

Monitoring Discounted Sum Properties

Filip Cano ✉

Institute of Science and Technology Austria

Thomas A. Henzinger ✉

Institute of Science and Technology Austria

Konstantin Kueffner ✉

Institute of Science and Technology Austria

N. Ege Saraç ✉

CISPA Helmholtz Center for Information Security, Germany

Abstract

Runtime monitoring of quantitative signals faces a fundamental trade-off between volatility and over-aggregation: instantaneous observations are noisy, while long-run averages obscure local structure. Localisation measures such as discounted averages offer a principled middle ground, yet remain poorly understood in runtime verification. This paper studies discounted sums from a monitoring perspective, in both deterministic and stochastic settings. We formalize the discounted monitoring problem and show that exact, sound monitoring of discounted sums cannot be achieved with finite memory. To overcome this impossibility, we introduce ε -approximately sound monitoring, deriving explicit bounds on memory and observation requirements. We then extend the framework to stochastic processes via expected discounted sums, defining pointwise and uniform (ε, δ) -soundness notions, establishing statistical optimality, and proving impossibility beyond a precision threshold. We also formalize the resource complexity of deterministic discounted monitoring via affine register machines and prove a tight worst-case lower bound. Finally, we present a specification language for arithmetic expressions over multiple discounted sums with synchronous and asynchronous semantics, and evaluate our approach on practical scenarios including algorithmic fairness.

2012 ACM Subject Classification Software and its engineering → Dynamic analysis; Mathematics of computing → Stochastic processes; Software and its engineering → Specification languages; Theory of computation → Quantitative automata

Keywords and phrases Runtime Verification, Probabilistic Systems, Quantitative Verification, Approximate Monitoring

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2026.32

Related Version *Full Version*: <https://doi.org/10.48550/arXiv.2606.25979>

Supplementary Material *Software (Source Code)*: <https://github.com/filipcano/monitoring-discounted-sum-properties>

Funding This work has been supported by the European Research Council under Grant No.: ERC-2020-AdG 101020093 and ERC-2021-AdG 101055412.

Filip Cano: ERC-2020-AdG 101020093

Thomas A. Henzinger: ERC-2020-AdG 101020093

Konstantin Kueffner: ERC-2020-AdG 101020093

N. Ege Saraç: ERC-2021-AdG 101055412

1 Introduction

Runtime monitoring is a central problem in modern computer science, with applications across systems, networks, machine learning, security, and cyber-physical systems [5, 34]. Traditionally, it has focused either on instantaneous properties or long-run behaviors. However, for quantitative signals, such as CPU temperature, model accuracy, or algorithmic fairness,



© Filip Cano, Thomas A. Henzinger, Konstantin Kueffner, N. Ege Saraç;
licensed under Creative Commons License CC-BY 4.0

37th International Conference on Concurrency Theory (CONCUR 2026).

Editors: Ana Sokolova and Patrick Totzke; Article No. 32; pp. 32:1–32:45

Leibniz International Proceedings in Informatics



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

there is a tension: single observations are noisy, while long-run aggregates can obscure local structure. We illustrate this trade-off with the following motivating examples.

- (i) *CPU temperature*: monitoring the temperature of a CPU is essential to prevent hardware damage; the conundrum: a single spike that dissipates quickly may raise an undue alarm, yet aggregating over long stretches of time can hide sustained overheating.
- (ii) *Model accuracy*: monitoring a model's accuracy is required to ensure sustained performance as distributions shift; the conundrum: one bad prediction is not indicative of the model's performance, yet a long-run average may fail to adapt to distribution shifts.
- (iii) *Algorithmic fairness*: monitoring the fairness of an arbiter may uncover uneven issuing of grants between two clients: in isolation, every grant issued to one client is maximally unfair to the other, yet a limit average can hide the deliberate starvation of a client.

A natural way to resolve this conundrum is to use *localisation measures*, such as the window average, discounted average, or, more generally, kernel-weighted averages. These are essentially convolutions that smooth volatile signals by weighting values around a given time point inversely with their distance. While common in time series analysis, signal processing, and image classification, localisation measures are underexplored in runtime verification. In this paper, we study a (seemingly) trivial and widely used localisation measure—the discounted average—from a runtime verification perspective. Discounted monitoring faces two fundamentally different sources of uncertainty: epistemic uncertainty from unseen past and future values, and statistical uncertainty from noisy observations. This paper treats both explicitly and shows that each leads to distinct impossibility thresholds.

Deterministic Discounted Monitoring. We assume a bi-infinite stream of numerical values $w = (x_i)_{i \in \mathbb{Z}}$. Centered at time index $t \in \mathbb{Z}$, we consider a bi-directional discounted sum, with discount factor $r \in [0, 1)$ for the past and $s \in [0, 1)$ for the future:

$$S_t^{r,s}(w) := \underbrace{\sum_{i=1}^{\infty} r^i \cdot x_{t-i}}_{\text{past}} + \underbrace{x_t}_{\text{present}} + \underbrace{\sum_{i=1}^{\infty} s^i \cdot x_{t+i}}_{\text{future}}. \quad (1)$$

► **Example 1 (Decision Fairness).** The stream w is a sequence of binary decisions, where $x_t = 1$ indicates that client 1 was granted access to the resource at time t , while $x_t = 0$ indicates the same for client 0. Suppose we are at time t and client 0 was granted the resource (so $x_t = 0$). Normalising the discounted sum by $\lambda^{r,s} := \frac{r}{1-r} + \frac{s}{1-s} + 1$, we obtain the discounted average. Then $S_t^{r,s}(w)/\lambda^{r,s}$ quantifies how fairly the arbiter treated the clients around time t , where 0.5 indicates perfect fairness. Setting $s = 0$ yields a purely backward-looking notion of fairness, whereas setting $r = 0$ yields a forward-looking notion. In particular, for forward-looking fairness, if the arbiter initially favours client 1, the discount factor dictates how much it must compensate later to remain fair, if possible at all.

In discounted monitoring, time progresses and we slide the discounted sum across the value stream, yielding the sequence of discounted property values. Without loss of generality, our monitors start observing the stream at time 0. After each new observation at time n , the monitor must assess, for every past index $t \in [0, n]$, whether the discounted sum lies inside or outside a target interval $\mathcal{I}$:

$$\begin{array}{c} \cdots, S_{t-2}^{r,s}(w), S_{t-1}^{r,s}(w), S_t^{r,s}(w), S_{t+1}^{r,s}(w), S_{t+2}^{r,s}(w), \cdots \\ \quad \quad \quad \uparrow \text{bound error} \\ w_0, w_1, \cdots, w_{n-1}, w_n \rightarrow \text{Monitor.} \rightarrow (\in \mathcal{I}?)\text{-verdict} \end{array}$$

The fundamental challenge is *epistemic uncertainty*: the monitor must reason about a function over a bi-infinite stream after observing only n values. Within this setting, we propose a monitor that maintains, for each time index t , an *uncertainty set* containing all possible values the full discounted sum at time t could take. Although the uncertainty sets shrink as more information becomes available, we show that the time it takes to shrink enough to be completely inside or outside the target interval is unbounded. As a solution, we study monitors with a tolerance of ε in their verdict, and obtain an upper bound on the number of observations required to reach a decisive verdict, enabling monitoring with bounded resources at the cost of precision. We call these ε -approximate monitors.

Statistical Discounted Monitoring. In many applications, the value stream is a realisation of a stochastic process $W = (X_i)_{i \in \mathbb{Z}}$. In such settings, we are often interested in expected quantities rather than noisy observations, e.g., the latent temperature or the arbiter's underlying propensity. This leads to our second object of study: the *expected discounted sum*, defined as the bi-directional discounted sum evaluated over conditional expectations,

$$\mathbb{S}_t^{r,s}(W) := \underbrace{\sum_{i=1}^{\infty} r^i \mathbb{E}_{t-i-1}(X_{t-i})}_{\text{past}} + \underbrace{\mathbb{E}_{t-1}(X_t)}_{\text{present}} + \underbrace{\sum_{i=1}^{\infty} s^i \mathbb{E}_{t+i-1}(X_{t+i})}_{\text{future}}. \quad (2)$$

► **Example 2 (Bias Fairness).** Consider a stochastic arbiter that generates its decision X_t at time t by tossing a coin with probability P_t , where P_t is chosen based on the history up to time t . This yields two processes: the hidden bias process $U = (P_t)_{t \in \mathbb{Z}}$ and the observed outcome process $W = (X_t)_{t \in \mathbb{Z}}$, with $\mathbb{E}_{t-1}(X_t) = P_t$. If we want to assess the arbiter around time t independently of chance, we should evaluate discounted sums (or averages) with respect to the bias process, e.g., $\mathbb{S}_t^{r,s}(U) = \mathbb{S}_t^{r,s}(W)$ rather than $S_t^{r,s}(W)$.

In statistical discounted monitoring, we slide both the expected and observed discounted sums over the stochastic stream, yielding expected and observed discounted property values. As before, monitors start observing the realised sequence at time 0. After each new observation at time n , the monitor must assess, for every past index $t \in [0; n]$, whether the *expected* discounted sum lies inside or outside a target interval $\mathcal{I}$ with high probability.

$$\begin{array}{ccccccc} \cdots, & \mathbb{S}_{t-2}^{r,s}(W), & \mathbb{S}_{t-1}^{r,s}(W), & \mathbb{S}_t^{r,s}(W), & \mathbb{S}_{t+1}^{r,s}(W), & \mathbb{S}_{t+2}^{r,s}(W), & \cdots \\ & & & \uparrow \text{bound deviation} & & & \\ \cdots, & S_{t-2}^{r,s}(W), & S_{t-1}^{r,s}(W), & S_t^{r,s}(W), & S_{t+1}^{r,s}(W), & S_{t+2}^{r,s}(W), & \cdots \\ & & & \uparrow \text{bound error} & & & \\ W_0, W_1, \cdots, W_{n-1}, W_n & \rightarrow \text{Monitor.} & \rightarrow (\in \mathcal{I}?)\text{-verdict} \end{array}$$

Apart from epistemic uncertainty, in this setting the extra challenge is *statistical inference*: the monitor has to estimate also the deviation between the observed and expected discounted sums. Because of the probabilistic nature of the monitored quantities, the monitor's verdict is accompanied by an error probability δ . Within this setting, we extend ε -approximate monitoring to the probabilistic setting, and define three notions of (ε, δ) -soundness with increasing strength of guarantees: pointwise, local, and uniform; depending on whether the probability of an error is given for each point after a fixed release time (pointwise), for each point with flexible release time (local), or for the whole execution (uniform). As the strength of the guarantee increases, the monitor's precision decreases. We show that the deviation bounds used by our pointwise sound monitors are minimax optimal, which establishes that statistical discounted monitoring is impossible beyond a certain precision.

Generalization to Arithmetic Expressions. Many relevant quantitative properties are naturally expressed as relations between discounted sums rather than as single aggregates. To support discounted versions of such properties, we introduce a richer specification language in which a monitor observes multiple outcome streams and monitors an arithmetic expression over their discounted sums. We distinguish between synchronous and asynchronous semantics for combining discounted sums.

► **Example 3 (Demographic Parity).** So far we have assumed that both clients always request access to the resource. We model the more general setting using the binary sequences representing: w^{r_1} the requests of client 1, w^{g_1} the grants client 1, w^{r_2} the requests client 2, and w^{g_2} the grants client 2. All indexed by the same global clock. Then we can express the discounted acceptance rate, i.e., a discounted form of demographic parity [15], as the difference between the ratios $S_t^{r,s}(w^{g_1})/S_t^{r,s}(w^{r_1})$ and $S_t^{r,s}(w^{g_2})/S_t^{r,s}(w^{r_2})$.

Register Complexity. The ε -approximate monitor from the deterministic setting yields a finite observation horizon; let τ denote the number of additional observations after which each monitored position receives a sound verdict up to the ε boundary region. This horizon is first a delay bound: a verdict for time t may require waiting until time $t + \tau$. It also serves as a memory bound: during this waiting period, several positions may remain unresolved simultaneously, and the monitor maintains a separate running discounted sum for each of them. The delay bound, however, only says when verdicts are guaranteed to arrive; it does not characterize how much information must be stored before those verdicts are reached. This leads to the register-complexity question: can the overlapping discounted sums in general be represented more compactly than by one running sum per pending position?

To make this question precise, we formalize monitors as *affine register machines (ARMs)*: finite-state monitors equipped with real-valued registers, affine updates, and linear guards. In this model, we answer this question negatively by proving a worst-case lower bound already for future-only monitoring. For every k , there is a future-only monitor with horizon k that cannot be monitored by any ARM with fewer than k registers.

Experimental Evaluation. Finally, we implement our monitors on scenarios mirroring the motivating examples, and investigate empirically how actual register use compares to our theoretical upper bounds, and how register usage differs between synchronous and asynchronous interpretations of arithmetic properties.

Contributions. The main contributions of this paper are as follows:

1. A formalization of the discounted monitoring problem, and a proof that exact, sound discounted monitoring cannot be achieved with finite memory.
2. An ε -approximately sound notion of discounted monitoring, accompanied by explicit bounds on the required number of observations and registers.
3. An extension of discounted monitoring to stochastic processes via expected discounted sums, including pointwise and uniform (ε, δ) -soundness notions, as well as results on statistical optimality and impossibility beyond a precision threshold.
4. A specification language for arithmetic expressions over multiple discounted sums, supporting both synchronous and asynchronous semantics.
5. A formalization of monitors as affine register machines and a tight bound on the number of registers required for sound ε -approximate monitoring in the future-only case.
6. An implementation and empirical evaluation of the proposed monitors on realistic case studies, analyzing performance and efficiency relative to the proven worst-case guarantees.

2 Monitoring Discounted Sums

2.1 Discounted Sum Property

The main object of study in this paper is the discounted sum operator $S_t^{r,s}$, where $r, s \in [0, 1]$ are the past and future discount factors and $t \in \mathbb{Z}$ is a time index. It acts on bi-infinite sequences $w = (x_t)_{t \in \mathbb{Z}}$ taking values in a nonempty bounded set $\mathcal{R} \subset \mathbb{R}$ of diameter $d_{\mathcal{R}} := \sup \mathcal{R} - \inf \mathcal{R}$. We call a tuple $(\mathcal{R}, r, s)$ an *input setting*. Intuitively, $S_t^{r,s}$ is centered at t , discounting past values ($i < t$) by r and future values ($j > t$) by s . Formally, for $n_l \leq t \leq n_u$, let $w_{n_l:n_u} = (x_{n_l}, \dots, x_{n_u})$ and define $S_t^{r,s}(w_{n_l:n_u}) := \sum_{i=1}^{t-n_l} r^i x_{t-i} + x_t + \sum_{i=1}^{n_u-t} s^i x_{t+i}$. For a bi-infinite sequence w , define $S_t^{r,s}(w) := \sum_{i=1}^{\infty} r^i x_{t-i} + x_t + \sum_{i=1}^{\infty} s^i x_{t+i}$.

Basic properties. We prove some basic properties of discounted sums. First, the discounted sum converges, and consecutive sums differ by at most $d_{\mathcal{R}}$. This follows from the definition.

► **Lemma 4.** *Let $(\mathcal{R}, r, s)$ be an input setting, $w \in \mathcal{R}^{\mathbb{Z}}$, and $t \in \mathbb{Z}$. Then $S_t^{r,s}(w)$ exists and $|S_t^{r,s}(w) - S_{t-1}^{r,s}(w)| \leq d_{\mathcal{R}}$.*

Second, we characterize the values a discounted sum can take. Formally, let $\text{Sums}_{\mathcal{R}}^{r,s} := \{S_t^{r,s}(w) : w \in \mathcal{R}^{\mathbb{Z}}\}$; note that this set does not depend on t . Its shape depends on whether $\mathcal{R}$ has gaps that are too large relative to the discount factors.

► **Definition 5 (Proper input setting).** *Let $(\mathcal{R}, r, s)$ be an input setting. The normalization factor $\lambda^{r,s}$ and the maximum gap $\Delta_{\mathcal{R}}$ are defined as*

$$\lambda^{r,s} := 1 + \frac{r}{1-r} + \frac{s}{1-s}, \quad \Delta_{\mathcal{R}} := \sup \left(\{0\} \cup \{b-a : [a, b] \subseteq [\inf \mathcal{R}, \sup \mathcal{R}] \setminus \mathcal{R}\} \right). \quad (3)$$

We say that $(\mathcal{R}, r, s)$ is proper if $\Delta_{\mathcal{R}} \leq \min \left\{ \frac{r}{1-r}, \frac{s}{1-s} \right\} d_{\mathcal{R}}$.

The convention $\sup \{0\} = 0$ ensures $\Delta_{\mathcal{R}} = 0$ whenever $\mathcal{R}$ has no gaps in its interval hull, so interval-valued domains are automatically proper. For non-interval domains such as $\mathcal{R} = \{0, 1\}$, properness is a substantive condition: it requires the discount factors to be large enough to interpolate across the gaps.

► **Lemma 6.** *Let $(\mathcal{R}, r, s)$ be an input setting.*

- (1) $\text{Sums}_{\mathcal{R}}^{r,s} \subseteq \lambda^{r,s} \cdot [\inf \mathcal{R}, \sup \mathcal{R}]$.
- (2) If $(\mathcal{R}, r, s)$ is proper, then $\text{Closure}(\text{Sums}_{\mathcal{R}}^{r,s}) = \lambda^{r,s} \cdot [\inf \mathcal{R}, \sup \mathcal{R}]$.
- (3) If $(\mathcal{R}, r, s)$ is proper and $\mathcal{R}$ is compact, then $\text{Sums}_{\mathcal{R}}^{r,s} = \lambda^{r,s} \cdot [\inf \mathcal{R}, \sup \mathcal{R}]$.

The three statements isolate the role of each assumption. Boundedness of $\mathcal{R}$ gives the interval over-approximation in (1). Properness ensures that the closure of the set of sums fills this interval, as stated in (2). Compactness lifts closure equality to equality in (3). For a non-closed domain, the endpoints need not be attainable; under compactness they are attained by the constant sequences at $\inf \mathcal{R}$ and $\sup \mathcal{R}$.

Sums vs. Averages. In some settings it is more natural to think of a discounted average than a discounted sum. A weighted average is a weighted sum whose weights are non-negative and sum to 1. Since the weights of a discounted sum add up to $\lambda^{r,s}$, an input setting $(\mathcal{R}, r, s)$ can be reframed for averages by considering the equivalent setting $(\mathcal{R}/\lambda^{r,s}, r, s)$ and rescaling each sequence $w = (x_t)_{t \in \mathbb{Z}} \in \mathcal{R}^{\mathbb{Z}}$ to $w' = (x_t/\lambda^{r,s})_{t \in \mathbb{Z}} \in (\mathcal{R}/\lambda^{r,s})^{\mathbb{Z}}$.

2.2 Monitoring Problem

A monitor observes a bi-infinite stream of values $w = (x_i)_{i \in \mathbb{Z}}$ from time 0 onward. After observing a finite prefix $w_{0:n}$, it is queried only on observed positions $t \in \{0, \dots, n\}$. Formally, a monitor computes a function $\mathcal{M}: \mathcal{R}^* \times \mathbb{N} \rightarrow \{\perp, ?, \top\}$. Only pairs (u, t) with $t < |u|$ are queried. The verdicts $\perp$ and $\top$ indicate that $S_t^{r,s}(w)$ lies outside or inside the target interval $\mathcal{I}$, respectively, and $?$ indicates that the verdict is inconclusive.

2.2.1 Sound Monitors

► **Definition 7** (Sound monitor). *Let $(\mathcal{R}, r, s)$ be an input setting and let $\mathcal{I}$ be a target interval. A monitor is sound if, for every $w \in \mathcal{R}^{\mathbb{Z}}$, $n \in \mathbb{N}$, and $t \in \{0, \dots, n\}$,*

$$\mathcal{M}(w_{0:n}, t) = \perp \implies S_t^{r,s}(w) \notin \mathcal{I} \quad \text{and} \quad \mathcal{M}(w_{0:n}, t) = \top \implies S_t^{r,s}(w) \in \mathcal{I}. \quad (4)$$

Uncertainty. Since the monitor receives values incrementally, it typically returns an inconclusive verdict ($?$) until the observed prefix contains enough information to commit to a final $\top$ or $\perp$. For $0 \leq t \leq n$, the *uncertainty set* of a finite prefix $w_{0:n}$ collects the values $S_t^{r,s}$ can take across all extensions of $w_{0:n}$: $U_t(w_{0:n}) := \{S_t^{r,s}(w') : w' \in \mathcal{R}^{\mathbb{Z}}, w'_{0:n} = w_{0:n}\}$. A sound monitor therefore satisfies $\mathcal{M}(w_{0:n}, t) = \top \implies U_t(w_{0:n}) \subseteq \mathcal{I}$ and $\mathcal{M}(w_{0:n}, t) = \perp \implies U_t(w_{0:n}) \cap \mathcal{I} = \emptyset$. Similarly as in Lemma 6, for proper input settings the closure of the uncertainty set is an interval; the uncertainty set itself is that interval when $\mathcal{R}$ is compact.

