Monitoring Algorithmic Fairness[*]

Thomas A. Henzinger[0000−0002−2985−7724],
Mahyar Karimi[0009−0005−0820−1696],
Konstantin Kueffner[0000−0001−8974−2542], and
Kaushik Mallik[0000−0001−9864−7475]

Institute of Science and Technology Austria (ISTA), Klosterneuburg, Austria.
{tah, mahyar.karimi, konstantin.kueffner, kaushik.mallik}@ist.ac.at

Abstract. Machine-learned systems are in widespread use for making decisions about humans, and it is important that they are *fair*, i.e., not biased against individuals based on sensitive attributes. We present runtime verification of algorithmic fairness for systems whose models are unknown, but are assumed to have a Markov chain structure. We introduce a specification language that can model many common algorithmic fairness properties, such as demographic parity, equal opportunity, and social burden. We build monitors that observe a long sequence of events as generated by a given system, and output, after each observation, a quantitative estimate of how fair or biased the system was on that run until that point in time. The estimate is proven to be correct modulo a variable error bound and a given confidence level, where the error bound gets tighter as the observed sequence gets longer. Our monitors are of two types, and use, respectively, frequentist and Bayesian statistical inference techniques. While the frequentist monitors compute estimates that are objectively correct with respect to the ground truth, the Bayesian monitors compute estimates that are correct subject to a given prior belief about the system's model. Using a prototype implementation, we show how we can monitor if a bank is fair in giving loans to applicants from different social backgrounds, and if a college is fair in admitting students while maintaining a reasonable financial burden on the society. Although they exhibit different theoretical complexities in certain cases, in our experiments, both frequentist and Bayesian monitors took less than a millisecond to update their verdicts after each observation.

1 Introduction

Runtime verification complements traditional static verification techniques, by offering lightweight solutions for checking properties based on a single, possibly long execution trace of a given system [8]. We present new runtime verification techniques for the problem of bias detection in decision-making software. The use of software for making critical decisions about humans is a growing trend;

---

example areas include judiciary [13,20], policing [23,48], banking [47], etc. It is important that these software systems are unbiased towards the protected attributes of humans, like gender, ethnicity, etc. However, they have often shown biases in their decisions in the past [20,46,54,56,57]. While there are many approaches for mitigating biases before deployment [20,46,54,56,57], recent runtime verification approaches [3,33] offer a new complementary tool to oversee *algorithmic fairness* in AI and machine-learned decision makers during deployment.

To verify algorithmic fairness at runtime, the given decision-maker is treated as a *generator* of events with an unknown model. The goal is to algorithmically design lightweight but rigorous *runtime monitors* against quantitative formal specifications. The monitors observe a long stream of events and, after each observation, output a quantitative, statistically sound estimate of how fair or biased the generator was until that point in time. While the existing approaches [3,33] considered only sequential decision making models and built monitors from the frequentist viewpoint in statistics, we allow the richer class of Markov chain models and present monitors from both the frequentist and the Bayesian statistical viewpoints.

Monitoring algorithmic fairness involves on-the-fly statistical estimations, a feature that has not been well-explored in the traditional runtime verification literature. As far as the algorithmic fairness literature is concerned, the existing works are mostly *model-based*, and either minimize decision biases of machine-learned systems at *design-time* (i.e., pre-processing) [40,65,11,64], or verify their absence at *inspection-time* (i.e., post-processing) [32]. In contrast, we verify algorithmic fairness at *runtime*, and do not require an explicit model of the generator. On one hand, the model-independence makes the monitors trustworthy, and on the other hand, it complements the existing model-based static analyses and design techniques, which are often insufficient due to partially unknown or imprecise models of systems in real-world environments.

We assume that the sequences of events generated by the generator can be modeled as sequences of states visited by a finite unknown Markov chain. This implies that the generator is well-behaved and the events follow each other according to some fixed probability distributions. Not only is this assumption satisfied by many machine-learned systems (see Sec. 1.1 for examples), it also provides just enough structure to lay the bare-bones foundations for runtime verification of algorithmic fairness properties. We emphasize that we do not require knowledge of the transition probabilities of the underlying Markov chain.

We propose a new specification language, called the Probabilistic Specification Expressions (PSEs), which can formalize a majority of the existing algorithmic fairness properties in the literature, including demographic parity [21], equal opportunity [32], disparate impact [25], etc. Let Q be the set of events. Syntactically, a PSE is a restricted arithmetic expression over the (unknown) transition probabilities of a Markov chain with the state space Q. Semantically, a PSE φ over Q is a function that maps every Markov chain M with the state space Q to a real number, and the value $\varphi(M)$ represents the degree of fairness or bias (with respect to φ) in the generator M. Our monitors observe a long

sequence of events from Q, and after each observation, compute a statistically rigorous estimate of $\varphi(M)$ with a PAC-style error bound for a given confidence level. As the observed sequence gets longer, the error bound gets tighter.

Algorithmic fairness properties that are expressible using PSEs are quantitative refinements of the traditional qualitative fairness properties studied in formal methods. For example, a qualitative fairness property may require that if a certain event A occurs infinitely often, then another event B should follow infinitely often. In particular, a coin is qualitatively fair if infinitely many coin tosses contain both infinitely many heads and infinitely many tails. In contrast, the coin will be algorithmically fair (i.e., unbiased) if approximately half of the tosses come up heads. Technically, while qualitative weak and strong fairness properties are ω-regular, the algorithmic fairness properties are statistical and require counting. Moreover, for a qualitative fairness property, the satisfaction or violation cannot be established based on a finite prefix of the observed sequence. In contrast, for any given finite prefix of observations, the value of an algorithmic fairness property can be estimated using statistical techniques, assuming the future behaves statistically like the past (the Markov assumption).

As our main contribution, we present two different monitoring algorithms, using tools from frequentist and Bayesian statistics, respectively. The central idea of the *frequentist monitor* is that the probability of every transition of the monitored Markov chain M can be estimated using the fraction of times the transition is taken per visit to its source vertex. Building on this, we present a practical implementation of the frequentist monitor that can estimate the value of a given PSE from an observed finite sequence of states. For the coin example, after every new toss, the frequentist monitor will update its estimate of probability of seeing heads by computing the fraction of times the coin came up heads so far, and then by using concentration bounds to find a tight error bound for a given confidence level. On the other hand, the central idea of the *Bayesian monitor* is that we begin with a prior belief about the transition probabilities of M, and having seen a finite sequence of observations, we can obtain an updated posterior belief about M. For a given confidence level, the output of the monitor is computed by applying concentration inequalities to find a tight error bound around the mean of the posterior belief. For the coin example, the Bayesian monitor will begin with a prior belief about the degree of fairness, and, after observing the outcome of each new toss, will compute a new posterior belief. If the prior belief agrees with the true model with a high probability, then the Bayesian monitor's output converges to the true value of the PSE more quickly than the frequentist monitor. In general, both monitors can efficiently estimate more complicated PSEs, such as the ratio and the squared difference of the probabilities of heads of two different coins. The choice of the monitor for a particular application depends on whether an objective or a subjective evaluation, with respect to a given prior, is desired.

Both frequentist and Bayesian monitors use registers (and counters as a restricted class of registers) to keep counts of the relevant events and store the intermediate results. If the size of the given PSE is n, then, in theory, the fre-

quentist monitor uses $\mathcal{O}(n^4 2^n)$ registers and computes its output in $\mathcal{O}(n^4 2^n)$ time after each new observation, whereas the Bayesian monitor uses $\mathcal{O}(n^2 2^n)$ registers and computes its output in $\mathcal{O}(n^2 2^n)$ time after each new observation. The computation time and the required number of registers get drastically reduced to $O(n^2)$ for the frequentist monitor with PSEs that contain up to one division operator, and for the Bayesian monitor with polynomial PSEs (possibly having negative exponents in the monomials). This shows that under given circumstances, one or the other type of the monitor can be favorable computation-wise. These special, efficient cases cover many algorithmic fairness properties of interest, such as demographic parity and equal opportunity.

Our experiments confirm that our monitors are fast in practice. Using a prototype implementation in Rust, we monitored a couple of decision-making systems adapted from the literature. In particular, we monitor if a bank is fair in lending money to applicants from different demographic groups [47], and if a college is fair in admitting students without creating an unreasonable financial burden on the society [53]. In our experiments, both monitors took, on an average, less than a millisecond to update their verdicts after each observation, and only used tens of internal registers to operate, thereby demonstrating their practical usability at runtime.

In short, we advocate that runtime verification introduces a new set of tools in the area of algorithmic fairness, using which we can monitor biases of deployed AI and machine-learned systems in real-time. While existing monitoring approaches only support sequential decision making problems and use only the frequentist statistical viewpoint, we present monitors for the more general class of Markov chain system models using both frequentist and Bayesian statistical viewpoints.

All proofs can be found in the appendix.

1.1   Motivating Examples

We first present two real-world examples from the algorithmic fairness literature to motivate the problem; these examples will later be used to illustrate the technical developments.

The lending problem [47]: Suppose a bank lends money to individuals based on certain attributes, like credit score, age group, etc. The bank wants to maximize profit by lending money to only those who will repay the loan in time—called the "true individuals." There is a sensitive attribute (e.g., ethnicity) classifying the population into two groups g and $\overline{g}$. The bank will be considered fair (in lending money) if its lending policy is independent of an individual's membership in g or $\overline{g}$. Several *group fairness* metrics from the literature are relevant in this context. *Disparate impact* [25] quantifies the *ratio* of the probability of an individual from g getting the loan to the probability of an individual from $\overline{g}$ getting the loan, which should be close to 1 for the bank to be considered fair. *Demographic parity* [21] quantifies the *difference* between the probability of an individual from g getting the loan and the probability of an individual from $\overline{g}$ getting the loan, which should be close to 0 for the bank to be considered fair.

Equal opportunity [32] quantifies the *difference* between the probability of a *true* individual from g getting the loan and the probability of a *true* individual from $\bar{g}$ getting the loan, which should be close to 0 for the bank to be considered fair. A discussion on the relative merit of various different algorithmic fairness notions is out of scope of this paper, but can be found in the literature [61,42,15,22]. We show how we can monitor whether a given group fairness criteria is fulfilled by the bank, by observing a sequence of lending decisions.

The college admission problem [53]: Consider a college that announces a cutoff of grades for admitting students through an entrance examination. Based on the merit, every truly qualified student belongs to group g, and the rest to group $\bar{g}$. Knowing the cutoff, every student can choose to invest a sum of money—proportional to the gap between the cutoff and their true merit—to be able to reach the cutoff, e.g., by taking private tuition classes. On the other hand, the college's utility is in minimizing admission of students from $\bar{g}$, which can be accomplished by raising the cutoff to a level that is too expensive to be achieved by the students from $\bar{g}$ and yet easy to be achieved by the students from g. The *social burden* associated to the college's cutoff choice is the expected expense of every student from g, which should be close to 0 for the college to be considered fair (towards the society). We show how we can monitor the social burden, by observing a sequence of investment decisions made by the students from g.

1.2 Related Work

There has been a plethora of work on algorithmic fairness from the machine learning standpoint [51,21,32,45,41,58,10,62,12,65,37,44]. In general, these works improve algorithmic fairness through de-biasing the training dataset (pre-processing), or through incentivizing the learning algorithm to make fair decisions (in-processing), or through eliminating biases from the output of the machine-learned model (post-processing). All of these are interventions in the design of the system, whereas our monitors treat the system as already deployed.

Recently, formal methods-inspired techniques have been used to guarantee algorithmic fairness through the verification of a learned model [2,9,60,29,52], and enforcement of robustness [38,6,30]. All of these works verify or enforce algorithmic fairness *statically* on all runs of the system with high probability. This requires certain knowledge about the system model, which may not be always available. Our runtime monitor dynamically verifies whether the current run of an opaque system is fair.

Our frequentist monitor is closely related to the novel work of Albarghouthi et al. [3], where the authors build a programming framework that allows runtime monitoring of algorithmic fairness properties on programs. Their monitor evaluates the algorithmic fairness of repeated "single-shot" decisions made by machine-learned functions on a sequence of samples drawn from an underlying unknown but fixed distribution, which is a special case of our more general Markov chain model of the generator. They do not consider the Bayesian point of view. Moreover, we argue and empirically show in Sec. 4 that our frequentist

approach produces significantly tighter statistical estimates than their approach on most PSEs. On the flip side, their specification language is more expressive, in that they allow atomic variables for expected values of events, which is useful for specifying individual fairness criteria [21]. We only consider group fairness, and leave individual fairness as part of future research. Also, they allow logical operators (like boolean connectives) in their specification language. However, we obtain tighter statistical estimates for the core arithmetic part of algorithmic fairness properties (through PSEs), and point out that we can deal with logical operators just like they do in a straightforward manner.

Shortly after the first manuscript of this paper was written, we published a separate work for monitoring long-run fairness in sequential decision making problems, where the feature distribution of the population may dynamically change due to the actions of the individuals [33]. Although this other work generalizes our current paper in some aspects (support for dynamic changes in the model), it only allows sequential decision making models (instead of Markov chains) and does not consider the Bayesian monitoring perspective.

There is a large body of research on monitoring, though the considered properties are mainly temporal [59,39,24,49,19,7,5]. Unfortunately, these techniques do not directly extend to monitoring algorithmic fairness, since checking algorithmic fairness requires statistical methods, which is beyond the limit of finite automata-based monitors used by the classical techniques. Although there are works on quantitative monitoring that use richer types of monitors (with counters/registers like us) [28,34,55,35], the considered specifications do not easily extend to statistical properties like algorithmic fairness. One exception is the work by Ferrère et al. [26], which monitors certain statistical properties, like mode and median of a given sequence of events. Firstly, they do not consider algorithmic fairness properties. Secondly, their monitors' outputs are correct only as the length of the observed sequence approaches infinity (asymptotic guarantee), whereas our monitors' outputs are *always* correct with high confidence (finite-sample guarantee), and the precision gets better for longer sequences.

Although our work uses similar tools as used in statistical verification [4,63,14,17,1], the goals are different. In traditional statistical verification, the system's runs are chosen probabilistically, and it is verified if any run of the system satisfies a boolean property with a certain probability. For us, the run is given as input to the monitor, and it is this run that is verified against a quantitative algorithmic fairness property with statistical error bounds. To the best of our knowledge, existing works on statistical verification do not consider algorithmic fairness properties.

2   Preliminaries

For any alphabet Σ, the notation Σ^* represents the set of all finite words over Σ. We write $\mathbb{R}$, $\mathbb{N}$, and $\mathbb{N}^+$ to denote the sets of real numbers, natural numbers (including zero), and positive integers, respectively. For a pair of real (natural) numbers a, b with $a < b$, we write $[a, b]$ ($[a \mathbin{.\,.} b]$) to denote the set of all real

(natural) numbers between and including a and b. For a given $c, r \in \mathbb{R}$, we write $[c \pm r]$ to denote the set $[c - r, c + r]$. For simpler notation, we will use $|\cdot|$ to denote both the cardinality of a set and the absolute value of a real number, whenever the intended use is clear.

