Statistical Model Checking Beyond Means: Quantiles, CVaR, and the DKW Inequality* (extended version)

Carlos E. Budde¹, Arnd Hartmanns², Tobias Meggendorfer³, Maximilian Weininger⁴, and Patrick Wienhöft^{5,6}

- ¹ Technical University of Denmark, Lyngby, Denmark
 - ² University of Twente, Enschede, The Netherlands
 - ³ Lancaster University Leipzig, Leipzig, Germany
 - ⁴ Ruhr University Bochum, Bochum, Germany
 - ⁵ Dresden University of Technology, Germany
- ⁶ Centre for Tactile Internet with Human-in-the-Loop (CeTI), Dresden, Germany

Abstract. Statistical model checking (SMC) randomly samples probabilistic models to approximate quantities of interest with statistical error guarantees. It is traditionally used to estimate probabilities and expected rewards, i.e. means of different random variables on paths. In this paper, we develop methods using the Dvoretzky-Kiefer-Wolfowitz-Massart inequality (DKW) to extend SMC beyond means to compute quantities such as quantiles, conditional value-at-risk, and entropic risk. The DKW provides confidence bounds on the random variable's entire cumulative distribution function, a much more versatile guarantee compared to the statistical methods prevalent in SMC today. We have implemented support for computing new quantities via the DKW in the MODES simulator of the MODEST TOOLSET. We highlight the implementation and its versatility on benchmarks from the quantitative verification literature.

1 Introduction

Statistical model checking (SMC) [1,31,33,47] avoids the state space explosion problem of classic probabilistic model checking approaches (PMC) [4,5] that explore and numerically analyse a model's entire state space [27]: SMC instead samples k random paths from the model to estimate the value of the quantity of interest. As a simulation-based approach, it applies to any effectively executable model, including non-Markovian [19] and hybrid [22,39] ones. An SMC result

^{*} This work was supported by the DFG through the Cluster of Excellence EXC 2050/1 (CeTI, project ID 390696704, as part of Germany's Excellence Strategy) and the TRR 248 (see perspicuous-computing.science, project ID 389792660), by the ERC Starting Grant DEUCE (101077178), by the European Union's Horizon 2020 research and innovation programme under Marie Skłodowska-Curie grant agreement 101008233 (MISSION), by the Interreg North Sea project STORM_SAFE, and by NWO VIDI grant VI.Vidi.223.110 (TruSTy).

comes with a *statistical* correctness guarantee, often expressed as a *confidence* interval [l, u] that contains the true result $(1 - \delta) \cdot 100\%$ of the times [15].

The most fundamental quantities computed by PMC and SMC are reachability probabilities and expected rewards [15, 27]. SMC estimates these quantities using statistical methods like the Clopper-Pearson confidence interval [16] for probabilities (i.e. binomial proportions) and Hoeffding's inequality [29] for means of bounded distributions, or compares them to each other [20] or to thresholds using Wald's sequential probability ratio test [45]. In the past decade, PMC has been extended to compute several other quantities of interest, such as quantiles/percentiles/value-at-risk [30,41,44], conditional value at risk [32], and entropic risk [6]. However, the application of SMC has so far been limited to probabilities and expected rewards, i.e. only the means of distributions associated to different random variables on sampled paths. To the best of our knowledge, no SMC approaches or tools support quantities other than means yet.

In this paper, we show how to extend SMC to estimate non-mean quantities using the Dvoretzky-Kiefer-Wolfowitz-Massart inequality (DKW) [21, 37]. The DKW provides a sound simultaneous confidence band around the cumulative distribution function (cdf), i.e. upper and lower bound functions completely enveloping the (unknown) cdf $(1 - \delta) \cdot 100\%$ of the times (see Fig. 1). This is a stronger statement compared to the currently-used statistical methods for estimation mentioned above, as it applies to the entire cdf rather than a single point or pointwise. From the DKW, we can again derive a confidence interval for the mean [3, 15], but equally (and simultaneously) obtain confidence intervals on other quantities as well. We show how to do so in particular for higher moments, quantiles, conditional value-at-risk, and entropic risk. We have implemented these DKW-based computations in the MODES statistical model checker [13], part of the MODEST TOOLSET [26]. MODES can now also export the empirical cdf and DKW confidence band for plotting and further analysis by the user. We highlight our implementation and its versatility using several models from the Quantitative Verification Benchmark Set (QVBS) [28] in Sec. 4.

2 Preliminaries

A probability distribution over a non-empty, countable set S is a function $\mu \colon S \to [0,1]$ such that $\sum_{s \in S} \mu(s) = 1$. The set of all distributions over S is denoted by $\mathcal{D}(S)$. The cumulative distribution function (cdf) of a random variable X is given by $F_X(x) \stackrel{\text{def}}{=} \mathbb{P}(X \leq x)$. A random variable X stochastically dominates another random variable Y, written $Y \lesssim_{SD} X$, if $F_Y(x) \geq F_X(x)$ for all x (i.e. for any x, obtaining a value less than or equal to x is more likely for Y than for X; intuitively, X yields larger values). If $Y \lesssim_{SD} X$, then $\mathbb{E}(Y) \leq \mathbb{E}(X)$.

Definition 1. A discrete-time Markov chain (DTMC) is a tuple $\langle S, R, T, s_I \rangle$ of a finite set of states S, a reward function $R: S \to \mathbb{R}_{\geq 0}$, an initial state $s_I \in S$, and a transition function $T: S \to \mathcal{D}(S)$ mapping each state to a probability distribution over successor states. A (finite) path π is (a prefix of) an infinite sequence $\pi = s_0 s_1 \ldots \in S^{\omega}$ such that $s_0 = s_I$ and $\forall i: T(s_i)(s_{i+1}) > 0$.




Fig. 1. Example of a DTMC (left) together with the probability distribution over possible reward outcomes (middle) and the corresponding cdf (right, solid line). The states in the DTMC are (only) labelled by their rewards. The right figure also includes an empirical cdf (dashed) and corresponding confidence band (gray) obtained from the DKW inequality (with $\delta = 0.1$ and k = 50).

See Fig. 1 (left) for an example of a DTMC. A DTMC induces a unique probability measure \mathbb{P} over sets of paths that, intuitively, corresponds to multiplying the probabilities along the path (see e.g. [7, Chapter 10]).

