Strategies, Credences, and Shannon Entropy: Reasoning about Strategic Uncertainty in Stochastic Environments

Wojciech Jamroga[1,2] , **Michał Tomasz Godziszewski**[3,4] , **Aniello Murano**[3]
[1] Institute of Computer Science, Polish Academy of Sciences
[2]Nicolaus Copernicus University, Toruń, Poland
[3]University of Naples Federico II, Italy
[4]University of Lodz, Poland
mtgodziszewski@gmail.com, nello.murano@gmail.com, wojciech.jamroga@ipipan.waw.pl

Abstract

A central focus in multi-agent systems is the agents' ability to achieve their goals, often involving epistemic objectives like acquiring knowledge about a crucial fact φ. Many such properties can be expressed using PATLK, an extension of probabilistic alternating-time temporal logic (PATL) with knowledge operators, or PATLC that extends PATL with probabilistic beliefs.

In many scenarios, however, the goal of the players is not to achieve high confidence about φ being true, but rather to reduce their uncertainty about φ (be it true or false). To capture such properties, we introduce PATLH, a logic extending PATL with information-theoretic modalities based on Shannon entropy. As technical results, we compare the epistemic and information-theoretic extensions of PATL with respect to their expressiveness, succinctness, and complexity of model checking.

1 Introduction

Motivation. In the field of formal reasoning about multi-agent systems (MAS), strategic uncertainty plays a critical role in decision-making processes (see [Bulling *et al.*, 2015; Ågotnes *et al.*, 2015]). Agents within such systems often need to reason about the outcomes of actions in environments that are not fully observable or deterministic. Therefore, strategic reasoning of individual agents and their groups must account for both incomplete information about their environment and the stochastic nature of this environment when formulating strategic decisions and plans. This requires the development of formal models that allow agents to make strategic decisions despite various layers of uncertainty, and take into account the stochastic behavior of MAS.

Knowledge and information play a vital role in interactions, especially with the rise of the Internet and social networks. Information is both a key resource for strategy formation and often the primary goal of interaction. Agents may act to learn new information or to protect secrets. For example, people seek economic trends, fashion updates, or even the quality of workplace coffee. Such scenarios inherently involve both knowledge and uncertainty. Analyzing probabilistic information is essential in these contexts. By computing the likelihood of outcomes and integrating uncertainty into decision-making, agents can make informed choices in environments where complete knowledge is unattainable.

Logics for strategic uncertainty. Alternating-time temporal logic **ATL** [Alur *et al.*, 2002]is a widely recognized framework for reasoning about strategic capabilities in MAS. **ATL** allows for the expression of strategic abilities of groups of agents, where certain outcomes can be guaranteed through cooperative strategies. A probabilistic extension of the logic, **PATL** [Chen and Lu, 2007; Huang *et al.*, 2012; Belardinelli *et al.*, 2024], enhanced **ATL** by probabilistic reasoning about the likelihood of achieving certain outcomes, rather than deterministic guarantees.

An important step towards reasoning about strategic uncertainty was made with **ATLH** [Tabatabaei and Jamroga, 2023]. **ATLH** extended **ATL** with modalities based on the Hartley measure [Hartley, 1928a], that can be seen as a non-probabilistic restriction of Shannon entropy [Shannon, 1948]. However, in scenarios where agents *do* have beliefs about the likelihoods of events, the information gain from revealing certain facts is intimately related to those likelihoods.

Contribution. To fill this major gap, we consider "**PATL** with Credences" (**PATLC**), which is in fact a more general version of PATEL [Huang and Luo, 2013]. **PATLC** extends Probabilistic **ATL** by graded epistemic operators $\mathcal{K}_a^{\geq q}\varphi$ for probabilistic beliefs (or *credences*), with the intended meaning: "agent a believes that φ holds with probability at least q." In contrast to **PATL** with binary knowledge, this allows us to incorporate the stochastic behavior of the environment and its influence on the epistemic aspects of strategic reasoning.

Even more importantly, we introduce "**PATL** with Shannon uncertainty" (**PATLH**). **PATLH** is an information-theoretic extension of **PATL** that adds quantitative uncertainty modalities based on Shannon entropy – a foundational concept in information theory. By incorporating entropy-based reasoning, **PATLH** allows agents to reason about both their strategic options and their epistemic uncertainty. This is particularly useful in scenarios where agents must either reduce the uncertainty to gain knowledge about specific facts or, conversely, maintain a high level of uncertainty to keep sensitive information hidden from adversaries. The added

layer of expressiveness makes **PATLH** suited, e.g., for reasoning about security and privacy in MAS, where the control of information flow is as important as the control of outcomes. To our best knowledge, this is the first logical formalism that incorporates stochastic aspects of the system directly into strategic reasoning about uncertainty of information flow.

In terms of technical results, we prove that (1) **PATLC** and **PATLH** have the same model checking complexity; (2) **PATLC** is strictly more expressive than **PATLH**; and conjecture that (3) **PATLH** is exponentially more succinct than **PATLH**. Thus, **PATLC** allows for expressing more properties, but **PATLH** may offer exponentially shorter encodings for those that it can express. In consequence, the *practical* verification of **PATLH** would be often exponentially easier, which makes both logics interesting for applications that require both strategic and information-theoretic reasoning.

Related Work Strategic-epistemic reasoning gained significant attention in the early 2000s, particularly within frameworks like ATEL [van der Hoek and Wooldridge, 2003; Ågotnes, 2006; Jamroga and Ågotnes, 2007] and DEL [van Ditmarsch *et al.*, 2007; Ågotnes and van Ditmarsch, 2008], To the best of our knowledge, the only works that combine logical approaches to strategic reasoning with information-theoretic concepts are [Jamroga and Tabatabaei, 2013; Tabatabaei and Jamroga, 2023], and both concern non-probabilistic interaction. MAS with stochastic interactions has been extensively studied, see e.g. [Hansson and Jonsson, 1994; Chen and Lu, 2007; Huang *et al.*, 2012; Huang and Luo, 2013; Belardinelli *et al.*, 2024]. However, these works did not focus on epistemic or information-theoretic aspects. A notable exception is [Huang and Luo, 2013] which proposed PATEL, i.e., Probabilistic **ATL** extended with probabilistic belief operators, and investigated its model checking complexity. Metrics of uncertainty and information gain were introduced by Hartley [Hartley, 1928a] and extended by Shannon [Shannon, 1948]. The latter is commonly used in information security to measure the amount of information leaking to the intruder [Alvim *et al.*, 2020]. An extension of **ATL** with modalities based on Hartley measure was recently introduced in [Tabatabaei and Jamroga, 2023]. In this paper, we extend this framework to probabilistic concepts and models.

