

Complete ω -Regular Supermartingale Certificates

Alessandro Abate 

University of Oxford, UK

Mirco Giacobbe 

University of Birmingham, UK

Sergey Ichtchenko 

University of Oxford, UK

Diptarko Roy 

University of Birmingham, UK

Abstract

We introduce a general methodology for the construction of sound and complete proof rules for the almost-sure and quantitative acceptance of reactivity properties on time-homogeneous Markov chains with general state spaces. Reactivity captures ω -regular properties and subsumes linear temporal logic. Our core technical result establishes that every reactivity property admits decomposition into multiple obligations of almost-sure termination into absorbing regions, and that appropriate absorbing regions always exist on general state spaces. This enables the extension of every complete proof rule for almost-sure termination into a proof rule for reactivity that is complete in the almost-sure case, and complete up to an arbitrarily small ε -approximation in the quantitative case. We apply our new methodology to recent results on sound and complete supermartingale certificates for almost-sure termination in the special case of countably infinite state spaces, alongside standard results on quantitative safety. As a result, we obtain the first sound and complete supermartingale certificates for almost-sure ω -regular properties and the first sound and ε -complete supermartingale certificates for quantitative ω -regular properties on time-homogeneous Markov chains with countably infinite state spaces.

2012 ACM Subject Classification Theory of computation $\rightarrow$ Logic and verification; Theory of computation $\rightarrow$ Probabilistic computation

Keywords and phrases Probabilistic Model Checking, Markov Chains on Measurable State Spaces, Omega-Regular Properties, Martingale Theory

Digital Object Identifier 10.4230/LIPIcs.LICS.2026.3

Funding This work was supported in part by the Advanced Research and Invention Agency under the Safeguarded AI programme.

1 Introduction

Supermartingale certificates are essential components of widely used proof rules for the formal verification of probabilistic systems whose semantics correspond to time-homogeneous Markov chains over general state spaces. Examples include probabilistic programs used to model randomised algorithms and inference protocols, which typically operate over countably infinite state spaces, as well as stochastic difference equations used to model discrete-time dynamical systems, which usually evolve over continuous state spaces. In many real-world domains, the finite-state assumptions underlying traditional probabilistic model-checking algorithms are inappropriate, making it crucial to model and reason about probabilistic systems over general measurable state spaces.

Deductive approaches based on martingale theory for reasoning about probabilistic systems bring principles from stochastic analysis into program analysis and control theory. Supermartingales are real-valued stochastic processes whose values satisfy specific convergence



© Alessandro Abate, Mirco Giacobbe, Sergey Ichtchenko, and Diptarko Roy;
licensed under Creative Commons License CC-BY 4.0
41st Annual Symposium on Logic in Computer Science (LICS 2026).

Editors: Claudia Faggian and Joost-Pieter Katoen; Article No. 3; pp. 3:1–3:28



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

properties; supermartingale certificates map a system to a supermartingale and are essentially the probabilistic analogues of ranking functions and inductive invariants in Hoare logic, as well as Lyapunov functions in stability analysis [11, 38]. In this paper, we bring together stochastic analysis and formal language theory, providing the technical elements for the construction of sound and complete supermartingale certificates for reactivity properties.

Reactivity is the temporal property stating that either an event A occurs only finitely many times or an event B occurs infinitely often or, more generally, it is a finite conjunction of such properties. It is a standard result that reactivity characterises the ω -regular languages, thereby subsuming linear temporal logic and generalising properties such as recurrence, persistence, safety (or invariance), and termination (or reachability) [41, 55]. Proof rules based on supermartingale certificates for the almost-sure and quantitative acceptance of these classes of properties – including reactivity – have been proposed previously [3, 43, 60]; however, for reactivity, these proof rules provide only sufficient conditions, meaning they may fail to certify systems that nonetheless satisfy the property.

The completeness of a proof rule involving supermartingale certificates concerns whether *every* stochastic process that satisfies the corresponding temporal property admits a certificate – typically expressed as a value function – that fulfils the associated proof rule. Complete proof rules for reactivity [37, Definition 4.4] have been presented recently, although under the assumption that value functions may express arbitrary distributions. While this enables reasoning about the distribution of the number of visits to states, it also means that it is hard to automate the construction of such proof certificates [37, p. 11]. In fact, the existing certificate construction methods require that proof certificates are scalar-valued functions. In this work, we focus on the setting of certificates that are (i) complete and (ii) scalar-valued. Prior to the present work, proof rules satisfying both criteria were known only for almost-sure termination [40, 43]; completeness for reactivity and therefore ω -regular properties remained an open problem.

We introduce a general methodology for characterising sound and complete proof rules for reactivity properties. Our core technical result shows that every reactivity property admits a decomposition into (i) multiple obligations of qualitative safety for each *absorbing region*, (ii) multiple obligations of almost-sure termination into the absorbing regions, and (iii) one obligation of quantitative safety for the overall system. In other words, our methodology reduces the problem of establishing almost-sure or quantitative reactivity to a collection of almost-sure termination and qualitative or quantitative safety problems. Crucially, we show that every time-homogeneous Markov chain on a general state space that satisfies a reactivity property admits suitable absorbing regions that may be used to approximate the probability of acceptance up to any arbitrary ε -approximation. Moreover, when the reactivity property holds almost surely, we show that there exist suitable absorbing regions that exactly characterise almost-sure acceptance.

Our result introduces the first sound and complete proof rules yielding supermartingale certificates for reactivity, and thus for ω -regular properties. We present two such proof rules, obtained by applying our methodology to complete proof rules for almost-sure termination on the special case of countably infinite state spaces [40, 43], alongside standard results on quantitative safety [38, 60]. As an immediate consequence of the completeness results underlying our absorbing-region decomposition, these new proof rules are complete on countably infinite state spaces up to arbitrary ε -approximation for quantitative reactivity, and fully complete when the reactivity property holds almost surely.

Our contribution is twofold. First, we show that every reactivity property admits a decomposition into multiple obligations of almost-sure termination and quantitative safety with respect to absorbing regions. Our decomposition is ε -complete in the general case and

fully complete in the almost-sure case for any time-homogeneous Markov chain on a general state space. Second, we instantiate our decomposition using existing complete proof rules for almost-sure termination over countably infinite state spaces. As a result, we obtain the first supermartingale certificates that are ε -complete in the quantitative case and fully complete in the almost-sure case for ω -regular properties on countably infinite state spaces.

2 Markov Chains and Probabilistic Reactivity

We consider stochastic processes on a general state space (S, Σ) , where S denotes the set of states and the associated σ -algebra Σ denotes the set of measurable regions. We treat the problem of determining the acceptance of infinite time horizon properties measured over $(\Omega, \mathcal{F})$, where the set of outcomes $\Omega = S^\omega$ are the infinite trajectories of a process and the set of events $\mathcal{F} = \bigotimes_{i \in \omega} \Sigma_i$ (with $\Sigma_i = \Sigma$) are the measurable properties [45, Chapter 3.4]. A filtration $\mathfrak{F} = (\mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_n, \dots)$ is an increasing sequence of sub- σ -algebras of $\mathcal{F}$ where every $L \in \mathcal{F}_n$ is an event whose membership is determined by the prefix of length n . A sequence of random variables $\Phi = (\Phi_0, \Phi_1, \dots, \Phi_n, \dots)$ is a stochastic process adapted to $\mathfrak{F}$ precisely if Φ_n is $\mathcal{F}_n$ -measurable for every $n \in \mathbb{N}$.

Throughout the paper, we use the notation $\mathbf{1}\{\pi\} : S^\omega \rightarrow \{0, 1\}$ for the indicator function of the proposition π , we use $\{\pi\}$ to denote the subset of S^ω consisting of trajectories that satisfy proposition π , and use $\delta_s : \Sigma \rightarrow [0, 1]$ for the Dirac measure $\delta_s(A) = \mathbf{1}\{s \in A\}$ at $s \in S$. We say that two events are *almost-surely equivalent* under a probability measure precisely if their symmetric difference is a probability zero event.

We restrict our attention to time-homogeneous Markov chains on general state spaces. As is standard in stochastic analysis, we rely on the result that an initial probability measure and a transition probability kernel define a unique probability measure over the measurable space of infinite trajectories and properties.

► **Theorem 2.1** ([45, Theorem 3.4.1]). *Let $\mu : \Sigma \rightarrow [0, 1]$ be an initial probability measure and $P : S \times \Sigma \rightarrow [0, 1]$ be a transition probability kernel. Then, there exists a stochastic process $\Phi = (\Phi_0, \Phi_1, \dots)$ adapted to $\mathfrak{F}$ and a probability measure $P_\mu : \mathcal{F} \rightarrow [0, 1]$ where $P_\mu(\Phi \in L)$ is the probability that Φ satisfies the property $L \in \mathcal{F}$ and, for every $n \in \mathbb{N}$ and $A_0 \in \Sigma, \dots, A_n \in \Sigma$,*


$$P_\mu(\Phi_0 \in A_0 \wedge \dots \wedge \Phi_n \in A_n) = \int_{s_0 \in A_0} \dots \int_{s_{n-1} \in A_{n-1}} \mu(ds_0) P(s_0, ds_1) \dots P(s_{n-1}, A_n). \quad (1)$$

► **Definition 2.2** (Time-Homogeneous Markov Chain). *A time-homogeneous Markov chain defined by an initial probability measure $\mu : \Sigma \rightarrow [0, 1]$ and the transition probability kernel $P : S \times \Sigma \rightarrow [0, 1]$ is a stochastic process satisfying Equation (1).*

We consider ω -regular properties which we define in terms of reactivity, which in turn we define in terms of visiting a region A finitely or infinitely many times, respectively, as follows:

$$\text{Fin}(A) = \bigcup_{n \in \mathbb{N}} \bigcap_{m \geq n} \{\Phi_m \notin A\}, \quad \text{Inf}(A) = \bigcap_{n \in \mathbb{N}} \bigcup_{m \geq n} \{\Phi_m \in A\}. \quad (2)$$

Using these components, we define reactivity as an intersection of events of the form $\text{Fin}(A_i) \cup \text{Inf}(B_i)$ for $i = 1, \dots, k$ on a finite sequence of pairs $(A_1, B_1), \dots, (A_k, B_k) \in \Sigma^2$, which we refer to as a Streett acceptance condition.



$$\begin{aligned} \text{Profit} &= \{w : w > 0\} \\ \text{Solvency} &= \{w : w \geq 0\} \\ \text{Debt} &= \{w : w < 0\} \end{aligned}$$

■ **Figure 1** The Lending Casino.

► **Definition 2.3** (Reactivity). A reactivity property defined by the Streett acceptance condition $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ is the event

$$\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i). \quad (3)$$

► **Remark 2.4** (ω -Regular Properties). Reactivity conditions are sufficient to encode ω -regular properties. Every ω -regular property φ over a finite set of atomic propositions Π may be represented as the set of accepting runs of a corresponding deterministic Streett automaton (DSA) [55]. A DSA $\mathcal{A}$ consists of a finite set of states Q , an initial state $q_0 \in Q$, a transition function $T: Q \times 2^\Pi \rightarrow Q$, and an acceptance condition $(F_1, G_1), \dots, (F_k, G_k)$ where $F_i, G_i \subseteq Q$ for $i = 1, \dots, k$. An infinite input trace $(p_0, p_1, p_2, \dots) \in (2^\Pi)^\omega$ is accepted if there exists an infinite run $(q_0, q_1, q_2, \dots) \in Q^\omega$ such that $q_{n+1} = T(q_n, p_n)$ for every $n \in \mathbb{N}$ and, for every $i = 1, \dots, k$, either $\sum_{n=0}^\infty \mathbf{1}\{q_n \in F_i\} < \infty$ or $\sum_{n=0}^\infty \mathbf{1}\{q_n \in G_i\} = \infty$.