Properties. Properties typically consist of two parts: First, a random variable X assigning a value to each path. For our results, the choice of X is largely irrelevant; we only require it to yield non-negative finite values. Concretely, we consider *total* and *reachability rewards*, i.e. $\mathsf{TR}(\pi) = \sum_{i=0}^{\infty} R(\pi_i)$ and the same sum cut off at the first goal state, respectively; see [15, Sec. 2] for details, and Fig. 1 (middle/right) for the distribution and cdf of TR on the example DTMC.

Second, a property comes with an aggregation function to "summarize" X into a single value, traditionally the expected value/mean $\mathbb{E}(X)$ (w.r.t. \mathbb{P}). Recently, alternative aggregations have gained popularity, for example

- higher moments (around 0), which are of the form $\mathbb{E}(X^n)$ for n > 1;
- the *t-quantile* (a.k.a. the *value-at-risk*) for $t \in (0,1)$, which is the smallest value v such that X is less than or equal to v with probability t [30,41]:

$$Q_t(X) \stackrel{\text{def}}{=} \inf \{ v \mid \mathbb{P}(X \leq v) \geq t \};$$

- the *conditional value-at-risk* (a.k.a. expected shortfall, expected tail loss, average value-at-risk), which is the expectation over the t-quantile, i.e.

$$CVaR_t(X) \stackrel{\text{def}}{=} \frac{1}{t} (P \cdot \mathbb{E}[X \mid X < v] + (t - P) \cdot v)$$

where $t \in (0,1)$, $v = Q_t(X)$ and $P = \mathbb{P}(X < v)$ [32,42]; and

- the *entropic risk*, which with $\gamma > 0$ is [6, 23]

$$\operatorname{ERisk}_{\gamma}(X) \stackrel{\text{def}}{=} -\frac{1}{\gamma} \log(\mathbb{E}(e^{-\gamma X})).$$

We illustrate these for the DTMC of Fig. 1 in Sec. A. Additionally, as in [15], we distinguish whether X has a known upper bound (i.e. some U such that $\mathbb{P}(X \leq U) = 1$, the bounded case) or not (the general case).

Statistical model checking is, at its core, Monte Carlo simulation for formal models and properties: randomly generate a (predetermined) number k of paths, or *simulations*, from the model that give rise to samples X_1, \ldots, X_k of the

random variable X; and from that draw statistical conclusions on the property. While PMC approaches exist for all of the aforementioned properties, SMC so far exclusively focused on means as follows: compute the empirical mean

$$\hat{X} \stackrel{\text{def}}{=} \frac{1}{k} \sum_{i=1}^{k} X_i,$$

and perform a statistical evaluation to obtain a confidence interval $I = [l, u] \ni \hat{X}$ at a predetermined confidence level δ , so that with (a priori) probability $1 - \delta$ we have $\mathbb{E}(X) \in I$. That is, if we repeat the SMC procedure m times to obtain confidence intervals I_1, \ldots, I_m , we may find some of them (up to $\delta \cdot 100\%$ on average) incorrect, i.e. $\mathbb{E}(X) \notin I_i$ for some i. Occasionally obtaining an "incorrect" result is the nature of a statistical approach based on sampling. In this work, we develop statistical methods for other aggregations beyond the mean.

3 Statistical Guarantees Beyond Means

Before we introduce our approach, we formalise the exact kind of guarantees we aim to give. Observe that simply returning confidence intervals $[0, \infty]$ is always sound. However, we also want SMC procedures to yield "small" intervals. To formalize this requirement, we say a procedure yields *effective bounds* if (i) it produces correct intervals with high confidence, and (ii) for a large enough number k of samples, the intervals produced by the procedure are smaller than any ε and still correct with high confidence; see Definition 2 for the formal definition. We note that this is related to the notion of *consistent* estimators [2] from statistics, as the mid-point of effective intervals is a consistent estimator. However, we pose a stronger requirement since we require correct bounds to be produced.

Definition 2 (effective bounds). Let X be a random variable and \mathcal{F} an aggregator, mapping random variables to real numbers. An SMC procedure \mathcal{A} yields effective bounds on $\mathcal{F}(X)$ if, for any confidence $\delta > 0$, the following two conditions hold: (i) For a collection of independent samples Ξ drawn from X, we have $\mathbb{P}(\mathcal{F}(X) \in \mathcal{A}(\Xi, \delta)) \geq 1 - \delta$. (ii) For any precision $\varepsilon > 0$, there exists a threshold k_0 such that for a collection of independent samples Ξ drawn from X with $|\Xi| \geq k_0$, we have $\mathbb{P}(\mathcal{F}(X) \in \mathcal{A}(\Xi, \delta) \wedge |\mathcal{A}(\Xi, \delta)| \leq \varepsilon) \geq 1 - \delta$.

Remark 1. Some works consider the dual problem of gathering enough samples until a given precision is reached. They seek so-called probably approximately correct (PAC) guarantees: Given confidence level δ and precision ε , gather enough samples to return I with $|I| \leq 2\varepsilon$. We focus on deriving intervals given a fixed k, and in Sec. B describe how our methods extend to the dual problem.

As already observed in [15], obtaining two-sided bounds sometimes is infeasible (depending on the nature of the DTMC, random variable X, and aggregator \mathcal{F}). However, we may still be able to derive statistically sound, "converging" lower bounds. Thus, we extend the definition of "limit-PAC" from [15, Def. 3] and say an SMC procedure yields *effective lower bounds* if the value it produces is, with

high confidence, (i) always a lower bound and (ii) close to the true value if given enough samples. Formally:

Definition 3 (effective lower bounds). Let X be a random variable and \mathcal{F} an aggregator, mapping random variables to real numbers. An SMC procedure \mathcal{A} yields effective lower bounds on $\mathcal{F}(X)$ if, for any confidence $\delta > 0$, the following two conditions hold: (i) For a collection of independent samples Ξ drawn from X, we have $\mathbb{P}(\mathcal{A}(\Xi, \delta) \leq \mathcal{F}(X)) \geq 1 - \delta$. (ii) For any precision $\varepsilon > 0$, there exists a threshold k_0 such that for a collection of independent samples Ξ drawn from X with $|\Xi| \geq k_0$, we have $\mathbb{P}(\mathcal{F}(X) - \varepsilon \leq \mathcal{A}(\Xi, \delta) \leq \mathcal{F}(X)) \geq 1 - \delta$.

