
Gambling-Based Confidence Sequences for Bounded Random Vectors

J. Jon Ryu¹ Gregory W. Wornell¹

Abstract

A confidence sequence (CS) is a sequence of confidence sets that contains a target parameter of an underlying stochastic process at any time step with high probability. This paper proposes a new approach to constructing CSs for means of bounded multivariate stochastic processes using a general gambling framework, extending the recently established coin toss framework for bounded random processes. The proposed gambling framework provides a general recipe for constructing CSs for categorical and probability-vector-valued observations, as well as for general bounded multidimensional observations through a simple reduction. This paper specifically explores the use of the mixture portfolio, akin to Cover’s universal portfolio, in the proposed framework and investigates the properties of the resulting CSs. Simulations demonstrate the tightness of these confidence sequences compared to existing methods. When applied to the sampling without-replacement setting for finite categorical data, it is shown that the resulting CS based on a universal gambling strategy is provably tighter than that of the posterior-prior ratio martingale proposed by Waudby-Smith and Ramdas.

1. Introduction

Time-uniform (or anytime-valid) confidence sets, which is often called a *confidence sequence* (CS), is a sequence of sets that contains a target parameter of the underlying stochastic process *at any time* with high probability. The time-uniformity guarantee implies that a practitioner can make a decision by continuously acquiring new data in a sequential manner, without the concern of the so-called p-hacking problem with the classical (non-time-uniform) confidence sets (Ramdas et al., 2020). The notion of time-

uniform guarantee dates back to Hoeffding (1963); Darling & Robbins (1967); Lai (1976), but the idea has regained attention only recently with the new modern applications such as A/B testing (Evidently, 2023) and election audits (Waudby-Smith et al., 2021). As tight confidence sequences can directly lead to saving various resources in practical sequential decision making applications, there has been a surge of interest in the research community, e.g., (Waudby-Smith & Ramdas, 2024; Howard et al., 2021; Orabona & Jun, 2024), to name a few.

While the field is rapidly evolving with new theories and constructions for CS (Waudby-Smith & Ramdas, 2024; Howard et al., 2021; Jun & Orabona, 2019; Orabona & Jun, 2024; Manole & Ramdas, 2023; Shekhar & Ramdas, 2023), we remark that most of the development has been especially for scalar-valued processes until recently; a few exceptions include (Waudby-Smith & Ramdas, 2020) for categorical observation in without replacement sampling, and (Whitehouse et al., 2023; Chugg et al., 2023) for unbounded vectors.

The goal of this paper is to develop a set of new methods for constructing CSs for *bounded, vector-valued* stochastic processes. Note that there exist natural bounded vector-valued data in practice, a canonical example of which is a categorical data. For bounded scalar-valued processes, Waudby-Smith & Ramdas (2024) proposed a coin-toss-based framework, studying several gambling strategies and properties of the resulting CSs. In the same framework, Orabona & Jun (2024) considered applying the celebrated algorithm of Cover’s universal portfolio (Cover, 1991) in information theory, proposed its tight outer approximations, and established analytical properties of their algorithms. A subsequent work of Ryu & Bhatt (2022) further explored the idea of universal gambling, arguing that the coin-toss framework can be understood as the continuous two-horse race, which admits more natural interpretation and closes the conceptual gap between CS and universal gambling. Shekhar & Ramdas (2023) recently established near-optimality of gambling-based CSs.

Given the elegance of the coin-toss framework for bounded, scalar-valued processes and its empirical tightness, especially in the small-sample regime (Orabona & Jun, 2024), it is tempting to ask whether an analogous gambling frame-

¹Department of EECS, MIT, Cambridge, Massachusetts, USA. Correspondence to: J. Jon Ryu <jongha@mit.edu>.

work exists for bounded, vector-valued processes. In this paper, we provide an affirmative answer: starting from the two-horse race formulation of [Ryu & Bhatt \(2022\)](#), we extend the existing coin-toss framework ([Waudby-Smith & Ramdas, 2024](#)) to a general gambling framework for bounded, vector-valued processes. The key observation that leads to the multivariate extension is that the coin-toss framework essentially constructs confidence sequences for binary probability vectors, each of which happens to be characterized by a single number. Going forward, we explore the resulting CSs from this framework when we use a universal gambling strategy such as [Cover \(1991\)](#)'s universal portfolio as in ([Orabona & Jun, 2024](#); [Ryu & Bhatt, 2022](#)) for scalar-valued processes, and study their theoretical properties and empirical performance.

The rest of the paper is organized as follows. In [Section 2](#), we gradually build up the idea for the general gambling framework, by starting from the standard definitions, basic tools, and the coin-toss framework for scalar-valued processes. In the next two sections, we study the resulting CSs for the categorical observation ([Section 3](#)) and probability-valued observation ([Section 4](#)), showing the nice properties of CSs resulting from universal gambling strategies. In [Section 4.2](#), we briefly describe how the general, bounded vector processes can be handled as probability-vector valued observations. We conclude with remarks in [Section 5](#).

2. General Gambling Framework

In this section, we illustrate the key idea of the paper on how to construct a confidence sequence for vector-valued observations via gambling.

2.1. Preliminaries

We start with a few definitions and notation. We use a boldface variable, e.g., \mathbf{x} , to denote a vector-valued observation. We use a capital letter X_t (or \mathbf{X}_t) to denote a random variable (or a random vector) and a small letter x_t (or \mathbf{x}_t) for its realization. We use the shorthand notation $\mathbf{x}_i^j := (x_i, \dots, x_j)$ and $\mathbf{x}^t := \mathbf{x}_1^t$. For an underlying stochastic process \mathbf{Y}^t , let $\mathcal{F}_t := \sigma(\mathbf{Y}^t)$ be the sigma-field generated by \mathbf{Y}^t , where \mathcal{F}_0 is the trivial sigma-field. We call the sequence $(\mathcal{F}_t)_{t=0}^\infty$ the *canonical filtration*.

We assume that there exists an underlying, bounded vector-valued stochastic process $(\mathbf{Y}_t)_{t=1}^\infty$ such that $\mathbb{E}[\mathbf{Y}_t | \mathcal{F}_{t-1}] = \boldsymbol{\mu}$ for any $t \geq 1$ for some $\boldsymbol{\mu} \in \mathbb{R}^K$. Our goal is to construct *time-uniform* confidence sets, so that the target mean parameter $\boldsymbol{\mu}$ is to be contained in the sequence at any time, with high probability. Formally, we wish to find a sequence of set-valued functions, each of which maps the observations \mathbf{y}^t until then to a *convex* subset $\mathcal{C}(\mathbf{y}^t; \delta) \subset \mathbb{R}^K$, such that

$$\mathbb{P}\{\boldsymbol{\mu} \in \mathcal{C}(\mathbf{Y}^t; \delta), \forall t \geq 1\} \geq 1 - \delta,$$

for any given desired parameter $\delta \in (0, 1)$.

Our development follows the standard technique based on the following inequality by [Ville \(1939\)](#) for supermartingales, which is a cornerstone of the most known confidence sequences. Different methods plug in different supermartingales given their own problem settings and considerations. A stochastic process $(W_t)_{t=0}^\infty$ is said to be *supermartingale* if $(W_t)_{t=0}^\infty$ is adapted to the canonical filtration and satisfy

$$\mathbb{E}[W_t | \mathcal{F}_{t-1}] \leq W_{t-1}$$

for every $t \geq 1$. It is said to be *martingale*, if the inequality holds with equality for any $t \geq 1$.

Theorem 1 (Ville's inequality). *For a nonnegative supermartingale sequence $(W_t)_{t=0}^\infty$ with $W_0 > 0$, for any $\delta > 0$,*

$$\mathbb{P}\left\{\sup_{t \geq 1} \frac{W_t}{W_0} \geq \frac{1}{\delta}\right\} \leq \delta.$$

A general recipe to derive confidence sequences based on Ville's inequality can be briefly summarized as follows. First, construct a sequence $(W_t(\mathbf{y}^t; \mathbf{m}))_{t=1}^\infty$ as a function of \mathbf{y}^t and a candidate parameter \mathbf{m} at each time $t \geq 1$. Define

$$\mathcal{C}(\mathbf{y}^t; \delta) := \left\{ \mathbf{m} : \frac{W_t(\mathbf{y}^t; \mathbf{m})}{W_0(\emptyset; \mathbf{m})} < \frac{1}{\delta} \right\}. \quad (1)$$

If the induced stochastic process $(W_t(\mathbf{Y}^t; \mathbf{m}))_{t=1}^\infty$ whose randomness comes from \mathbf{Y}^t is (super)martingale for $\mathbf{m} = \boldsymbol{\mu}$, then it readily follows from Ville's inequality that $(\mathcal{C}(\mathbf{Y}^t; \delta))_{t=1}^\infty$ is a confidence sequence for $\boldsymbol{\mu}$ of coverage $1 - \delta$.¹ Since we expect the confidence sequence to be shrinking to the singleton that contains $\boldsymbol{\mu}$ as fast as possible while maintaining the desired coverage, the sequence $(W_t(\mathbf{y}^t; \mathbf{m}))_{t=1}^\infty$ needs to be carefully constructed. Ideally, the stochastic process $(W_t(\mathbf{y}^t; \mathbf{m}))_{t=1}^\infty$ should diverge quickly for all $\mathbf{m} \neq \boldsymbol{\mu}$, so that any $\mathbf{m} \neq \boldsymbol{\mu}$ can be excluded from $(\mathcal{C}(\mathbf{Y}^t; \delta))_{t=1}^\infty$, eventually letting it converge to the singleton $\{\boldsymbol{\mu}\}$. In summary, the (super)martingale property of the process for $\mathbf{m} = \boldsymbol{\mu}$ guarantees the constructed sequence to be a valid confidence sequence, and the speed of the divergence of the stochastic process for $\mathbf{m} \neq \boldsymbol{\mu}$ will govern the speed of concentration to the target parameter.