For a given vector $v \in \mathbb{R}^n$ and a given $m \times n$ real matrix M, for some m, n, we write v_i to denote the i-th element of v and write M_{ij} to denote the element at the i-th row and the j-th column of M. For a given $n \in \mathbb{N}^+$, a *simplex* is the set of vectors $\Delta(n) := \{x \in [0, 1]^{n+1} \mid \sum_{i=1}^{n+1} x_i = 1\}$. Notice that the dimension of $\Delta(n)$ is $n + 1$ (and not n), a convention that is standard due to the interpretation of $\Delta(n)$ as the $n + 1$ vertices of an n-dimensional polytope. A *stochastic matrix* of dimension $m \times m$ is a matrix whose every row is in $\Delta(m - 1)$, i.e. $M \in \Delta(m - 1)^m$. Random variables (see App. A.1) will be denoted using uppercase symbols from the Latin alphabet (e.g. X), while the associated outcomes will be denoted using lowercase font of the same symbol (x is an outcome of X). We will interchangeably use the expected value $\mathbb{E}(X)$ and the mean μ_X of X. For a given set S, define $\mathcal{D}(S)$ as the set of every random variable—called a *probability distribution*[1]—with set of outcomes being 2^S. A Bernoulli random variable that produces "1" (the alternative is "0") with probability p is written as $Bernoulli(p)$.

2.1   Markov chains as randomized generators of events

We use finite Markov chains as sequential randomized generators of events. A (finite) Markov chain $\mathcal{M}$ is a triple (Q, M, π), where $Q = [1 \mathinner{.\,.} N]$ is a set of states for a finite N, $M \in \Delta(N-1)^N$ is a stochastic matrix called the transition probability matrix, and $\pi \in \mathcal{D}(Q)$ is the distribution over initial states. We often refer to a pair of states $(i, j) \in Q \times Q$ as an *edge*. The Markov chain $\mathcal{M}$ generates an infinite sequence of random variables $X_0 = \pi, X_1, \ldots$, with $X_i \in \mathcal{D}(Q)$ for every i, such that the Markov property is satisfied: $\mathbb{P}(X_{n+1} = i_{n+1} \mid X_0 = i_0, \ldots, X_n = i_n) = \mathbb{P}(X_{n+1} = i_{n+1} \mid X_n = i_n)$, which is $M_{i_n i_{n+1}}$ in our case. A finite *path* $\vec{x} = x_0, \ldots, x_n$ of $\mathcal{M}$ is a finite word over Q such that for every $t \in [0; n]$, $\mathbb{P}(X_t = x_t) > 0$. Let $Paths(\mathcal{M})$ be the set of every finite path of $\mathcal{M}$.

We use Markov chains to model the probabilistic interaction between a machine-learned decision maker with its environment. Intuitively, the Markov assumption on the model puts the restriction that the decision maker does not change over time, e.g., due to retraining.

In Fig. 1 we show the Markov chains for the lending and the college admission examples from Sec. 1.1. The Markov chain for the lending example captures the sequence of loan-related probabilistic events, namely, that a loan applicant is randomly sampled and the group information (g or $\bar{g}$) is revealed, a probabilistic decision is made by the decision-maker and either the loan was granted (gy or $\bar{g}y$, depending on the group) or refused ($\bar{y}$), and if the loan is granted then with some probabilities it either gets repaid (z) or defaulted ($\bar{z}$). The Markov chain

---

[1] An alternate commonly used definition of probability distribution is directly in terms of the probability measure induced over S, instead of through the random variable.

for the college admission example captures the sequence of admission events, namely, that a candidate is randomly sampled and the group is revealed $(g, \overline{g})$, and when the candidate is from group g (truly qualified) then the amount of money invested for admission is also revealed.
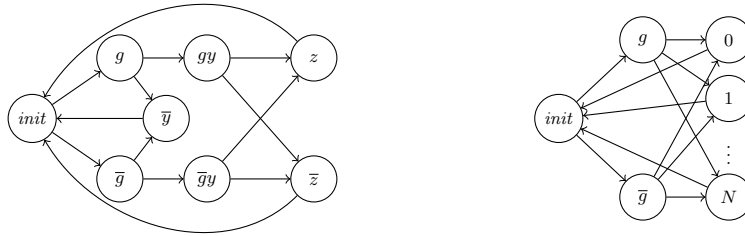

Fig. 1: Markov chains for the lending and the college-admission examples. (**left**) The lending example: The state *init* denotes the initiation of the sampling, and the rest represent the selected individual, namely, g and $\overline{g}$ denote the two groups, (gy) and $(\overline{g}y)$ denote that the individual is respectively from group g and group $\overline{g}$ and the loan was granted, $\overline{y}$ denotes that the loan was refused, and z and $\overline{z}$ denote whether the loan was repaid or not. (**right**) The college admission example: The state *init* denotes the initiation of the sampling, the states $g, \overline{g}$ represent the group identity of the selected candidate, and the states $\{0, \ldots, N\}$ represent the amount of money invested by a truly eligible candidate.

2.2   Randomized register monitors

Randomized register monitors, or simply monitors, are adapted from the (deterministic) polynomial monitors of Ferrère et al. [27]. Let R be a finite set of integer variables called registers. A function $v \colon R \to \mathbb{N}$ assigning concrete value to every register in R is called a valuation of R. Let $\mathbb{N}^R$ denote the set of all valuations of R. Registers can be read and written according to relations in the signature $S = \langle 0, 1, +, -, \times, \div, \leq \rangle$. We consider two basic operations on registers:

- A *test* is a conjunction of atomic formulas over S and their negation;
- An *update* is a mapping from variables to terms over S.

We use $\Phi(R)$ and $\Gamma(R)$ to respectively denote the set of tests and updates over R. *Counters* are special registers with a restricted signature $S = \langle 0, 1, +, -, \leq \rangle$.

Definition 1 (Randomized register monitor). *A randomized register monitor is a tuple $(\Sigma, \Lambda, R, \lambda, T)$ where Σ is a finite input alphabet, Λ is an output alphabet, R is a finite set of registers, $\lambda \colon \mathbb{N}^R \to \Lambda$ is an output function, and $T \colon \Sigma \times \Phi(R) \to \mathcal{D}(\Gamma(R))$ is the randomized transition function such that for every $\sigma \in \Sigma$ and for every valuation $v \in \mathbb{N}^R$, there exists a unique $\phi \in \Phi(R)$ with $v \models \phi$ and $T(\sigma, \phi) \in \mathcal{D}(\Gamma(R))$. A deterministic register monitor is a randomized register monitor for which $T(\sigma, \phi)$ is a Dirac delta distribution, if it is defined.*

A *state* of a monitor $\mathcal{A}$ is a valuation of its registers $v \in \mathbb{N}^R$. The monitor $\mathcal{A}$ *transitions* from state v to a *distribution* over states given by the random variable $Y = T(\sigma, \phi)$ on input $\sigma \in \Sigma$ if there exists ϕ such that $v \models \phi$. Let γ be an outcome of Y with $\mathbb{P}(Y = \gamma) > 0$, in which case the registers are updated as $v'(x) = v(\gamma(x))$ for every $x \in R$, and the respective concrete transition is written as $v \xrightarrow{\sigma} v'$. A *run* of $\mathcal{A}$ on a word $w_0 \ldots w_n \in \Sigma^*$ is a sequence of concrete transitions $v_0 \xrightarrow{w_0} v_1 \xrightarrow{w_1} \ldots \xrightarrow{w_n} v_{n+1}$. The probabilistic transitions of $\mathcal{A}$ induce a probability distribution over the sample space of finite runs of the monitor, denoted $\widehat{\mathbb{P}}(\cdot)$. For a given finite word $w \in \Sigma^*$, the *semantics* of the monitor $\mathcal{A}$ is given by a random variable $[\![\mathcal{A}]\!](w) := \lambda(Y)$ inducing the probability measure $\mathbb{P}_{\mathcal{A}}$, where Y is the random variable representing the distribution over the final state in a run of $\mathcal{A}$ on the word w, i.e., $\mathbb{P}_{\mathcal{A}}(Y = v) := \widehat{\mathbb{P}}(\{r = r_0 \ldots r_m \in \Sigma^* \mid r \text{ is a run of } \mathcal{A} \text{ on } w \text{ and } r_m = v\})$.

Example: A monitor for detecting the (unknown) bias of a coin. We present a simple deterministic monitor that computes a PAC estimate of the bias of an unknown coin from a sequence of toss outcomes, where the outcomes are denoted as "h" for heads and "t" for tails. The input alphabet is the set of toss outcomes, i.e., $\Sigma = \{h, t\}$, the output alphabet is the set of every bias intervals, i.e., $\Gamma = \{[a, b] \mid 0 \le a < b \le 1\}$, the set of registers is $R = \{r_n, r_h\}$, where r_n and r_h are counters counting the total number of tosses and the number of heads, respectively, and the output function λ maps every valuation of r_n, r_h to an interval estimate of the bias that has the form $\lambda \equiv v(r_h)/v(r_n) \pm \varepsilon(r_n, \delta)$, where $\delta \in [0, 1]$ is a given upper bound on the probability of an incorrect estimate and $\varepsilon(r_n, \delta)$ is the estimation error computed using PAC analysis. For instance, after observing a sequence of 67 tosses with 36 heads, the values of the registers will be $v(r_n) = 67$ and $v(r_h) = 36$, and the output of the monitor will be $\lambda(67, 36) = 36/67 \pm \varepsilon(n, \delta)$ for some appropriate $\varepsilon(\cdot)$. Now, suppose the next input to the monitor is h, in which case the monitor's transition is given as $T(h, \cdot) = (r_n + 1, r_h + 1)$, which updates the registers to the new values $v'(r_n) = 67 + 1 = 68$ and $v'(r_h) = 36 + 1 = 37$. For this example, the tests $\Phi(R)$ over the registers are redundant, but they can be used to construct monitors for more complex properties.

3   Algorithmic Fairness Specifications and Problem Formulation

3.1   Probabilistic Specification Expressions

To formalize algorithmic fairness properties, like the ones in Sec. 1.1, we introduce *probabilistic specification expressions* (PSE). A PSE φ over a given finite set Q is an algebraic expression with some restricted set of operations that uses variables labeled v_{ij} with $i, j \in Q$ and whose domains are the real interval $[0, 1]$. The syntax of φ is:

$$\xi ::= v \in \{v_{ij}\}_{i,j \in Q} \mid \xi \cdot \xi \mid 1 \div \xi, \tag{1a}$$

$$\varphi ::= \kappa \in \mathbb{R} \mid \xi \mid \varphi + \varphi \mid \varphi - \varphi \mid \varphi \cdot \varphi \mid (\varphi), \tag{1b}$$

where $\{v_{ij}\}_{i,j \in Q}$ are the variables with domain $[0,1]$ and κ is a constant. The expression ξ in (1a) is called a *monomial* and is simply a product of powers of variables with integer exponents. A *polynomial* is a weighted sum of monomials with constant weights.[2] Syntactically, polynomials form a strict subclass of the expressions definable using (1b), because the product of two polynomials is not a polynomial, but is a valid expression according to (1b). A PSE φ is *division-free* if there is no division operator involved in φ. The *size* of an expression φ is the total number of arithmatic operators (i.e. $+, -, \cdot, \div$) in φ. We use V_φ to denote the set of variables appearing in the expression φ, and for every $V \subseteq V_\varphi$ we define $Dom(V) := \{i \in Q \mid \exists v_{ij} \in V \vee \exists v_{ki} \in V\}$ as the set containing any state of the Markov chain that is involved in some variable in V.

The semantics of a PSE φ is interpreted *statically* on the unknown Markov chain M: we write $\varphi(M)$ to denote the evaluation or the value of φ by substituting every variable v_{ij} in φ with M_{ij}. E.g., for a Markov chain with state space $\{1,2\}$ and transition probabilities $M_{11} = 0.2$, $M_{12} = 0.8$, $M_{21} = 0.4$, and $M_{22} = 0.6$, the expression $\varphi = v_{11} - v_{21}$ has the evaluation $\varphi(M) = 0.2 - 0.4 = -0.2$. We will assume that for every expression $(1 \div \xi)$, $\xi(M) \neq 0$.

Example: Group fairness. Using PSEs, we can express the group fairness properties for the lending example described in Sec. 1.1, with the help of the Markov chain in the left subfigure of Fig. 1:

> **Disparate impact** [25]: $\qquad v_{gy} \div v_{\bar{g}y}$

> **Demographic parity** [21]: $\qquad v_{gy} - v_{\bar{g}y}$

The equal opportunity criterion requires the following probability to be close to zero: $p = \mathbb{P}(y \mid g, z) - \mathbb{P}(y \mid \bar{g}, z)$, which is tricky to monitor as p contains the counter-factual probabilities representing "the probability that an individual from a group would repay had the loan been granted." We apply Bayes' rule, and turn p into the following equivalent form: $p' = \frac{\mathbb{P}(z|g,y) \cdot \mathbb{P}(y|g)}{\mathbb{P}(z|g)} - \frac{\mathbb{P}(z|\bar{g},y) \cdot \mathbb{P}(y|\bar{g})}{\mathbb{P}(z|\bar{g})}$. Assuming $\mathbb{P}(z \mid g) = c_1$ and $\mathbb{P}(z \mid \bar{g}) = c_2$, where c_1 and c_2 are known constants, the property p' can be encoded as a PSE as below:

> **Equal opportunity** [32]: $\qquad (v_{(gy)z} \cdot v_{gy}) \div c_1 - (v_{(\bar{g}y)z} \cdot v_{\bar{g}y}) \div c_2.$

Example: Social burden. Using PSEs, we can express the social burden of the college admission example described in Sec. 1.1, with the help of the Markov chain depicted in the right subfigure of Fig. 1:

> **Social burden** [53]: $\qquad 1 \cdot v_{g1} + \ldots + N \cdot v_{gN}.$

3.2   The Monitoring Problem

Informally, our goal is to build monitors that observe a single long path of a Markov chain and, after each observation, output a new estimate for the value

---

[2] Although monomials and polynomials usually only have positive exponents, we take the liberty to use the terminologies even when negative exponents are present.

of the PSE. Since the monitor's estimate is based on statistics collected from a finite path, the output may be incorrect with some probability, where the source of this probability is different between the frequentist and the Bayesian approaches. In the frequentist approach, the underlying Markov chain is fixed (but unknown), and the randomness stems from the sampling of the observed path. In the Bayesian approach, the observed path is fixed, and the randomness stems from the uncertainty about a prior specifying the Markov chain's parameters. The commonality is that, in both cases, we want our monitors to estimate the value of the PSE up to an error with a fixed probabilistic confidence.

We formalize the monitoring problem separately for the two approaches. A *problem instance* is a triple (Q, φ, δ), where $Q = [1 \mathinner{.\,.} N]$ is a set of states, φ is a PSE over Q, and $\delta \in [0, 1]$ is a constant. In the frequentist approach, we use $\mathbb{P}_s$ to denote the probability measure induced by *sampling* of paths, and in the Bayesian approach we use $\mathbb{P}_\theta$ to denote the probability measure induced by the *prior* probability density function $p_\theta \colon \Delta(n-1)^n \to \mathbb{R} \cup \{\infty\}$ over the transition matrix of the Markov chain. In both cases, the output alphabets of the monitors contain every real interval.

Problem 1 (Frequentist monitor) *Suppose (Q, φ, δ) is a problem instance given as input. Design a monitor $\mathcal{A}$ such that for every Markov chain $\mathcal{M}$ with transition probability matrix M and for every finite path $\vec{x} \in Paths(\mathcal{M})$:*

$$\mathbb{P}_{s, \mathcal{A}} \left(\varphi(M) \in [\![\mathcal{A}]\!](\vec{x}) \right) \geq 1 - \delta, \tag{2}$$

where $\mathbb{P}_{s, \mathcal{A}}$ is the joint probability measure of $\mathbb{P}_s$ and $\mathbb{P}_{\mathcal{A}}$.