3.1 DKW: The Dvoretzky-Kiefer-Wolfowitz-Massart Inequality

Our key to obtain effective bounds is the *Dvoretzky-Kiefer-Wolfowitz-Massart* inequality (DKW), which relates the cdf of the unknown distribution of X to the empirical cdf $\hat{F}(x) = \frac{1}{k} |\{X_i \mid X_i \leq x\}|$ by

$$\mathbb{P}\Big(\sup_{x\in\mathbb{R}}|\hat{F}(x)-F_X(x)|>\Delta\Big)\leq\delta\quad\text{where }\Delta=\sqrt{\log(\delta/2)/(-2k)}.$$

Note that F_X is fixed but unknown, while \hat{F} depends on the samples drawn from X. Intuitively, the DKW gives a confidence band in which the true cdf lies with high probability; see Fig. 1 (right) for an illustration. There, \hat{F} is drawn dashed, and the gray area around \hat{F} shows the confidence band (with width 2Δ). We refer to the bounds of this band as $\underline{F}(x) \stackrel{\text{def}}{=} \min \{\hat{F}(x) + \Delta, 1\}$ and $\overline{F}(x) \stackrel{\text{def}}{=} \max \{0, \hat{F}(x) - \Delta\}$, respectively. We denote the random variables that $\underline{F}, \hat{F}, \hat{F$

3.2 Obtaining Effective Bounds

As it turns out, computing the aggregations for \underline{X} (and \overline{X}) gives effective (lower) bounds for all considered properties. We implicitly assume that the DKW condition holds and prove (below) that we then obtain correct (and converging) estimates. This means in general we get such estimates with high confidence. Moreover, as all results only depend on the DKW condition holding, we can give guarantees on *all* aggregations simultaneously, without splitting the confidence budget, which is particularly useful for e.g. multi-objective queries [30, 32, 41].

Moments. For higher-order moments, note that $Y \stackrel{\text{def}}{=} X^n$ is non-negative and has finite expectation if X satisfies these assumptions. Thus, the results of [15, Thm. 1] are directly applicable, which state that then the DKW yields effective lower bounds in the general case. In the bounded case, we naturally obtain effective bounds by direct application of the DKW (see Sec. 3.1).

Quantiles. By their definition, quantiles are monotone w.r.t. stochastic dominance, i.e. if $Y \lesssim_{SD} X$, then $Q_t(Y) \leq Q_t(X)$. Thus, we also have $Q_t(\underline{X}) \leq$ $Q_t(X) \leq Q_t(\overline{X})$, ensuring correctness of the computed values. While we can always obtain lower and upper bounds, even in the general case (by choosing k so that $t \in (\Delta, 1 - \Delta]$, only the lower bounds may be effective: Consider an X with distribution $\{1 \mapsto \frac{1}{2}, 2 \mapsto \frac{1}{2}\}$. We have $Q_{0.5}(X) = 1$, but any sound statistical upper bound on the cdf of X will have $\overline{F}(1) < F_X(1) = 0.5$, and thus always yield a 0.5-quantile of 2. This is a fundamental property of quantiles: they are not continuous w.r.t. small changes in the distribution. This already happens for the simple example in Fig. 1, as we illustrate in Sec. A. Thus, in general we cannot provide effective bounds. X is always discrete for DTMC as per Definition 1 and X = TR; we can have non-discrete X if we allow e.g. continuously-distributed random rewards, use other models like continuous-time Markov chains (CTMCs), or other properties. Then, if X is continuous or, at least, if F_X is continuous at $Q_t(X)$, we get effective bounds.

Conditional value-at-risk. CVaR is a distortion risk measure (as is Q_t), which are monotone w.r.t. stochastic dominance [46]. Thus we again immediately get $\text{CVaR}_t(\underline{X}) \leq \text{CVaR}_t(X) \leq \text{CVaR}_t(\overline{X})$. In contrast to general expectations, the bounded and general case do not differ: By assumption, we have $X < \infty$, hence there exists T such that $F_X(T) > 1 - \frac{t}{2}$. For a large enough k, we have $\Delta < \frac{t}{2}$, and $\overline{F}(T) \geq 1 - t$. Then, we know (with high confidence) that $X \leq T$ with probability t, i.e. $Q_t(\overline{X}) \leq T$ and therefore $CVaR_t(\overline{X}) \leq T < \infty$. Thus, we can directly bound $|\text{CVaR}_t(X) - \text{CVaR}_t(\overline{X})|$ by $T \cdot 2\Delta$, which goes to 0 for large enough k. Together, we obtain effective bounds in the general case.

Entropic risk. First, observe that if $Y \lesssim_{SD} X$, then $e^{-\gamma X} \lesssim_{SD} e^{-\gamma Y}$ (the order reverses as $e^{-\gamma x}$ is decreasing). Consequently, $\mathbb{E}(e^{-\gamma X}) \leq \mathbb{E}(e^{-\gamma Y})$, and thus $\operatorname{ERisk}_{\gamma}(Y) < \operatorname{ERisk}_{\gamma}(X)$ (recall that $\operatorname{ERisk} = -1/\gamma \cdot \ldots$). Hence, we get $\operatorname{ERisk}_{\gamma}(\underline{X}) \leq \operatorname{ERisk}_{\gamma}(X) \leq \operatorname{ERisk}_{\gamma}(\overline{X})$. While there is no strict "cut-off" as for CVaR, we argue that we can still bound the overall difference between X and the bounds \underline{X} and \overline{X} in general. We have

$$\operatorname{ERisk}_{\gamma}(\overline{X}) - \operatorname{ERisk}_{\gamma}(X) = -1/\gamma \cdot \log(\mathbb{E}(e^{-\gamma \overline{X}})/\mathbb{E}(e^{-\gamma X})).$$