2.2. Coin-Toss Framework for $[0, 1]$ -Valued Processes

One natural and elegant way of constructing a martingale sequence is to consider the wealth process of *gambling*. The high-level idea is that if a gambling is *subfair* (i.e., statistically not favorable to gamblers), then a gambler's

¹We note in passing that the running intersection $(\cap_{i=1}^t \mathcal{C}_i(\mathbf{y}^i; \delta))_{i=1}^\infty$ is always tighter or equal to the original CS and also a valid CS, and thus this operation is always implicitly assumed.

wealth as a stochastic process must be supermartingale, that is, the gambler cannot earn money out of the gambling. The idea of using gambling has been extensively studied by a line of recent works for bounded, scalar-valued stochastic processes (Waudby-Smith & Ramdas, 2024; Orabona & Jun, 2024; Ryu & Bhatt, 2022; Shekhar & Ramdas, 2023). Below, we briefly overview the existing formulation and motivate our idea for the extension to multivariate processes.

We start from the standard betting framework introduced in (Waudby-Smith & Ramdas, 2024) for a conditionally mean- μ stochastic process $(Y_t)_{t=1}^\infty$ for some $\mu \in (0, 1)$. Waudby-Smith & Ramdas (2024) introduced the definition of a *capital process* $K_t(m)$ for any $m \in (0, 1)$ as

$$K_t(y^t; m) := \prod_{i=1}^t (1 + \lambda_i(m) \cdot (y_i - m)), \quad (2)$$

where $(\lambda_t(m))_{t=1}^\infty$ is a $[-1/(1-m), 1/m]$ -valued predictable sequence, i.e., $\lambda_t(m)$ only depends on Y^{t-1} , but not on the future Y_t^∞ . It is clear that the capital process is martingale for $m = \mu$, since $E[K_t(Y^t; \mu) | \mathcal{F}_{t-1}] = K_t(Y^{t-1}; \mu) \cdot (1 + \lambda_t(m) \cdot (E[Y_t | \mathcal{F}_{t-1}] - \mu)) = K_t(Y^{t-1}; \mu)$. Here, $(Y_t)_{t=1}^\infty$ can be viewed as the stochastic outcomes of coin betting, and the predictable sequence $(\lambda_t(m))_{t=1}^\infty$ can be understood as a causal betting strategy. Waudby-Smith & Ramdas (2024) studied various different betting strategies and their properties in terms of the resulting confidence sequences. We remark, however, that it is nontrivial to extend this betting framework to vector-valued processes, as also pointed out in the discussion of Li et al. (2023). One critical complication in this convention is that the range of the bet $\lambda_t(m)$ depends on the candidate parameter m .

More recently, Ryu & Bhatt (2022) proposed an equivalent, yet alternative formulation of the betting framework, connecting it to the standard language of gambling in information theory (Cover & Thomas, 2006, Ch. 6, 16). The alternative formulation is based on the (continuous) two-horse race interpretation of the capital process (2). The rule of the standard horse race is as follows. There are two horses in the race, say horse 1 and horse 2, associated with odds $o_1 > 0$ and $o_2 > 0$, respectively. These numbers indicate that if the horse i wins, a \$1 bet on the horse results in a payoff of $\$o_i$. At each round $t \geq 1$, a gambler bets over the (yet unseen) outcome of the horse race with two horses. The bet is characterized by a single number $b_t \in [0, 1]$, which is to distribute the current wealth W_t of the gambler to the two horses as $\$b_t W_t$ and $\$(1 - b_t)W_t$. After the race, the per-round multiplicative gain of the gambler's wealth can be written as

$$(o_1 b_t)^{Y_t} (o_2 (1 - b_t))^{1 - Y_t}, \quad (3)$$

where $Y_t := \mathbb{I}\{Z_t = 1\} \in \{0, 1\}$ and $Z_t \in \{1, 2\}$ denotes the index of the winning horse. The continuous variation

of the horse race allows the outcome Y_t to be $[0, 1]$ -valued (hence continuous) and defines the per-round gain as

$$o_1 b_t Y_t + o_2 (1 - b_t) (1 - Y_t). \quad (4)$$

Note that (4) is equivalent to (3) for discrete outcomes $Y_t \in \{0, 1\}$. Ryu & Bhatt (2022) noted that the capital process (2) can be viewed as the wealth process from the continuous two-horse race setting with odds $o_1 = \frac{1}{m}$, $o_2 = \frac{1}{1-m}$ and a strategy $(b_t)_{t=1}^\infty$, as the per-round gain can be written as

$$\frac{1}{m} b_t y_t + \frac{1}{1-m} (1 - b_t) (1 - y_t) = 1 + \lambda_t(m) (y_t - m),$$

if we define the *scaled bet*

$$\lambda_t(m) := \frac{b_t}{m(1-m)} - \frac{1}{1-m} \in \left[-\frac{1}{1-m}, \frac{1}{m}\right]. \quad (5)$$

Compared to (2), this reformulation separates out the amount of bet b_t from m and thus the range of the bet is simply $[0, 1]$, being m -independent.

2.3. Extension for Multivariate Processes: Overview

We are now ready to introduce an extension of the betting framework for vector-valued processes following the convention of Ryu & Bhatt (2022). To proceed, we first introduce a general abstract formulation of the gambling as a multiplicative game. Without loss of generality, we assume a unit dollar for the initial wealth of a gambler, i.e., $W_0 = \$1$. We assume that $\mathcal{M} \subseteq \mathbb{R}_+^K$ is given as a set of *odds vectors*. For each round $t \geq 1$, a gambler chooses a bet $\mathbf{b}_t = \mathbf{b}_t(\mathbf{x}^{t-1}) \in \Delta^{K-1}$ as a function of the previous odds vectors $\mathbf{x}^{t-1} \in \mathcal{M}^{t-1}$, where we use $\Delta^{K-1} := \{\mathbf{p} \in \mathbb{R}_{\geq 0}^K : p_1 + \dots + p_K = 1\}$ to denote the $(K-1)$ -dimensional probability simplex in \mathbb{R}^K . After each round, the odds vector $\mathbf{x}_t \in \mathcal{M}$ is revealed, and the gambler's wealth is multiplied by $\mathbf{b}_t^\top \mathbf{x}_t$, and thus the multiplicative gain after round t is

$$W_t(\mathbf{x}^t) = \prod_{i=1}^t \mathbf{b}_i^\top \mathbf{x}_i.$$

As alluded to earlier, it is easy to show that the wealth process $(W_t)_{t=1}^\infty$ attained by any causal strategy is a supermartingale, if the sequence of odds $(\mathbf{X}_t)_{t=1}^\infty$ *subfair*, i.e., $E[\mathbf{X}_t | \mathcal{F}_{t-1}] \leq \mathbf{1}_K$ for any $t \geq 1$, where the inequality holds coordinatewise for $\mathbf{1}_K := [1, \dots, 1] \in \mathbb{R}^K$.

We remark that this abstract setting subsumes the two-horse race HR_2 and continuous two-horse race CHR_2 as special cases, which correspond to certain restricted sets of odds vectors $\mathcal{M} = \{o_1 \mathbf{e}_1, o_2 \mathbf{e}_2\}$ and $\mathcal{M} = \{o_1 y \mathbf{e}_1 + o_2 (1 - y) \mathbf{e}_2 : y \in [0, 1]\}$, respectively. For $K \geq 2$ horses, we can consider the K -horse race HR_K and continuous K -horse race CHR_K , where $\mathcal{M} = \{o_1 \mathbf{e}_1, \dots, o_K \mathbf{e}_K\}$ and

$\mathcal{M} = \{o_1 y_1 \mathbf{e}_1 + \dots + o_K y_K \mathbf{e}_K : \mathbf{y} \in \Delta_K\}$ are assumed, respectively. The most general scenario is when the odds vector \mathbf{x}_t can take any value from the positive orthant, i.e., $\mathcal{M} = \mathbb{R}_{>0}^K$. In this case, the gambling becomes equivalent to the *stock market investment* (Cover, 1991).

