The Annals of Applied Probability
2023, Vol. 33, No. 5, 3987–4007
https://doi.org/10.1214/22-AAP1913

MULTIPARAMETER BERNOULLI FACTORIES

BY RENATO PAES LEME[a] AND JON SCHNEIDER[b]

Google Research, NYC, [a]*renatoppl@google.com,* [b]*jschnei@google.com*

We consider the problem of computing with many coins of unknown bias. We are given access to samples of n coins with *unknown* biases $p_1, \dots, p_n$ and are asked to sample from a coin with bias $f(p_1, \dots, p_n)$ for a given function $f : [0, 1]^n \to [0, 1]$. We give a complete characterization of the functions f for which this is possible. As a consequence, we show how to extend various combinatorial sampling procedures (most notably, the classic Sampford sampling for k-subsets) to the boundary of the hypercube.

1. Introduction. The Bernoulli factory problem was formally introduced by Keane and O'Brien (1994), inspired by earlier work by Von Neumann (1951) and Asmussen, Glynn and Thorisson (1992). While their initial goal was to design methods to exactly simulate certain stochastic processes, this tool later found applications in many different fields such as mechanism design (Dughmi et al. (2017), Cai et al. (2021)), quantum physics (Dale, Jennings and Rudolph (2015), Yuan et al. (2016)), Markov chain Monte Carlo (MCMC) methods (Flegal and Herbei (2012)), and exact Bayesian inference (Gonçalves, Łatuszyński and Roberts (2017), Herbei and Berliner (2014)).

The original problem can be best described as manufacturing new (random) coins from old ones. One is given a Bernoulli variable of *unknown* bias p, or for short, a p-coin. Even though we do not know the bias, we can flip the coin as many times as we need obtaining i.i.d. samples from it. The goal is to produce an $f(p)$-coin for a given function $f : [0, 1] \to [0, 1]$.

To give an example, consider $f(p) = p^2 - p^3$. A simple way to sample from an $f(p)$-coin is to flip the p-coin three times obtaining samples $X_1, X_2, X_3 \in \{0, 1\}$. Now we output 1 if $X_1 = X_2 = 1$ and $X_3 = 0$. The probability of outputting 1 is $p^2(1 - p) = f(p)$.

Keane and O'Brien (1994) gave necessary and sufficient conditions for a function $f : [0, 1] \to [0, 1]$ to be implementable. The first condition is that the function f must be continuous. The second condition says that either f is the constant function $f(x) = 0$, the constant function $f(x) = 1$, or there exists some integer $m > 0$ such that for all $p \in [0, 1]$:

$$\min(p, 1 - p)^m \le f(p) \le 1 - \min(p, 1 - p)^m. \tag{1}$$

Furthermore, their proof is algorithmic: given a function satisfying the conditions above, they give a procedure for sampling from $f(p)$.

1.1. *Why exact sampling?* An important aspect of Bernoulli factories is that they ask for *exact* sampling. The original motivation for the Bernoulli factory problem was to perform exact simulations of stochastic processes. In these simulations, small sampling errors quickly compound, sometimes exponentially—hence the need for exact sampling. A similar situation arises in Bayesian inference, where sampling is a sub-routine in an iterative procedure.

Finally, in mechanism design the fact that sampling is exact allows us to design black-box-reductions that are Bayesian-incentive compatible. Before the introduction of this machinery, the known reduction in the general case was ϵ-Bayesian-incentive-compatible, that is, agents

Received February 2022; revised July 2022.
MSC2020 subject classifications. 60-08.
Key words and phrases. Bernoulli factory, exact sampling.

had still a small incentive to deviate from truth-telling, which results in a much weaker game-theoretical guarantee.

This discussion is to motivate why in certain situations *approximately* simulating an $f(p)$-coin is not enough. Approximately sampling can be easily done by the following method: let $X_1, \dots, X_t$ be t draws from the p-coin and define its empirical average as $\bar{X}_t = (X_1 + \cdots + X_t)/t$. We know by the Chernoff bound that for $n = \Omega(\epsilon^{-2} \log(1/\delta))$ we have $\mathbb{P}[|p - \bar{X}_t| > \epsilon] < \delta$. Hence if f is continuous, estimating p by $\bar{X}_t$ and using external randomness to sample from a $f(\bar{X}_t)$-coin produces a reasonable approximation of the $f(p)$-coin.

1.2. *Multiparameter factories.* In this paper we study the multiparameter version of this problem: given a compact set $K \subseteq [0, 1]^n$ and n coins with *unknown* biases $(p_1, \dots, p_n) \in K$, how to sample from a $f(p_1, \dots, p_n)$-coin for a multivariate function $f : K \to [0, 1]$. The sampling procedure is represented by a possibly infinite tree with coins in the internal nodes (either one of the n coins of unknown bias or an auxiliary coin with known bias) and outcomes in the leaves:

DEFINITION 1.1. A Bernoulli factory $\mathcal{F}$ with input $(p_1, \dots, p_n)$ corresponds to a (possibly infinite) rooted binary tree $\mathcal{T}$ where each node of $\mathcal{T}$ has either 2 children (an internal node) or 0 (a leaf). While the tree is allowed to be infinite, each node has finite depth. Each internal node is labelled either with a variable p_i or with a constant $c \in (0, 1)$.

To execute the factory with coins $(p_1, \dots, p_n)$ we start from the root and at each node we flip the coin given the label of that node (either one of the p_i-coins of unknown bias or a c-coin of known bias). Based on the outcome, we either take the left edge (0) or the right edge (1). If we reach a leaf, we output its label. A factory $\mathcal{F}$ is valid if for any input $p = (p_1, \dots, p_n) \in [0, 1]^n$ it reaches a leaf almost surely. A factory $\mathcal{F}$ is valid for a set $K \subseteq [0, 1]^n$ if it is valid for all inputs $p \in K$.

In Figure 1 we give an example of a Bernoulli factory that given coins with biases p_1 and p_2, samples from a $f(p)$-coin for $f(p) = \frac{1}{3}(1 - p_1) + p_1(1 - p_2)$. We will assume throughout the paper that we have access to auxiliary coins with known biases.[1]

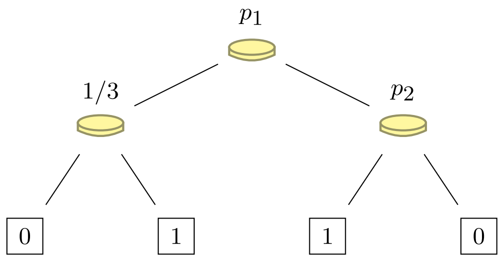

FIG. 1. *Example of a Bernoulli factory with input coins* p_1, p_2 *samples from a coin with bias* $\frac{1}{3}(1 - p_1) + p_1(1 - p_2)$. *Note that the factory uses an auxiliary coin with known bias* 1/3.

[1] The assumption that we have access to auxiliary coins of known bias can be removed in certain cases. Keane and O'Brien (1994) observe that given access to a p-coin of unknown bias $p \in (0, 1)$, we can always sample from any c-coin for $c \in (0, 1)$ using the procedure by Von Neumann (1951). For example, to sample a 1/2-coin using a p-coin with unknown bias $p \in (0, 1)$ we obtain two samples of the p-coin. If the samples are 01 we output 0, if the samples are 10 we output 1. In the remaining cases (00 and 11) we re-try. It is not difficult to see this finishes almost surely and outputs 0 and 1 with equal probabilities. This can be extended to sample to any real number $c \in [0, 1]$. This procedure does not terminate if $p \in \{0, 1\}$. Since in our case we are particularly concerned with the boundary of the hypercube, it is more convenient to simply assume we have access to coins of known bias.

Here we investigate how to design factories that terminate almost surely for every $p \in K$ and establish necessary and sufficient conditions for implementability. To describe our conditions, we need a few definitions. We will use $[n]$ to denote $\{1, 2, \dots, n\}$. Given a vector $p \in [0, 1]^n$ and $S \subseteq [n]$ we will denote

$$p^S = \prod_{i \in S} p_i \quad \text{and} \quad (1-p)^S = \prod_{i \in S} (1 - p_i).$$

Given a function $f : [0, 1]^n \to [0, 1]$, subset $S \subseteq [0, 1]^n$ and a constant $c \in [0, 1]$ we will write $f|_S \equiv c$ to denote that $f(p) = c$ for all $p \in S$. If we write $f \equiv c$ (omitting S) it means f is the constant function on the entire hypercube.

Given a partition of $[n]$ into three sets A, S and B, we define the open face $F_{A,S,B}$ of the hypercube $[0, 1]^n$ as

$$F_{A,S,B} := \left\{ p \in [0,1]^n; \begin{array}{rl} p_i = 0, & i \in A \\ 0 < p_i < 1, & i \in S \\ p_i = 1, & i \in B \end{array} \right\}.$$

Now we are ready to define the notion of a polynomially bounded function:

DEFINITION 1.2 (Polynomially bounded function). A function $f : [0, 1]^n \to [0, 1]$ is polynomially bounded if there is an integer $m \geq 0$ and a real constant $c > 0$ such that for each open face $F_{A,S,B}$ of the hypercube the following condition holds:

$$f|_{F_{A,S,B}} \not\equiv 0 \quad \Rightarrow \quad f(p) \geq c\big((1-p)^A \cdot p^S (1-p)^S \cdot p^B\big)^m \quad \forall p \in [0, 1]^n. \tag{2}$$

Now we are ready to state the main theorem for the hypercube (later in Section 5 we generalize this theorem to any compact subdomain):

THEOREM 1.3. *A function $f : [0, 1]^n \to [0, 1]$ can be implemented by a Bernoulli factory if and only if it is continuous and both f and $1 - f$ are polynomially bounded.*

In summary, the inequality condition of Keane and O'Brien (1994) for one variable become combinatorial for multiple parameters: one imposed by each open face of the hypercube. If the function is nonzero at an open face, it must be lower bounded by a polynomial associated with that face. Similarly, if the function is nonone at a face, it must be upper bounded by a polynomial associated with that face.

We show that these conditions also turn out to be sufficient and exhibit an algorithm to sample from it. The difficulty of designing such algorithm is to make sure it works near the faces of the hypercube and dealing with the interaction of multiple combinatorial constraints when the faces meet. The heart of the proof is a new concentration argument in Section 4.3. We study a random vector $\bar{X}_t$ where each coordinate is an average of Bernoulli variables drawn from the coins of unknown bias $p = (p_1, \dots, p_n)$. We relate the probability of a large deviation in a subset T of the coordinates to the combinatorial structure of the hypercube, in particular to the polynomials associated with the faces where coordinates in T are free. By doing so, we can bound the probability that $f(\bar{X}_t) \geq 1/2$ in terms of $f(p)$ in a way which holds *uniformly* everywhere within the hypercube (even on the boundary).