We now apply two useful general facts about cdfs, namely that (i) $\mathbb{E}(X)$ = $\begin{array}{l} \int_x (1-F_X(x)) \cdot x \, dx \text{ and (ii) for a positive, continuous, strictly decreasing function} \\ f \text{ we have } F_{f(X)}(x) = 1 - F_X(f^{-1}(x)). \text{ We get} \\ \mathbb{E}(e^{-\gamma \overline{X}})/\mathbb{E}(e^{-\gamma X}) = \int \overline{F}(-\frac{1}{\gamma}\log(x))e^{-\gamma x} \, dx/\mathbb{E}(e^{-\gamma X}). \end{array}$

$$\mathbb{E}(e^{-\gamma \overline{X}})/\mathbb{E}(e^{-\gamma X}) = \int \overline{F}(-\frac{1}{\gamma}\log(x))e^{-\gamma x} \, dx/\mathbb{E}(e^{-\gamma X}).$$

Recall that $F_X(x) - \Delta \leq \overline{F}(x)$. Hence,

$$\begin{split} \int \overline{F}(-\frac{1}{\gamma}\log(x))e^{-\gamma x} \, dx / \mathbb{E}(e^{-\gamma X}) &\geq \int (F_X(-\frac{1}{\gamma}\log(x)) - \Delta)e^{-\gamma x} \, dx / \mathbb{E}(e^{-\gamma X}) \\ &= 1 - \Delta \int e^{-\gamma x} \, dx / \mathbb{E}(e^{-\gamma X}). \end{split}$$

Consequently, for $\Delta \to 0$ this expression converges to 1, and thus $\mathrm{ERisk}_{\gamma}(X)$ – $\operatorname{ERisk}_{\gamma}(X) = -\frac{1}{\gamma}\log(\ldots) \to 0$. The proof for \underline{X} is analogous.

		expected value		0.3-quantile		$\text{CVaR}_{0.3}$	
example	k	\hat{X}	conf. int.	est.	conf. int.	est.	conf. int.
coupon	100	13.11	$[10.08, \infty)$	10	[9, 11]	8.90	[4.56, 10.14]
	1000	12.90	$[11.82, \infty)$	10	[10, 11]	8.71	[7.28, 9.21]
$leader_sync$	100	1.17	$[0.89, \infty)$	1	[1,1]	1.00	[0.55, 1.00]
embedded	100	0.35	n/a	0.13	[0.10, 0.20]	0.35	n/a
	1000	0.33	n/a	0.17	[0.15, 0.19]	0.33	n/a

Table 1. Estimates and DKW confidence intervals for the examples' properties.

4 Tool Implementation

The Modes SMC tool [13] was recently extended with sound statistical methods for estimating means, including the DKW [15]. Now, in version 3.1.287, we added syntax for quantile and CVaR properties to the parsers for its input languages, Modest [9, 25] and Jani [14], and extended its implementation of the DKW to estimate and provide bounds for such properties. Additionally, empirical cdfs can be exported to CSV and Excel files for plotting and further analysis.

To demonstrate the new tool features, we use three examples from the QVBS selected for diversity in cdfs: (1) the *coupon* model with parameters N=15, DRAWS = 4, B = 5 (a DTMC of 17 billion states, to which SMC is agnostic) and the random variable underlying property exp_draws ; (2) $leader_sync$ with N=5, K=4 (DTMC, 4244 states) and time; and (3) embedded with MAX_COUNT = 8, T=12 (a continuous-time Markov chain of 8548 states) and $danger_time$. The original properties query for expected reachability rewards; we add properties querying for the 0.3-quantile and $CVaR_{0.3}$ of the same reward specification, i.e. the same random variable on paths. We run MODES on each example with k=100 simulations, and on coupon and embedded additionally with k=1000. In addition to obtaining DKW-based confidence intervals, we use MODES' new --cdf parameter to export empirical CDFs with DKW confidence bands.

In Table 1, we show the results that MODES obtains for the properties. As reachability rewards fall into the general case, we can only obtain lower bounds for the expected values [15]. For quantiles and CVaR, as per Sec. 3.2, the DKW allows us to obtain (for CVaR effective) lower and upper bounds. On *embedded*, MODES cannot apply the DKW to expectation and CVaR because its syntactic procedure to find a lower bound for the rewards fails as they are encoded via an unbounded real-valued variable. For quantiles, the absence of bounds on the distribution is no hindrance. The DKW can produce rather asymmetric confidence intervals, which we see for *leader sync*'s CVaR property.

We plot the empirical cdfs and associated DKW confidence bounds MODES delivered in Fig. 2. Graphically, the confidence interval for the quantile is the p=0.3 line's segment between the bound curves, while the CVaR estimate and confidence interval stem from the curves cut off at that line. The DTMCs' reward distributions are necessarily discrete: $leader_sync$ has only 3 possible outcomes




Fig. 2. DKW CDF confidence bands obtained on the example benchmarks.

of non-negligible probability (we thus omit k = 1000); increasing k for coupon does not smoothen the curve much, because most elements of the distribution's support were already sampled at k = 100—only the confidence band gets much thinner. For the CTMC *embedded*, the distribution is visibly continuous.

We do not report runtimes because the overhead of using the DKW—for collecting all samples to finally compute the intervals, instead of incremental averaging plus evaluation based on k and δ only as for traditional methods for the mean—was negligible in these experiments. Aside from DTMCs and CTMCs, MODES also supports more complex and expressive formalisms up to stochastic hybrid automata [24]; the new methods to check and bound quantiles and CVaRs work independent of the model type. They equally combine orthogonally with MODES' features for rare event simulation [12] as well as learning and scheduler sampling for nondeterministic models [17, 18, 38].

5 Conclusion

In this work, we have shown how the DKW inequality can be used to derive bounds on various aggregation functions beyond the classical expectation/mean, closing a significant gap of SMC compared to traditional verification. Moreover, as all our estimations are based on the DKW inequality, our methods can estimate all values simultaneously. Our experimental evaluation confirms the effectiveness of our methods, allowing for scalable estimation of such aggregation values for large systems. For future work, we believe that our approach should also be applicable to other risk measures such as variance [10,35], variance-penalized expected payoff [8,34,36,40], or cumulative prospect theory [11,43].

Data availability. The MODES tool is available at modestchecker.net. An artifact for this paper—a reproduction package with the models and commands for the experiments in Sec. 4—is available at DOI 10.5281/zenodo.15286509.