Problem 2 (Bayesian monitor) *Suppose (Q, φ, δ) is a problem instance and p_θ is a prior density function, both given as inputs. Design a monitor $\mathcal{A}$ such that for every Markov chain $\mathcal{M}$ with transition probability matrix M and for every finite path $\vec{x} \in Paths(\mathcal{M})$:*

$$\mathbb{P}_{\theta, \mathcal{A}} \left(\varphi(M) \in [\![\mathcal{A}]\!](\vec{x}) \mid \vec{x} \right) \geq 1 - \delta, \tag{3}$$

where $\mathbb{P}_{\theta, \mathcal{A}}$ is the joint probability measure of $\mathbb{P}_\theta$ and $\mathbb{P}_{\mathcal{A}}$.

Notice that the state space of the Markov chain and the input alphabet of the monitor are the same, and so, many times, we refer to observed states as (input) symbols, and vice versa. The estimate $[l, u] = [\![\mathcal{A}]\!](\vec{x})$ is called the $(1 - \delta) \cdot 100\%$ *confidence interval* for $\varphi(M)$.[3] The radius, given by $\varepsilon = 0.5 \cdot (u - l)$, is called the *estimation error*, and the quantity $1 - \delta$ is called the *confidence*. The estimate gets more precise as the error gets smaller and the confidence gets higher.

In many situations, we are interested in a *qualitative* question of the form "is $\varphi(M) \leq c$?" for some constant c. We point out that, once the quantitative problem is solved, the qualitative questions can be answered using standard procedures by setting up a hypothesis test [43, p. 380].

---

[3] While in the Bayesian setting *credible intervals* would be more appropriate, we use confidence intervals due to uniformity and the relative ease of computation. To relate the two, our confidence intervals are over-approximations of credible intervals (non-unique) that are centered around the posterior mean.

4   Frequentist Monitoring

Suppose the given PSE is only a single variable $\varphi = v_{ij}$, i.e., we are monitoring the probability of going from state i to another state j. The frequentist monitor $\mathcal{A}$ for φ can be constructed in two steps: (1) empirically compute the average number of times the edge (i, j) was taken per visit to the state i on the observed path of the Markov chain, and (2) compute the $(1 - \delta) \cdot 100\%$ confidence interval using statistical concentration inequalities.

Now consider a slightly more complex PSE $\varphi' = v_{ij} + v_{ik}$. One approach to monitor φ', proposed by Albarghouthi et al. [3], would be to first compute the $(1 - \delta) \cdot 100\%$ confidence intervals $[l_1, u_1]$ and $[l_2, u_2]$ separately for the two constituent variables v_{ij} and v_{ik}, respectively. Then, the $(1-2\delta)\cdot100\%$ confidence interval for φ' would be given by the sum of the two intervals $[l_1, u_1]$ and $[l_2, u_2]$, i.e., $[l_1 + l_2, u_1 + u_2]$; notice the drop in overall confidence due to the union bound. The drop in the confidence level and the additional error introduced by the interval arithmetic accumulate quickly for larger PSEs, making the estimate unusable. Furthermore, we lose all the advantages of having any dependence between

Fig. 2: Variation of ratio of the est. error using the existing approach [3] to est. error using our approach, w.r.t. the size of the chosen PSE.

the terms in the PSE. For instance, by observing that v_{ij} and v_{ik} correspond to the mutually exclusive transitions i to j and i to k, we know that $\varphi'(M)$ is always less than 1, a feature that will be lost if we use plain merging of individual confidence intervals for v_{ij} and v_{ik}. We overcome these issues by estimating the value of the PSE as a whole as much as possible. In Fig. 2, we demonstrate how the ratio between the estimation errors from the two approaches vary as the number of summands (i.e., n) in the PSE $\varphi = \sum_{i=1}^{n} v_{1n}$ changes; in both cases we fixed the overall δ to 0.05 (95% confidence). The ratio remains the same for different observation lengths. Our approach is always at least as accurate as their approach [3], and is significantly better for larger PSEs.
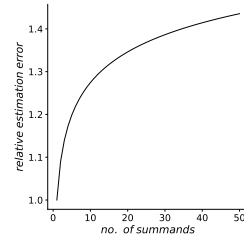
4.1   The Main Principle

We first explain the idea for division-free PSEs, i.e., PSEs that do not involve any division operator; later we extend our approach to the general case.

Divison-free PSEs: In our algorithm, for every variable $v_{ij} \in V_\varphi$, we introduce a *Bernoulli*(M_{ij}) random variable Y^{ij} with the mean M_{ij} unknown to us. We make an observation y_p^{ij} for every p-th visit to the state i on a run, and if j follows immediately afterwards then record $y_p^{ij} = 1$ else record $y_p^{ij} = 0$. This gives us a sequence of observations $\overrightarrow{y}^{ij} = y_1^{ij}, y_2^{ij}, \ldots$ corresponding to the sequence of i.i.d. random variables $\overrightarrow{Y}^{ij} = Y_1^{ij}, Y_2^{ij}, \ldots$. For instance, for the run 121123 we obtain $\overrightarrow{y}^{12} = 1, 0, 1$ for the variable v_{12}.

The heart of our algorithm is an aggregation procedure of every sequence of random variable $\{\overrightarrow{Y}^{ij}\}_{v_{ij} \in V_\varphi}$ to a single i.i.d. sequence $\overrightarrow{W}$ of an auxiliary random variable W, such that the mean of W is $\mu_W = \mathbb{E}(W) = \varphi(M)$. We can then use known concentration inequalities on the sequence $\overrightarrow{W}$ to estimate μ_W. Since μ_W exactly equals $\varphi(M)$ by design, we obtain a tight concentration bound on $\varphi(M)$. We informally explain the main idea of constructing $\overrightarrow{W}$ using simple examples; the details can be found in Alg. 2.

Sum and difference: Let $\varphi = v_{ij} + v_{kl}$. We simply combine $\overrightarrow{Y}^{ij}$ and $\overrightarrow{Y}^{kl}$ as $W_p = Y_p^{ij} + Y_p^{kl}$, so that $w_p = y_p^{ij} + y_p^{kl}$ is the corresponding observation of W_p. Then $\mu_{W_p} = \varphi(M)$ holds, because $\mu_{W_p} = \mathbb{E}(W_p) = \mathbb{E}(Y_p^{ij} + Y_p^{kl}) = \mathbb{E}(Y_p^{ij}) + \mathbb{E}(Y_p^{kl}) = M_{ij} + M_{kl}$. Similar approach works for $\varphi = v_{ij} - v_{kl}$.

Multiplication: For multiplications, the same linearity principle will not always work, since for random variables A and B, $\mathbb{E}(A \cdot B) = \mathbb{E}(A) \cdot \mathbb{E}(B)$ *only if* A and B are statistically independent, which will not be true for specifications of the form $\varphi = v_{ij} \cdot v_{ik}$. In this case, the respective Bernoulli random variables Y_p^{ij} and Y_p^{ik} are dependent: $\mathbb{P}(Y_p^{ij} = 1) \cdot \mathbb{P}(Y_p^{ik} = 1) = M_{ij} \cdot M_{ik}$, but $\mathbb{P}(Y_p^{ij} = 1 \wedge Y_p^{ik} = 1)$ is always 0 (since *both* j and k cannot be visited following the p-th visit to i).

To benefit from independence once again, we temporally shift one of the random variables by defining $W_p = Y_{2p}^{ij} \cdot Y_{2p+1}^{ik}$, with $w_p = y_{2p}^{ij} \cdot y_{2p+1}^{ik}$. Since the random variables Y_{2p}^{ij} and Y_{2p+1}^{ik} are independent, as they use separate visits of state i, hence we obtain $\mu_{W_p} = M_{ij} \cdot M_{ik}$. For independent multiplications of the form $\varphi = v_{ij} \cdot v_{kl}$ with $i \neq k$, we can simply use $W_p = Y_p^{ij} \cdot Y_p^{ik}$.

In general, we use the ideas of aggregation and temporal shift on the syntax tree of the PSE φ, inductively. With an aggregated sequence of observations for the auxiliary variable W for φ, we can find an estimate for $\varphi(M)$ using the Hoeffding's inequality. We present the detailed algorithm of this monitor, namely `FreqMonitorDivFree`, in Alg. 1.

The general case (PSEs with division operators): We observe that every arbitrary PSE φ of size n can be transformed into a semantically equivalent PSE of the form $\varphi_a + \frac{\varphi_b}{\varphi_c}$ of size $\mathcal{O}(n^2 2^n)$, where φ_a, φ_b, and φ_c are all division-free. Once in this form, we can employ three different `FreqMonitorDivFree` monitors from Alg. 1 to obtain separate interval estimates for φ_a, φ_b, and φ_c, which are then combined using standard interval arithmetic and the resulting confidence of the estimate is obtained through the union bound. The steps for constructing the (general-case) `FrequentistMonitor` are shown in Alg. 2, and the detailed analysis can be found in the proof of Thm. 1.

Bounding memory: Consider a PSE $\varphi = v_{ij} + v_{kl}$. The outcome w_p for φ can only be computed when both the Bernoulli outcomes y_p^{ij} and y_p^{kl} are available. If at any point only one of the two is available, then we need to store the available one so that it can be used later when the other one gets available. It can be shown that the storage of "unmatched" outcomes may need unbounded memory.

To bound the memory, we use the insight that a *random reshuffling* of the i.i.d. sequence y_p^{ij} would still be i.i.d. with the same distribution, so that we do not need to store the exact order in which the outcomes appeared. Instead, for

every $v_{ij} \in V_\varphi$, we only store the number of times we have seen the state i and the edge (i, j) in counters c_i and c_{ij}, respectively. Observe that $c_i \geq \sum_{v_{ik} \in V_\varphi} c_{ik}$, where the possible difference accounts for the visits to irrelevant states, denoted as a dummy state $\top$. Given $\{c_{ik}\}_k$, whenever needed, we generate in x_i a *random reshuffling* of the sequence of states, together with $\top$, seen after the past visits to i. From the sequence stored in x_i, for every $v_{ik} \in V_\varphi$, we can consistently determine the value of y_p^{ik} (consistency dictates $y_p^{ik} = 1 \Rightarrow y_p^{ij} = 0$). Moreover, we reuse space by resetting x_i whenever the sequence stored in x_i is no longer needed. It is shown in the proof of Thm. 2 in App. A.2 that the size of every x_i can be at most the size of the expression. This random reshuffling of the observation sequences is the cause of the probabilistic transitions of the frequenitst monitor.

4.2   Implementation of the Frequentist Monitor

Fix a problem instance (Q, φ, δ), with size of φ being n. Let φ be transformed into φ^l by relabeling duplicate occurrences of v_{ij} using distinct labels $v_{ij}^1, v_{ij}^2, \ldots$. The set of labeled variables in φ^l is V_φ^l, and $|V_\varphi^l| = \mathcal{O}(n)$. Let $SubExpr(\varphi)$ denote the set of every subexpression in the expression φ, and use $[l_\varphi, u_\varphi]$ to denote the range of values the expression φ can take for every valuation of every variable as per the domain $[0, 1]$. Let $Dep(\varphi) = \{i \mid \exists v_{ij} \in V_\varphi\}$, and every subexpression $\varphi_1 \cdot \varphi_2$ with $Dep(\varphi_1) \cap Dep(\varphi_2) \neq \emptyset$ is called a *dependent multiplication*.

Implementation of `FreqMonitorDivFree` in Alg. 1 has two main functions. *Init* initializes the registers. *Next* implements the transition function of the monitor, which attempts to compute a new observation w for $\overrightarrow{W}$ (Line 4) after observing a new input σ', and if successful it updates the output of the monitor by invoking the *UpdateEst* function. In addition to the registers in *Init* and *Next* labeled in the pseudocode, following registers are used internally:

- x_i, $i \in Dom(V_\varphi)$: reshuffled sequence of states that followed i.
- t_{ij}^l: the index of x_i that was used to obtain the latest outcome of v_{ij}^l.

Now, we summarize the main results for the frequentist monitor.

Theorem 1 (Correctness). *Let (Q, φ, δ) be a problem instance. Alg. 2 implements a monitor for (Q, φ, δ) that solves Prob. 1.*

Theorem 2 (Computational resources). *Let (Q, φ, δ) be a problem instance and $\mathcal{A}$ be the monitor implemented using the `FrequentistMonitor` routine of Alg. 2. Suppose the size of φ is n. The monitor $\mathcal{A}$ requires $\mathcal{O}(n^4 2^{2n})$ registers, and takes $\mathcal{O}(n^4 2^{2n})$ time to update its output after receiving a new input symbol. For the special case of φ containing at most one division operator (division by constant does not count), $\mathcal{A}$ requires only $\mathcal{O}(n^2)$ registers, and takes only $\mathcal{O}(n^2)$ time to update its output after receiving a new input symbol.*

There is a tradeoff between the estimation error, the confidence, and the length of the observed sequence of input symbols. For instance, for a fixed confidence, the longer the observed sequence is, the smaller is the estimation error. The following theorem establishes a lower bound on the length of the sequence for a given upper bound on the estimation error and a fixed confidence.

Algorithm 1 FreqMonitorDivFree

Parameters: Q, φ, δ

Output: Λ

1: **function** $Init(\sigma)$

2:    $\varphi^l \xleftarrow{\text{unique labeling}} \varphi$

3:    **for all** $v_{ij} \in V_\varphi$ **do**

4:       $c_{ij} \leftarrow 0$ ▷# of (i,j)

5:       $c_i \leftarrow 0$ ▷# of i

6:    $n \leftarrow 0$ ▷length of $\vec{w}$

7:    $\sigma \leftarrow \sigma$ ▷prev. symbol

8:    $\mu_\Lambda \leftarrow \bot$ ▷est. mean

9:    $\varepsilon_\Lambda \leftarrow \bot$ ▷est. error

10:    $ResetX()$ ▷reset x_i-s

11:    Compute l_φ, u_φ ▷int. arith.

1: **function** $Next(\sigma')$

2:    $c_\sigma \leftarrow c_\sigma + 1$ ▷update counters

3:    $c_{\sigma\sigma'} \leftarrow c_{\sigma\sigma'} + 1$

4:    $w \leftarrow Eval(\varphi^l)$

5:    **if** $w \neq \bot$ **then**

6:       $n \leftarrow n + 1$

7:       $\Lambda \leftarrow UpdateEst(w, n)$

8:       $ResetX()$

9:    $\sigma \leftarrow \sigma'$

10:    **return** Λ

1: **function** $Eval(\varphi^l)$

2:    **if** $r_{\varphi^l} = \bot$ **then**

3:       **if** $\varphi^l \equiv \varphi_1^l + \varphi_2^l$ **then**

4:          $r_{\varphi^l} \leftarrow Eval(\varphi_1^l) + Eval(\varphi_2^l)$

5:       **else if** $\varphi^l \equiv \varphi_1^l - \varphi_2^l$ **then**

6:          $r_{\varphi^l} \leftarrow Eval(\varphi_1^l) - Eval(\varphi_2^l)$

7:       **else if** $\varphi^l \equiv \varphi_1^l \cdot \varphi_2^l$ **then**

8:          **if** $Dep(V_{\varphi_1}^l) \cap Dep(V_{\varphi_2}^l) = \emptyset$ **then**

9:             $r_{\varphi^l} \leftarrow Eval(\varphi_1^l) \cdot Eval(\varphi_2^l)$

10:          **else**   ▷dep. mult.

11:             **for** $v_{ij}^l \in V_{\varphi_2}^l \cap Dep(V_{\varphi_1}^l)$ **do**

12:                $t_{ij}^l \leftarrow \max(\{t_{ik}^m \mid v_{ik}^m \in V_{\varphi_1}^l\})$

13:                $t_{ij}^l \leftarrow t_{ij}^l + 1$ ▷make indep.