1.3. *The case of Sampford sampling.* A particularly curious Bernoulli factory is the procedure due to Sampford (1967). Sampford sampling actually predates the notion of a Bernoulli factory by 25 years and is commonly used throughout the statistics literature for sampling k-subsets with "unequal probabilities of selection." Formally, the problem is the following: given probabilities $(p_1, \dots, p_n)$ such that $\sum_i p_i = k$, sample a subset S of size k such that $\mathbb{P}[i \in S] = p_i$. Sampford's solution just requires sample access to the coins.

A natural but incorrect solution is the following: sample each coin i once and let $X_i \in \{0, 1\}$ be the outcome. Output $S = \{i \in [n]; X_i = 1\}$ if $|S| = k$. If not, retry. To see that this does not work, execute this procedure with coins with biases $(1 - \epsilon, 1 - \epsilon, 2\epsilon)$ and observe that the last element is chosen with $O(\epsilon^2)$ probability. There is a simple (but ingenious) fix to this algorithm: first we obtain S as before (retrying until $|S| = k$). We then choose a coin in $[n] \setminus S$ uniformly at random and flip it again. If this coin comes up 1, we output S. If not, we resample S and try again.

For the previous procedure to terminate, we need at least one coin with $0 < p_i < 1$, since one of the coins that came up 0 initially must be re-flipped and needs to come up 1. If all the coins are deterministic, for example, $p = (1, 1, 0)$, the above procedure never terminates.

A natural open question is whether there exists an alternative Bernoulli factory for this problem that terminates for every input in $p \in K := \{p \in [0, 1]^n; \sum_i p_i = k\}$. Niazadeh, Paes Leme and Schneider (2021) show the following negative result: there is no *exponentially converging* factory for k-subset that terminates for all $p \in K$. A factory is said to be exponentially converging if for every p, there is a rate $r(p)$ such that the probability that the procedure has not terminated after flipping t coins is at most $r(p)^t$. Note that Sampford sampling is exponentially converging for every $p \in K \cap (0, 1)^n$.

The negative result by Niazadeh, Paes Leme and Schneider (2021) excludes techniques based on Bernstein-rational functions, which are the only known ideas for designing factories that terminate a.s. at the boundary but all lead to exponentially converging factories.

Despite this negative evidence, our new algorithm produces a factory for k-subset that terminates a.s. for every $p \in K$. More generally, it solves a wider class of problems introduced in Niazadeh, Paes Leme and Schneider (2021): given a polytope $\mathcal{P} \subseteq [0, 1]^n$ and n coins with biases $p = (p_1, \dots, p_n) \in \mathcal{P}$, sample a random vertex v of $\mathcal{P}$ such that $\mathbb{E}[v] = p$. Niazadeh, Paes Leme and Schneider (2021) show that is is possible to construct factories for $\mathcal{P} \cap (0, 1)^n$ only when $\mathcal{P}$ is the intersection of the hypercube $[0, 1]^n$ and an affine subspace.

These factories, however, suffer from the same problem as Sampford sampling: they diverge for certain points in the boundary of the hypercube. Using the techniques developed in this paper, we exhibit alternative factories that terminate a.s. for all points in $\mathcal{P}$. For example, consider the matching polytope: we are given coins p_{ij} forming a doubly stochastic matrix and asked to sample a matching M in the complete bipartite graph such that $\mathbb{P}[(i, j) \in M] = p_{ij}$. The previous factory required $p_{ij} > 0$ for all edges (i, j). The alternative factories constructed in this paper no longer have this restriction.

1.4. *Previous results on multiparameter factories.*

Previous approaches to this problem either are restricted to rational functions (Mossel and Peres (2005), Morina et al. (2022) and Niazadeh, Paes Leme and Schneider (2021)) or assume that the vector of coins $(p_1, \dots, p_n)$ lies away from the boundary of the hypercube (Nacu and Peres (2005) and Morina (2021)). Mossel and Peres (2005) and Morina et al. (2022) show how to design factories for Bernstein-rational functions, that is, rational functions of the type $f(p) = a(p)/[a(p) + b(p)]$ where $a(p)$ and $b(p)$ are of the form $\sum_i c_i \prod_{j \in [n]} p_i^{a_{ij}} (1 - p_i)^{b_{ij}}$ with $c_i > 0$. Recently, Niazadeh, Paes Leme and Schneider (2021) showed how to design factories for certain combinatorial objects (such as matchings and flows) using Bernstein-rational functions.

Beyond rational functions, Nacu and Peres (2005) give a procedure for sampling from any continuous function $f : [\epsilon, 1 - \epsilon]^n \to (0, 1)$. Their procedure is based on Bernstein's proof the Weierstrass approximation theorem. Their result is originally written for $n = 1$ but there is nothing particular about one dimension in their construction.

Using a very clever idea, Morina (2021) shows in Chapter 3 of his PhD thesis how to combine factories defined in $[\epsilon, 1 - \epsilon]^n$ for decreasing values of ϵ into a single factory defined on the interior of the hypercube $(0,1)^n$. The condition required for the factories to be combined is a single inequality which is a generalization of the condition of Keane and O'Brien (1994). The result is stated for the open simplex $\Delta_n^0 = \{p \in (0,1)^n; \sum_{i=1}^n p_i = 1\}$. It shows that a function $f : \Delta_n^0 \to (0,1)$ is implementable by a Bernoulli factory if and only if it is continuous and for some integer m it holds that

$$\Bigl(\prod_i p_i\Bigr)^m < f(p) < 1 - \Bigl(\prod_i p_i\Bigr)^m \quad \forall p \in \Delta_n^0.$$

A restriction of Morina's factory is that it diverges by construction at the boundary and its running time blows up when we get close to it. The first step in its construction (Lemma 3.13 in Morina (2021)) is to keep sampling all of the coins until each coin comes up 1 at least $\Omega(mn)$ times and 0 at least $\Omega(mn)$ times. At a high level, it uses the coin outcomes to pick a value of ϵ and then it uses a factory for $[\epsilon, 1-\epsilon]^n$ domain. The process never terminates if the input has coins with $p_i \in \{0,1\}$.

2. Preliminaries.

2.1. *Multiparameter factory.* A multiparameter Bernoulli factory (following the model of Niazadeh, Paes Leme and Schneider (2021)) is defined in the introduction (Definition 1.1). From that definition, a valid factory can be seen as a random variable $\mathcal{F}$ taking values in $\{0,1\}$. The distribution of $\mathcal{F}$ will naturally depend on the input coins p. For that reason it is convenient to use the notation $\mathbb{P}_p[\cdot]$ and $\mathbb{E}_p[\cdot]$ to denote the probability measure and expectation of random variables when $(p_1, \dots, p_n)$ coins are used.

We will say that a factory is finite if the tree $\mathcal{T}$ contains finitely many nodes (and thus, the factory is guaranteed to terminate after a finite number of coin flips).

2.2. *Concentration bounds.* We will use $X_t \in \{0,1\}^n$ to denote a random vector distributed according to the product Bernoulli distribution $\mathrm{Ber}(p_1) \times \cdots \times \mathrm{Ber}(p_n)$ corresponding to the input coin flips. For each $t = 1, 2, \dots$ and $i \in [n]$, the variable $X_{t,i}$ is an independent Bernoulli variable with bias p_i. We will also let $\bar{X}_t$ be a random variable equal to the average of the first t flips of all n coins:

$$\bar{X}_t := \frac{X_1 + \cdots + X_t}{t} \in [0,1]^n. \tag{3}$$

We will write $\bar{X}_{t,i}$ to denote the ith component of $\bar{X}_t$. Next, we state two well-known concentration bounds. The first is the Hoeffding bound:

$$\mathbb{P}_p\bigl[|\bar{X}_{t,i} - p_i| > \delta\bigr] \le 2\exp(-2\delta^2 t). \tag{4}$$

The second is the sharper Chernoff bound:

$$\mathbb{P}_p[\bar{X}_{t,i} - p_i > \delta] \le \left(\left(\frac{p_i}{p_i+\delta}\right)^{p_i+\delta}\left(\frac{1-p_i}{1-p_i-\delta}\right)^{1-p_i+\delta}\right)^t. \tag{5}$$

For values of p_i that are closer to zero (say $p_i < 1/2$ and $\delta < 1/4$) we can bound the second term by a constant. Hence for such p_i we can write

$$\mathbb{P}_p[\bar{X}_{t,i} - p_i > \delta] \le \bigl(c_\delta \cdot p_i^\delta\bigr)^t, \tag{6}$$

where c_δ is some constant depending on δ.

2.3. *Real topology.* Here we recall some elementary facts and definitions from real topology. For $p \in \mathbb{R}^n$ and a real $r > 0$, we denote the ℓ_∞-ball around p of radius r by

$$\mathsf{B}_\infty(p; r) := \{x \in \mathbb{R}^n; \|x - p\|_\infty < r\}.$$

Given any set $S \subseteq \mathbb{R}^n$ we denote

$$\mathsf{B}_\infty(S; r) := \bigcup_{p \in S} \mathsf{B}_\infty(p; r) = \{x \in \mathbb{R}^n; \exists p \in S \text{ s.t. } \|x - p\|_\infty < r\}.$$

A set $U \subset \mathbb{R}^n$ is open if for every $p \in U$ there is an $r > 0$ such that $\mathsf{B}_\infty(p; r) \subseteq U$. An *open cover* of a set S consists of a collection of open sets U_i for $i \in I$ such that $\bigcup_{i \in I} U_i \supseteq S$. The index set I is possibly infinite and uncountable. We say that the cover $\bigcup_{i \in I} U_i \supseteq S$ admits a *finite subcover* if there is a finite set $I_0 \subset I$ such that $\bigcup_{i \in I_0} U_i \supseteq S$.

We will make extensive use of the following elementary fact.

LEMMA 2.1 (Heine–Borel). *A set K is a compact set (i.e., every cover admits a finite subcover) iff it is topologically closed and bounded.*

2.4. *Decomposing the hypercube.* We recall the decomposition of the hypercube $[0, 1]^n$ into 3^n disjoint regions that we will refer as open faces. Each open face will correspond to a partition of $[n]$ into three sets A, S and B. We define the open face $F_{A,S,B}$ as

$$F_{A,S,B} := \left\{ p \in [0,1]^n; \begin{array}{rl} p_i = 0, & i \in A \\ 0 < p_i < 1, & i \in S \\ p_i = 1, & i \in B \end{array} \right\}.$$

For example, the square $[0, 1]^2$ is the union of nine disjoint open faces: the interior $(0, 1)^2$, the four edges $\{0\} \times (0, 1)$, $\{1\} \times (0, 1)$, $(0, 1) \times \{0\}$, $(0, 1) \times \{1\}$ and the four vertices $\{(0, 0)\}$, $\{(0, 1)\}$, $\{(1, 0)\}$, $\{(1, 1)\}$.