The key idea of this paper is to observe that one can consider HR_K and CHR_K to derive CSs for K -categorical data and Δ^{K-1} -valued stochastic processes. Recall that HR_2 and CHR_2 can be used to derive CSs of the mean parameter for $\{0, 1\}$ -valued and $[0, 1]$ -valued stochastic processes, respectively. The extension is natural, once a single number Y_t is mapped into the form of a binary probability vector $[Y_t, 1 - Y_t]$ lying in the probability simplex Δ^1 . In general, for a Δ^{K-1} -valued process $(\mathbf{Y}_t)_{t=1}^\infty$ such that $\mathbb{E}[\mathbf{Y}_t | \mathcal{F}_{t-1}] = \boldsymbol{\mu}$, the gambling with odds vector sequence $(\mathbf{X}_t)_{t=1}^\infty$, which is defined as

$$\mathbf{X}_t := \frac{\mathbf{Y}}{\boldsymbol{\mu}} := \left(\frac{Y_{t1}}{\mu_1}, \dots, \frac{Y_{tK}}{\mu_K} \right), \quad (6)$$

is fair, since $\mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}] = \mathbb{1}$ by construction. Note that the odds vector (6) corresponds to HR_K if $\mathbf{Y}_t \in \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$ (i.e., multinomial, or categorical observations) and CHR_K if $\mathbf{Y}_t \in \Delta_K$ in general. The embedding of the mean vector parameter in (6) is the essence of the multivariate extension.

In the rest of this paper, we build up this idea more concretely. In Section 3, we start by categorical observations, where our key ideas and the application of universal gambling strategies are illustrated with the simple setting of HR_K . In this special case, we also show that when reduced to the sampling without replacement setting, the derived CS yields a provably tighter CS than that of (Waudby-Smith & Ramdas, 2020). We then present its natural continuous extension for probability-vector-valued observations in Section 4, where the underlying gambling becomes CHR_K , a special case of the stock market investment. We also present how a general $[0, 1]^{K-1}$ -valued observations can be handled by a simple reduction to CHR_K . All proofs are deferred to Appendix A.

For a concise treatment of the confidence sequences to be developed, we define a few additional notations. Let $\mathcal{V}_K := \{\mathbf{e}_1, \dots, \mathbf{e}_K\}$ denote the set of vertices of the probability simplex Δ^{K-1} , i.e., the set of one-hot vectors. Let $\mathcal{G}_{K,t} := \{\mathbf{k} \in \mathbb{Z}_{\geq 0}^K : k_1 + \dots + k_K = t\}$ for $t \geq 0$ denote the set of all nonnegative-integer-valued vectors each of whose sum of the coordinates is equal to t . For example, $\mathcal{G}_{K,0} = \{\mathbf{0}_K\}$ and $\mathcal{G}_{K,1} = \mathcal{V}_K$. Since these objects will be frequently referred, we summarize them in Table 1.

3. Categorical Observations

We first consider obtaining a confidence sequence for multinomial (or discrete categorical) observations by considering

Table 1: Summary of notations.

Notation	Definition
\mathcal{V}_K	$\{\mathbf{e}_1, \dots, \mathbf{e}_K\}$
$\mathcal{G}_{K,t}$	$\{\mathbf{k} \in \mathbb{Z}_{\geq 0}^K : k_1 + \dots + k_K = t\}$

the standard K -horse race HR_K . Hereafter, we refer to the horse race with odds $\mathbf{o} := (o_1, \dots, o_K)$ by $\text{HR}_K(\mathbf{o})$.

For a K -categorical stochastic process $(Z_t)_{t=1}^\infty$ such that $\mathbb{E}[\mathbb{1}\{Z_t = j\} | \mathcal{F}_{t-1}] = \mu_j$ for each $j \in [K]$ for some $\boldsymbol{\mu} \in \Delta^{K-1}$, we can define its vector equivalent $\mathbf{Z}_t \in \mathcal{V}_K \subset \Delta_K$, where $Z_{tj} \leftarrow \mathbb{1}\{Z_t = j\}$ for each $j \in [K]$. As hinted in the previous section, we consider the gambling $\text{HR}_K(\mathbf{m}^{-1})$, where the odds vector corresponds to

$$\mathbf{X}_t := \mathbf{m}^{-1} \odot \mathbf{Z}_t = \left(\frac{\mathbb{1}\{Z_{t1} = 1\}}{m_1}, \dots, \frac{\mathbb{1}\{Z_{tK} = 1\}}{m_K} \right).$$

Let $W(\mathbf{Z}^t; \mathbf{o})$ be the cumulative wealth from the horse race $\text{HR}_K(\mathbf{o})$. Since $\mathbb{E}[Z_t | \mathcal{F}_{t-1}] = \boldsymbol{\mu}$, $\mathbb{E}[\mathbf{X}_t | \mathcal{F}_{t-1}] = \mathbb{1}$ for $\mathbf{m} = \boldsymbol{\mu}$, and thus the resulting wealth process is martingale. For $0 < \delta < 1$, define

$$\mathcal{C}(\mathbf{Z}^t; \delta) := \left\{ \mathbf{m} \in \Delta_K : W(\mathbf{Z}^t; \mathbf{m}^{-1}) < \frac{1}{\delta} \right\}. \quad (7)$$

The reasoning in the previous section readily implies:

Theorem 2. *For any causal gambling strategy, $(\mathcal{C}(\mathbf{Z}^t; \delta))_{t=1}^\infty$ is a CS with level $1 - \delta$.*

3.1. Krichevsky–Trofimov CS

Since any causal gambling strategy results in a valid confidence sequence, a natural question is then which gambling strategy should be used to derive tight confidence sequences. In this section, we consider applying a celebrated yet simple *mixture* betting strategy from universal gambling. The resulting CS is easy to compute, naturally convex, and always nonvacuous even after observing one sample. As a shorthand notation, for two vectors \mathbf{a} and \mathbf{b} of same dimension K , we denote $\mathbf{a}^{\mathbf{b}} := \prod_{j=1}^K a_j^{b_j}$.

The celebrated idea of universal gambling from information theory is essentially to utilize a continuous mixture of the wealth processes of constant bettors. Consider the cumulative wealth attained by a constant bettor \mathbf{b} from $\text{HR}_K(\mathbf{o})$:

$$W^{\mathbf{b}}(\mathbf{z}^t; \mathbf{o}) := \prod_{i=1}^t \prod_{j=1}^K (o_j b_j)^{\mathbb{1}\{z_{ti}=1\}} = (\mathbf{o} \odot \mathbf{b})^{\mathbf{k}(\mathbf{z}^t)}.$$

Here, we define the cumulative count vector as

$$\mathbf{k}(\mathbf{z}^t) := \left(\sum_{i=1}^t \mathbb{1}\{z_{ij} = 1\} \right)_{j \in [K]} = \sum_{i=1}^t \mathbf{z}_i.$$

While any mixture of this wealth process induces a valid wealth process, we specifically consider the so-called Krichevsky–Trofimov (KT) strategy (Krichevsky & Trofimov, 1981), which takes a continuous mixture of the wealth of constant bettors with respect to the Dirichlet distribution $\text{Dir}(\mathbf{b}; \boldsymbol{\alpha}) := \frac{1}{B(\boldsymbol{\alpha})} \prod_{i=1}^K b_i^{\alpha_i - 1}$. Here, $B(\boldsymbol{\alpha}) := (\prod_{i=1}^K \Gamma(\alpha_i)) / \Gamma(\sum_{i=1}^K \alpha_i)$ denotes the multivariate beta function and $\Gamma(\alpha)$ denotes the gamma function. That is,

$$\begin{aligned} W^{\text{KT}}(\mathbf{z}^t; \mathbf{m}^{-1}) &:= \int_{\Delta^{K-1}} W^{\mathbf{b}}(\mathbf{z}^t; \mathbf{m}^{-1}) \text{Dir}(\mathbf{b}; \boldsymbol{\alpha}) d\mathbf{b} \\ &= \mathbf{m}^{-\mathbf{k}(\mathbf{z}^t)} \frac{B(\mathbf{k}(\mathbf{z}^t) + \boldsymbol{\alpha})}{B(\boldsymbol{\alpha})}. \end{aligned} \quad (8)$$

The KT mixture specifically sets $\boldsymbol{\alpha} = \frac{\mathbf{1}_K}{2}$. This specific choice is the Jeffreys prior for multinomial distributions (Yang & Berger, 1996). Krichevsky & Trofimov (1981) proved, in the equivalent language of universal coding, that the KT mixture wealth can achieve the optimal wealth achieved by the best constant bettor in hindsight with a minimax optimal rate in the regret in log wealth ratios under a stochastic market assumption. It is also later shown that the optimality even holds for individual sequences (Xie & Barron, 2000).

This wealth process is easy to compute for any given \mathbf{m} , since it is only a function of the count vector $\mathbf{k}(\mathbf{z}^t)$, which can be viewed as the sufficient statistics. We can further manipulate this expression to better understand the behavior of this wealth process. First, we define the empirical mean vector $\hat{\boldsymbol{\mu}}_t := \frac{1}{t} \sum_{i=1}^t \mathbf{z}_i = \frac{\mathbf{k}(\mathbf{z}^t)}{t}$. Second, we define $q^{\text{KT}}(\mathbf{x}) := \frac{B(\mathbf{x} + \boldsymbol{\alpha})}{B(\boldsymbol{\alpha})}$ for $\mathbf{x} \in \mathbb{R}_{\geq 0}^K$. Then, the corresponding confidence set $\mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta)$ can be written as

$$\begin{aligned} \mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta) &= \left\{ \mathbf{m}: \hat{\boldsymbol{\mu}}_t^\top \log \frac{1}{\mathbf{m}} < \frac{1}{t} \log \frac{1}{\delta} + \frac{1}{t} \log \frac{1}{q^{\text{KT}}(t\hat{\boldsymbol{\mu}}_t)} \right\} \\ &= \left\{ \mathbf{m}: D(\hat{\boldsymbol{\mu}}_t \| \mathbf{m}) < \frac{1}{t} \log \frac{1}{\delta} + \frac{1}{t} \log \frac{e^{-tH(\hat{\boldsymbol{\mu}}_t)}}{q^{\text{KT}}(t\hat{\boldsymbol{\mu}}_t)} \right\}, \end{aligned}$$

where $D(\mathbf{p} \| \mathbf{q}) := \mathbf{p}^\top \log \frac{\mathbf{p}}{\mathbf{q}}$ denotes the Kullback–Leibler divergence between $\mathbf{p}, \mathbf{q} \in \Delta^{K-1}$ and $H(\mathbf{p}) := -\mathbf{p}^\top \log \mathbf{p}$ denotes the Shannon entropy for a probability vector $\mathbf{p} \in \Delta^{K-1}$. The following properties are now immediate.