14:             $r_{\varphi^l} \leftarrow Eval(\varphi_1^l) \cdot Eval(\varphi_2^l)$

15:       **else if** $\varphi^l \equiv v_{ij}^l$ **then**

16:          **if** $x_i[t_{ij}^l + 1] = \bot$ **then**

17:             $ExtractOutcome(x_i, t_{ij}^l + 1)$

18:          **if** $x_i[t_{ij}^l + 1] = j \neq \bot$ **then**

19:             $r_{\varphi^l} \leftarrow 1$

20:          **else**

21:             $r_{\varphi^l} \leftarrow 0$

22:       **else if** $\varphi^l \equiv c$ **then**

23:          $r_{\varphi^l} \leftarrow c$

24:    **return** r_{φ^l}

1: **function** $UpdateEst(w, n)$

2:    $\mu_\Lambda \leftarrow \frac{\mu_\Lambda \cdot (n-1) + w}{n}$

3:    $\varepsilon_\Lambda \leftarrow \sqrt{-\frac{(u_\varphi - l_\varphi)^2}{2n} \cdot \ln\left(\frac{\delta}{2}\right)}$

4:    **return** $[\mu_\Lambda \pm \varepsilon_\Lambda]$

1: **function** $ExtractOutcome(x_i, t)$

   ▷generate a shuffled sequence of symbols seen after i so that $|x_i| = t$

2:    Let $U \leftarrow \{j \in Q \mid v_{ij} \in V_\varphi\}$

3:    **for** $p = |x_i| + 1, \ldots, t$ **do**

4:       $q \leftarrow \forall u \in U$.

          pick u w/ prob. $\frac{c_{iu}}{c_i}$,

          pick $\top$ w/ prob. $\frac{(c_i - \sum_j c_{ij})}{c_i}$

5:       $c_i \leftarrow c_i - 1$

6:       **if** $q \neq \top$ **then**

7:          $c_{iq} \leftarrow c_{iq} - 1$

8:       $x_i[|x_i| + 1] \leftarrow q$

1: **function** $ResetX()$

2:    **for all** $i \in Dom(V_\varphi)$ **do**

3:       $x_i \leftarrow \emptyset$

4:    **for all** $v_{ij}^l \in V_\varphi^l$ **do**

5:       $t_{ij}^l \leftarrow 0$

Algorithm 2 `FrequentistMonitor`

Parameters: Q, φ, δ
Output: Λ

1: **function** $Init(\sigma)$

2:     $\varphi_a + \frac{\varphi_b}{\varphi_c} \xleftarrow{\text{change form}} \varphi^l \xleftarrow{\text{labeling}} \varphi$

3:     $\mathcal{A}_a \leftarrow \texttt{FreqMonitorDivFree}(Q, \varphi_a, \delta/3)$

4:     $\mathcal{A}_b \leftarrow \texttt{FreqMonitorDivFree}(Q, \varphi_b, \delta/3)$

5:     $\mathcal{A}_c \leftarrow \texttt{FreqMonitorDivFree}(Q, \varphi_c, \delta/3)$

6:     $\mathcal{A}_a.Init(\sigma)$

7:     $\mathcal{A}_b.Init(\sigma)$

8:     $\mathcal{A}_c.Init(\sigma)$

1: **function** $Next(\sigma')$

2:     $[\mu_a \pm \varepsilon_a] \leftarrow \mathcal{A}_a.Next(\sigma')$

3:     $[\mu_b \pm \varepsilon_b] \leftarrow \mathcal{A}_b.Next(\sigma')$

4:     $[\mu_c \pm \varepsilon_c] \leftarrow \mathcal{A}_c.Next(\sigma')$

5:     **if** $\mu_a \neq \perp \wedge \mu_b \neq \perp \wedge \mu_c \neq \perp$ **then**

6:         $[\mu_\Lambda \pm \varepsilon_\Lambda] \leftarrow [\mu_a \pm \varepsilon_a] + \frac{[\mu_b \pm \varepsilon_b]}{[\mu_c \pm \varepsilon_c]}$

7:     **return** $[\mu_\Lambda \pm \varepsilon_\Lambda]$

Theorem 3 (Convergence speed). *Let (Q, φ, δ) be a problem instance where φ does not contain any division operator, and let $\mathcal{A}$ be the monitor computed using Alg. 2. Suppose the size of φ is n. For a given upper bound on estimation error $\bar{\varepsilon} \in \mathbb{R}$, the minimum number of visits to every state in $Dom(V_\varphi)$ for obtaining an output with error at most $\bar{\varepsilon}$ and confidence at least $1 - \delta$ on any path is given by:*

$$-\frac{(u_\varphi - l_\varphi)^2 \ln\left(\frac{\delta}{2}\right) n}{2\bar{\varepsilon}^2}, \tag{4}$$

where $[l_\varphi, u_\varphi]$ is the set of possible values of φ for every valuation of every variable (having domain $[0,1]$) in φ.

The bound follows from the Hoeffding's inequality, together with the fact that every dependent multiplication increments the required number of samples by 1. A similar bound for the general case with division is left open.

5   Bayesian Monitoring

Fix a problem instance $(Q = [1 .. N], \varphi, \delta)$. Let $\mathbb{M} = \Delta(N-1)^N$ be the shorthand notation for the set of transition probability matrices of the Markov chains with state space Q. Let $p_\theta \colon \mathbb{M} \to [0,1]$ be the prior probability density function over $\mathbb{M}$, which is assumed to be specified using the matrix beta distribution (the definition is standard [36, pp. 280] and has been included in App. A.3 for completeness). Let $\mathbb{1}$ be a matrix, with its size dependent on the context, whose every element is 1. We make the following common assumption [31,36, p. 50]:

Assumption 1 (Prior) *We are given a parameter matrix $\theta \geq \mathbb{1}$, and p_θ is specified using the matrix beta distribution with parameter θ. Moreover, the initial state of the Markov chain is fixed.*

When $\theta = \mathbb{1}$, then p_θ is the uniform density function over $\mathbb{M}$. After observing a path $\vec{x}$, using Bayes' rule we obtain the *posterior* density function $p_\theta(\cdot \mid \vec{x})$, which is known to be efficiently computable due to the so-called conjugacy property that holds due to Assump. 1 (see App. A.3 for details). From

the posterior density, we obtain the expected posterior semantic value of φ as: $\mathbb{E}_\theta(\varphi(M) \mid \vec{x}) \coloneqq \int_\mathbb{M} \varphi(M) \cdot p_\theta(M \mid \vec{x}) dM$. The heart of our Bayesian monitor is an efficient incremental computation of $\mathbb{E}_\theta(\varphi(M) \mid \vec{x})$—free from numerical integration. Once we can compute $\mathbb{E}_\theta(\varphi(M) \mid \vec{x})$, we can also compute the posterior variance S^2 of $\varphi(M)$ using the known expression $S^2 = \mathbb{E}_\theta(\varphi^2(M) \mid \vec{x}) - \mathbb{E}_\theta(\varphi(M) \mid \vec{x})$, which enables us to compute a confidence interval for $\varphi(M)$ using the Chebyshev's inequality (see Prop. 1 in App. A.3). In the following, we summarize our procedure for estimating $\mathbb{E}_\theta(\varphi(M) \mid \vec{x})$.

5.1   The Main Principle

The incremental computation of $\mathbb{E}_\theta(\varphi(M) \mid \vec{x})$ is implemented in `BayesExpMonitor`. We first transform the expression φ into the polynomial form $\varphi' = \sum_l \kappa_l \xi_l$, where $\{\kappa_l\}_l$ are the weights and $\{\xi_l\}_l$ are monomials. If the size of φ is n then the size of φ' is $\mathcal{O}(n2^{\frac{n}{2}})$. Then we can use linearity to compute the overall expectation as the weighted sum of expectations of the individual monomials: $\mathbb{E}_\theta(\varphi(M) \mid \vec{x}) = \mathbb{E}_\theta(\varphi'(M) \mid \vec{x}) = \sum_l \kappa_l \mathbb{E}_\theta(\xi_l(M) \mid \vec{x})$. In the following, we summarize the procedure for estimating $\mathbb{E}_\theta(\xi(M) \mid \vec{x})$ for every monomial ξ.

Let ξ be a monomial, and let $\vec{x}ab \in Q^*$ be a sequence of states. We use d_{ij} to store the exponent of the variable v_{ij} in the monomial ξ, and define $d_a \coloneqq \sum_{j \in [1..N]} d_{aj}$. Also, we record the sets of (i,j)-s and i-s with positive and negative d_{ij} and d_i entries: $D_i^+ \coloneqq \{j \mid d_{ij} > 0\}$, $D_i^- \coloneqq \{j \mid d_{ij} < 0\}$, $D^+ \coloneqq \{i \mid d_i > 0\}$, and $D^- \coloneqq \{i \mid d_i < 0\}$.

For any given word $\vec{w} \in Q^*$, let $c_{ij}(\vec{w})$ denote the number of ij-s in $\vec{w}$ and let $c_i(\vec{w}) \coloneqq \sum_{j \in Q} c_{ij}(\vec{w})$. Define $\bar{c}_i(\vec{w}) \coloneqq c_i(\vec{w}) + \sum_{j \in [1..N]} \theta_{ij}$ and $\bar{c}_{ij}(\vec{w}) \coloneqq c_{ij}(\vec{w}) + \theta_{ij}$. Let $\mathcal{H} \colon Q^* \to \mathbb{R}$ be defined as:

$$\mathcal{H}(\vec{w}) \coloneqq \frac{\prod_{i=1}^{N} \prod_{j \in D_i^+} {}^{(\bar{c}_{ij}(\vec{w})-1)+|d_{ij}|}P_{|d_{ij}|}}{\prod_{i \in D^+} {}^{(\bar{c}_i(\vec{w})-1)+|d_i|}P_{|d_i|}} \cdot \frac{\prod_{i \in D^-} {}^{(\bar{c}_i(\vec{w})-1)}P_{|d_i|}}{\prod_{i=1}^{N} \prod_{j \in D_i^-} {}^{(\bar{c}_{ij}(\vec{w})-1)}P_{|d_{ij}|}}, \quad (5)$$

where ${}^nP_k \coloneqq \frac{n!}{(n-k)!}$ is the number of permutations of $k > 0$ items from $n > 0$ objects, for $k \leq n$, and we use the convention that for $S = \emptyset$, $\prod_{s \in S} \ldots = 1$. Below, in Lem. 1, we establish that $\mathbb{E}_\theta(\xi(M) \mid \vec{w}) = \mathcal{H}(\vec{w})$, and present an efficient incremental scheme to compute $\mathbb{E}_\theta(\xi(M) \mid \vec{x}ab)$ from $\mathbb{E}_\theta(\xi(M) \mid \vec{x}a)$.

Lemma 1 (Incremental computation of $\mathbb{E}(\cdot \mid \cdot)$). *If the following consistency condition*

$$\forall i, j \in [1..N] . \bar{c}_{ij}(\vec{w}) + d_{ij} > 0 \quad (6)$$

is met, then the following holds:

$$\mathbb{E}(\xi(M) \mid \vec{x}ab) = \mathcal{H}(\vec{x}ab) = \mathcal{H}(\vec{x}a) \cdot \frac{\bar{c}_{ab}(\vec{x}) + d_{ab}}{\bar{c}_{ab}(\vec{x})} \cdot \frac{\bar{c}_a(\vec{x})}{\bar{c}_a(\vec{x}) + d_a}. \quad (7)$$

Cond. (6) guarantees that the permutations in (5) are well-defined. The first equality in (7) follows from Marchal et al. [50], and the rest uses the conjugacy of the prior (proof in App. A.3). Lem. 1 forms the basis of the efficient update of our

---

Algorithm 3 `BayesExpMonitor`

Parameters: $Q, \varphi = \sum_{l=1}^{p} \kappa_l \xi_l, \theta$
Output: E

1: **function** $Init(\sigma = 1)$
2:     **for** $v_{ij} \in V_\varphi$ **do**
3:         $\bar{c}_{ij} \leftarrow \theta_{ij}$
4:         $\bar{c}_i \leftarrow \sum_{j \in [1..N]} \theta_{ij}$
5:         $m_{ij} \leftarrow \min_{l \in [1..p]} d_{ij}^l$ ▷cache
6:     active $\leftarrow$ *false*   ▷eq. 6 not true
7:     $\sigma \leftarrow \sigma$ ▷prev. state
8:     $E \leftarrow \bot$ ▷expect. val.

1: **function** $Next(\sigma')$
2:     $\bar{c}_\sigma \leftarrow \bar{c}_\sigma + 1$ ▷update counters
3:     $\bar{c}_{\sigma\sigma'} \leftarrow \bar{c}_{\sigma\sigma'} + 1$
4:     **if** active $=$ *false* **then**
5:         **if** $(\forall v_{ij} \in V_\varphi \, . \, \bar{c}_{ij} + m_{ij} > 0)$ **then**
6:             active $\leftarrow$ *true*   ▷eq. 6 is true
7:             **for** $l \in [1..p]$ **do**   ▷eq. 5
8:                 $h^l \leftarrow \mathcal{H}^l\left(\{\bar{c}_{ij}\}_{i,j}, \{\bar{c}_i\}_i\right)$
9:     **else**
10:         **for** $l \in [1..p]$ **do**   ▷eq. 7
11:             $h^l \leftarrow h^l \cdot \frac{\bar{c}_{\sigma\sigma'}-1+d_{\sigma\sigma'}^l}{\bar{c}_{\sigma\sigma'}-1} \cdot \frac{\bar{c}_\sigma-1}{\bar{c}_\sigma-1+d_\sigma^l}$
12:     **if** active $=$ *true* **then**
13:         $E \leftarrow \sum_{l=1}^{p} \kappa_l \cdot h^l$ ▷overall expect.
14:     $\sigma \leftarrow \sigma'$
15:     **return** E

Bayesian monitor. Observe that on any given path, once (6) holds, it continues to hold forever. Thus, initially the monitor keeps updating $\mathcal{H}$ internally without outputting anything. Once (6) holds, it keeps outputting $\mathcal{H}$ from then on.

5.2   Implementation of the Bayesian Monitor

We present the Bayesian monitor implementation in `BayesConfIntMonitor` (Alg. 4), which invokes `BayesExpMonitor` (Alg. 3) as subroutine. `BayesExpMonitor` computes the expected semantic value of an expression φ in polynomial form, by computing the individual expected value of each monomial using Prop. 1, and combining them using the linearity property. We drop the arguments from $\bar{c}_i(\cdot)$ and $\bar{c}_{ij}(\cdot)$ and simply write $\bar{c}_i$ and $\bar{c}_{ij}$ as constants associated to appropriate words. The symbol m_{ij} in Line 5 of *Init* is used as a bookkeeping variable for quickly checking the consistency condition (Eq. 6) in Line 5 of *Next*. In `BayesConfIntMonitor`, we compute the expected value and the variance of φ, by invoking `BayesExpMonitor` on φ and φ^2 respectively, and then compute the confidence interval using the Chebyshev's inequality (Thm. 7 in App. A.2). It can be observed in the *Next* subroutines of `BayesConfIntMonitor` and `BayesExpMonitor` that a deterministic transition function suffices for the Bayesian monitors.