Our main technical results concern the comparison of expressiveness between epistemic and information-theoretic modalities, and the succinctness of information-theoretic modalities based on Shannon entropy. The study of the succinctness of logical representations dates back to the early 1970s [Stockmeyer, 1972]. The relative succinctness of branching-time logics was explored in [Wilke, 1999; Adler and Immerman, 2001; Markey, 2003], while the succinctness of **ATL*** with past-time operators was studied in [Bozzelli *et al.*, 2020]. The method of proving succinctness using *formula size games* was first introduced in [Adler and Immerman, 2001], and later generalized in [French *et al.*, 2013b]. In this work, we build on the latter approach, and follow [Tabatabaei and Jamroga, 2023] in establishing our result.

2 Preliminaries

Distributions and Markov Chains. Let X be a finite nonempty set. A probability distribution over X is a function $p : X \to [0, 1]$ such that $\sum_{x \in X} p(x) = 1$, and $Dist(X)$ is the set of probability distributions over X. A *Markov Chain* M is a tuple (St, p) where St is a nonempty set of states and $p \in Dist(St \times St)$ is a distribution. The values $p(s, t)$ are transition probabilities of M denoting the likelihood of a transition from s to t. A path is an infinite sequence of states.

Stochastic iCGS. *Stochastic Imperfect Information Concurrent Game Structures* (*Stochastic iCGS*, or *SiCGS*) extend the well-known Concurrent Game Structures [Alur *et al.*, 2002] to allow for both imperfect information and stochastic transitions. Formally, a Stochastic iCGS is a tuple $\mathcal{G} = (\mathbb{A}\mathrm{gt}, St, Act, d, \delta, \{\sim_a\}_{a \in \mathbb{A}\mathrm{gt}})$, where $\mathbb{A}\mathrm{gt}$ is a finite set of agents, St is a finite set of states, Act is a finite set of actions, $d : \mathbb{A}\mathrm{gt} \times St \to 2^{Act} \setminus \{\emptyset\}$ is a function defining the available actions for each agent in each state, $\delta : St \times Act^n \to \mathcal{D}(St)$ is a stochastic transition function (where $\mathcal{D}(St)$ denotes the set of probability distributions over St) that gives the (conditional) probability $\delta(s, \mathbf{c})$ of a transition from state s for all $s' \in St$ if each player $a \in \mathbb{A}\mathrm{gt}$ plays the action $\mathbf{c_a}$ (we also write this probability as $\delta(s, \mathbf{c})(\mathbf{s}')$ to emphasize that $\delta(s, \mathbf{c})$ is a probability distribution on St), and $\sim_a \subseteq St \times St$ for each $a \in \mathbb{A}\mathrm{gt}$ is an equivalence relation capturing the indistinguishability of states for agent a.

Strategies and Probabilistic Outcomes. For the interplay between a coalition $C \subseteq \mathbb{A}\mathrm{gt}$ and its opponents $\overline{C} = \mathbb{A}\mathrm{gt} \setminus C$, we follow [Belardinelli *et al.*, 2023a] and assume that C play *uniform deterministic memoryless strategies* (ir-strategies in short), whereas $\overline{C}$ can respond with any pattern of behavior (possibly probabilistic and history-based). Formally, an *ir-strategy* for agent a is a function $\sigma_a : St \to Act$in which: (i) for each s, we have $\boldsymbol{\sigma}_a(s) \in d(s, a)$; and (ii) if $s \sim_a s'$ then $\boldsymbol{\sigma}_a(s) = \boldsymbol{\sigma}_a(s')$. Moreover, a *general strategy* for agent $a \in \mathbb{A}\mathrm{gt}$ is represented by $\boldsymbol{\sigma}_a : St^+ \to \mathcal{D}(Act)$ that maps each finite history to a probability distribution over the agent's actions, such that $\boldsymbol{\sigma}_a(h) \in d(last(h), a)$. A *collective strategy* $\boldsymbol{\sigma}_C$ for C is a tuple of strategies $\boldsymbol{\sigma}_a$, one per agent $a \in C$.

The *outcome* of strategy $\boldsymbol{\sigma}_C$ from state s is the set $out_C(\boldsymbol{\sigma}_C, s)$ of probability distributions over infinite paths in the model, consistent with C's choices prescribed by $\boldsymbol{\sigma}_C$. Each distribution $\mu_{\boldsymbol{\sigma}_C, s} \in out_C(\boldsymbol{\sigma}_C, s)$ is obtained from the Markov chain that combines the Stochastic iCGS M with $\boldsymbol{\sigma}_C$ and a possible general strategy of $\overline{C}$. We refer to [Belardinelli *et al.*, 2023a] for the detailed construction.

Probabilistic Alternating-Time Logic PATL. We now introduce Probabilistic Alternating-Time The language of **PATL*** is defined as follows:

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \varphi \mathcal{U} \varphi \mid \bigcirc \varphi \mid \langle\!\langle C \rangle\!\rangle^{\propto p} \varphi,$$

where $\langle\!\langle C \rangle\!\rangle^{\propto p} \varphi$ means that there exists a strategy for the coalition C to collaboratively enforce φ with a probability in relation $\propto$ with constant p, where $\propto \in \{=, \neq, >, <, \geq, \leq\}$.

The language of **PATL** is a restriction of the language of **PATL*** and is defined by the following grammar:

$$\varphi ::= \mathsf{p} \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\!\langle C \rangle\!\rangle^{\propto p} \bigcirc \varphi \mid \langle\!\langle C \rangle\!\rangle^{\propto p} \varphi \mathcal{U} \varphi,$$

C to collaboratively enforce φ with a probability in relation $\propto$ with constant p.

An iSCGS together with a set of atomic propositions AP and a valuation function $V : AP \to 2^{St}$ is called a stochastic concurrent game model (SCGM). The formulas of **PATL** are interpreted over pairs (M, π), consisting of an SCGM and an infinite path in it.

Satisfaction relation. The semantics of **PATL** extends the semantics of **ATL** and is defined by the following clauses:

$M, \pi \models p$ iff $p \in V(\pi_0)$;

$M, \pi \models \neg\varphi$ iff $M, \pi \not\models \varphi$;

$M, \pi \models \varphi \wedge \psi$ iff $M, \pi \models \varphi$ and $M, \pi \models \psi$;

$M, \pi \models \langle\langle C\rangle\rangle^{\propto p}\varphi$ iff there is an ir-strategy $\boldsymbol{\sigma}_C$ s.t. for all $\mu_{\boldsymbol{\sigma}_C,\pi_0} \in out_C(\boldsymbol{\sigma}_C, \pi_0)$ we have $\mu_{\boldsymbol{\sigma}_C,\pi_0}(\{\rho : M, \rho \models \varphi\}) \propto p$;

$M, \pi \models \bigcirc\varphi$ iff $M, \pi_{\geq 1} \models \varphi$;

$M, \pi \models \varphi\mathcal{U}\psi$ iff there is a k s.t. $M, \pi_{\geq k} \models \psi$ and for all $j < k$ we have $M, \pi_{\geq j} \models \varphi$.

We define $M, s \models \varphi$ iff $M, \pi \models \varphi$ holds for any π starting in s. $Sat(M, \varphi) = \{s \in St \mid M, s \models \varphi\}$ is the subset of states in M satisfying φ. Moreover, *pointed models* are pairs (M, s) consisting of an SCGM and a state in it. For a subset of pointed models A, we use $A \models \varphi$ to mean that for all $(M, s) \in A$ it holds that $M, s \models \varphi$. We also use $\hat{M}$ to denote the set of all pointed models in M.

