) > \frac{1}{2}$ if and only if $\Pr_{\mathcal{P}'}(u) > \frac{1}{2}$.

900 Let $\mathcal{P} = (S_{\mathcal{P}}, \Sigma_{\mathcal{P}}, M_{\mathcal{P}}, s_{0\mathcal{P}}, \eta)$ be the given PA. We will construct a DFA \mathcal{A} and an MC \mathcal{M} over the alphabet $\Sigma := \Sigma_{\mathcal{P}} \cup \{0, 1\}$, where 0, 1 are fresh letters. Define \mathcal{A} to be a DFA that accepts $L = 0(\Sigma\Sigma\Sigma\Sigma\Sigma)^*\Sigma\Sigma\Sigma\Sigma\Sigma 1\Sigma^\omega$. Ignoring the very first letter 0, we might characterize L by saying that 1 appears on a position that is divisible by 5.

We construct an MC $\mathcal{M} = (S, \Sigma, M, s_0)$ with

$$S := \{s_0\} \cup (S_{\mathcal{P}} \times \{0, 1, 2, 3\}) \cup \tilde{S},$$

905 such that s_0 and the states in \tilde{S} are fresh states. The MC \mathcal{M} initially splits randomly into a “0-copy” (with states in $S \times \{0, 2\}$) and a “1-copy” (with states in $S \times \{1, 3\}$), in either case emitting the letter 0. Formally, $M(0)(s_0, (s_{0\mathcal{P}}, 0)) = M(0)(s_0, (s_{0\mathcal{P}}, 1)) = \frac{1}{2}$. The MC \mathcal{M} is constructed such that if it goes into the 0-copy then surely it emits an infinite word that is not in L ; and if it goes into
910 the 1-copy then almost surely it emits an infinite word in L . Therefore, for an observation policy it suffices to identify which copy \mathcal{M} has entered.

The transitions in \mathcal{M} depend on $M_{\mathcal{P}}$ and η . In Figure A.4 we illustrate this dependence with an example, where $\Sigma_{\mathcal{P}} = \{a, b\}$.

Formally, for all $(s, i) \in S_{\mathcal{P}} \times \{0, 1\}$ we define transitions (using states in \tilde{S}