References

- Agha, G., Palmskog, K.: A survey of statistical model checking. ACM Trans. Model. Comput. Simul. 28(1), 6:1–6:39 (2018). https://doi.org/10.1145/3158668
- 2. Amemiya, T.: Advanced econometrics. Harvard university press (1985)
- 3. Anderson, T.W.: Confidence limits for the value of an arbitrary bounded random variable with a continuous distribution function. Bulletin of The International and Statistical Institute 43, 249–251 (1969)
- Baier, C.: Probabilistic model checking. In: Esparza, J., Grumberg, O., Sickert, S. (eds.) Dependable Software Systems Engineering, NATO Science for Peace and Security Series D: Information and Communication Security, vol. 45, pp. 1–23. IOS Press (2016). https://doi.org/10.3233/978-1-61499-627-9-1
- Baier, C., de Alfaro, L., Forejt, V., Kwiatkowska, M.: Model checking probabilistic systems. In: Clarke, E.M., Henzinger, T.A., Veith, H., Bloem, R. (eds.) Handbook of Model Checking, pp. 963–999. Springer (2018). https://doi.org/10. 1007/978-3-319-10575-8_28
- Baier, C., Chatterjee, K., Meggendorfer, T., Piribauer, J.: Entropic risk for turn-based stochastic games. Inf. Comput. 301, 105214 (2024). https://doi.org/10.1016/J.IC.2024.105214, https://doi.org/10.1016/j.ic.2024.105214
- 7. Baier, C., Katoen, J.P.: Principles of model checking. MIT Press (2008), https://mitpress.mit.edu/books/principles-model-checking
- Baier, C., Piribauer, J., Starke, M.: Risk-averse optimization of total rewards in Markovian models using deviation measures. In: CONCUR. LIPIcs, vol. 311, pp. 9:1–9:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2024)
- Bohnenkamp, H.C., D'Argenio, P.R., Hermanns, H., Katoen, J.P.: MoDeST: A compositional modeling formalism for hard and softly timed systems. IEEE Trans. Software Eng. 32(10), 812–830 (2006). https://doi.org/10.1109/TSE.2006.104
- Brázdil, T., Chatterjee, K., Forejt, V., Kucera, A.: Trading performance for stability in Markov decision processes. J. Comput. Syst. Sci. 84, 144–170 (2017)
- 11. Brihaye, T., Chatterjee, K., Mohr, S., Weininger, M.: Risk-aware Markov decision processes using cumulative prospect theory. In: LICS (to appear, preprint available at https://arxiv.org/abs/2505.09514) (2025)
- Budde, C.E., D'Argenio, P.R., Hartmanns, A.: Automated compositional importance splitting. Sci. Comput. Program. 174, 90–108 (2019). https://doi.org/10.1016/J.SCICO.2019.01.006
- Budde, C.E., D'Argenio, P.R., Hartmanns, A., Sedwards, S.: An efficient statistical model checker for nondeterminism and rare events. Int. J. Softw. Tools Technol. Transf. 22(6), 759–780 (2020). https://doi.org/10.1007/S10009-020-00563-2
- Budde, C.E., Dehnert, C., Hahn, E.M., Hartmanns, A., Junges, S., Turrini, A.: JANI: Quantitative model and tool interaction. In: Legay, A., Margaria, T. (eds.)
 23rd International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). Lecture Notes in Computer Science, vol. 10206, pp. 151–168 (2017). https://doi.org/10.1007/978-3-662-54580-5_9
- 15. Budde, C.E., Hartmanns, A., Meggendorfer, T., Weininger, M., Wienhöft, P.: Sound statistical model checking for probabilities and expected rewards. In:

- TACAS (1). Lecture Notes in Computer Science, vol. 15696, pp. 167–190. Springer (2025). https://doi.org/10.1007/978-3-031-90643-5_9, https://doi.org/10.1007/978-3-031-90643-5_9
- Clopper, C., Pearson, E.: The use of confidence or fiducial limits illustrated in the case of the binomial. Biometrika 26(4), 404–413 (1934). https://doi.org/10. 1093/biomet/26.4.404
- D'Argenio, P.R., Fraire, J.A., Hartmanns, A., Raverta, F.D.: Comparing statistical, analytical, and learning-based routing approaches for delay-tolerant networks. ACM Trans. Model. Comput. Simul. 35(2) (2025). https://doi.org/10.1145/3665927
- D'Argenio, P.R., Hartmanns, A., Sedwards, S.: Lightweight statistical model checking in nondeterministic continuous time. In: Margaria, T., Steffen, B. (eds.) 8th International Symposium on Leveraging Applications of Formal Methods, Verification and Validation (ISoLA). Lecture Notes in Computer Science, vol. 11245, pp. 336–353. Springer (2018). https://doi.org/10.1007/978-3-030-03421-4_22
- D'Argenio, P.R., Monti, R.E.: Input/output stochastic automata with urgency: Confluence and weak determinism. In: Fischer, B., Uustalu, T. (eds.) 15th International Colloquium on Theoretical Aspects of Computing (ICTAC). Lecture Notes in Computer Science, vol. 11187, pp. 132–152. Springer (2018). https://doi.org/10.1007/978-3-030-02508-3_8
- David, A., Larsen, K.G., Legay, A., Mikucionis, M., Wang, Z.: Time for statistical model checking of real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) 23rd International Conference on Computer Aided Verification (CAV). Lecture Notes in Computer Science, vol. 6806, pp. 349–355. Springer (2011). https://doi.org/ 10.1007/978-3-642-22110-1_27
- Dvoretzky, A., Kiefer, J., Wolfowitz, J.: Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. The Annals of Mathematical Statistics 27(3), 642–669 (1956). https://doi.org/10.1214/aoms/1177728174
- Ellen, C., Gerwinn, S., Fränzle, M.: Statistical model checking for stochastic hybrid systems involving nondeterminism over continuous domains. Int. J. Softw. Tools Technol. Transf. 17(4), 485–504 (2015). https://doi.org/10.1007/S10009-014-0329-Y
- Föllmer, H., Schied, A.: Convex measures of risk and trading constraints. Finance and stochastics 6, 429–447 (2002)
- Fränzle, M., Hahn, E.M., Hermanns, H., Wolovick, N., Zhang, L.: Measurability and safety verification for stochastic hybrid systems. In: Caccamo, M., Frazzoli, E., Grosu, R. (eds.) 14th ACM International Conference on Hybrid Systems: Computation and Control (HSCC). pp. 43–52. ACM (2011). https://doi.org/10.1145/ 1967701.1967710
- Hahn, E.M., Hartmanns, A., Hermanns, H., Katoen, J.P.: A compositional modelling and analysis framework for stochastic hybrid systems. Formal Methods Syst. Des. 43(2), 191–232 (2013). https://doi.org/10.1007/S10703-012-0167-Z
- 26. Hartmanns, A., Hermanns, H.: The Modest Toolset: An integrated environment for quantitative modelling and verification. In: Ábrahám, E., Havelund, K. (eds.) 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). Lecture Notes in Computer Science, vol. 8413, pp. 593–598. Springer (2014). https://doi.org/10.1007/978-3-642-54862-8_51
- 27. Hartmanns, A., Junges, S., Quatmann, T., Weininger, M.: A practitioner's guide to MDP model checking algorithms. In: Sankaranarayanan, S., Sharygina, N. (eds.)