3 Adding Knowledge and Uncertainty

In this section we introduce probabilistic extensions of **ATLK** [van der Hoek and Wooldridge, 2003] and **ATLH** [Tabatabaei and Jamroga, 2023] that can be simultaneously seen as epistemic and information-theoretic extensions of **PATL** [Belardinelli *et al.*, 2023b].

Knowledge and Credences In multi-agent epistemic logic, *knowledge* of agents is formalized by epistemic formulas $\mathcal{K}_a\varphi$, stating "agent a knows that φ holds." They are interpreted by the following clause:

$M, s \models \mathcal{K}_a\varphi$ iff, for every state t such that $s \sim_a t$, we have that $M, t \models \varphi$,

where $\sim_a \subseteq St \times St$ is an *epistemic equivalence relation* connecting states that are indistinguishable to a.

Analogously, **PATL** can be extended by operators $K_a^{\propto q}$ for probabilistic knowledge/beliefs, or *credences* (where $\propto \in \{<, >, \leq, \geq, =\}$). To provide semantics, we augment SiGS by a probabilistic observation function $obs_a : St \to Dist(St)$, with the idea that $obs_a(s)(t)$ gives the subjective probability with which agent a believes that the current state is t, provided that the actual current state is, in fact, s. We will often write $obs_a(t|s)$ instead of $obs_a(s)(t)$ to make this even clearer. Additionally, we lift the notation to subsets of states by defining $obs_a(T|s) = \sum_{t\in T} obs_a(t|s)$.

We require that epistemic indistinguishability, captured by $\sim_a$, is consistent with obs_a as follows: $s \sim_a t$ (when in s, the agent considers t as possible) iff $obs_a(t|s) > 0$ (when in s, the agent considers t with nonzero probability). Since $\sim_a$ is an equivalence, this implies the following requirements: $obs(s|s) > 0$ (due to reflexivity of $\sim_a$); if $obs(t|s) > 0$ then $obs(s|t) > 0$ (due to symmetry); if $obs(t|s) > 0$ and $obs(s|w) > 0$ then $obs(t|w) > 0$ (due to transitivity).

An SiCGS extended by observation functions obs_a is called a Stochastic Observational CGS (SOCGS). An SOCGS together with a set of atomic propositions AP and a valuation function $V : AP \to 2^{St}$ is called a Stochastic Observational CGM (SOCGM). The semantics of **PATLC** extends **PATL** by the clause:

$$M, s \models \mathcal{K}_a^{\propto q}\varphi \text{ iff } obs_a(Sat(M, \varphi)|s) \propto q.$$

We note in passing that, for finite models, we can express classical knowledge by credences with $\mathcal{K}_a\varphi \equiv \mathcal{K}_a^{=1}\varphi$.

Probabilistic Uncertainty Operators We propose modal operators $\mathcal{H}_a^{\propto m}\Phi$, based on the fundamental notion of *information entropy*, due to [Shannon, 1948]. Let $X = \{x_1, \ldots, x_n\}$ be a countable set of *possible outcomes* (typically, values of a given random variable), and $Pr \in Dist(X)$ a probability distribution over X. Then, its Shannon entropy is defined as $H_S(X) = -\sum_{i=1}^{n} Pr(x_i) \log Pr(x_i)$. Clearly, Shannon entropy is minimal (and equal to 0) if p is a Dirac distribution, i.e., we are certain with probability 1 which outcome is the right one. Conversely, H_S is maximal when Pr is uniform, i.e., the subjective randomness of the system is highest. In that case, Shannon entropy coincides with Hartley uncertainty, defined as $H(X) = \log(|X|)$ [Hartley, 1928b].

Syntax. **PATLH** extends **PATL** with a family of operators $\mathcal{H}_a^{\propto m}\Phi$, where $a \in \mathbb{A}\mathrm{gt}$, Φ is a finite nonempty subset of **PATLH** formulas, and $\propto \in \{<, \leq, >, \geq, =\}$. The reading of $\mathcal{H}_a^{\propto m}\{\varphi_1, \ldots, \varphi_n\}$ is "the uncertainty of a about the actual values of $\varphi_1, \ldots, \varphi_n$ is at least (at most, equal to, etc.) m."

Semantics. We define the semantics of **PATLH** over SOCGMs, as for **PATLK** and **PATLC**. We follow the approach of [Tabatabaei and Jamroga, 2023]. The idea is that, in order to measure agent a's uncertainty about formulas $\Phi = \{\varphi_1, \ldots, \varphi_n\}$, we take the different valuations of Φ as the "possible outcomes." Then, we take a's subjective probabilities about each possible valuation, and compute Shannon entropy for the resulting distribution.

Formally, we start by defining relation $\sim^{\varphi} \in St \times St$ that connects states with the same valuation of φ, i.e.: $s \sim^{\varphi} t$ iff $M, s \models \varphi \Leftrightarrow M, t \models \varphi$. This can be lifted to indiscernibility of states w.r.t. a set of formulas Φ in a natural way: $\sim^{\Phi} = \bigcap_{\varphi\in\Phi}^{n} \sim^{\varphi}$. Moreover, we combine syntactic and semantic indistinguishability through relation $\sim_a^{\Phi} = \sim_a \cap \sim^{\Phi}$. In other words, $s \sim_a^{\Phi} t$ iff s and t are epistemically indistinguishable and no formula in Φ can distinguish between them. Clearly, $\sim_a^{\Phi}$ is an equivalence relation.

Next, we define the relevant outcomes as the abstraction classes of $\sim_a^{\Phi}$ that are contained in the current epistemic class due to $\sim_a$, i.e.: $R_{a,s}(\Phi) = \{[t]_{\sim_a^{\Phi}} \mid t \sim_a s\}$. Then, we construct $Pr_{a,s,\Phi} \in Dist(R_{a,s}(\Phi))$ with $Pr_{a,s,\Phi}([t])$ being the (normalized) aggregate probability that a associates with the states in $[t]$ when the real state of the system is s, i.e.:

$$Pr_{a,s,\Phi}([t]_{\sim_a^{\Phi}}) = \sum_{t'\in[t]_{\sim_a^{\Phi}}} obs_a(t'|s).$$

Finally, we can define the semantics of Shannon uncertainty modalities via the following clause:

$$M, s \models \mathcal{H}_a^{\propto m}\Phi \text{ iff } \big(-\sum_{[t]\in R_{a,s}(\Phi)} Pr_{a,s,\Phi}([t]) \log Pr_{a,s,\Phi}([t])\big) \propto m.$$

Simple Example Consider the following scenario: a voter v decides to vote for or against A, by choosing (at the initial state s_0) whether to proceed to state s_1 (with proposition V_A) or to s_2 (with $\neg V_A$). After v has voted, the coercer c learns the value of the vote with 70% confidence, represented by: $obs_c(s_1|s_1) = obs_c(s_2|s_2) = 0.7, obs_c(s_2|s_1) = obs_c(s_1|s_2) = 0.3$. The following **PATLC** formula holds in s_0:

$\langle\langle v\rangle\rangle\Diamond(V_A \wedge \neg\mathcal{K}_c^{\geq 0.9} V_A) \wedge \langle\langle v\rangle\rangle\Diamond(\neg V_A \wedge \neg\mathcal{K}_c^{\geq 0.9}\neg V_A)$,

expressing that the voter can vote for or against A in such a way that the coercer does not learn about that with confidence 90% or more. Note that the analogous formula for confidence level of 0.7 would not be true anymore. Moreover, the following **PATLH** formula is also true in s_0: $\langle\langle v\rangle\rangle\Box H_c^{\geq 0.88}\{V_A, \neg V_A\}$, saying that the voter can keep the coercer's uncertainty about v's vote at a reasonable level of at least 0.88 bits.