Theorem 4 (Correctness). *Let (Q, φ, δ) be a problem instance, and p_θ be given as the prior distribution which satisfies Assump. 1. Alg. 4 produces a monitor for (Q, φ, δ) that solves Prob. 2.*

Theorem 5 (Computational resources). *Let (Q, φ, δ) be a problem instance and $\mathcal{A}$ be the monitor computed using the `BayesConfIntMonitor` routine of Alg. 4. Suppose the size of φ is n. The monitor $\mathcal{A}$ requires $\mathcal{O}(n^2 2^n)$ registers, and takes $\mathcal{O}(n^2 2^n)$ time to update its output after receiving a new input symbol.*

Algorithm 4 `BayesConfIntMonitor`

Parameters: Q, φ, θ

Output: Λ

1: **function** $Init(\sigma = 1)$

2:     $\overline{\varphi} \xleftarrow{\text{polyn.}} \varphi, \overline{\varphi^2} \xleftarrow{\text{polyn.}} \varphi^2$ ▷polyn. form

3:     $EXP \leftarrow$ `BayesExpMonitor`$(Q, \overline{\varphi}, \theta)$

4:     $EXP2 \leftarrow$ `BayesExpMonitor`$(Q, \overline{\varphi^2}, \theta)$

5:     $EXP.Init(\sigma)$

6:     $EXP2.Init(\sigma)$

7:     $\Lambda \leftarrow \perp$

1: **function** $Next(\sigma')$

2:     $E \leftarrow EXP.Next(\sigma')$

3:     $E2 \leftarrow EXP2.Next(\sigma')$

4:     **if** $E \neq \perp$ and $E2 \neq \perp$ **then**

5:         $S \leftarrow E2 - E^2$ ▷variance

6:         $\Lambda \leftarrow \left[E \pm \sqrt{\frac{S}{\delta}} \right]$ ▷Chebysh.

7:     **return** Λ

For the special case of φ being in polynomial form, $\mathcal{A}$ requires only $\mathcal{O}(n^2)$ registers, and takes only $\mathcal{O}(n^2)$ time to update its output after receiving a new input symbol.
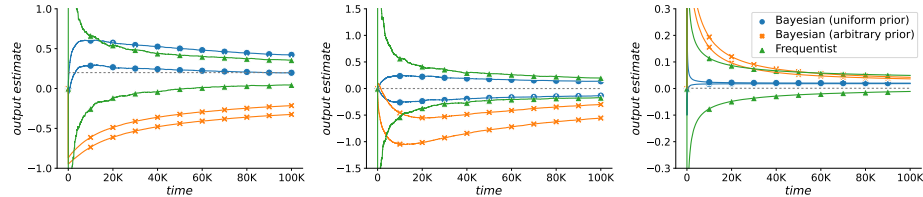
A bound on the convergence speed of the Bayesian monitor is left open. This would require a bound on the change in variance with respect to the length of the observed path, which is not known for the general case of PSEs. Note that the efficient (quadratic) cases are different for the frequentist and Bayesian monitors, suggesting the use of different monitors for different specifications.

6   Experiments

We implemented our frequentist and Bayesian monitors in a tool written in Rust, and used the tool to design monitors for the lending and the college admission examples taken from the literature [47,53] (described in Sec. 1.1). The generators are modeled as Markov chains (see Fig. 1)—unknown to the monitors—capturing the sequential interactions between the decision-makers (i.e., the bank or the college) and their respective environments (i.e., the loan applicants or the students), as described by D'Amour et al. [16]. The setup of the experiments is as follows: We created a multi-threaded wrapper program, where one thread simulates one long run of the Markov chain, and a different thread executes the monitor. Every time a new state is visited by the Markov chain on the first thread, the information gets transmitted to the monitor on the second thread, which then updates the output. The experiments were run on a Macbook Pro 2017 equipped with a 2,3 GHz Dual-Core Intel Core i5 processor and 8GB RAM.

We summarize the experimental results in Fig 3, and, from the table, observe that both monitors are extremely lightweight: they take less than a millisecond per update and small numbers of registers to operate. From the plots, we observe that the frequentist monitors' outputs are always centered around the ground truth values of the properties, empirically showing that they are always objectively correct. On the other hand, the Bayesian monitors' outputs can vary drastically for different choices of the prior, empirically showing that the correctness of outputs is subjective. It may be misleading that the outputs of the Bayesian monitors are wrong as they often do not contain the ground truth values. We reiterate that from the Bayesian perspective, the ground truth does not

Fig. 3: The plots show the 95% confidence intervals estimated by the monitors over time, averaged over 10 different sample paths, for the lending with demographic parity (left), lending with equalized opportunity (middle), and the college admission with social burden (right) problems. The horizontal dotted lines are the ground truth values of the properties, obtained by analyzing the Markov chains used to model the systems (unknown to the monitors). The table summarizes various performance metrics.

Scenario	Size of expression	Av. comp. time/step		# registers	
		Freq.	Bayes.	Freq.	Bayes.
Lending (bias) + dem. par.	1	13.0μs	29.3μs	15	17
Lending (fair) + eq. opp.	5	21.6μs	31.0μs	29	27
Admission + soc. burden	19	53.8μs	184.6μs	46	102

exist. Instead, we only have a probability distribution over the true values that gets updated after observing the generated sequence of events. The choice of the type of monitor ultimately depends on the application requirements.

7 Conclusion

We showed how to monitor algorithmic fairness properties on a Markov chain with unknown transition probabilities. Two separate algorithms are presented, using the frequentist and the Bayesian approaches to statistics. The performances of both approaches are demonstrated, both theoretically and empirically.

Several future directions exist. Firstly, more expressive classes of properties need to be investigated to cover a broader range of algorithmic fairness criteria. We believe that boolean logical connectives, as well as min and max operators can be incorporated straightforwardly using ideas from the related literature [3]. This also adds support for absolute values, since $|x| = \max\{x, -x\}$. On the other hand, properties that require estimating how often a state is visited would require more information about the dynamics of the Markov chain, including its mixing time. Monitoring statistical hyperproperties [18] is another important direction, which will allow us to encode individual fairness properties [21]. Secondly, more liberal assumptions on the system model will be crucial for certain practical applications. In particular, hidden Markov models, time-inhomogeneous Markov models, Markov decision processes, etc., are examples of system models with widespread use in real-world applications. Finally, better error bounds tailored for specific algorithmic fairness properties can be developed through a deeper

mathematical analysis of the underlying statistics, which will sharpen the conservative bounds obtained through off-the-shelf concentration inequalities.

References

1. Agha, G., Palmskog, K.: A survey of statistical model checking. ACM Transactions on Modeling and Computer Simulation (TOMACS) **28**(1), 1–39 (2018)
2. Albarghouthi, A., D'Antoni, L., Drews, S., Nori, A.V.: Fairsquare: probabilistic verification of program fairness. Proceedings of the ACM on Programming Languages **1**(OOPSLA), 1–30 (2017)
3. Albarghouthi, A., Vinitsky, S.: Fairness-aware programming. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. pp. 211–219 (2019)
4. Ashok, P., Křetínský, J., Weininger, M.: Pac statistical model checking for markov decision processes and stochastic games. In: International Conference on Computer Aided Verification. pp. 497–519. Springer (2019)
5. Baier, C., Haverkort, B., Hermanns, H., Katoen, J.P.: Model-checking algorithms for continuous-time markov chains. IEEE Transactions on Software Engineering **29**(6), 524–541 (2003). `https://doi.org/10.1109/TSE.2003.1205180`
6. Balunovic, M., Ruoss, A., Vechev, M.: Fair normalizing flows. In: International Conference on Learning Representations (2021)
7. Bartocci, E., Deshmukh, J., Donzé, A., Fainekos, G., Maler, O., Ničković, D., Sankaranarayanan, S.: Specification-based monitoring of cyber-physical systems: a survey on theory, tools and applications. In: Lectures on Runtime Verification, pp. 135–175. Springer (2018)
8. Bartocci, E., Falcone, Y.: Lectures on Runtime Verification. Springer (2018)
9. Bastani, O., Zhang, X., Solar-Lezama, A.: Probabilistic verification of fairness properties via concentration. Proceedings of the ACM on Programming Languages **3**(OOPSLA), 1–27 (2019)
10. Bellamy, R.K., Dey, K., Hind, M., Hoffman, S.C., Houde, S., Kannan, K., Lohia, P., Martino, J., Mehta, S., Mojsilović, A., et al.: Ai fairness 360: An extensible toolkit for detecting and mitigating algorithmic bias. IBM Journal of Research and Development **63**(4/5), 4–1 (2019)
11. Berk, R., Heidari, H., Jabbari, S., Joseph, M., Kearns, M., Morgenstern, J., Neel, S., Roth, A.: A convex framework for fair regression. arXiv preprint arXiv:1706.02409 (2017)
12. Bird, S., Dudik, M., Edgar, R., Horn, B., Lutz, R., Milan, V., Sameki, M., Wallach, H., Walker, K.: Fairlearn: A toolkit for assessing and improving fairness in ai. Microsoft, Tech. Rep. MSR-TR-2020-32 (2020)
13. Chouldechova, A.: Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. Big data **5**(2), 153–163 (2017)
14. Clarke, E.M., Zuliani, P.: Statistical model checking for cyber-physical systems. In: International symposium on automated technology for verification and analysis. pp. 1–12. Springer (2011)
15. Corbett-Davies, S., Pierson, E., Feller, A., Goel, S., Huq, A.: Algorithmic decision making and the cost of fairness. In: Proceedings of the 23rd acm sigkdd international conference on knowledge discovery and data mining. pp. 797–806 (2017)
16. D'Amour, A., Srinivasan, H., Atwood, J., Baljekar, P., Sculley, D., Halpern, Y.: Fairness is not static: Deeper understanding of long term fairness via simulation studies. In: Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency. p. 525–534. FAT* '20 (2020)

17. David, A., Du, D., Guldstrand Larsen, K., Legay, A., Mikučionis, M.: Optimizing control strategy using statistical model checking. In: NASA Formal Methods Symposium. pp. 352–367. Springer (2013)
18. Dimitrova, R., Finkbeiner, B., Torfah, H.: Probabilistic hyperproperties of markov decision processes (2020). `https://doi.org/10.48550/ARXIV.2005.03362`, `https://arxiv.org/abs/2005.03362`
19. Donzé, A., Maler, O.: Robust satisfaction of temporal logic over real-valued signals. In: International Conference on Formal Modeling and Analysis of Timed Systems. pp. 92–106. Springer (2010)
20. Dressel, J., Farid, H.: The accuracy, fairness, and limits of predicting recidivism. Science advances 4(1), eaao5580 (2018)
21. Dwork, C., Hardt, M., Pitassi, T., Reingold, O., Zemel, R.: Fairness through awareness. In: Proceedings of the 3rd innovations in theoretical computer science conference. pp. 214–226 (2012)
22. Dwork, C., Ilvento, C.: Individual fairness under composition. Proceedings of Fairness, Accountability, Transparency in Machine Learning (2018)
23. Ensign, D., Friedler, S.A., Neville, S., Scheidegger, C., Venkatasubramanian, S.: Runaway feedback loops in predictive policing. In: Conference on Fairness, Accountability and Transparency. pp. 160–171. PMLR (2018)
24. Faymonville, P., Finkbeiner, B., Schwenger, M., Torfah, H.: Real-time stream-based monitoring. arXiv preprint arXiv:1711.03829 (2017)
25. Feldman, M., Friedler, S.A., Moeller, J., Scheidegger, C., Venkatasubramanian, S.: Certifying and removing disparate impact. In: proceedings of the 21th ACM SIGKDD international conference on knowledge discovery and data mining. pp. 259–268 (2015)
26. Ferrère, T., Henzinger, T.A., Kragl, B.: Monitoring Event Frequencies. In: Fernández, M., Muscholl, A. (eds.) 28th EACSL Annual Conference on Computer Science Logic (CSL 2020). Leibniz International Proceedings in Informatics (LIPIcs), vol. 152, pp. 20:1–20:16. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2020). `https://doi.org/10.4230/LIPIcs.CSL.2020.20`, `https://drops.dagstuhl.de/opus/volltexte/2020/11663`
27. Ferrère, T., Henzinger, T.A., Saraç, N.E.: A theory of register monitors. In: Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science. pp. 394–403 (2018)
28. Finkbeiner, B., Sankaranarayanan, S., Sipma, H.: Collecting statistics over runtime executions. Electronic Notes in Theoretical Computer Science 70(4), 36–54 (2002)
29. Ghosh, B., Basu, D., Meel, K.S.: Justicia: A stochastic sat approach to formally verify fairness. arXiv preprint arXiv:2009.06516 (2020)
30. Ghosh, B., Basu, D., Meel, K.S.: Algorithmic fairness verification with graphical models. arXiv preprint arXiv:2109.09447 (2021)
31. Gómez-Corral, A., Insua, D.R., Ruggeri, F., Wiper, M.: Bayesian inference of markov processes. Wiley StatsRef: Statistics Reference Online pp. 1–15 (2014)
32. Hardt, M., Price, E., Srebro, N.: Equality of opportunity in supervised learning. Advances in neural information processing systems 29 (2016)
33. Henzinger, T.A., Karimi, M., Kueffner, K., Mallik, K.: Runtime monitoring of dynamic fairness properties. arXiv preprint arXiv:2305.04699 (2023)
34. Henzinger, T.A., Saraç, N.E.: Monitorability under assumptions. In: International Conference on Runtime Verification. pp. 3–18. Springer (2020)
35. Henzinger, T.A., Saraç, N.E.: Quantitative and approximate monitoring. In: 2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS). pp. 1–14. IEEE (2021)

36. Insua, D., Ruggeri, F., Wiper, M.: Bayesian analysis of stochastic process models. John Wiley & Sons (2012)
37. Jagielski, M., Kearns, M., Mao, J., Oprea, A., Roth, A., Sharifi-Malvajerdi, S., Ullman, J.: Differentially private fair learning. In: International Conference on Machine Learning. pp. 3000–3008. PMLR (2019)
38. John, P.G., Vijaykeerthy, D., Saha, D.: Verifying individual fairness in machine learning models. In: Conference on Uncertainty in Artificial Intelligence. pp. 749–758. PMLR (2020)
39. Junges, S., Torfah, H., Seshia, S.A.: Runtime monitors for markov decision processes. In: International Conference on Computer Aided Verification. pp. 553–576. Springer (2021)
40. Kamiran, F., Calders, T.: Data preprocessing techniques for classification without discrimination. Knowledge and information systems **33**(1), 1–33 (2012)
41. Kearns, M., Neel, S., Roth, A., Wu, Z.S.: Preventing fairness gerrymandering: Auditing and learning for subgroup fairness. In: International Conference on Machine Learning. pp. 2564–2572. PMLR (2018)
42. Kleinberg, J., Mullainathan, S., Raghavan, M.: Inherent Trade-Offs in the Fair Determination of Risk Scores. In: Papadimitriou, C.H. (ed.) 8th Innovations in Theoretical Computer Science Conference (ITCS 2017). Leibniz International Proceedings in Informatics (LIPIcs), vol. 67, pp. 43:1–43:23. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany (2017). https://doi.org/10.4230/LIPIcs.ITCS.2017.43, http://drops.dagstuhl.de/opus/volltexte/2017/8156
43. Knight, K.: Mathematical statistics. CRC Press (1999)
44. Konstantinov, N.H., Lampert, C.: Fairness-aware pac learning from corrupted data. Journal of Machine Learning Research **23** (2022)
45. Kusner, M.J., Loftus, J., Russell, C., Silva, R.: Counterfactual fairness. Advances in neural information processing systems **30** (2017)
46. Lahoti, P., Gummadi, K.P., Weikum, G.: ifair: Learning individually fair data representations for algorithmic decision making. In: 2019 ieee 35th international conference on data engineering (icde). pp. 1334–1345. IEEE (2019)
47. Liu, L.T., Dean, S., Rolf, E., Simchowitz, M., Hardt, M.: Delayed impact of fair machine learning. In: International Conference on Machine Learning. pp. 3150–3158. PMLR (2018)
48. Lum, K., Isaac, W.: To predict and serve? Significance **13**(5), 14–19 (2016)
49. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Formal Techniques, Modelling and Analysis of Timed and Fault-Tolerant Systems, pp. 152–166. Springer (2004)
50. Marchal, O., Arbel, J.: On the sub-gaussianity of the beta and dirichlet distributions. Electronic Communications in Probability **22**, 1–14 (2017)
51. Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., Galstyan, A.: A survey on bias and fairness in machine learning. ACM Computing Surveys (CSUR) **54**(6), 1–35 (2021)
52. Meyer, A., Albarghouthi, A., D'Antoni, L.: Certifying robustness to programmable data bias in decision trees. Advances in Neural Information Processing Systems **34**, 26276–26288 (2021)
53. Milli, S., Miller, J., Dragan, A.D., Hardt, M.: The social cost of strategic classification. In: Proceedings of the Conference on Fairness, Accountability, and Transparency. pp. 230–239 (2019)
54. Obermeyer, Z., Powers, B., Vogeli, C., Mullainathan, S.: Dissecting racial bias in an algorithm used to manage the health of populations. Science **366**(6464), 447–453 (2019)