Theorem 3 (KT CS). *Pick any $\delta \in (0, 1)$.*

- (a) (Convexity) $\mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta) \subset \Delta^{K-1}$ is always convex.
- (b) (Always-non-vacuous guarantee) If $\mathbf{m} \in \mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta)$,

$$m_j > \frac{k_j(\mathbf{z}^t)}{t} (\delta q^{\text{KT}}(\mathbf{k}(\mathbf{z}^t)))^{\frac{1}{t}}$$

for any $j \geq 1$. In particular, for any $j \in [K]$, $\mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta)$ is strictly bounded away from the hyperplane $\{\mathbf{m}: m_j = 0\}$ as soon as $\min_{i \in [t]} z_{ij} = 1$.

- (c) (Large-sample behavior) For t sufficiently large, $\mathcal{C}^{\text{KT}}(\mathbf{Z}^t; \delta)$ asymptotically behaves as

$$\left\{ \mathbf{m}: D(\hat{\boldsymbol{\mu}}_t \| \mathbf{m}) < \frac{1}{t} \log \frac{1}{\delta} + \frac{K-1}{2t} \log t + o(1) \right\},$$

where $o(1)$ is a vanishing term as $t \rightarrow \infty$.

3.1.1. BASELINES

Instead of the multi-horse-race approach, one can construct a confidence set for categorical data by properly aggregating confidence intervals constructed for each coordinate, e.g., using the existing gambling techniques for scalar-valued processes. Below, we empirically demonstrate the benefit of the proposed approach compared to such baselines.

More concretely, for each coordinate $j \in [K]$, we can construct a confidence set with level $1 - \delta'$ for μ_j , say $\mathcal{C}_{\delta'}^{(j)}$, based on a binary sequence $(\mathbb{I}\{z_{ij} = 1\})_{i=1}^t$. To construct a single confidence set for $\boldsymbol{\mu}$ from such K confidence sets, one straightforward choice is the well-known *Bonferroni correction*, which takes the intersection of K CSs constructed with level $1 - \delta'$ with $\delta' = \frac{\delta}{K}$ for each coordinate:

$$\mathcal{C}_\delta^{\text{bonf}} := \bigcap_{j=1}^K \mathcal{C}_{\frac{\delta}{K}}^{(j)}.$$

By a union bound, it readily follows that the resulting CS is a CS with level $1 - \delta$.

For gambling-based (or martingale-based) CSs, there exists yet another natural way to aggregate CSs from different wealth sequences. Suppose that we are given K wealth sequences $(W_t^{(j)})_{t=0}^\infty$ for $j \in [K]$, each of which induces a level- $(1 - \delta')$ CS as

$$\mathcal{C}_{\delta'}^{(j)} := \left\{ \mathbf{m}: W_t^{(j)}(\mathbf{m}) \leq \frac{1}{\delta'} \right\}.$$

We can then aggregate the K sets by a *mixture wealth*, i.e.,

$$\mathcal{C}_\delta^{\text{mix}} := \left\{ \mathbf{m}: \frac{1}{K} \sum_{j=1}^K W_t^{(j)}(\mathbf{m}) \leq \frac{1}{\delta} \right\}.$$

Then, it is straightforward that the Bonferroni correction is provably not smaller than the mixture-based aggregation:

Proposition 4. $\mathcal{C}_\delta^{\text{mix}} \subseteq \mathcal{C}_\delta^{\text{bonf}}$.

Proof. If $\mathbf{m} \in \mathcal{C}_\delta^{\text{mix}}$, then $W_t^{(j)}(\mathbf{m}) \leq \sum_{j'=1}^K W_t^{(j')}(\mathbf{m}) \leq \frac{K}{\delta}$ for every $j \in [K]$, which implies $\mathbf{m} \in \mathcal{C}_\delta^{\text{bonf}}$. \square

We note that a concurrent work by Cho et al. (2024, Lemma 4) essentially discusses these two baselines and argues the tightness of the mixture-based aggregation, while we independently rediscovered this during the revision of our manuscript.

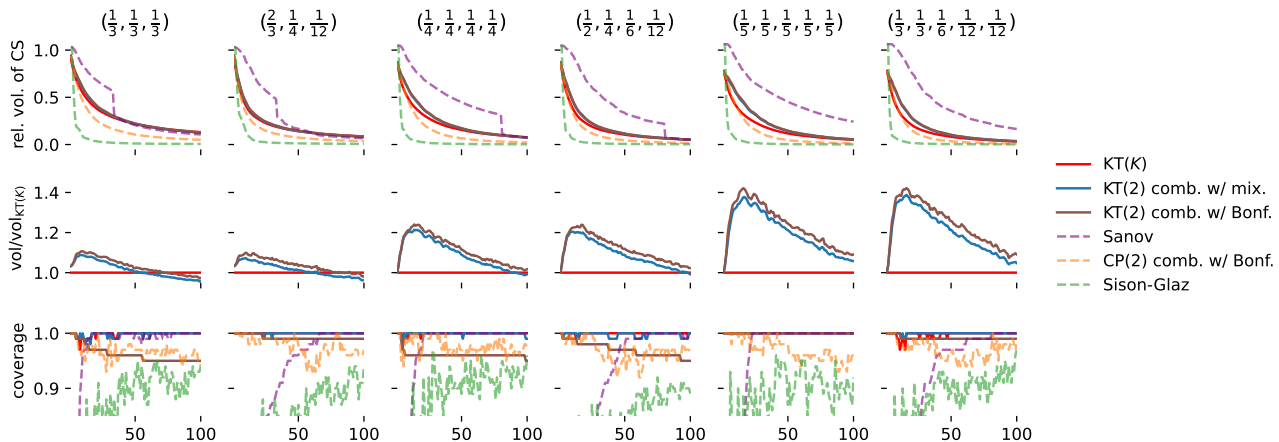


Figure 1: Performance of time-uniform confidence sets (solid lines) and non-time-uniform confidence sets (dashed lines) with respect to i.i.d. categorical data with mean vector $\mu \in \Delta^{K-1}$ (indicated as the titles), in terms of the relative volume of the confidence sets (first row) and the per-time-step coverage (third row). The second row visualizes how the baselines KT(2) w/ mix. and KT(2) w/ Bonf. are compared to KT(K).

3.1.2. SIMULATION²

To demonstrate the performance of the KT CS (denoted as $\text{KT}(K)$), we run simulations with i.i.d. categorical observations of length $T = 100$, with different mean vectors of dimension $K \in \{3, 4, 5\}$; see Fig. 1, where the title of each column indicates the underlying mean. Each experiment was run with 100 random realizations and $\delta = 0.05$ was used. In the first row, we plot the relative volume of the CS over Δ^{K-1} . As another CS baseline, we consider applying the mixture aggregation and Bonferroni correction to combine $\text{KT}(2)$'s induced by each coordinate $j \in [K]$ with $\delta' = \delta/K$; check $\text{KT}(2)$ w/ mix. and $\text{KT}(2)$ w/ Bonf. in the figure. Note that, as argued above, mixture-based aggregations are tighter than Bonferroni, though marginally. While the difference is not significant in $K = 3$, we observe that the multivariate version $\text{KT}(K)$ exhibits much tighter behavior as K increases, as visualized in the second row. Interestingly, however, the baselines seem to outperform in a long run.

Exploring further, we compare the performance of the KT CS to the non-time-uniform confidence intervals (shown in dashed lines with transparency). We first note that the asymptotic behavior of KT CS resembles Sanov's theorem (Cover & Thomas, 2006, Thm. 11.4.1), which can be proved by the method of types: for each $t \geq 1$,

$$\mathbb{P}\left\{D(\hat{\mu}_t \parallel \mu) < \frac{1}{t} \log \frac{1}{\delta} + \frac{K}{t} \log(t+1)\right\} \geq 1 - \delta. \quad (9)$$

Note that it has a worse redundant term than KT CS. [Mardia](#)

²We have open-sourced the implementation of the proposed confidence sequences and provided the codes to reproduce the simulation results online: <https://github.com/jongharyu/confidence-sequence-via-gambling>.

et al. (2020) proposed several improved bounds of a similar flavor; for example, for each $t \geq 8\pi(K/e)^3$,

$$\mathbb{P}\left\{D(\hat{\mu}_t \parallel \mu) < \frac{K-1}{t} \log \frac{2(K-1)}{\delta}\right\} \geq 1 - \delta, \quad (10)$$

which removes the $O(\frac{1}{t} \log t)$ term. Note that the required number of samples for this improved bound to kick in is $t = \Omega(K^3)$ and grows fast in K . As shown in the simulation above, KT is almost always tighter than Sanov and on par after the improved bound kicks in, while KT CS provides a much stronger time-uniform guarantee.