The above formulas allow us to pinpoint two distinct flavors of information leakage: one based on the coercer's increased confidence in the actual value of the vote, and the other on the reduction of his subjective information entropy.

4 Model Checking

In this section we show that adding probabilistic knowledge and uncertainty does not increase the complexity of verification. More precisely, model checking of all the considered logics (**PATLH**, **PATLK**, and **PATLC**) is $\mathbf{\Delta_2^P}$-complete for uniform deterministic strategies of the coalition, and thus no worse than for $\mathbf{ATL}_{ir}$ and $\mathbf{PATL}_{ir}$. Since the input to model checking should be finite, we only consider finite state and action spaces, and rational probability values.

Theorem 1. *Model checking of **PATLK** with deterministic ir-strategies is $\mathbf{\Delta_2^P}$-complete.*

Proof. We recall that model checking of epistemic logic is in $\mathbf{P}$ [Halpern and Vardi, 1991], and model checking **PATL** with deterministic ir-strategies is $\mathbf{\Delta_2^P}$-complete w.r.t. the size of the model and the length of the formula [Belardinelli *et al.*, 2023b]. Moreover, model checking **PATL** formulas of type $\langle\langle C\rangle\rangle\gamma$ with no nested strategic operators is $\mathbf{NP}$-complete.

The lower bound for **PATLK** is immediate from the fact that **PATLK** subsumes **PATL**. The upper bound can be obtained via the standard recursive algorithm that starts from the simplest subformulas (i.e., ones with no nested strategic nor epistemic modalities), computes their extensions (i.e., subsets of states that satisfy the formula), and replaces each of them with a fresh atomic proposition with exactly the same extension. Clearly, the complexity of the algorithm is $\mathbf{P}^{\mathbf{P}\cup\mathbf{NP}} = \mathbf{\Delta_2^P}$. □

The next results are proved analogously.

Theorem 2. *Model checking of $\mathbf{PATLC}_{ir}$ with deterministic ir-strategies is $\mathbf{\Delta_2^P}$-complete*

Proof. The lower bound for **PATLC** is immediate as **PATLC** subsumes **PATL**. For the upper bound, it suffices to prove that the extension of $K_a^{\propto q}\varphi$ in model M, for φ containing no nested modal operators, can be computed in polynomial time.

To do this, for each state $s \in St$, we compute $Bel(s, a, \varphi) = \sum_{t\in Sat(M,\varphi)} obs_a(t|s)$, and check if $Bel(a, s, \varphi) \propto q$. Clearly, the procedure runs in polynomial time w.r.t. the number of states, transitions, and representation of probabilities in the model and the formula.

To model-check an arbitrary formula of **PATLC**, we apply the same recursive algorithm as in Theorem 1. □

Theorem 3. *Model checking of **PATLH** with deterministic ir-strategies is $\mathbf{\Delta_2^P}$-complete.*

Proof. The lower bound is immediate as **PATLH** subsumes **PATL**. For the upper bound, we prove that the extension of $\mathcal{H}_a^{\propto m}\Phi$, for $\Phi = \{\varphi_1, \ldots, \varphi_n\}$ containing no nested modal operators, can be computed in polynomial time.

To do this, for each state $s \in St$, we construct the set of equivalence classes $R_{a,s}(\Phi)$ by checking the values of $\varphi_1, \ldots, \varphi_n$ in all states $t \in [s]_{\sim_a}$. Then, we compute $Pr_{a,s,\Phi}(x)$ for every $x \in R_{a,s}(\Phi)$. Finally, we calculate the Shannon entropy $H_S(R_{a,s}(\Phi))$ and compare it with m according to operator $\propto$. The procedure runs in polynomial time w.r.t. the number of states, transitions, representation of probabilities, and the precision with which logarithms are computed. For arbitrary formulas, we proceed recursively. □

5 Expressiveness

We start by recalling the semantic concepts of comparative expressiveness [Wang and Dechesne, 2009]. Then, we prove that **PATLC** is strictly more expressive than **PATLH**.

Definition 1 (Distinguishing and expressive power)**.** *Let $L_1 = (\mathcal{L}_1, \models_1)$ and $L_2 = (\mathcal{L}_2, \models_2)$ be two logical systems with sets of formulas $\mathcal{L}_1, \mathcal{L}_2$ and semantic relations $\models_1, \models_2$ interpreted over the same class of models $\mathcal{M}$. By $Sat(\phi) = \{(M, q) \mid M, q \models \phi\}$, we denote the class of pointed models that satisfy ϕ in the semantics given by $\models$. Likewise, $Sat(M, \phi) = \{q \mid M, q \models \phi\}$ is the set of states (or, equivalently, pointed models) that satisfy ϕ in a given structure M.*

L_2 is at least as expressive as L_1 ($L_1 \preceq_e L_2$) iff for every $\phi_1 \in \mathcal{L}_1$ there is $\phi_2 \in \mathcal{L}_2$ such that $Sat_{L_1}(\phi_1) = Sat_{L_2}(\phi_2)$.

L_2 is at least as distinguishing as L_1 ($L_1 \preceq_d L_2$) iff for every model M and formula $\phi_1 \in \mathcal{L}_1$ there exists $\phi_2 \in \mathcal{L}_2$ such that $Sat_{L_1}(\phi_1, M) = Sat_{L_2}(\phi_2, M)$. Equivalently: every pair of pointed models that can be distinguished by some $\phi_1 \in \mathcal{L}_1$ can be also distinguished by some $\phi_2 \in \mathcal{L}_2$.

Note that $L_1 \preceq_e L_2$ implies $L_1 \preceq_d L_2$. By transposition, it also holds that $L_1 \not\preceq_d L_2$ implies $L_1 \not\preceq_e L_2$.

Theorem 4. ***PATLH** covers neither the expressive nor the distinguishing power of **PATLC**.*

Proof. We show that **PATLC** $\not\preceq_d$ **PATLH**, and thus also **PATLC** $\not\preceq_e$ **PATLH**.