Given a time-homogeneous Markov chain $\hat{\Phi}$ defined by $(\hat{\mu}, \hat{P})$ with state space $(\hat{S}, \hat{\Sigma})$, we define the synchronous composition of $\hat{\Phi}$ and the automaton $\mathcal{A}$ as another time-homogeneous Markov chain Φ with state space $(\hat{S} \times Q, \hat{\Sigma} \otimes 2^Q)$, defined by the initial distribution $\mu(A) = \int_{s \in \hat{S}} \mathbf{1}\{(s, q_0) \in A\} \hat{\mu}(ds)$, and the transition probability kernel

$$P((s, q), A) = \int_{(s', q') \in A} \mathbf{1}\{T(q, \langle\langle s \rangle\rangle) = q'\} \hat{P}(s, ds'), \quad (4)$$

where $\langle\langle \cdot \rangle\rangle: \hat{S} \rightarrow 2^\Pi$ denotes the observation (or labelling) function of $\hat{\Phi}$, yielding the set of propositions $\langle\langle s \rangle\rangle$ that hold true in s . We then define the Streett acceptance condition $(A_1, B_1), \dots, (A_k, B_k)$ where $A_i = \hat{S} \times F_i$ and $B_i = \hat{S} \times G_i$ for $i = 1, \dots, k$. Then, this construction yields an equivalence between the probabilities of the original process $\hat{\Phi}$ satisfying the ω -regular property φ and the synchronous product Φ satisfying the derived reactivity property:

$$\mathbb{P}_{\hat{\mu}}(\hat{\Phi} \in \varphi) = \mathbb{P}_{\mu}\left(\Phi \in \bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i)\right). \quad (5)$$

For this reason, henceforth we consider reactivity properties, without loss of generality for ω -regular properties. $\lrcorner$

► **Example 2.5** (The Lending Casino, Figure 1). A casino decides on the following business strategy. Each gambler is provided \$1 of initial wealth. A fair coin is tossed, and the gambler's wealth either increases or decreases by \$1 depending upon the outcome of the toss.

If the gambler's wealth becomes negative, then they are given the option of borrowing from the casino, but while they are in debt, the coin is biased such that the probability of winning \$1 is $\frac{1-\varepsilon}{2}$ for some $\varepsilon > 0$. Notably, there is no upper limit to either the gambler's wealth or the size of their overdraft, and there is no upper bound to the duration of the game.

The casino wishes to answer the following question regarding infinite trajectories: *will a gambler eventually fall into debt and remain there forever?* This corresponds to the requirement that the property $\text{Inf}(\text{Solvency})$ holds with probability zero, or equivalently $\text{Fin}(\text{Solvency})$ holds with probability one. This is the reactivity property defined by the Streett acceptance condition given by the single pair $(\text{Solvency}, \emptyset)$. $\lrcorner$

A stopping time $\zeta : \Omega \rightarrow \mathbb{N} \cup \{\infty\}$ is a random variable for which the event $\{\zeta = n\}$ is an element of $\mathcal{F}_n$, for every $n \in \mathbb{N}$. A special case of a stopping time is the first hitting time of a region, which is key to defining the notion of return times in Section 3 and of termination in Section 4.

► **Definition 2.6** (First Hitting Time). *The first hitting time τ_A of a region $A \in \Sigma$ by the stochastic process Φ is*

$$\tau_A = \inf\{n \geq 0 : \Phi_n \in A\}. \quad (6)$$

► **Example 2.7** (Profit Leads to Debt, and Debt Doesn't Always Lead to Solvency). We use the notion of first hitting time to describe the event that the gambler hits *Debt* as $\tau_{\text{Debt}} < \infty$. It is a classic result [26, p.182] that a symmetric random walk on the integers visits every integer with probability 1, which implies that from any initial positive wealth $w \in \text{Profit}$ that the gambler eventually enters debt:

$$\forall w \in \text{Profit}. P_{\delta_w}(\tau_{\text{Debt}} < \infty) = 1, \quad (7)$$

and $\forall w \in \text{Profit}. E_{\delta_w}[\tau_{\text{Debt}}] = \infty$ [26, p.304]. However, when the gambler reaches a wealth of $-\$1$ or smaller, their probability of returning to *Solvency* is less than 1:

$$\sup_{w \in \text{Debt}} P_{\delta_w}(\tau_{\text{Solvency}} < \infty) = \sup_{w \in \text{Debt}} \left(\frac{1-\varepsilon}{1+\varepsilon} \right)^{|w|} = \left(\frac{1-\varepsilon}{1+\varepsilon} \right) < 1. \quad \lrcorner \quad (8)$$

In order to reason about the interaction of stopping times with the Markov property, we recall the (time) shift operator $\theta : \Omega \rightarrow \Omega$ [25, Definition 3.1.8, p.58] defined by

$$\theta(s_0, s_1, \dots, s_n, \dots) = (s_1, s_2, \dots, s_{n+1}, \dots) \quad (9)$$

that deletes the first element of a trajectory. Given $n \in \mathbb{N}$, we can apply the shift operator n times to delete a prefix of length n , denoted as θ^n in the usual way. The shift operator gives rise to two fundamental concepts in stochastic analysis: shift invariance and the Markov property. Our technical results on completeness for quantitative reactivity relies on the fact that reactivity is shift-invariant, i.e., invariant to addition and removal of arbitrary prefixes.

► **Theorem 2.8** (Reactivity is Shift-Invariant). *Let $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ be a Streett acceptance condition, and let L be the corresponding reactivity property defined in Equation (3). Then,*

$$\theta^{-1}L = L. \quad (10)$$

Proof. The event of visiting a region A finitely many times $\text{Fin}(A)$, as well as the event of visiting a region B infinitely many times $\text{Inf}(B)$ both satisfy shift invariance:

$$\begin{aligned}\theta^{-1} \text{Fin}(A) &= \{w \in \Omega: \exists m \geq 1. \forall n \geq m. \Phi_n(w) \notin A\} \\ &= \{w \in \Omega: \exists m \geq 0. \forall n \geq m. \Phi_n(w) \notin A\} = \text{Fin}(A),\end{aligned}\tag{11}$$

$$\begin{aligned}\theta^{-1} \text{Inf}(B) &= \{w \in \Omega: \forall m \geq 1. \exists n \geq m. \Phi_n(w) \in B\} \\ &= \{w \in \Omega: \forall m \geq 0. \exists n \geq m. \Phi_n(w) \in B\} = \text{Inf}(B).\end{aligned}\tag{12}$$

The class of events L for which Equation (10) holds is closed under countable unions and intersections ([25, Proposition 5.1.5(i), p.99]). Since the reactivity property Equation (3) is a finite Boolean combination of Fin - and Inf -events, it satisfies Equation (10) as well. $\blacktriangleleft$

Our technical results on completeness for Streett acceptance (Theorems 3.12 and 3.14) build upon the result that every time-homogeneous Markov chain satisfies the Markov property.

► **Theorem 2.9** (Markov Property [45, p. 66]). *Let Φ be a time-homogeneous Markov chain. Then, for any bounded random variable H on $(\Omega, \mathcal{F}, \mathbb{P}_\mu)$, and any time $n \in \mathbb{N}$:*

$$\mathbb{E}_\mu [H \circ \theta^n \mid \mathcal{F}_n] = \mathbb{E}_{\delta_{\Phi_n}} [H] \quad \mathbb{P}_\mu\text{-a.s.}\tag{13}$$

3 Absorbing-Region Decomposition for Almost-Sure and Quantitative Reactivity

We show that every reactivity property admits a decomposition into multiple independent obligations, one for each Streett pair, where each obligation consists of two interdependent components: quantitative safety (i.e., safety with positive probability) and almost-sure termination into appropriate regions of the state space, termed absorbing regions.

We first recall the notion of first return time, which is the first hitting time that occurs strictly after the initial step of the process.

► **Definition 3.1** (First Return Time). *The first return time σ_A of a region $A \in \Sigma$ by the stochastic process Φ is*

$$\sigma_A = 1 + \tau_A \circ \theta,\tag{14}$$

where τ_A is the first hitting time of the region A (Definition 2.6).


We observe that, as time tends to infinity, the probability of returning to a region A converges almost surely to the probability of visiting A infinitely often, and that this limit necessarily takes a value of either 0 or 1 for every initial distribution.

► **Lemma 3.2** ([25, Equation 4.2.5]). *Let $A \in \Sigma$ be a region. Then*

$$\lim_{n \rightarrow \infty} \mathbb{P}_{\delta_{\Phi_n}} (\sigma_A < \infty) = \mathbf{1}\{\Phi \in \text{Inf}(A)\},\tag{15}$$

holds $\mathbb{P}_\xi$ -a.s. for any initial probability distribution ξ on (S, Σ) .

► **Example 3.3** (Simulating Lemma 3.2). Figure 2 empirically illustrates this standard result using 10 stochastic simulations of the Lending Casino from Figure 1. The left hand plot shows that, over long runs, the process tends to remain in *Debt* states (negative values)



■ **Figure 2** Simulation of Return Probabilities.

and stops visiting *Solvency* states (non-negative values). The right hand plot, based on the same simulations, shows that the probability of returning to *Solvency* disappears from the plot, because it tends to 0. This supports the following almost-sure property of the Lending Casino:

$$\mathbf{1}\{\Phi \in \text{Inf}(\text{Solvency})\} = 0 = \lim_{n \rightarrow \infty} P_{\delta_{\Phi_n}}(\sigma_{\text{Solvency}} < \infty). \quad \lrcorner \quad (16)$$

Lemma 3.2 leads to the following standard result – Orey’s Theorem – which provides a sufficient condition for the reactivity property defined by a single Streett pair.


► **Theorem 3.4** (Orey’s Theorem [25, Theorem 4.2.6]). *Let $A, B \in \Sigma$ be regions such that*

$$\inf_{s \in A} P_{\delta_s}(\sigma_B < \infty) > 0. \quad (17)$$


Then $P_\mu(\text{Fin}(A) \cup \text{Inf}(B)) = 1$.

Intuitively, the theorem states that if every state in A leads to a trajectory that eventually returns to B with positive probability, then any process that visits A infinitely often must almost surely also visit B infinitely often. In other words, it provides a sufficient condition for the event $\text{Fin}(A) \cup \text{Inf}(B)$, that is, the acceptance of a single Streett pair (A, B) . However, Orey’s theorem is incomplete, because a Markov chain can almost-surely satisfy a Streett pair, yet fail to satisfy the condition (17). We will now present two counterexamples to its completeness and introduce strategies to address them. Our three main results in Theorems 3.7, 3.12, and 3.14 then generalise these two strategies, leading to our framework for the construction of sound and complete proof rules for reactivity properties on general state spaces.

► **Example 3.5** (Orey’s Theorem is Incomplete for Almost-Sure Reactivity). We show that it is possible for a Markov chain to satisfy a Streett pair almost-surely without satisfying the premise of Orey’s Theorem. Observe the finite-state system given in Figure 3, and consider the Streett pair (A, B) where $A = \{s_1, s_3\}$ are the states labelled by a , and $B = \{s_2\}$ are the states labelled by b . It is easy to see that the system satisfies the Streett condition (A, B) almost surely, as we either alternate between s_1 and s_2 , visiting B infinitely often, or visit s_3 once and then stay in s_4 , visiting A at most finitely many times. However, it is also easy to see that the premise of Orey’s Theorem (Equation (17)) does not hold: consider the problematic state s_3 . For this state, $s_3 \in A$, but it is impossible to reach a state in B , making the minimum probability to reach B from A nil and violating Equation (17). We address the incompleteness of this example by extending the set of states that the system must return to infinitely often. In other words, we extend B with the additional set $J = \{s_4\}$ which, in this example, consists of all states outside of A from which the probability of returning



■ Figure 3 Almost-Sure Reactivity.



■ Figure 4 Quantitative Reactivity.

to A is no greater than $1/2$. We then observe that while Equation (17) does not hold for the original pair (A, B) , it does hold on the extended pair $(A, B \cup J)$. Furthermore, we can intuitively see that it is sound to use this new extended Streett pair, because in this example, it is impossible to alternate between A and $J \setminus B$ infinitely many times. Therefore, visiting $J \setminus B$ infinitely often implies visiting A finitely many times. $\lrcorner$

► **Example 3.6** (Orey’s Theorem is Incomplete for Quantitative Reactivity). Consider the Markov chain given in Figure 4, which extends the previous process with the additional problematic state s_5 . This state fails to satisfy the reactivity property for the Streett condition (A, B) , where $A = \{s_1, s_3, s_5\}$ and $B = \{s_2\}$, but such trajectories account only for $1/3$ of the total probability. The remainder of the system satisfies the property and these trajectories constitute the remaining $2/3$ of the probability mass. However, the strategy of Example 3.5 alone fails to capture this quantitative distinction. In fact, the system does not satisfy the updated condition $\inf_{s \in A} \mathbb{P}_{\delta_s}(\sigma_{B \cup J} < \infty) > 0$, because $s_5 \in A$ has no path that can reach $B \cup J$.