55. Otop, J., Henzinger, T.A., Chatterjee, K.: Quantitative automata under probabilistic semantics. Logical Methods in Computer Science **15** (2019)
56. Scheuerman, M.K., Paul, J.M., Brubaker, J.R.: How computers see gender: An evaluation of gender classification in commercial facial analysis services. Proceedings of the ACM on Human-Computer Interaction **3**(CSCW), 1–33 (2019)
57. Seyyed-Kalantari, L., Liu, G., McDermott, M., Chen, I.Y., Ghassemi, M.: Chexclusion: Fairness gaps in deep chest x-ray classifiers. In: BIOCOMPUTING 2021: proceedings of the Pacific symposium. pp. 232–243. World Scientific (2020)
58. Sharifi-Malvajerdi, S., Kearns, M., Roth, A.: Average individual fairness: Algorithms, generalization and experiments. Advances in Neural Information Processing Systems **32** (2019)
59. Stoller, S.D., Bartocci, E., Seyster, J., Grosu, R., Havelund, K., Smolka, S.A., Zadok, E.: Runtime verification with state estimation. In: International conference on runtime verification. pp. 193–207. Springer (2011)
60. Sun, B., Sun, J., Dai, T., Zhang, L.: Probabilistic verification of neural networks against group fairness. In: International Symposium on Formal Methods. pp. 83–102. Springer (2021)
61. Wachter, S., Mittelstadt, B., Russell, C.: Bias preservation in machine learning: the legality of fairness metrics under eu non-discrimination law. W. Va. L. Rev. **123**, 735 (2020)
62. Wexler, J., Pushkarna, M., Bolukbasi, T., Wattenberg, M., Viégas, F., Wilson, J.: The what-if tool: Interactive probing of machine learning models. IEEE transactions on visualization and computer graphics **26**(1), 56–65 (2019)
63. Younes, H.L., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: International Conference on Computer Aided Verification. pp. 223–235. Springer (2002)
64. Zafar, M.B., Valera, I., Gomez-Rodriguez, M., Gummadi, K.P.: Fairness constraints: A flexible approach for fair classification. The Journal of Machine Learning Research **20**(1), 2737–2778 (2019)
65. Zemel, R., Wu, Y., Swersky, K., Pitassi, T., Dwork, C.: Learning fair representations. In: International conference on machine learning. pp. 325–333. PMLR (2013)

A   Appendices

A.1   Probability Spaces and Random Variables

A probability space is a triple $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$, where Ω is a set of outcomes, $\mathcal{F}_\Omega \subseteq 2^\Omega$ is a σ-algebra over the set Ω, and $\mathbb{P}_\Omega \colon \mathcal{F} \to [0,1]$ is a probability measure over $\mathcal{F}_\Omega$. A $((S, \mathcal{F}_S)$-valued) random variable is a measurable function $X \colon (\Omega, \mathcal{F}_\Omega) \to (S, \mathcal{F}_S)$ where $\mathcal{F}_S$ is a σ-algebra on the co-domain S. An *observation* of the random variable X is an element x in $\mathcal{F}_S$, which will always be written using the lower-case font of the same character used to denote the random variable. The random variable X induces a probability measure on the space $(S, \mathcal{F}_S)$ as follows: For every $A \in \mathcal{F}_S$, $\mathbb{P}(X \in A) = \mathbb{P}_\Omega\left(\{\omega \in \Omega \mid X(\omega) \in A\}\right)$. Often we discuss probability measure using a random variable X directly over the space $(S, \mathcal{F}_S)$ without mentioning the underlying probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$. Moreover, we often simplify notation by writing $\mathbb{P}(A)$ instead of $\mathbb{P}(X \in A)$, if the random variable X is clear from the context. For a given random variable

X, we often interchangably use the mean μ_X and the expected value $\mathbb{E}(X)$ of X.

For a given set S, define $\mathcal{D}(S)$ as the set of every random variable—called a *probability distribution*[4]—with domain $(S, 2^S)$.

Let $X, Y \in \mathcal{D}(\mathbb{R})$ be a pair of real valued random variables defined using the same underlying probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$. For $\odot \in \{+, -, *, \div\}$, define $Z = X \odot Y$ as the random variable $Z \colon (\Omega, \mathcal{F}_\Omega) \to (\mathbb{R}, 2^{\mathbb{R}})$ such that for every $A \in 2^{\mathbb{R}}$, $\mathbb{P}(Z \in A) \coloneqq \mathbb{P}_\Omega(\{\omega \in \Omega \mid X(\omega) \odot Y(\omega) \in A\})$, provided $X(\omega) \odot Y(\omega)$ is defined for every ω.

Let $X, Y \in \mathcal{D}(\mathbb{R})$ be a pair of real valued random variables defined using the same underlying probability space $(\Omega, \mathcal{F}_\Omega, \mathbb{P}_\Omega)$. Suppose $X \colon (\Omega, \mathcal{F}_\Omega) \to (S_X, \mathcal{F}_X)$ and $Y \colon (\Omega, \mathcal{F}_\Omega) \to (S_Y, \mathcal{F}_Y)$, and denote $\mathbb{P}_X$ and $\mathbb{P}_Y$ as the probability measures induced by the random variables X and Y respectively. The joint probability measure $\mathbb{P}_{X,Y}(\cdot)$ of $\mathbb{P}_X$ and $\mathbb{P}_Y$ is the unique probability measure on the space $(S_X \times S_Y, \mathcal{F}_X \otimes \mathcal{F}_Y)$, defined as:

$$\mathbb{P}_{X,Y}(E) = \mathbb{P}_\Omega((X, Y) \in E).$$

A.2   Frequentist Monitor: Additional Technical Preliminaries and Detailed Proofs

Preliminaries: Confidence Intervals We summarize three well-known statistical concentration inequalities.

Theorem 6 (Hoeffding's inequality). *Suppose $\overrightarrow{X} = X_0, \ldots, X_n$ is a sequence of i.i.d. random variables such that every $X_i \in [a, b]$ almost surely, and every X_i has mean μ_X. Let $\widehat{\mu}(\overrightarrow{X}) = \frac{\sum_{k=1}^n X_k}{n}$ be the sample mean (which is a random variable), and $\varepsilon \in \mathbb{R}$ be a constant. Then the following holds:*

$$\mathbb{P}(|\mu_X - \widehat{\mu}(\overrightarrow{X})| \leq \varepsilon) \geq 1 - 2e^{-\frac{2\varepsilon^2 n}{(b-a)^2}}. \tag{8}$$

Hoeffding's inequality requires a bound $[a, b]$ on the values of the random variable, which is overcome by the Chebyshev's inequality.

Theorem 7 (Chebyshev's inequality). *Suppose X is a random variable with finite mean μ_X and finite non-zero variance σ^2. For any given $\varepsilon \in \mathbb{R}$, the following holds:*

$$\mathbb{P}(|X - \mu_X| \leq \varepsilon) \geq 1 - \frac{\sigma^2}{\varepsilon^2}. \tag{9}$$

Detailed Proof of Technical Claims Proof of Thm. 1. Fix a problem instance (Q, φ, δ), and assume that φ is division-free so that the routine `FreqMonitorDivFree`

---

[4] In literature, a probability distribution over S is often defined as a probability measure over S, instead of the random variable like us that induces the measure.

of Alg. 1 gives us the desired monitor; we will handle the case with division separately. For a given concrete finite path $\vec{x} \in Q^*$ of the Markov chain, Alg. 1 computes a sequence $\vec{w}$ (of possibly shorter length), so that if $\vec{x}$ is a concrete sample of the evolution $\vec{X}$ of the Markov chain then $\vec{w}$ is a sample of a sequence $\vec{W} = W_1, W_2, \ldots$ of i.i.d. random variables such that $\mu_W = \varphi(M)$. Then given the sequence $\vec{w}$, we can estimate the mean μ_W using the Hoeffding's inequality. Consequently, for soundness, we need to prove the claim that (A) every member W_p of $\vec{W}$ has mean $\mathbb{E}(W_p) = \mu_W = \varphi(M)$, and moreover (B) every pair of W_p, W_k for $p \neq k$ are i.i.d. The proof is inductive over the structure of the formula φ.

Base case: If φ is a variable $v_{ij} \in V$, then the sequence $\vec{W}$ is the same as a *uniformly random reordering* of the sequence of independent Bernoulli random variables $\vec{Y}^{ij}$: that W_i-s are Bernoulli follows from Line 18 of Subr. *Eval*(v_{ij}), and that the uniformly random reordering happens follows from the invocation of the Subr. *ExtractOutcomes*() in Line 17 of *Eval*(v_{ij}). Since $\vec{Y}^{ij}$ is i.i.d., hence a uniform random reordering of $\vec{Y}^{ij}$ is also i.i.d. with the same distribution. On the other hand, if φ is a constant $c \in \mathbb{R}$, then for every p, $\mathbb{P}(W_p = c) = 1$ (Line 23 of Subr. *Eval*). It follows that both (A) and (B) hold in both cases.

Induction hypothesis: If φ is neither a variable nor a constant, then it is of the form $\varphi \equiv \varphi_1 \odot \varphi_2$ (recall that φ is assumed to be division-free), where φ_1 and φ_2 are two subformulas over variables V_1 and V_2 respectively (with $V_1 \cup V_2 \subseteq V_\varphi$) and $\odot \in \{+, -, \cdot\}$. Suppose $\vec{U} = U_1, U_2, \ldots$ and $\vec{R} = R_1, R_1, \ldots$ are sequences of independent random variables internally generated by Alg. 1 for the subformulas φ_1 and φ_2 respectively. Let (A) and (B) hold for both $\vec{U}$ and $\vec{R}$.

Induction step: Given $\varphi = \varphi_1 \odot \varphi_2$ as defined above, we have the following possibilities:

Case $\varphi \equiv \varphi_1 + \varphi_2$**:** It follows from Line 4 of Subr. *Eval* that for every p, we have $W_p := U_p + R_p$. Then using linearity of expectation it follows that $\mathbb{E}(W_p) = \mathbb{E}(U_p + R_p) = \mathbb{E}(U_p) + \mathbb{E}(R_p) = \mu_U + \mu_R = \varphi(M)$, i.e., (A) holds.

Case $\varphi \equiv \varphi_1 - \varphi_2$**:** It follows from Line 6 of Subr. *Eval* that for every p, we have $W_p := U_p - R_p$. Then using linearity of expectation it follows that $\mathbb{E}(W_p) = \mathbb{E}(U_p - R_p) = \mathbb{E}(U_p) - \mathbb{E}(R_p) = \mu_U - \mu_R = \varphi(M)$, i.e., (A) holds.

Case $\varphi \equiv \varphi_1 \cdot \varphi_2$**:** We distinguish between two cases:

    Independent multiplication: It follows from Line 9 of Subr. *Eval* that if $Dep(R_1) \cap Dep(R_2) = \emptyset$, then for every p, we have $W_p := U_p \cdot R_p$. Then $\mathbb{E}(W_p) = \mathbb{E}(U_p \cdot R_p) = \mathbb{E}(U_p) \cdot \mathbb{E}(R_p) = \mu_U \cdot \mu_R = \varphi(M)$, since U_p and R_p are independent. Hence (A) holds.

    Dependent multiplication: It follows from Line 14 of Subr. *Eval* that if $Dep(R_1) \cap Dep(R_2) \neq \emptyset$, then for every p, we have $W_p := U_{2p} \cdot R_{2p+1}$. Then $\mathbb{E}(W_p) = \mathbb{E}(U_{2p} \cdot R_{2p+1}) = \mathbb{E}(U_{2p}) \cdot \mathbb{E}(R_{2p+1}) = \mu_U \cdot \mu_R = \varphi(M)$, since U_{2p} and R_{2p+1} are independent. Hence (A) holds.

Claim (B) follows in all the above cases because the elements of $\vec{W}$ are all i.i.d. as $\vec{U}$ and $\vec{R}$ are i.i.d sequences. This completes the proof.

Now suppose φ has at least one division operator, in which case we will need to use the routine `FreqMonitor` from Alg. 2. After assigning distinct labels to the repeatedly occurring variables in φ to form φ^l, we convert φ^l to the form $\varphi_a + \frac{\varphi_b}{\varphi_c}$, where φ_a, φ_b, and φ_c are division-free. We employ the monitors $\mathcal{A}_a$, $\mathcal{A}_b$, and $\mathcal{A}_c$ to estimate the values $\varphi_a(M)$, $\varphi_b(M)$, and $\varphi_c(M)$, respectively, and the correctness of the outputs of the respective monitors follow from the argument made in the first part of this proof for division-free formulas. Note that the interval estimates $\varphi_a(M)$, $\varphi_b(M)$, and $\varphi_c(M)$ are each with confidence $\delta/3$. The claim that the output of `FreqMonitor` follows from the interval arithmetic used to estimate the interval $[\mu_\Lambda \pm \varepsilon_\Lambda]$, and the union bound used to estimate the overall estimate which is the sum of the individual estimates and equals δ; details can be found in the paper by Albarghouthi et al. [3].

Proof of Thm. 2. First, let us assume that the PSE φ is division-free, so that effectively `FreqMonitor` reduces to `FreqMonitorDivFree`. In this case, the number of registers for $\{c_{ij}\}$, $\{c_i\}$, $\{r_\varphi\}$, $\{t_{ij}^l\}$, $\{b_{\varphi'}\}$ can be at most $\mathcal{O}(n)$, where n is the number of terms in the formula φ. The total number of registers is dominated by the total space occupied by all the x_i-s (each location of the array x_i is interpreted as a register). We first argue that every x_i can grow up to size at most $\mathcal{O}(n)$. Moreover, the most amount of registers in x_i are required when the operation involved is a dependent multiplication. Observe that for every dependent multiplication $\varphi = \varphi_1 \cdot \varphi_2$ with $i \in Dep(\varphi_1) \cap Dep(\varphi_2)$, if φ_1 and φ_2 need $\mathcal{O}(m_1)$ and $\mathcal{O}(m_2)$ samples of x_i, then φ needs $m_1 + m_2$ samples of x_i. As a result, the size of x_i can be at most $\mathcal{O}(n)$, and hence the total space occupied by all the x_i registers will be $\mathcal{O}(n^2)$.