Take two models M_1, M_2, each with a single agent a, two states s_1, s_2, only self-loops as transitions, and a sole proposition p holding only in s_2. The observations in (M_1, s_1) are:

$obs(s_1|s_1) = \frac{1}{4}$, $obs(s_2|s_1) = \frac{3}{4}$; in (M_2, s_1): $obs(s_1|s_1) = \frac{3}{4}$, $obs(s_2|s_1) = \frac{1}{4}$. Notice that: (1) (M_1, s_1) and (M_2, s_1) satisfy the same formulas of **PATL** (obvious), and the only nontrivial uncertainty formulas are $\mathcal{H}_a^{\propto q}\{p\}$ (or, equivalently, $\mathcal{H}_a^{\propto q}\{\neg p\}$. But the Shannon entropy for $\Phi = \{p\}$ is the same in (M_1, s_1) and (M_2, s_1). So, they satisfy exactly the same formulas of **PATLH**. (2) The models can be discerned by the **PATLC** formula $\mathcal{K}_a^{\geq 3/4} p$, which holds in (M_1, s_1) but not in (M_2, s_1). Hence, we have that **PATLC** $\not\preceq_d$ **PATLH**. □

Theorem 5. ***PATLC*** *covers both the expressive and the distinguishing power of* ***PATLH****.*

Proof. We show a translation of formulas $\mathcal{H}_a^{\propto q}\{\Phi\}$ into **PATLC** that preserves the set of satisfying pointed models.

Let $\Phi = \{\varphi_1, \dots, \varphi_n\}$, and let V_Φ be the set of all possible valuations for Φ (there are 2^n such valuations). We write $\varphi_i \in v$ iff $v(\varphi_i) = true$. Clearly, each $v \in V_\Phi$ can be characterized by a conjunction $\psi_v \equiv \big(\bigwedge_{\varphi_i \in v} \varphi_i\big) \wedge \big(\bigwedge_{\varphi_i \notin v} \neg\varphi_i\big)$. Moreover, given a pointed model (M, s), each $v \in V_\Phi$ corresponds to a different outcome $[t_v] \in R_{a,s}(\Phi)$, assigned probability $p_v = Pr_{a,s,\Phi}([t_v])$.[1]

Now, the condition for $M, s \models \mathcal{H}_a^{\propto q}\{\Phi\}$ can be written as $-\sum_{v \in V_\Phi} p_v \log p_v \ \propto \ q$ (*). The set of probability distributions $(p_v)_{v \in V_\Phi}$ satisfying (*) is a finite union of hypercubes $H_1, \dots, H_m \in [0,1]^{|V_\Phi|}$, obtained as the solution of inequality (*), together with constraints (**) $p_v \geq 0, v \in V_\Phi$, and (***) $\sum_{v \in V_\Phi} p_v = 1$. Moreover, each hypercube H_i can be characterized by a conjunction of conditions $p_v \propto_v^i q_v^i$, $v \in V_\Phi$,[2] and hence equivalently by $\bigwedge_{v \in V_\Phi} Pr_{a,s,\Phi}([t_v]) \propto_v^i q_v^i$ which is in turn equivalent to $M, s \models \bigwedge_{v \in V_\Phi} \mathcal{K}_a^{\propto_v^i q_v^i} \varphi_v$. Putting it all together, we obtain

$$M, s \models \mathcal{H}_a^{\propto q}\{\Phi\} \quad \text{iff} \quad M, s \models \bigvee_i^m \bigwedge_{v \in V_\Phi} \mathcal{K}_a^{\propto_v^i q_v^i} \varphi_v.$$

Thus, $\bigvee_i^m \bigwedge_{v \in V_\Phi} \mathcal{K}_a^{\propto_v^i q_v^i} \varphi_v$ is an extension-preserving translation of $\mathcal{H}_a^{\propto q}\{\Phi\}$, which concludes the proof. □

Corollary 1. ***PATLC*** *has strictly more expressive and distinguishing power than* ***PATLC****.*

So, formulas based on probabilistic knowledge allow for expressing a strictly larger class of properties than uncertainty-based ones. This is a big surprise, and makes probabilistic logics distinctly different from non probabilistic ones (where both kinds of modalities were equally expressive). On the other hand, note that the **PATLC** translation of $\mathcal{H}_a^{\propto q}\{\Phi\}$, used in Theorem 5, is quite complicated and at least exponentially longer than its **PATLH** counterpart. As we show in the next section, this is no coincidence.

6 Succinctness

The concept of *succinctness* focuses on whether there is a substantial difference in the length of encodings provided by logics L_1, L_2 for some scalable property [Stockmeyer, 1972; Wilke, 1999; Adler and Immerman, 2001]). In this section, we show that **PATLH** is exponentially more succinct than **PATLC**. To prove this, we follow [Tabatabaei and Jamroga, 2023] and use *formula size games (FSG)* introduced in [French *et al.*, 2013b]. Specifically, we demonstrate that there is a sequence of **PATLH** formulas $(\varphi_n)_{n \in \mathbb{N}}$ with length $O(n)$, such that any **PATLC** formula ψ_n with the extension as φ_n has a parse tree with at least 2^n distinct vertices, and thus that the length of ψ_n must be at least $O(2^n)$.[3]

Definition 2 (Succinctness, [Adler and Immerman, 2001])**.** *Let $L_1 = (\mathcal{L}_1, \models_1)$ and $L_2 = (\mathcal{L}_2, \models_2)$ be two logical systems as in Definition 1. Further, suppose $f, g : \mathbb{N} \to \mathbb{N}$ are two functions such that $f(n) = O(g(n))$ is a strictly increasing function.*

L_1 is exponentially more succinct than L_2 ($L_1 \not\preceq_M^{subexp} L_2$) iff for each $n \in \mathbb{N}$ there are formulas $\varphi_n \in \mathcal{L}_1$ and $\psi_n \in \mathcal{L}_2$ with: (1) $|\varphi_n| = f(n)$, (2) $|\psi_n| = 2^{g(n)}$, (3) ψ_n is the shortest formula in $\mathcal{L}_2$, equivalent to φ_n.

We now adapt the one-person formula size games of [French *et al.*, 2013a], where the player, called the spoiler tries to synthesize a formula to discern between sets of models A and B, i.e., some $\varphi \in$ **PATLC** such that $A \models \varphi$ and $B \models \neg\varphi$.