To resolve this, we identify an additional invariant region I , such that any trajectory that remains entirely within I satisfies the Streett acceptance condition $(A, B \cup J)$ almost surely. This is equivalent to identifying a sufficiently large region for which either I is exited or the extended Streett pair $(A, B \cup J)$ is satisfied. As a result, in this example, the probability of exiting I serves to upper-bound the probability of violating the original reactivity condition. $\lrcorner$

Our technical results formalise the intuition by which adding absorbing regions J and invariant regions I provides completeness while retaining soundness, as we illustrated in Examples 3.5 and 3.6. Crucially, we show that this approach generalises to every time-homogeneous Markov chain on a general state space.

Our arguments build on the scaffolding illustrated in Figure 5. We present our decomposition in full generality and show its soundness in Theorem 3.7. We then prove in Lemma 3.9 that every Markov chain that either violates an invariant I or satisfies a reactivity property necessarily admits appropriate absorbing regions J . This provides the basis for establishing the completeness of our decomposition from both the quantitative and the qualitative (almost-sure) perspectives. In the general quantitative setting, we show that every Markov chain admits a sufficiently strong invariant I , whose probability of invariance lower-bounds the probability of the reactivity property with any arbitrary $\varepsilon > 0$ gap; this is shown in Lemma 3.11 which in turn establishes our ε -completeness result of Theorem 3.12.

Quantitative	Almost-Sure
Theorem 3.12 (ε -Completeness)	Theorem 3.14 (Completeness)
Lemma 3.11: $\exists I. P(I^\omega) \geq P(L) - \varepsilon$	Lemma 3.13: $\exists I. P(I^\omega) = P(L) = 1$
Lemma 3.9: $P((I^\omega)^c \cup L) = 1 \implies \exists J. \text{Eqs. (18)} \wedge (19)$	
Theorem 3.7 (Soundness): $\forall I, J. \text{Eqs. (18)} \wedge (19) \Rightarrow P(L) \geq P(I^\omega)$	

■ **Figure 5** Structure of our Technical Results.

In the special case that the reactivity property is satisfied almost surely, we show that this gap is zero, which in turn establishes our completeness result for almost-sure reactivity of Theorem 3.14.

► **Theorem 3.7 (Absorbing-Region Decomposition).** *Let $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ be a Streett acceptance condition. Suppose there exists a region $I \in \Sigma$ and k regions $J_1 \in \Sigma, \dots, J_k \in \Sigma$ such that, for every $i = 1, \dots, k$, $J_i \subseteq I \setminus A_i$ and*

$$\sup_{s \in J_i} P_{\delta_s}(\sigma_{A_i} < \infty) < 1, \quad (18)$$

$$\inf_{s \in I} P_{\delta_s}(\sigma_{B_i \cup J_i \cup I^c} < \infty) = 1. \quad (19)$$

Then

$$P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) \geq P_\mu(I^\omega). \quad (20)$$

Proof. Let i be arbitrary among $\{1, \dots, k\}$. We apply Orey's Theorem (Theorem 3.4) to a consequence of Equation (19) obtained by restricting the infimum to a smaller set $A_i \cap I \subseteq I$:

$$\inf_{s \in A_i \cap I} P_{\delta_s}(\sigma_{B_i \cup J_i \cup I^c} < \infty) \geq \inf_{s \in I} P_{\delta_s}(\sigma_{B_i \cup J_i \cup I^c} < \infty) = 1, \quad (21)$$

and Orey's Theorem applied to (21) yields:

$$P_\mu(\text{Fin}(A_i \cap I) \cup \text{Inf}(B_i \cup J_i \cup I^c)) = 1. \quad (22)$$

We now establish the inclusion

$$\text{Fin}(A_i \cap I) \subseteq (I^\omega)^c \cup (\text{Fin}(A_i) \cap I^\omega) \quad (23)$$

so that we can replace $\text{Fin}(A_i \cap I)$ in Equation (22) with $(I^\omega)^c \cup (\text{Fin}(A_i) \cap I^\omega)$ to obtain Equation (24). To this end, we first note that $\text{Fin}(A_i \cap I) \cap I^\omega = \text{Fin}(A_i) \cap I^\omega$, because for any trajectory $(s_0, s_1, \dots, s_n, \dots) \in I^\omega$, and any time $n \in \mathbb{N}$, we have $s_n \in A_i \cap I \Leftrightarrow s_n \in A_i$, of which Equation (23) is a consequence. We thus obtain:

$$P_\mu((I^\omega)^c \cup (\text{Fin}(A_i) \cap I^\omega) \cup \text{Inf}(B_i \cup J_i \cup I^c)) = 1 \quad (24)$$

Using the fact that $\text{Inf}(B_i \cup J_i \cup I^c) = \text{Inf}(B_i) \cup \text{Inf}(J_i) \cup \text{Inf}(I^c)$, and that any trajectory performing infinitely many visits to I^c must necessarily leave I , namely:

$$\text{Inf}(I^c) \subseteq (I^\omega)^c, \quad (25)$$

3:10 Complete ω -Regular Supermartingale Certificates

we rewrite Equation (24) into Equation (26):

$$\mathbf{P}_\mu((I^\omega)^\complement \cup (\text{Fin}(A_i) \cap I^\omega) \cup \text{Inf}(B_i) \cup \text{Inf}(J_i)) = 1. \quad (26)$$

By set-theoretic reasoning, we replace $(I^\omega)^\complement \cup (\text{Fin}(A_i) \cap I^\omega)$ in Equation (26) by $(I^\omega)^\complement \cup \text{Fin}(A_i)$, since

$$\overbrace{(I^\omega)^\complement \cup (\text{Fin}(A_i) \cap (I^\omega)^\complement)}^{(I^\omega)^\complement} \cup (\text{Fin}(A_i) \cap I^\omega) = (I^\omega)^\complement \cup \text{Fin}(A_i) \cap \overbrace{((I^\omega)^\complement \cup I^\omega)}^\Omega. \quad (27)$$

Therefore we obtain:

$$\mathbf{P}_\mu((I^\omega)^\complement \cup \text{Fin}(A_i) \cup \text{Inf}(B_i) \cup \text{Inf}(J_i)) = 1. \quad (28)$$

Our next objective is to remove $\text{Inf}(J_i)$ from Equation (28). We first note that by Equation (18) that there exists $\gamma < 1$ for which

$$\forall s \in J_i. \mathbf{P}_{\delta_s}(\sigma_{A_i} < \infty) \leq \gamma. \quad (29)$$

Therefore, in any trajectory that visits J_i infinitely many times, the probability of returning to A_i must drop below γ infinitely often. Formally, using the definition of $\text{Inf}(J_i)$ in Equation (2) and Equation (29) we derive:

$$\Phi \in \text{Inf}(J_i) \iff \forall n \exists m \geq n. \Phi_m \in J_i \quad (30)$$

$$\implies \forall n \exists m \geq n. \mathbf{P}_{\delta_{\Phi_m}}(\sigma_{A_i} < \infty) \leq \gamma. \quad (31)$$

By Lemma 3.2, the limit $\lim_{n \rightarrow \infty} \mathbf{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty)$ exists and takes value either zero or one, $\mathbf{P}_\mu$ -almost surely. We note that (31) implies that the limit cannot be one, meaning that the limit must be equal to zero, $\mathbf{P}_\mu$ -almost surely. Summarising, we conclude that there exists a $\mathbf{P}_\mu$ -negligible set $N_3 \in \mathcal{F}$ for which

$$\Phi \in \text{Inf}(J_i) \setminus N_3 \implies \lim_{n \rightarrow \infty} \mathbf{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty) = 0. \quad (32)$$

Finally, since (by Lemma 3.2) $\lim_{n \rightarrow \infty} \mathbf{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty) = 0$ and $\text{Fin}(A_i)$ differ by a $\mathbf{P}_\mu$ -negligible set, we conclude for some $\mathbf{P}_\mu$ -negligible set N_4 , that:

$$\text{Inf}(J_i) \setminus N_4 \subseteq \text{Fin}(A_i). \quad (33)$$

This means we can replace $\text{Fin}(A_i) \cup \text{Inf}(J_i)$ in Equation (28) by $\text{Fin}(A_i)$ to obtain:

$$\mathbf{P}_\mu((I^\omega)^\complement \cup \text{Fin}(A_i) \cup \text{Inf}(B_i)) = 1. \quad (34)$$


Since this holds for arbitrary $i = 1, \dots, k$, we conclude the result of the theorem by noting that the intersection of finitely many probability 1 events has probability 1:

$$\mathbf{P}_\mu \left((I^\omega)^\complement \cup \bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1, \quad (35)$$

and by applying a union bound we obtain

$$\mathbf{P}_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) \geq 1 - \mathbf{P}_\mu((I^\omega)^\complement) = \mathbf{P}_\mu(I^\omega) \quad (36)$$

which is the desired result. $\blacktriangleleft$



■ **Figure 6** Intuition for Absorbing-Region Decomposition.

► **Example 3.8** (Intuition for Theorem 3.7). Consider Figure 6, which displays a Markov chain that satisfies the requirements of Theorem 3.7 and therefore satisfies the Streett condition with probability $\frac{2}{3}$. In the previous examples, we could not return from the set J to A at all. However, in this example, we can go from the state s_6 to s_1 with probability $\frac{1}{2}$.

We observe that any path that visits the set A infinitely often must either exit the set I , as before, or visit the set $B \cup J$ infinitely often. If B is visited infinitely often, we satisfy Streett. On the other hand, we see that for the choice of $J = \{s_4, s_6\}$ that the probability of returning to A from it is no greater than $\frac{1}{2}$, which means that the probability of returning to A from it infinitely often is zero. In our example, this corresponds to the fact that, in order to reach A infinitely often, the Markov chain must either go to s_5 , which is outside the set I , or take the transition from s_6 to s_1 infinitely often, the probability of which is zero. $\lrcorner$

Theorem 3.7 establishes that if a Markov chain admits an invariant and absorbing regions then it satisfies the reactivity property with probability at least $P_\mu(I^\omega)$. We next show, in Lemma 3.9, that given a sufficiently strong invariant region (namely, a region $I \in \Sigma$ satisfying (37)), there exist suitable absorbing regions, such that the invariant and the absorbing regions jointly satisfy requirements (18) and (19).

► **Lemma 3.9.** *Let $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ be a Streett acceptance condition. Let $I \in \Sigma$ be a region for which*

$$\forall s \in S. P_{\delta_s} \left((I^\omega)^c \cup \bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1. \quad (37)$$

Then, there exist regions $J_i \in \Sigma$ with $J_i \subseteq I \setminus A_i$ for $i = 1, \dots, k$ such that Equations (18) and (19) hold.

Proof. For every $i \in \{1, \dots, k\}$ we define the region

$$J_i = \{s \in I \setminus A_i : P_{\delta_s}(\sigma_{A_i} < \infty) \leq 0.5\}, \quad (38)$$

with 0.5 being an arbitrarily chosen constant in $(0, 1)$. The region J_i of (38) satisfies (18) by construction, since $\sup_{s \in J_i} P_{\delta_s}(\sigma_{A_i} < \infty) \leq 0.5$.