We also compare with the combined Clopper–Pearson intervals (Clopper & Pearson, 1934) with Bonferroni correction (CP(2) + Bonf) and the Sison–Glaz method (Sison & Glaz, 1995). We note that CP(2) + Bonf provides consistently tighter confidence sets with guaranteed coverage, but without the time-uniform guarantee³. While the Sison–Glaz method outputs extremely small confidence sets, the coverage is much below than the desired threshold, since the method was designed based on the asymptotic approximation. This demonstrates that the performance of KT CS is even comparable to the existing non-time-uniform bounds.

3.2. CS for Sampling Without Replacement

In this section, we specifically consider the sampling *without replacement* (WoR) setup, where we aim to construct a tight, vanishing confidence sequence to the mean of finite samples with sampling WoR. The main result in this section is that if one adapts the KT CS for this setting via a general reduction

³The *coverage* (or *coverage probability*) is the probability that the confidence set contains the target parameter. In the experiments, we compute the empirical coverage at each time instance, i.e., non-time-uniform coverage, over the repeated experiments.

to be described below, the resulting CS is always tighter than the posterior-prior martingale proposed by [Waudby-Smith & Ramdas \(2020\)](#).

Suppose that there exists a finite set of N multinomial observations $\{\mathbf{z}_1, \dots, \mathbf{z}_N\} \subset \mathcal{V}_K$, which are fixed and nonrandom. The underlying population is now fixed, and the goal here is to estimate the population mean $\boldsymbol{\mu} := \frac{1}{N} \sum_{i=1}^N \mathbf{z}_i$ by sequentially drawing an observation from the population uniformly at random WoR. Formally, the probabilistic model for the sampling WoR procedure is

$$\mathbf{Z}_t | \mathbf{Z}^{t-1} \sim \text{Unif}(\{\mathbf{z}_1, \dots, \mathbf{z}_N\} \setminus \{\mathbf{Z}_1, \dots, \mathbf{Z}_{t-1}\}).$$

We note that this setup has important practical applications; for example, a tight WoR confidence sequence can lead to a fast decision in the real-world auditing for votes ([Waudby-Smith et al., 2021](#)). As argued in ([Waudby-Smith & Ramdas, 2024](#)), there exists a general reduction that maps any CS for bounded random processes to CS for sampling WoR. For each $t \geq 1$, after having observed \mathbf{Z}^{t-1} , the mean of the remaining samples is

$$\boldsymbol{\mu}_{N,t}^{\text{WoR}}(\mathbf{Z}^{t-1}) := \mathbb{E}[\mathbf{Z}_t | \mathcal{F}^{t-1}] = \frac{N\boldsymbol{\mu} - \sum_{i=1}^{t-1} \mathbf{Z}_i}{N - (t-1)}.$$

For any $\mathbf{m} \in \Delta^{K-1}$, we define an analogous transform

$$\mathbf{m}_{N,t}^{\text{WoR}}(\mathbf{Z}^{t-1}) := \frac{N\mathbf{m} - \sum_{i=1}^{t-1} \mathbf{Z}_i}{N - (t-1)}.$$

To testing if the underlying population mean is \mathbf{m} , we can plug-in $\mathbf{m}_{N,t+1}^{\text{WoR}}(\mathbf{Z}^t)$ in place of \mathbf{m} to the definition of any valid confidence sequence $(\mathcal{C}(\mathbf{z}^t; \delta))_{t=1}^\infty$ for bounded vectors; we defer the statement Proposition 10 to Appendix B.

Concretely, we can adapt the KT CS to this setup by only plugging-in $\mathbf{m}_{N,t+1}^{\text{WoR}}(\mathbf{Z}^t)$ in place of \mathbf{m} in the definition of the wealth process (8), which leads to

$$\begin{aligned} W_{\text{WoR}}^{\text{KT}}(\mathbf{Z}^t; \mathbf{m}) &:= W^{\text{KT}}(\mathbf{Z}^t; \mathbf{m}_{N,t+1}^{\text{WoR}}(\mathbf{Z}^t)) \\ &= \frac{B(\mathbf{K}_t + \boldsymbol{\alpha})}{B(\boldsymbol{\alpha})} \left(\frac{N\mathbf{m} - \mathbf{K}_t}{N - t} \right)^{-\mathbf{K}_t}, \end{aligned}$$

where we define $\mathbf{K}_t := \mathbf{k}(\mathbf{Z}^t) = \sum_{i=1}^t \mathbf{Z}_i$ as a shorthand. Here, we allow $\boldsymbol{\alpha} \in \mathbb{R}_{>0}^K$ to be a hyperparameter.

[Waudby-Smith & Ramdas \(2020\)](#) proposed a CS for categorical observations based on the process

$$R_{\text{WoR}}(\mathbf{Z}^t; \mathbf{m}) := \frac{B(\mathbf{K}_t + \boldsymbol{\alpha})}{B(\boldsymbol{\alpha})} \frac{\binom{N}{N\mathbf{m}}}{\binom{N-t}{N\mathbf{m} - \mathbf{K}_t}} \quad (11)$$

$\mathbf{m} \in \mathcal{G}_{K,N}$, i.e., $\mathbf{m} \in \mathbb{Z}_{\geq 0}^K$ and $\mathbf{m}/N \in \Delta^{K-1}$, where $\boldsymbol{\alpha} \in \mathbb{R}_{>0}^K$ is again a hyperparameter. This sequence is derived as a posterior-prior ratio (PPR) under a hypergeometric

observation model, and thus we refer to the resulting CS as PPR CS. Interestingly, KT CS is always tighter than or equal to PPR CS with probability 1 as implied by the following theorem. Recall that a CS becomes tighter as the underlying sequence diverges faster.

Theorem 5. *Pick any $0 \leq t \leq N$ and $\tilde{\mathbf{z}}^t \in \mathcal{V}_K^t$. For any $\mathbf{m} \in \mathcal{G}_{K,N}$, we have*

$$W_{\text{WoR}}^{\text{KT}}(\tilde{\mathbf{z}}^t; \mathbf{m}) \geq R_{\text{WoR}}(\tilde{\mathbf{z}}^t; \mathbf{m}).$$

3.2.1. SIMULATION

We demonstrate the tightness of the (WoR-) KT CS compared to the PPR CS in Fig. 2. We consider a finite population of 1000 balls consisting of 600 red, 250 green, and 150 blue balls, and report the result averaged over 1000 random permutations. While both KT and PPR CSs naturally converge to the singleton containing the true population after observing all balls, KT can be relatively much tighter than PPR, as shown in the second panel. If one wishes to decide the rank of the colors by acquiring balls sequentially, the stopping time, which is defined as the first time when all colors can be distinguished from all the others with certainty, is the practical measure for the quality of CS. In the third panel, we visualize the histogram of the ratio of the stopping times between KT and PPR; on average, KT CS can decide the rank about 26% faster than PPR CS.

4. Probability-Vector-Valued Observations

We now assume that $(\mathbf{Y}_t)_{t=1}^\infty$ is a Δ^{K-1} -valued stochastic process such that $\mathbb{E}[\mathbf{Y}_t | \mathcal{F}_{t-1}] \equiv \boldsymbol{\mu}$ for some $\boldsymbol{\mu} \in \Delta^{K-1}$. The development is in parallel Section 3 for categorical observations. We only need to replace $\text{HR}_K(\mathbf{m}^{-1})$ with $\text{CHR}_K(\mathbf{m}^{-1})$, where the odds vector $\mathbf{X}_t := \mathbf{m}^{-1} \odot \mathbf{Y}_t$ now can take any value in the convex hull of $\{\frac{\mathbf{e}_1}{m_1}, \dots, \frac{\mathbf{e}_K}{m_K}\}$. With a slight abuse of notation, we also denote the cumulative wealth from the horse race $\text{CHR}_K(\mathbf{o})$ by $W(\mathbf{Y}^t; \mathbf{o})$. It then readily follows that the same definition in (7) is a $(1 - \delta)$ -CS as shown in Theorem 2:

Theorem 6. *For any causal gambling strategy, $(\mathcal{C}(\mathbf{Y}^t; \delta))_{t=1}^\infty$ is a CS with level $1 - \delta$.*

4.1. Universal-Portfolio CS

As previously for categorical observations and the HR-based CS construction, any causal gambling strategy leads to a valid CS. In the previous section, we demonstrated some nice properties of the KT CS, which is induced by a universal gambling strategy based on the KT mixture. Analogously, we will also explore the idea of universal gambling for probability-vector-valued observations. The idea of universal gambling in this continuous case is same: aiming to achieve the best wealth attained by any constant bettor, consider a mixture of the wealth of all constant bettors. If