Definition 3 (FSG for Credences)**.** *The game is defined for two given sets of pointed models $A, B \subseteq \hat{M}$, and played according to the following rules. The player, called* the spoiler*, constructs a game tree in such a way that each vertex is labeled with a pair (C, D) of subsets of pointed models. The possible moves for the spoiler on each vertex of the tree are $\{\{p\}_{p \in AP}, \neg, \vee, \{\mathcal{K}_a^{\propto q}\}_{a \in \mathbb{A}\text{gt}; q \in \mathbb{Q} \cap [0,1]}\}$. A vertex can be open or closed, and once it becomes closed, no further move can be played from it. The rules for each of possible moves are defined as follows:*

1. *(Atomic move ($p \in AP$)): the spoiler chooses $p \in AP$ s.t. $C \models p$ and $D \models \neg p$. The vertex becomes closed.*
2. *($\neg$ move): A new vertex (D, C) is added to the tree.*
3. *($\vee$ move): two vertices (C_1, D) and (C_2, D) are added to the tree, where $C_1 \cup C_2 = C$ are chosen by the spoiler.*
4. *($\mathcal{K}_a^{\propto q}$ move): For each $(M, s) \in C$ the spoiler chooses a subset of pointed models $C_s \subseteq \hat{M}$ such that $obs_a(C_s|s) \propto q$. Moreover, for each $(M, s) \in D$, the spoiler chooses a subset of pointed models $D_s \subseteq \hat{M}$ such that $obs_a(C_s|s) \widehat{\propto} 1 - q$, where $\widehat{=}$ denotes $\neq$, $\widehat{\geq}$ is $>$, $\widehat{\leq}$ is $<$, and vice versa. Then a new vertex $(\bigcup_{s \in C} C_s, \bigcup_{s \in D} D_s)$ is added to the tree.*

We say that the spoiler wins the FSG starting at (A, B) in n moves iff there exists a game tree T with the root (A, B) and precisely n vertices such that every leaf of T is closed.

We can now state the following result, which is an adaptation of a classical theorem in [French *et al.*, 2013b] to probabilistic modal operators.

[1] Possibly with probability $p_v = 0$.

[2] Note that the comparison operators $\propto_v^i$ do not have to be the same, and in particular can be different from the operator $\propto$ in the translated uncertainty formula.

[3] The length of formula φ is defined as the total number of symbols in φ, including atomic propositions, boolean, modal and arithmetic operators, and numbers (treated each as a single symbol).

Theorem 6. *Spoiler can win FSG starting at (A, B) in n moves iff there exists a **PATLC** formula φ (without strategic operators) where $A \models \varphi$, $B \models \neg\varphi$, and $|\varphi| = n$.*

Proof. The proof is analogous to [French *et al.*, 2013b, Theorem 3.2], using credence moves instead of epistemic ones. Each credence move corresponds to a $K_a^{\propto q}$ operator in φ. □

The purpose of Theorem 6 is to use Formula Size Games for proving that **PATLH** is exponentially more succinct than **PATLC**. Analogous results were presented in [French *et al.*, 2013b] for mutual knowledge vs. standard epistemic logic, and in [Tabatabaei and Jamroga, 2023] for Hartley uncertainty vs. standard knowledge. Unfortunately, those proofs cannot be easily adapted to the case of Shannon uncertainty vs. probabilistic beliefs. We are still convinced that the analogous relationship holds, but for now we only formulate it as a conjecture, and leave its ultimate proof for future work.

Conjecture 1. ***PATLH*** $\not\trianglerighteq^{subexp}$ ***PATLC***.

Proof idea. For every natural number $n \in \mathbb{N}$, we define a formula φ_n in the language of **PATLH** and consider the shortest formula ψ_n in the language of **PATLC** that is expressively equivalent to φ_n. The existence of such ψ_n is guaranteed by Theorem 5. Further, we construct two sets of pointed models A_n and B_n, such that $A_n \models_{\mathcal{H}} \varphi_n$ and $B_n \models_{\mathcal{H}} \neg\varphi_n$,

The idea is to use $A_n = \{M\}$ with two agents $\{a, b\}$ and the state space divided into n "layers," where obs_a is defined within each layer according to the geometric probability distribution (a.k.a. Furry distribution), and obs_b leads to states in the subsequent layer. Moreover, B_n is the set of possible variants of M obtained by removal of an arbitrary state within each layer. Then, we take: $\phi_1 \equiv \mathcal{H}_a^{\geq q}\{\mathsf{p}_1, \ldots \mathsf{p}_k\}$, $\phi_2 \equiv \mathcal{H}_a^{\geq q}\{\mathcal{H}_b^{=0}\mathcal{H}_a^{\geq q}\{\mathsf{p}_1, \ldots \mathsf{p}_k\}, \ldots \mathsf{p}_k\}$, and so on, for an appropriately chosen q.

We will thus have that $A_n \models \psi_n$ and $B_n \models \neg\psi_n$. Hence, in the FSG with the starting point at the vertex labelled by (A_n, B_n), the spoiler has a winning strategy given by playing the moves in accordance with the syntactic structure of ψ_n. Still, the last step remains to demonstrate that any such winning strategy requires that the length of the play of the game (starting at (A_n, B_n)) is no shorter than 2^n. □

7 Case Studies

Scenario 1: Probabilistic Analysis of Voting ThreeBallot [Rivest, 2006], is an end-to-end verifiable voting protocol designed to ensure both vote privacy and coercion resistance. In this protocol, each voter fills out three ballots per issue. To cast a **'yes'** vote for an issue, the voter fills two of the three corresponding fields, leaving one blank. Conversely, for a **'no'** vote, the voter fills one field and leaves the other two blank. The voter then separates the three columns to create three individual ballots, retains one as a receipt, and submits the rest to the ballot box. After tallying, all ballots are published on a public bulletin board for verification.

A key property of ThreeBallot is *coercion resistance*. A coercer, even with access to a voter's receipt, cannot deduce the voter's complete voting choice due to the indistinguishability sets formed by the combinations of filled and blank spaces.

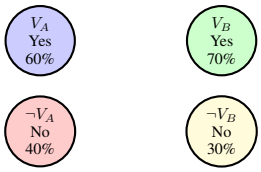

Figure 1: Opinion poll probabilities for the two-issue referendum. Each circle represents a possible vote outcome with its probability.

An example ThreeBallot configuration for a two-issue referendum is shown in Table 1. The voter votes 'yes' on Issue A and 'no' on Issue B, resulting in the following ballots:

- **Ballot 1**: Filled for Issue A, blank for Issue B.
- **Ballot 2**: Filled for both Issues A and B.
- **Ballot 3**: Blank for Issue A, filled for Issue B.

Issue A	Issue B	Resulting Ballots
×	□	Ballot 1: $\{F, B\}$
×	×	Ballot 2: $\{F, F\}$
□	×	Ballot 3: $\{B, F\}$

Table 1: Example of a ThreeBallot vote. Here, 'F' denotes a filled space and 'B' a blank space. The voter votes 'yes' on Issue A and 'no' on Issue B.

In contrast to [Tabatabaei and Jamroga, 2023], we consider a probabilistic model of ThreeBallot, with the probabilities coming from opinion polls about voters' preferences. Based on that, we will demonstrate how **PATLH** and **PATLC** capture coercion resistance in stochastic models of voting.

Probabilistic model of voting. Consider a two-issue referendum with propositions A and B. The voters' preferences are influenced by opinion polls, which provide the following probabilities for the outcomes: $p(V_A) = 0.6$, $p(\neg V_A) = 0.4$, $p(V_B) = 0.7$, $p(\neg V_B) = 0.3$. Here, V_A and V_B represent votes in favor of A and B, while $\neg V_A$ and $\neg V_B$ represent votes against them. The ballots are filled following the ThreeBallot rules: voters fill two fields for a "yes" vote and one for a "no" vote for each issue. After filling the ballots, voters retain one as a receipt and submit the others to the ballot box.