The premise (37) is equivalent to

$$\forall s \in S. \forall i = 1, \dots, k. P_{\delta_s}((I^\omega)^c \cup \text{Fin}(A_i) \cup \text{Inf}(B_i)) = 1, \quad (39)$$

and by Equation (27) we replace $(I^\omega)^c \cup \text{Fin}(A_i)$ by $(I^\omega)^c \cup (\text{Fin}(A_i) \cap I^\omega)$ in Equation (39) to obtain:

$$\forall s \in S. \forall i = 1, \dots, k. P_{\delta_s}((I^\omega)^c \cup (I^\omega \cap \text{Fin}(A_i)) \cup \text{Inf}(B_i)) = 1. \quad (40)$$

3:12 Complete ω -Regular Supermartingale Certificates

We now show that $I^\omega \cap \text{Fin}(A_i)$ in Equation (40) may be replaced by $\text{Inf}(J_i)$ to obtain Equation (48). To this end, we invoke Lemma 3.2, which says that the events $\text{Fin}(A_i)$ and $\{\lim_{n \rightarrow \infty} \mathbb{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty) = 0\}$ are $\mathbb{P}_{\delta_s}$ -equivalent for every $s \in S$. This means that for every $s \in S$ there exists a $\mathbb{P}_{\delta_s}$ -negligible event $N_s \in \mathcal{F}$ for which $\text{Fin}(A_i) \setminus N_s \subseteq \{\lim_{n \rightarrow \infty} \mathbb{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty) = 0\}$. We extend this to the following inclusions among events, which hold for every $s \in S$:

$$I^\omega \cap \text{Fin}(A_i) \setminus N_s \subseteq I^\omega \cap \text{Fin}(A_i) \cap \left\{ \lim_{n \rightarrow \infty} \mathbb{P}_{\delta_{\Phi_n}}(\sigma_{A_i} < \infty) = 0 \right\} \quad (41)$$

$$= I^\omega \cap \text{Fin}(A_i) \cap \left\{ \forall \varepsilon_1 > 0 \exists n \forall m \geq n. \mathbb{P}_{\delta_{\Phi_m}}(\sigma_{A_i} < \infty) \leq \varepsilon_1 \right\} \quad (42)$$

$$\subseteq I^\omega \cap \text{Fin}(A_i) \cap \left\{ \exists n \forall m \geq n. \mathbb{P}_{\delta_{\Phi_m}}(\sigma_{A_i} < \infty) \leq 0.5 \right\}. \quad (43)$$

We next observe that any trajectory that always remains within I and visits A_i finitely often must eventually always remain inside $I \setminus A_i$:

$$\Phi \in I^\omega \cap \text{Fin}(A_i) \iff \forall n \in \mathbb{N}. \Phi_n \in I \wedge \exists n \forall m \geq n. \Phi_m \notin A_i \quad (44)$$

$$\implies \exists n \forall m \geq n. \Phi_m \in I \setminus A_i. \quad (45)$$

Equations (41) to (45), combined with the definition of J_i in (38), shows that $I^\omega \cap \text{Fin}(A_i) \setminus N_s$ is a subset of trajectories that visit J_i infinitely often, namely:

$$I^\omega \cap \text{Fin}(A_i) \setminus N_s \subseteq \left\{ \exists n \forall m \geq n. \Phi_m \in I \setminus A_i \wedge \mathbb{P}_{\delta_{\Phi_m}}(\sigma_{A_i} < \infty) \leq 0.5 \right\} \quad (46)$$

$$= \left\{ \exists n \forall m \geq n. \Phi_m \in J_i \right\} = \text{Fin}(J_i^c) \subseteq \text{Inf}(J_i). \quad (47)$$

Using inclusions (46) and (47) to replace the event $I^\omega \cap \text{Fin}(A_i)$ in Equation (40) by the larger event $\text{Inf}(J_i)$ we obtain:

$$\forall s \in S. \forall i = 1, \dots, k. \mathbb{P}_{\delta_s}((I^\omega)^c \cup \text{Inf}(J_i) \cup \text{Inf}(B_i)) = 1. \quad (48)$$

We next seek to show that Equation (48) implies Equation (19), by analysing its consequences upon the return times to the regions B_i , J_i , and I^c . We observe that to visit a region $X \in \Sigma$ infinitely often a trajectory must return to it at least once:

$$\Phi \in \text{Inf}(X) \iff [\forall n \exists m \geq n. \Phi_m \in X] \implies [\exists m \geq 1. \Phi_m \in X] \iff \sigma_X < \infty, \quad (49)$$

and therefore,

$$\text{Inf}(J_i) \subseteq \{\sigma_{J_i} < \infty\}, \quad \text{Inf}(B_i) \subseteq \{\sigma_{B_i} < \infty\}. \quad (50)$$

Furthermore, the trajectories that satisfy $(I^\omega)^c$ are precisely those for which $\tau_{I^c} < \infty$:

$$\Phi \in (I^\omega)^c \iff \neg[\forall n. \Phi_n \in I] \iff [\exists n \geq 0. \Phi_n \in I^c] \iff \tau_{I^c} < \infty. \quad (51)$$

Using (50), and (51) we see that (48) implies

$$\forall s \in S. \forall i = 1, \dots, k. \mathbb{P}_{\delta_s}(\tau_{I^c} < \infty \vee \sigma_{J_i} < \infty \vee \sigma_{B_i} < \infty) = 1. \quad (52)$$

Next, we observe that for every $s \in I$ that the events $\{\tau_{I^c} < \infty\}$ and $\{\sigma_{I^c} < \infty\}$ are $\mathbb{P}_{\delta_s}$ -equivalent, because for any trajectory whose initial state s belongs to I , the hitting time to I^c coincides with the return time to I^c :

$$\Phi_0 \in I \implies \underbrace{(\exists n \geq 0. \Phi_n \in I^c)}_{\tau_{I^c} < \infty} \iff \underbrace{(\exists n \geq 1. \Phi_n \in I^c)}_{\sigma_{I^c} < \infty} \quad (53)$$

holds pointwise. This allows us to rewrite Equation (52) into

$$\forall s \in I. \forall i = 1, \dots, k. \mathbf{P}_{\delta_s} (\sigma_{I^c} < \infty \vee \sigma_{J_i} < \infty \vee \sigma_{B_i} < \infty) = 1. \quad (54)$$

by restricting the quantifier to $s \in I$. We observe that

$$\sigma_{I^c} < \infty \vee \sigma_{J_i} < \infty \vee \sigma_{B_i} < \infty \iff \sigma_{B_i \cup J_i \cup I^c} < \infty \quad (55)$$

which allows us to rewrite Equation (54) into

$$\forall s \in I. \forall i = 1, \dots, k. \mathbf{P}_{\delta_s} (\sigma_{B_i \cup J_i \cup I^c} < \infty) = 1, \quad (56)$$

from which Equation (19) follows. $\blacktriangleleft$

The premise of Lemma 3.9 is that the invariant I is sufficiently strong, in the sense of condition (37), which requires that almost every trajectory that remains forever within I also satisfies the reactivity property $\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i)$. Equivalently, the probability of the reactivity property conditional upon I^ω must equal 1 (cf. [15, Theorem 2], [4, Theorem 6]). Given any such invariant, Lemma 3.9 guarantees the existence of absorbing regions that, together with I , satisfy Equations (18) and (19).

► **Example 3.10** (Necessity of Quantitative Safety). In the quantitative setting, where the reactivity property fails to hold almost surely, any sufficiently strong invariant must be exited with positive probability. For instance, in Figure 6, no invariant satisfying Equation (37) may contain s_5 , since otherwise the trajectory $s_0(s_5)^\omega$ would remain within I while failing the reactivity property. Since the Markov chain reaches s_5 with probability $1/3$, any such I must be exited with probability at least $1/3$. The joint conditions on the absorbing regions and the invariant accommodate this precisely, by permitting trajectories to exit I in Equation (19).

It remains to show that a sufficiently strong invariant always exists. To this end, we rely on the shift-invariance of reactivity properties. In Lemma 3.11, we show that there always exists a region whose probability of invariance can be made arbitrarily close to that of the shift-invariant property in question.

► **Lemma 3.11.** *Suppose $L \in \mathcal{F}$ is shift-invariant. Then, for every arbitrary $\varepsilon > 0$ there exists a region $I \in \Sigma$ such that*

$$\mathbf{P}_\mu(I^\omega) \geq \mathbf{P}_\mu(L) - \varepsilon, \quad (57)$$

$$\forall s \in S. \mathbf{P}_{\delta_s}((I^\omega)^c \cup L) = 1. \quad (58)$$

Proof. We consider the event $\{\inf_{n \in \mathbb{N}} \mathbf{P}_{\delta_{\Phi_n}}(L) > 0\}$ that the satisfaction probability of L , when evaluated along a trajectory, is uniformly bounded away from zero. We note that if, for some $k \in \mathbb{N}$, the satisfaction probability is always above $\frac{1}{k+1}$, then it is uniformly bounded away from zero. That is, we note

$$\left\{ \inf_{n \in \mathbb{N}} \mathbf{P}_{\delta_{\Phi_n}}(L) \geq \frac{1}{k+1} \right\} \subseteq \left\{ \inf_{n \in \mathbb{N}} \mathbf{P}_{\delta_{\Phi_n}}(L) > 0 \right\} \quad (59)$$

for every $k \in \mathbb{N}$, and that

$$\bigcup_{k \in \mathbb{N}} \left\{ \inf_{n \in \mathbb{N}} \mathbf{P}_{\delta_{\Phi_n}}(L) \geq \frac{1}{k+1} \right\} = \left\{ \inf_{n \in \mathbb{N}} \mathbf{P}_{\delta_{\Phi_n}}(L) > 0 \right\}. \quad (60)$$

3:14 Complete ω -Regular Supermartingale Certificates

We define, for each $k \in \mathbb{N}$, the region

$$I_k = \left\{ s \in S : P_{\delta_s}(L) \geq \frac{1}{k+1} \right\}, \quad (61)$$

which has the property that the events I_k^ω and $\{\inf_{n \in \mathbb{N}} P_{\delta_{\Phi_n}}(L) \geq \frac{1}{k+1}\}$ are identical.

By [4, Theorem 3], the events $\{\inf_{n \in \mathbb{N}} P_{\delta_{\Phi_n}}(L) > 0\}$ and L are P_ξ -equivalent for every initial distribution ξ . Using this, and applying the Monotone Convergence Theorem ([9, Theorem 2.59]) to Equation (60) yields $\lim_{k \rightarrow \infty} P_\mu(I_k^\omega) = P_\mu(L)$, so there must exist $k_0 \in \mathbb{N}$ for which $P_\mu(I_{k_0}^\omega) \geq P_\mu(L) - \varepsilon$, i.e. Equation (57). Since $I_{k_0}^\omega \subseteq \{\inf_{n \in \mathbb{N}} P_{\delta_{\Phi_n}}(L) > 0\}$ and $\{\inf_{n \in \mathbb{N}} P_{\delta_{\Phi_n}}(L) > 0\}$ is (for every ξ) P_ξ -almost surely equivalent to L , we conclude that $(I_{k_0}^\omega)^c \cup L$ holds P_{δ_s} -almost surely for every $s \in S$, i.e. Equation (58). $\blacktriangleleft$

A consequence of Lemma 3.11 is that, for reactivity properties, a sufficiently strong invariant always exists, and can be chosen so that its probability of invariance lower-bounds the probability of reactivity up to an arbitrarily small gap (Theorem 3.12).

► **Theorem 3.12** (ε -Completeness of Absorbing-Region Decomposition). *Let $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ be a Streett acceptance condition. Suppose*

$$P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) > 0. \quad (62)$$

Then for every arbitrary $\varepsilon > 0$, there exists a region $I \in \Sigma$, and regions $J_i \in \Sigma$ with $J_i \subseteq I \setminus A_i$ for $i = 1, \dots, k$ such that Equations (18) and (19) hold and

$$P_\mu(I^\omega) \geq P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) - \varepsilon. \quad (63)$$

Proof. The reactivity property corresponding to the Streett acceptance condition is shift-invariant, and therefore by Lemma 3.11 there exists a region I for which Equation (63) holds and for which

$$\forall s \in S. P_{\delta_s} \left((I^\omega)^c \cup \bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1. \quad (64)$$

Using this equality, we can invoke Lemma 3.9 to show that there exist regions $J_i \in \Sigma$ with $J_i \subseteq I \setminus A_i$ for $i = 1, \dots, k$ such that Equations (18) and (19) hold, which completes the proof of the theorem. $\blacktriangleleft$

By contrast, in the almost-sure setting, there exists a region whose probability of invariance matches the probability of the shift-invariant property exactly (Lemma 3.13).

► **Lemma 3.13.** *Suppose $L \in \mathcal{F}$ is shift-invariant and $P_\mu(L) = 1$. Then there exists a region $I \in \Sigma$ such that*

$$P_\mu(I^\omega) = 1, \quad (65)$$

$$\forall s \in S. P_{\delta_s}((I^\omega)^c \cup L) = 1. \quad (66)$$

Proof. We define $I = \{s \in S : P_{\delta_s}(L) = 1\}$. We prove that I satisfies Equation (65) (which also proves I to be non-empty) and Equation (66).