► **Lemma 8.** *Let $(\mathcal{R}, r, s)$ be an input setting, let $0 \leq t \leq n$, and let $\gamma_{t,n}^{r,s} := \frac{r^{t+1}}{1-r} + \frac{s^{n-t+1}}{1-s}$. Then, $U_t^{r,s}(w_{0:n}) \subseteq S_t^{r,s}(w_{0:n}) + \gamma_{t,n}^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. If $(\mathcal{R}, r, s)$ is proper, then $\text{Closure}(U_t^{r,s}(w_{0:n})) = S_t^{r,s}(w_{0:n}) + \gamma_{t,n}^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. If, additionally, $\mathcal{R}$ is compact, then equality holds without taking closures.*

A natural quantity is then how long a sound monitor must wait before it can commit to a final verdict at a given time.

► **Definition 9.** *Let $\mathcal{I}$ be a target interval, let $t \in \mathbb{N}$, and let $w \in \mathcal{R}^{\mathbb{Z}}$. The minimum required observation time is*

$$\mu_{\mathcal{I}}(t, w) := \min(\{n \in \mathbb{N} : n \geq t \text{ and } (U_t(w_{0:n}) \subseteq \mathcal{I} \text{ or } U_t(w_{0:n}) \cap \mathcal{I} = \emptyset)\} \cup \{\infty\}).$$

For a fixed time t , future observations eliminate the unobserved future contribution, but they do not reveal the values before time 0. Therefore, an irreducible uncertainty of diameter $\frac{r^{t+1}}{1-r}d_{\mathcal{R}}$ may remain, and an exact verdict need not ever become possible. Even when exact verdicts are possible, their delay has no uniform finite bound in general.

► **Theorem 10** (No uniform exact verdict delay). *Let $(\mathcal{R}, r, s)$ be a proper input setting such that $\mathcal{R}$ is compact and $r + s > 0$, and let $J := \lambda^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. Let $\mathcal{I}$ be a target interval satisfying $J \cap \mathcal{I} \neq \emptyset$ and $J \setminus \mathcal{I} \neq \emptyset$. Then, for every $t \in \mathbb{N}$ and every $n \geq t$, there exists $w \in \mathcal{R}^{\mathbb{Z}}$ such that $\mu_{\mathcal{I}}(t, w) > n$.*

The proof shows that for any proposed finite delay, one can construct a prefix whose remaining uncertainty crosses a boundary of the target interval: one completion yields a discounted sum in the target, while another yields one outside it. A sound exact monitor must therefore remain inconclusive, which motivates approximate soundness.

2.2.2 Approximately Sound Monitors

► **Definition 11** (ε -approximately sound monitor). *Let $(\mathcal{R}, r, s)$ be an input setting, let $\mathcal{I}$ be a target interval, and let $\varepsilon > 0$. A monitor $\mathcal{M}$ is ε -approximately sound if, for every $w \in \mathcal{R}^{\mathbb{Z}}$, $n \in \mathbb{N}$, and $t \in \{0, \dots, n\}$,*

$$\mathcal{M}(w_{0:n}, t) = \perp \implies S_t^{r,s}(w) \notin \mathcal{I}_{-\varepsilon} \quad \text{and} \quad \mathcal{M}(w_{0:n}, t) = \top \implies S_t^{r,s}(w) \in \mathcal{I}_{+\varepsilon}. \quad (5)$$

where, for $\mathcal{I} = (L, U)$, $\mathcal{I}_{+\varepsilon} := (L - \varepsilon, U + \varepsilon)$ and $\mathcal{I}_{-\varepsilon} := (L + \varepsilon, U - \varepsilon)$.

Although $\mathcal{I}$, $\mathcal{I}_{+\varepsilon}$, and $\mathcal{I}_{-\varepsilon}$ are open, we note that the results of this section hold for arbitrary intervals. We next give a tight bound on the minimum required observation time for ε -approximately sound monitors, and show how to realise them with finite resources.

► **Theorem 12.** *Let $(\mathcal{R}, r, s)$ be an input setting, $\mathcal{I}$ a target interval, $\varepsilon > 0$, and $T \in \mathbb{N}$. With*

$$\tau^*(\mathcal{R}, r, s, \varepsilon, T) := \min \left(\left\{ \tau \in \mathbb{N} : d_{\mathcal{R}} \left(\frac{r^{T+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right) \leq 2\varepsilon \right\} \cup \{\infty\} \right). \quad (6)$$

1. *If $\tau^* < \infty$, there exists an ε -approximately sound monitor that produces a verdict for every $t \geq T$ within τ^* steps.*
2. *Suppose that $\mathcal{R} = [m, M]$ with $m < M$, and write $J = \lambda^{r,s}[m, M]$. If $J \cap \mathcal{I}_{-\varepsilon} \neq \emptyset$ and $J \setminus \mathcal{I}_{+\varepsilon} \neq \emptyset$, then, for every ε -approximately sound monitor and every $\tau < \tau^*$, there exists $w \in \mathcal{R}^{\mathbb{Z}}$ on which the monitor does not produce a verdict for time T after τ steps.*

For the usual case $d_{\mathcal{R}} > 0$ and $0 < s < 1$, let $B_T := \left(\frac{2\varepsilon}{d_{\mathcal{R}}} - \frac{r^{T+1}}{1-r} \right) (1-s)$. If $B_T > 0$, then $\tau^* = \max\{0, \lceil \log_s B_T \rceil - 1\}$. If $B_T \leq 0$, then $\tau^* = \infty$. The semantic definition in Eq. (6) also covers the exceptional cases: $d_{\mathcal{R}} = 0$ gives $\tau^* = 0$, while for $s = 0$ we have $\tau^* = 0$ if $d_{\mathcal{R}} \frac{r^{T+1}}{1-r} \leq 2\varepsilon$ and $\tau^* = \infty$ otherwise. Unless stated otherwise, we assume $0 < \tau^* < \infty$.

For every bounded input domain, Lemma 8 provides an interval enclosure whose diameter at time t , after τ further observations, is $d_{\mathcal{R}} \left(\frac{r^{t+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right)$. A diameter of at most 2ε forces an approximately sound verdict, which proves the upper bound. For interval-valued input domains, the finite observed contribution can be positioned continuously; under the two nontriviality conditions in point (2), this gives matching accepting and rejecting completions for every smaller delay.

2.3 Monitor Construction

Algorithm 1 constructs an ε -approximately sound monitor. With $\tau = \tau^*(\mathcal{R}, r, s, \varepsilon, T)$ from Eq. (6), it maintains:

- a global variable **Psum** storing the present-inclusive discounted sum of all past observations;
- an array of τ tuples (**RunningSum** _{i} , **pos** _{i}), one per active monitor;
- a Boolean array **in_use**[1: τ] marking which tuples are active.

Each active tuple tracks a candidate time **pos** _{i} whose discounted sum is being monitored. The invariant is that after observing x_n , every active tuple with $p = \text{pos}_j$ satisfies **RunningSum** _{j} = $S_p^{r,s}(w_{0:n}) = \sum_{i=1}^p r^i x_{p-i} + \sum_{i=0}^{n-p} s^i x_{p+i}$.

Algorithm 1 ε -Approximately Sound Discounted-Sum Monitor

Require: Input setting $(\mathcal{R}, r, s)$, tolerance ε , target interval $\mathcal{I}$, start time T

```

1:  $\tau \leftarrow \tau^*(\mathcal{R}, r, s, \varepsilon, T)$  ▷ Eq. (6)
2:  $\text{Psum} \leftarrow 0$ 
3: for  $i = 1$  to  $\tau$  do
4:    $\text{in\_use}[i] \leftarrow \perp$ 
5: end for
6: for  $t = 0, 1, 2, \dots$  do
7:   observe  $x_t$ ;  $\text{Psum} \leftarrow x_t + r \cdot \text{Psum}$  ▷ now  $\text{Psum} = x_t + \sum_{i=1}^t r^i x_{t-i}$ 
8:   for  $j = 1$  to  $\tau$  do ▷ Update all active registers
9:     if  $\text{in\_use}[j] = \top$  then
10:       $\text{RunningSum}_j \leftarrow \text{RunningSum}_j + x_t \cdot s^{t-\text{pos}_j}$  ▷ future contribution for time  $\text{pos}_j$ 
11:      compute  $C_{\text{pos}_j}^{r,s} = \text{RunningSum}_j + \left( \frac{s^{\text{pos}_j+1}}{1-r} + \frac{s^{t-\text{pos}_j+1}}{1-s} \right) \cdot [\inf \mathcal{R}, \sup \mathcal{R}]$ 
12:      if  $C_{\text{pos}_j}^{r,s} \subseteq \mathcal{I}_{+\varepsilon}$  then
13:        output verdict  $\top$  for time  $\text{pos}_j$ ;  $\text{in\_use}[j] \leftarrow \perp$ 
14:      else if  $C_{\text{pos}_j}^{r,s} \cap \mathcal{I}_{-\varepsilon} = \emptyset$  then
15:        output verdict  $\perp$  for time  $\text{pos}_j$ ;  $\text{in\_use}[j] \leftarrow \perp$ 
16:      end if
17:    end if
18:  end for
19:  if  $t \geq T$  then ▷ Instantiate new register for current  $t$ 
20:    select  $i$  with  $\text{in\_use}[i] = \perp$  ▷ Guaranteed to exist by Thm. 12
21:     $\text{RunningSum}_i \leftarrow \text{Psum}$ ; ▷ initially  $S_t^{r,s}(w_{0:t})$   $\text{pos}_i \leftarrow t$ ;  $\text{in\_use}[i] \leftarrow \top$ 
22:  end if
23: end for

```

Special cases. For future-only sums ($r = 0$), the past uncertainty vanishes. For past-only sums ($s = 0$), later observations cannot refine the uncertainty of an earlier position. Accordingly, the uniform bound in Eq. (6) is either 0 or ∞ : each position is tested once at creation time, and if its enclosure is still inconclusive, no later observation can change that.

General bounded input domains. Neither the upper bound in Theorem 12(1) nor Algorithm 1 requires properness. Both use only the interval enclosure from Lemma 8. The interval-domain assumption in Theorem 12(2) is used only to position the observed finite contribution continuously. For input domains with gaps, τ^* remains a valid upper bound but need not be tight for a fixed target interval.

3 Statistical Discounted Monitor

In statistical discounted monitoring, we assume that the observation sequence is generated by a stochastic process and the monitors objective is to estimate the value of the expected discounted sum. The statistical monitors are similar to the monitors of Section 2, but for an additional statistical error term which impacts the release condition (and time τ^*).

Setting. Assume that the value sequence w is a realisation of the stochastic process $W = (X_t)_{t \in \mathbb{Z}}$ in $\mathcal{R}$. Let $\mathcal{F} := (\mathcal{F}_t)_{t \in \mathbb{Z}}$ be its canonical filtration, i.e., $\mathcal{F}_t$ is the sigma-algebra generated by $(X_i)_{i \leq t}$. No further assumptions are placed on W .

Expected discounted sum. We are interested in the expected discounted sum, in infinite and finite form. We define the conditional expectation for the integrable random variable Y as $\mathbb{E}_t(Y) := \mathbb{E}(Y \mid \mathcal{F}_t)$. We define $\mathbb{S}_t^{r,s}(W) := \lim_{n_l \rightarrow -\infty} \lim_{n_u \rightarrow +\infty} \mathbb{S}_t^{r,s}(W_{n_l:n_u})$

$$\mathbb{S}_t^{r,s}(W_{n_l:n_u}) := \sum_{i=1}^{t-n_l} r^i \mathbb{E}_{t-i-1}(X_{t-i}) + \sum_{i=0}^{n_u-t} s^i \mathbb{E}_{t+i-1}(X_{t+i}) \text{ for } n_l \leq t \leq n_u, r, s \in [0, 1].$$

3.1 Monitoring

Statistical monitors compute the observed discounted sum and bound its deviation from the expected discounted sum. This permits three natural notions of statistical soundness, which differ in how many possible verdicts are protected by the same probability guarantee.

► **Definition 13** (Statistical soundness). For $(\mathcal{R}, r, s)$, let $\mathcal{I} = (L, U)$ be a target interval, $\varepsilon \geq 0$ a precision, and $\delta \in (0, 1)$ an error probability. For $t \leq n \in \mathbb{N}$ define the success event

$$E_t(n) := \{\mathcal{M}(W_{0:n}, t) = \perp \Rightarrow \mathbb{S}_t^{r,s}(W) \notin \mathcal{I}_{-\varepsilon} \wedge \mathcal{M}(W_{0:n}, t) = \top \Rightarrow \mathbb{S}_t^{r,s}(W) \in \mathcal{I}_{+\varepsilon}\},$$

where $\mathcal{I}_{+\varepsilon} := (L - \varepsilon, U + \varepsilon)$ and $\mathcal{I}_{-\varepsilon} := (L + \varepsilon, U - \varepsilon)$. A statistical monitor is

- (i) *pointwise* (ε, δ) -approximately sound if $\forall t \in \mathbb{N} \forall n \geq t: \mathbb{P}(E_t(n)) \geq 1 - \delta$;
- (ii) *locally* (ε, δ) -approximately sound if $\forall t \in \mathbb{N}: \mathbb{P}(\forall n \geq t: E_t(n)) \geq 1 - \delta$;
- (iii) *uniformly* (ε, δ) -approximately sound if $\mathbb{P}(\forall t \in \mathbb{N} \forall n \geq t: E_t(n)) \geq 1 - \delta$.

The three notions express increasingly stronger guarantees. *Pointwise soundness* protects one fixed verdict at a fixed time t and observation horizon n . Hence, flexible register release is not possible and an error probability of δ must be tolerated for the verdict at each t . *Local soundness* protects one fixed monitored time t for all observation horizons $n \geq t$. Hence, flexible register release is possible, but an error probability of δ remains for the verdict at each t . *Uniform soundness* protects the entire run, i.e., the invariant “every verdict issued for every time index and at every observation horizon is correct” holds with probability at least $1 - \delta$. This is the right notion when the correctness of verdicts is imperative.

3.2 Monitor Construction

We modify the deterministic monitor by adding a statistical error term to the uncertainty interval (which affects the register release bound τ^*). The resulting statistical uncertainty interval for the expected discounted sum for time t and observations n is

$$\mathcal{U}_t^{r,s}(\delta; W_{0:n}) := \mathcal{I}_t^{r,s}(\delta; W_{0:n}) + \gamma_{t,n}^{r,s} \cdot \Gamma \quad \text{where} \quad \mathcal{I}_t^{r,s}(\delta; W_{0:n}) := \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta). \quad (7)$$

where $\beta_{t,n}^{r,s}(\delta)$ is a statistical error term and $\gamma_{t,n}^{r,s}$ is a tail error as in Lemma 8, and $\Gamma := [\inf \mathcal{R}, \sup \mathcal{R}]$. To obtain soundness it suffices to ensure that $\mathcal{I}_t^{r,s}(\delta; W_{0:n})$ covers the finite expected sum $\mathbb{S}_t^{r,s}(W_{0:n})$ with the desired probability guarantee. Then the deterministic tail term $\gamma_{t,n}^{r,s} \cdot \Gamma$ lifts this to coverage of the infinite expected sum.

Error bounds. For simplicity we focus on data-independent deviation bounds, but the results can be extended to variance adaptive bounds [24]. Hence, the width of the statistical error term is dictated by the squared discount factors, i.e., for every $n \in \mathbb{N}$ and $t \leq n$ as

$$\omega_{t,n}^{r,s} := \sum_{i=1}^t r^{2i} + \sum_{i=0}^{n-t} s^{2i} = \frac{r^2(1 - r^{2t})}{1 - r^2} + \frac{1 - s^{2(n-t+1)}}{1 - s^2},$$

32:10 Monitoring Discounted Sum Properties

and a uniform upper bound σ on the conditional sub-Gaussian norm σ_t of $X_t - \mathbb{E}_{t-1}(X_t)$, which is trivially given by $d_{\mathcal{R}}/2$, i.e., $\sigma_t \leq \sigma \leq d_{\mathcal{R}}/2$ [38]. We define the pointwise, local, and uniform error bound, respectively, for $\delta \in (0, 1)$ as

$$\begin{aligned} \text{PE}_{t,n}^{r,s}(\delta) &:= \sqrt{2\sigma^2 \omega_{t,n}^{r,s} \log(2/\delta)} \quad \text{LE}_{t,n}^{r,s}(\delta) := k_1 \sqrt{V_{t,n}^{r,s} \left(2 \log(\log_2(V_{t,n}^{r,s}) + 1) + \log\left(\frac{2\pi^2}{6\delta}\right) \right)} \\ \text{UE}_{t,n}^{r,s}(\delta) &:= \text{LE}_{t,n}^{r,s} \left(\frac{6\delta}{\pi^2(t+1)^2} \right) \quad \text{where } V_{t,n}^{r,s} := \max(1, \sigma^2 \omega_{t,n}^{r,s}) \text{ and } k_1 := 2^{1/4} + 2^{-1/4}/\sqrt{2}. \end{aligned}$$

The pointwise sound monitor leverages the fixed-sample Hoeffding–Azuma deviation bound [3]. The locally sound monitor leverages the anytime-valid deviation bounds from Howard et al. [24]. The uniformly sound monitor leverages the locally sound statistical error bounds with error level $6\delta/(\pi^2(t+1)^2)$ for time t , because a union bound over all monitored times results in a uniform guarantee, i.e., $\sum_{t=0}^{\infty} 6\delta/(\pi^2(t+1)^2) = \delta$. Unsurprisingly, we can observe that as the guarantees become stronger, the interval width becomes wider.

► **Theorem 14.** *Let $(\mathcal{R}, r, s)$ be an input setting, $\mathcal{I} = (L, U)$ a target interval, $\varepsilon \geq 0$, and $\delta \in (0, 1)$. Consider Algorithm 1 where a register for time t is released at a fixed (n) or at a flexible $(\mathcal{U}_t^{r,s}(\delta; W_{0:n}) \subseteq \mathcal{I}_{+\varepsilon} \text{ or } \mathcal{U}_t^{r,s}(\delta; W_{0:n}) \cap \mathcal{I}_{-\varepsilon} = \emptyset)$ observation horizon. Then*

- (i) *if $\beta_{t,n}^{r,s}(\delta) = \text{PE}_{t,n}^{r,s}(\delta)$ the monitor is pointwise (ε, δ) -approx. sound for fixed release;*
- (ii) *if $\beta_{t,n}^{r,s}(\delta) = \text{LE}_{t,n}^{r,s}(\delta)$, the monitor is locally (ε, δ) -approx. sound for flexible release;*
- (iii) *if $\beta_{t,n}^{r,s}(\delta) = \text{UE}_{t,n}^{r,s}(\delta)$, the monitor is uniformly (ε, δ) -approx. sound for flexible release.*

Impossibility. If we consider the statistical uncertainty intervals as $n \rightarrow \infty$, we observe that the tail error $\gamma_{t,n}^{r,s}$ converges to 0 from above and the statistical error $\beta_{t,n}^{r,s}(\delta)$ converges to a non-zero constant from below. In particular, for the pointwise bound we have

$$\text{PE}_{t,n}^{r,s}(\delta) \rightarrow \sqrt{2\sigma^2 \left(\frac{r^2(1-r^{2t})}{1-r^2} + \frac{1}{1-s^2} \right) \log(2/\delta)} \quad \text{and} \quad d_{\mathcal{R}} \gamma_{t,n}^{r,s} \rightarrow d_{\mathcal{R}} \frac{r^{t+1}}{1-r} \quad \text{as } n \rightarrow \infty.$$

Without additional assumptions this error is unavoidable. Specifically, the pointwise statistical term (the smallest among the soundness notions) is minimax optimal. Intuitively, this is because the variance of the discounted average converges to a positive constant and to 0 only as the discount factors approach 1.