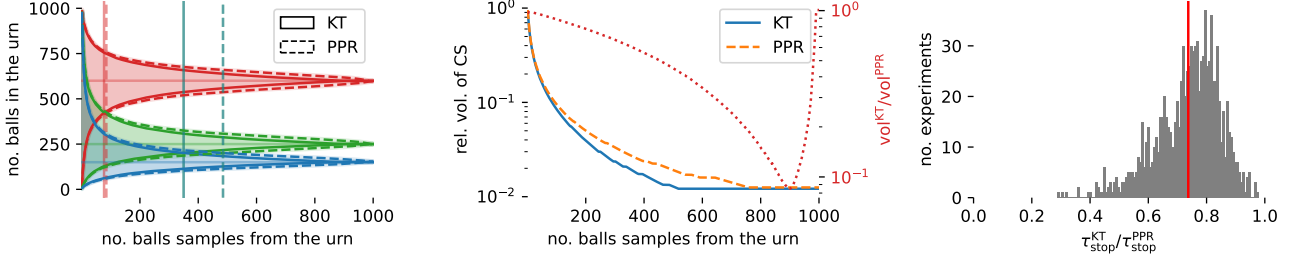


Figure 2: Simulation of KT and PPR CSs for sampling without replacement (WoR). The underlying population consists of 1000 balls (600 red, 250 green, and 150 blue). The plotted results are averaged over random experiments with 1000 random permutations. The vertical lines in the first panel visualize the stopping time for each color when the confidence interval for the number of the balls of the particular color becomes disjoint from all the other confidence intervals. Note, however, that the curves shown here are averaged over random trials and the stopping times indicated here are only for an illustrative purpose. The second panel visualizes the relative volume of the CSs at each time step. The red, dotted line in the second panel shows the ratio of the volume of KT CS to that of PPR CS. The last panel presents a histogram of the ratios of the final stopping time of KT to PPR.

we consider the KT mixture with $\text{Dir}(\mathbf{b}; \boldsymbol{\alpha})$ with $\boldsymbol{\alpha} = \frac{\mathbb{1}_K}{2}$, the universal gambling strategy is Cover’s universal portfolio (UP) algorithm (Cover, 1991; Cover & Ordentlich, 1996). For the scalar case ($K = 2$), constructing a CS using UP was studied and its excellent empirical performance especially in a small-sample regime was demonstrated by (Orabona & Jun, 2024; Ryu & Bhatt, 2022). Here, we demonstrate a similar property in the multidimensional case.

While the conceptual development is exactly the same, the difference arises in the definition of the wealth. We start by considering the cumulative wealth attained by a constant bettor \mathbf{b} from $\text{CHR}_K(\mathbf{o})$ for a Δ_K -valued sequence \mathbf{y}^t :

$$W^{\mathbf{b}}(\mathbf{y}^t; \mathbf{o}) := \prod_{i=1}^t \sum_{j=1}^K o_j b_j y_{tj}. \quad (12)$$

$$\begin{aligned} &\stackrel{(\star)}{=} \sum_{\boldsymbol{\eta}^t \in [K]^t} \prod_{i=1}^t (o_{\eta_i} b_{\eta_i}) y_{i\eta_i} \\ &= \sum_{\mathbf{z}^t \in \mathcal{V}_K^t} \prod_{i=1}^t (\mathbf{o} \odot \mathbf{b} \odot \mathbf{y}_i)^{\mathbf{z}_i} \\ &= \sum_{\mathbf{k} \in \mathcal{G}_{K,t}} \mathbf{y}^t[\mathbf{k}] (\mathbf{o} \odot \mathbf{b})^{\mathbf{k}}. \end{aligned} \quad (13)$$

Here, noting that the wealth (12) is a product of summations, we rewrite it as a summation of products by applying the distributive law in (\star) . We then further simplify the expression with the notations \mathcal{V}_K and $\mathcal{G}_{K,t}$ defined in Table 1. For $\mathbf{k} \in \mathcal{G}_{K,t}$, we also define in (13)

$$\mathbf{y}^t[\mathbf{k}] := \sum_{\mathbf{z}^t \in \mathcal{V}_K^t: \mathbf{k}(\mathbf{z}^t) = \mathbf{k}} \prod_{i=1}^t y_i^{\mathbf{z}_i}.$$

Note that \mathbf{z}_i ’s are one-hot vectors and thus $\mathbf{y}_i^{\mathbf{z}_i} = \prod_{j=1}^K y_{i,z_{ij}}$ simply reads out the coordinate of \mathbf{y}_i where \mathbf{z}_i has 1. As a sanity check, if $\mathbf{y} \in \mathcal{V}_K$, the wealth simplifies to $W^{\mathbf{b}}(\mathbf{y}^t; \mathbf{o}) = (\mathbf{o} \odot \mathbf{b})^{\mathbf{k}(\mathbf{y}^t)}$ as in $\text{HR}_K(\mathbf{o})$.

Now, if we consider the mixture of the wealth of constant bettors \mathbf{b} with respect to a Dirichlet distribution $\text{Dir}(\mathbf{b}; \boldsymbol{\alpha})$, the resulting wealth is the wealth achieved by Cover’s UP:

$$\begin{aligned} W^{\text{UP}}(\mathbf{y}^t; \mathbf{o}) &:= \int_{\Delta_{K-1}} W_t^{\mathbf{b}}(\mathbf{y}^t; \mathbf{o}) \text{Dir}(\mathbf{b}; \boldsymbol{\alpha}) d\mathbf{b} \\ &= \sum_{\mathbf{k} \in \mathcal{G}_{K,t}} \mathbf{y}^t[\mathbf{k}] \mathbf{o}^{\mathbf{k}} \frac{B(\mathbf{k} + \boldsymbol{\alpha})}{B(\boldsymbol{\alpha})}, \end{aligned}$$

since $\int_{\Delta_{K-1}} \mathbf{b}^{\mathbf{k}} \text{Dir}(\mathbf{b}; \boldsymbol{\alpha}) d\mathbf{b} = B(\mathbf{k} + \boldsymbol{\alpha})/B(\boldsymbol{\alpha})$. Note that the summation is over $\mathcal{G}_{K,t}$ which has cardinality $\binom{t+K-1}{t} = O(t^{K-1})$, and thus the summation can be computed in $O(t^{K-1})$ with memory complexity $O(t^{K-1})$ at each round t , by maintaining and updating the statistics $\mathbf{y}^t[\mathbf{k}]$ for $\mathbf{k} \in \mathcal{G}_{K,t}$ by

$$\mathbf{y}^t[\mathbf{k}] = \sum_{j=1}^K y_{tj} \mathbf{y}^{t-1}[\mathbf{k} - \mathbf{e}_j] \mathbb{I}\{k_j \geq 1\}. \quad (14)$$

after observing \mathbf{y}_t at round $t \geq 1$. We note that this dynamic programming argument was essentially discussed in Cover & Ordentlich (1996, Section VI).

Let $(\mathcal{C}^{\text{UP}}(\mathbf{Y}^t; \delta))_{t=1}^{\infty}$ denote the resulting $(1 - \delta)$ -CS from the UP wealth $W^{\text{UP}}(\mathbf{y}^t; \mathbf{m}^{-1})$, i.e.,

$$\mathcal{C}^{\text{UP}}(\mathbf{y}^t; \delta) := \left\{ \mathbf{m} \in \Delta_K : W(\mathbf{y}^t; \mathbf{m}^{-1}) < \frac{1}{\delta} \right\}. \quad (15)$$

We formally state the properties of UP CS as in Theorem 3:

Theorem 7 (UP CS). *Pick any $\delta \in (0, 1)$.*

- (a) (Convexity) $\mathcal{C}_t^{\text{UP}}(\mathbf{y}^t; \delta) \subset \Delta^{K-1}$ is always convex.
- (b) (Always-non-vacuous guarantee) For any $j \in [K]$, $\mathcal{C}^{\text{UP}}(\mathbf{y}^t; \delta)$ is strictly bounded away from the hyperplane $\{\mathbf{m} : m_j = 0\}$ as soon as $\max_{i \in [t]} y_{ij} > 0$.
- (c) $\hat{\boldsymbol{\mu}}_t \in \mathcal{C}_t^{\text{UP}}(\mathbf{Y}^t)$ with probability 1.




Figure 3: A similar experiment to Fig. 1 for i.i.d. Dirichlet observations for $K \in \{3, 4\}$.

Orabona & Jun (2024) established a similar property to (b) for UP CS in the scalar-valued case; (c) is a generalization of Ryu & Bhatt (2022, Theorem 4.3(c)).

4.1.1. SIMULATION

We conducted a simulation similar to Fig. 1, but for i.i.d. Dirichlet observations of dimension $K \in \{3, 4\}$ with different concentration parameters. As baselines, we tested the aggregated UP(2) CSs with mixture and Bonferroni correction, as well as $KT(K)$. Compared to $KT(K)$, $UP(K)$ is much tighter in general, as the distribution gets more concentrated (small variance). Similar to the observation in Fig. 1, the gap between UP(2) w/ mix. and $UP(K)$ increases as the dimension K increases.