The crux of our proof is that the stochastic process $n \mapsto P_{\delta_{\Phi_n}}(L)$ is a martingale adapted to the filtered probability space $(\Omega, \mathcal{F}, \mathfrak{F}, P_\xi)$ for every initial distribution ξ on (S, Σ) , which in concrete terms means that

$$\forall n \in \mathbb{N}. E_\xi \left(P_{\delta_{\Phi_{n+1}}}(L) \mid \mathcal{F}_n \right) = P_{\delta_{\Phi_n}}(L) \quad P_\xi\text{-a.s.}, \quad (67)$$

which we now prove. First, we let ξ be an arbitrary initial distribution on (S, Σ) , and we recall that L is shift-invariant, meaning that the following equality

$$\mathbf{1}\{\Phi \in L\} = \mathbf{1}\{\Phi \circ \theta^n \in L\}, \quad (68)$$

holds pointwise for every $n \in \mathbb{N}$. Applying the conditional expectation operator $E_\xi(\cdot \mid \mathcal{F}_n)$ to both sides:

$$\forall n \in \mathbb{N}. E_\xi (\mathbf{1}\{\Phi \in L\} \mid \mathcal{F}_n) = E_\xi (\mathbf{1}\{\Phi \circ \theta^n \in L\} \mid \mathcal{F}_n) \quad P_\xi\text{-a.s.} \quad (69)$$

The Markov property (Equation (13)) states

$$E_\xi (\mathbf{1}\{\Phi \circ \theta^n \in L\} \mid \mathcal{F}_n) = P_{\delta_{\Phi_n}}(L) \quad P_\xi\text{-a.s.}, \quad (70)$$

and applying this to Equation (69) yields:

$$\forall n \in \mathbb{N}. E_\xi (\mathbf{1}\{\Phi \in L\} \mid \mathcal{F}_n) = P_{\delta_{\Phi_n}}(L) \quad P_\xi\text{-a.s.} \quad (71)$$

Lévy's 0-1 Law [26, Theorem 5.5.8] states

$$\lim_{n \rightarrow \infty} E_\xi (\mathbf{1}\{\Phi \in L\} \mid \mathcal{F}_n) = \mathbf{1}\{\Phi \in L\} \quad P_\xi\text{-a.s.} \quad (72)$$

Combining Equations (71) and (72), we have:

$$\mathbf{1}\{\Phi \in L\} = \lim_{n \rightarrow \infty} P_{\delta_{\Phi_n}}(L) \quad P_\xi\text{-a.s.} \quad (73)$$

Therefore, for some P_ξ -negligible event $N \in \mathcal{F}$, we have:

$$\left\{ \lim_{n \rightarrow \infty} P_{\delta_{\Phi_n}}(L) = 1 \right\} \setminus N \subseteq L. \quad (74)$$

We note that $I^\omega \setminus N \subseteq \left\{ \lim_{n \rightarrow \infty} P_{\delta_{\Phi_n}}(L) = 1 \right\}$ as:

$$I^\omega \setminus N = \{ \forall n \in \mathbb{N}. P_{\delta_{\Phi_n}}(L) = 1 \} \setminus N \subseteq \left\{ \lim_{n \rightarrow \infty} P_{\delta_{\Phi_n}}(L) = 1 \right\} \setminus N \subseteq L, \quad (75)$$

which proves that $P_\xi((I^\omega)^c \cup L) = 1$. Since ξ was arbitrary, setting ξ to be δ_s for every arbitrary initial state $s \in S$ concludes the proof of Equation (66).

We next prove Equation (65), noting that since

$$P_\mu(I^\omega) = P_\mu(\forall n \in \mathbb{N}. \Phi_n \in I) = P_\mu(\cap_{n \in \mathbb{N}} \{\Phi_n \in I\}), \quad (76)$$

it suffices to establish

$$\forall n \in \mathbb{N}. P_\mu(\Phi_n \in I) = 1, \quad (77)$$

which we now prove by induction.

Base case. By Equation (1) we observe that:

$$P_\mu(\Phi_0 \in I) = \int_{s_0 \in I} \mu(ds) = \mu(I) \quad (78)$$

and so it suffices to show $\mu(I) = 1$. By [25, Proposition 3.1.3(ii)] and the premise $P_\mu(L) = 1$, we obtain:

$$\int_{s_0 \in S} P_{\delta_{s_0}}(L) \mu(ds_0) = P_\mu(L) = 1. \quad (79)$$

The term $\int_{s_0 \in S} P_{\delta_{s_0}}(L) \mu(ds_0)$ is the expectation of the $[0, 1]$ -valued random variable $s_0 \mapsto P_{\delta_{s_0}}(L)$ with respect to the probability measure μ ; since its expected value is 1, the integrand must μ -almost surely equal 1, that is $\mu(\{s \in S : P_{\delta_s}(L) = 1\}) = 1$ which is precisely $\mu(I) = 1$ (cf. [52, Lemma 26(iv)]). This concludes the proof of the base case.

Inductive Step. We prove that $P_{\delta_{\Phi_{n+1}}}(L) = 1$ holds P_μ -a.s. under the assumption that $P_{\delta_{\Phi_n}}(L) = 1$ holds P_μ -a.s. We recall that the martingale property (67) of the process $n \mapsto P_{\delta_{\Phi_n}}(L)$, instantiating the initial distribution ξ to μ and applying the (unconditional) expectation operator $E_\mu(\cdot)$ to both sides, yields:

$$E_\mu \left(P_{\delta_{\Phi_{n+1}}}(L) \right) \stackrel{[26, \text{Thm. 5.1.6}]}{=} E_\mu \left(E_\mu \left(P_{\delta_{\Phi_{n+1}}}(L) \mid \mathcal{F}_n \right) \right) \stackrel{\text{Eq. (67)}}{=} E_\mu \left(P_{\delta_{\Phi_n}}(L) \right) \stackrel{\text{IH}}{=} 1 \quad (80)$$

where we note that the expression $E_\mu(P_{\delta_{\Phi_n}}(L))$ is equal to one since it is the expected value of a random variable that is P_μ -a.s. equal to one (by the inductive hypothesis). Equation (80) proves $E_\mu(P_{\delta_{\Phi_{n+1}}}(L)) = 1$, implying that $P_{\delta_{\Phi_{n+1}}}(L) = 1$ holds P_μ -a.s, which concludes the inductive step, and the proof of the theorem. $\blacktriangleleft$

Combining Lemma 3.13 with Lemma 3.9, we conclude that whenever a Markov chain satisfies a reactivity property almost surely, there exists an almost-sure invariant and corresponding absorbing regions satisfying our decomposition (Theorem 3.14).

► Theorem 3.14 (Completeness of Absorbing-Region Decomposition for Almost-Sure Reactivity). *Let $(A_1, B_1) \in \Sigma^2, \dots, (A_k, B_k) \in \Sigma^2$ be a Streett acceptance condition. Suppose that*


$$P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1. \quad (81)$$

Then, there exists a region $I \in \Sigma$, and regions $J_i \in \Sigma$ with $J_i \subseteq I \setminus A_i$ for $i = 1, \dots, k$ such that $P_\mu(I^\omega) = 1$ and Equations (18) and (19) hold.

Proof. Using the fact that reactivity is shift-invariant (Theorem 2.8), and that the reactivity property in Equation (81) holds P_μ -a.s., we invoke Lemma 3.13 to obtain an invariant $I \in \Sigma$ for which $P_\mu(I^\omega) = 1$ and

$$\forall s \in S. P_{\delta_s} \left((I^\omega)^c \cup \bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1. \quad (82)$$

But now we can simply apply Lemma 3.9 to the latter equation to show the existence of regions $J_i \subseteq I \setminus A_i$, for $i = 1, \dots, k$, for which Equations (18) and (19) hold, which concludes the proof of the theorem. $\blacktriangleleft$



■ **Figure 7** Methodology Underlying our Supermartingale Certificates for Quantitative Reactivity.

4 Complete Supermartingale Certificates for Almost-Sure and Quantitative Reactivity

Our technical results of Section 3 form the basis for complete proof rules for probabilistic reactivity. They show that reactivity can be established by identifying suitable absorbing regions together with a global invariant, satisfying three conditions: (i) the qualitative (non-almost-sure) safety property of Equation (18), (ii) the almost-sure termination property of Equation (19), and (iii) the quantitative safety property used to lower-bound the right-hand side of Equation (20), which in turn lower-bounds the probability of satisfying the reactivity property. Figure 7 summarises our framework, which lifts every complete proof rule for almost-sure termination into an ε -complete proof rule for quantitative reactivity.

In this section, we show that quantitative and qualitative safety admit sound and complete proof rules in Theorems 4.1 and 4.2, respectively, over general state spaces. It then suffices to pick a sound and complete proof rule for almost-sure termination and substitute it into this framework to obtain a sound and complete proof rule for Streett conditions. Below, we carry out this instantiation with two pre-existing supermartingale proof rules for almost-sure termination (Proof Rules #1 and #2), and establish their soundness and completeness in the theorems that follow, with the proofs following from our results in Section 3.

Supermartingale certificates provide proof rules for temporal properties of stochastic processes and are expressed as non-negative, real-valued measurable functions $V: S \rightarrow \mathbb{R}_{\geq 0}$ over the state space S . Such certificates are typically required to satisfy conditions on their expected values with respect to the transition kernel P and the initial distribution μ of a time-homogeneous Markov chain.

Two standard tools that we will need in this setting are the post-expectation $(PV): S \rightarrow \mathbb{R}_{\geq 0}$ and the init-expectation $(\mu V) \in \mathbb{R}_{\geq 0}$ of V , respectively defined as follows:

$$PV(s) = \int_{s' \in S} V(s') P(s, ds'), \quad \mu V = \int_{s \in S} V(s) \mu(ds). \quad (83)$$

Intuitively, the post-expectation is just an expected value formula and gives the expected value of V at the next state given that we are in a state s . On the other hand, the init-expectation is the expected value of V at the first state given an initial distribution μ .

We instantiate the safety obligations of our framework using a standard result for quantitative safety which requires the identification of a value function to upper bound the probability that the stochastic process under consideration leaves a given region.

► **Theorem 4.1** ([25, Corollary 4.4.7]). *Let $X \in \Sigma$ be a region.*

- i) *Suppose there exists a measurable function $V : S \rightarrow \mathbb{R}_{\geq 0}$ for which the following conditions hold*

$$\forall s \in X. PV(s) \leq V(s), \quad (84)$$

$$\forall s \notin X. V(s) \geq 1. \quad (85)$$

Then, $P_\xi(X^\omega) \geq 1 - \xi V$ for every initial distribution ξ on (S, Σ) .

- ii) *Suppose ξ is a probability distribution on (S, Σ) for which $P_\xi(X^\omega) \geq p$, for some $p \in [0, 1]$. Then, there exists a measurable function $V : S \rightarrow \mathbb{R}_{\geq 0}$ for which Equations (84) and (85) hold, and for which $\xi V \leq 1 - p$.*

Theorem 4.1 provides a sound and complete proof rule for quantitative safety with respect to an invariant I . Combined with Equation (20), it yields a lower bound on the probability of reactivity, provided that reactivity holds almost surely within I , which in turn is ensured by Equations (18) and (19). It remains to give a proof rule for the qualitative safety obligation of Equation (18). To this end, we apply Theorem 4.1 to prove Theorem 4.2, which provides a sound and complete supermartingale certificate for establishing that J is an absorbing region with respect to A , that is, that the probability of returning from J to A is uniformly bounded below one.

► **Theorem 4.2.** *Suppose $A, J \in \Sigma$ are regions with A disjoint from J . Then, there exists a function $V : S \rightarrow \mathbb{R}_{\geq 0}$ for which*

$$\forall s \notin A. PV(s) \leq V(s), \quad (86)$$

$$\forall s \in A. V(s) \geq 1, \quad (87)$$

$$\exists \gamma > 0 \forall s \in J. V(s) \leq 1 - \gamma, \quad (88)$$

if and only if $\sup_{s \in J} P_{\delta_s}(\sigma_A < \infty) < 1$.