- 29th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). Lecture Notes in Computer Science, vol. 13993, pp. 469–488. Springer (2023). https://doi.org/10.1007/978-3-031-30823-9_24
- Hartmanns, A., Klauck, M., Parker, D., Quatmann, T., Ruijters, E.: The quantitative verification benchmark set. In: Vojnar, T., Zhang, L. (eds.) 25th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS). Lecture Notes in Computer Science, vol. 11427, pp. 344–350. Springer (2019). https://doi.org/10.1007/978-3-030-17462-0_20
- 29. Hoeffding, W.: Probability inequalities for sums of bounded random variables. Journal of the American Statistical Association 58(301), 13–30 (1963). https://doi.org/10.1080/01621459.1963.10500830
- Krähmann, D., Schubert, J., Baier, C., Dubslaff, C.: Ratio and weight quantiles. In: Italiano, G.F., Pighizzini, G., Sannella, D. (eds.) 40th International Symposium on Mathematical Foundations of Computer Science 2015 (MFCS). Lecture Notes in Computer Science, vol. 9234, pp. 344–356. Springer (2015). https://doi.org/10.1007/978-3-662-48057-1_27
- Kretínský, J.: Survey of statistical verification of linear unbounded properties: Model checking and distances. In: ISoLA (1). Lecture Notes in Computer Science, vol. 9952, pp. 27–45 (2016). https://doi.org/10.1007/978-3-319-47166-2_3
- Kretínský, J., Meggendorfer, T.: Conditional value-at-risk for reachability and mean payoff in markov decision processes. In: Dawar, A., Grädel, E. (eds.) 33rd Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). pp. 609–618. ACM (2018). https://doi.org/10.1145/3209108.3209176
- Legay, A., Lukina, A., Traonouez, L.M., Yang, J., Smolka, S.A., Grosu, R.: Statistical model checking. In: Steffen, B., Woeginger, G.J. (eds.) Computing and Software Science State of the Art and Perspectives, Lecture Notes in Computer Science, vol. 10000, pp. 478–504. Springer (2019). https://doi.org/10.1007/978-3-319-91908-9_23
- Ma, S., Ma, X., Xia, L.: A unified algorithm framework for mean-variance optimization in discounted Markov decision processes. Eur. J. Oper. Res. 311(3), 1057–1067 (2023)
- Mannor, S., Tsitsiklis, J.N.: Mean-variance optimization in Markov decision processes. In: ICML. pp. 177–184. Omnipress (2011)
- Markowitz, H.M.: Foundations of portfolio theory. The journal of finance 46(2), 469–477 (1991)
- Massart, P.: The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. The Annals of Probability 18(3), 1269–1283 (1990). https://doi.org/10.1214/aop/ 1176990746
- Niehage, M., Hartmanns, A., Remke, A.: Learning optimal decisions for stochastic hybrid systems. In: Arun-Kumar, S., Méry, D., Saha, I., Zhang, L. (eds.) 19th ACM-IEEE International Conference on Formal Methods and Models for System Design (MEMOCODE). pp. 44–55. ACM (2021). https://doi.org/10.1145/3487212. 3487339
- 39. Pilch, C., Remke, A.: Statistical model checking for hybrid Petri nets with multiple general transitions. In: 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN). pp. 475–486. IEEE Computer Society (2017). https://doi.org/10.1109/DSN.2017.41
- Piribauer, J., Sankur, O., Baier, C.: The variance-penalized stochastic shortest path problem. In: ICALP. LIPIcs, vol. 229, pp. 129:1–129:19. Schloss Dagstuhl -Leibniz-Zentrum für Informatik (2022)

- 41. Randour, M., Raskin, J.F., Sankur, O.: Percentile queries in multi-dimensional Markov decision processes. Formal Methods Syst. Des. **50**(2-3), 207–248 (2017). https://doi.org/10.1007/S10703-016-0262-7
- Rockafellar, R.T., Uryasev, S.: Conditional value-at-risk for general loss distributions. Journal of banking & finance 26(7), 1443–1471 (2002)
- Tversky, A., Kahneman, D.: Advances in prospect theory: Cumulative representation of uncertainty. Journal of Risk and uncertainty 5(4), 297–323 (1992)
- 44. Ummels, M., Baier, C.: Computing quantiles in markov reward models. In: Pfenning, F. (ed.) Foundations of Software Science and Computation Structures 16th International Conference, FOSSACS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7794, pp. 353–368. Springer (2013). https://doi.org/10.1007/978-3-642-37075-5_23
- 45. Wald, A.: Sequential tests of statistical hypotheses. The Annals of Mathematical Statistics 16(2), 117–186 (1945). https://doi.org/10.1214/aoms/1177731118
- Wirch, J.L., Hardy, M.R.: Distortion risk measures: Coherence and stochastic dominance. In: International congress on insurance: Mathematics and economics. pp. 15–17. Citeseer (2001)
- 47. Younes, H.L.S., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Brinksma, E., Larsen, K.G. (eds.) 14th International Conference on Computer Aided Verification (CAV). Lecture Notes in Computer Science, vol. 2404, pp. 223–235. Springer (2002). https://doi.org/10.1007/3-540-45657-0_17




Fig. 3. Example of a DTMC (left) together with the probability distribution over possible reward outcomes (middle) and the corresponding cdf (right, solid line). The states in the DTMC are (only) labelled by their rewards. The right figure also includes an empirical cdf (dashed) and corresponding confidence band (gray) obtained from the DKW inequality (with $\delta = 0.1$ and k = 50).