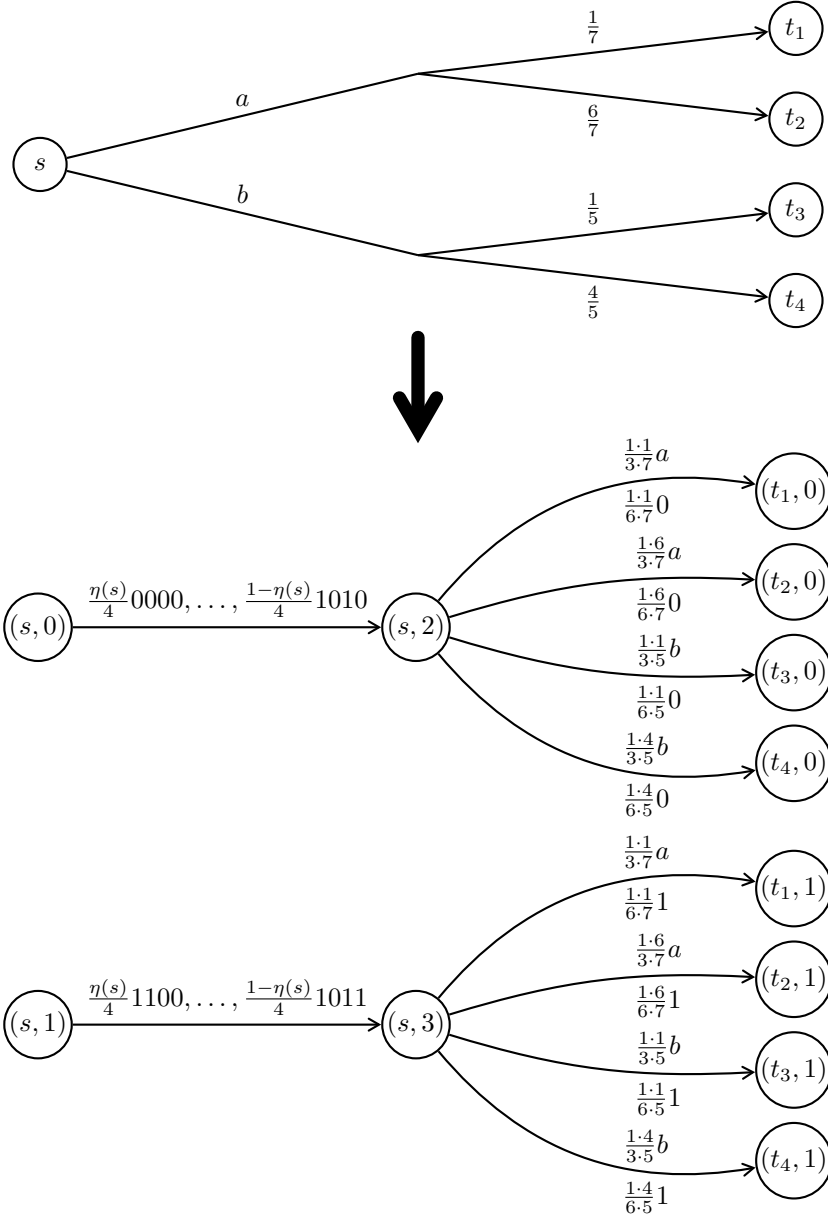


Figure A.4: Illustration of the reduction. The PA at the top has a state s whose outgoing transitions are such that $M(a)(s, t_1) = \frac{1}{7}$, $M(a)(s, t_2) = \frac{6}{7}$, $M(b)(s, t_3) = \frac{1}{5}$, $M(b)(s, t_4) = \frac{4}{5}$. The resulting MC has two corresponding states, $(s, 0)$ and $(s, 1)$. From such a state (s, i) it first emits, with a probability depending on $\eta(s)$, a 4-bit sequence (“block”), and then emits, with probability $\frac{1}{3}$, the letter i (thus giving away in which copy the MC is), or, with probability $\frac{2}{3} \cdot \frac{1}{|\Sigma_P|} = \frac{1}{3}$ each, the letters a or b .

that are not further specified) such that:

$$\begin{aligned}
M(0000)((s, 0), (s, 2)) &= \eta(s)/4 & M(0101)((s, 0), (s, 2)) &= (1 - \eta(s))/4 \\
M(0001)((s, 0), (s, 2)) &= \eta(s)/4 & M(0110)((s, 0), (s, 2)) &= (1 - \eta(s))/4 \\
M(0010)((s, 0), (s, 2)) &= \eta(s)/4 & M(1001)((s, 0), (s, 2)) &= (1 - \eta(s))/4 \\
M(0011)((s, 0), (s, 2)) &= \eta(s)/4 & M(1010)((s, 0), (s, 2)) &= (1 - \eta(s))/4 \\
\\
M(1100)((s, 1), (s, 3)) &= \eta(s)/4 & M(0100)((s, 1), (s, 3)) &= (1 - \eta(s))/4 \\
M(1101)((s, 1), (s, 3)) &= \eta(s)/4 & M(0111)((s, 1), (s, 3)) &= (1 - \eta(s))/4 \\
M(1110)((s, 1), (s, 3)) &= \eta(s)/4 & M(1000)((s, 1), (s, 3)) &= (1 - \eta(s))/4 \\
M(1111)((s, 1), (s, 3)) &= \eta(s)/4 & M(1011)((s, 1), (s, 3)) &= (1 - \eta(s))/4
\end{aligned} \tag{A.1}$$

These 4-bit sequences (henceforth “blocks”) are chosen such that if the first two bits agree (say they are both $b \in \{0, 1\}$) then \mathcal{M} is in the b -copy; if the first two bits of a block do not agree then the xor of all four bits (equivalently, the parity of the sum of the four bits) identifies in which copy \mathcal{M} is. States (s, i) with higher acceptance probability $\eta(s)$ have a higher chance to emit a block where the first two bits agree. The intention of this construction is to make it cost-efficient for an observation policy to observe (some) letters of a block when the policy believes that it is in a state (s, i) with high $\eta(s)$. We will show later that this is beneficial on some runs if and only if there is u with $\Pr_{\mathcal{P}}(u) > \frac{1}{2}$.