How to specify coercion resistance. The coercer attempts to deduce the voter's choices using the receipt and published data. To ensure coercion resistance, we evaluate the uncertainty of the coercer about the voter's choices using **PATLH** and **PATLC**. In what follows we use $(V_1 = V_i)$ in as an atomic proposition that is evaluated as true if V_1 is equal to V_i. Condition $(V_1 \neq V_2)$ is interpreted in an analogous manner.

The coercion resistance property can be expressed in **PATLH** as

$$\bigwedge_{V_i \in Votes} \neg\langle\!\langle v, c\rangle\!\rangle\Diamond(V_1 = V_i \wedge V_2 \neq V_1 \rightarrow H_c^{<1.8}\{V_A, V_B\}),$$

where $Votes = \{V_A V_B, V_A \neg V_B, \neg V_A V_B, \neg V_A \neg V_B\}$ represents all possible combinations of votes. Here, H_c quantifies the coercer's entropy regarding the voter's choices. By maintaining an entropy of at least 1.8 bits, the coercer's maximal

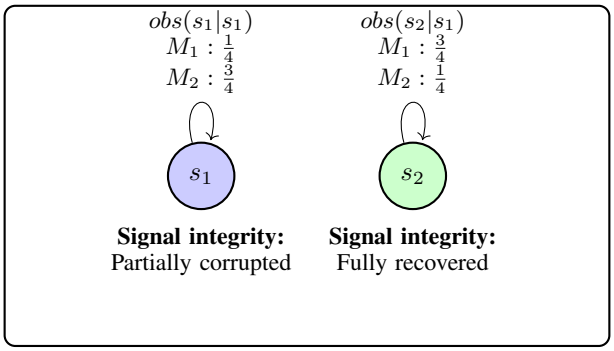

Figure 2: Signal integrity: states and observation probabilities

uncertainty about the votes is guaranteed. This formula ensures that no strategy allows the coercer to achieve a lower level of uncertainty.

Following the translation from Theorem 5, we can express the same property in **PATLC** as:

$$\bigwedge_{V_i \in Votes} \neg\langle\langle v, c\rangle\rangle\Diamond\big(V_1 = V_i \wedge V_1 \neq V_2 \rightarrow \bigvee_{j}^{m} \bigwedge_{v \in V_\Phi} \mathcal{K}_c^{\propto_v^j 1.8_v^j}\varphi_v\big),$$

where $\Phi = \{V_A, V_B\}$, V_Φ is the set of valuations for Φ, and 1.8_v^j are the numbers corresponding to the hypercubes in $[0,1]^{|V_\Phi|}$, as in the proof of Theorem 5. This formula ensures that no strategy allows the coercer to achieve a degree of belief exceeding $\max\{1.8_v^j : v \in V_\Phi, j = 1, \ldots, m\}$ for any specific vote combination. Clearly, using **PATLH**, we express the same property more succinctly.

Analysis and comparison. The **PATLH** representation leverages entropy to directly capture the coercer's uncertainty, avoiding the need to enumerate all vote combinations explicitly. This succinctness highlights **PATLH**'s advantage in scenarios with probabilistic uncertainty.

In contrast, **PATLC** relies on probabilistic beliefs, which, while expressive, require detailed enumeration of strategies and outcomes. Both approaches demonstrate coercion resistance, but **PATLH** provides a more compact formalization.

Scenario 2: Signal Integrity To illustrate the differences between the expressive power of **PATLC** and **PATLH**, consider a story of a cybersecurity analyst, Alice, investigating potential vulnerabilities in a communication system. Alice is analyzing two models of data transmission security, M_1 and M_2, to assess their ability to handle interference. Both models describe a device operating in two states: s_1 represents a state where data is partially corrupted but recoverable, and s_2 represents a state of complete recovery. The device transmits a key signal σ, which holds specific integrity information: in s_2, the signal σ is fully verified as correct (p_σ), while in s_1, it is not ($\neg p_\sigma$).

The device observations of its signal integrity are modeled as follows. In M_1, the probabilities are $obs(s_1|s_1) = \frac{1}{4}$ and $obs(s_2|s_1) = \frac{3}{4}$; in M_2, the probabilities are $obs(s_1|s_1) = \frac{3}{4}$ and $obs(s_2|s_1) = \frac{1}{4}$, see also Figure 2. We emphasize that agent a in the model represents the entity making observations about the signal integrity and reasoning about its states, i.e., the device. Alice is an external entity using **PATLH** and **PATLC** to evaluate the properties of the system.

Alice first uses **PATLH** to evaluate the system's uncertainty. By calculating the Shannon entropy for $\Phi = \{p_\sigma\}$, she finds that the entropy is identical in both models, and thus they satisfy the same formulas of **PATLH**.

However, Alice may turn to **PATLC** to gain a more granular understanding. In particular, she can check formula $\mathcal{K}_a^{\geq 3/4} p_\sigma$, which is satisfied in (M_1, s_1) but not in (M_2, s_1). This allows Alice to identify M_1 as a system where the device has higher probabilistic confidence in its recovery signal, a critical feature in detecting and managing interference.

8 Summary

This work introduces **PATLH**, a groundbreaking extension of probabilistic alternating-time temporal logic. **PATLH** incorporates Shannon entropy to provide a formal framework for reasoning about uncertainty and information gain. By combining strategic reasoning with a measure of probabilistic uncertainty, **PATLH** fills a significant gap in the study of stochastic MAS where agents operate under imperfect information. The logic enables a nuanced approach to both reducing and maintaining uncertainty, addressing real-world needs in scenarios like privacy protection and information security, where agents aim to control the flow of sensitive information.

The introduction of entropy-based reasoning is particularly innovative, as it allows for the modeling of agents' strategic capabilities in controlling uncertainty, something that previous epistemic logics have struggled to express succinctly. **PATLH** 's application to security and privacy concerns in multi-agent systems sets a new standard for addressing these challenges, providing a richer, more flexible framework for researchers and practitioners alike. Moreover, the work also extends the capabilities of existing probabilistic logics by introducing **PATLC**, which facilitates reasoning about probabilistic beliefs (credences) instead of full knowledge, and turns out to be strictly more expressive than **PATLH**.

Future work. First and foremost, we plan to complete the proof of the succinctness claim. As for longer-term plans, the study opens the door to further extensions in strategic logics. One particularly promising avenue for future research is extending **PATLH** to more expressive frameworks, such as Strategy Logic (SL, see e.g., [Berthon *et al.*, 2017]). These extensions would enable reasoning about complex solution concepts like Nash equilibria, subgame-perfect equilibria, and other cooperative or adversarial strategies in multi-agent systems. SL, which allows explicit quantification over strategies, would benefit from the integration of entropy-based reasoning, enabling the representation of complex objectives related to information control and uncertainty.