► **Theorem 15.** *Given an input setting $(\mathcal{R}, r, s)$ s.t. $\mathcal{R}$ contains an interval of length at least 2σ . Fix $n \in \mathbb{N}$ and $t \leq n$. Consider the class $\mathcal{D}_{\sigma}$ of product measures on $\mathcal{R}^{n+1}$ s.t. $X_i - \mathbb{E}_{i-1}(X_i)$ is conditionally σ -sub-Gaussian. We define: the set of all $(1 - \delta)$ confidence intervals $\text{CI}_{\sigma}(\delta)$ for $\mathbb{S}_t^{r,s}(W_{0:n})$, i.e., all $I : \mathcal{R}^{n+1} \rightarrow \text{Interval}(\mathbb{R})$ satisfying $\sup_{P \in \mathcal{D}_{\sigma}} \mathbb{P}_{W \sim P}(\mathbb{S}_t^{r,s}(W_{0:n}) \notin I(W_{0:n})) \leq \delta$, $|I| = \sup I - \inf I$ as interval length, and $\eta_{t,n}^{r,s} = \sum_{i=1}^t r^i + \sum_{i=0}^{n-t} s^i$. There exist universal constants $0 < c_0 < C_0 < \infty$ s.t. $\delta \in (0, 1/4)$*

$$c_0 \cdot \min(\sigma \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta)) \leq \inf_{I \in \text{CI}_{\sigma}(\delta)} \sup_{P \in \mathcal{D}_{\sigma}} \mathbb{E}_P[|I(W_{0:n})|] \leq C_0 \cdot \min(d_{\mathcal{R}} \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta)).$$

4 General Discounted Properties

Many quantitative properties involve arithmetic relations between multiple discounted aggregates. A canonical example is (group) fairness, where one compares *acceptance rates* across groups; such rates are quotients of discounted counts of accepted and total events. This section shows how to use the monitors of Section 2 as building blocks for monitoring multi-aggregate expressions.

Setting. Let $\mathcal{N} := [n]$, $n \in \mathbb{N}^+$, be the set of event types, and let $(\mathcal{R}, r, s)$ be an input setting. Let $\mathcal{R}_\square := \mathcal{R} \cup \{\square\}$ extend $\mathcal{R}$ with a distinguished empty symbol $\square$. An *event* is a vector $e \in \mathcal{E} := (\mathcal{R}_\square)^n$, where the k -th component $e^{(k)} \in \mathcal{R}_\square$ is the *atomic event of type k* . We consider bi-infinite event streams $w = (e_t)_{t \in \mathbb{Z}} \in \mathcal{E}^{\mathbb{Z}}$.

4.1 Specification Language

We introduce a simple expression language whose atoms are (bi-directional) discounted sums over individual event types and whose connectives are the standard arithmetic operations. We give two interpretations—*synchronous* and *asynchronous*—differing in how empty atomic events affect discounting.

Syntax. Formulas are built from discounted-sum atoms by scalar multiplication, addition, multiplication, and division:

$$\varphi ::= \mathcal{S}_{r,s}^{(k)} \mid c \cdot \varphi \mid \varphi + \varphi \mid \varphi \cdot \varphi \mid \varphi \div \varphi, \quad \text{where } k \in \mathcal{N}, c \in \mathbb{R}. \quad (8)$$

Semantics: atoms. Consider an event stream $w = (e_t)_{t \in \mathbb{Z}}$ and time $t \in \mathbb{Z}$. We map empty atomic events to zero via $x_t^{(k)} := \mathbb{1}[e_t^{(k)} \neq \square] \cdot e_t^{(k)}$ for $t \in \mathbb{Z}$, $k \in \mathcal{N}$. The value of $\mathcal{S}_{r,s}^{(k)}$ at time t is a bi-directional discounted sum over the k -typed atomic values, with discount exponents depending on the interpretation:

$$\llbracket \mathcal{S}_{r,s}^{(k)} \rrbracket_{sync/async}(w, t) := \underbrace{\sum_{i=1}^{\infty} r^{\tau_{t,-i}^{(k)}} x_{t-i}^{(k)}}_{\text{past}} + \underbrace{x_t^{(k)}}_{\text{present}} + \underbrace{\sum_{i=1}^{\infty} s^{\tau_{t,+i}^{(k)}} x_{t+i}^{(k)}}_{\text{future}}. \quad (9)$$

The discount exponents $\tau_{t,\pm i}^{(k)}$ are defined as $\tau_{t,\pm i}^{(k)} := i$ in the synchronous interpretation, and $\tau_{t,\pm i}^{(k)} := \sum_{j=1}^i \mathbb{1}[e_{t\pm j}^{(k)} \neq \square]$ in the asynchronous one. Thus, in the synchronous interpretation, every time step advances discounting, even if the k -typed atomic event is empty. In the asynchronous interpretation, only *non-empty* atomic events advance discounting for type k ; empty entries are ignored for the purpose of discount progression.

Note that, if $0 \in \mathcal{R}$, any stream $w \in \mathcal{E}^{\mathbb{Z}}$ can be converted to $w' \in \mathcal{E}^{\mathbb{Z}}$ such that $\llbracket \mathcal{S}_{r,s}^{(k)} \rrbracket_{sync}(w, t) = \llbracket \mathcal{S}_{r,s}^{(k)} \rrbracket_{async}(w', t)$.

Semantics: expressions. The semantics extends naturally to arithmetic operations: for a stream w , time t , and expressions ψ, χ ,

$$\llbracket c \cdot \psi \rrbracket(w, t) = c \cdot \llbracket \psi \rrbracket(w, t), \quad \llbracket \psi \circ \chi \rrbracket(w, t) = \llbracket \psi \rrbracket(w, t) \circ \llbracket \chi \rrbracket(w, t) \quad (\circ \in \{+, \cdot, \div\}).$$

Note that division is undefined when the denominator evaluates to 0.

4.2 Uncertainty Propagation Through Interval Arithmetic

Algorithm 1 naturally extends to arbitrary formulas by instantiating a register for each atomic discounted sum and propagating uncertainty intervals via standard interval arithmetic.

Theorem 12 bounds the observation time needed for an ε -approximately sound verdict on a single discounted sum. A natural question is whether an analogous bound holds for arbitrary formulas. The proof of Theorem 12 relies on the fact that, for fixed t , the uncertainty interval of a discounted sum has length $\gamma_{t,t+\tau}^{r,s}(\sup \mathcal{R} - \inf \mathcal{R})$ and depends only on the observation count τ , allowing to identify the earliest τ at which uncertainty falls below 2ε .

32:12 Monitoring Discounted Sum Properties

The first observation is that this type of reasoning can only be applied for the synchronous interpretation, as in the asynchronous interpretation, a stream can contain arbitrary many empty events for each atom, which make it impossible to give a minimum required observation time. For the synchronous interpretation, we have to develop a theory for different fragments of the whole language. In a nutshell, analogous results to Thm. 12 (1) exist for expressions that contain no divisions, by considering how the uncertainty intervals of sums and products grow through interval arithmetic. This avoids the undefined divisions by 0. When $\mathcal{R}$ is an interval and φ contains only sums and scalar products, the analogous result to Thm. 12 (2) also holds, as monitoring such an expression is equivalent to monitoring an atom on a larger input setting.

We start by defining set arithmetic as usual: for sets S_1, S_2 and operator $\circ \in \{+, -, \cdot, \div\}$, $S_1 \circ S_2 := \{s_1 \circ s_2 : s_1 \in S_1, s_2 \in S_2\}$. For a scalar $c \in \mathbb{R}$ and a set $S \subseteq \mathbb{R}$, we use $c \cdot S$ to denote $\{c\} \cdot S$. Notably, for $I_1 = [L_1, U_1]$, and $I_2 = [L_2, U_2]$, and $c \in \mathbb{R}$:

- $c \cdot I_1 = [cL_1, cU_1]$ if $c \geq 0$, $c \cdot I_1 = [cU_1, cL_1]$ if $c < 0$.
- $I_1 + I_2 = [L_1 + L_2, U_1 + U_2]$.
- $I_1 \cdot I_2 = [\min\{L_1 \cdot L_2, L_1 \cdot U_2, U_1 \cdot L_2, U_1 \cdot U_2\}, \max\{L_1 \cdot L_2, L_1 \cdot U_2, U_1 \cdot L_2, U_1 \cdot U_2\}]$.
- If $0 \notin I_2$, then $I_1/I_2 = [\min\{L_1/L_2, L_1/U_2, U_1/L_2, U_1/U_2\}, \max\{L_1/L_2, L_1/U_2, U_1/L_2, U_1/U_2\}]$.

► **Definition 16** (Linear and multiplicative properties). *Let φ be a general discounted property as defined in Eq. 8. We say that φ is linear if it contains only atoms, additions, and scalar multiplications. We say that φ is multiplicative if it contains only atoms, additions, scalar multiplication, and products.*

We define the spread of a formula as the range of values it can take at each timestep.

► **Definition 17** (Spread). *Let $(\mathcal{R}, r, s)$ be an input setting. The spread of a formula φ is defined recursively as:*

- If φ is an atom, $\text{Spread}_{\mathcal{R}}(\varphi) = \mathcal{R}$.
- If $\varphi = c \cdot \psi$ for $c \in \mathbb{R}$, $\text{Spread}_{\mathcal{R}}(\varphi) = c \cdot \text{Spread}_{\mathcal{R}}(\psi)$.
- If $\varphi = \psi \circ \chi$, then $\text{Spread}_{\mathcal{R}}(\varphi) = \text{Spread}_{\mathcal{R}}(\psi) \circ \text{Spread}_{\mathcal{R}}(\chi)$, for $\circ \in \{+, \cdot, \div\}$.
If, in any of the divisions, the denominator interval contains 0, the spread is undefined.

Note that the spread is always defined for multiplicative properties. For properties containing divisions, it may not be defined.

► **Remark 18.** Because of the interval arithmetic, the expression of $\text{Spread}(\varphi)$ is generally convoluted. A special case is when φ is linear. If φ is linear it can be written as

$$\varphi = \sum_{i=1}^{m^-} \alpha_i^- \cdot x_i^- + \sum_{i=1}^{m^+} \alpha_i^+ \cdot x_i^+.$$

Therefore, for $\mathcal{R} = [\inf \mathcal{R}, \sup \mathcal{R}]$, its spread can be written as

$$\text{Spread}_{\mathcal{R}}(\varphi) = \left[\sup \mathcal{R} \cdot \sum_{i=1}^{m^-} \alpha_i^- + \inf \mathcal{R} \cdot \sum_{i=1}^{m^+} \alpha_i^+, \inf \mathcal{R} \cdot \sum_{i=1}^{m^-} \alpha_i^- + \sup \mathcal{R} \cdot \sum_{i=1}^{m^+} \alpha_i^+ \right].$$

Using synchronous semantics, monitoring a linear expression φ on an input setting $(\mathcal{R}, r, s)$ is equivalent to monitoring an atom on the input setting $(\mathcal{R}_{\varphi}, r, s)$ where $\mathcal{R}_{\varphi} = \text{Spread}_{\mathcal{R} \cup \{0\}}(\varphi)$. Therefore, we have the following corollary to Theorem 12.

► **Corollary 19.** *Let $(\mathcal{R}, r, s)$ be an input setting, $\mathcal{I}$ a target interval, $\varepsilon > 0$, $T \in \mathbb{N}$, and φ a linear formula. Let $\mathcal{R}_{\varphi} = \text{Spread}_{\mathcal{R} \cup \{0\}}(\varphi)$. With*

$$\tau^*(\mathcal{R}, r, s, \varepsilon, T, \varphi) := \min \left(\left\{ \tau \in \mathbb{N} : d_{\mathcal{R}_{\varphi}} \left(\frac{r^{T+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right) \leq 2\varepsilon \right\} \cup \{\infty\} \right). \quad (10)$$

1. If $\tau^* < \infty$, there exists an ε -approximately sound monitor that produces a verdict for every $t \geq T$ within τ^* steps.
2. Suppose that $\mathcal{R}_\varphi = [m, M]$ with $m < M$, and write $J = \lambda^{r,s}[m, M]$. If $J \cap \mathcal{I}_{-\varepsilon} \neq \emptyset$ and $J \setminus \mathcal{I}_{+\varepsilon} \neq \emptyset$, then, for every ε -approximately sound monitor of φ and every $\tau < \tau^*$, there exists $w \in \mathcal{R}^\omega$ on which the monitor does not produce a verdict for time T after τ steps.

► **Remark 20.** Corollary 19 cannot be extended expressions with the asynchronous interpretation. If φ is interpreted asynchronously, then it can contain arbitrary many empty events for each atom, which make it impossible to give a minimum required observation time.

5 Register Complexity of Discounted-Sum Monitoring

The quantity τ plays two roles in Algorithm 1: it bounds both the *verdict delay*, i.e., the number of steps required to resolve the uncertainty interval for a position, and the *register count*, i.e., the number of running sums maintained in parallel. These roles have different origins. The delay is inherent to discounted sums: contributions from unobserved values decay geometrically, and τ is the first horizon where this is resolved for sure. The register count, by contrast, is an implementation cost. Since the running sums are related, it is natural to ask whether we can track the τ pending positions with less memory.

We show that no such reduction is possible in the worst case in the future-only setting. To make the question precise, we formalize monitors as *affine register machines* (ARMs): finite-state machines equipped with real-valued registers, affine updates, and strict affine guards [27, 2, 17]. By Theorem 12, every monitored position admits an ε -approximately sound verdict by its horizon, so we encode these deadline verdicts as a safety language over $\mathcal{R}^\omega$. We build, for every $\tau \geq 1$, a monitoring instance with horizon τ whose language cannot be recognized by any ARM with fewer than τ registers.

Safety Formulation of Approximate Monitoring. We first recast approximate monitoring as a safety-language recognition problem. Throughout this section we work over the normalized input domain $\mathcal{R} = [0, 1]$. This is without loss of generality for interval domains: any bounded interval $[m, M]$ can be mapped to $[0, 1]$ by an affine transformation, and discounted sums, target intervals, and tolerances rescale accordingly.

A *monitoring instance* is a tuple $\mathcal{P} = (\mathcal{R}, r, s, \mathcal{I}, \varepsilon)$, where $(\mathcal{R}, r, s)$ is a proper input setting with $\mathcal{R} = [0, 1]$ and $r, s \in [0, 1)$, where $\mathcal{I} = (L, U)$ is a target interval and $\varepsilon > 0$. The *horizon* of $\mathcal{P}$ is $\tau(\mathcal{P}) := \tau^*(\mathcal{R}, r, s, \varepsilon, 0)$, as in Theorem 12. We say that $\mathcal{P}$ is *future only* if $r = 0$ and $s > 0$, and *past only* if $r > 0$ and $s = 0$.

► **Definition 21 (Monitoring Language).** The language of a monitoring instance $\mathcal{P} = (\mathcal{R}, r, s, \mathcal{I}, \varepsilon)$ is $L(\mathcal{P}) := \{x \in \mathcal{R}^\omega \mid \forall t \geq 0 : U_t^{r,s}(x_{0:t+\tau(\mathcal{P})}) \subseteq \mathcal{I}_{+\varepsilon}\}$.

By Lemma 8 and Theorem 12, the uncertainty interval at time $t + \tau(\mathcal{P})$ has diameter at most 2ε . Hence, if it is not contained in $\mathcal{I}_{+\varepsilon}$, then it is disjoint from $\mathcal{I}_{-\varepsilon}$. Therefore, the rule that returns $\top$ exactly when $U_t^{r,s}(x_{0:t+\tau(\mathcal{P})}) \subseteq \mathcal{I}_{+\varepsilon}$ and returns $\perp$ otherwise is always defined and ε -approximately sound.

Since non-membership in $L(\mathcal{P})$ is witnessed by a finite prefix, $L(\mathcal{P})$ is a safety language. Indeed, $U_t^{r,s}(x_{0:t+\tau(\mathcal{P})})$ depends only on $x_{0:t+\tau(\mathcal{P})}$, so any violation at time t is shared by every extension of this prefix. By Lemma 8 and since $\mathcal{R} = [0, 1]$, the condition $U_t^{r,s}(x_{0:t+\tau(\mathcal{P})}) \subseteq \mathcal{I}_{+\varepsilon}$ is equivalent to $L - \varepsilon < S_t^{r,s}(x_{0:t+\tau(\mathcal{P})}) < U + \varepsilon - \gamma_{t,t+\tau(\mathcal{P})}^{r,s}$.

Affine Register Machines. Let $Y = \{y_1, \dots, y_k\}$ be a finite set of registers. A *valuation* is a map $\nu : Y \rightarrow \mathbb{R}$. An *update* over Y is a parallel assignment $y_i \leftarrow \sum_{j=1}^k a_{i,j} y_j + b_i \varsigma + c_i$ for all $1 \leq i \leq k$, where $\varsigma \in \mathcal{R}$ is the current input and all coefficients are real. We write $\Gamma(Y)$ for the set of such updates. A *guard* over Y is a finite conjunction of strict affine inequalities $\sum_{j=1}^k a_{ij} y_j + b_i \varsigma \bowtie c_i$, where $\bowtie \in \{<, >\}$. We write $(\nu, \varsigma) \models \varphi$ if the guard φ holds under valuation ν and input ς , and we write $\Phi(Y)$ for the set of guards over Y .

► **Definition 22** (Affine Register Machine). A (*deterministic*) affine register machine (ARM) is a tuple $\mathcal{M} = (\mathcal{R}, Y, Q, q_0, \nu_0, \Delta)$ where $\mathcal{R} \subset \mathbb{R}$ is a bounded input set, Y is a finite set of registers, Q is a finite set of control locations, $q_0 \in Q$ is the initial location, $\nu_0 : Y \rightarrow \mathbb{R}$ is the initial valuation, and $\Delta \subseteq Q \times \Phi(Y) \times \Gamma(Y) \times Q$ is a finite transition relation. We write transitions as $(q, \varphi, \gamma, q') \in \Delta$ and require the following determinism condition: for every location $q \in Q$, every valuation ν , and every input $\varsigma \in \mathcal{R}$, there is at most one transition $(q, \varphi, \gamma, q') \in \Delta$ such that $(\nu, \varsigma) \models \varphi$.

A configuration is a pair (q, ν) . On input ς , the machine follows the unique enabled transition (q, φ, γ, q') to $(q', \gamma(\nu, \varsigma))$. If no transition is enabled, the run terminates. A run on $x \in \mathcal{R}^\omega$ is *accepting* if infinite. The language $L(\mathcal{M})$ is the set of streams with an accepting run; hence every ARM recognizes a safety language.

Sufficiency: An ARM Construction. We show that $L(\mathcal{P})$ can be recognized by an ARM with $\tau(\mathcal{P}) + 1$ registers. The construction follows the monitor from Algorithm 1. The ARM keeps one register for each pending position and updates all pending truncated sums in parallel. If a register y_a stores the truncated sum for the pending position of age a , then on input x the affine updates are $y_0 \leftarrow r y_0 + x$ and $y_a \leftarrow y_{a-1} + s^a x$ for all $1 \leq a \leq \tau(\mathcal{P}) - 1$. When a position reaches age $\tau(\mathcal{P})$, its completed truncated sum is $y_{\tau(\mathcal{P})-1} + s^{\tau(\mathcal{P})} x$, which the machine checks by a guard before applying these updates. One additional register tracks the geometric factor needed to express the time-dependent past residual $r^{t+1}/(1-r)$.

► **Theorem 23** (General Upper Bound). For every monitoring instance $\mathcal{P}$, there is an ARM $\mathcal{M}$ with $\tau(\mathcal{P}) + 1$ registers such that $L(\mathcal{M}) = L(\mathcal{P})$.

In the future-only case the past residual is zero. The auxiliary register is therefore unnecessary, and the same construction uses exactly one register per pending future window.

► **Corollary 24** (Future-Only Upper Bound). For every future-only monitoring instance $\mathcal{P}$, there is an ARM $\mathcal{M}$ with $\tau(\mathcal{P})$ registers such that $L(\mathcal{M}) = L(\mathcal{P})$.

Necessity: A Register Lower Bound. We now show that the future-only upper bound is optimal in the worst case. For each $k \geq 1$, consider the future-only monitoring instance $\mathcal{P}_k := ([0, 1], 0, \frac{1}{2}, (0, 1), 2^{-(k+1)})$. For $x \in [0, 1]^\omega$ and $t \geq 0$, write $V_t^k(x) := \sum_{i=0}^k 2^{-i} x_{t+i}$ and $T_k := 1 - 2^{-(k+1)}$. For this family, the safety condition has a particularly simple form. The horizon is k , the residual future uncertainty is 2^{-k} , and the lower bound is vacuous because all inputs are nonnegative. Thus, the monitoring language is exactly the set of streams whose discounted windows of length $k+1$ remain below the threshold T_k .

► **Lemma 25** (Characterization of the hard instances). For every $k \geq 1$, the future-only monitoring instance $\mathcal{P}_k$ has horizon $\tau(\mathcal{P}_k) = k$, and its monitoring language is $L(\mathcal{P}_k) = \{x \in [0, 1]^\omega \mid \forall t \geq 0 : \sum_{i=0}^k 2^{-i} x_{t+i} < 1 - 2^{-(k+1)}\}$.

The lower-bound proof is a finite-dimensional linear-algebra argument. Suppose an ARM with only $k-1$ registers recognizes the language above. Consider the first k transitions taken

on a suitable accepted anchor word. On all length- k prefixes inducing this same transition sequence, the register valuation after k steps is an affine map from a k -dimensional prefix space to $\mathbb{R}^{k-1}$. Therefore, this map has a nonzero kernel direction: a perturbation that is invisible to the registers after the first k inputs.