The transition function of the monitor is implemented by the Subr. *Next* and the output function is implemented by the Subr. *UpdEst*. The computation time of the transition function is dominated by the $Eval(\varphi^l)$ operation in Line 4. Observe that computation time of $Eval(\cdot)$ is dominated by the computation time of dependent multiplications, where every dependent multiplication requires $\mathcal{O}(n)$ operations to shift every t_{ij}^l by one place (there are $\mathcal{O}(n)$-many t_{ij}^l-s). Thus, in the worst case there will be $\mathcal{O}(n)$ dependent multiplications, giving us the $\mathcal{O}(n^2)$ bound on the computation time. The Subr. *UpdEst* requires constant amount of memory and runs in constant time, which can be easily observed from the pseudocode, giving us the overall quadratic bounds on the computation time and memory. This proves the last part of the theorem.

When φ contains division, then first φ is converted to the form $\varphi_a + \frac{\varphi_b}{\varphi_c}$, where φ_a, φ_b, and φ_c are all division-free. We will argue that if the size of φ is n, then the sizes of φ_a and φ_c are each $\mathcal{O}(n2^{\frac{n}{2}})$, and the size of φ_b is $\mathcal{O}(n^2 2^n)$. Therefore the computation in `FreqMonitor` will be dominated by the invocation of `FreqMonitorDivFree` on the sub-expression φ_b (Line 4 of Alg. 2). First, observe that any arbitrary PSE φ can be translated into a semantically equivalent polynomial PSE φ' of size $\mathcal{O}(n2^{\frac{n}{2}})$; a formal treatment of this claim can be found in Lem. 6. Given the polynomial PSE φ', we can collect all the division-free monomials as a sum of monomials and use it as our φ_a, whose size will be at most the size of φ', which is $\mathcal{O}(n2^{\frac{n}{2}})$. The rest of the monomials of φ',

the ones which contain divisions, have only single variables in the denominator (because of the syntax of PSEs). Hence, when we combine them in the form of a single ratio $\frac{\varphi_b}{\varphi_c}$, the denominator φ_c is a single monomial, whose size can be at most the size of the PSE φ', which is $\mathcal{O}(n2^{\frac{n}{2}})$. The numerator φ_b, on the other hand, is a sum (or difference) of $\mathcal{O}(n2^{\frac{n}{2}})$-many monomials, and every monomial can be at most $\mathcal{O}(n2^{\frac{n}{2}})$ large (because in the worst case they are of the form $\varphi_d \cdot \varphi_c$, where φ_d is some division-free term and the size of the product can be at most the size of the formula φ'). Therefore, the size of φ_b can be at most $\mathcal{O}(n^2 2^n)$, and the invocation of `FreqMonitorDivFree` dominates the memory and the computation time in `FreqMonitor`. Since `FreqMonitorDivFree` takes $\mathcal{O}(n^2)$ time and $\mathcal{O}(n^2)$ registers for its computation for an input PSE of size n, hence, for the input PSE φ_b of size $\mathcal{O}(n^2 2^n)$, `FreqMonitorDivFree` would take $\mathcal{O}(n^4 2^{2n})$ time and $\mathcal{O}(n^4 2^{2n})$ registers.

Proof of Thm. 3. Observe that every dependent multiplication requires one additional observation sample from the path, and thus in the worst case we will need n observations (n is the size of the expression) for obtaining one observation of $\overrightarrow{W}$. On the other hand, from the Hoeffding's inequality, it follows that the minimum number of samples of $\overrightarrow{W}$ required for an estimate with error at most $\overline{\varepsilon}$ is:

$$-\frac{(u_\varphi - l_\varphi)^2 \ln\left(\frac{\delta}{2}\right)}{2\overline{\varepsilon}^2}. \tag{10}$$

Thus the bound follows.

A.3   Bayesian Monitor: Additional Technical Preliminaries and Detailed Proofs

Preliminaries on Bayesian Inference and Conjugate Priors Let $\mathbb{M}$ be the set of all Markov chains with state space Q (characterized by the set of transition matrices of size $N \times N$) and initial state 1. Let be $p\colon \mathbb{M} \to \mathbb{R} \cup \{\infty\}$ the prior density function. Suppose we are given a sequence of states $\overrightarrow{w} \in Q^*$. We apply Bayes' theorem to compute the *posterior density function* from the given prior density $p(\cdot)$ as:

$$p(\mathcal{M} \mid \overrightarrow{w}) = \frac{\mathbb{P}(\overrightarrow{w} \mid \mathcal{M}) \cdot p(\mathcal{M})}{\mathbb{P}(\overrightarrow{w})} = \frac{\mathbb{P}(\overrightarrow{w} \mid \mathcal{M}) \cdot p(\mathcal{M})}{\int_{\mathbb{M}} \mathbb{P}(\overrightarrow{w} \mid \mathcal{M}') \cdot (\mathcal{M}') d\mathcal{M}'},$$

which forms the core equation in Bayesian inference.

In general, computing the posterior density is quite challenging. However, for certain likelihood function there exist priors for which the posterior distribution belongs to the same distribution family as the prior. The *likelihood function* for $M \in \mathbb{M}$ is

$$\mathcal{L}_M((c_{ij})_{i,j \in Q}) = \prod_{i=1}^{N} \prod_{j=1}^{N} M_{ij}^{c_{ij}(\overrightarrow{w})}.$$

The conjugate prior for this likelihood function is the matrix beta distribution.

Definition of the Matrix Beta Distribution The matrix beta distribution $p_\theta \colon \mathbb{M} \to \mathbb{R} \cup \{\infty\}$ is a parameterized probability distribution. Its parameter θ is an $N \times N$ matrix containing only positive entries. The density is given by $p_\theta(\mathcal{M}) = \mathcal{N}(\theta) \cdot \mathcal{L}_M(\theta - \mathbb{1})$, where $\mathcal{N}(\cdot)$ is the so-called *normalization constant* defined as:

$$\mathcal{N}(\theta) := \prod_{i=1}^{N} \frac{\Gamma\left(\sum_{j=1}^{N} \theta_{ij}\right)}{\prod_{j=1}^{N} \Gamma(\theta_{ij})},$$

with $\Gamma(\cdot)$ being the gamma function [31,36, p. 50]. For our purposes it suffices to consider $\theta \in \mathbb{N}^{N \times N}$ as a positive integer matrix.

Derivation of $\mathcal{H}$ and its Update For convenience we use the following more elaborate notation in the proof. Let $n \in \mathbb{N}$ and let $K \in \mathbb{Z}$ s.t. $n - |K| > 0$, we define $n!_K^+ = \prod_{k=0}^{|K|-1} n + k$ and $n!_K^- = \prod_{k=0}^{|K|-1} n - (k+1)$. Notice that

$$n!_K^+ = {}^{n+|K|-1}P_{|K|} \qquad \text{and} \qquad n!_K^- = {}^{n-1}P_{|K|}$$

Moreover, observe that

$$n!_K^+ = n!_K^+ \cdot \frac{n + |K|}{n} \quad \text{and} \quad (n+1)!_K^- = n!_K^- \cdot \frac{n}{n - |K|}$$

Let $A \in \mathbb{N}^{N \times N}$, let $B \in \mathbb{Z}^{N \times N}$ such that $A + B > 0$, let $\star \in \{+, -, 0\}$ and $\circ \in \{+, -\}$ we define

$$\mathcal{H}_B^{\star \circ}(A) := \prod_{i \in I_B^\star} \prod_{j \in B_i^\circ} A_{ij}!_{B_{ij}}^\circ \quad \text{and} \quad \mathcal{H}_B^\star(A) := \prod_{i \in I_B^\star} \langle A_i \rangle!_{\langle B_i \rangle}^\star$$

Where $\langle B_i \rangle$ denotes the sum of the row vector B_i. To make the derivation cleaner we define $\mathcal{H}$ slightly different (but equivalent) to Equation 5.

$$\mathcal{H}_B(A) := \left(\frac{\prod_{i \in B^-} \langle A_i \rangle!_{\langle B_i \rangle}^-}{\prod_{i \in B^+} \langle A_i \rangle!_{\langle B_i \rangle}^+} \right) \cdot \left(\prod_{i=1}^{N} \frac{\prod_{j \in J_B^+(i)} A_{ij}!_{B_{ij}}^+}{\prod_{j \in J_B^-(i)} A_{ij}!_{B_{ij}}^-} \right)$$

For any matrix $D \in \mathbb{R}^{N \times N}$ we define $D_i^+ := \{j \mid D_{ij} > 0\}$, $D_i^- := \{j \mid D_{ij} < 0\}$, $D^+ := \{i \mid \langle D_i \rangle > 0\}$, and $D^- := \{i \mid \langle D_i \rangle < 0\}$. Additionally, let $D_i^- := \{j \mid D_{ij} = 0\}$ and $D^- := \{j \mid \langle D_i \rangle = 0\}$

Lemma 2. *Let $A \in \mathbb{N}^{N \times N}$, let $B \in \mathbb{Z}^{N \times N}$ such that $A + B > 0$.*

$$\mathcal{N}(A + B) = \mathcal{N}(A) \cdot \mathcal{H}_B(A)^{-1}$$

Proof. From the definition we know that

$$\mathcal{N}(A + B) = \prod_{i=1}^{N} \frac{\Gamma\left(\sum_{j=1}^{N} A_{ij} + B_{ij}\right)}{\prod_{j=1}^{N} \Gamma(A_{ij} + B_{ij})}$$

We can split $\mathcal{N}(A+B)$ into

$$\left(\prod_{i\in B^+} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right)\right) \cdot \left(\prod_{i\in B^+} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

$$\cdot \left(\prod_{i\in B^-} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right)\right) \cdot \left(\prod_{i\in B^-} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

$$\cdot \left(\prod_{i\in B^0} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right)\right) \cdot \left(\prod_{i\in B^0} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

We know that for all $n \in \mathbb{N}$ the recursion $\Gamma(n+1) = n\Gamma(n)$ holds. We observe that

$$\prod_{i\in B^+} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right) = \left(\prod_{i\in B^+} \Gamma(\langle A_i\rangle)\right) \cdot \left(\prod_{i\in B^+} \prod_{k=0}^{\langle B_i\rangle-1} (\langle A_i\rangle + k)\right)$$

and

$$\prod_{i\in B^-} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right) = \left(\prod_{i\in B^-} \Gamma(\langle A_i\rangle)\right) \cdot \left(\prod_{i\in B^-} \prod_{k=0}^{\langle B_i\rangle-1} (\langle A_i\rangle - k - 1)\right)^{-1}$$

and

$$\prod_{i\in B^0} \Gamma\left(\sum_{j=1}^N A_{ij} + B_{ij}\right) = \prod_{i\in B^0} \Gamma(\langle A_i\rangle)$$

Hence, we obtain

$$\left(\prod_{i\in B^+} \Gamma(\langle A_i\rangle)\right) \cdot \mathcal{H}_B^+(A) \cdot \left(\prod_{i\in B^+} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

$$\left(\prod_{i\in B^-} \Gamma(\langle A_i\rangle)\right) \cdot \frac{1}{\mathcal{H}_B^-(A)} \cdot \left(\prod_{i\in B^-} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

$$\left(\prod_{i\in I_{Bx}^0} \Gamma(\langle A_i\rangle)\right) \cdot \left(\prod_{i\in B^0} \prod_{i=1}^N \Gamma(A_{ij} + B_{ij})\right)^{-1}$$

For any $i \in [0..N]$ we can split $\prod_{i=1}^{N} \Gamma(A_{ij} + B_{ij})$ into

$$\left(\prod_{j \in J_B^+(i)} \Gamma(A_{ij} + B_{ij}) \right) \cdot \left(\prod_{j \in B_i^-} \Gamma(A_{ij} + B_{ij}) \right) \cdot \left(\prod_{j \in B_i^0} \Gamma(A_{ij} + B_{ij}) \right)$$

Now notice,

$$\prod_{j \in B_i^+} \Gamma(A_{ij} + B_{ij}) = \prod_{j \in B_i^+} \Gamma(A_{ij}) \prod_{k=0}^{B_{ij}-1} (A_{ij} + k)$$

$$= \left(\prod_{j \in B_i^+} \Gamma(A_{ij}) \right) \cdot \left(\prod_{j \in B_i^+} \prod_{k=0}^{B_{ij}-1} (A_{ij} + k) \right)$$

and

$$\prod_{j \in B_i^-} \Gamma(A_{ij} + B_{ij}) = \prod_{j \in B_i^-} \Gamma(A_{ij}) \left(\prod_{k=0}^{|B_{ij}|-1} (A_{ij} - k - 1) \right)^{-1}$$

$$= \left(\prod_{j \in B_i^-} \Gamma(A_{ij}) \right) \cdot \left(\prod_{j \in B_i^-} \prod_{k=0}^{|B_{ij}|-1} (A_{ij} - k - 1) \right)^{-1}$$

Hence, $\prod_{i=1}^{N} \Gamma(A_{ij} + B_{ij})$ is equivalent to

$$\left(\prod_{i=1}^{N} \Gamma(A_{ij}) \right) \cdot \left(\prod_{j \in B_i^+} \prod_{k=0}^{B_{ij}-1} (A_{ij} + k) \right) \cdot \left(\prod_{j \in B_i^-} \prod_{k=0}^{|B_{ij}|-1} (A_{ij} - k - 1) \right)^{-1}$$

Combining everything we obtain

$$\left(\prod_{i \in B^+} \Gamma(\langle A_i \rangle) \right) \cdot \mathcal{H}_B^+(A) \cdot \left(\prod_{i \in B^+} \prod_{i=1}^{N} \Gamma(A_{ij}) \right)^{-1} \cdot \frac{\mathcal{H}_B^{+-}(A)}{\mathcal{H}_B^{++}(A)}$$

$$\left(\prod_{i \in B^-} \Gamma(\langle A_i \rangle) \right) \cdot \frac{1}{\mathcal{H}_B^-(A)} \cdot \left(\prod_{i \in B^-} \prod_{i=1}^{N} \Gamma(A_{ij}) \right)^{-1} \cdot \frac{\mathcal{H}_B^{--}(A)}{\mathcal{H}_B^{-+}(A)}$$

$$\left(\prod_{i \in B^0} \Gamma(\langle A_i \rangle) \right) \left(\prod_{i \in B^0} \prod_{i=1}^{N} \Gamma(A_{ij}) \right)^{-1} \cdot \frac{\mathcal{H}_B^{0-}(A)}{\mathcal{H}_B^{0+}(A)}$$

Now by rearranging the products we obtain

$$c(A) \cdot \frac{\mathcal{H}_B^+(A) \cdot \mathcal{H}_B^{+-}(A)}{\mathcal{H}_B^{++}(A)} \cdot \frac{\mathcal{H}_B^{--}(A)}{\mathcal{H}_B^-(A) \cdot \mathcal{H}_B^{-+}(A)} \cdot \frac{\mathcal{H}_B^{0-}(A)}{\mathcal{H}_B^{0+}(A)}$$

The rest follows from the definitions of $\mathcal{H}_B^{\star \circ}(A)$ and $\mathcal{H}_B^{\star}(A)$

For $i, j \in Q$, we define δ^{ij} be a $N \times N$ matrix which is 1 at entry (i, j) and 0 otherwise.