4.2. Reduction for Bounded Vectors

Finally, the technique developed so far can be applied to construct a CS for $[0, 1]^{K-1}$ -valued observations for $K \geq 2$ by a simple reduction described below. Suppose that we observe $\mathbf{Y}_t = (Y_{t1}, \dots, Y_{t,K-1}) \in [0, 1]^{K-1}$ for each $t \geq 1$, which satisfies $\mathbb{E}[\mathbf{Y}_t | \mathcal{F}_{t-1}] \equiv \boldsymbol{\mu}$ for some $\boldsymbol{\mu} \in [0, 1]^{K-1}$. To construct a confidence sequence for $\boldsymbol{\mu}$, we can consider its probability-vector representation $\tilde{\mathbf{Y}}_t = (\tilde{Y}_{t1}, \dots, \tilde{Y}_{tK}) \in \Delta^K$, which is defined as

$$\tilde{Y}_{tK} := 1 - \frac{1}{K-1}(Y_{t1} + \dots + Y_{t,K-1}) \quad (16)$$

and $\tilde{Y}_{tj} := Y_{tj}/(K-1)$ for $j \in [K-1]$. Note that $\tilde{\mathbf{Y}}_t$ lies on the K -dimensional probability simplex Δ^{K-1} by construction and now its conditional mean is

$$\mathbb{E}[\tilde{\mathbf{Y}}_t | \mathcal{F}_{t-1}] = \begin{bmatrix} \frac{\boldsymbol{\mu}}{K-1} \\ 1 - \frac{1}{K-1}\boldsymbol{\mu} \end{bmatrix} \in \Delta^{K-1}.$$

Now we can apply the UP CS on the sequence $(\tilde{\mathbf{Y}}_t)_{t=1}^\infty$ to construct a confidence sequence for $\boldsymbol{\mu}$.

5. Concluding Remarks

In this paper, we presented a general gambling framework for constructing confidence sequences and established some

theoretical properties of CSs from universal gambling strategies. Our simulations show that compared to the baseline CSs constructed with Bonferroni correction and averaged wealth, the multivariate gambling CSs are more conservative (i.e., of higher coverage for each time instance), provide tighter confidence sets in the small-sample regime, and the performance margin increases as the dimensionality increases. Such a tight confidence set construction may lead to faster decisions in sequential decision-making problems, i.e., as we demonstrated in Section 3.2. We remark that one can consider applying the KT CS for multi-arm bandits as demonstrated in (Malloy et al., 2020, Section 4.4).

While UP CS provides extremely tight confidence sets with guaranteed coverage even with small samples, the prohibitive computational complexity $O(T^K)$ for K -dimensional probability vectors of length T makes UP CS less applicable in practice for high-dimensional data. Efficient approximation of UP has been an open problem in information theory and theoretical computer science for decades (van Erven et al., 2020). For example, (Kalai & Vempala, 2002) proposed efficient algorithms for approximate computation of UP based on Monte Carlo Markov chain methods, and several developments have been made recently, including (Luo et al., 2018; van Erven et al., 2020; Mhammedi & Rakhlin, 2022; Zimmert et al., 2022; Jézéquel et al., 2022; Tsai et al., 2024). The exploration of efficient approximate UP methods and the computational-statistical trade-off is left for future research. In the context of CSs, finding an efficient yet tight approximation of UP in the small-sample regime would be of great practical importance.

An experienced reader may wonder if the proposed gambling technique is applicable to sequential kernel regression (or kernel bandit) (Vakili et al., 2021), or even to simpler linear bandits (Abbasi-Yadkori et al., 2011). While this extension is intriguing, nontrivial technicalities arise. First, the proposed methods in this paper only apply to bounded vectors, requiring a suitable problem setting in such an extension. Second, the UP-based confidence sets are implicitly defined, needing a tight analytical outer bound or an efficient numerical procedure for implementation. We leave this direction for future work.

Acknowledgements

The authors appreciate insightful feedback and suggestions from anonymous reviewers to improve the manuscript. This work was supported, in part, by the MIT-IBM Watson AI Lab under Agreement No. W1771646, and by MIT Lincoln Laboratory.

Impact Statement

This paper presents work whose goal is to advance the field of Sequential Decision Making. There are many potential societal consequences of our work, none of which we feel must be specifically highlighted here.

References

- Abbasi-Yadkori, Y., Pál, D., and Szepesvári, C. Improved algorithms for linear stochastic bandits. In *Adv. Neural Inf. Proc. Syst.*, volume 24, 2011. 9
- Cho, B., Gan, K., and Kallus, N. Peeking with PEAK: Sequential, nonparametric composite hypothesis tests for means of multiple data streams. *arXiv preprint arXiv:2402.06122*, 2024. 5
- Chugg, B., Wang, H., and Ramdas, A. Time-uniform confidence spheres for means of random vectors. *arXiv preprint arXiv:2311.08168*, 2023. 1
- Clopper, C. J. and Pearson, E. S. The use of confidence or fiducial limits illustrated in the case of the binomial. *Biometrika*, 26(4):404–413, 1934. 6
- Cover, T. M. Universal portfolios. *Math. Financ.*, 1(1):1–29, 1991. 1, 2, 4, 8
- Cover, T. M. and Ordentlich, E. Universal portfolios with side information. *IEEE Trans. Inf. Theory*, 42(2):348–363, 1996. 8
- Cover, T. M. and Thomas, J. A. *Elements of information theory*. John Wiley & Sons, 2006. 3, 6
- Darling, D. A. and Robbins, H. Confidence sequences for mean, variance, and median. *Proc. Natl. Acad. Sci. U. S. A.*, 58(1):66, 1967. 1
- Evidently. How evidently calculates results, 2023. URL <https://docs.aws.amazon.com/AmazonCloudWatch/latest/monitoring/CloudWatch-Evidently-calculate-results.html>. 1
- Hoeffding, W. Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.*, 58(301):13–30, 1963. 1
- Howard, S. R., Ramdas, A., McAuliffe, J., and Sekhon, J. Time-uniform, nonparametric, nonasymptotic confidence sequences. *Ann. Statist.*, 49(2):1055–1080, 2021. 1
- Jézéquel, R., Ostrovskii, D. M., and Gaillard, P. Efficient and near-optimal online portfolio selection. *arXiv e-prints*, pp. arXiv–2209, 2022. 9
- Jun, K.-S. and Orabona, F. Parameter-free online convex optimization with sub-exponential noise. In *Conf. Learn. Theory*, pp. 1802–1823. PMLR, 2019. 1
- Kalai, A. T. and Vempala, S. Efficient algorithms for universal portfolios. *J. Mach. Learn. Res.*, pp. 423–440, 2002. 9
- Krichevsky, R. and Trofimov, V. The performance of universal encoding. *IEEE Trans. Inf. Theory*, 27(2):199–207, 1981. 5
- Lai, T. L. On confidence sequences. *Ann. Statist.*, 4(2):265–280, 1976. 1
- Li, J., Li, Y., and Dai, X. Li, Li, and Dai’s Contribution to the Discussion of “Estimating Means of Bounded Random Variables by Betting” by Waudby-Smith and Ramdas. *J. R. Stat. Soc. Series B Stat. Methodol.*, pp. qkad111, October 2023. ISSN 1369-7412. doi: 10.1093/jrsssb/qkad111. 3
- Luo, H., Wei, C.-Y., and Zheng, K. Efficient online portfolio with logarithmic regret. In *Adv. Neural Inf. Proc. Syst.*, volume 31, 2018. 9
- Malloy, M. L., Tripathy, A., and Nowak, R. D. Optimal Confidence Regions for the Multinomial Parameter. *arXiv*, February 2020. 9
- Manole, T. and Ramdas, A. Martingale methods for sequential estimation of convex functionals and divergences. *IEEE Trans. Inf. Theory*, 69(7):4641–4658, 2023. doi: 10.1109/TIT.2023.3250099. 1
- Mardia, J., Jiao, J., Tónczos, E., Nowak, R. D., and Weissman, T. Concentration inequalities for the empirical distribution of discrete distributions: beyond the method of types. *Inf. Inference*, 9(4):813–850, 2020. 6
- Mhammedi, Z. and Rakhlin, A. Damped online Newton step for portfolio selection. *arXiv preprint arXiv:2202.07574*, 2022. 9
- Orabona, F. and Jun, K.-S. Tight concentrations and confidence sequences from the regret of universal portfolio. *IEEE Trans. Inf. Theory*, 70(1):436–455, 2024. doi: 10.1109/TIT.2023.3330187. arXiv:2110.14099. 1, 2, 3, 8, 9