Acknowledgments

M.T. Godziszewski and A. Murano's research was supported by the PRIN project RIPER (No. 20203FFYLK), and the PNRR MUR projects PE0000013-FAIR and INFANT (No. E23C24000390006). W. Jamroga's research was supported by NCBR Poland under the project POLLUX-XI/14/SpaceVote/2023.

References

[Adler and Immerman, 2001] Micah Adler and Neil Immerman. An n! lower bound on formula size. In *Proceedings of IEEE Symposium on Logic in Computer Science*, pages 197–206, 2001.

[Ågotnes and van Ditmarsch, 2008] T. Ågotnes and H. van Ditmarsch. Coalitions and announcements. In *Proceedings of International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 673–680, 2008.

[Ågotnes *et al.*, 2015] T. Ågotnes, V. Goranko, W. Jamroga, and M. Wooldridge. Knowledge and ability. In *Handbook of Epistemic Logic*, pages 543–589. 2015.

[Ågotnes, 2006] T. Ågotnes. Action and knowledge in alternating-time temporal logic. *Synthese*, 149(2):377–409, 2006.

[Alur *et al.*, 2002] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time Temporal Logic. *Journal of the ACM*, 49:672–713, 2002.

[Alvim *et al.*, 2020] Mário S. Alvim, Konstantinos Chatzikokolakis, Annabelle McIver, Carroll Morgan, Catuscia Palamidessi, and Geoffrey Smith. *The Science of Quantitative Information Flow*. Information Security and Cryptography. Springer, 2020.

[Belardinelli *et al.*, 2023a] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. Strategic abilities of forgetful agents in stochastic environments. In *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR*, pages 726–731, 2023.

[Belardinelli *et al.*, 2023b] Francesco Belardinelli, Wojciech Jamroga, Munyque Mittelmann, and Aniello Murano. Strategic abilities of forgetful agents in stochastic environments. In Pierre Marquis, Tran Cao Son, and Gabriele Kern-Isberner, editors, *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning, KR 2023, Rhodes, Greece, September 2-8, 2023*, pages 726–731, 2023.

[Belardinelli *et al.*, 2024] Francesco Belardinelli, Wojtek Jamroga, Munyque Mittelmann, and Aniello Murano. Verification of stochastic multi-agent systems with forgetful strategies. In *Proceedings of AAMAS*, pages 160–169. IFAAMAS/ACM, 2024.

[Berthon *et al.*, 2017] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Y. Vardi. Strategy logic with imperfect information. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12, 2017.

[Bozzelli *et al.*, 2020] Laura Bozzelli, Aniello Murano, and Loredana Sorrentino. Alternating-time temporal logics with linear past. *Theor. Comput. Sci.*, 813:199–217, 2020.

[Bulling *et al.*, 2015] N. Bulling, V. Goranko, and W. Jamroga. Logics for reasoning about strategic abilities in multi-player games. In *Models of Strategic Reasoning. Logics, Games, and Communities*, volume 8972 of *Lecture Notes in Computer Science*, pages 93–136. 2015.

[Chen and Lu, 2007] Taolue Chen and Jian Lu. Probabilistic alternating-time temporal logic and model checking algorithm. In *Proceedings of FSKD, Volume 2*, pages 35–39, 2007.

[French *et al.*, 2013a] Tim French, Wiebe Van Der Hoek, Petar Iliev, and Barteld Kooi. On the succinctness of some modal logics. *Artificial Intelligence*, 197:56–85, 2013.

[French *et al.*, 2013b] Tim French, Wiebe van der Hoek, Petar Iliev, and Barteld P. Kooi. On the succinctness of some modal logics. *Artificial Intelligence*, 197:56–85, 2013.

[Halpern and Vardi, 1991] Joseph Y Halpern and Moshe Y Vardi. Model checking vs. theorem proving: a manifesto. *Artificial intelligence and mathematical theory of computation*, 212:151–176, 1991.

[Hansson and Jonsson, 1994] H. Hansson and B. Jonsson. A logic for reasoning about time and reliability. *Formal Aspects of Computing*, 6(5):512–535, 1994.

[Hartley, 1928a] R.V.L. Hartley. Transmission of information. *Bell Systems Technical Journal*, 7:535, 1928.

[Hartley, 1928b] R.V.L. Hartley. Transmission of information. *The Bell System Technical Journal*, 7(3):535–563, 1928.

[Huang and Luo, 2013] Xiaowei Huang and Cheng Luo. A logic of probabilistic knowledge and strategy. In *Proceedings of AAMAS*, pages 845–852. IFAAMAS, 2013.

[Huang *et al.*, 2012] X. Huang, K. Su, and C. Zhang. Probabilistic alternating-time temporal logic of incomplete information and synchronous perfect recall. In *Proceedings of AAAI-12*, 2012.

[Jamroga and Ågotnes, 2007] W. Jamroga and T. Ågotnes. Constructive knowledge: What agents can achieve under incomplete information. *Journal of Applied Non-Classical Logics*, 17(4):423–475, 2007.

[Jamroga and Tabatabaei, 2013] W. Jamroga and M. Tabatabaei. Accumulative knowledge under bounded resources. In *Proceedings of CLIMA XIV*, volume 8143 of *Lecture Notes in Computer Science*, pages 206–222, 2013.

[Markey, 2003] Nicolas Markey. Temporal logic with past is exponentially more succinct. *Bull. EATCS*, 79:122–128, 2003.

[Rivest, 2006] R. Rivest. The threeballot voting system. Available online at http://theory.csail.mit.edu/~rivest/Rivest-TheThreeBallotVotingSystem.pdf, 2006.

[Shannon, 1948] Claude E. Shannon. A mathematical theory of communication. *Bell Syst. Tech. J.*, 27(3):379–423, 1948.

[Stockmeyer, 1972] Larry J. Stockmeyer. *The Complexity of Decision Problems in Automata Theory and Logic*. PhD thesis, Massachusetts Institute of Technology, 1972.

[Tabatabaei and Jamroga, 2023] Masoud Tabatabaei and Wojciech Jamroga. Playing to learn, or to keep secret: Alternating-time logic meets information theory. In Noa Agmon, Bo An, Alessandro Ricci, and William Yeoh, editors, *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2023, London, United Kingdom, 29 May 2023 - 2 June 2023*, pages 766–774. ACM, 2023.

[van der Hoek and Wooldridge, 2003] W. van der Hoek and M. Wooldridge. Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications. *Studia Logica*, 75(1):125–157, 2003.

[van Ditmarsch *et al.*, 2007] H. van Ditmarsch, W. van der Hoek, and B. Kooi. *Dynamic Epistemic Logic*. 2007.

[Wang and Dechesne, 2009] Y. Wang and F. Dechesne. On expressive power and class invariance. *CoRR*, abs/0905.4332, 2009.

[Wilke, 1999] Thomas Wilke. CTL+ is exponentially more succinct than CTL. In *Proceedings of FSTTCS*, volume 1738 of *Lecture Notes in Computer Science*, pages 110–121, 1999.