Lemma 3. *Let $A \in \mathbb{N}^{N \times N}$, let $B \in \mathbb{Z}^{N \times N}$, for any $a, b \in [1..N]$*

$$\mathcal{H}_B \left(A + \delta^{ab} \right) = \mathcal{H}_B \left(A \right) \cdot \left(\frac{\langle A_a \rangle + \langle B_a \rangle}{\langle A_a \rangle} \right) \cdot \left(\frac{A_{ab}}{A_{ab} + B_{ab}} \right)$$

Proof. For $\star \in \{+, -\}$ we observe

$$\prod_{i \in B^\star} \langle A_i + \delta_i^{ab} \rangle!^\star_{\langle B_i \rangle} = \left(\prod_{i \in B^\star \backslash \{a\}} \langle A_i \rangle!^\star_{\langle B_i \rangle} \right) \cdot \left((\langle A_a \rangle + 1)!^\star_{\langle B_a \rangle} \right)$$

$$= \left(\prod_{i \in B^\star} \langle A_i \rangle!^\star_{\langle B_i \rangle} \right) \cdot \left(\frac{\langle A_a \rangle + \langle B_a \rangle}{\langle A_a \rangle} \right)^{\star 1}$$

For $\circ \in \{+, -\}$ we observe

$$\prod_{j \in J_B^\circ(a)} (A_{aj} + \delta_{aj}^{ab})!^\circ_{B_{aj}} = \left(\prod_{j \in J_B^\circ(a) \backslash \{b\}} A_{aj}!^\circ_{B_{aj}} \right) \cdot \left((A_{ab} + 1)!^\circ_{B_{ab}} \right)$$

$$= \left(\prod_{j \in J_B^\circ(a)} A_{aj}!^\circ_{B_{aj}} \right) \cdot \left(\frac{A_{ab} + B_{ab}}{A_{ab}} \right)^{\circ 1}$$

Consider the values of $\mathcal{H}_B \left(A + \delta^{ab} \right)$. If $\star = +$. If $\circ = +$ then

$$\frac{\mathcal{H}_B^{++} (A) \cdot \left(\frac{A_{ab} + B_{ab}}{A_{ab}} \right)}{\mathcal{H}_B^+ (A) \cdot \mathcal{H}_B^{+-} (A) \cdot \left(\frac{\langle A_a \rangle + \langle B_a \rangle}{\langle A_a \rangle} \right)} \cdot \frac{\mathcal{H}_B^- (A) \cdot \mathcal{H}_B^{-+} (A)}{\mathcal{H}_B^{--} (A)} \cdot \frac{\mathcal{H}_B^{0+} (A)}{\mathcal{H}_B^{0-} (A)}$$

$$= \mathcal{H}_B \left(A + \delta^{ab} \right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}} \cdot \frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle}$$

If $\circ = -$ then

$$\frac{\mathcal{H}_B^{++} (A)}{\mathcal{H}_B^+ (A) \cdot \mathcal{H}_B^{+-} (A) \cdot \left(\frac{\langle A_a \rangle + \langle B_a \rangle}{\langle A_a \rangle} \right) \cdot \left(\frac{A_{ab}}{A_{ab} + B_{ab}} \right)} \cdot \frac{\mathcal{H}_B^- (A) \cdot \mathcal{H}_B^{-+} (A)}{\mathcal{H}_B^{--} (A)} \cdot \frac{\mathcal{H}_B^{0+} (A)}{\mathcal{H}_B^{0-} (A)}$$

$$= \mathcal{H}_B \left(A + \delta^{ab} \right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}} \cdot \frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle}$$

If $\star = -$. If $\circ = +$ then

$$\frac{\mathcal{H}_B^{++} (A)}{\mathcal{H}_B^+ (A) \cdot \mathcal{H}_B^{+-} (A)} \cdot \frac{\mathcal{H}_B^- (A) \cdot \mathcal{H}_B^{-+} (A) \cdot \left(\frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle} \right) \cdot \left(\frac{A_{ab} + B_{ab}}{A_{ab}} \right)}{\mathcal{H}_B^{--} (A)} \cdot \frac{\mathcal{H}_B^{0+} (A)}{\mathcal{H}_B^{0-} (A)}$$

$$= \mathcal{H}_B \left(A + \delta^{ab} \right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}} \cdot \frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle}$$

If $\circ = -$ then

$$\frac{\mathcal{H}_B^{++}(A)}{\mathcal{H}_B^+(A) \cdot \mathcal{H}_B^{+-}(A)} \cdot \frac{\mathcal{H}_B^-(A) \cdot \mathcal{H}_B^{-+}(A) \cdot \left(\frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle}\right)}{\mathcal{H}_B^{--}(A) \cdot \left(\frac{A_{ab}}{A_{ab} + B_{ab}}\right)} \cdot \frac{\mathcal{H}_B^{0+}(A)}{\mathcal{H}_B^{0-}(A)}$$

$$= \mathcal{H}_B\left(A + \delta^{ab}\right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}} \cdot \frac{\langle A_a \rangle}{\langle A_a \rangle + \langle B_a \rangle}$$

If $\star = 0$. If $\circ = +$ then

$$\frac{\mathcal{H}_B^{++}(A)}{\mathcal{H}_B^+(A) \cdot \mathcal{H}_B^{+-}(A)} \cdot \frac{\mathcal{H}_B^-(A) \cdot \mathcal{H}_B^{-+}(A)}{\mathcal{H}_B^{--}(A)} \cdot \frac{\mathcal{H}_B^{0+}(A) \cdot \left(\frac{A_{ab} + B_{ab}}{A_{ab}}\right)}{\mathcal{H}_B^{0-}(A)}$$

$$= \mathcal{H}_B\left(A + \delta^{ab}\right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}}$$

If $\circ = -$ then

$$\frac{\mathcal{H}_B^{++}(A)}{\mathcal{H}_B^+(A) \cdot \mathcal{H}_B^{+-}(A)} \cdot \frac{\mathcal{H}_B^-(A) \cdot \mathcal{H}_B^{-+}(A)}{\mathcal{H}_B^{--}(A)} \cdot \frac{\mathcal{H}_B^{0+}(A)}{\mathcal{H}_B^{0-}(A)} \cdot \left(\frac{A_{ab}}{A_{ab} + B_{ab}}\right)$$

$$= \mathcal{H}_B\left(A + \delta^{ab}\right) \cdot \frac{A_{ab} + B_{ab}}{A_{ab}}$$

Polynomial Form

Lemma 4. *Any PSE containing only divisions of the form $\frac{1}{v_{ij}}$ can be transformed into a polynomial.*

Proof. Let φ and φ' be two polynomials. Then $\varphi + \psi$ is a polynomial, i.e.

$$\varphi + \varphi' = \sum_{k=1}^p \kappa_k \prod_{i=1,j=1}^N v_{ij}^{d_{ij}^k} + \sum_{l=1}^q \kappa_l' \prod_{i=1,j=1}^N v_{ij}^{d'_{ij}^l}$$

Then $\varphi \cdot \psi$ is a polynomial, i.e.

$$\varphi \cdot \varphi' = \sum_{k=1}^p \kappa_k \prod_{i=1,j=1}^N v_{ij}^{d_{ij}^k} \cdot \sum_{l=0}^q \kappa_l' \prod_{i=1,j=1}^N v_{ij}^{d'_{ij}^l}$$

$$= \sum_{k=1}^p \sum_{l=1}^q \kappa_k \kappa_l' \prod_{i=1,j=1}^N v_{ij}^{d_{ij}^k} \cdot \prod_{i=1,j=1}^N v_{ij}^{d'_{ij}^l}$$

Trivially the leafs, i.e. v_{ij} or $1 \div v_{ij}$ of the formula tree are polynomials. Hence, by starting from the leafs and propagating the transformations upwards we obtain a formula in polynomial form.

Lemma 5. *Let $m \in \mathbb{N}$ s.t. $m \geq 2$, let*

$$\varphi_m := \prod_{i=0}^{m-1} (q_{2i} + q_{2i+1})$$

containing $2m$ unique variables then its polynomial form is of size $2^{2m+1} - 1$.

Proof. For some $m \in \mathbb{N}$, we show by induction that the polynomial form of φ_m is $\sum_{i=0}^{2^m-1} \prod_{j=0}^{m-1} q_{x_{ij}}$ where $x_{ij} \in [0..m]$. First,

$$\varphi_2 = (q_0 + q_1) \cdot (q_2 + q_3) = q_0 q_2 + q_0 q_3 + q_1 q_2 + q_1 q_3 = \sum_{i=0}^{2^2-1} \prod_{j=0}^{2-1} q_{x_{ij}}$$

Second, by IH

$$\varphi_{k+1} = \varphi_k \cdot (q_{2k} + q_{2k+1}) = \left(\sum_{i=0}^{2^k} \prod_{j=0}^{k-1} q_{x_{ij}} \right) \cdot (q_{2k} + q_{2k+1}) =$$

$$= \sum_{i=0}^{2^k-1} q_{2k} \prod_{j=0}^{k-1} q_{x_{ij}} + \sum_{i=0}^{2^k-1} q_{2k+1} \prod_{j=0}^{k-1} q_{x_{ij}} = \sum_{i=0}^{2^k} \prod_{j=0}^{k} q_{x_{ij}}$$

Therefore, the sum consists of $2^m - 1$ additions symbols and 2^m products, with each product containing m variable symbols and $m - 1$ product symbols. Thus we obtain $2^m(m + m - 1) + 2^m - 1 = 2^{m+1}m - 1$.

Lemma 6. *Every PSE φ can be reduced to a PSE ψ that is in polynomial form, such that φ and ψ are semantically equivalent. If the size of φ is n, then the size of ψ is bounded by $\mathcal{O}\left(n2^{\frac{n}{2}}\right)$.*

Proof. Choose m s.t. $4(m-1) - 1 \leq n \leq 4m - 1$. Then $|\varphi_{m-1}| < \varphi \leq |\varphi_m|$ and from Lemma 5 that φ_m's polynomial form is smaller than $2^{m+1}m - 1$.

Soundness of Algorithm 4 To compute the expected value of a PSE φ w.r.t. to a matrix beta distribution we utilise results about the Dirichlet distribution. For some $N \in \mathbb{N}^+$, $a \in (\mathbb{R}^+)^N$ the probability density function of the Dirichlet distribution is defined as

$$p_a(v) := \frac{\Gamma\left(\sum_{i=1}^N a_i\right)}{\prod_{i=1}^N \Gamma(a_i)} \cdot \prod_{j=1}^N m_i^{a_i-1}$$

for any $m \in \Delta(N-1)$. Moreover, let $b \in \mathbb{R}^N$ then we can compute the expectation [50]

$$\mathbb{E}_a \left(\prod_{i=1}^N m_i^{b_i} \right) = \frac{\Gamma\left(\sum_{i=1}^N a_i\right)}{\prod_{i=1}^N \Gamma(a_i)} \cdot \frac{\prod_{i=1}^N \Gamma(a_i + b_i)}{\Gamma\left(\sum_{i=1}^N a_i + b_i\right)}$$

Notice that the probability density function of the matrix beta distribution is the product of N different Dirichlet probability density functions. Hence, this result can easily be extended to matrix beta distributions.

That is, for some set of parameter θ and some matrix $A \in \mathbb{N}^{N \times N}$ we obtain

$$\mathbb{E}_\theta \left(\prod_{i=1}^{N} \prod_{j=1}^{N} M_{ij}^{A_{ij}} \right) = \frac{\mathcal{N}(\theta)}{\mathcal{N}(\theta + A)}$$

by independence. Moreover, notice that if $\theta = \mathbb{1}$, i.e. the uniform prior, we obtain

$$\int_{\Delta(N,N)} \mathcal{L}_M(A) \, dM = \mathbb{E}_1 \left(\prod_{i=1}^{N} \prod_{j=1}^{N} M_{ij}^{A_{ij}} \right) = \frac{1}{\mathcal{N}(A) + \mathbb{1}}$$

Moreover, we extend this result further to allow for limited negative powers.

From now on we use $C = (c_{ij})_{i,j \in Q}$ for the count matrix and $D = (d_{ij})_{i,j \in Q}$ for the exponent matrix of a string $\vec{w} \in Q$.

Lemma 7. *Let θ be matrix beta parameter and let $D\varphi = \kappa_1 \cdot \eta_1$ be a monomial.*

$$\mathbb{E}_\theta \left(\varphi(M) \mid \vec{w} \right) = \kappa_1 \cdot \mathcal{H}_D \left(\theta + C \right)$$

if $\theta + C + D \geq 0$,

Proof. By definition and Bayes' Theorem and linearity we obtain

$$\mathbb{E}_\theta \left(\varphi(M) \mid \vec{w} \right) = \int_{\Delta(N,N)} \varphi(M) \cdot \mathbb{P}_\theta(M \mid \vec{w}) \, dM$$

$$= \int_{\Delta(N,N)} \varphi(M) \cdot \mathbb{P}_\theta(M \mid C) \, dM$$

$$= \int_{\Delta(N,N)} \kappa_1 \cdot \prod_{i,j \in Q} M_{ij}^{D_{ij}} \cdot \mathcal{N}(\theta + C) \cdot \mathcal{L}_M(\theta + C - \mathbb{1}) \, dM$$

$$= \kappa_1 \cdot \mathcal{N}(\theta + C) \cdot \int_{\Delta(N,N)} \mathcal{L}_M(\theta + D + C - \mathbb{1}) \, dM$$

Hence we obtain

$$\mathbb{E}_\theta \left(\varphi(M) \mid \vec{w} \right) = \frac{\mathcal{N}(\theta + C)}{\mathcal{N}(\theta + D + C)}$$

From Lemma 2 we obtain

$$\frac{\mathcal{N}(\theta + C)}{\mathcal{N}(\theta + D + C)} = \frac{\mathcal{N}(\theta + C)}{\mathcal{N}(\theta + C) \cdot \mathcal{H}_D \left(\theta + C \right)^{-1}} = \mathcal{H}_D \left(\theta + C \right)$$

The following Theorem demonstrates that we are able to compute the expected value of a probability property.

Theorem 8. *Let $\vec{w} \in Q^*$, let $\theta \in \mathbb{N}^{N \times N}$ be a matrix beta parameter, and let $\varphi = \sum_{k=1}^{p} \kappa_k \eta_k$ be a probability expression in polynomial form. Then*

$$\mathbb{E}_\theta \left(\varphi(M) \mid \vec{w} \right) = \sum_{k=1}^{p} \kappa_k \mathcal{H}_{D^k} \left(\theta + C \right)$$

if $\forall k \in [1..p].\theta + C + D^k \geq 0$.

Proof. This follows directly from Lemma 7 and the linearity of expectation.

Proposition 1 (Confidence interval.). *Given the problem instance (Q, φ, δ) and the parameter matrix θ for the prior distribution, and given a sequence of states $\vec{w} \in Q^*$, the following is a $(1 - \delta)\%$ confidence interval of $\varphi(M)$:*

$$\left[\mathbb{E}_\theta(\varphi(M) \mid \vec{w}) \pm \sqrt{\frac{\mathbb{E}_\theta(\varphi^2(M) \mid \vec{w}) - \mathbb{E}_\theta(\varphi(M) \mid \vec{w})^2}{\delta}} \right]. \tag{11}$$

Proof. A direct consequence of Chebyshev's inequality (Thm. 7).

Resource and Time bounds of Algorithm 3

Lemma 8. *Let φ be a PSE in polynomial form containing p monomials. Then Algorithm 3 requires less than $V_\varphi + Dom(V_\varphi) + 2p$ counters.*

Proof. We need a counter for each variable that occurs in φ they correspond to c_{ij}. We need a counter for the number of visit to a particular state, this corresponds to c_i, the number of which is smaller than $|Dom(V_\varphi)|$. And we need 2 counter to store each h^l for $l \in [1..p]$ resulting in $2p$ counters.

Lemma 9. *Let φ be a PSE in polynomial form containing p monomials. Then Algorithm 3 requires updates its verdict in $\mathcal{O}(p)$ time.*

Proof. We need to update h^l for each $l \in [1..l]$. To do so at most two additions, two multiplications and two divisions are required. To compute E we need to multiply each h_l with some constant κ_l and sum them up. Resulting in $\mathcal{O}(p)$.