We will denote by $\bar{F}_{A,S,B}$ the topological closure of $F_{A,S,B}$ which can be written as

$$\bar{F}_{A,S,B} = \bigcup_{A' \supseteq A, B' \supseteq B} F_{A',S',B'}.$$

For example, the closure of the open face $\{0\} \times (0, 1)$ of the square is: $\{0\} \times [0, 1]$ which is the union of three open faces: one representing that edge and two representing the vertices.

3. Main theorem.

Our main result is to identify a condition called polynomially boundedness (Definition 1.2) which together with continuity is necessary and sufficient for the existence of a factory. In this section we prove our main result (Theorem 1.3).

3.1. *Sanity check.* It is useful to check that when we set $n = 1$ we recover the result by Keane and O'Brien. For $n = 1$ there are three open faces: $\{0\}$, $\{1\}$ and $(0, 1)$. For $(0, 1)$ the condition that f is polynomially bounded means that

$$f \not\equiv 0 \quad \Rightarrow \quad f(p) \geq c\bigl(p(1 - p)\bigr)^m.$$

Observe that $\min(p, 1 - p) \leq \frac{1}{2}$, so if we take $k = \lceil \log_2 c \rceil$, then $cp(1 - p)^m \geq \min(p, 1 - p)^{k+2m}$. Similarly the condition that $1 - f$ is polynomially bounded implies that

$$f \not\equiv 1 \quad \Rightarrow \quad f(p) \leq 1 - c\bigl(p(1 - p)\bigr)^m.$$

Hence if $f \not\equiv 0$ and $f \not\equiv 1$ then

$$\min(p, 1 - p)^{k+2m} \leq f(p) \leq 1 - \min(p, 1 - p)^{k+2m},$$

which is the one-dimensional condition (1). Notice that the conditions for the open faces $\{0\}$ and $\{1\}$ are superfluous here. For example, for $\{0\}$ our condition says that

$$f(0) > 0 \quad \Rightarrow \quad f(p) \geq c(1-p).$$

Note that this is implied by continuity in a neighborhood of 0 and by the condition (1) elsewhere.

3.2. *Necessary conditions.* The next lemmas show that every function f that is implementable by a factory has f and $1-f$ polynomially bounded.

LEMMA 3.1. *If a function $f : [0,1]^n \to [0,1]$ can be implemented by a Bernoulli factory and for some open face $F_{A,S,B}$ we have $f|_{F_{A,S,B}} \not\equiv 0$ then there exists an integer m and a constant c such that $f(p) \geq c((1-p)^A \cdot p^S(1-p)^S \cdot p^B)^m$, $\forall p \in [0,1]^n$.*

PROOF. Consider a rooted binary tree (as in Definition 1.1) implementing f. Since $f(p) > 0$ for some coins in $p \in F_{A,S,B}$, there must be a path in the tree reaching a leaf labelled 1 that always takes the 0-edge when we flip a coin with an index in A and always takes the 1-edge when we flip a coin with an index in B, or else this path would never be reachable using the coins p. Since each leaf is at a finite depth, this path is finite. Each path in the tree corresponds to a polynomial of the form $c \cdot \prod_{i\in[n]} p_i^{g_i}(1-p_i)^{h_i}$ (a "Bernstein monomial"), where c is the product of the helper coins flipped along the path, g_i is the number of 1-edges (right) takes after a p_i-flip and h_i is the number of 0-edges (left) takes after a p_i-flip. By the observation above $g_i = 0$ for $i \in A$ and $h_i = 0$ for $i \in B$. Taking $m = \max_i \max(g_i, h_i)$ we obtain the inequality in the statement of the lemma. □

LEMMA 3.2. *If a function $f : [0,1]^n \to [0,1]$ can be implemented by a Bernoulli factory and for some open face $F_{A,S,B}$ we have $f|_{F_{A,S,B}} \not\equiv 1$ then there exists an integer m and a constant c such that $1 - f(p) \geq c((1-p)^A \cdot p^S(1-p)^S \cdot p^B)^m$, $\forall p \in [0,1]^n$.*

PROOF. Same proof as the previous lemma swapping 0 and 1. □

Note that Lemmas 3.1 and 3.2 show constants c_F, m_F for each face F, while Definition 1.2 asks for constants c, m that hold uniformly over the faces. To translate from the face-dependent constants to uniform constants, observe that the expression $c((1-p)^A p^S (1-p)^S p^B)^m$ is monotone nonincreasing in c and monotone nondecreasing in m since $(1-p)^A p^S (1-p)^S p^B \in [0,1]$. Since there are finitely many faces, it is enough to take uniform constants $c = \min_F c_F$ and $m = \max_F m_F$.

The continuity condition is more intuitive: if two vectors of coins $p', p'' \in [0,1]^n$ are close, the finite sequence of coin flips generated by them will also be close in total variation distance. Since the output only depends on the sequence of coins flips observed, the distribution of outputs must also be close. This intuition is formalized by the following lemma.

LEMMA 3.3. *If a function $f : [0,1]^n \to [0,1]$ can be implemented by a Bernoulli factory then it is continuous.*

PROOF. Consider an implementation of f by a Bernoulli factory $\mathcal{F}$ and fix a point $a \in [0,1]^n$ in the domain of f. We want to show that for every $\epsilon > 0$, there is δ such that if $\|p - a\| < \delta$ then $|f(p) - f(a)| < \epsilon$.

To show that, let T be a random variable showing the number of coins flipped before the output if the decision tree is executed using an a-coin (this is equal to the depth of the output node reached in the tree). Note that all probabilities of random events depend on both the factory $\mathcal{F}$ and the vector of coins used. Since we will use the same factory $\mathcal{F}$ throughout the proof, we will omit it from the subscripts, but we will keep using $\mathbb{P}_a[\cdot]$ to denote which vector of coins is used.

Now fix t such that $\mathbb{P}_a[T > t] < \epsilon/4$. Represent a possible realization of the first t coin flips of each coin by a tuple $x = (x_1, \dots, x_t)$ for $x_i \in \{0,1\}^n$ we define function $F(x) \in \{0, 1, \varnothing\}$ indicating whether the decision tree outputs 0, 1 or does not yet terminate after seeing inputs $x_1, \dots, x_t$. Also, let $X = (X_1, \dots, X_t) \in \{0,1\}^{nt}$ be the random output of the coins. With that, we can rewrite $\mathbb{P}_a[T > t] < \epsilon/4$ as

$$\sum_{x \in \{0,1\}^{nt}; F(x) = \varnothing} \mathbb{P}_a[X = x] \le \frac{\epsilon}{4}. \tag{7}$$

Now choose δ small enough such that the total variation distance between the sequences $X = (X_1, \dots, X_t)$ generated under p and a is at most $\epsilon/3$ for any $\|p - a\| < \delta$. More formally,

$$\sum_{x \in \{0,1\}^{nt}} \left|\mathbb{P}_a[X = x] - \mathbb{P}_p[X = x]\right| < \frac{\epsilon}{4} \quad \forall p \in \mathsf{B}_\infty(a, \delta). \tag{8}$$

Now we can bound $f(a)$ and $f(p)$ for $\|p - a\| < \delta$ as follows:

$$\left| f(a) - \sum_{x \in \{0,1\}^{nt}; F(x) \in \{0,1\}} F(x)\mathbb{P}_a[X = x] \right| \le \sum_{x \in \{0,1\}^{nt}; F(x) = \varnothing} \mathbb{P}_a[X = x] < \frac{\epsilon}{4},$$

and similarly:

$$\left| f(p) - \sum_{x \in \{0,1\}^{nt}; F(x) \in \{0,1\}} F(x)\mathbb{P}_p[X = x] \right| \le \sum_{x \in \{0,1\}^{nt}; F(x) = \varnothing} \mathbb{P}_p[X = x] < \frac{\epsilon}{2},$$

where the last bound follows from combining equations (7) and (8). Now, taking it all together, we have

$$|f(a) - f(p)| \le \left| \sum_{x \in \{0,1\}^{nt}; F(x) \in \{0,1\}} F(x)\big(\mathbb{P}_a[X = x] - \mathbb{P}_p[X = x]\big) \right| + \frac{3\epsilon}{4} < \frac{\epsilon}{4} + \frac{3\epsilon}{4} = \epsilon. \qquad \square$$

3.3. *Sufficient conditions.* To prove that continuous and polynomially bounded are sufficient conditions for implementability, we will use the following lemma (Lemma 3.4), which is the main technical lemma of the paper. We will prove it in the next section. Before we do this, however, we will assume it is true and use it to prove Theorem 1.3. Recall the definition of $\bar{X}_t$ in equation (3).

LEMMA 3.4. *Let f be a continuous and polynomially bounded function. Then there is an integer t_0 such that for $t \ge t_0$ it holds that*

$$f(p) - \frac{1}{4} \cdot \mathbb{P}_p\left[f(\bar{X}_t) \ge \frac{1}{2} \right] \ge \frac{1}{8} f(p) \quad \forall p \in [0,1]^n. \tag{9}$$

Lemma 3.4 will allow us to decompose $f(p)$ into two smaller functions: one which we can simulate with a finite Bernoulli factory (this will be $\mathbb{P}_p[f(\bar{X}_t) \ge \frac{1}{2}]$), and a remaining piece with probability mass at most 3/4 that we can decompose recursively.

LEMMA 3.5. *Let $f : [0,1]^n \to [0,1]$ be a continuous function such that f and $1-f$ are polynomially bounded. Then there exists a function $g : [0,1]^n \to [0,1]$ that is implementable by a finite Bernoulli factory such that $\tilde{f}$ defined as follows*:

$$\tilde{f}(p) = \frac{4}{3}\Big(f(p) - \frac{1}{4}g(p)\Big)$$

is such that $\tilde{f}([0,1]^n) \subseteq [0,1]$, $\tilde{f}$ is continuous and $\tilde{f}$ and $1-\tilde{f}$ are polynomially bounded.

PROOF. We start by applying Lemma 3.4 to both f and $1-f$. We know that there are integers t_0 and t_1 such that

$$f(p) - \frac{1}{4}\cdot\mathbb{P}_p\Big[f(\bar{X}_t) \ge \frac{1}{2}\Big] \ge \frac{1}{8}f(p) \quad \forall t \ge t_0, p \in [0,1]^n,$$

$$1 - f(p) - \frac{1}{4}\cdot\mathbb{P}_p\Big[1 - f(\bar{X}_t) \ge \frac{1}{2}\Big] \ge \frac{1}{8}[1-f(p)] \quad \forall t \ge t_1, p \in [0,1]^n.$$

Note that we can rewrite

$$\mathbb{P}_p\Big[1-f(\bar{X}_t) \ge \frac{1}{2}\Big] = \mathbb{P}_p\Big[f(\bar{X}_t) \le \frac{1}{2}\Big] = 1 - \mathbb{P}_p\Big[f(\bar{X}_t) > \frac{1}{2}\Big] \ge 1 - \mathbb{P}_p\Big[f(\bar{X}_t) \ge \frac{1}{2}\Big].$$

Replacing it in the previous expression we obtain:

$$f(p) - \frac{1}{4}\mathbb{P}_p\Big[f(\bar{X}_t) \ge \frac{1}{2}\Big] \le \frac{3}{4} - \frac{1}{8}[1-f(p)] \quad \forall t \ge t_1, p \in [0,1]^n.$$

Now if we set $t = \max(t_0, t_1)$ and define a function $g : [0,1]^n \to [0,1]$ as

$$g(p) = \mathbb{P}_p\Big[f(\bar{X}_t) \ge \frac{1}{2}\Big]$$

then we have that

$$\frac{1}{6}f(p) \le \frac{4}{3}\Big(f(p) - \frac{1}{4}g(p)\Big) \le 1 - \frac{1}{6}[1-f(p)]. \tag{10}$$

First observe that (10) implies that $\tilde{f}(p) \in [0,1]$.