The other transitions of \mathcal{M} are as follows: For all $s, t \in \mathcal{S}_{\mathcal{P}}$ and both $i \in \{0, 1\}$ we have the following transitions:

$$\begin{aligned}
M(a)((s, i + 2), (t, i)) &:= \frac{2}{3} \cdot \frac{1}{|\Sigma_{\mathcal{P}}|} \cdot M_{\mathcal{P}}(a)(s, t) & \forall a \in \Sigma_{\mathcal{P}} \\
M(i)((s, i + 2), (t, i)) &:= \frac{1}{3} \cdot \frac{1}{|\Sigma_{\mathcal{P}}|} \cdot \sum_{a \in \Sigma_{\mathcal{P}}} M_{\mathcal{P}}(a)(s, t)
\end{aligned}$$

Note that any state $(s, i + 2)$ emits, with probability $\frac{1}{3}$, the letter i (thus giving away in which copy \mathcal{M} is), or, with probability $\frac{2}{3}$, a letter in $\Sigma_{\mathcal{P}}$ that is chosen uniformly at random. (Thus, if $|\Sigma_{\mathcal{P}}| = 2$, all letters in $\Sigma_{\mathcal{P}} \cup \{i\}$ have probability $\frac{1}{3}$.) In more detail, one can view the behaviour of \mathcal{M} in a state $(s, i + 2)$ as follows: \mathcal{M} first samples a letter a from $\Sigma_{\mathcal{P}}$ uniformly at random, then samples a successor state from $S \times \{i\}$ (each (t, i) with probability $M_{\mathcal{P}}(a)(s, t)$), and then, upon transitioning to (t, i) , emits, with probability $\frac{2}{3}$, the sampled letter a , or, with probability $\frac{1}{3}$, the letter i . This completes the description of the construction.

There exists a diagnoser: Consider the policy $\bar{\rho}$, which, after the initial 0, observes every 5th letter; each time this letter will be either from $\Sigma_{\mathcal{P}}$ (with probability $\frac{2}{3}$) or from $\{0, 1\}$ (with probability $\frac{1}{3}$). With probability 1 the latter case occurs eventually. Once it has occurred, the copy has been revealed; thus the observation prefix becomes deciding and $\bar{\rho}$ stops making any observations. The required number of observations under $\bar{\rho}$ is geometrically distributed, with expectation $1/\frac{1}{3} = 3$. Hence $\text{Ex}(C_{\bar{\rho}}) = 3$ and thus $c_{inf} \leq 3$.

It remains to show that $c_{inf} < 3$ if and only if there is u such that $\Pr_{\mathcal{P}}(u) > \frac{1}{2}$.
 940 Suppose u_0 is such that $\Pr_{\mathcal{P}}(u_0) > \frac{1}{2}$. Consider the following policy ρ_0 :

- do not observe the initial 0 and then observe every 5th letter until either the prefix is deciding (because 0 or 1 has been observed) or the observation prefix becomes equal to u_0 ;
- in the former case: stop observing forever;
- 945 • in the latter case: observe the first two letters of the following block;
 - if they agree then the observation prefix has become deciding, so stop observing;
 - if they do not agree, observe also the next two letters; this produces a deciding prefix, so stop observing.

950 This policy coincides mostly with the diagnoser $\bar{\rho}$ described above. It only deviates if and when u_0 has been observed. We show that this deviation improves the expected cost.

To argue in more formal terms, we extend (for the PA \mathcal{P}) the mapping $M_{\mathcal{P}} : \Sigma_{\mathcal{P}}^* \rightarrow [0, 1]^{S_{\mathcal{P}} \times S_{\mathcal{P}}}$ to $M_{\mathcal{P}} : (\Sigma_{\mathcal{P}} \cup \{\perp\})^* \rightarrow [0, 1]^{S_{\mathcal{P}} \times S_{\mathcal{P}}}$ by defining

$$M_{\mathcal{P}}(v) := \sum_{u \sim v} M_{\mathcal{P}}(u) / |\{u \mid u \sim v\}| \quad \text{and define}$$

$$\mu_{\mathcal{P}}(v) := e_{s_0 \mathcal{P}} M_{\mathcal{P}}(v).$$

One can view the distribution $\mu_{\mathcal{P}}(v)$ as the expected distribution after having fed the PA \mathcal{P} with a randomly sampled $u \sim v$. Similarly, we extend (for the MC \mathcal{M}) the mapping $M : \Sigma^* \rightarrow [0, 1]^{S \times S}$ to $M : \Sigma_{\perp}^* \rightarrow [0, 1]^{S \times S}$ by defining