Proof. ($\Rightarrow$) Suppose $V : S \rightarrow \mathbb{R}_{\geq 0}$ is a function for which Equations (86)–(88) hold, then by Theorem 4.1(i) we conclude

$$\forall s \in S. P_{\delta_s}(\tau_A < \infty) = 1 - P_{\delta_s}((A^c)^\omega) \leq \delta_s V = V(s). \quad (89)$$

Since A and J are disjoint, we note that for any trajectory initialised in J , the first hitting time to A and the first return time to A coincide:

$$\forall s \in J. P_{\delta_s}(\sigma_A = \tau_A) = 1. \quad (90)$$

This allows us to replace τ_A by σ_A in Equation (89) once we restrict the quantifier to the subset J which is disjoint from A :

$$\forall s \in J. P_{\delta_s}(\sigma_A < \infty) \leq V(s). \quad (91)$$

Applying Equation (88) to Equation (91) we obtain that the return probability to A is upper bounded by $1 - \gamma$ when the initial state belongs to J :


$$\forall s \in J. P_{\delta_s}(\sigma_A < \infty) \leq 1 - \gamma, \quad (92)$$

which implies $\sup_{s \in J} P_{\delta_s}(\sigma_A < \infty) < 1$, concluding the proof of the forward direction.

($\Leftarrow$) We define $V(s) = \mathbb{P}_{\delta_s}(\tau_A < \infty)$ which by [25, Corollary 4.4.7] satisfies conditions (86) and (87). Moreover, we note that the premise $\sup_{s \in J} \mathbb{P}_{\delta_s}(\sigma_A < \infty) < 1$ is equivalent to the existence of a constant $\gamma > 0$ such that

$$\forall s \in J. \mathbb{P}_{\delta_s}(\sigma_A < \infty) \leq 1 - \gamma, \quad (93)$$

which (recalling Equation (90)) is precisely Equation (88) for our choice of V , which completes the reverse direction of the proof. $\blacktriangleleft$



■ **Figure 8** Debt is an Absorbing Region.

► **Example 4.3** (Debt is an Absorbing Region). We illustrate our absorbing-region decomposition in the context of the Lending Casino (Figure 1) for the property $\text{Fin}(\text{Solvency})$, which corresponds to the Streett pair $(\text{Solvency}, \emptyset)$, as introduced in Example 2.5. Throughout, we take the global invariant $I = \mathbb{Z}$ to be the entire state space and the absorbing region $J_1 = \text{Debt}$, and verify that this choice satisfies the premises of Theorem 3.7. Since Debt is disjoint from Solvency , the return time to Solvency from any $w \in \text{Debt}$ coincides with the hitting time computed in Equation (8), which establishes Equation (18). For Equation (19), because the states in Solvency form a symmetric one-dimensional random walk, the state -1 is reached almost surely from any non-negative starting point; therefore the chain returns to $B_1 \cup J_1 \cup I^c = \emptyset \cup \text{Debt} \cup \emptyset$ almost surely from every state in I . Since $\mathbb{P}_\mu(I^\omega) = 1$ trivially, Theorem 3.7 yields $\mathbb{P}_\mu(\text{Fin}(\text{Solvency})) = 1$.

We now establish the same facts via supermartingale certificates. To prove that Debt is an absorbing region with respect to Solvency , we invoke Theorem 4.2 for the certificate V_1 shown in Figure 8. By construction, $V_1(s)$ gives the exact probability of returning to Solvency from $s \in \text{Debt}$, and its supremum over Debt is $\left(\frac{1-\varepsilon}{1+\varepsilon}\right) < 1$ (cf. Equation (8)). This discharges the qualitative safety obligation of Equation (18) and confirms that $J_1 = \text{Debt}$ is absorbing with respect to Solvency . It remains to certify the almost-sure termination obligation Equation (19), which we resolve in the rest of the section by applying pre-existing proof rules for almost-sure termination (cf. Examples 4.7 and 4.11). $\dashv$

Theorems 4.1 and 4.2 discharge the safety obligations of our decomposition using sound supermartingale certificates that are complete over general state spaces. It remains to certify the almost-sure termination obligation of Equation (19) generated by our absorbing-region decomposition, that requires termination into $B_i \cup J_i \cup I^c$ from every state in I , for each Streett pair (A_i, B_i) . This obligation may be discharged using any proof rule for almost-sure termination. Furthermore, if the chosen rule is sound and complete over a specified class of systems, our framework lifts it to a sound and ε -complete proof rule for quantitative reactivity over that class, and to a sound and fully complete one when the reactivity property holds almost surely.

We illustrate our methodology using two existing supermartingale proof rules for almost-sure termination that are known to be complete over countably infinite state spaces, yielding the first sound and complete supermartingale proof rules for ω -regular properties on this class of systems.

Proof Rule #1

We discharge the almost-sure termination obligation (Equation (19)) of our absorbing-region decomposition using the sound and complete proof rule of [40, Proof Rule 3.2]. Let $(A_1, B_1), \dots, (A_k, B_k)$ be a Streett acceptance condition, with $A_i, B_i \subseteq S$ for $i = 1, \dots, k$, over the countable state space $(S, 2^S)$ of a time-homogeneous Markov chain with transition kernel P . Then, *Proof Rule #1* requires the existence of

- **a supporting invariant** $I \subseteq S$ with $V_0: S \rightarrow \mathbb{R}_{\geq 0}$,
- **absorbing regions** $J_1 \subseteq I \setminus A_1, \dots, J_k \subseteq I \setminus A_k$,
- **supermartingale certificates** $V_1, W_1, \dots, V_k, W_k: S \rightarrow \mathbb{R}_{\geq 0}$,
- **ranking functions** $U_1, \dots, U_k: S \rightarrow \mathbb{N}$,

such that the following conditions hold:

Qualitative safety

For every $i = 1, \dots, k$:

$$\forall s \notin A_i. PV_i(s) \leq V_i(s) \quad (94)$$

$$\forall s \in A_i. V_i(s) \geq 1 \quad (95)$$

$$\exists \gamma_i > 0 \forall s \in J_i. V_i(s) \leq 1 - \gamma_i \quad (96)$$

Almost-sure termination

For every $i = 1, \dots, k$:

$$\forall s \in I \setminus (B_i \cup J_i). PW_i(s) \leq W_i(s) \quad (97)$$

$$\forall s \in I \setminus (B_i \cup J_i). W_i(s) > 0 \quad (98)$$

$$\forall s \in I \cap (B_i \cup J_i). W_i(s) = U_i(s) = 0 \quad (99)$$

$$\forall r \in \mathbb{R}_{\geq 0} \exists \varepsilon_r \in \mathbb{R}_{> 0} \forall s \in I \setminus (B_i \cup J_i). \\ W_i(s) \leq r \implies P(s, \{s' : U_i(s') < U_i(s)\}) \geq \varepsilon_r \quad (100)$$

$$\forall r \in \mathbb{R}_{\geq 0}. U_i[\{s \in I : W_i(s) \leq r\}] \subseteq \mathbb{N} \text{ is bounded.} \quad (101)$$

Quantitative safety

$$\forall s \in I. PV_0(s) \leq V_0(s) \quad (102)$$

$$\forall s \notin I. V_0(s) \geq 1 \quad (103)$$

We apply the results of Section 3 to establish that, over countable Markov chains, *Proof Rule #1* is sound and ε -complete for quantitative reactivity properties, and complete for almost-sure reactivity.

► **Theorem 4.4** (Proof Rule #1 is Sound). *Suppose conditions (94) to (103) hold.*

Then $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) \geq 1 - \mu V_0$.

Proof. We note that *Proof Rule #1* is an instantiation of the proof rules in Theorems 4.1 and 4.2 along with the sound proof rule for almost-sure termination from [40, Proof Rule 3.2 & Lemma 3.4] which are sufficient conditions for the premises of Theorem 3.7. ◀

► **Theorem 4.5** (Proof Rule #1 is ε -Complete). *Suppose that (μ, P) is a Markov chain on a countable state-space with $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = p$ for some $p \in (0, 1]$. Then for every $\varepsilon > 0$, conditions (94) to (103) hold for a suitable supporting invariant and absorbing regions, supermartingale certificates, and ranking functions, with $\mu V_0 \leq 1 - p + \varepsilon$.*

Proof. By Theorem 3.12 there exists a suitable supporting invariant and absorbing regions, such that the almost-sure termination obligation (19) holds, and by the completeness result of [40, Lemma 3.4] the result follows. ◀

► **Theorem 4.6** (Proof Rule #1 is Complete for Almost-Sure Reactivity). *Suppose that (μ, P) is a Markov chain on a countable state-space with $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1$. Then, conditions (94) to (103) hold for suitable supporting invariant, absorbing regions, supermartingale certificates, and ranking functions, with $\mu V_0 = 0$.*

Proof. We invoke an analogous argument to Theorem 4.5, except we invoke Theorem 3.14 instead of Theorem 3.12. ◀

► **Example 4.7** (Proof Rule #1 applied to the Lending Casino). We continue Example 4.3, where we established that $I = \mathbb{Z}$ and $J_1 = \text{Debt}$ form a valid invariant and absorbing region, witnessed by the certificates V_0 and V_1 of Figure 8. To discharge the almost-sure termination obligation (19), Proof Rule #1 requires a supermartingale certificate $W_1: \mathbb{Z} \rightarrow \mathbb{R}_{\geq 0}$ and a ranking function $U_1: \mathbb{Z} \rightarrow \mathbb{N}$ satisfying (97)–(101) with $J_1 = \text{Debt}$ and $B_1 = \emptyset$, for which we exhibit $U_1(s) = \max\{s + 1, 0\}$ and $W_1(s) = \max\{s + 1, 0\}$ (cf. [40, Example 3.5]). ▮

Proof Rule #2

We next instantiate our absorbing-region decomposition using the proof rule for almost-sure termination introduced by [43, Theorem 4.1] (equivalently, [40, Proof Rule 3.3]), and which is known to be complete over countable-state Markov chains [40, Theorem 3.11]. Suppose $(A_1, B_1), \dots, (A_k, B_k)$ is a Streett acceptance condition with $A_i, B_i \subseteq S$ for $i = 1, \dots, k$ on the countable state space $(S, 2^S)$ of a time-homogeneous Markov chain defined by (μ, P) . *Proof Rule #2* requires the existence of

- a supporting invariant $I \subseteq S$ with $V_0: S \rightarrow \mathbb{R}_{\geq 0}$
- absorbing regions $J_1 \subseteq I \setminus A_1, \dots, J_k \subseteq I \setminus A_k$
- supermartingale certificates $V_1, W_1, \dots, V_k, W_k: S \rightarrow \mathbb{R}_{\geq 0}$
- decrease functions $d_1, \dots, d_k: \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{> 0}$, and
- probability functions $p_1, \dots, p_k: \mathbb{R}_{\geq 0} \rightarrow (0, 1]$

such that the following conditions hold:

Qualitative safety

For every $i = 1, \dots, k$:

$$\forall s \notin A_i. PV_i(s) \leq V_i(s) \quad (104)$$

$$\forall s \in A_i. V_i(s) \geq 1 \quad (105)$$

$$\exists \gamma_i > 0 \forall s \in J_i. V_i(s) \leq 1 - \gamma_i \quad (106)$$

Almost-sure termination

For every $i = 1, \dots, k$:

$$\forall s \in I \setminus (B_i \cup J_i). W_i(s) > 0 \quad (107)$$

$$\forall s \in I \cap (B_i \cup J_i). W_i(s) = 0 \quad (108)$$

$$\forall s \in I \setminus (B_i \cup J_i). PW_i(s) \leq W_i(s) \quad (109)$$

$$\forall r \in \mathbb{R}_{>0} \forall s \in I \setminus (B_i \cup J_i).$$

$$(W_i(s) = r) \implies P(s, \{s' : W_i(s') \leq r - d_i(r)\}) \geq p_i(r) \quad (110)$$

$$\forall r_1, r_2 \in \mathbb{R}_{>0}. r_1 \leq r_2 \implies p_i(r_1) \geq p_i(r_2) \quad (111)$$

$$\forall r_1, r_2 \in \mathbb{R}_{>0}. r_1 \leq r_2 \implies d_i(r_1) \geq d_i(r_2) \quad (112)$$

Quantitative safety

$$\forall s \in I. PV_0(s) \leq V_0(s) \quad (113)$$

$$\forall s \notin I. V_0(s) \geq 1 \quad (114)$$

We employ the results of Section 3 to delineate the properties of Proof Rule #2.