Now let's argue that $\tilde{f}$ is polynomially bounded. First observe that if $f(p) = 0$ at some point p then we must have $\tilde{f}(p) = 0$ since the first inequality implies that $0 \le -\frac{1}{4}g(p)$. Since $g(p) \ge 0$ we must have $g(p) = 0$ and hence $\tilde{f}(p) = \frac{4}{3}f(p) = 0$. Conversely, if $\tilde{f}(p) = 0$, then by the first inequality of equation (10) it also holds that $f(p) = 0$.

Therefore, for any open face $F_{A,S,B}$ we have $f|_{F_{A,S,B}} \equiv 0$ iff $\tilde{f}|_{F_{A,S,B}} \equiv 0$. If $\tilde{f}|_{F_{A,S,B}} \not\equiv 0$ then by the first inequality in (10) and the fact that f is polynomially bounded, we have

$$\tilde{f}(p) \ge \frac{1}{6}f(p) \ge \frac{c}{6}((1-p)^A p^S (1-p)^S p^B)^m.$$

The same argument can be repeated with $1-f$ instead of f to argue this function is also polynomially bounded. First observe that if $f(p) = 1$ at some point p we must have $\tilde{f}(p) = 1$ since the last inequality would imply that $1 - \frac{1}{4}g(p) \le \frac{3}{4}$. Since $g(p) \le 1$ this must imply that $g(p) = 1$ and hence $\tilde{f}(p) = \frac{4}{3}(1 - \frac{1}{4}) = 1$. Therefore $f|_{F_{A,S,B}} \equiv 1$ iff $\tilde{f}|_{F_{A,S,B}} \equiv 1$. Conversely, if $\tilde{f}(p) = 1$, then by the second inequality of equation (10) it also holds that $f(p) = 1$.

If $\tilde{f}|_{F_{A,S,B}} \not\equiv 1$ then by the second inequality in (10) and the fact that $1-f$ is polynomially bounded, we have

$$1 - \tilde{f}(p) \ge \frac{1}{6}[1-f(p)] \ge \frac{c}{6}((1-p)^A p^S (1-p)^S p^B)^m.$$

The only part left to argue is that g is implementable by a finite Bernoulli factory, but this follows by the definition of g: one can implement it by taking t samples of each coin, building $\bar{X}_t$ and checking whether $f(\bar{X}_t) \geq \frac{1}{2}$. Note that after sampling each of the coins of unknown bias t times, the vector $\bar{X}_t$ is constructing by summing the outcomes and dividing by t. No extra sampling is involved. Also checking the functional value of f at $\bar{X}_t$ requires no additional sampling. □

We now can derive the proof of Theorem 1.3 by recursively applying the previous lemma.

PROOF OF THEOREM 1.3. We will define a sequence of functions $f_1(p), f_2(p), \dots$ as follows. First we define $f_1(p) = f(p)$. Then for every $k \geq 1$ let g_k correspond to the g function in Lemma 3.5 obtained from f_k. Then define $f_{k+1}(p) = \frac{4}{3}(f_k(p) - \frac{1}{4}g_k(p))$. Unrolling the recursion we get

$$f(p) = \left(\frac{3}{4}\right)^k f_{k+1}(p) + \sum_{s=1}^{k} \frac{1}{4}\left(\frac{3}{4}\right)^{s-1} g_s(p).$$

Since $f_{k+1}(p) \in [0,1]$, it means that $0 \leq f(p) - \sum_{s=1}^{k} \frac{1}{4}(\frac{3}{4})^{s-1} g_s(p) \leq (\frac{3}{4})^k$, $\forall p \in [0,1]^n$ or in other words, the series $\sum_{k=1}^{\infty} \frac{1}{4}(\frac{3}{4})^{k-1} g_k(p)$ converges uniformly to $f(p)$. This observation gives a natural algorithm for sampling from $f(p)$: first sample an index $k \in \mathbb{Z}_{>0}$ with probability $\frac{1}{4}(\frac{3}{4})^{k-1}$. Then use the Bernoulli factory for $g_k(p)$ constructed in Lemma 3.5 to sample 1 with that probability. □

4. Proof of Lemma 3.4.

The heart of the argument is establishing Lemma 3.4. Note that as $t \to \infty$ we know that

$$\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2}\right] \to \mathbf{1}\left\{f(p) > \frac{1}{2}\right\}$$

if $f(p) \neq \frac{1}{2}$, so clearly it holds that for each $p \in [0,1]^n$ there is a large enough t such that the inequality in the lemma holds. The difficulty in the argument is to show a single t holds for all points p simultaneously.

The argument will proceed as follows: first let us define the region where the values of f are small:

$$L = \left\{p \in [0,1]^n; f(p) \leq \frac{3}{8}\right\}.$$

It is simple to see that if $p \notin L$ then for any value of t it holds that

$$f(p) - \frac{1}{4}\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2}\right] \geq \frac{3}{8} - \frac{1}{4} = \frac{1}{8} \geq \frac{1}{8} f(p).$$

Our proof strategy will be to argue that there exists some t such that equation (9) holds for all $q \in L$, we will prove the following claim.

CLAIM 4.1. *For every $q \in L$ there is a radius $r_q > 0$ and an integer t_q such that equation (9) in the statement of Lemma 3.4 holds for all $p \in L \cap \mathsf{B}_\infty(q; r_q)$ and $t \geq t_q$.*

If Claim 4.1 is true, then it provides us with an open cover $\bigcup_{q \in L} \mathsf{B}_\infty(q; r_q)$ of L. Hence we can use Lemma 2.1 to argue it must have a finite subcover, that is, there is a finite set of points $Q \subset L$, $|Q| < \infty$ such that $L \subset \bigcup_{q \in Q} \mathsf{B}_\infty(q; r_q)$. Hence if we take $t_0 = \max_{q \in Q} t_q$, then equation (9) in Lemma 3.4 holds for all $t \geq t_0$ and $p \in L$.

4.1. *A safe distance from the boundary.* We will start by making two observations about the geometry of set L.

CLAIM 4.2. *There is some positive constant $\delta > 0$ such that*

$$f(q) < \frac{1}{2} \quad \forall q \in \mathsf{B}_\infty(p, \delta) \text{ and } p \in L.$$

PROOF. Define $H = \{p \in [0,1]^n; f(p) \geq \frac{1}{2}\}$. By the continuity of f, the sets H and L are compact. Since H and L are disjoint compact sets there is a constant δ such that $\|p - q\|_\infty > \delta, \forall p \in H, q \in L$. □

With that observation, for any $p \in L$ we will bound $\mathbb{P}_p[f(\bar{X}_t) \geq \frac{1}{2}] \leq \mathbb{P}_p[\|\bar{X}_t - p\|_\infty > \delta]$, which allows us to apply concentration bounds. This will give us a good enough argument whenever $f(p) > 0$. For the case $f(p) = 0$, however, we need a more detailed understanding of how f behaves close to the boundary. For that, we will use the following two claims.

CLAIM 4.3. *If $p \in F_{A,S,B}$ and $f(p) = 0$ then $f|_{F_{A,S,B}} \equiv 0$.*

PROOF. If $f|_{F_{A,S,B}} \not\equiv 0$ then $f(q) \geq c \cdot ((1-q)^A q^S (1-q)^S q^B)^m$ which contradicts the fact that $f(p) = 0$ since $(1-p)^A = 1$, $p^B = 1$ and $p^S(1-p)^S > 0$. □

The second claim establishes that if a function is zero on an open face, then there exists a safe distance δ such that any point q within distance δ of this open face satisfies $f(q) < 1/2$.

CLAIM 4.4. *If $f|_{F_{A,S,B}} \equiv 0$ then $f(q) < 1/2$ for all $q \in [0,1]^n \cap B_\infty(F_{A,S,B}, \delta)$ for the constant δ in Claim 4.2.*

PROOF. Follows directly from Claim 4.2 and the fact that $F_{A,S,B} \subset L$. □

4.2. *Proof of Claim 4.1 for $f(q) > 0$.* By continuity, there is a radius r small enough such that

$$2f(q) \geq f(p) \geq \frac{1}{2} f(q) \quad \forall p \in L \cap \mathsf{B}_\infty(q, r).$$

By Claim 4.2 and the Hoeffding bound (equation (4)) we know that

$$\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2}\right] \leq \mathbb{P}_p[\|\bar{X}_t - p\|_\infty > \delta] \leq 2n \exp(-2\delta^2 t) \leq \frac{1}{4} f(q)$$

for $t \geq t_q := \lceil -\frac{1}{2\delta^2} \log(\frac{1}{8n} f(q)) \rceil$ and $p \in L$. Therefore, we have

$$f(p) - \frac{1}{4}\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2}\right] \geq \frac{1}{2} f(q) - \frac{1}{4} f(q) = \frac{1}{4} f(q) \geq \frac{1}{8} f(p),$$

therefore establishing equation (9) for $p \in L \cap \mathsf{B}_\infty(q, r)$ and $t \geq t_q$.

4.3. *Proof of Claim 4.1 for $f(q) = 0$.* Let $F_{A,S,B}$ be the open face containing q. By Claim 4.3 we know that $f|_{F_{A,S,B}} \equiv 0$. We start by taking a small radius r such that $r < \delta/2$ and $q_i - r > r > 0$ and $1 - q_i - r > r > 0$ for all $i \in S$. Through the course of the proof, we may decrease r further if necessary.

Our goal is to bound the probability $\mathbb{P}_p[f(\bar{X}_t) \geq 1/2]$ for $p \in \mathsf{B}_\infty(q,r) \cap L$. We will decompose this probability by looking at the coordinates (specifically, the coordinates in $A \cup B$) on which $\bar{X}_t$ significantly deviates from p:

$$\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2}\right] = \sum_{T \subseteq A \cup B} \mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2} \text{ and } T = \left\{i \in A \cup B; |\bar{X}_{t,i} - p_i| > \frac{\delta}{2}\right\}\right].$$

We will show that for each $T \subseteq A \cup B$ there is r small enough such that

$$\mathbb{P}_p\left[f(\bar{X}_t) \geq \frac{1}{2} \text{ and } T = \left\{i \in A \cup B; |\bar{X}_{t,i} - p_i| > \frac{\delta}{2}\right\}\right] \leq \frac{1}{2^{n+2}} f(p). \tag{11}$$

If we can establish this, then it would imply that $\mathbb{P}_p[f(\bar{X}_t) \geq \frac{1}{2}] \leq \frac{1}{4} f(p)$ which directly implies equation (9). We will consider two cases depending on the value of f in the open face $F_{A\setminus T, S\cup T, B\setminus T}$.