The separating prefixes are chosen from an explicit one-parameter family. For θ just below T_k , the word $u(\theta) \cdot \alpha \cdot 0^\omega$ is constructed so that its first k discounted windows all have value θ . Moving slightly in the invisible kernel direction preserves the first k transitions, the register valuation after k steps, and the entire suffix. However, by choosing the last nonzero coordinate of the kernel direction, one discounted window is pushed above T_k . Determinism then forces the ARM to treat two identical configurations with the same suffix identically, although one stream satisfies the monitoring condition and the other violates it. Note that $\mathcal{R}$ being an interval is crucial here, as it ensures that the kernel direction (an arbitrary vector in $\mathbb{R}^k$) can be realized as a small perturbation inside $\mathcal{R}^k$.

► **Theorem 26** (Future-Only Lower Bound). *For every $k \geq 1$, there is a future-only monitoring instance $\mathcal{P}_k$ such that for every ARM $\mathcal{M}$ with $\tau(\mathcal{P}_k) - 1$ registers we have $L(\mathcal{M}) \neq L(\mathcal{P}_k)$.*

Special Cases. For the past-only case ($s = 0$ and $r > 0$), it suffices to use just two registers: one for the observed past, which follows the rule $p_t = x_t + rp_{t-1}$, and one for the uncertainty term $r^{t+1}/(1-r)$. The register complexity for a general $s > 0$ and $r > 0$ remains open.


Over finite alphabets, the future-only lower bound disappears. For $r = 0$, monitoring reduces to checking a sliding window $x_t, \dots, x_{t+\tau}$, so a finite-state monitor suffices, at the cost of exponential state space in τ . With past discounting, however, this is no longer possible: although a finite alphabet collapses the τ parallel future registers into finite control, the recurrence $p_t = x_t + rp_{t-1}$ still can take infinitely many values, which may require an infinite-state monitor.

6 Experiments


We evaluate our monitors on a collection of real-world and synthetic scenarios with the objective of demonstrating the practical effectiveness of the proposed monitors and to examine how tight the theoretical bounds on resource usage and accuracy are in practice. In the main paper our evaluation is guided by the research questions: **(RQ1)** How closely does the observed register usage of Algorithm 1 match the theoretical bound of Thm. 12?, and **(RQ2)** How does register usage vary under synchronous and asynchronous semantics? We address research questions beyond (RQ1) and (RQ2) in the full version.

We evaluate our monitors on traces from two benchmarks: **PowerData**, a dataset of power usage traces from Google data centers [33]; and **Adult**, demographic parity traces produced by a fairness-aware neural network trained on the Adult dataset [7].

Register Usage on PowerData. To address RQ1, we study register usage on traces from **PowerData**. We report *normalized register usage*, defined as the number of active registers divided by the theoretical upper bound τ^* from Thm. 12. The input range is $\mathcal{R} = [0, 1]$ and we fix the discount factors to $r = s = 0.9$. In Fig. 1 (a), we vary the tolerance ε logarithmically between 0.5 and $5 \cdot 10^{-4}$, while fixing the target interval to $\mathcal{I} = [0.703, 0.728]$, centered at the empirical mean with width equal to one standard deviation of the dataset. In Fig. 1 (b), we fix $\varepsilon = 5 \cdot 10^{-4}$ and vary the width of $\mathcal{I}$ between 10 and 0.01 times the standard deviation, again centered at the mean. In all experiments, the initial monitoring time T is chosen such that $r^{T+1}/(1-r)$ is closest to $\varepsilon/(\sup \mathcal{R} - \inf \mathcal{R})$. In both Fig. 1 (a)



■ **Figure 1** Register usage in terms of normalized active registers, **PowerData** dataset.



■ **Figure 2** Active registers when demographic parity, variable ϵ . **Adult** dataset.

and (b), the left graph represents the cumulative distribution of normalized active monitors, with one line representing each parameter setting – ϵ in (a), $|I|$ in (b) –, while the right graph represents each distribution as a boxplot. Each boxplot corresponds to a parameter setting, color matched with the corresponding left graph.

Across all configurations, most verdicts are produced with fewer than 60% of the theoretically available registers active, and register usage never exceeds 80%. This result indicates that we could actually implement a monitor with, for example, 70% of the theoretically required registers, and still output a sound verdict most of the time.

Monitoring Demographic Parity. To investigate RQ2, we monitor demographic parity on the **Adult** dataset. The input set is $\mathcal{R} = \{0, 1\}$ and we again fix $r = s = 0.95$. We consider decreasing tolerance values with a fixed target interval $\mathcal{I} = [-0.1, 0.1]$. We report the results in Fig. 2, in the same format as the previous figure, for the demographic parity property (Ex. 3) with the synchronous (blue) and asynchronous (green) interpretation. In contrast to the previous experiments, we report absolute register counts rather than normalized usage. This is because Thm. 12 does not provide a uniform bound in this setting, as the denominators can be arbitrarily close to zero. As expected, synchronous interpretations are significantly more efficient than asynchronous ones, since discounting is applied uniformly across streams rather than independently.

7 Related Work

Discounted-sum objectives are well studied in verification, automata, games, and logic as quantitative specifications where recent events outweigh distant ones [10, 11, 8, 36, 18, 1]. The nontrivial structure of achievable values underlies the difficulty of the target discounted-sum problem [9], motivating approximation: nondeterministic discounted-sum automata cannot be determinized exactly, but approximate determinization is always possible [8] and approximate inclusion is decidable [4]. We study online monitoring under bounded memory

and show a parallel dichotomy: exact sound monitoring requires unbounded memory, whereas approximate monitoring admits tight bounds.

Runtime verification provides lightweight formal guarantees by checking specifications against executions online [5]. Quantitative monitoring replaces boolean verdicts with real-valued measures; STL robustness semantics, for instance, quantify how strongly a signal satisfies a specification [30, 16, 14, 13, 12]. Recent work [23] establishes precision-cost tradeoffs for quantitative monitoring. We focus on discounted-sum monitoring and extend the framework to stochastic processes via confidence sequences.

Discounted sums can be viewed as convolution with an exponential kernel: the exponentially weighted moving average (EWMA) computes a normalized discounted sum [32], equivalent to a first-order IIR filter [35]. EWMA control charts detect mean shifts under i.i.d. inputs [29, 31], and exponential smoothing in forecasting balances responsiveness and noise reduction [19, 25]. We complement these by studying memory requirements and approximation guarantees for arbitrary input streams.

Standard fairness metrics such as demographic parity and equalized odds [21, 15] are static population statistics. Existing fairness monitors use only cumulative statistics [22] or stream-based specifications [6]. Recent work highlights temporal aspects of fairness [20]: discounting appears in fair reinforcement learning [26], decision making [37], and resource allocation [28], but targets policy design rather than monitoring with formal guarantees.

References

- 1 Shaull Almagor, Udi Boker, and Orna Kupferman. Discounting in ltl. In *International conference on tools and algorithms for the construction and analysis of systems*, pages 424–439. Springer, 2014.
- 2 Rajeev Alur, Loris D’Antoni, Jyotirmoy V. Deshmukh, Mukund Raghothaman, and Yifei Yuan. Regular functions and cost register automata. In *28th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2013, New Orleans, LA, USA, June 25–28, 2013*, pages 13–22. IEEE Computer Society, 2013. doi:10.1109/LICS.2013.65.
- 3 Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series*, 19(3):357–367, 1967.
- 4 Suguman Bansal, Swarat Chaudhuri, and Moshe Y Vardi. Comparator automata in quantitative verification. *Logical Methods in Computer Science*, 18, 2022.
- 5 Ezio Bartocci, Yliès Falcone, Adrian Francalanza, and Giles Reger. Introduction to runtime verification. In *Lectures on Runtime Verification: Introductory and Advanced Topics*, pages 1–33. Springer, 2018.
- 6 Jan Baumeister, Bernd Finkbeiner, Frederik Scheerer, Julian Siber, and Tobias Wagenpfeil. Stream-based monitoring of algorithmic fairness. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 60–81. Springer, 2025.
- 7 Barry Becker and Ronny Kohavi. Adult. UCI Machine Learning Repository, 1996. DOI: <https://doi.org/10.24432/C5XW20>.
- 8 Udi Boker and Thomas A Henzinger. Approximate determinization of quantitative automata. *Leibniz International Proceedings in Informatics*, 18, 2012.
- 9 Udi Boker, Thomas A Henzinger, and Jan Otop. The target discounted-sum problem. In *2015 30th Annual ACM/IEEE Symposium on Logic in Computer Science*, pages 750–761. IEEE, 2015.
- 10 Krishnendu Chatterjee, Laurent Doyen, and Thomas A Henzinger. Quantitative languages. *ACM Transactions on Computational Logic (TOCL)*, 11(4):1–38, 2010.
- 11 Luca De Alfaro, Thomas A Henzinger, and Rupak Majumdar. Discounting the future in systems theory. In *International Colloquium on Automata, Languages, and Programming*, pages 1022–1037. Springer, 2003.

- 12 Jyotirmoy V Deshmukh, Alexandre Donzé, Shromona Ghosh, Xiaoqing Jin, Garvit Juniwal, and Sanjit A Seshia. Robust online monitoring of signal temporal logic. *Formal Methods in System Design*, 51(1):5–30, 2017.
- 13 Alexandre Donzé, Thomas Ferrere, and Oded Maler. Efficient robust monitoring for stl. In *International conference on computer aided verification*, pages 264–279. Springer, 2013.
- 14 Alexandre Donzé and Oded Maler. Robust satisfaction of temporal logic over real-valued signals. In *International conference on formal modeling and analysis of timed systems*, pages 92–106. Springer, 2010.
- 15 Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*, pages 214–226, 2012.
- 16 Georgios E Fainekos and George J Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science*, 410(42):4262–4291, 2009.
- 17 Thomas Ferrère, Thomas A. Henzinger, and N. Ege Saraç. A theory of register monitors. In Anuj Dawar and Erich Grädel, editors, *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*, pages 394–403. ACM, 2018. doi:10.1145/3209108.3209194.
- 18 Jerzy Filar and Koos Vrieze. *Competitive Markov decision processes*. Springer Science & Business Media, 2012.
- 19 Everette S Gardner Jr. Exponential smoothing: The state of the art—part ii. *International journal of forecasting*, 22(4):637–666, 2006.
- 20 Usman Gohar, Zeyu Tang, Jialu Wang, Kun Zhang, Peter Spirtes, Yang Liu, and Lu Cheng. Long-term fairness inquiries and pursuits in machine learning: A survey of notions, methods, and challenges. *Trans. Mach. Learn. Res.*, 2025.
- 21 Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
- 22 Thomas A Henzinger, Mahyar Karimi, Konstantin Kueffner, and Kaushik Mallik. Monitoring algorithmic fairness. In *International Conference on Computer Aided Verification*, pages 358–382. Springer, 2023.
- 23 Thomas A Henzinger and N Ege Saraç. Quantitative and approximate monitoring. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–14. IEEE, 2021.
- 24 Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 49(2):1055–1080, 2021.
- 25 Rob J Hyndman and Yeasmin Khandakar. Automatic time series forecasting: the forecast package for r. *Journal of statistical software*, 27:1–22, 2008.
- 26 Shahin Jabbari, Matthew Joseph, Michael Kearns, Jamie Morgenstern, and Aaron Roth. Fairness in reinforcement learning. In *International conference on machine learning*, pages 1617–1626. PMLR, 2017.
- 27 Michael Kaminski and Nissim Francez. Finite-memory automata. *Theor. Comput. Sci.*, 134(2):329–363, 1994. doi:10.1016/0304-3975(94)90242-9.
- 28 Ashwin Kumar and William Yeoh. Remember, but also, forget: Bridging myopic and perfect recall fairness with past-discounting. *arXiv preprint arXiv:2504.01154*, 2025.
- 29 James M Lucas and Michael S Saccucci. Exponentially weighted moving average control schemes: properties and enhancements. *Technometrics*, 32(1):1–12, 1990.
- 30 Oded Maler and Dejan Nickovic. Monitoring temporal properties of continuous signals. In *International symposium on formal techniques in real-time and fault-tolerant systems*, pages 152–166. Springer, 2004.
- 31 Douglas C Montgomery. *Introduction to statistical quality control*. John Wiley & sons, 2020.
- 32 Stuart W Roberts. Control chart tests based on geometric moving averages. *Technometrics*, 42(1):97–101, 2000.

- 33 Varun Sakalkar, Vasileios Kontorinis, David Landhuis, Shaohong Li, Darren De Ronde, Thomas Blooming, Anand Ramesh, James Kennedy, Christopher Malone, Jimmy Clidas, and Parthasarathy Ranganathan. Data center power oversubscription with a medium voltage power plane and priority-aware capping. In *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, page 497–511, 2020. Data available: <https://github.com/google/cluster-data/blob/master/PowerData2019.md>.
- 34 César Sánchez, Gerardo Schneider, Wolfgang Ahrendt, Ezio Bartocci, Domenico Bianculli, Christian Colombo, Yliès Falcone, Adrian Francalanza, Srđan Krstić, João M Lourenço, et al. A survey of challenges for runtime verification from advanced application domains (beyond software). *Formal Methods in System Design*, 54(3):279–335, 2019.
- 35 Ronald William Schafer and Alan V Oppenheim. *Discrete-time signal processing*, volume 5. Prentice Hall Englewood Cliffs, New Jersey, 1989.
- 36 Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- 37 Manuel R. Torres, Parisa Zehtabi, Michael Cashmore, Daniele Magazzeni, and Manuela Veloso. Temporal fairness in decision making problems. In *European Conference of Artificial Intelligence (ECAI)*, Frontiers in Artificial Intelligence and Applications, pages 1132–1139. IOS Press, 2024.
- 38 Roman Vershynin. *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press, 2018.

A Discussion: Alternative Local Measures

In this paper we study discounted sums as a method to obtain a local measure for global property. Discounted sums are a principled method of aggregation, where instances in the future and in the past contribute to the local measure according to their value and the distance from one event to the other.

In this section, we discuss alternative localization measures, and how do they compare to discounted sums. We mainly discuss window averages and kernel-weighted averages.

A.1 Window average

A natural localization measure, alternative to the discounted average, is the rolling window average of length $2l + 1$, defined as

$$\text{WA}_t^l(w) = \frac{1}{2l + 1} \sum_{i=-l}^l x_{t+i}.$$

The parameter of window length serves a similar purpose as the pair of discount factors (r, s) . While there is no canonical correspondance, a natural one is to consider symmetric discounting ($r = s$) and that the weight of the central element in both averages is the same. This yields

$$\frac{1}{2l + 1} = \frac{1}{1 + \frac{2r}{1-r}} \iff l = \frac{r}{1-r} \iff r = \frac{l}{1+l}.$$

Under these conditions, we have a similar bounded difference as in Lemma 4, that is $|\text{WA}_{t+1}^l(w) - \text{WA}_t^l(w)| = d_{\mathcal{R}}$. We can also compute uncertainty sets in a similar fashion as Lemma 8, although the uncertainty sets may not be intervals (i.e., they may have gaps) if $\mathcal{R}$ has gaps of any size, due to the window averages expressing a finite sum. Sound monitors that give a verdict in at most l observations trivially exist, and they can be trivially made to use only $2l + 1$ registers.

32:20 Monitoring Discounted Sum Properties

In essence, window averages are computationally easier to manage, but have two main undesired effects:

- There is a sharp jump between the importance of a single observation for one step and its consecutive. In contrast, in discounted averages, the importance associated with one step falls gradually as it becomes further away from the point of interest. For many applications, like monitoring fairness, this is a more principled approach than setting a hard cutoff distance in time.
- For properties expressed as quotients, even if $0 \in \mathcal{R}$, as long as it does not contain both positive and negative elements at the same time, the discounted sum of the denominator will be well defined from the first instance of a non-zero value. In contrast, when localizing with a rolling window average, as soon as we have a sequence of $2l+1$ instances of zero, the value of the property is not well defined. This is particularly important for synchronous properties, as non-appearances are in effect substituted by zero.

A.2 Kernel-Weighted Average

Both rolling window and discounted averages are a particular case of kernel weighted averages. A kernel is a weight function $\kappa: \mathbb{Z} \rightarrow \mathbb{R}_{\geq 0}$ that has a finite sum, i.e., $\sum_{i \in \mathbb{Z}} \kappa(i) < \infty$. To have the effect of a local measure, we ask for the kernel to be non-decreasing in $\mathbb{Z}_{\leq 0}$ and non-increasing for $\mathbb{Z}_{\geq 0}$. Kernels generalize both discounted and rolling window averages:

- The window average corresponds to the kernel $\kappa(i) = \frac{1}{2l+1} \mathbb{1}[-l \leq i \leq l]$.
- The discounted sum corresponds to the kernel $\kappa(i) = r^{\max\{0, -i\}} \cdot s^{\max\{0, i\}}$.

While kernel-weighted averages are conceptually appealing, it is not clear how we could efficiently monitor a property defined by an arbitrary kernel. For discounted sums, we take advantage of the fact that the value can be computed recursively, so we can efficiently monitor into the future while forgetting past values (see the update, for example, in line 6 of Alg. 1). For many popular kernels, like Gaussian kernels, this is not possible.

B Monitoring Discounted Sums

Lemma 4. *Let $(\mathcal{R}, r, s)$ be an input setting and $w \in \mathcal{R}^{\mathbb{Z}}$ be a sequence of values, and $t \in \mathbb{Z}$ a point in time. Then S_t exists and $|S_t^{r,s}(w) - S_{t-1}^{r,s}(w)| \leq d_{\mathcal{R}}$.*

Proof. Since S_t is defined as an infinite sum, we first check convergence. Let $a = \inf \mathcal{R}$ and $b = \sup \mathcal{R}$. Then $|x_{t \pm i}| \leq \max\{|a|, |b|\}$ for all $i \geq 0$, and hence the past and future parts are dominated by geometric series with ratios r and s , respectively. Since $r, s \in [0, 1)$, both series converge absolutely, so S_t exists.

Next we show the Lipschitz property. Write

$$S_t = \sum_{i=1}^{\infty} r^i x_{t-i} + x_t + \sum_{i=1}^{\infty} s^i x_{t+i},$$

$$S_{t-1} = \sum_{i=1}^{\infty} r^i x_{t-1-i} + x_{t-1} + \sum_{i=1}^{\infty} s^i x_{t-1+i}.$$

By re-indexing,

$$\sum_{i=1}^{\infty} r^i x_{t-1-i} + x_{t-1} = \sum_{j=0}^{\infty} r^j x_{t-1-j}, \quad x_t + \sum_{i=1}^{\infty} s^i x_{t+i} = \sum_{j=0}^{\infty} s^j x_{t+j},$$

and similarly

$$\sum_{i=1}^{\infty} r^i x_{t-i} = r \sum_{j=0}^{\infty} r^j x_{t-1-j}, \quad \sum_{i=1}^{\infty} s^i x_{t-1+i} = s \sum_{j=0}^{\infty} s^j x_{t+j}. \quad (11)$$

Therefore,

$$\begin{aligned} S_{t-1} - S_t &= \left(\sum_{j=0}^{\infty} r^j x_{t-1-j} - r \sum_{j=0}^{\infty} r^j x_{t-1-j} \right) + \left(s \sum_{j=0}^{\infty} s^j x_{t+j} - \sum_{j=0}^{\infty} s^j x_{t+j} \right) \\ &= (1-r) \sum_{j=0}^{\infty} r^j x_{t-1-j} - (1-s) \sum_{j=0}^{\infty} s^j x_{t+j}. \end{aligned}$$

Since $a \leq x_s \leq b$ for all s , we have

$$\frac{a}{1-r} \leq \sum_{j=0}^{\infty} r^j x_{t-1-j} \leq \frac{b}{1-r}, \quad \frac{a}{1-s} \leq \sum_{j=0}^{\infty} s^j x_{t+j} \leq \frac{b}{1-s}.$$

Multiplying by $(1-r)$ and $(1-s)$, respectively, yields

$$a \leq (1-r) \sum_{j=0}^{\infty} r^j x_{t-1-j} \leq b, \quad a \leq (1-s) \sum_{j=0}^{\infty} s^j x_{t+j} \leq b.$$

Hence

$$a - b \leq S_{t-1} - S_t \leq b - a, \quad \text{which implies } |S_{t-1} - S_t| \leq b - a = d_{\mathcal{R}}.$$

◀

Lemma 6. *Let $(\mathcal{R}, r, s)$ be an input setting. Then the following hold.*

- (1) $\text{Sums}_{\mathcal{R}}^{r,s} \subseteq [\lambda m, \lambda M]$.
- (2) If $(\mathcal{R}, r, s)$ is proper, then $\text{Closure}(\text{Sums}_{\mathcal{R}}^{r,s}) = [\lambda m, \lambda M]$.
- (3) If $(\mathcal{R}, r, s)$ is proper and compact, then $\text{Sums}_{\mathcal{R}}^{r,s} = [\lambda m, \lambda M]$. Where $m = \inf \mathcal{R}$ and $M = \sup \mathcal{R}$.