► **Theorem 4.8** (Proof Rule #2 is Sound). *Suppose conditions (104) to (114) hold. Then $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) \geq 1 - \mu V_0$.*

Proof. In the same manner as Theorem 4.4, we have invoked an existing proof rule [43, Theorem 4.1] for condition (19), and the proof follows by Theorem 3.7. ◀

► **Theorem 4.9** (Proof Rule #2 is ε -Complete). *Suppose that (μ, P) is a Markov chain on a countable state-space for which $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = p$ for some $p \in (0, 1]$. Then for every $\varepsilon > 0$, conditions (104) to (114) hold for suitable supporting invariant, absorbing regions, supermartingale certificates, decrease and probability functions, with $\mu V_0 \leq 1 - p + \varepsilon$.*

Proof. We apply the same argument as Theorem 4.5, except relying on [43, Theorem 4.1] for the soundness of conditions (107) to (112). ◀

► **Theorem 4.10** (Proof Rule #2 is Complete for Almost-Sure Reactivity). *Suppose that $P_\mu \left(\bigcap_{i=1}^k \text{Fin}(A_i) \cup \text{Inf}(B_i) \right) = 1$. Then, conditions (104) to (114) hold for suitable supporting invariant, absorbing regions, supermartingale certificates, decrease and probability functions, with $\mu V_0 = 0$.*

Proof. Similarly, we apply the same argument as Theorem 4.6, except relying on [40, Theorem 3.11] for the completeness of the conditions (107) to (112). ◀

► **Example 4.11** (Proof Rule #2 applied to the Lending Casino). Continuing Example 4.7 and using $I = \mathbb{Z}$, $J_1 = \text{Debt}$, and the safety certificates of Figure 8, we additionally exhibit $W_1(s) = \max\{s + 1, 0\}$ for conditions (107) to (112) within Proof Rule #2, paired with the constant decrease and probability functions $d_1(r) = 1$ and $p_1(r) = \frac{1}{2}$ (cf. [43, p.8]). ▮

► **Remark 4.12** (Prior Supermartingale Certificates for ω -Regular Properties are Incomplete). We note that all prior supermartingale certificates for ω -regular properties in the literature, which meet the requirement that they are defined in terms of value functions whose codomain is real-valued [3, 4, 12, 32], are inapplicable to proving the almost-sure satisfaction of the reactivity property $\text{Fin}(\text{Solvency})$ explored in Example 4.3, due to the fact that the time to reach *Debt* from any state in *Solvency* has infinite expected value, which is beyond the reach of prior proof rules (cf. [37, Example 4.12]). ▮

5 Related Work

Supermartingale proof rules have been developed as tools to analyse transience and recurrence on aperiodic and irreducible Markov chains, and applied to draw conclusions about their long-term behaviour, such as their stationary distributions (cf. [28, 50] and [45, Theorem 11.0.1]). For countable-state Markov chains that are aperiodic and irreducible, [44] introduced sound and complete supermartingale proof rules for almost-sure recurrence [44, Theorem 1] (cf. [25, Theorem 7.5.2]) and almost-sure transience [44, Theorem 2] (cf. [25, Theorem 7.5.1]). Note that for irreducible aperiodic chains on $\mathbb{N}$, recurrence corresponds to visiting $\{0\}$ infinitely often, and transience to visiting it finitely often.

Although transience and recurrence are closely related to the events $\text{Fin}(A)$ and $\text{Inf}(B)$ that constitute reactivity properties (Definition 2.3), we identify two major reasons why these classic results on necessary and sufficient conditions for almost-sure transience and almost-sure recurrence do not easily transfer to reactivity verification on countably infinite Markov chains. Firstly, they require irreducibility, meaning that a finite path exists between any pair of states. The Markov chain induced by synchronous composition (cf. Remark 2.4) with a deterministic Streett automaton $\mathcal{A}$ fails to be irreducible, because $\mathcal{A}$ in general contains more than one strongly connected component, regardless of the irreducibility of the underlying Markov chain $\hat{\Phi}$. Despite the fact that countable Markov chains admit a partitioning into transient and recurrent components ([45, Section 4.5], [25, Theorem 7.3.3] and [24]), the transient (and thus potentially non-irreducible) component may encompass the entire chain. Secondly, every reactivity property is a Boolean combination of Fin - and Inf -events, and can therefore be satisfied almost surely without any of its constituent Fin - or Inf -events being satisfied almost surely. Proof rules for almost-sure transience and almost-sure recurrence, applied in isolation, are therefore too conservative for reactivity properties represented by deterministic Streett and parity automata [3, Figure 1].

Subsequent works have relaxed the irreducibility assumption of [44] for special cases of almost-sure reactivity. In the case of the reactivity property defined by a single Streett pair $(A, \emptyset)$, [45, Proposition 8.3.1(iv)] defines a proof rule for uniform transience on Markov chains without assuming irreducibility. While this condition is similar in spirit to our requirement in Equation (18), almost-sure uniform transience is strictly stronger than requiring $\text{Fin}(A)$ to hold with probability 1. For instance, the region *Solvency* in the Lending Casino (Figure 1) is not uniformly transient, even though the event $\text{Fin}(\text{Solvency})$ holds almost surely (see Example 4.3). Dually, by considering the reactivity property defined by the Streett pair (S, Halt) – where every state in $\text{Halt} \subseteq S$ is made absorbing – one recovers the property of probabilistic termination. Under this instantiation, our framework subsumes the known result that quantitative termination is equivalent (up to ε -approximation) to a combination of almost-sure termination and quantitative safety obligations (cf. [40, Lemma 4.6], [16], and [4, Example 4]).

While we consider reactivity properties defined upon Markov chains, prior work has treated questions of completeness for recurrence and transience properties defined upon countable MDPs under adversarial and controllable interpretations of non-determinism, as well as for stochastic games [33, 36]. In the case of controllable non-determinism, prior results demonstrate the necessary existence of policies for the attainment of reactivity conditions [34, 35] and show that these may be constructed from optimal policies for simpler obligations (such as reachability, safety, and recurrence). It is an open question as to whether our decomposition results (Section 3) extend to MDPs in any form.

The algorithmic synthesis of supermartingale certificates draws on several research areas: ranking functions and inductive invariants in program analysis [21, 27, 42, 46, 47, 49, 57, 60], Lyapunov functions and barrier certificates in control theory [51, 53, 54, 56, 58], and martingale-like value functions underpinning reinforcement learning [59]. Across these domains, supermartingale certificates are typically expressed as real-valued functions and synthesized either by deductive symbolic reasoning or by inductive learning. Symbolic methods compute linear or polynomial certificates via Farkas' Lemma [5, 11, 14, 20, 22, 23] and Positivstellensatz results [8, 13, 15, 17, 18, 31, 61, 64], reducing synthesis to decision problems in the first-order theory of the reals. Learning-based techniques instead infer certificates from parametric families via counterexample-guided inductive synthesis [6, 10], often using neural networks as their representation [1, 2, 7, 19, 29, 30, 39, 48, 62, 63].

6 Conclusion

We have established that every reactivity property defined upon a Markov chain over a general state space admits a decomposition into multiple absorbing regions and a single global invariant. This decomposition, in turn, yields a strategy for proving probabilistic reactivity that reduces the problem to solving obligations of qualitative safety, quantitative safety, and almost-sure termination. Our technical results in Theorems 3.7 and 3.12 establish that this strategy is sound and ε -complete in the general case, while Theorem 3.14 shows that it is fully complete when the reactivity property holds almost surely.

Our strategy yields the first proof rules involving scalar-valued supermartingale certificates that are both sound and complete for reactivity properties (Section 4); this provides a stepping stone towards automation through the adaptation of existing methods for scalar-valued certificates (Section 5), which remains an open question and a subject for future work. These new supermartingale certificates for reactivity are obtained by instantiating our strategy with existing proof rules for almost-sure termination that are sound and complete in the special case of countably infinite state spaces.

Our decomposition results (Theorems 3.12 and 3.14 and Figure 7) apply to general state spaces, whereas the concrete supermartingale certificates for reactivity we present currently possess completeness guarantees in the setting of countably infinite state spaces. Extending these results to uncountable state spaces is an interesting direction for future work. Our results show that addressing this challenge reduces to that of certifying almost-sure termination over general state spaces in a complete fashion, from which a complete proof rule for reactivity over general state spaces would follow (cf. Figure 7). Moreover, our strategy assumes time homogeneity; extending its scope of applicability to Markov decision processes over general state spaces is another promising direction for future research.

References

- 1 Alessandro Abate, Alec Edwards, Mirco Giacobbe, Hashan Punchihewa, and Diptarko Roy. Quantitative verification with neural networks. *Log. Methods Comput. Sci.*, 22(2), 2026. doi:10.46298/LMCS-22(2:4)2026.
- 2 Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Learning probabilistic termination proofs. In *CAV (2)*, volume 12760 of *Lecture Notes in Computer Science*, pages 3–26. Springer, 2021. doi:10.1007/978-3-030-81688-9_1.
- 3 Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Stochastic omega-regular verification and control with supermartingales. In *CAV (3)*, volume 14683 of *Lecture Notes in Computer Science*, pages 395–419. Springer, 2024. doi:10.1007/978-3-031-65633-0_18.

- 4 Alessandro Abate, Mirco Giacobbe, and Diptarko Roy. Quantitative supermartingale certificates. In *CAV (2)*, volume 15932 of *Lecture Notes in Computer Science*, pages 3–28. Springer, 2025. doi:10.1007/978-3-031-98679-6_1.
- 5 Sheshansh Agrawal, Krishnendu Chatterjee, and Petr Novotný. Lexicographic ranking supermartingales: an efficient approach to termination of probabilistic programs. *Proc. ACM Program. Lang.*, 2(POPL):34:1–34:32, 2018. doi:10.1145/3158122.
- 6 Daniele Ahmed, Andrea Peruffo, and Alessandro Abate. Automated and sound synthesis of lyapunov functions with SMT solvers. In *TACAS (1)*, volume 12078 of *Lecture Notes in Computer Science*, pages 97–114. Springer, 2020. doi:10.1007/978-3-030-45190-5_6.
- 7 Matin Ansaripour, Krishnendu Chatterjee, Thomas A. Henzinger, Mathias Lechner, and Đorđe Žikelić. Learning provably stabilizing neural controllers for discrete-time stochastic systems. In *ATVA (1)*, volume 14215 of *Lecture Notes in Computer Science*, pages 357–379. Springer, 2023. doi:10.1007/978-3-031-45329-8_17.
- 8 Ali Asadi, Krishnendu Chatterjee, Hongfei Fu, Amir Kafshdar Goharshady, and Mohammad Mahdavi. Polynomial reachability witnesses via stellensätze. In *PLDI*, pages 772–787. ACM, 2021. doi:10.1145/3453483.3454076.
- 9 S. Axler. *Measure, Integration & Real Analysis*. Graduate Texts in Mathematics. Springer International Publishing, 2019.
- 10 Kevin Batz, Mingshuai Chen, Sebastian Junges, Benjamin Lucien Kaminski, Joost-Pieter Katoen, and Christoph Matheja. Probabilistic program verification via inductive synthesis of inductive invariants. In *TACAS (2)*, volume 13994 of *Lecture Notes in Computer Science*, pages 410–429. Springer, 2023. doi:10.1007/978-3-031-30820-8_25.
- 11 Aleksandar Chakarov and Sriram Sankaranarayanan. Probabilistic program analysis with martingales. In *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 511–526. Springer, 2013. doi:10.1007/978-3-642-39799-8_34.
- 12 Aleksandar Chakarov, Yuen-Lam Voronin, and Sriram Sankaranarayanan. Deductive proofs of almost sure persistence and recurrence properties. In *TACAS*, volume 9636 of *Lecture Notes in Computer Science*, pages 260–279. Springer, 2016. doi:10.1007/978-3-662-49674-9_15.
- 13 Krishnendu Chatterjee, Hongfei Fu, and Amir Kafshdar Goharshady. Termination analysis of probabilistic programs through positivstellensatz’s. In *CAV (1)*, volume 9779 of *Lecture Notes in Computer Science*, pages 3–22. Springer, 2016. doi:10.1007/978-3-319-41528-4_1.
- 14 Krishnendu Chatterjee, Hongfei Fu, Petr Novotný, and Rouzbeh Hasheminezhad. Algorithmic analysis of qualitative and quantitative termination problems for affine probabilistic programs. *ACM Trans. Program. Lang. Syst.*, 40(2):7:1–7:45, 2018. doi:10.1145/3174800.
- 15 Krishnendu Chatterjee, Amir Kafshdar Goharshady, Ehsan Kafshdar Goharshady, Mehrdad Karrabi, and Đorđe Žikelić. Sound and complete witnesses for template-based verification of LTL properties on polynomial programs. In *FM (1)*, volume 14933 of *Lecture Notes in Computer Science*, pages 600–619. Springer, 2024. doi:10.1007/978-3-031-71162-6_31.
- 16 Krishnendu Chatterjee, Amir Kafshdar Goharshady, Tobias Meggendorfer, and Đorđe Žikelić. Sound and complete certificates for quantitative termination analysis of probabilistic programs. In *CAV (1)*, *Lecture Notes in Computer Science*, pages 55–78. Springer, 2022. doi:10.1007/978-3-031-13185-1_4.
- 17 Krishnendu Chatterjee, Ehsan Kafshdar Goharshady, Mehrdad Karrabi, Harshit J. Motwani, Maximilian Seeliger, and Đorđe Žikelić. Quantified linear and polynomial arithmetic satisfiability via template-based skolemization. In *AAAI*, pages 7326–7336. AAAI Press, 2025.
- 18 Krishnendu Chatterjee, Ehsan Kafshdar Goharshady, Petr Novotný, and Đorđe Žikelić. Equivalence and similarity refutation for probabilistic programs. *Proc. ACM Program. Lang.*, 8(PLDI):2098–2122, 2024. doi:10.1145/3656462.
- 19 Krishnendu Chatterjee, Thomas A. Henzinger, Mathias Lechner, and Đorđe Žikelić. A learner-verifier framework for neural network controllers and certificates of stochastic systems. In *TACAS (1)*, volume 13993 of *Lecture Notes in Computer Science*, pages 3–25. Springer, 2023. doi:10.1007/978-3-031-30823-9_1.