A Example: Computing Quantities of Interest

For convenience, we repeat Fig. 1 here in Fig. 3. In this appendix, we show how compute the quantities of interest listed in Sec. 2 for a given cdf. We highlight that while the example uses a discrete random variable, the aggregators are also applicable in the general case; in the same vein, our theory developed in Sec. 3 also applies to general random variables (i.e. discrete, continuous, or mixed). We denote by X the random variable and by dom(X) its domain; in the example it is (a subset of) the natural numbers. For $x \in dom(X)$ we write p(x) for $\mathbb{P}(X = x)$.

Mean Recall that $\mathbb{E}(X) = \int X(\omega) d\mathbb{P}(\omega)$. As our example comprises a discrete random variable, we simply compute the weighted sum $\mathbb{E}(X) = \sum_{x \in dom(X)} x \cdot p(x) = 0.5 \cdot 1 + \sum_{i=1}^{\infty} 2i \cdot (\frac{1}{2})^{i+1} = 2.5$.

Higher moments are computed as the mean of X^n . Essentially, we rescale every outcome by taking its *n*-th power. For example, with n=2 we get $\mathbb{E}(X^2) = \sum_{x \in dom(X)} x^2 \cdot p(x) = 0.5 \cdot 1^2 + \sum_{i=1}^{\infty} (2i)^2 \cdot (\frac{1}{2})^{i+1} = 12.5$.

Every outcome by taking its n-th power. For example, with n=2 we get $\mathbb{E}(X^2) = \sum_{x \in dom(X)} x^2 \cdot p(x) = 0.5 \cdot 1^2 + \sum_{i=1}^{\infty} (2i)^2 \cdot (\frac{1}{2})^{i+1} = 12.5$. Quantiles Recall that $Q_t(X) = \inf\{v \mid \mathbb{P}(X \leq v) \geq t\}$ for $t \in (0,1)$. Intuitively, we take the highest value of the worst $\frac{t}{100}$ % outcomes. We note that quantiles are sometimes also defined by "partitioning" the outcomes into n different blocks and then taking the k-the value, for example "the second 20% quantile", which is equal to t = 0.4 in our definition, i.e. $t = \frac{k}{n}$ in general.

Choosing $t \in (0,0.5]$, we get $Q_t(X) = 1$, since that is the smallest value with a positive probability. Then we observe the non-continuity of quantiles that is mentioned as a complication in Sec. 3.2: Choosing $t \in (0.5,0.75]$, we obtain $Q_t(X) = 2$. Quantiles can easily be read off from the cdf by finding the smallest (leftmost) value x such that the cdf is above t.

Conditional Value-at-Risk Formally, CVaR is defined as: $\text{CVaR}_t(X) = \frac{1}{t}(P \cdot \mathbb{E}[X \mid X < v] + (t - P) \cdot v)$, where $t \in (0, 1)$, $v = Q_t(X)$ and $P = \mathbb{P}(X < v)$. Intuitively, this means taking the expectation over the t-quantile. The reason this definition looks surprisingly complicated is that it has to account for the probability mass exactly at v potentially only being partially included.

In our example, for t = 0.75, we have v = 2 and P = 0.5, yielding $\text{CVaR}_{0.75}(X) = \frac{1}{0.75} \cdot (0.5 \cdot 1 + 0.25 \cdot 2)$. However, choosing t = 0.7, observe that we only

want to consider 0.2 of the outcome 2, even though p(2) = 0.25, reflected by $(t - P) \cdot v = 0.2 \cdot 2$.

Entropic Risk Entropic risk is defined as $\mathrm{ERisk}_{\gamma}(X) = -\frac{1}{\gamma}\log(\mathbb{E}(e^{-\gamma X}))$, with $\gamma > 0$. This function first rescales the obtained reward using the exponential function $e^{-\gamma X}$, takes the expectation of the exponentially rescaled rewards, and then re-normalizes the value by $-\gamma\log(\dots)$. We refer to [6] for a more detailed explanation.

In our example, for $\gamma = 2$, we get

ERisk₁(X) =
$$-\frac{1}{2}\log(\sum_{x \in dom(X)} e^{-2x} \cdot p(x))$$

= $-\frac{1}{2}\log(e^{-2} \cdot 0.5 + \sum_{i=1}^{\infty} e^{-4i} \cdot (\frac{1}{2})^{i+1}) \approx 1.35.$

B Sequential DKW

In the main body, we considered the problem of deriving "as good as possible" bounds given a sample budget (or an already gathered set of samples). However, as mentioned in Remark 1, sometimes we are also interested in the *sequential* setting, where the goal is to gather samples until a certain precision can be guaranteed. As a naive approach, we could try to derive an a-priori upper bound on the number of samples required. For example, [15] notes that such bounds for estimating the mean can be derived by using the fact that DKW reduces to Hoeffding's inequality in the worst case. However, far fewer samples might be sufficient in case the variance of the sample data is low.

Additionally, the "DKW-Lower" procedure introduced in [15] gives a lower bound for the expected value of an unbounded reward, again for a fixed set of samples. It is shown that DKW-Lower "converges" in the sense that the lower bound can be arbitrarily close to the true expected value if a large enough sample set is chosen.

Intuitively, one might think it is possible to derive a simple sequential procedure by repeatedly computing bounds via the DKW inequality until a satisfactory confidence interval is achieved. However, this simple approach to build sequential procedures is no longer sound: The probability that any confidence interval is incorrect accumulates and is no longer guaranteed to be $\leq \delta$. Even if we are only interested in the final confidence interval in the sequence, the guarantees of the DKW inequality no longer apply since the stopping condition is not independent of the outcome of the sampling process, as discussed in [15].

In this section, we now aim to define a sequential procedure that retains soundness while converging to the true expected value. The core idea is to define certain stages at which we build confidence intervals and split the confidence budget δ over all stages in such a way that the probability of all stages yielding a correct confidence interval is $\geq 1 - \delta$ while at the same time ensuring that the sample count per stage grows fast enough to counteract the growth of the confidence intervals due to decreasing δ . Further, we show that this sequential procedure extends to a very general class of objectives beyond expected rewards.