Proof. (1) Since $m \leq x_i \leq M$ for all $i \in \mathbb{Z}$ and all coefficients r^i, s^i are nonnegative, we obtain

$$S_{r,s}(x) = x_0 + \sum_{i=1}^{\infty} x_{-i} r^i + \sum_{i=1}^{\infty} x_i s^i \leq M \left(1 + \sum_{i=1}^{\infty} r^i + \sum_{i=1}^{\infty} s^i \right).$$

Similarly,

$$S_{r,s}(x) \geq m \left(1 + \sum_{i=1}^{\infty} r^i + \sum_{i=1}^{\infty} s^i \right).$$

Using the geometric series identities $\sum_{i=1}^{\infty} r^i = \frac{r}{1-r}$ and $\sum_{i=1}^{\infty} s^i = \frac{s}{1-s}$ yields

$$\lambda m \leq S_{r,s}(x) \leq \lambda M.$$

Since x was arbitrary, this proves $\text{Sums}_{\mathcal{R}}^{r,s} \subseteq [\lambda m, \lambda M]$.

32:22 Monitoring Discounted Sum Properties

(3) First, define the one-sided discounted sets

$$A_r(\mathcal{R}) := \left\{ \sum_{i=1}^{\infty} x_i r^i : x_i \in \mathcal{R} \right\}, \quad A_s(\mathcal{R}) := \left\{ \sum_{i=1}^{\infty} x_i s^i : x_i \in \mathcal{R} \right\}.$$

Then

$$\text{Sums}_{r,s}^{\mathcal{R}} = \mathcal{R} + A_r(\mathcal{R}) + A_s(\mathcal{R})$$

as a Minkowski sum. We split the proof into three steps.

Step 1: each one-sided set is an interval. Fix $r \in [0, 1)$ and consider the space $\mathcal{K}$ of nonempty compact subsets of $\mathbb{R}$ equipped with the Hausdorff metric d_H , defined as usual as

$$d_H(X, Y) = \max \left\{ \sup_{x \in X} \inf_{y \in Y} |x - y|, \sup_{y \in Y} \inf_{x \in X} |x - y| \right\}.$$

Define the operator

$$F_r : \mathcal{K} \rightarrow 2^{\mathbb{R}}, \quad F_r(K) := \bigcup_{a \in \mathcal{R}} (ra + rK).$$

1. F_r maps $\mathcal{K}$ to $\mathcal{K}$. $F_r(K) = r \cdot (\mathcal{R} + K)$ (as a Minkowski sum). Since the sum is a continuous function in the product topology, the image of a compact set is a compact set.
2. F_r is a contraction. For $K, L \in \mathcal{K}$, let $x = r(a + k) \in F_r(K)$, where $a \in \mathcal{R}$ and $k \in K$. Choose $\ell \in L$ such that $|k - \ell| \leq d_H(K, L)$. Then $r(a + \ell) \in F_r(L)$ and $|r(a + k) - r(a + \ell)| \leq rd_H(K, L)$. The symmetric argument, exchanging K and L , gives $d_H(F_r(K), F_r(L)) \leq rd_H(K, L)$. Thus, F_r is a contraction with constant r .
3. $A_r(\mathcal{R})$ is the unique fixed point of F_r . Since F_r is a contraction, it has a unique fixed point (Banach's fixed-point theorem). We first need to show that $A_r(\mathcal{R})$ is indeed in the domain of F_r . First note the set equation

$$A_r(\mathcal{R}) = \bigcup_{a \in \mathcal{R}} (ra + rA_r(\mathcal{R})) = F_r(A_r(\mathcal{R})), \quad (12)$$

- which follows by separating the first digit of the series $\sum_{i=1}^{\infty} a_i r^i = ra_1 + r \sum_{i=2}^{\infty} a_i r^{i-1}$. Because $\mathcal{R}$ is compact, $\mathcal{R}^{\mathbb{N}}$ is compact (Thychonoff's theorem). The map $\eta : \mathcal{R}^{\mathbb{N}} \rightarrow \mathbb{R}$, defined as $(x_i)_{i \in \mathbb{N}} \mapsto \sum_{i=1}^{\infty} x_i r^i$ is continuous, so $A_r(\mathcal{R})$ is the image of a compact set through a continuous map, therefore compact and hence an element of $\mathcal{K}$. By the Banach fixed point theorem, F_r has a unique fixed point in $\mathcal{K}$; by (12), that fixed point is $A_r(\mathcal{R})$.
4. A particular interval is also a fixed point under the gap condition. Define

$$I_r := \left[\frac{rm}{1-r}, \frac{rM}{1-r} \right].$$

For each $a \in \mathcal{R}$,

$$ra + rI_r = \left[ra + \frac{r^2 m}{1-r}, ra + \frac{r^2 M}{1-r} \right],$$

so

$$F_r(I_r) = \bigcup_{a \in \mathcal{R}} \left[ra + \frac{r^2 m}{1-r}, ra + \frac{r^2 M}{1-r} \right].$$

The leftmost endpoint occurs at $a = m$ and equals $\frac{rm}{1-r}$, and the rightmost endpoint occurs at $a = M$ and equals $\frac{rM}{1-r}$. Moreover, if $a < b$ are such that there is no point of $\mathcal{R}$ in (a, b) , then the corresponding two intervals overlap provided

$$ra + \frac{r^2M}{1-r} \geq rb + \frac{r^2m}{1-r} \iff r(b-a) \leq \frac{r^2}{1-r}(M-m).$$

Since the largest such gap is Δ , the condition

$$\Delta \leq \min\left\{\frac{r}{1-r}, \frac{s}{1-s}\right\}(M-m) \quad (13)$$

implies all adjacent pieces overlap, hence $F_r(I_r)$ is a connected compact set with the same endpoints as I_r , i.e. $F_r(I_r) = I_r$. Therefore I_r is a fixed point of F_r .

5. *Conclude $A_r(\mathcal{R}) = I_r$.* Finally, by uniqueness of the fixed point of F_r in $\mathcal{K}$ and since both $A_r(\mathcal{R})$ and I_r are fixed points, we get $A_r(\mathcal{R}) = I_r$.

Applying the same argument with s in place of r , and using (13) (which implies $\Delta \leq \frac{r}{1-r}(M-m)$ and $\Delta \leq \frac{s}{1-s}(M-m)$), we obtain

$$A_r(\mathcal{R}) = \left[\frac{rm}{1-r}, \frac{rM}{1-r}\right], \quad A_s(\mathcal{R}) = \left[\frac{sm}{1-s}, \frac{sM}{1-s}\right]. \quad (14)$$

Step 2: the two-sided tail set is an interval.

$$I_{\text{tail}} := A_r(\mathcal{R}) + A_s(\mathcal{R}).$$

Since the Minkowski sum of intervals is an interval, we have

$$I_{\text{tail}} = \left[\frac{rm}{1-r} + \frac{sm}{1-s}, \frac{rM}{1-r} + \frac{sM}{1-s}\right].$$

Step 3: adding the center digit yields $[\lambda m, \lambda M]$. Using $\text{Sums}_{r,s}^{\mathcal{R}} = \mathcal{R} + I_{\text{tail}}$, we can write

$$\text{Sums}_{r,s}^{\mathcal{R}} = \bigcup_{a \in \mathcal{R}} (a + I_{\text{tail}}).$$

Each set $a + I_{\text{tail}}$ is an interval of length

$$|I_{\text{tail}}| = \left(\frac{r}{1-r} + \frac{s}{1-s}\right)(M-m),$$

so adjacent pieces overlap whenever the gap between adjacent points of $\mathcal{R}$ is at most $|I_{\text{tail}}|$. But by (13),

$$\Delta \leq \min\left\{\frac{r}{1-r}, \frac{s}{1-s}\right\}(M-m) \leq \left(\frac{r}{1-r} + \frac{s}{1-s}\right)(M-m) = |I_{\text{tail}}|,$$

hence all adjacent intervals overlap and the union is a single interval.

The left endpoint is obtained by taking $a = m$ and the left endpoint of I_{tail} :

$$m + \frac{rm}{1-r} + \frac{sm}{1-s} = \lambda m,$$

32:24 Monitoring Discounted Sum Properties

and similarly the right endpoint is

$$M + \frac{rM}{1-r} + \frac{sM}{1-s} = \lambda M.$$

Therefore $\text{Sums}_{r,s}^{\mathcal{R}} = [\lambda m, \lambda M]$.

(2) The set $\text{Closure}(\mathcal{R})$ is compact, has the same infimum and supremum as $\mathcal{R}$, and satisfies

$$\Delta_{\text{Closure}(\mathcal{R})} \leq \Delta_{\mathcal{R}}.$$

Hence, $(\text{Closure}(\mathcal{R}), r, s)$ is proper, and part (3) gives

$$\text{Sums}_{\text{Closure}(\mathcal{R})}^{r,s} = [\lambda m, \lambda M].$$

Moreover, $\mathcal{R}^{\mathbb{Z}}$ is dense in $\text{Closure}(\mathcal{R})^{\mathbb{Z}}$, and the discounted-sum function is continuous because its geometric tails converge uniformly on the bounded input domain. Therefore,

$$\text{Closure}(\text{Sums}_{\mathcal{R}}^{r,s}) = \text{Sums}_{\text{Closure}(\mathcal{R})}^{r,s} = [\lambda m, \lambda M].$$

◀

Lemma 8. *Let $(\mathcal{R}, r, s)$ be an input setting, let $0 \leq t \leq n$, and let $\gamma_{t,n}^{r,s} := \frac{r^{t+1}}{1-r} + \frac{s^{n-t+1}}{1-s}$. Then, $U_t^{r,s}(w_{0:n}) \subseteq S_t^{r,s}(w_{0:n}) + \gamma_{t,n}^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. If $(\mathcal{R}, r, s)$ is proper, then $\text{Closure}(U_t^{r,s}(w_{0:n})) = S_t^{r,s}(w_{0:n}) + \gamma_{t,n}^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. If, additionally, $\mathcal{R}$ is compact, then equality holds without taking closures.*

Proof. Let $m = \inf \mathcal{R}$ and $M = \sup \mathcal{R}$. For every extension of $w_{0:n}$, the unobserved contribution to the discounted sum at time t is

$$\sum_{i=t+1}^{\infty} r^i x_{t-i} + \sum_{i=n-t+1}^{\infty} s^i x_{t+i}.$$

The sums of the corresponding coefficients are

$$\sum_{i=t+1}^{\infty} r^i = \frac{r^{t+1}}{1-r}$$

and

$$\sum_{i=n-t+1}^{\infty} s^i = \frac{s^{n-t+1}}{1-s}.$$

Since every unobserved value lies in $[m, M]$, this proves

$$U_t^{r,s}(w_{0:n}) \subseteq S_t^{r,s}(w_{0:n}) + \gamma_{t,n}^{r,s}[m, M].$$

Under properness, the closures of the two one-sided tail sets are, by the argument used in Lemma 6,

$$\frac{r^{t+1}}{1-r}[m, M]$$

and

$$\frac{s^{n-t+1}}{1-s}[m, M].$$

Taking their Minkowski sum proves the stated closure equality. If $\mathcal{R}$ is compact, the two tail sets are compact and already equal to these intervals, so equality holds without closure. ◀

Theorem 10. *Let $(\mathcal{R}, r, s)$ be a proper input setting such that $\mathcal{R}$ is compact and $r + s > 0$, and let $J := \lambda^{r,s}[\inf \mathcal{R}, \sup \mathcal{R}]$. Let $\mathcal{I}$ be a target interval satisfying $J \cap \mathcal{I} \neq \emptyset$ and $J \setminus \mathcal{I} \neq \emptyset$. Then, for every $t \in \mathbb{N}$ and every $n \geq t$, there exists $w \in \mathcal{R}^{\mathbb{Z}}$ such that $\mu_{\mathcal{I}}(t, w) > n$.*

Proof. Let $m = \inf \mathcal{R}$, $M = \sup \mathcal{R}$, and fix $t \in \mathbb{N}$ and $n \geq t$. Set

$$\Gamma := \frac{r^{t+1}}{1-r} + \frac{s^{n-t+1}}{1-s}.$$

The assumptions on the target imply $m < M$, and $r + s > 0$ implies $\Gamma > 0$. For each prefix $u \in \mathcal{R}^{n+1}$, Lemma 8 and compactness give

$$U_t^{r,s}(u) = S_t^{r,s}(u) + \Gamma[m, M].$$

These are closed intervals of the same positive length $\Gamma(M - m)$. As u ranges over all prefixes, their union is exactly

$$\text{Sums}_{\mathcal{R}}^{r,s} = J.$$

Choose a boundary point $c \in J$ of $\mathcal{I}$ at which membership in $\mathcal{I}$ changes relative to J . Choose points $y_k \in J$ converging to c from the side whose membership differs from that of c . For each k , choose a prefix u_k such that $y_k \in U_t^{r,s}(u_k)$. If some $U_t^{r,s}(u_k)$ also contains c , it contains one accepted and one rejected feasible value, and we are done. Otherwise these fixed-length intervals lie strictly on the side of y_k . Their relevant endpoints converge to c . Because $\mathcal{R}^{n+1}$ is compact and the interval endpoints depend continuously on the prefix, a limiting prefix has an uncertainty interval with endpoint c and with positive length on the side of the y_k . This interval again contains both an accepted and a rejected feasible sum. Hence the resulting prefix admits two completions with opposite exact verdicts. A sound monitor must therefore remain inconclusive after observing through time n , so $\mu_{\mathcal{I}}(t, w) > n$ for any completion w of this prefix. ◀

Theorem 12. *Let $(\mathcal{R}, r, s)$ be an input setting, let $\mathcal{I}$ be a target interval, let $\varepsilon > 0$, and let $T \in \mathbb{N}$. Define*

$$\tau^*(\mathcal{R}, r, s, \varepsilon, T) := \min \left(\left\{ \tau \in \mathbb{N} : d_{\mathcal{R}} \left(\frac{r^{T+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right) \leq 2\varepsilon \right\} \cup \{\infty\} \right).$$

Then:

1. *If $\tau^* < \infty$, there exists an ε -approximately sound monitor that produces a verdict for every $t \geq T$ within τ^* steps.*
2. *Suppose that $\mathcal{R} = [m, M]$ with $m < M$, and write $J = \lambda^{r,s}[m, M]$. If $J \cap \mathcal{I}_{-\varepsilon} \neq \emptyset$ and $J \setminus \mathcal{I}_{+\varepsilon} \neq \emptyset$, then, for every ε -approximately sound monitor and every $\tau < \tau^*$, there exists $w \in \mathcal{R}^{\mathbb{Z}}$ on which the monitor does not produce a verdict for time T after τ steps.*

Proof. Let $m = \inf \mathcal{R}$, $M = \sup \mathcal{R}$, and $d_{\mathcal{R}} = M - m$.

(1) Upper bound. For $0 \leq t \leq n$, define the interval enclosure

$$C_t(w_{0:n}) := S_t^{r,s}(w_{0:n}) + \left(\frac{r^{t+1}}{1-r} + \frac{s^{n-t+1}}{1-s} \right) [m, M].$$

By Lemma 8,

$$U_t^{r,s}(w_{0:n}) \subseteq C_t(w_{0:n}).$$

32:26 Monitoring Discounted Sum Properties

Consider the monitor that returns $\top$ if $C_t(w_{0:n}) \subseteq \mathcal{I}_{+\varepsilon}$, returns $\perp$ if $C_t(w_{0:n}) \cap \mathcal{I}_{-\varepsilon} = \emptyset$, and returns $?$ otherwise. The inclusion above immediately gives approximate soundness. Suppose neither decisive condition holds. Then there are

$$y \in C_t(w_{0:n}) \cap \mathcal{I}_{-\varepsilon}$$

and

$$z \in C_t(w_{0:n}) \setminus \mathcal{I}_{+\varepsilon}.$$

Since $\mathcal{I}_{-\varepsilon} = (L + \varepsilon, U - \varepsilon)$ and $\mathcal{I}_{+\varepsilon} = (L - \varepsilon, U + \varepsilon)$, we have $|y - z| > 2\varepsilon$. Consequently, every enclosing interval of diameter at most 2ε is decisive. At observation time $n = t + \tau$, the diameter of the enclosure is

$$d_{\mathcal{R}} \left(\frac{r^{t+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right) \leq d_{\mathcal{R}} \left(\frac{r^{T+1}}{1-r} + \frac{s^{\tau+1}}{1-s} \right)$$

for every $t \geq T$. The definition of τ^* therefore proves point (1). When $d_{\mathcal{R}} > 0$ and $0 < s < 1$, the defining inequality is equivalent to

$$s^{\tau+1} \leq B_T := \left(\frac{2\varepsilon}{d_{\mathcal{R}}} - \frac{r^{T+1}}{1-r} \right) (1-s).$$

If $B_T \leq 0$, no finite τ satisfies it. If $B_T > 0$, the least nonnegative integer solution is

$$\max \{0, \lceil \log_s B_T \rceil - 1\}.$$

The cases $d_{\mathcal{R}} = 0$ and $s = 0$ follow directly from the defining inequality.

(2) Lower bound. By affine rescaling, it suffices to prove the result for $\mathcal{R} = [0, 1]$; the target interval and ε are rescaled by the same affine map. Write $\lambda = \lambda^{r,s}$ and fix $\tau < \tau^*$. Set

$$\Gamma := \frac{r^{T+1}}{1-r} + \frac{s^{\tau+1}}{1-s}.$$

Set

$$A := \lambda - \Gamma.$$

By the definition of τ^* ,

$$\Gamma > 2\varepsilon.$$

As $w_{0:T+\tau}$ ranges over $[0, 1]^{T+\tau+1}$, its observed discounted contribution at time T ranges over the full interval $[0, A]$. For a prefix having observed contribution $p \in [0, A]$, compactness and Lemma 8 give

$$U_T(w_{0:T+\tau}) = [p, p + \Gamma].$$

Let $\mathcal{I} = (L, U)$. Since

$$[0, \lambda] \setminus \mathcal{I}_{+\varepsilon} \neq \emptyset,$$

there is a feasible definitely rejecting value either to the left of the target or to its right. Suppose first that there is one on the left. Then $L - \varepsilon \geq 0$. The condition $[0, \lambda] \cap \mathcal{I}_{-\varepsilon} \neq \emptyset$ also gives $L + \varepsilon < \lambda$. Since $\Gamma > 2\varepsilon$, the interval

$$(L + \varepsilon - \Gamma, L - \varepsilon] \cap [0, A]$$

is nonempty. Choose p in this intersection. Then $[p, p + \Gamma]$ contains a point outside $\mathcal{I}_{+\varepsilon}$ and a point in $\mathcal{I}_{-\varepsilon}$. If the definitely rejecting value lies to the right, then $U + \varepsilon \leq \lambda$ and $U - \varepsilon > 0$. In this case

$$[U + \varepsilon - \Gamma, U - \varepsilon) \cap [0, A]$$

is nonempty. Choosing p in this intersection gives the same conclusion. Thus in either case there is a prefix with two completions, one whose discounted sum belongs to $\mathcal{I}_{-\varepsilon}$ and one whose discounted sum lies outside $\mathcal{I}_{+\varepsilon}$. Approximate soundness rules out both $\top$ and $\perp$ on this prefix. Hence no approximately sound monitor can guarantee a verdict after $\tau < \tau^*$ steps. $\blacktriangleleft$

C Statistical Discounted Monitor

C.1 Remarks

► **Remark 27.** We emphasise that $\mathbb{E}_{t-1}(X_t)$ is a *random variable* (it depends on the realised history up to time $t - 1$), whereas $\mathbb{E}(X_t) = \mathbb{E}(\mathbb{E}_{t-1}(X_t))$ is deterministic. This is the relevant notion in monitoring: at time t we evaluate a decision maker relative to the realised stream (via $\mathcal{F}_{t-1}$), rather than relative to a hypothetical “future” distribution.

► **Remark 28.** The deterministic register bound τ^* no longer applies automatically, because the statistical interval contains the additional term $\beta_{t,n}^{r,s}(\delta)$. A register can be safely released at time n whenever $\mathcal{U}_t^{r,s}(\delta; W_{0:n}) \subseteq \mathcal{I}_{+\varepsilon}$ or $\mathcal{U}_t^{r,s}(\delta; W_{0:n}) \cap \mathcal{I}_{-\varepsilon} = \emptyset$. A deterministic sufficient condition for release is $2\beta_{t,n}^{r,s}(\delta) + d_{\mathcal{R}}\gamma_{t,n}^{r,s} \leq 2\varepsilon$. For pointwise bounds this condition may only be used at deterministic horizons; for data-dependent release one must use the local or uniform bounds.

► **Remark 29.** The uniform bounds can be tightened up to constant factors by leveraging the Lipschitz constant of discounted sum properties Lemma 4 to construct a grid of time points on which to perform the union bounds.