$$M(v) := \sum_{u \sim v} M(u) \quad \text{and define}$$

$$\mu(v) := \frac{e_{s_0} M(v)}{e_{s_0} M(v) \mathbf{1}^T}.$$

One can view the distribution $\mu(v)$ as the expected distribution after having observed v . Finally, for any finite word $v = o_1 o_2 \cdots o_n \in (\Sigma_{\mathcal{P}} \cup \{\perp\})^*$ define the following padding:

$$\hat{v} := \begin{cases} \perp & \text{if } n = 0 \\ \perp \perp \perp \perp \perp o_1 \perp \perp \perp \perp o_2 \perp \perp \perp \perp o_3 \cdots \perp \perp \perp \perp o_n & \text{if } n \geq 1 \end{cases}$$

With these definitions, it is straightforward to check that we have:

$$\frac{1}{2} \mu_{\mathcal{P}}(v)(s) = \mu(\hat{v})((s, 0)) = \mu(\hat{v})((s, 1)) \quad \forall v \in (\Sigma_{\mathcal{P}} \cup \{\perp\})^* \quad \forall s \in S_{\mathcal{P}} \quad (\text{A.2})$$

For the word u_0 from above we have $\mu_{\mathcal{P}}(u_0) \eta^T = \Pr_{\mathcal{P}}(u_0) > \frac{1}{2}$. It follows from (A.2) and (A.1) that, conditioned under prefix \hat{u}_0 , the conditional probability

of emitting 00 or 11 at the beginning of the following block is greater than $\frac{1}{2}$; formally:

$$\Pr(\{u00, u11\}\Sigma^\omega \mid u \sim \widehat{u}_0 \mid \{w \succeq \widehat{u}_0\}) = \mu_{\mathcal{P}}(u_0)\eta^\top > \frac{1}{2}$$

Recall that ρ_0 is defined so that once it has observed \widehat{u}_0 , it makes either exactly 2 or exactly 4 further observations. From the previous inequality it follows for ρ_0 that, conditioned under observing \widehat{u}_0 , the conditional probability to make exactly 2 further observations is greater than the conditional probability to make exactly 4 further observations. Hence, the conditional expected number of observations after having observed u_0 is less than 3.

More formally, for any policy ρ and an observation prefix $v = o_1 \cdots o_k$ define a random variable C_ρ^v by

$$C_\rho^v(w) := C_\rho(w) - \sum_{k=0}^{|v|-1} \rho(o_1 \dots o_k) \quad \text{for all } w \succeq v,$$

i.e., C_ρ^v is the number of observations that ρ makes after v . By the argument above we have $\text{Ex}(C_{\rho_0}^{\widehat{u}_0} \mid \{w \succeq \widehat{u}_0\}) < 3$. For all u that are not prefixes of u_0 we have $\text{Ex}(C_{\rho_0}^{\widehat{u}} \mid \{w \succeq \widehat{u}\}) = 3$. It follows $\text{Ex}(C_{\rho_0}) < 3$. Hence $c_{inf} < 3$.

It remains to show the converse. To this end, suppose all u satisfy $\Pr_{\mathcal{P}}(u) \leq \frac{1}{2}$. It suffices to show that all feasible policies ρ satisfy $\text{Ex}(C_\rho) \geq 3$.

In the following we use regular expressions to describe observation prefixes. For improved readability we may indicate the borders of a block with a dot. For instance, the regular expression

$$\perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot 0\perp\perp\perp\perp.$$

indicates observation prefixes of the following form: first all observations (if any) are made in non-blocks and are in $\Sigma_{\mathcal{P}}$ (i.e., are not 0 or 1), and then 0 is observed at the beginning of a block, and the other three letters in the block are not observed.

We call $k \in (0+1+\perp)^4$ *block-deciding* when $k \in (0+1)^4 + (00+11)(0+1+\perp)^2$. We have the following lemma:

Lemma 29. *Let v be an observation prefix that is not deciding and satisfies $|v| = 1 + 5n$ for some $n \in \mathbb{N}$. Let $k \in (0+1+\perp)^4$. Then $v.k$ is deciding if and only if k is block-deciding.*

PROOF (OF THE LEMMA). Let v and k be as in the statement. Suppose k is block-deciding. Then, by inspecting (A.1), it follows that $v.k$ is deciding.