Case 1: $f_{F_{A\setminus T,S\cup T,B\setminus T}} \equiv 0$. In this case, by Claim 4.4 we have that all points p at a ℓ_∞ distance at most δ from $F_{A\setminus T,S\cup T,B\setminus T}$ have $f(p) < 1/2$. Now observe that all the coordinates in which p differs from $\bar{X}_t$ by at least $\delta/2$ belong to $S \cup T$. Hence, the ℓ_∞ distance between $\bar{X}_t$ and $F_{A\setminus T,S\cup T,B\setminus T}$ is at most $\delta/2$, and hence $\bar{X}_t \in L$, which in turn means that $f(\bar{X}_t) < \frac{1}{2}$. So the left-hand side of equation (11) is zero.

Case 2: $f_{F_{A\setminus T,S\cup T,B\setminus T}} \not\equiv 0$. In this case, by the fact that f is polynomially bounded we have

$$f(p) \geq c \cdot \left((1-p)^{A\setminus T} \cdot p^{S\cup T}(1-p)^{S\cup T} \cdot p^{B\setminus T}\right)^m.$$

For $p \in \mathsf{B}_\infty(q,r)$ we have $1 - p_i > r$ for $i \in A \cup S$ and $p_i > r$ for $i \in S \cup B$ by the definition of r. Since r is a constant we can write

$$f(p) \geq C \cdot \left(p^{T\cap A} \cdot (1-p)^{T\cap B}\right)^m.$$

Now note that the left-hand side of equation (11) is at most

$$\mathbb{P}_p\left[\bar{X}_{t,i} - p_i > \frac{\delta}{2}, \forall i \in T \cap A \text{ and } p_i - \bar{X}_{t,i} > \frac{\delta}{2}, \forall i \in T \cap B\right]$$

since for $i \in A$ we have $p_i - \bar{X}_{t,i} \leq p_i < r < \delta/2$ and for $i \in B$ we have $\bar{X}_{t,i} - p_i \leq 1 - p_i \leq r < \delta/2$. By the Chernoff bound (equation (5)) this can be bounded by

$$\prod_{i\in A\cap T} \left(C' p_i^{\delta/2}\right)^t \cdot \prod_{i \in B\cap T} \left(C'(1-p_i)^{\delta/2}\right)^t = (C')^{|T|t} \cdot \left(p^{A\cap T}(1-p)^{B\cap T}\right)^{\delta t/2}$$

for some constant C'. Choose $t = \lceil 4m/\delta \rceil$. Then for a constant C'' we can bound the expression above by

$$C'' \cdot \left(p^{A\cap T}(1-p)^{B\cap T}\right)^{2m}.$$

We know that $|T| \geq 1$ since $f|_{A,S,B} \equiv 0$, so $p^{A\cap T}(1-p)^{B\cap T} \leq r$. It follows that, for sufficiency small values of r,

$$r \leq \left(\frac{C}{2^n C''}\right)^{1/m}$$

and therefore equation (11) holds for all $p \in L \cap \mathsf{B}_\infty(q,r)$.

5. Factories on subdomains of the hypercube. Next we characterize which functions defined on a compact subdomain of the hypercube are implementable. Given a compact set $K \subseteq [0,1]^n$ and a continuous function $f : K \to [0,1]$, our first instinct is to check whether we can extend f to the entire hypercube satisfying the properties in Theorem 1.3. The problem with this approach is that it is possible to construct functions that are implementable on a subset K but cannot be extended to an implementable function on the entire hypercube. Here is an example:

EXAMPLE 5.1. Let K be the convex hull of $\{(\frac{1}{2}, 0), (0, \frac{1}{2})\}$ and define $f : K \to [0,1]$ as $f(p) = p_1/(p_1 + p_2)$. We first observe that f is implementable by the following procedure: choose a coin $i \in \{1,2\}$ uniformly at random and flip it. If the coin comes up 1, then output 1 if $i = 1$ and 0 if $i = 2$. If the flipped coin comes up 0 we retry. Each time we do it, we output 1 with probability $\frac{p_1}{2}$, we output 0 with probability $\frac{p_2}{2}$ and retry with probability $1 - \frac{p_1+p_2}{2}$. The total probability of outputting one is therefore $\sum_{k=0}^{\infty}(1 - \frac{p_1+p_2}{2})^k \frac{p_1}{2} = \frac{p_1}{p_1+p_2}$.

However, f does not have a continuous and polynomially bounded extension to the hypercube. To see that observe that since $f(\frac{1}{2}, 0) = 1$, then it must be 1 on the open face $(0,1) \times \{0\}$ by the fact it is polynomially bounded. Similarly, since $f(0, \frac{1}{2}) = 0$ then it must be 0 on the open face $\{0\} \times (0,1)$. Such a function cannot be continuous on $[0,1]^2$ since there are sequences approaching $(0,0)$ with different limits.

Instead of trying to extend the function we will adapt the proof in the previous section to deal with any domain K. First we say that a function $f : K \to [0,1]$ is polynomially bounded on K if there is an integer m and a real constant $c > 0$ such that for every open face $F_{A,S,B}$ of the hypercube, it holds that

$$\exists q \in K \cap F_{A,S,B}, \quad f(q) > 0 \quad \Rightarrow \quad f(p) \geq c\big((1-p)^A \cdot p^S \cdot (1-p)^S \cdot p^B\big)^m \quad \forall p \in K.$$

With this extended definition we state the following.

THEOREM 5.2 (Extension of Theorem 1.3 to subdomains). *For a compact $K \subseteq [0,1]^n$, a function $f : K \to [0,1]$ is implementable by a Bernoulli factory if and only if it is continuous and f and $1 - f$ are polynomially bounded on K.*

The necessary conditions follow from the exact same arguments as in Lemmas 3.1, 3.2 and 3.3. To prove it is sufficient, we need to deal with the following difficulty in extending the proof: the sampled average $\bar{X}_t$ may not be in the domain K so $f(\bar{X}_t)$ is not well defined. Therefore we cannot use Lemma 3.4 directly.

To address this, we need two new ideas. The first new idea is to project the sampled point $\bar{X}_t$ to the domain K. One difficulty in simply projecting the sampled average is that if Y is the projection of $\bar{X}_t$ to K then a large deviation in one coordinate (say $|\bar{X}_{t,i} - p_i|$ is large) can cause a large deviation for possibly many other coordinates $|Y_j - p_j|$ in the projection. This makes it hard to apply the argument in Section 4.3 since we need to reason about large deviations of subsets of coordinates.

The second new idea seeks to address this point: we will only project if the sampled average $\bar{X}_t$ is close enough to the domain K. If not, we will resample $\bar{X}_t$. By doing this, we can guarantee that the coordinates will not move too much.

5.1. *Project if close enough.* The previous discussion motivates the definition a new random variable $Z_{t,\epsilon}$. First, we will define the projection operator $\Pi_K : [0,1]^n \to K$ which is a function such that

$$\|\Pi_K(p) - p\|_\infty \leq \|q - p\|_\infty \quad \forall p \in [0,1]^n, q \in K.$$

Note that there may be many choices for Π_K, in which case we may choose arbitrarily. Also observe that Π_K is not necessarily continuous.

Now define $Y_{t,\epsilon}$ as a random variable taking values in $[0,1]^n$ distributed according to the conditional distribution of $\bar{X}_t$ given that $\bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)$:

$$\mathbb{P}_p[Y_{t,\epsilon} \in A] = \mathbb{P}_p\big[\bar{X}_t \in A \mid \bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)\big] \quad \forall \text{ measurable } A \subseteq [0,1]^n.$$

This variable can be sampled as follows: first sample $\bar{X}_t$. If $\bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)$ then set $Y_{t,\epsilon} = \bar{X}_t$. If not, resample $\bar{X}_t$ and try again until $\bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)$. Now define

$$Z_{t,\epsilon} = \Pi_K(Y_{t,\epsilon}).$$

LEMMA 5.3. *If $t \geq \log(8n)/(2\epsilon^2)$ then for any $p \in K$ and any measurable set A it holds that*

$$\mathbb{P}_p[Y_{t,\epsilon} \in A] \leq 2 \cdot \mathbb{P}_p[\bar{X}_{t,\epsilon} \in A].$$

PROOF. By the definition of $Y_{t,\epsilon}$ we can write

$$\mathbb{P}_p[Y_{t,\epsilon} \in A] = \mathbb{P}_p\big[\bar{X}_{t,\epsilon} \in A \mid \bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)\big] = \frac{\mathbb{P}_p[\bar{X}_{t,\epsilon} \in A \text{ and } \bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)]}{\mathbb{P}_p[\bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)]}.$$

The numerator of the last expression is clearly bounded above by $\mathbb{P}_p[\bar{X}_{t,\epsilon} \in A]$. For the denominator, notice that for $p \in K$ we have $\mathsf{B}_\infty(p,\epsilon) \subseteq \mathsf{B}_\infty(K,\epsilon)$ hence

$$\mathbb{P}_p\big[\bar{X}_t \in \mathsf{B}_\infty(K,\epsilon)\big] \geq \mathbb{P}_p\big[\bar{X}_t \in \mathsf{B}_\infty(p,\epsilon)\big] = 1 - \mathbb{P}_p\big[\|\bar{X}_t - p\|_\infty \geq \epsilon\big].$$

By the Hoeffding bound, we have $P_p[\|\bar{X}_t - p\|_\infty \geq \epsilon] \leq 2n\exp(-2\epsilon^2 t) \leq \frac{1}{2}$ for $t \geq \log(8n)/(2\epsilon^2)$. Putting it all together, we obtain the result in the statement. □

5.2. *Extension of Lemma 3.4 to subdomains.*

LEMMA 5.4. *Let $f : K \to [0,1]$ be a continuous and polynomially bounded function on K. Then there is some $\epsilon > 0$ and an integer t_0 such that for $t \geq t_0$ it holds that*

$$f(p) - \frac{1}{4} \cdot \mathbb{P}_p\left[f(Z_{t,\epsilon}) \geq \frac{1}{2}\right] \geq \frac{1}{8} f(p) \quad \forall p \in [0,1]^n. \tag{12}$$

With this lemma, the proof of Theorem 5.2 follows from exact the same arguments used in Section 3.3 to prove Theorem 1.3. The only new thing to note is that for any fixed t and ϵ the function $g_{t,\epsilon}(p) = \mathbb{P}_p[f(Z_{t,\epsilon}) \geq \frac{1}{2}]$ can be implemented by a Bernoulli factory since $Z_{t,\epsilon}$ can be sampled with only sample access to the p_i-coins.