- 20 Krishnendu Chatterjee, Petr Novotný, and Đorđe Źikelić. Stochastic invariants for probabilistic termination. In *POPL*, pages 145–160. ACM, 2017. doi:10.1145/3009837.3009873.
- 21 Michael Colón, Sriram Sankaranarayanan, and Henny Sipma. Linear invariant generation using non-linear constraint solving. In *CAV*, volume 2725 of *Lecture Notes in Computer Science*, pages 420–432. Springer, 2003. doi:10.1007/978-3-540-45069-6_39.
- 22 Michael Colón and Henny Sipma. Synthesis of linear ranking functions. In *TACAS*, volume 2031 of *Lecture Notes in Computer Science*, pages 67–81. Springer, 2001. doi:10.1007/3-540-45319-9_6.
- 23 Michael Colón and Henny Sipma. Practical methods for proving program termination. In *CAV*, volume 2404 of *Lecture Notes in Computer Science*, pages 442–454. Springer, 2002. doi:10.1007/3-540-45657-0_36.
- 24 W. Doblin. éléments d’une théorie générale des chaînes simples constantes de Markoff. *Annales scientifiques de l’École Normale Supérieure*, 3e série, 57:61–111, 1940. doi:10.24033/asens.883.
- 25 R. Douc, E. Moulines, P. Priouret, and P. Soulier. *Markov Chains*. Springer Series in Operations Research and Financial Engineering. Springer, 2018.
- 26 Rick Durrett. *Probability: Theory and Examples, 4th Edition*. Cambridge University Press, 2010. doi:10.1017/CB09780511779398.
- 27 Michael D. Ernst, Jeff H. Perkins, Philip J. Guo, Stephen McCamant, Carlos Pacheco, Matthew S. Tschantz, and Chen Xiao. The daikon system for dynamic detection of likely invariants. *Sci. Comput. Program.*, 69(1-3):35–45, 2007. doi:10.1016/J.SCICO.2007.01.015.
- 28 F. G. Foster. On the Stochastic Matrices Associated with Certain Queuing Processes. *The Annals of Mathematical Statistics*, 24(3):355–360, 1953. doi:10.1214/aoms/1177728976.
- 29 Mirco Giacobbe, Daniel Kroening, Abhinandan Pal, and Michael Tautschnig. Neural model checking. In *NeurIPS*, 2024.
- 30 Mirco Giacobbe, Daniel Kroening, Abhinandan Pal, and Michael Tautschnig. Let a neural network be your invariant. In *NeurIPS*, 2025.
- 31 David Handelman. Representing polynomials by positive linear functions on compact convex polyhedra. *Pacific Journal of Mathematics*, 132, March 1988.
- 32 Thomas A. Henzinger, Kaushik Mallik, Pouya Sadeghi, and Đorđe Źikelić. Supermartingale certificates for quantitative omega-regular verification and control. In *CAV (2)*, volume 15932 of *Lecture Notes in Computer Science*, pages 29–55. Springer, 2025. doi:10.1007/978-3-031-98679-6_2.
- 33 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Büchi objectives in countable mdps. In *ICALP*, volume 132 of *LIPIcs*, pages 119:1–119:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.119.
- 34 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Strategy complexity of parity objectives in countable MDPs. In *CONCUR*, volume 171 of *LIPIcs*, pages 39:1–39:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CONCUR.2020.39.
- 35 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Transience in countable mdps. In *CONCUR*, volume 203 of *LIPIcs*, pages 11:1–11:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPIcs.CONCUR.2021.11.
- 36 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Strategy complexity of Büchi and transience objectives in concurrent stochastic games. In *EC*, pages 390–409. ACM, 2025. doi:10.1145/3736252.3742560.
- 37 Satoshi Kura and Hiroshi Unno. A hierarchy of supermartingales for ω -regular verification. *CoRR*, abs/2512.00270, 2025. doi:10.48550/arXiv.2512.00270.
- 38 H.J. Kushner. *Stochastic Stability and Control*. Mathematics in science and engineering. Academic Press, 1967.

- 39 Mathias Lechner, Đorđe Žikelić, Krishnendu Chatterjee, and Thomas A. Henzinger. Stability verification in stochastic control systems via neural network supermartingales. In *AAAI*, pages 7326–7336. AAAI Press, 2022. doi:10.1609/AAAI.V36I7.20695.
- 40 Rupak Majumdar and V. R. Sathiyarayanan. Sound and complete proof rules for probabilistic termination. *Proc. ACM Program. Lang.*, 9(POPL):1871–1902, 2025. doi:10.1145/3704899.
- 41 Zohar Manna and Amir Pnueli. A hierarchy of temporal properties. In *PODC*, pages 377–410. ACM, 1990. doi:10.1145/93385.93442.
- 42 Zohar Manna and Amir Pnueli. *Temporal verification of reactive systems - safety*. Springer, 1995.
- 43 Annabelle McIver, Carroll Morgan, Benjamin Lucien Kaminski, and Joost-Pieter Katoen. A new proof rule for almost-sure termination. *Proc. ACM Program. Lang.*, 2(POPL):33:1–33:28, 2018. doi:10.1145/3158121.
- 44 Jean-François Mertens, Ester Samuel-Cahn, and Shmuel Zamir. Necessary and sufficient conditions for recurrence and transience of markov chains, in terms of inequalities. *Journal of Applied Probability*, 15(4):848–851, 1978. URL: <http://www.jstor.org/stable/3213440>.
- 45 Sean Meyn and Richard L. Tweedie. *Markov Chains and Stochastic Stability*. Cambridge Mathematical Library. Cambridge University Press, 2 edition, 2009.
- 46 Marcel Moosbrugger, Ezio Bartocci, Joost-Pieter Katoen, and Laura Kovács. The probabilistic termination tool amber. In Marieke Huisman, Corina Păsăreanu, and Naijun Zhan, editors, *Formal Methods*, pages 667–675, Cham, 2021. Springer International Publishing. doi:10.1007/978-3-030-90870-6_36.
- 47 Marcel Moosbrugger, Miroslav Stankovic, Ezio Bartocci, and Laura Kovács. This is the moment for probabilistic loops. *Proc. ACM Program. Lang.*, 6(OOPSLA2):1497–1525, 2022. doi:10.1145/3563341.
- 48 Grigory Neustroev, Mirco Giacobbe, and Anna Lukina. Neural continuous-time supermartingale certificates. In *AAAI*. AAAI Press, 2025. doi:10.1609/AAAI.V39I26.34966.
- 49 ThanhVu Nguyen, Deepak Kapur, Westley Weimer, and Stephanie Forrest. DIG: A dynamic invariant generator for polynomial and array invariants. *ACM Trans. Softw. Eng. Methodol.*, 23(4):30:1–30:30, 2014. doi:10.1145/2556782.
- 50 A. G. Pakes. Some conditions for ergodicity and recurrence of markov chains. *Operations Research*, 17(6):1058–1061, 1969. doi:10.1287/OPRE.17.6.1058.
- 51 Antonis Papachristodoulou and Stephen Prajna. On the construction of lyapunov functions using the sum of squares decomposition. In *CDC*, pages 3482–3487. IEEE, 2002. doi:10.1109/CDC.2002.1184414.
- 52 David Pollard. *A User’s Guide to Measure Theoretic Probability*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2001.
- 53 Stephen Prajna, Ali Jadbabaie, and George J. Pappas. Stochastic safety verification using barrier certificates. In *CDC*, pages 929–934. IEEE, 2004. doi:10.1109/CDC.2004.1428804.
- 54 Stephen Prajna, Antonis Papachristodoulou, and Pablo A. Parrilo. Introducing SOSTOOLS: a general purpose sum of squares programming solver. In *CDC*, pages 741–746. IEEE, 2002. doi:10.1109/CDC.2002.1184594.
- 55 Shmuel Safra. On the complexity of omega-automata. In *FOCS*, pages 319–327. IEEE Computer Society, 1988. doi:10.1109/SFCS.1988.21948.
- 56 Sriram Sankaranarayanan, Xin Chen, and Erika Ábrahám. Lyapunov Function Synthesis Using Handelman Representations. *IFAC Proceedings Volumes*, 46(23):576–581, 2013. doi:10.3182/20130904-3-FR-2041.00198.
- 57 Sriram Sankaranarayanan, Henny B. Sipma, and Zohar Manna. Constraint-based linear-relations analysis. In *SAS*, volume 3148 of *Lecture Notes in Computer Science*, pages 53–68. Springer, 2004. doi:10.1007/978-3-540-27864-1_7.
- 58 Zhikun She, Haoyang Li, Bai Xue, Zhiming Zheng, and Bican Xia. Discovering Polynomial Lyapunov Functions for Continuous Dynamical Systems. *Journal of Symbolic Computation*, 58:41–63, November 2013. doi:10.1016/J.JSC.2013.06.003.

- 59 Richard S. Sutton and Andrew G. Barto. *Reinforcement learning - an introduction*. Adaptive computation and machine learning. MIT Press, 1998.
- 60 Toru Takisaka, Yuichiro Oyabu, Natsuki Urabe, and Ichiro Hasuo. Ranking and repulsing supermartingales for reachability in randomized programs. *ACM Trans. Program. Lang. Syst.*, 43(2):5:1–5:46, 2021. doi:10.1145/3450967.
- 61 Đorđe Žikelić, Bor-Yuh Evan Chang, Pauline Bolignano, and Franco Raimondi. Differential cost analysis with simultaneous potentials and anti-potentials. In *PLDI*, pages 442–457. ACM, 2022. doi:10.1145/3519939.3523435.
- 62 Đorđe Žikelić, Mathias Lechner, Thomas A. Henzinger, and Krishnendu Chatterjee. Learning control policies for stochastic systems with reach-avoid guarantees. In *AAAI*, pages 11926–11935. AAAI Press, 2023. doi:10.1609/AAAI.V37I10.26407.
- 63 Đorđe Žikelić, Mathias Lechner, Abhinav Verma, Krishnendu Chatterjee, and Thomas A. Henzinger. Compositional policy learning in stochastic control systems with formal guarantees. In *NeurIPS*, 2023.
- 64 Peixin Wang, Tengshun Yang, Hongfei Fu, Guanyan Li, and C.-H. Luke Ong. Static posterior inference of bayesian probabilistic programming via polynomial solving. *Proc. ACM Program. Lang.*, 8(PLDI):1361–1386, 2024. doi:10.1145/3656432.