Conversely, suppose k is not block-deciding. From the assumption made about \mathcal{P} in the beginning we obtain that there is $s \in S_{\mathcal{P}}$ with $0 < \eta(s) < 1$ such that $M(v)(s_0, (s, 0)) > 0$ and $M(v)(s_0, (s, 1)) > 0$. By inspecting (A.1), it follows that $v.k$ is not deciding. \square

Let ρ be any feasible policy, and let v be an observation prefix that ρ produces with positive probability. Recall that we aim to show $\text{Ex}(C_\rho) \geq 3$. With
980 distinguish between several possible forms of v ; in each case we give a lower bound on $\text{Ex}(C_\rho^v)$.

1. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot k(\cdot\perp^5)^*(\Sigma_{\mathcal{P}} + \perp)(\cdot\perp^5)^* \cdot \ell$ where $k \in (0 + 1 + \perp)^4$ is not block-deciding, and $\ell \in (0 + 1 + \perp)^{\leq 4}$ has exactly one observation.
985 By [Lemma 29](#) the prefix up to and including k is not deciding. Another application of [Lemma 29](#) shows that v is not deciding, so $\text{Ex}(C_\rho^v) \geq 1$.
2. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot k(\cdot\perp^5)^* \Sigma_{\mathcal{P}}$ where $k \in (0 + 1 + \perp)^4$ is not block-deciding. Similarly as before, it follows that v is not deciding.

- Suppose the next observation that ρ makes is in a block. Then, by
990 [item 1](#), we have $\text{Ex}(C_\rho^v) \geq 1 + 1 = 2$.
- Otherwise the next observation is in a non-block. If this observation yields 0 or 1 (which happens with probability $\frac{1}{3}$) then no further observation is needed; otherwise, a further observation will be needed. So $\text{Ex}(C_\rho^v) \geq 1 + \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot 1 \geq \frac{5}{3}$.

995 Thus, in either case we have $\text{Ex}(C_\rho^v) \geq \frac{5}{3}$.

3. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot k$, where $k \in (0 + 1 + \perp)^4$ is not block-deciding. Similarly as before, it follows that v is not deciding.
 - Suppose the next observation that ρ makes is in a block. Then, by
1000 [item 1](#), we have $\text{Ex}(C_\rho^v) \geq 1 + 1 = 2$.
 - Otherwise the next observation is in a non-block. If this observation yields 0 or 1 (which happens with probability $\frac{1}{3}$) then no further observation is needed; otherwise, the resulting observation prefix will have the form from [item 2](#). It follows from [item 2](#) that $\text{Ex}(C_\rho^v) \geq 1 + \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot \frac{5}{3} > 2$.

1005 Thus, in either case we have $\text{Ex}(C_\rho^v) \geq 2$.

4. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot \perp\perp b$ for $b \in \{0, 1\}$. Then v is not deciding.
 - Suppose the next observation is not in the same block. Then the resulting observation prefix has the form from either [item 1](#) or [item 2](#). It follows from these items that $\text{Ex}(C_\rho^v) \geq 1 + \min\{1, \frac{5}{3}\} = 2$.
 - Otherwise the next observation follows immediately and yields $b' \in \{0, 1\}$. The word $\perp\perp bb'$ is not block-deciding, hence, by [Lemma 29](#), vb' is not deciding, so $\text{Ex}(C_\rho^v) \geq 2$.

1010 Thus, in either case we have $\text{Ex}(C_\rho^v) \geq 2$.

5. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^* \cdot bb'$ for $bb' \in (0 + 1 + \perp)^2$ and $bb' \notin 00 + 11$.
1015 Then v is not deciding.
 - Suppose the next observation is not in the same block. Then the resulting observation prefix has the form from either [item 1](#) or [item 2](#). It follows from these items that $\text{Ex}(C_\rho^v) \geq 1 + \min\{1, \frac{5}{3}\} = 2$.

- Otherwise the next observation results in an observation prefix vv' with $v' \in \{0, 1, \perp 0, \perp 1\}$. No word in $\{bb'\}\{0\perp, 1\perp, \perp 0, \perp 1\}$ is block-deciding. Hence, by [Lemma 29](#), vv' is not deciding, so $\text{Ex}(C_\rho^v) \geq 2$.