5.3. *Proof of Lemma 5.4.* We are now left to prove Lemma 5.4, for which we will need a slight modification in the arguments. As before we can define

$$L = \left\{p \in K; f(p) \leq \frac{3}{8}\right\}.$$

For $p \notin L$, the statement of the lemma is once again trivial. For $p \in L$ we will follow the strategy in Claim 4.1. First, observe that Claim 4.2 still holds since L and $\{p \in K; f(p) \geq \frac{1}{2}\}$ are disjoint compact sets. Next we strengthen Claims 4.3 and 4.4.

CLAIM 5.5. *Let $\bar{F}_{A,S,B}$ be the closure of $F_{A,S,B}$. If $p \in F_{A,S,B}$ and $f(p) = 0$ then $f|_{K \cap \bar{F}_{A,S,B}} \equiv 0$.*

PROOF. The closure of $F_{A,S,B}$ is the union of all open faces $F_{A',S',B'}$ where $A \subseteq A'$ and $B \subseteq B'$. Now there is some point $q \in K \cap F_{A',S',B'}$ such that $f(q) > 0$ then $f(q) \geq c \cdot ((1-q)^{A'} q^{S'} (1-q)^{S'} q^{B'})^m$ which contradicts the fact that $f(p) = 0$ since $p_i < 1$ for $i \in A'$ since $A' \subseteq A \cup S$, $p_i > 0$ for $i \in B'$ since $B' \subseteq B \cup S$ and $0 < p_i < 1$ for $i \in S' \subseteq S$. □

CLAIM 5.6. *Let $\bar{F}_{A,S,B}$ be the closure of $F_{A,S,B}$. If $f|_{K \cap \bar{F}_{A,S,B}} \equiv 0$, there is some δ such that if $f(q) < \frac{1}{2}$ for all $q \in K \cap \mathsf{B}_\infty(\bar{F}_{A,S,B}, \delta)$.*

PROOF. For the second part, if no such δ exists, then there must be a sequence of points q_t with $q_t \in K \cap \mathsf{B}_\infty(\bar{F}_{A,S,B}, \frac{1}{t})$ such that $f(q_t) \geq \frac{1}{2}$. Since K is compact, there must be a subsequence of q_t that converges to a point $q^* \in K \cap \bar{F}_{A,S,B}$. Since $f(q_t) \geq \frac{1}{2}$ we have also $f(q^*) \geq \frac{1}{2}$. But this contradicts the previous paragraph, which shows that $f(q^*) = 0$. □

Now fix $\delta > 0$ small enough such that δ satisfies Claims 4.2 and Claim 5.6. We will mirror Sections 4.2 and 4.3 and prove Claim 4.1 first for $f(q) > 0$ and then for $f(q) = 0$.

5.3.1. *Claim for $f(q) > 0$.* We will set $\epsilon = \delta/2$. Observe that $\|Z_{t,\epsilon} - Y_{t,\epsilon}\|_\infty \leq \epsilon$ so $\|Y_{t,\epsilon} - p\|_\infty \leq \|Y_{t,\epsilon} - Z_{t,\epsilon}\|_\infty + \|Z_{t,\epsilon} - p\|_\infty \leq \epsilon + \|Z_{t,\epsilon} - p\|_\infty$, hence

$$\mathbb{P}_p\left[f(Z_{t,\epsilon}) \geq \frac{1}{2}\right] \leq \mathbb{P}_p\left[\|Z_{t,\epsilon} - p\|_\infty > \delta\right] \leq \mathbb{P}_p\left[\|Y_{t,\epsilon} - p\|_\infty > \frac{\delta}{2}\right].$$

By Lemma 5.3, the last probability is at most $2 \cdot \mathbb{P}_p[\|\bar{X}_t - p\|_\infty > \frac{\delta}{2}]$ for large enough t. From this point on, the proof is exactly the same as in Section 4.2.

5.3.2. *Claim for $f(q) = 0$.* We set $\epsilon = \delta/4$ and as in Section 4.3 we split the probability of $\mathbb{P}_p[f(Z_{t,\epsilon}) \geq \frac{1}{2}]$ depending on which components have a large deviation in $Y_{t,\epsilon}$:

$$\mathbb{P}_p\left[f(Z_{t,\epsilon}) \geq \frac{1}{2}\right] = \sum_{T \subseteq A \cup B} \mathbb{P}_p\left[f(Z_{t,\epsilon}) \geq \frac{1}{2} \text{ and } T = \left\{i \in A \cup B; \left|(Y_{t,\epsilon})_i - p_i\right| > \frac{\delta}{2}\right\}\right].$$

As before, we argue that for all points p in a small enough ball around q we have

$$\mathbb{P}_p\left[f(Z_{t,\epsilon}) \geq \frac{1}{2} \text{ and } T = \left\{i \in A \cup B; \left|(Y_{t,\epsilon})_i - p_i\right| > \frac{\delta}{2}\right\}\right] \leq \frac{1}{2^{n+2}} f(p). \tag{13}$$

For each $T \subseteq A \cup B$ we will consider two cases depending on the value of f on $K \cap \bar{F}_{A \setminus T, S \cup T, B \setminus T}$. A first observation is that the closure $F_{A \setminus T, S \cup T, B \setminus T}$ contains $F_{A,S,B}$ and hence $K \cap \bar{F}_{A \setminus T, S \cup T, B \setminus T} \neq \varnothing$.

Case 1: $f|_{K \cap \bar{F}_{A \setminus T, S \cup T, B \setminus T}} \equiv 0$. By Claim 5.6 we have $f(q) < 1/2$ for all $q \in \mathsf{B}_\infty(\bar{F}_{A \setminus T, S \cup T, B \setminus T}, \delta)$. Since the distance between $Z_{t,\epsilon}$ and that face is at most $\max_{i \in A \cup B \setminus T} |(Z_{t,\epsilon})_i - (Y_{t,\epsilon})_i| + |(Y_{t,\epsilon})_i - p_i| + r \leq \frac{3\delta}{4} + r < \delta$ for a radius $r < \frac{\delta}{4}$. Hence the probability in equation (13) is zero.

Case 2: $f|_{K \cap \bar{F}_{A \setminus T, S \cup T, B \setminus T}} \not\equiv 0$. In that case, there is an open face $F_{A',S',B'}$ in the closure $\bar{F}_{A \setminus T, S \cup T, B \setminus T}$ such that $K \cap F_{A',S',B'} \neq \varnothing$ and $f|_{F_{A',S',B'}} \not\equiv 0$. Since f is polynomially bounded on K, we have

$$f(p) \geq c \cdot \left((1-p)^{A'} p^{S'} (1-p)^{S'} p^{B'}\right)^m \geq c \cdot \left((1-p)^{A \setminus T} p^{S \cup T} (1-p)^{S'} p^{B \setminus T}\right)^m,$$

where the inequality follows from the fact that $A \setminus T \subseteq A'$ and $B \setminus T \subseteq B'$ since $F_{A',S',B'} \subseteq \bar{F}_{A \setminus T, S \cup T, B \setminus T}$. From this point on, the proof is exactly the same as in Section 4.2, using Lemma 5.3 to translate statements about $Y_{t,\epsilon}$ to $\bar{X}_t$.

6. Extending Sampford sampling to the boundary.

Given probabilities $(p_1, p_2, \dots, p_n)$ with $\sum_i p_i = k$ (for some integer $1 \le k < n$), we wish to sample a k-element sized subset U of $\{1, 2, \dots, n\}$ with the property that $\mathbb{P}[i \in U] = p_i$. Sampford sampling is a method for accomplishing this given only sample access to coins with these probabilities. Sampford sampling proceeds as follows:

- Sample each coin i (with probability p_i) once and let $X_i \in \{0, 1\}$ be the outcome.
- Let $U = \{i; X_i = 1\}$ be the set of coins that came up heads. If $|U| \ne k$, go back to step 1.
- Choose a uniform random coin in $[n] \setminus U$, and flip it. If it comes up heads, output U. Otherwise, go back to step 1.

For each $U \subset [n]$ with $|U| = k$, define

$$g_U(x) = \frac{1}{n-k}\Bigl(\prod_{i \in U} x_i\Bigr) \cdot \Bigl(\prod_{i \notin U} (1 - x_i)\Bigr) \cdot \Bigl(\sum_{i \notin U} x_i\Bigr).$$

Note that $g_U(p)$ is exactly the probability that we output a specific set U for one individual trial of the above procedure (i.e., without restarting the procedure). It follows that the above procedure samples a subset U with probability

$$f_U(x) = \frac{g_U(x)}{\sum_{V \subset [n], |V|=k} g_U(x)}.$$

Although $f_U(x)$ is defined on the interior $[0, 1]^n$, $f_U(x)$ is undefined for some points on the boundary of $[0, 1]^n$ (and even for some points within the subset $K = \{p \in [0, 1]^n; \sum p_i = k\}$). For example, consider the point p with $p_i = 1$ for $1 \le i \le k$ and $p_i = 0$ for $k + 1 \le i \le n$. Although this value of p satisfies $\sum_i p_i = k$, $f_U(p)$ is undefined at this point; in particular, $g_U(p) = 0$ for every single subset U. Indeed, for this set of probabilities, it's easy to verify that the procedure described above can never terminate: every round we will sample the set $U = \{1, 2, \dots, k\}$, and then immediately fail the subsequent check in step 3.

In this section we will show that it is indeed possible to construct a multiparameter Bernoulli factory for this problem that terminates almost surely for all valid sets of coins:

THEOREM 6.1. *There exists a multiparameter factory for Sampford sampling that terminates everywhere in the set $K = \{p \in [0, 1]^n; \sum p_i = k\}$.*

To prove this theorem we will show that there exists a continuous completion of $f_U(x)$ that satisfies the constraints of Theorem 5.2. Then we will apply the Bernoulli race construction by Dughmi et al. (2017).

LEMMA 6.2 (Bernoulli race (Dughmi et al. (2017))). *If functions $f_1, \dots, f_k : K \to [0, 1]$ can be implemented by a Bernoulli factory and $\sum_{j=1}^k f_j(p) > 0$ for all $p \in K$, then there exists a sampling algorithm that for each $p \in K$ samples an index $i \in [k]$ with probability proportional to $f_i(p)/(\sum_{j=1}^k f_j(p))$.*

PROOF. We sample an index $i \in [k]$ uniformly at random and then sample from the $f_i(p)$-coin. If it comes up 1 we return index i. Otherwise we retry. The procedure terminates a.s. since it has a positive probability of outputting for each retry. Since each trial outputs index i with probability proportional to $f_i(p)/k$, the overall procedure outputs index i with probability $f_i(p)/(\sum_{j=1}^k f_j(p))$. □

Given a set $U \subset [n]$, let $e_U \in [0, 1]^n$ be the point such that $(e_U)_i = 1$ if $i \in U$ and $(e_U)_i = 0$ otherwise. Note that each e_U with $|U| = k$ lies in K. We show that these are the only points of discontinuity of $f_U(x)$ within K, and that these discontinuities can be resolved.