► **Remark 30.** For the other soundness notions this pointwise optimality definition does not apply directly, because local and uniform soundness require validity over stretches of time. This allows one to trade tightness at one horizon against additional slack at another. Hence, Theorem 15 is restricted to pointwise bounds. We only remark that the additional log log term in local and uniform bounds is known to be necessary in classical undiscounted settings where $V_n \rightarrow \infty$, due to the law of the iterated logarithm [24].

► **Remark 31.** The implications of the impossibility result are particularly visible for the *expected discounted average*, where according to Theorem 15 the statistical error for the two-sided limit

$$\frac{\mathbb{S}_t^{r,s}}{\frac{1-rs}{(1-r)(1-s)}} \text{ is (up to constants) } O\left(\sqrt{\frac{1}{2}\sigma^2 \frac{(1-r)(1-s)(1+rs)}{(1+r)(1+s)(1-rs)} \log(2/\delta)}\right).$$

And as $r, s \uparrow 1$, the normalized weights spread over an increasingly long effective window approaching the limit average, and thus approaching 0.

C.2 Proofs

Basics. Here we define the distribution class, bounded range, and summable weight sequence used throughout the appendix. Let $\mathcal{D}$ be the set of all probability distributions over $\mathcal{R}^{\mathbb{Z}}$,

32:28 Monitoring Discounted Sum Properties

where $\mathcal{R} \subset \mathbb{R}$ is bounded and closed with $a = \inf \mathcal{R}$, $b = \sup \mathcal{R}$, and $d_{\mathcal{R}} = b - a$. Let $\mathbf{w} \in \mathbb{R}^{\mathbb{Z}}$ be a weight sequence such that $\|\mathbf{w}\|_1 := \sum_{i \in \mathbb{Z}} |w_i| < \infty$; hence also $\|\mathbf{w}\|_2^2 := \sum_{i \in \mathbb{Z}} w_i^2 < \infty$. Fix $P \in \mathcal{D}$ and let $\mathbf{X} = (X_i)_{i \in \mathbb{Z}}$ be the process with law P , adapted to the canonical past filtration $(\mathcal{F}_i)_{i \in \mathbb{Z}}$.

Expectations. We define the predictable target and the corresponding plug-in estimator. We write $\mathbb{E}_t(\cdot) = \mathbb{E}(\cdot \mid \mathcal{F}_t)$ for the expectation conditioned on the past and define the predictable target and its plug-in estimator as

$$\mu(P) := \sum_{i \in \mathbb{Z}} w_i \mathbb{E}_{i-1}(X_i) \quad \text{and} \quad \hat{\mu}(\mathbf{X}) := \sum_{i \in \mathbb{Z}} w_i X_i.$$

Both sums are well-defined since $X_i \in [a, b]$ and $\|\mathbf{w}\|_1 < \infty$.

Confidence intervals. We define the admissible confidence intervals and the minimax length criterion. A confidence interval is a function $I: \mathcal{R}^{\mathbb{Z}} \rightarrow \text{Interval}(\mathbb{R})$. For $\delta \in (0, 1)$ define

$$\text{Cl}_{\sigma}(\delta) := \left\{ I \mid \sup_{P \in \mathcal{D}_{\sigma}} \mathbb{P}_P(\mu(P) \notin I(\mathbf{X})) \leq \delta \right\}.$$

For $\mathbf{x} \in \mathcal{R}^{\mathbb{Z}}$, let $|I(\mathbf{x})| := \sup I(\mathbf{x}) - \inf I(\mathbf{x})$ and define the worst-case length $\ell(I) := \sup_{\mathbf{x} \in \mathcal{R}^{\mathbb{Z}}} |I(\mathbf{x})|$. Our objective is to bound the minimax length $\inf_{I \in \text{Cl}_{\sigma}(\delta)} \ell(I)$.

KL-Divergence bound. This lemma gives a quadratic upper bound on the KL divergence between two symmetric Bernoulli parameters.

► **Lemma 32.** *Let $x \in (-1/2, 1/2)$. Then*

$$D_{\text{KL}}\left(\frac{1}{2} + x \parallel \frac{1}{2} - x\right) \leq \frac{8x^2}{1 - 4x^2}.$$

In particular, if $|x| \leq 1/4$, then

$$D_{\text{KL}}\left(\frac{1}{2} + x \parallel \frac{1}{2} - x\right) \leq \frac{32}{3}x^2.$$

Proof. A direct calculation gives

$$D_{\text{KL}}\left(\frac{1}{2} + x \parallel \frac{1}{2} - x\right) = \left(\frac{1}{2} + x\right) \log \frac{1+2x}{1-2x} + \left(\frac{1}{2} - x\right) \log \frac{1-2x}{1+2x} = 2x \log \frac{1+2x}{1-2x}.$$

Because we have for $x \geq 0$,

$$\log \frac{1+2x}{1-2x} = \int_{-2x}^{2x} \frac{1}{1+u} du \leq \frac{4x}{1-2|x|} \leq \frac{4x}{1-4x^2}$$

we can bound the KL divergence as follows

$$D_{\text{KL}}\left(\frac{1}{2} + x \parallel \frac{1}{2} - x\right) \leq \frac{8x^2}{1 - 4x^2}.$$

If $|x| \leq 1/4$, then $1 - 4x^2 \geq 3/4$, which gives the second claim. ◀

Weighted sub-Gaussian lower bound. This lemma proves that every valid confidence interval must have length at least the minimum of the range and sub-Gaussian scales.

► **Lemma 33** (Bounded sub-Gaussian two-point lower bound). *Assume that there exist $x_-, x_+ \in \mathcal{R}$ with $x_+ - x_- = 2\sigma$. Then, for every $\delta \in (0, 1/4)$ and an universal constant $c > 0$.*

$$\inf_{I \in \text{Cl}_\sigma(\delta)} \ell(I) \geq c \min \left(\sigma \|\mathbf{w}\|_1, \sigma \|\mathbf{w}\|_2 \sqrt{\log(2/\delta)} \right)$$

Proof. If $\|\mathbf{w}\|_2 = 0$, the claim is trivial. Hence assume $\|\mathbf{w}\|_2 > 0$. It suffices to consider product measures supported on $\{x_-, x_+\}$. Let $L := \log(1/(4\delta))$, $B := 1 \vee L$, $a := \|\mathbf{w}\|_2/\sqrt{B}$, and $\kappa := 1/8$. Define

$$u_i := \kappa \operatorname{sgn}(w_i) \min \left(\frac{|w_i|}{a}, 1 \right), \quad p_i := \frac{1}{2} + u_i, \quad q_i := \frac{1}{2} - u_i.$$

Then $|u_i| \leq \kappa \leq 1/4$ and

$$\sum_i u_i^2 \leq \kappa^2 \sum_i \frac{w_i^2}{a^2} = \kappa^2 B.$$

Let P_p and P_q be the product measures under which $X_i = x_+$ with probability p_i and q_i , respectively, and $X_i = x_-$ otherwise. For every coordinate, the centred variable is supported in an interval of length 2σ , and hence is σ -sub-Gaussian by Hoeffding's lemma. Thus $P_p, P_q \in \mathcal{D}_\sigma$. Writing $\mu_p := \mu(P_p)$ and $\mu_q := \mu(P_q)$, we get

$$|\mu_p - \mu_q| = 2\sigma \left| \sum_i w_i(p_i - q_i) \right| = 4\sigma \sum_i w_i u_i = 4\sigma \kappa \sum_i |w_i| \min \left(\frac{|w_i|}{a}, 1 \right).$$

Set $A := \|\mathbf{w}\|_1$ and $C := \|\mathbf{w}\|_2 \sqrt{B} = \|\mathbf{w}\|_2^2/a$. Using the elementary inequality

$$\sum_i \min \left(|w_i|, \frac{w_i^2}{a} \right) \geq \frac{AC}{A+C} \geq \frac{1}{2} \min(A, C),$$

we obtain

$$\Delta := |\mu_p - \mu_q| \geq 2\sigma \kappa \min \left(\|\mathbf{w}\|_1, \|\mathbf{w}\|_2 \sqrt{B} \right). \quad (15)$$

Next, by Lemma 32,

$$D_{\text{KL}}(P_p \| P_q) = \sum_i D_{\text{KL}} \left(\frac{1}{2} + u_i \parallel \frac{1}{2} - u_i \right) \leq \frac{32}{3} \sum_i u_i^2 \leq \frac{32}{3} \kappa^2 B = \frac{1}{6} B.$$

Suppose, for contradiction, that $\ell(I) < \Delta$ for some $I \in \text{Cl}_\sigma(\delta)$. Define the test $\varphi(\mathbf{X}) := \mathbf{1}[\mu_p \in I(\mathbf{X})]$. Since an interval of length less than Δ cannot contain both μ_p and μ_q , coverage of I under P_p and P_q implies $\mathbb{P}_{P_p}(\varphi = 0) \leq \delta$ and $\mathbb{P}_{P_q}(\varphi = 1) \leq \delta$. If $L \geq 1$, then $B = L$ and the Bretagnolle–Huber inequality gives

$$2\delta \geq \mathbb{P}_{P_p}(\varphi = 0) + \mathbb{P}_{P_q}(\varphi = 1) \geq \frac{1}{2} \exp(-D_{\text{KL}}(P_p \| P_q)) > \frac{1}{2} e^{-L} = 2\delta,$$

a contradiction. If $L < 1$, then $B = 1$ and Pinsker's inequality gives

$$\text{TV}(P_p, P_q) \leq \sqrt{\frac{D_{\text{KL}}(P_p \| P_q)}{2}} \leq \sqrt{\frac{1}{12}} < \frac{1}{2}.$$

32:30 Monitoring Discounted Sum Properties

Thus every test satisfies

$$\mathbb{P}_{P_p}(\varphi = 0) + \mathbb{P}_{P_q}(\varphi = 1) \geq 1 - \text{TV}(P_p, P_q) > \frac{1}{2} > 2\delta,$$

again a contradiction. Hence $\ell(I) \geq \Delta$ for every $I \in \text{Cl}_\sigma(\delta)$. Finally, since $B = 1 \vee \log(1/(4\delta))$ is comparable to $\log(2/\delta)$ on $\delta \in (0, 1/4)$, Eq. (15) yields the stated bound after absorbing constants. $\blacktriangleleft$

Sub-Gaussian MGF bound. This lemma shows that the weighted estimation error is itself sub-Gaussian with variance proxy $\sigma^2 \|\mathbf{w}\|_2^2$.

► **Lemma 34** (Sub-Gaussian MGF bound). *For every $P \in \mathcal{D}_\sigma$ and every $\lambda \in \mathbb{R}$,*

$$\mathbb{E}_P [\exp(\lambda(\hat{\mu}(\mathbf{X}) - \mu(P)))] \leq \exp\left(\frac{\lambda^2}{2} \sigma^2 \|\mathbf{w}\|_2^2\right).$$

Proof. For $m \in \mathbb{N}$ define the truncations

$$\mu_m(P) := \sum_{i=-m}^m w_i \mathbb{E}_{i-1}(X_i), \quad \hat{\mu}_m(\mathbf{X}) := \sum_{i=-m}^m w_i X_i, \quad S_m := \hat{\mu}_m(\mathbf{X}) - \mu_m(P).$$

Then $S_m = \sum_{i=-m}^m w_i (X_i - \mathbb{E}_{i-1}(X_i))$. Since $X_i - \mathbb{E}_{i-1}(X_i)$ is conditionally σ -sub-Gaussian,

$$\mathbb{E} [\exp(\lambda w_i (X_i - \mathbb{E}_{i-1}(X_i))) | \mathcal{F}_{i-1}] \leq \exp\left(\frac{\lambda^2 w_i^2 \sigma^2}{2}\right) \quad \text{a.s.}$$

Iterating the tower property gives

$$\mathbb{E}[e^{\lambda S_m}] \leq \exp\left(\frac{\lambda^2 \sigma^2}{2} \sum_{i=-m}^m w_i^2\right).$$

Since $\|\mathbf{w}\|_1 < \infty$, we have $S_m \rightarrow \hat{\mu}(\mathbf{X}) - \mu(P)$ almost surely. By applying Fatou's lemma yields

$$\mathbb{E} [e^{\lambda(\hat{\mu}(\mathbf{X}) - \mu(P))}] \leq \exp\left(\frac{\lambda^2 \sigma^2}{2} \|\mathbf{w}\|_2^2\right).$$

$\blacktriangleleft$

Sub-Gaussian variance bound. This lemma converts the sub-Gaussian MGF bound into a variance bound.

► **Lemma 35** (Variance bound). *For every $P \in \mathcal{D}$,*

$$\text{Var}_P(\hat{\mu}(\mathbf{X}) - \mu(P)) \leq \frac{d_{\mathcal{R}}^2}{4} \|\mathbf{w}\|_2^2.$$

Proof. By Lemma 34, the random variable $Y := \hat{\mu}(\mathbf{X}) - \mu(P)$ is sub-Gaussian with proxy variance $(d_{\mathcal{R}} \|\mathbf{w}\|_2/2)^2$. For any centred sub-Gaussian Y with $\mathbb{E}[e^{\lambda Y}] \leq \exp(\lambda^2 \sigma^2/2)$, differentiating at $\lambda = 0$ gives $\text{Var}(Y) \leq \sigma^2$. Applying this with $\sigma^2 = (d_{\mathcal{R}} \|\mathbf{w}\|_2/2)^2$ proves the claim. $\blacktriangleleft$

Sub-Gaussian upper bound. This lemma constructs confidence intervals whose lengths match the lower bound up to constants.

► **Lemma 36** (Upper bound). *For every $\delta \in (0, 1)$,*

$$\inf_{I \in \text{CI}_\sigma(\delta)} \ell(I) \leq C \min \left(d_{\mathcal{R}} \|\mathbf{w}\|_1, \sigma \|\mathbf{w}\|_2 \sqrt{\log(2/\delta)} \right)$$

for a universal constant $C > 0$.

Proof. First, the deterministic range bound gives the interval

$$I_{\text{rng}}(\mathbf{x}) := \left[\sum_{i:w_i \geq 0} w_i a + \sum_{i:w_i < 0} w_i b, \sum_{i:w_i \geq 0} w_i b + \sum_{i:w_i < 0} w_i a \right].$$

Since $\mathbb{E}_{i-1}(X_i) \in [a, b]$ almost surely, this interval covers $\mu(P)$ for every $P \in \mathcal{D}_\sigma$. Its length is $d_{\mathcal{R}} \|\mathbf{w}\|_1$. Second, by Chernoff's method and Lemma 34,

$$\sup_{P \in \mathcal{D}_\sigma} \mathbb{P}_P(|\hat{\mu}(\mathbf{X}) - \mu(P)| \geq \varepsilon) \leq 2 \exp \left(-\frac{\varepsilon^2}{2\sigma^2 \|\mathbf{w}\|_2^2} \right).$$

Choosing $\varepsilon = \sqrt{2\sigma^2 \|\mathbf{w}\|_2^2 \log(2/\delta)}$ yields the interval

$$I_H(\mathbf{x}) := \hat{\mu}(\mathbf{x}) \pm \sqrt{2\sigma^2 \|\mathbf{w}\|_2^2 \log \frac{2}{\delta}} \in \text{CI}_\sigma(\delta),$$

whose length is $2\sqrt{2\sigma^2 \|\mathbf{w}\|_2^2 \log(2/\delta)}$. Taking the better of I_{rng} and I_H proves the claim. ◀

Generic minimax theorem. This theorem combines the upper and lower bounds into a minimax characterisation for weighted sums.

► **Theorem 37** (Minimax confidence interval length). *For every $\delta \in (0, 1/4)$,*

$$c \min \left(\sigma \|\mathbf{w}\|_1, \sigma \|\mathbf{w}\|_2 \sqrt{\log(2/\delta)} \right) \leq \inf_{I \in \text{CI}_\sigma(\delta)} \ell(I) \leq C \min \left(d_{\mathcal{R}} \|\mathbf{w}\|_1, \sigma \|\mathbf{w}\|_2 \sqrt{\log(2/\delta)} \right)$$

for universal constants $0 < c < C < \infty$.

Proof. Combine Lemma 33 and Lemma 36. ◀

Discounted sum minimax theorem. This theorem specialises the generic weighted-sum minimax bound to finite discounted-sum weights.

► **Theorem 38** (Specialisation to discounted sums). *Fix $r, s \in [0, 1)$, a center time $t \in \{0, \dots, n\}$, and a horizon $n \geq t$. Consider the finite discounted-sum weights*

$$w_{t-i} = r^i \ (i = 1, \dots, t), \quad w_t = 1, \quad w_{t+i} = s^i \ (i = 1, \dots, n-t), \quad w_k = 0 \text{ otherwise.}$$

Then

$$\|\mathbf{w}\|_2^2 = \sum_{i=1}^t r^{2i} + \sum_{i=0}^{n-t} s^{2i} = \omega_{t,n}^{r,s}, \quad \|\mathbf{w}\|_1 = \sum_{i=1}^t r^i + \sum_{i=0}^{n-t} s^i = \eta_{t,n}^{r,s}.$$

Consequently,

$$c_0 \cdot \min \left(\sigma \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta) \right) \leq \inf_{I \in \text{CI}_\sigma(\delta)} \sup_{P \in \mathcal{D}_\sigma} \mathbb{E}_P[|I(W_{0:n})|] \leq C_0 \cdot \min \left(d_{\mathcal{R}} \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta) \right)$$

for universal constants $0 < c_0 < C_0 < \infty$.

32:32 Monitoring Discounted Sum Properties

Proof. The identities for $\|\mathbf{w}\|_1$ and $\|\mathbf{w}\|_2^2$ follow by geometric-series summation. For these weights, $\hat{\mu}(\mathbf{X}) = \mathbb{S}_t^{r,s}(W_{0:n})$ and $\mu(P) = \mathbb{S}_t^{r,s}(W_{0:n})$. Moreover,

$$\sigma \|\mathbf{w}\|_2 \sqrt{\log(2/\delta)} \asymp \sqrt{2\sigma^2 \omega_{t,n}^{r,s} \log(2/\delta)} = \text{PE}_{t,n}^{r,s}(\delta),$$

where the constants are universal. The proof concludes by applying Theorem 37 and then Lemma 39. $\blacktriangleleft$

Paper minimax theorem. This paragraph derives the paper's minimax theorem from the generic weighted-sum result.

► **Theorem 15.** *Given an input setting $(\mathcal{R}, r, s)$ s.t. $\mathcal{R}$ contains an interval of length at least 2σ . Fix $n \in \mathbb{N}$ and $t \leq n$. Consider the class $\mathcal{D}_\sigma$ of product measures on $\mathcal{R}^{n+1}$ s.t. $X_i - \mathbb{E}_{i-1}(X_i)$ is conditionally σ -sub-Gaussian. We define: the set of all $(1 - \delta)$ confidence intervals $\text{Cl}_\sigma(\delta)$ for $\mathbb{S}_t^{r,s}(W_{0:n})$, i.e., all $I : \mathcal{R}^{n+1} \rightarrow \text{Interval}(\mathbb{R})$ satisfying $\sup_{P \in \mathcal{D}_\sigma} \mathbb{P}_{W \sim P}(\mathbb{S}_t^{r,s}(W_{0:n}) \notin I(W_{0:n})) \leq \delta$, $|I| = \sup I - \inf I$ as interval length, and $\eta_{t,n}^{r,s} = \sum_{i=1}^t r^i + \sum_{i=0}^{n-t} s^i$. There exist universal constants $0 < c_0 < C_0 < \infty$ s.t. $\delta \in (0, 1/4)$*

$$c_0 \cdot \min(\sigma \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta)) \leq \inf_{I \in \text{Cl}_\sigma(\delta)} \sup_{P \in \mathcal{D}_\sigma} \mathbb{E}_P[|I(W_{0:n})|] \leq C_0 \cdot \min(d_{\mathcal{R}} \eta_{t,n}^{r,s}, \text{PE}_{t,n}^{r,s}(\delta)).$$

Proof of Theorem 15. This is a direct instantiation of Theorem 37 (the weighted-sum minimax theorem), with the finite-horizon discounted-sum weights

$$w_{t-i} = r^i, \quad w_t = 1, \quad w_{t+i} = s^i, \quad (i \geq 1), \quad w_k = 0 \text{ otherwise on } \{0, \dots, n\}.$$

For these weights,

$$\|\mathbf{w}\|_2^2 = \sum_{i=1}^t r^{2i} + \sum_{i=0}^{n-t} s^{2i} = \omega_{t,n}^{r,s}, \quad \|\mathbf{w}\|_1 = \sum_{i=1}^t r^i + \sum_{i=0}^{n-t} s^i = \eta_{t,n}^{r,s},$$

and $\hat{\mu}(\mathbf{X}) = \mathbb{S}_t^{r,s}(W_{0:n})$, $\mu(P) = \mathbb{S}_t^{r,s}(W_{0:n})$. Substituting these identities into Theorem 37 yields the claimed bound. Finally, Lemma 39 identifies worst-case length with worst-case expected length over the stated class, so the minimax statements match. $\blacktriangleleft$

Expected and worst-case length. This lemma shows that worst-case expected length and worst-case pointwise length coincide when the class contains all Dirac measures.