Thus, in either case we have $\text{Ex}(C_\rho^v) \geq 2$.

6. Suppose $v \in \perp(\cdot\perp^4(\Sigma_{\mathcal{P}} + \perp))^*b$ for $b \in \{0, 1\}$. Then v is not deciding.

- Suppose the next observation is not in the same block. Then the resulting observation prefix has the form from either [item 1](#) or [item 2](#). It follows from these items that $\text{Ex}(C_\rho^v) \geq 1 + \min\{1, \frac{5}{3}\} = 2$.
- Suppose the next observation is in the same block but does not follow immediately. This results in an observation prefix $v\perp v'$ with $v' \in \{0, 1, \perp 0, \perp 1\}$. No word in $\{b\perp\}\{0\perp, 1\perp, \perp 0, \perp 1\}$ is block-deciding. Hence, by [Lemma 29](#), vv' is not deciding, so $\text{Ex}(C_\rho^v) \geq 2$.
- Otherwise the next observation follows immediately. Let v_0 be such that $v = v_0b$. Then there exists $v_1 \in (\Sigma_{\mathcal{P}} \cup \{\perp\})^*$ with $v_0 = \hat{v}_1$. Since all $u \in \Sigma_{\mathcal{P}}^*$ satisfy $\mu_{\mathcal{P}}(u)\eta^{\text{T}} = \Pr_{\mathcal{P}}(u) \leq \frac{1}{2}$, we also have $\mu_{\mathcal{P}}(v_1)\eta^{\text{T}} \leq \frac{1}{2}$. It follows from [\(A.2\)](#) and [\(A.1\)](#) that, conditioned under prefix v_0 , the conditional probability of emitting 00 or 11 is at most $\frac{1}{2}$; formally:

$$p := \Pr(\{u00, u11\}\Sigma^\omega \mid u \sim v_0 \mid \{w \succcurlyeq v_0\}) = \mu_{\mathcal{P}}(v_1)\eta^{\text{T}} \leq \frac{1}{2}$$

For symmetry reasons we must have:

$$\Pr(\{ub\}\Sigma^\omega \mid u \sim v \mid \{w \succcurlyeq v\}) = p \leq \frac{1}{2}$$

In words, the conditional probability that the next observation equals b is at most $\frac{1}{2}$. If this happens then no further observation is needed, as $bb\perp\perp$ is block-deciding; otherwise, the resulting observation prefix has the form from [item 5](#). It follows from [item 5](#) that $\text{Ex}(C_\rho^v) \geq 1 + p \cdot 0 + (1 - p) \cdot 2 \geq 1 + \frac{1}{2} \cdot 2 = 2$.

Thus, in all three cases we have $\text{Ex}(C_\rho^v) \geq 2$.

Let ρ be feasible. Call an observation prefix v *conventional* when

- v is not deciding;
- ρ produces v with positive probability;
- and v does not contain observations in blocks.

Towards a contradiction, assume that $\text{Ex}(C_\rho) < 3$. Then there exist a conventional observation prefix v and $x > 0$ such that:

- $\text{Ex}(C_\rho^v) \leq 3 - x$;
- and for all conventional observation prefixes v' we have $\text{Ex}(C_\rho^{v'}) > 3 - \frac{3}{2}x$.

Since ρ is feasible, there must be a non- \perp observation (at some point) after v .

- Suppose this next observation is in a non-block, resulting in an observation prefix v' . Then v' is either deciding (with probability $\frac{1}{3}$) or conventional (with probability $\frac{2}{3}$). It follows:

$$3 - x \geq \text{Ex}(C_\rho^v) \geq 1 + \frac{1}{3} \cdot 0 + \frac{2}{3} \cdot \text{Ex}(C_\rho^{v'}) > 1 + \frac{2}{3} \cdot \left(3 - \frac{3}{2}x\right) = 3 - x,$$

which is a contradiction.

- Otherwise, the next observation is in a block, resulting in an observation prefix v' . Then v' is not deciding and has the form of items 3–6 above. It follows from these items that $\text{Ex}(C_\rho^v) = 1 + \text{Ex}(C_\rho^{v'}) \geq 1 + 2 = 3$, contradicting that $\text{Ex}(C_\rho^v) \leq 3 - x$.

1050

In either case we have a contradiction. Hence $\text{Ex}(C_\rho) \geq 3$. □