LEMMA 6.3. *Define $\overline{f}_U(x) : K \to [0, 1]$ as*

$$\overline{f}_U(x) = \begin{cases} f_U(x) & \text{for } x \in K, x \neq e_U, \\ 1 & \text{for } x = e_U, \\ 0 & \text{for } x = e_V, V \neq U. \end{cases}$$

Then $\overline{f}_U(x)$ is a continuous function defined on all of K.

PROOF. We first show that $f_U(x)$ is defined for all points in K not of the form e_U (for any subset $U \subset [n]$ of size k). To see this, fix an $x \in K$ and let $I(x) = \{i \in [n]; x_i > 0\}$ be the set of nonzero coordinates of x. Since $x \in K$, $\sum x_i = k$ and therefore $|I(x)| \geq k$ – moreover, if $|I(x)| = k$, then we must have $x_i = e_{I(x)}$. It follows that if x is not of the form e_U, then $|I(x)| \geq k + 1$.

Now choose any subset U' of $I(x)$ of size k that contains all indices i such that $x_i = 1$ (since $\sum x_i = k$, there are at most k such indices, and they all must belong to $I(x)$). Note that each of the three terms of $g_{U'}(x)$ are positive, so $g_{U'}(x) > 0$. It follows that the denominator of $f_U(x)$ is positive, and therefore $f_U(x)$ is well-defined for all such points (and therefore $\overline{f}_U$ is well-defined for all points in K).

It remains to show $\overline{f}_U$ is continuous on K. It suffices to check continuity at the points $e_{U'}$. To see this, observe that $f_U(p)$ satisfies $\sum_{U;|U|=k} f_U(p)e_U = p$ for all $p \in K \setminus \{e_U; |U| = k\}$. Now fix a sequence $p_t \to e_{U'}$ and some subset $U'' \neq U'$. There is some coordinate $i \in U'' \setminus U'$. Then looking at the ith coordinate we have that $f_{U''}(p_t) \leq (p_t)_i \to (e_{U'})_i = 0$. Since $f_{U''} \geq 0$ we must have $f_{U''}(p_t) \to 0$. Hence $\overline{f}_{U''}$ is continuous at all points $e_{U'}$ with $U' \neq U''$. To check continuity when $U' = U''$ take any coordinate $i \in U'$ and $p_t \to e_{U'}$. Then $f_{U'}(p_t) = (p_t)_i - \sum_{U'' \neq U'} f_{U''}(p_t)(e_{U''})_i \to 1$ by the previous observation. Hence $\overline{f}_{U'}$ is also continuous at $e_{U'}$. □

We will now show that this function $\overline{f}_U$ satisfies the constraints of Theorem 5.2 on the set K (and hence we can construct a multivariate Bernoulli factory for the function $\overline{f}_U$ that terminates a.s. for all points in K).

LEMMA 6.4. *The function $\overline{f}_U$ is polynomially bounded on K.*

PROOF. Let us begin by characterizing the faces $F_{A,S,B}$ where there exists a point $p \in F_{A,S,B} \cap K$ such that $\overline{f}_U(p) > 0$. In particular, we claim that if this happens, then $B \subseteq U \subseteq (S \cup B)$. To see why, note that if $i \in B$ then $p_i = 1$ for $p \in F_{A,S,B}$, so if we have a positive probability of outputting subset U, U must contain element i. Similarly, if $i \in A$, then $p_i = 0$ for $p \in F_{A,S,B}$, so if we have a positive probability of outputting subset U, U cannot contain element i (and thus must be contained in $S \cup B$).

We will now show that for all $p \in K$ and faces $F_{A,S,B}$ satisfying $B \subseteq U \subseteq (S \cup B)$, there exist constants $c, m > 0$ such that

$$\overline{f}_U(p) \geq c \cdot \big((1-p)^A \cdot (1-p)^S p^S \cdot p^B\big)^m. \tag{14}$$

Let $\overline{U} = [n] \setminus U$. Note that since $p, 1 - p \leq 1$, $U \subseteq (S \cup B)$, and $\overline{U} \subseteq (S \cup A)$, (14) is implied by the following inequality:

$$\overline{f}_U(p) \geq c\big((1-p)^{\overline{U}} \cdot p^U\big)^m. \tag{15}$$

We will prove (15). First, note that this holds for all p of the form e_V with $|V| = k$ (in particular, whenever $\overline{f}_U(e_V) = 0$, the RHS of (15) is also 0). It suffices to prove (15) on all other points of K. On these points $\overline{f}_U(p) = f_U(p)$, so by substituting in the definition of $f_U(p)$, it suffices to prove that

$$g_U(p) \geq c\big((1-p)^{\overline{U}} \cdot p^U\big)^m \sum_{|V|=k} g_V(p). \tag{16}$$

Note that

$$g_U(p) = \frac{1}{n-k} p^U (1-p)^{\overline{U}} \sum_{i \in U} p_i.$$

Inequality (16) thus reduces to

$$\sum_{i \notin U} p_i \geq c(n-k)\big((1-p)^{\overline{U}} \cdot p^U\big)^{m-1} \sum_{|V|=k} g_V(p). \tag{17}$$

We will now prove the following: for any $\varepsilon > 0$, if $\sum_{i \notin U} p_i = \varepsilon$, then $g_V(p) \leq k\varepsilon$ for each $V \subseteq [n]$ with $|V| = k$. Note that this implies (17) (in particular, it suffices to set $c = 1/((n-k)k\binom{n}{k})$ and $m = 1$).

To show the above claim, note that if $\sum_{i \notin U} p_i = \varepsilon$, then $p_i \leq \varepsilon$ for all $i \notin U$. Now note that for any $V \neq U$ with $|V| = |U| = k$, there must exist an index i^* belonging to V that does not belong to U. It follows that for $V \neq U$.

$$g_V(p) = \frac{1}{n-k} p^V (1-p)^{\overline{V}} \left(\sum_{i \in V} p_i\right) \leq k p_{i^*} \leq k\varepsilon.$$

For $V = U$, it immediately follows that $g_U(p) \leq \sum_{i \in U} p_i = \varepsilon$, concluding our proof. □

LEMMA 6.5. *If functions $f_1, \dots, f_k : K \to [0,1]$ are polynomially bounded, then so is their sum $f := f_1 + \cdots + f_k$.*

PROOF. If for a certain open face $F_{A,S,B}$ we have $f|_{K \cap F_{A,S,B}} \not\equiv 0$ then there must be one index i such that $f_i|_{K \cap F_{A,S,B}} \not\equiv 0$ and since f_i is polynomially bounded, we have $f(p) \geq f_i(p) \geq c((1-p)^A p^S (1-p)^S p^B)^m$ for some constant $c > 0$ and integer $m \geq 0$. □

Putting it all together we obtain a proof of Theorem 6.1.

PROOF OF THEOREM 6.1. We first argue that $\overline{f}_U$ satisfies the constraints of Theorem 5.2 for the set K, and therefore that we can construct a multivariate Bernoulli factory for $\overline{f}_U$ that terminates a.s. for all $p \in K$.

To show this, we must show that $\overline{f}_U$ is continuous on K, and that both $\overline{f}_U$ and $1 - \overline{f}_U$ are polynomially bounded on K. We have already shown that $\overline{f}_U$ is continuous on K (Lemma 6.3) and that $\overline{f}_U$ is polynomially bounded on K (Lemma 6.4). To see that $1 - \overline{f}_U$ is polynomially bounded on K, note that $1 - \overline{f}_U = \sum_{V \neq U} \overline{f}_V$. Since this is a sum of functions of each polynomially bounded on K, it follows that $1 - \overline{f}_U$ is polynomially bounded on K (Lemma 6.5).

Finally, we will use our factories that output a coin with probability $\overline{f}_U(p)$ to construct a factory that outputs an actual subset U using the Bernoulli race in Lemma 6.2. □

7. Combinatorial Bernoulli factories. Given a polytope $\mathcal{P} \subseteq [0,1]^n$ (with vertices $V(\mathcal{P})$), a *combinatorial Bernoulli factory* for $\mathcal{P}$ is an exact sampling procedure that, given coins $(p_1, p_2, \dots, p_n) \in \mathcal{P}$, outputs a vertex $v \in V(\mathcal{P})$ such that $\mathbb{E}_p[v] = p$. More formally, a combinatorial Bernoulli factory is a collection[2] of $|V(\mathcal{P})|$ multiparameter Bernoulli factories for functions $f_v(p)$ satisfying (for all $p \in \mathcal{P}$)

$$\sum_{v \in V(\mathcal{P})} f_v(p) = 1 \quad \text{and} \quad \sum_{v \in V(\mathcal{P})} v f_v(p) = p.$$

Combinatorial Bernoulli factories capture a wide range of combinatorial sampling problems. For example, the problem of Sampford sampling is equivalent to the problem of constructing a combinatorial Bernoulli factory for the polytope $\mathcal{P} = [0, 1]^n \cap \{p \in \mathbb{R}^n | \sum_i p_i = k\}$. Other problems captured by combinatorial Bernoulli factories include exact sampling of matchings and flows.

Niazadeh, Paes Leme and Schneider (2021) show that any polytope $\mathcal{P}$ that admits a combinatorial Bernoulli factory must be of the form $\mathcal{P} = [0, 1]^n \cap K$, where K is an affine subspace of $\mathbb{R}^n$. Moreover, they give a general method for constructing combinatorial Bernoulli factories for any such polytope; however, as with existing implementations of Sampford sampling, the factories they generate can fail to terminate at some points on the boundary of $[0, 1]^n$. In this section, we provide an alternate method for constructing combinatorial Bernoulli factories that works for all polytopes of the form $\mathcal{P} = [0, 1]^n \cap K$, *everywhere* in $\mathcal{P}$.

We begin by presenting our new construction. Given a d-dimensional polytope $\mathcal{P}$ and a vertex w of $\mathcal{P}$, we say that the *fan triangulation* $\mathcal{T}_w$ of $\mathcal{P}$ corresponding to vertex w is the division of $\mathcal{P}$ into simplices with disjoint interiors formed by connecting w to each facet of $\mathcal{P}$ that does not contain w (if a facet contains more than d vertices of $\mathcal{P}$, arbitrarily triangulate it into $(d-1)$-dimensional simplices first).