► **Lemma 39** (Expected vs. worst-case length). *Fix $n \in \mathbb{N}$ and a distribution class $\mathcal{D}$ on $\mathcal{R}^{\{0, \dots, n\}}$ that contains, for every $\mathbf{x} \in \mathcal{R}^{\{0, \dots, n\}}$, the degenerate (Dirac) product measure $P_{\mathbf{x}}$ with $P_{\mathbf{x}}(W_{0:n} = \mathbf{x}) = 1$. For any interval-valued map $I : \mathcal{R}^{\{0, \dots, n\}} \rightarrow \text{Interval}(\mathbb{R})$ define*

$$\ell(I) := \sup_{\mathbf{x} \in \mathcal{R}^{\{0, \dots, n\}}} |I(\mathbf{x})| \quad \text{then} \quad \sup_{P \in \mathcal{D}} \mathbb{E}_P[|I(W_{0:n})|] = \ell(I).$$

In particular, for any constraint set $\text{Cl}_\sigma(\delta)$,

$$\inf_{I \in \text{Cl}_\sigma(\delta)} \sup_{P \in \mathcal{D}} \mathbb{E}_P[|I(W_{0:n})|] = \inf_{I \in \text{Cl}_\sigma(\delta)} \ell(I).$$

Proof. Fix I .

(Upper bound) For any $P \in \mathcal{D}$ we have $|I(W_{0:n})| \leq \ell(I)$ almost surely, hence $\mathbb{E}_P[|I(W_{0:n})|] \leq \ell(I)$. Taking the supremum over P gives

$$\sup_{P \in \mathcal{D}} \mathbb{E}_P[|I(W_{0:n})|] \leq \ell(I).$$

(Lower bound) For every $\varepsilon > 0$ choose $\mathbf{x}_\varepsilon$ with $|I(\mathbf{x}_\varepsilon)| \geq \ell(I) - \varepsilon$ by definition of the supremum. Since $P_{\mathbf{x}_\varepsilon} \in \mathcal{D}$ and $W_{0:n} = \mathbf{x}_\varepsilon$ a.s. under $P_{\mathbf{x}_\varepsilon}$,

$$\mathbb{E}_{P_{\mathbf{x}_\varepsilon}}[|I(W_{0:n})|] = |I(\mathbf{x}_\varepsilon)| \geq \ell(I) - \varepsilon.$$

Thus $\sup_{P \in \mathcal{D}} \mathbb{E}_P[|I(W_{0:n})|] \geq \ell(I) - \varepsilon$ for all $\varepsilon > 0$, so letting $\varepsilon \downarrow 0$ yields

$$\sup_{P \in \mathcal{D}} \mathbb{E}_P[|I(W_{0:n})|] \geq \ell(I).$$

Combining both bounds gives equality. The minimax identity follows by taking $\inf_{I \in \text{CI}_\sigma(\delta)}$. ◀

Soundness of tail completion. This lemma shows that finite-sum coverage lifts to infinite-sum coverage after adding the deterministic discounted tail.

► **Lemma 40.** *Let $(\mathcal{R}, r, s)$ be an input setting and let $\Gamma := [\inf \mathcal{R}, \sup \mathcal{R}]$. For every $t \in \mathbb{N}$ and $n \geq t$ $\{\mathbb{S}_t^{r,s}(W_{0:n}) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n})\} \subseteq \{\mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n})\}$. Consequently,*

$$\begin{aligned} \mathbb{P}(\mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n})) &\geq \mathbb{P}(\mathbb{S}_t^{r,s}(W_{0:n}) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n})), \\ \mathbb{P}(\forall n \geq t : \mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n})) &\geq \mathbb{P}(\forall n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n})), \\ \mathbb{P}(\forall t \in \mathbb{N} \forall n \geq t : \mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n})) &\geq \mathbb{P}(\forall t \in \mathbb{N} \forall n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n})). \end{aligned}$$

Proof of Lemma 40. Fix $t \in \mathbb{N}$ and $n \geq t$. Write the infinite expected discounted sum as the finite observed part plus the unobserved tails:

$$\mathbb{S}_t^{r,s}(W) - \mathbb{S}_t^{r,s}(W_{0:n}) = \sum_{i=t+1}^{\infty} r^i \mathbb{E}_{t-i-1}(X_{t-i}) + \sum_{i=n-t+1}^{\infty} s^i \mathbb{E}_{t+i-1}(X_{t+i}).$$

Since $X_k \in \mathcal{R}$ almost surely, also $\mathbb{E}_{k-1}(X_k) \in \Gamma$ almost surely. Hence

$$\mathbb{S}_t^{r,s}(W) - \mathbb{S}_t^{r,s}(W_{0:n}) \in \left(\sum_{i=t+1}^{\infty} r^i + \sum_{i=n-t+1}^{\infty} s^i \right) \Gamma = \gamma_{t,n}^{r,s} \cdot \Gamma.$$

Therefore, whenever $\mathbb{S}_t^{r,s}(W_{0:n}) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n})$, we have

$$\mathbb{S}_t^{r,s}(W) \in \mathcal{I}_t^{r,s}(\delta; W_{0:n}) + \gamma_{t,n}^{r,s} \cdot \Gamma = \mathcal{U}_t^{r,s}(\delta; W_{0:n}).$$

This proves the event inclusion. The pointwise, local, and uniform inequalities follow by applying the same inclusion respectively for fixed (t, n) , for fixed t and all $n \geq t$, and for all $t \in \mathbb{N}$ and all $n \geq t$. ◀

Pointwise soundness. This lemma gives a fixed-time confidence interval for the finite expected discounted sum.

► **Lemma 41.** *Let $(\mathcal{R}, r, s)$ be a input setting. Then for every $\delta \in (0, 1)$ and all $n \in \mathbb{N}$, $t \in \{0, \dots, n\}$,*

$$\mathbb{P}(\mathbb{S}_t^{r,s}(W_{0:n}) \in \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) \geq 1 - \delta, \quad \beta_{t,n}^{r,s}(\delta) := \sqrt{2\sigma^2 \omega_{t,n}^{r,s} \log(2/\delta)}.$$

Proof of Lemma 41. Fix $n \in \mathbb{N}$ and $t \in \{0, \dots, n\}$. Define weights $(w_k)_{k=0}^n$ by

$$w_{t-i} = r^i \ (i = 1, \dots, t), \quad w_t = 1, \quad w_{t+i} = s^i \ (i = 1, \dots, n-t).$$

32:34 Monitoring Discounted Sum Properties

Then

$$S_t^{r,s}(W_{0:n}) - \mathbb{S}_t^{r,s}(W_{0:n}) = \sum_{k=0}^n w_k (X_k - \mathbb{E}_{k-1}(X_k)).$$

By Lemma 34 applied with sub-Gaussian proxy σ in place of $d_{\mathcal{R}}$, for all $\lambda \in \mathbb{R}$,

$$\mathbb{E} \left(\exp \left(\lambda \sum_{k=0}^n w_k (X_k - \mathbb{E}_{k-1}(X_k)) \right) \right) \leq \exp \left(\frac{\lambda^2}{2} \sigma^2 \sum_{k=0}^n w_k^2 \right).$$

Chernoff's method gives

$$\mathbb{P}(S_t^{r,s}(W_{0:n}) - \mathbb{S}_t^{r,s}(W_{0:n}) \geq \varepsilon) \leq \exp \left(-\frac{\varepsilon^2}{2\sigma^2 \sum_{k=0}^n w_k^2} \right),$$

and the same for the left tail. Using $\sum_{k=0}^n w_k^2 = \omega_{t,n}^{r,s}$ and setting $\varepsilon = \sqrt{2\sigma^2 \omega_{t,n}^{r,s} \log(2/\delta)} = \beta_{t,n}^{r,s}(\delta)$ results in $\mathbb{P}(|S - \mathbb{S}| \geq \beta_{t,n}^{r,s}(\delta)) \leq \delta$. $\blacktriangleleft$

Local soundness. This lemma gives an anytime-valid confidence interval over all horizons $n \geq t$ for a fixed center time t .

► **Lemma 42.** *Let $(\mathcal{R}, r, s)$ be a input setting, $t \in \mathbb{N}$ and $\delta \in (0, 1)$. Then:*

$$\begin{aligned} \mathbb{P}(\forall n \geq t : S_t^{r,s}(W_{0:n}) \in \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) &\geq 1 - \delta, \\ \beta_{t,n}^{r,s}(\delta) &:= k_1 \sqrt{V_{t,n}^{r,s} \left(2 \log(\log_2(V_{t,n}^{r,s}) + 1) + \log\left(\frac{2\pi^2}{6\delta}\right) \right)} \end{aligned}$$

where $V_{t,n}^{r,s} := \max(1, \sigma^2 \omega_{t,n}^{r,s})$ and $k_1 := 2^{1/4} + 2^{-1/4}/\sqrt{2}$.

Proof of Lemma 42. Let $S_n := S_t^{r,s}(W_{0:n}) - \mathbb{S}_t^{r,s}(W_{0:n})$. For fixed t , define weights $w_k^{(t)}$ by $w_{t-i}^{(t)} = r^i$ for $i = 1, \dots, t$ and $w_{t+i}^{(t)} = s^i$ for $i \geq 0$. For $m \geq 0$, let

$$S_m := \sum_{k=0}^m w_k^{(t)} (X_k - \mathbb{E}_{k-1}(X_k)).$$

Then $(S_m)_{m \geq 0}$ is a martingale with conditionally σ -sub-Gaussian increments, i.e.,

$$V_{t,n}^{r,s} := \max \left(1, \sigma^2 \sum_{k=0}^n w_k^2 \right) = \max(1, \sigma^2 \omega_{t,n}^{r,s}).$$

Then for every $n \geq t$, $S_n = S_t^{r,s}(W_{0:n}) - \mathbb{S}_t^{r,s}(W_{0:n})$, we apply the stitched sub-Gaussian uniform boundary of [24, Thm. 1] with $\eta = 2$, $m = 1$, and $h(j) = a(j+1)^2$ where $a = \pi^2/6$, and union bound the two one-sided boundaries, splitting δ into $\delta/2$, to obtain

$$\mathbb{P}(\forall n \geq t : |S_n| \leq \beta_{t,n}^{r,s}(\delta)) \geq 1 - \delta,$$

with $\beta_{t,n}^{r,s}(\delta)$ exactly as stated after substituting $V_{t,n}^{r,s}$. Finally, $|S_n| \leq \beta_{t,n}^{r,s}(\delta)$ is equivalent to $S_t^{r,s}(W_{0:n}) \in \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)$. $\blacktriangleleft$

Uniform soundness. This lemma obtains simultaneous coverage over all center times t and all horizons $n \geq t$ by a union bound over the local guarantees.

► **Lemma 43.** *Let $(\mathcal{R}, r, s)$ be a input setting and $\delta \in (0, 1)$. Then:*

$$\mathbb{P}(\forall t \in \mathbb{N} \forall n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \in \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) \geq 1 - \delta,$$

$$\beta_{t,n}^{r,s}(\delta) := k_1 \sqrt{V_{t,n}^{r,s} \left(2 \log(\log_2(V_{t,n}^{r,s}) + 1) + \log\left(\frac{2\pi^2}{6\delta}\right) \right)} \quad \text{where} \quad \delta_t := \frac{6\delta}{\pi^2(t+1)^2}.$$

where $V_{t,n}^{r,s} := \max(1, \sigma^2 \omega_{t,n}^{r,s})$ and $k_1 := 2^{1/4} + 2^{-1/4}/\sqrt{2}$.

Proof of Lemma 43. This follows directly from a simple union bound.

$$\begin{aligned} & \mathbb{P}(\forall t \in \mathbb{N} \forall n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \in \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) \\ &= 1 - \mathbb{P}(\exists t \in \mathbb{N} \exists n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \notin \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) \\ &\geq 1 - \sum_{t=0}^{\infty} \mathbb{P}(\exists n \geq t : \mathbb{S}_t^{r,s}(W_{0:n}) \notin \mathbb{S}_t^{r,s}(W_{0:n}) \pm \beta_{t,n}^{r,s}(\delta)) \\ &\geq 1 - \sum_{t=0}^{\infty} \delta_t = 1 - \sum_{t=0}^{\infty} \frac{6\delta}{\pi^2(t+1)^2} = 1 - \delta. \end{aligned}$$

◀

Paper soundness. This paragraph derives the paper's statistical soundness theorem from finite-sum coverage and tail completion.

► **Theorem 14.** *Let $(\mathcal{R}, r, s)$ be an input setting, $\mathcal{I} = (L, U)$ a target interval, $\varepsilon \geq 0$, and $\delta \in (0, 1)$. Consider Algorithm 1 where a register for time t is released at a fixed (n) or at a flexible ($\mathcal{U}_t^{r,s}(\delta; W_{0:n}) \subseteq \mathcal{I}_{+\varepsilon}$ or $\mathcal{U}_t^{r,s}(\delta; W_{0:n}) \cap \mathcal{I}_{-\varepsilon} = \emptyset$) observation horizon. Then*

- (i) *if $\beta_{t,n}^{r,s}(\delta) = \text{PE}_{t,n}^{r,s}(\delta)$ the monitor is pointwise (ε, δ) -approx. sound for fixed release;*
- (ii) *if $\beta_{t,n}^{r,s}(\delta) = \text{LE}_{t,n}^{r,s}(\delta)$, the monitor is locally (ε, δ) -approx. sound for flexible release;*
- (iii) *if $\beta_{t,n}^{r,s}(\delta) = \text{UE}_{t,n}^{r,s}(\delta)$, the monitor is uniformly (ε, δ) -approx. sound for flexible release.*

Proof. By Lemmas 41, 42, and 43, respectively, the intervals $\mathcal{I}_t^{r,s}(\delta; W_{0:n})$ cover the finite expected discounted sum $\mathbb{S}_t^{r,s}(W_{0:n})$ with the pointwise, local, and uniform guarantees. By Lemma 40, the corresponding tail-completed intervals $\mathcal{U}_t^{r,s}(\delta; W_{0:n})$ cover the infinite expected discounted sum $\mathbb{S}_t^{r,s}(W)$ with the same type of guarantee. On this coverage event, if the monitor outputs $\top$, then $\mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n}) \subseteq \mathcal{I}_{+\varepsilon}$. If the monitor outputs $\perp$, then

$$\mathbb{S}_t^{r,s}(W) \in \mathcal{U}_t^{r,s}(\delta; W_{0:n}) \quad \text{and} \quad \mathcal{U}_t^{r,s}(\delta; W_{0:n}) \cap \mathcal{I}_{-\varepsilon} = \emptyset,$$

and therefore $\mathbb{S}_t^{r,s}(W) \notin \mathcal{I}_{-\varepsilon}$. This is exactly the required statistical approximate soundness. The pointwise case applies only when the verdict horizon is fixed in advance. ◀

D General Discounted Properties

D.1 Synchronous vs Asynchronous example

In Table 1, we show an example of a sequence of events with two types, to illustrate the evolution of the values and the corresponding asynchronous discounting.

i	0	1	2	3	4	5	6	7	8 ...
$e_i^{(1)}$	1	1	□	0	1	0	0	□	1 ...
$e_i^{(2)}$	□	0	1	1	□	□	0	□	1 ...
$x_i^{(1)}$	1	1	0	0	1	0	0	0	1 ...
$x_i^{(2)}$	0	0	1	1	0	0	0	0	1 ...
$\llbracket \cdot \rrbracket_S \quad \tau_{0,+i}^{(k)}$	0	1	2	3	4	5	6	7	8 ...
$\llbracket \cdot \rrbracket_A \quad \tau_{0,+i}^{(1)}$	0	1	1	2	3	4	5	5	6 ...
$\llbracket \cdot \rrbracket_A \quad \tau_{0,+i}^{(2)}$	0	0	1	2	2	2	3	3	4 ...

■ **Table 1** Excerpt of an event sequence, with the corresponding values and discount factors.

E Register Complexity of Discounted-Sum Monitoring

Theorem 23. *For every monitoring instance $\mathcal{P}$, there is an ARM $\mathcal{M}$ with $\tau(\mathcal{P}) + 1$ registers such that $L(\mathcal{M}) = L(\mathcal{P})$.*

Proof. Let $\tau = \tau(\mathcal{P})$. We construct an ARM $\mathcal{M} = (\mathcal{R}, Y, Q, q_0, \nu_0, \Delta)$. Let $Y = \{y_0, \dots, y_{\tau-1}, \rho\}$, where $y_0, \dots, y_{\tau-1}$ store running sums for pending positions, and ρ tracks r^{n+1} before reading the n th input. Let $Q = \{q_c : 0 \leq c \leq \tau\}$, where c counts how many positions are currently pending, capped at τ . The initial location is q_0 with $\nu_0(y_i) = 0$ for all i and $\nu_0(\rho) = r$.

From location q_c on input $x \in \mathcal{R}$, $\mathcal{M}$ applies the parallel update $\rho \leftarrow r\rho$, $y_0 \leftarrow ry_0 + x$, and $y_a \leftarrow y_{a-1} + s^a x$ for $a \in \{1, \dots, \min\{c, \tau - 1\}\}$, and moves to $q_{\min\{c+1, \tau\}}$. In location q_τ , before applying the update above, $\mathcal{M}$ tests whether the oldest pending position t can be resolved using the current input x . At this step $t = n - \tau$, and by the maintained invariant $\rho = r^{n+1}$ and $y_{\tau-1} = S_t^{r,s}(x_{0:t+\tau-1})$. Thus the completed truncated value is $V := y_{\tau-1} + s^\tau x = S_t^{r,s}(x_{0:t+\tau})$. Writing $\delta_f := \frac{s^{\tau+1}}{1-s}$, we define $\kappa_{r,\tau} := \frac{r^{-\tau}}{1-r}$ if $r > 0$, and $\kappa_{r,\tau} := 0$ otherwise. Since $t = n - \tau$ and $\rho = r^{n+1}$, we have $\delta_p(t) = \frac{r^{t+1}}{1-r} = \kappa_{r,\tau}\rho$. Therefore, $\mathcal{M}$ includes in q_τ a self-loop guarded by $V > L - \varepsilon \wedge V + \kappa_{r,\tau}\rho < U + \varepsilon - \delta_f$, and blocks if the guard fails.

By induction on n , just before reading x_n we have $\rho = r^{n+1}$ and for every age $a < \min\{n, \tau\}$, $y_a = S_{n-1-a}^{r,s}(x_{0:n-1})$. In particular, in q_τ the guard is satisfied iff $U_t^{r,s}(x_{0:t+\tau}) \subseteq \mathcal{I}_{+\varepsilon}$ for $t = n - \tau$. Hence $\mathcal{M}$ blocks on an input stream exactly when some position violates the monitoring condition, and therefore $L(\mathcal{M}) = L(\mathcal{P})$. ◀

Lemma 25. *For every $k \geq 1$, the monitoring instance $\mathcal{P}_k$ satisfies $\tau(\mathcal{P}_k) = k$ and $L(\mathcal{P}_k) = \{x \in [0, 1]^\omega \mid \forall t \geq 0 : V_t^k(x) < T_k\}$.*

Proof. By Theorem 12, $\tau(\mathcal{P}_k) = \tau^*([0, 1], 0, \frac{1}{2}, 2^{-(k+1)}, 0) = \lceil \log_{1/2}(2^{-(k+1)}) \rceil - 1 = k$.

Since $r = 0$, for every $t \geq 0$ we have $S_t^{0,1/2}(x_{0:t+k}) = \sum_{i=0}^k 2^{-i} x_{t+i} = V_t^k(x)$. The residual uncertainty after observing k future symbols is $\gamma_{t,t+k}^{0,1/2} = \frac{(1/2)^{k+1}}{1-1/2} = 2^{-k}$. Hence the condition $U_t^{0,1/2}(x_{0:t+k}) \subseteq (0, 1)_{+\varepsilon}$ with $\varepsilon = 2^{-(k+1)}$ is equivalent to $-\varepsilon < V_t^k(x) < 1 + \varepsilon - 2^{-k}$. The left inequality is trivial because $V_t^k(x) \geq 0$, and the right inequality simplifies to $V_t^k(x) < 1 - 2^{-(k+1)} = T_k$. Therefore $x \in L(\mathcal{P}_k)$ iff $V_t^k(x) < T_k$ for all $t \geq 0$. ◀

Theorem 26. *For every $k \geq 1$, there is a future-only monitoring instance $\mathcal{P}_k$ such that for every ARM $\mathcal{M}$ with $\tau(\mathcal{P}_k) - 1$ registers we have $L(\mathcal{M}) \neq L(\mathcal{P}_k)$.*

Proof. Let $k \geq 1$ and $\mathcal{P}_k$ be the monitoring instance defined in Section 5. For brevity, write $V_t := V_t^k$ and $T := T_k$. Suppose toward contradiction that there exists an ARM $\mathcal{M}$ with $m := k - 1$ registers such that $L(\mathcal{M}) = L(\mathcal{P}_k)$.