Formally, let \mathcal{F} be an aggregation. (Technically, when we write $\mathcal{F}(F)$ for a cdf F, we refer to computing the aggregation of the associated random variable.) Like in the DKW inequality, for $\varepsilon \geq 0$ we define the confidence band

$$\hat{F}_{\varepsilon} = \left\{ F'(x) \mid \sup_{x \in \mathbb{R}_{\geqslant 0}} |F(x) - F'(x)| \le \varepsilon \right\}.$$

and derive the confidence interval C_{ε} as

$$C_{\varepsilon} = [\underline{C}_{\varepsilon}, \overline{C}_{\varepsilon}] = \left[\inf_{F \in \hat{F}_{\varepsilon}} \mathcal{F}(F), \sup_{F \in \hat{F}_{\varepsilon}} \mathcal{F}(F)\right].$$

Given an infinite stream of random variables $\mathcal{X}=X_1,X_2,\ldots$ and confidence level $1-\delta$, we define, we define the procedure "Sequential-DKW" as follows: We choose an $n\in\mathbb{N}$ and define $n_i=ni^2$ as well as $\delta_i=\frac{\delta}{2^i}$ for all $i\geq 0$. We also define $\Xi_i=\{X_1,\ldots,X_{n_i}\}$ with the corresponding eCDF F_i . Finally, we define the sequence of confidence intervals output by Sequential-DKW as $(C_i)_{i\geq 0}$ where $C_i=C_{\varepsilon_i}$ with $\varepsilon_i=\sqrt{\frac{\ln(2/\delta_i)}{2n_i}}$.

Theorem 1. For any aggregation \mathcal{F} , stream of random variables \mathcal{X} drawn from a distribution with cdf F_X and $\delta \in (0,1)$, the probability that all confidence intervals C_i produced by Sequential-DKW are correct is at least $1 - \delta$, i.e.

$$\mathbb{P}\left(\forall i > 0.\mathcal{F}(X) \in C_i\right) > 1 - \delta.$$

Further, if C_{ε} is always continuous at $\varepsilon = 0$, we have

$$\lim_{i\to\infty} \overline{C}_i = \lim_{i\to\infty} \underline{C}_i = \mathcal{F}(F_X).$$

Proof. By the DKW inequality, he have

$$\mathbb{P}\left(\mathcal{F}(F_X) \notin C_i\right) \leq \mathbb{P}\left(F_X \notin \hat{F}_{\varepsilon_i}\right)$$

$$= \mathbb{P}\left(\sup_{x \in \mathbb{R}_{\geq 0}} |F_X - F_i| \geq \sqrt{\ln(2/\delta_i)/(2n_i)}\right)$$

$$\leq 2e^{-2n_i\sqrt{\ln(2/\delta_i)/(2n_i)}^2}$$

$$= \delta_i$$

By definition of δ_i and the union bound we have

$$\mathbb{P}\left(\forall i \geq 0. \mathcal{F}(F_X) \in C_i\right) \leq \sum_{i=0}^{\infty} \mathbb{P}\left(\mathcal{F}(F_X) \notin C_i\right) \leq \sum_{i=0}^{\infty} \delta_i \leq \sum_{i=0}^{\infty} \frac{\delta}{2^i} = \delta$$

To show the second part of the theorem, note that for all $i \geq 0$, by definition of C_{ε_i} , there is a sequence $(\underline{F}_i^j)_{j\geq 0}$ for which all $\underline{F}_i^j \in \hat{F}_{\varepsilon_i}$ and $\lim_{j\to\infty} \mathcal{F}(\underline{F}_i^j) = \underline{C}_i$ and and analogous sequence $(\underline{F}_i^j)_{j\geq 0}$. Then, since $\underline{F}_i^j, \overline{F}_i^j \in \hat{F}_{\varepsilon_i}$ for all $j \geq 0$ and by definition of n_i we have

$$\begin{split} \sup_{x \in \mathbb{R}} \left(\lim_{j \to \infty} \overline{F}_i^j(x) - \lim_{j \to \infty} \underline{F}_i^j(x) \right) &\leq 2 \sqrt{\ln(2/\delta_i)/(2n_i)} \\ &= 2 \sqrt{\ln(2^{i+1}/\delta)/(2ni^2)} = 2 \sqrt{(i+1)\ln(2/\delta)/(2ni^2)} \end{split}$$

Thus, since $\lim_{i\to\infty} \varepsilon_i = 0$ and C_{ε} is continuous at $\varepsilon = 0$, we have

$$\lim_{i \to \infty} \overline{C}_i - \underline{C}_i = \lim_{i \to \infty} \left(\lim_{j \to \infty} \mathcal{F}(\overline{F}_i^j) - \lim_{j \to \infty} \mathcal{F}(\underline{F}_i^j) \right) = 0,$$

i.e. $\lim_{i\to\infty} \overline{C}_i = \lim_{i\to\infty} \underline{C}_i$. Finally, by the law of large numbers $\lim_{i\to\infty} F_i = F_X$ and thus again by $\lim_{i\to\infty} \varepsilon_i = 0$ and continuity of C_{ε} at $\varepsilon = 0$ we have $\lim_{i\to\infty} \overline{C}_i = \lim_{i\to\infty} \underline{C}_i = \mathcal{F}(F_X)$.

Remark 2. Sequential-DKW naturally extends to bounded random variables (e.g. positive expected reward with an a priori upper bound) where $a \le X \le b$ for all $X \in \mathcal{X}$ by additionally requiring F(a) = 0 and F(b) = 1 for all $F \in \hat{F}_{\varepsilon_i}$ for all $i \ge 0$.

Remark 3. If only \underline{C}_i is continuous at $\varepsilon=0$ we can still obtain convergence for the lower bound, i.e. $\lim_{i\to\infty}\underline{C}_i=\mathcal{F}(F_X)$, in the same way as for Theorem 1. The analogous statement holds for \overline{C}_i . This is for example relevant for unbounded expected rewards (where \overline{C}_ε is always finite for $\varepsilon=0$ but infinite for $\varepsilon>0$). There, $(\underline{C}_i)_{i\geq 0}$ as in Sequential-DKW is a sound sequence of lower bounds converging towards the true expected reward, but Sequential-DKW likely does not yield converging upper bounds since $\overline{C}_i=\infty$ for all $i\geq 0$. Similarly, for VaR \overline{C}_ε may not always be continuous at $\varepsilon=0$ for discrete distributions but \underline{C}_i is.