Note that given a simplex, there is a unique way to write a point in the simplex as a convex combination of its vertices. This implies that any triangulation $\mathcal{T}$ of $\mathcal{P}$ gives rise to a natural way to decompose a point $p \in \mathcal{P}$ as a convex combination of the vertices of $\mathcal{P}$: namely, find the simplex T of the triangulation that p belongs to, and write p as a convex combination of the vertices of T. Let $g_v^{(w)}(p)$ be the coefficient of vertex v in the decomposition stemming from the fan triangulation $\mathcal{T}_w$. Note that all these functions $g_v^{(w)} : \mathcal{P} \to [0, 1]$ are continuous since they are continuous on each simplex of the triangulation and agree on the common faces.

$$f_v(p) = \frac{1}{|V(\mathcal{P})|} \sum_{w \in V(\mathcal{P})} g_v^{(w)}(p). \tag{18}$$

By construction, it follows that $\sum_v f_v(p) = 1$ and $\sum_v v f_v(p) = p$ for all $p \in \mathcal{P}$. In the remainder of the section, we will show that if $\mathcal{P}$ is of the form $[0, 1]^n \cap K$ for some affine subspace K, then each $f_v(p)$ satisfies the conditions of Theorem 5.2 (and thus can be implemented by a multiparameter Bernoulli factory).

We will need the following lemma.

LEMMA 7.1. *Let $\mathcal{P}$ be a polytope of the form* $[0, 1]^n \cap K$, *where K is an affine subspace of $\mathbb{R}^n$. Let $v \in V(\mathcal{P})$ be a vertex of $\mathcal{P}$, and let F be a facet of $\mathcal{P}$ that does not contain v. Then there exists a coordinate i such that either $v_i > 0$ and $x_i = 0$ for all $x \in F$, or $v_i < 1$ and $x_i = 1$ for all $x \in F$.*

PROOF. Since K is an affine subspace, each facet of $\mathcal{P}$ can be written as the intersection of a facet of $[0, 1]^n$ with K. Each facet of $[0, 1]^n$ is given by a single constraint of the form $x_i = 0$ or $x_i = 1$.

[2] The Bernoulli race in Lemma 6.2 is used to convert this collection of factories into a procedure for sampling a vertex with the desired probability.

Assume that the facet F is equal to $\{x_i = 0\} \cap K$. Then if v is not contained in F, it must be the case that $v_i \neq 0$ (and thus $v_i > 0$ and the first condition of the theorem holds). Similarly, if the facet F is given by $\{x_i = 1\} \cap K$, then any vertex v not contained in F must satisfy $v_i < 1$, and the second condition of the theorem holds. □

Note that the previous lemma fails if $\mathcal{P}$ is not the intersection of the hypercube with an affine subspace. For example, if $\mathcal{P}$ is the convex hull of $(0, 0)$, $(1, 0)$, $(0, 1)$, the lemma fails for $v = (0, 0)$ and F the opposite edge. This is an important sanity check, as Niazadeh, Paes Leme and Schneider (2021) shows that no other polytope admits a combinatorial Bernoulli factory.

We can now show that $f_v(p)$ is polynomially bounded (and thus that we can construct combinatorial Bernoulli factories that terminate everywhere on $\mathcal{P}$).

LEMMA 7.2. *If $\mathcal{P}$ is a polytope of the form $[0, 1]^n \cap K$ where K is an affine subspace of $\mathbb{R}^n$, then the functions $f_v(p)$ are polynomially bounded on $\mathcal{P}$.*

PROOF. Fix a vertex v of $\mathcal{P}$. To begin, we'll argue that $g_v^{(v)}(p) > 0$ for exactly the points $p \in \mathcal{P}$ where $f_v(p) > 0$. To see this, note that if $g_v^{(v)}(p) > 0$, then $f_v(p) > 0$ (since $g_v^{(v)}(p)$ is a summand in $f_v(p)$). But conversely, by construction $g_v^{(v)}(p)$ only equals 0 on (closed) faces of $\mathcal{P}$ that do not contain v. Since $f_v(p)$ also forms a convex decomposition of p into the vertices of $\mathcal{P}$, $f_v(p)$ must also equal 0 on all these closed faces and it follows that $g_v(p) = 0$ implies that $f_v(p) = 0$.

We will now show that $g_v^{(v)}(p)$ is polynomially bounded on $\mathcal{P}$; it then follows from (18) that $f_v(p)$ is polynomially bounded on $\mathcal{P}$ (since $f_v(p) > 0$ implies $g_v^{(v)}(p) > 0$).

Recall that $g_v^{(v)}(p)$ is the decomposition induced by the fan triangulation $\mathcal{T}_v$. That is, to compute the value of $g_v^{(v)}(p)$, we first must identify the simplex of $\mathcal{T}_v$ that p belongs to, and (uniquely) write p as a convex combination of the vertices of that simplex.

Let us assume that p belongs to the simplex $T \in \mathcal{T}_v$. Since $\mathcal{T}_v$ is the fan triangulation for vertex v, T must be the convex hull of a facet F of $\mathcal{P}$ (not containing v) and v. By Lemma 7.1, there exists some coordinate i such that either $v_i > 0$ and $F \subset \{x; x_i = 0\}$ or $v_i < 1$ and $F \subset \{x; x_i = 1\}$.

In the first case, note that $g_v^{(v)}(p)$ must equal p_i/v_i (since v_i is the only vertex of T that contains a positive ith component). We claim that given this, $g_v^{(v)}(p)$ is polynomially bounded. To see why, note that if $f_v(p) > 0$ then $g_v^{(v)}(p) > 0$ and as a consequence $p_i > 0$ (since $v_i > 0$). It follows that if $g_v^{(v)}(p) > 0$ and p belongs to some open face $F_{A,S,B}$ of the hypercube, then $i \notin A$. But now note that if $i \notin A$, then there exist constants c and m such that

$$\frac{p_i}{v_i} \geq c \cdot \left((1-p)^A \cdot p^S (1-p)^S \cdot p^B\right)^m. \tag{19}$$

In particular, since i lies either in S or B, it suffices to take $m = 1$ and $c = 1/v_i$.

Similarly, in the second case $g_v^{(v)}(p)$ must equal $(1 - p_i)/(1 - v_i)$. A similar argument shows that $g_v^{(v)}(p)$ is polynomially bounded in this case (now we must have $i \notin B$, and a factor of $(1 - p_i)$ will appear on the RHS of the analogue of (19)). □

THEOREM 7.3. *If $\mathcal{P}$ is a polytope of the form $[0, 1]^n \cap K$, where K is an affine subspace $\mathbb{R}^n$, then there exists a combinatorial Bernoulli factory for $\mathcal{P}$ which terminates almost surely everywhere on the boundary.*

PROOF. It suffices to show that the functions $f_v(p)$ defined in (18) satisfy the conditions of Theorem 5.2. Namely, we must show that $f_v(p)$ are continuous, and that both $f_v(p)$ and $1 - f_v(p)$ are polynomially bounded on $\mathcal{P}$.

Since each $g_v^{(w)}(p)$ is continuous, $f_v(p)$ is continuous. By Lemma 7.2, each $f_v(p)$ is polynomially bounded on $\mathcal{P}$. Finally, note that $1 - f_v(p) = \sum_{w \neq v} f_w(p)$, so $1 - f_v(p)$ is polynomially bounded by Lemma 6.5. □

REFERENCES

ASMUSSEN, S., GLYNN, P. W. and THORISSON, H. (1992). Stationarity detection in the initial transient problem. *ACM Trans. Model. Comput. Simul.* **2** 130–157.

CAI, Y., OIKONOMOU, A., VELEGKAS, G. and ZHAO, M. (2021). An efficient ε-BIC to BIC transformation and its application to black-box reduction in revenue maximization. In *Proceedings of the* 2021 *ACM-SIAM Symposium on Discrete Algorithms* (*SODA*) 1337–1356. SIAM, Philadelphia, PA. MR4262514 https://doi.org/10.1137/1.9781611976465.81

DALE, H., JENNINGS, D. and RUDOLPH, T. (2015). Provable quantum advantage in randomness processing. *Nat. Commun.* **6** 1–4.

DUGHMI, S., HARTLINE, J. D., KLEINBERG, R. and NIAZADEH, R. (2017). Bernoulli factories and black-box reductions in mechanism design. In *STOC'*17*—Proceedings of the* 49*th Annual ACM SIGACT Symposium on Theory of Computing* 158–169. ACM, New York. MR3678179 https://doi.org/10.1145/3055399.3055492

FLEGAL, J. M. and HERBEI, R. (2012). Exact sampling for intractable probability distributions via a Bernoulli factory. *Electron. J. Stat.* **6** 10–37. MR2879671 https://doi.org/10.1214/11-EJS663

GONÇALVES, F. B., ŁATUSZYŃSKI, K. G. and ROBERTS, G. O. (2017). Exact Monte Carlo likelihood-based inference for jump-diffusion processes. Preprint. Available at arXiv:1707.00332.

HERBEI, R. and BERLINER, L. M. (2014). Estimating ocean circulation: An MCMC approach with approximated likelihoods via the Bernoulli factory. *J. Amer. Statist. Assoc.* **109** 944–954. MR3265667 https://doi.org/10.1080/01621459.2014.914439

KEANE, M. and O'BRIEN, G. L. (1994). A Bernoulli factory. *ACM Trans. Model. Comput. Simul.* **4** 213–219.

MORINA, G. (2021). Extending the Bernoulli factory to a dice enterprise. Ph.D. thesis, Univ. Warwick.

MORINA, G., ŁATUSZYŃSKI, K., NAYAR, P. and WENDLAND, A. (2022). From the Bernoulli factory to a dice enterprise via perfect sampling of Markov chains. *Ann. Appl. Probab.* **32** 327–359. MR4386529 https://doi.org/10.1214/21-aap1679

MOSSEL, E. and PERES, Y. (2005). New coins from old: Computing with unknown bias. *Combinatorica* **25** 707–724. MR2199432 https://doi.org/10.1007/s00493-005-0043-1

NACU, Ş. and PERES, Y. (2005). Fast simulation of new coins from old. *Ann. Appl. Probab.* **15** 93–115. MR2115037 https://doi.org/10.1214/105051604000000549

NIAZADEH, R., PAES LEME, R. and SCHNEIDER, J. (2021). Combinatorial Bernoulli factories: Matchings, flows, and other polytopes. In *STOC '*21*—Proceedings of the* 53*rd Annual ACM SIGACT Symposium on Theory of Computing* 833–846. ACM, New York. MR4398885 https://doi.org/10.1145/3406325.3451072

SAMPFORD, M. R. (1967). On sampling without replacement with unequal probabilities of selection. *Biometrika* **54** 499–513. MR0223051 https://doi.org/10.1093/biomet/54.3-4.499

VON NEUMANN, J. (1951). Various techniques used in connection with random digits. *Appl. Math. Ser.* **12** 5.

YUAN, X., LIU, K., XU, Y., WANG, W., MA, Y., ZHANG, F., YAN, Z., VIJAY, R., SUN, L. et al. (2016). Experimental quantum randomness processing using superconducting qubits. *Phys. Rev. Lett.* **117** 010502.