Let us record two identities that follow from the definition of the monitoring instance $\mathcal{P}_k$. For every $t \geq 0$,

$$V_t(x) = x_t + \frac{1}{2}V_{t+1}(x) - 2^{-(k+1)}x_{t+k+1}, \quad (16)$$

$$V_j(x') - V_j(x) = \sum_{i=j}^{k-1} 2^{-(i-j)}(x'_i - x_i) \quad (0 \leq j \leq k-1), \quad (17)$$

whenever $x_i = x'_i$ for all $i \geq k$.

Fixed control paths. For a length- k transition sequence π of $\mathcal{M}$, let $C_\pi \subseteq [0, 1]^k$ be the set of input prefixes $u = (u_0, \dots, u_{k-1})$ such that, during the first k steps while reading u , the machine executes exactly the transitions of π . Unrolling $\mathcal{M}$ along π shows that C_π is the intersection of $[0, 1]^k$ with finitely many strict affine half-spaces, hence is relatively open in $[0, 1]^k$. It also yields for each π a matrix $A_\pi \in \mathbb{R}^{m \times k}$ and a vector $b_\pi \in \mathbb{R}^m$ such that for every $u \in C_\pi$, the register valuation after k steps is $A_\pi u + b_\pi$. Since $m = k - 1 < k$, every matrix A_π has a nontrivial kernel.

Near-boundary prefixes and an anchor word. Let $\alpha := T/2$. For $\theta \in (3T/4, T]$, define $u(\theta) \in \mathbb{R}^k$ by setting $u_i(\theta) = \theta/2$ for $0 \leq i \leq k-2$ and $u_{k-1}(\theta) = \theta - \frac{T}{4}$. All coordinates of $u(\theta)$ lie in $(0, 1)$, so $u(\theta) \in (0, 1)^k$.


We first note that for every $\theta < T$ the word $w_\theta := u(\theta) \cdot \alpha \cdot 0^\omega$ belongs to $L(\mathcal{P}_k)$. Indeed, for every $j \in \{0, \dots, k-1\}$ we have $V_j(w_\theta) = \theta$. For $j = k-1$, the window contains $u_{k-1}(\theta)$ with weight 1 and α with weight $1/2$, so $V_{k-1}(w_\theta) = (\theta - \frac{T}{4}) + \frac{1}{2} \cdot \frac{T}{2} = \theta$. For $j \leq k-2$, the symbol at position $j+k+1$ is 0, so Equation (16) gives $V_j(w_\theta) = u_j(\theta) + \frac{1}{2}V_{j+1}(w_\theta) = \frac{\theta}{2} + \frac{1}{2}\theta = \theta$, by backward induction from the case $j = k-1$. Moreover, for $t \geq k$, the only possible nonzero symbol still visible is the single letter α at position k , hence $V_t(w_\theta) \leq \alpha = T/2 < T$. Therefore $w_\theta \in L(\mathcal{P}_k)$ whenever $\theta < T$.

Next, we use the prefix $u(T)$ to fix a control path. Let $\hat{w} := u(T) \cdot 0^\omega$. Again, $\hat{w} \in L(\mathcal{P}_k)$: we have $V_{k-1}(\hat{w}) = T - \frac{T}{4} = \frac{3T}{4} < T$, and for $j \leq k-2$, we have $V_j(\hat{w}) = u_j(T) + \frac{1}{2}V_{j+1}(\hat{w}) < \frac{T}{2} + \frac{1}{2}T = T$ by backward induction, while $V_t(\hat{w}) = 0$ for all $t \geq k$.

Since $\hat{w} \in L(\mathcal{P}_k)$ and $\mathcal{M}$ recognizes $L(\mathcal{P}_k)$, the run of $\mathcal{M}$ on $\hat{w}$ is infinite. In particular, while reading the prefix $u(T)$ during the first k steps, the machine follows some length- k transition sequence π^* .

An invisible perturbation inside one control cell. Consider the transition sequence π^* above. Let $O := C_{\pi^*} \cap (0, 1)^k$ be the set of length- k prefixes that force the first k transitions of $\mathcal{M}$ to be exactly π^* . Since O is open in $\mathbb{R}^k$ and $u(T) \in O$, there is a neighborhood of $u(T)$ in which the control path is stable: Let $\eta > 0$ be such that the open ball $B_\infty(u(T), \eta) := \{v \in \mathbb{R}^k : \|v - u(T)\|_\infty < \eta\}$ satisfies $B_\infty(u(T), \eta) \subseteq O$.

Choose any nonzero vector $d \in \ker(A_{\pi^*})$. Intuitively, d encodes an “invisible” direction for the register valuations after k following π^* . Let j be the largest index such that $d_j \neq 0$, and replace d by $-d$ if necessary so that $d_j > 0$. Let $\lambda_{\min} := \frac{\eta}{2\|d\|_\infty}$ and $\delta := \min\{\eta/2, T/4\}$. We pick λ such that $0 < \lambda < \min\{\lambda_{\min}, 2\delta/d_j\}$, and define $\theta := T - \frac{\lambda d_j}{2}$. Then, $\theta \in (T - \delta, T) \subseteq (3T/4, T)$, so $u(\theta)$ is well-defined.



■ **Figure 3** Normalized active registers and discounted values, MNIST dataset.

By the explicit form of $u(\cdot)$ we have $\|u(\theta) - u(T)\|_\infty = |T - \theta| = \frac{\lambda d_j}{2} < \delta \leq \eta/2$, and therefore $u(\theta) \in B_\infty(u(T), \eta/2)$. Also, $\|\lambda d\|_\infty < \lambda_{\min} \|d\|_\infty = \eta/2$, so $\|u(\theta) + \lambda d - u(T)\|_\infty \leq \|u(\theta) - u(T)\|_\infty + \|\lambda d\|_\infty < \eta$, which implies $u(\theta) + \lambda d \in B_\infty(u(T), \eta) \subseteq O$.

Let $u := u(\theta)$ and $u' := u + \lambda d$ and define $w := u \cdot \alpha \cdot 0^\omega$ and $w' := u' \cdot \alpha \cdot 0^\omega$. Because $u, u' \in C_{\pi^*}$, the runs of $\mathcal{M}$ on w and w' take the same first k transitions. Moreover, $u' - u = \lambda d$ and $d \in \ker(A_{\pi^*})$, so the register valuations after k steps coincide. Since the suffixes from position k onward are identical, determinism yields $w \in L(\mathcal{M})$ iff $w' \in L(\mathcal{M})$.

Contradiction. Since $\theta < T$, the construction above gives $w \in L(\mathcal{P}_k)$. On the other hand, w and w' agree at all positions $i \geq k$, so Equation (17) yields $V_j(w') - V_j(w) = \lambda \sum_{i=j}^{k-1} 2^{-(i-j)} d_i$. By the choice of j as the largest index with $d_j \neq 0$, we have $d_i = 0$ for all $i > j$, and therefore $V_j(w') - V_j(w) = \lambda d_j$. Since $V_j(w) = \theta$, it follows that $V_j(w') = \theta + \lambda d_j > T$ by the definition of θ . Hence $w' \notin L(\mathcal{P}_k)$. However, since we assumed $L(\mathcal{M}) = L(\mathcal{P}_k)$ and showed that $w \in L(\mathcal{M})$ iff $w' \in L(\mathcal{M})$, we obtain a contradiction. ◀

F Further Experimental Evaluation


Because of space restrictions, we include extra experiments supporting our evaluation in this appendix.

F.1 RQ1: Resource Usage

In addition to the experiments with **PowerData** traces, we also test our monitors for resource monitoring classification accuracy of a neural network trained on the MNIST digit dataset. In particular, we are interested in how the monitoring evolves under a shift in the underlying distribution.

For MNIST, we train a neural network on the training set and generate a trace of length 10000 by sequentially sampling test instances, recording whether each prediction is correct. Thus, $\mathcal{R} = \{0, 1\}$, with discount factors fixed to $r = s = 0.95$, target interval $\mathcal{I} = [0.96, 1]$, and tolerance $\varepsilon = 0.001$.

We induce distribution shift by adding Gaussian noise with time-increasing variance $\sigma(t)$ to the test inputs; the parameter `max_noise` controls the final corruption level. As shown in Fig. 3(a), increasing noise leads to reduced register usage. Figures 3(b) and (c) illustrate a representative execution, plotting the number of active registers and the monitored discounted accuracy, respectively. As accuracy degrades over time, the monitor can more quickly conclude violation, requiring fewer active registers. Consistent with the **PowerData** results, using only 70% of the theoretical register budget would still be mostly sound.



■ **Figure 4** Active registers when monitoring demographic parity with variable target interval width, `Adult` dataset.

F.2 RQ2: Monitoring Demographic Parity

To further investigate RQ2, we monitor demographic parity on the `Adult` dataset with a fixed tolerance $\varepsilon = 5 \cdot 10^{-5}$ and decreasing interval width. As before, the input set is $\mathcal{R} = \{0, 1\}$ and we again fix $r = s = 0.95$. We report the results in Fig. 4, with the same format Fig. 2, for the demographic parity property (Ex. 3) with the synchronous (blue) and asynchronous (green) interpretation. As expected, and validating the results on Fig. 2, synchronous interpretations are significantly more efficient than asynchronous ones, since discounting is applied uniformly across streams rather than independently.

F.3 RQ3: Statistical Monitoring Behaviour


Our objective is to assess the behaviour of the statistical monitor during deployment. Since the statistical uncertainty intervals do not converge to width zero, the main question is power, i.e., whether the monitor eventually produces a decisive verdict. To evaluate this question, we use a synthetic piecewise-stationary Beta process. We use synthetic data because the evaluation requires access to the ground-truth expected discounted average.

Setup. The Beta process consists of four phases, each spanning 150 time steps. At each time index, the observation X_t is sampled independently from a Beta distribution parametrised by (a, b) , where $(a, b) = (1, 9), (5, 5), (8, 2), (4, 6)$. The corresponding phase means are approximately 0.10, 0.50, 0.80, and 0.40. We use a uniform upper bound on the conditional sub-Gaussian norm of approximately $\sigma = 0.15$. The monitor observes only the realised samples X_t , while the latent process $P_t = \mathbb{E}_{t-1}(X_t)$ is used only for evaluation.

We monitor the expected discounted average $\mathbb{S}_t^{r,s}(W)/\lambda^{r,s}$ where $\lambda^{r,s} = 1 + \frac{r}{1-r} + \frac{s}{1-s}$. In the experiment, $r = s = 0.95$, so $\lambda^{r,s} \approx 39$. The target interval is $I = [0.4, 0.6]$, the tolerance is $\varepsilon = 0.05$, and the error probability is $\delta = 0.01$. Thus, positive verdicts are sound with respect to $I_{+\varepsilon} = [0.35, 0.65]$, while negative verdicts are sound with respect to the complement of $I_{-\varepsilon} = [0.45, 0.55]$.

We use the process depicted in Fig. 5 to evaluate the pointwise, local, and uniform statistical error bounds. Since $r = s = 0.95$, the discounted average is strongly smoothed: near phase boundaries, it blends past and future phases instead of following the latent mean instantaneously. The figure therefore illustrates the main purpose of discounted monitoring: it suppresses single-sample noise while still reacting to changes in the local behaviour.

Important: In the experiments below, we also use flexible release with the pointwise bound. This is for demonstration purposes only: the goal is to visualise how the



■ **Figure 5** Four-phase Beta process $\text{Beta}(a, b)$. The figure shows the realised observations X_t , the latent mean process P_t , the observed discounted average $S_t^{r,s}(W)/\lambda^{r,s}$, and the oracle expected discounted average $S_t^r(W)/\lambda^{r,s}$. The shaded background indicates the four Beta phases.

statistical uncertainty intervals evolve over time, independently of the choice of a fixed release horizon.

Monitor verdicts. We first assess how many observations are required before each monitored time index t receives a decisive verdict. Fig. 6 shows the resulting verdicts for the pointwise, local, and uniform statistical error bounds. The observed behaviour matches the phase structure of the process. Time indices centred in the low and high phases tend to receive negative verdicts, while time indices centred in the middle phases tend to receive positive verdicts. Around phase boundaries, verdicts are delayed or remain inconclusive because the discounted average mixes information from neighbouring phases.

The number of inconclusive verdicts increases with the strength of the soundness guarantee. This is also visible in Fig. 7, which overlays the statistical uncertainty intervals with the resulting verdicts. Pointwise intervals are the narrowest, local intervals are wider because they permit flexible release over all horizons $n \geq t$, and uniform intervals are widest because they protect all monitored time indices simultaneously. Consequently, pointwise verdicts tend to appear earliest, while uniform verdicts are the most conservative.

Statistical uncertainty interval evolution. We next decompose the half-width of the statistical uncertainty interval. Fig. 8 shows that the deterministic tail error decreases as the observation horizon n grows, because more future observations have been revealed. The statistical error term, in contrast, increases toward its limiting value, since more noisy observations enter the discounted estimate. Thus, the total uncertainty is governed by a tradeoff: waiting reduces deterministic tail error, but exposes the estimate to its limiting statistical error.

Fig. 9 shows the statistical uncertainty interval at the time of the first decisive verdict and at the end of the run, assuming that the register is never released. This highlights the tension between statistical certainty and resource consumption: delaying release may reduce the deterministic tail error, but it also keeps registers active for longer and cannot remove the limiting statistical error.

Monte Carlo violation rates. We use a Monte Carlo simulation with 1000 repetitions to evaluate the coverage of the statistical uncertainty intervals, the verdict frequency, the verdict




Figure 6 Monitor verdicts. Rows correspond to monitored time indices t , columns to observation horizons n , and colours encode the verdict: $\top$ for inside, $\perp$ for outside, and $?$ for inconclusive.

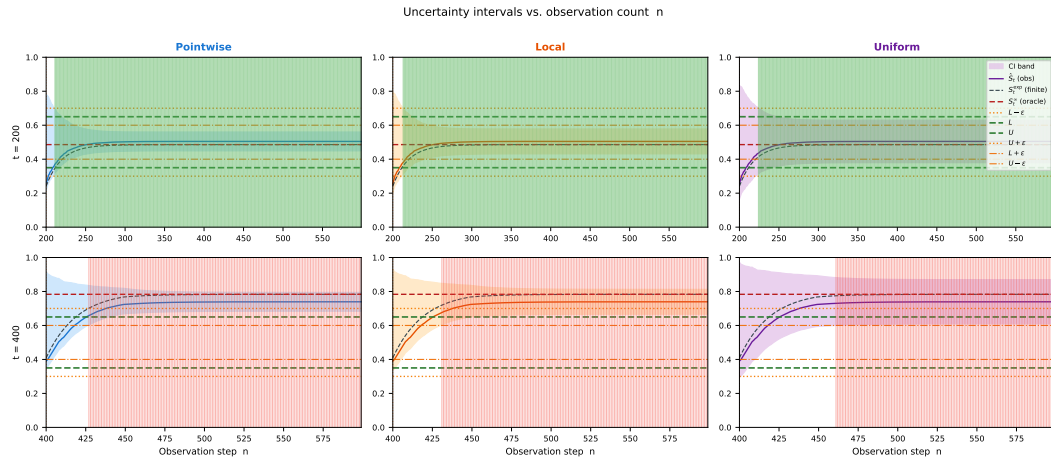



Figure 7 Evolution of pointwise, local, and uniform statistical uncertainty intervals. The figure fixes several monitored time indices t and shows how the uncertainty interval evolves as more observations $n \geq t$ become available. The solid line is the observed discounted average, the dashed line is the finite expected discounted average, and the horizontal red line is the oracle expected discounted average.



■ **Figure 8** Uncertainty half-width decomposition. The figure decomposes the half-width of the statistical uncertainty interval into the deterministic tail term $d_{\mathcal{R}}\gamma_{t,n}^{r,s}$ and the statistical error term $\beta_{t,n}^{r,s}(\delta)$.

delay, and the verdict correctness. Table 2 reports Monte Carlo averages with standard deviations over independent runs. The first block uses the setting of this section. The second block uses a lower-variance setting, where all Beta parameters are multiplied by 10, i.e., $\text{Beta}(10a, 10b)$, giving a sub-Gaussian norm upper bound of approximately $\sigma = 0.05$. We report the average statistical uncertainty interval violation rate, both pointwise over all (t, n) pairs and run-wise as the fraction of runs containing at least one interval violation. We also report the fraction of monitored time indices that eventually receive a decisive verdict, the average release delay $n - t$, the fraction of incorrect first verdicts among released verdicts, and the fraction of runs containing at least one incorrect verdict. The statistical uncertainty intervals are conservative in these experiments. Only the pointwise bound exceeds the nominal error probability in the run-wise interval violation metric, which is expected because pointwise soundness does not protect flexible release or whole-run events. Compared with the lower-variance setting, the higher-variance setting issues fewer verdicts and has larger release delays. Across all configurations, incorrect verdicts are rare.

F.4 RQ4: Statistical Uncertainty Intervals

Our objective is to assess how the parameter choices affect the half-width of the statistical uncertainty interval. For all experiments in this subsection, we set the past discount factor to zero, i.e., $r = 0$, and monitor the expected discounted average. If not explicitly state we set the error probability to $\delta = 0.01$.

Convergence behaviour. In Fig. 10, we study how the statistical error term and the deterministic tail error converge to their limiting values. We vary the future discount factor, the sub-Gaussian norm, and the soundness notion. Higher sub-Gaussian norms lead to larger statistical error terms and therefore wider statistical uncertainty intervals. Larger future discount factors spread the normalised weights over a longer effective window, which reduces the limiting statistical error for the discounted average. At the same time, larger future discount factors increase the deterministic tail error for finite observation horizons, so more observations are required before the tail error becomes negligible.





Figure 9 Statistical uncertainty intervals at the point of first decisive verdict and at the end of the run.

■ **Table 2** Aggregate statistics of the Monte Carlo experiments with 1000 repetitions. *Coverage:* Interval viol. is the fraction of statistical uncertainty interval violations; Any interval is the fraction of runs with at least one statistical uncertainty interval violation. *Release:* Released is the fraction of monitored time indices that receive a decisive verdict; Delay is the average release delay. *Verdicts:* Wrong rate is the fraction of incorrect first verdicts among released verdicts; Any wrong is the fraction of runs with at least one incorrect verdict.

(a, b) Sound.	Coverage		Release		Verdicts	
	Interval viol.	Any interval	Released	Delay	Wrong rate	Any wrong
Pointwise	0.000 ± 0.003	0.021 ± 0.143	0.924 ± 0.014	21.279 ± 0.920	0.000 ± 0.000	0.002 ± 0.045
Local	0.000 ± 0.000	0.000 ± 0.000	0.877 ± 0.025	24.204 ± 1.067	0.000 ± 0.000	0.000 ± 0.000
Uniform	0.000 ± 0.000	0.000 ± 0.000	0.662 ± 0.024	33.304 ± 1.436	0.000 ± 0.000	0.000 ± 0.000

$(10a, 10b)$ Sound.	Coverage		Release		Verdicts	
	Interval viol.	Any interval	Released	Delay	Wrong rate	Any wrong
Pointwise	0.000 ± 0.002	0.014 ± 0.118	0.958 ± 0.002	14.372 ± 0.178	0.000 ± 0.000	0.000 ± 0.000
Local	0.000 ± 0.000	0.000 ± 0.000	0.955 ± 0.002	15.277 ± 0.191	0.000 ± 0.000	0.000 ± 0.000
Uniform	0.000 ± 0.000	0.000 ± 0.000	0.944 ± 0.003	18.335 ± 0.243	0.000 ± 0.000	0.000 ± 0.000

Limit behaviour. In Fig. 11, we study the limiting statistical error bound as a function of the future discount factor. The error probability has a comparatively small effect relative to the sub-Gaussian norm. Decreasing the error probability increases the bound logarithmically, while increasing the sub-Gaussian norm increases the bound directly. The bounds become practically useful mainly for large discount factors close to 1, where the discounted average aggregates over a longer effective window.

Uniform soundness. In Fig. 12, we study the effect of the monitored time index t on the uniform statistical error bound. The uniform bound uses a time-dependent error budget, and therefore increases with t . As expected, the increase is logarithmic in the monitored time index.




Figure 10 Convergence of the error components to their limiting values. The figure shows different sub-Gaussian norms (rows) and future discount factors (hue) for each soundness notion (columns).




Figure 11 Effect of the future discount factor on the limiting statistical error bound for different error probabilities δ (hue) and sub-Gaussian norms σ (style). To visualise the effects, the axis scales vary between columns.




Figure 12 Effect of the monitored time index on the uniform statistical error bound. The left plot shows the effect on a log-log scale for different discount factors (hue) and sub-Gaussian norms (style). The right plot shows the same effect on a linear scale for one representative parameter choice.