- Ramdas, A., Ruf, J., Larsson, M., and Koolen, W. Admissible anytime-valid sequential inference must rely on non-negative martingales. *arXiv preprint arXiv:2009.03167*, September 2020. 1
- Ryu, J. J. and Bhatt, A. On confidence sequences for bounded random processes via universal gambling strategies. *arXiv preprint arXiv:2207.12382*, 2022. 1, 2, 3, 8, 9
- Shekhar, S. and Ramdas, A. On the near-optimality of betting confidence sets for bounded means. *arXiv preprint arXiv:2310.01547*, 2023. 1, 3
- Sison, C. P. and Glaz, J. Simultaneous confidence intervals and sample size determination for multinomial proportions. *J. Am. Statist. Assoc.*, 90(429):366–369, 1995. 6
- Tsai, C.-E., Lin, Y.-T., and Li, Y.-H. Data-dependent bounds for online portfolio selection without Lipschitzness and smoothness. In *Adv. Neural Inf. Proc. Syst.*, volume 36, 2024. 9
- Vakili, S., Scarlett, J., and Javidi, T. Open problem: Tight online confidence intervals for RKHS elements. In *Conf. Learn. Theory*, pp. 4647–4652. PMLR, 2021. 9
- van Erven, T., Van der Hoeven, D., Kotłowski, W., and Koolen, W. M. Open problem: Fast and optimal online portfolio selection. In *Conf. Learn. Theory*, pp. 3864–3869. PMLR, 2020. 9
- Ville, J. Etude critique de la notion de collectif. *Bull. Amer. Math. Soc.*, 45(11):824, 1939. 2
- Waudby-Smith, I. and Ramdas, A. Confidence sequences for sampling without replacement. In Larochele, H., Ranzato, M., Hadsell, R., Balcan, M., and Lin, H. (eds.), *Adv. Neural Inf. Proc. Syst.*, volume 33, pp. 20204–20214. Curran Associates, Inc., 2020. 1, 4, 7, 13, 14
- Waudby-Smith, I. and Ramdas, A. Estimating means of bounded random variables by betting. *J. R. Stat. Soc. B.*, 86(1):1–27, 2024. 1, 2, 3, 7
- Waudby-Smith, I., Stark, P., and Ramdas, A. RiLACS: Risk Limiting Audits via Confidence Sequences. *E-Vote-ID 2021*, pp. 130, 2021. 1, 7
- Whitehouse, J., Wu, Z. S., and Ramdas, A. Time-uniform self-normalized concentration for vector-valued processes. *arXiv preprint arXiv:2310.09100*, 2023. 1
- Xie, Q. and Barron, A. R. Asymptotic minimax regret for data compression, gambling, and prediction. *IEEE Trans. Inf. Theory*, 46(2):431–445, 2000. 5
- Yang, R. and Berger, J. O. *A catalog of noninformative priors*, volume 1018. Institute of Statistics and Decision Sciences, Duke University Durham, NC, USA, 1996. 5
- Zimmert, J., Agarwal, N., and Kale, S. Pushing the efficiency-regret Pareto frontier for online learning of portfolios and quantum states. *arXiv preprint arXiv:2202.02765*, 2022. 9

A. Deferred Proofs

A.1. Proof of Theorem 3

Proof. The convexity of the KT CS in (b) follows from the observation that $\mathbf{m} \mapsto \boldsymbol{\alpha}^\top \log \mathbf{m}^{-1}$ is convex for any $\boldsymbol{\alpha} \in \mathbb{R}_{\geq 0}^K$, and the sublevel set of a convex function is always convex. To prove (c), observe that

$$\begin{aligned} \mathcal{C}^{\text{KT}}(\mathbf{z}^t; \delta) &\stackrel{(i)}{\subset} \left\{ \mathbf{m} \in \Delta^{K-1} : \widehat{\boldsymbol{\mu}}_t^\top \mathbf{m}^{-1} < (\delta q^{\text{KT}}(\mathbf{k}(\mathbf{z}^t)))^{-\frac{1}{t}} \right\} \\ &\stackrel{(ii)}{\subset} \left\{ \mathbf{m} \in \Delta^{K-1} : m_j > \frac{k_j(\mathbf{z}^t)}{t} (\delta q^{\text{KT}}(\mathbf{k}(\mathbf{z}^t)))^{\frac{1}{t}} \right\}, \end{aligned}$$

where (i) follows from Jensen's inequality and (ii) holds by deploying the lower bound $\widehat{\boldsymbol{\mu}}_t^\top \mathbf{m}^{-1} \geq \frac{k_j(\mathbf{z}^t)}{t} \frac{1}{m_j}$ for any $j \in [K]$. The last property (d) follows from noting that by Stirling's formula, the second term in the definition of the CS behaves as

$$\frac{1}{t} \log \frac{e^{-tH(\widehat{\boldsymbol{\mu}}_t)}}{q^{\text{KT}}(t\widehat{\boldsymbol{\mu}}_t)} \sim \frac{K-1}{2t} \log t + o(1)$$

for t sufficiently large. \square

A.2. Proof of Theorem 5

The proof relies on a simple combinatorial inequality. It suffices to show that

$$\binom{K}{\mathbf{k}} \geq \binom{K+M}{\mathbf{k}+\mathbf{m}} \left(\frac{\mathbf{k}}{K}\right)^{\mathbf{m}}$$

for $K, M \in \mathbb{N}$, $\mathbf{k}, \mathbf{m} \in \mathbb{N}_+^m$ such that $\sum_{i=1}^m k_i = K$ and $\sum_{i=1}^m m_i = M$, since the desired inequality follows by plugging in $K \leftarrow N-t$, $\mathbf{k} \leftarrow \mathbf{n} - \mathbf{K}_t$, $M \leftarrow t$, $\mathbf{m} \leftarrow \mathbf{K}_t$. The desired inequality follows from Lemma 8, which is stated and proved below. \square

Lemma 8. *Let $k_1, \dots, k_d \geq 1$ and $m_1, \dots, m_d \geq 0$. Then,*

$$\frac{(\sum_{i=1}^d k_i)! (\sum_{i=1}^d k_i)^{\sum_{i=1}^d m_i}}{(\sum_{i=1}^d (k_i + m_i))!} \geq \prod_{i=1}^d \frac{k_i! k_i^{m_i}}{(k_i + m_i)!}.$$

Proof. It suffices to show for $d = 2$, i.e.,

$$\binom{k_1 + k_2}{k_1} \geq \binom{k_1 + k_2 + m_1 + m_2}{k_1 + m_1} \frac{k_1^{m_1} k_2^{m_2}}{(k_1 + k_2)^{m_1 + m_2}}. \quad (17)$$

Rearranging the terms, we can write

$$\binom{k_1 + k_2}{k_1} \left(\frac{k_1}{k_1 + k_2}\right)^{k_1} \left(\frac{k_2}{k_1 + k_2}\right)^{k_2} \geq \binom{k_1 + k_2 + m_1 + m_2}{k_1 + m_1} \left(\frac{k_1}{k_1 + k_2}\right)^{k_1 + m_1} \left(\frac{k_2}{k_1 + k_2}\right)^{k_2 + m_2}.$$

If we consider two independent binomial random variables $Z \sim \text{Binom}(k_1 + k_2, \frac{k_1}{k_1 + k_2})$ and $W \sim \text{Binom}(m_1 + m_2, \frac{k_1}{k_1 + k_2})$, then the inequality is equivalent to

$$\mathbb{P}(Z = k_1) \geq \mathbb{P}(Z + W = k_1 + m_1),$$

since $Z + W \sim \text{Binom}(k_1 + k_2 + m_1 + m_2, \frac{k_1}{k_1 + k_2})$. Finally, the inequality readily follows from the following chain of inequalities:

$$\begin{aligned} \mathbb{P}(Z + W = k_1 + m_1) &= \sum_{0 \leq z \leq \min\{k_1 + m_1, k_1 + k_2\}} \mathbb{P}(Z = z) \mathbb{P}(W = k_1 + m_1 - z) \\ &\leq \mathbb{P}(Z = k_1) \sum_{0 \leq z \leq \min\{k_1 + m_1, k_1 + k_2\}} \mathbb{P}(W = k_1 + m_1 - z) \\ &\leq \mathbb{P}(Z = k_1). \end{aligned}$$

Here, we used the fact that $\max_{0 \leq z \leq k_1 + k_2} \mathbb{P}(Z = z) = \mathbb{P}(Z = k_1)$ in the first inequality. This completes the proof. \square

A.3. Proof of Theorem 7

Proof. The convexity of UP CS follows from the fact that $\mathbf{m} \mapsto \mathbf{m}^{-\mathbf{k}}$ is log-convex for each $\mathbf{k} \in \mathcal{G}_{K,t}$, a sum of any log-convex functions is also log-convex. The proof for (b) is similar to that of Theorem 3(b) and thus omitted. We can prove (c) via the Lemma 9 stated below, which states that no constant bettor can earn any money from $\text{CHR}_K(\widehat{\boldsymbol{\mu}}_t^{-1})$ for any outcome sequence \mathbf{y}^t , *deterministically*. \square

Lemma 9. For $\mathbf{y}^t \in (\Delta^{K-1})^t$, let $\widehat{\boldsymbol{\mu}}_t = \frac{1}{t} \sum_{i=1}^t \mathbf{y}_i$. For any $\mathbf{b} \in \Delta^{K-1}$, we have

$$W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1}) \leq 1.$$

Proof. Note that, from (12), we can write

$$W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1}) = \sum_{\mathbf{k} \in \mathcal{G}_t} \frac{\mathbf{b}^{\mathbf{k}}}{\widehat{\boldsymbol{\mu}}_t^{\mathbf{k}}} \mathbf{y}^t[\mathbf{k}].$$

Since it is easy to verify that $W_t^{\widehat{\boldsymbol{\mu}}_t}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1}) = 1$, we only need to show that the global maximum of the function $\mathbf{b} \mapsto W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1})$ is attained when $\mathbf{b} = \widehat{\boldsymbol{\mu}}_t$. Note that the function $\mathbf{b} \mapsto W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1})$ is log-concave, since $\mathbf{b} \mapsto \mathbf{b}^{\mathbf{k}}$ is log-concave and a sum of any log-concave functions is also log-concave. Moreover, it is easy to check that the gradient of the function with respect to \mathbf{b} at $\mathbf{b} = \widehat{\boldsymbol{\mu}}_t$ is perpendicular to the probability simplex Δ_m . That is, we have

$$\frac{\partial}{\partial b_j} W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1})|_{\mathbf{b}=\widehat{\boldsymbol{\mu}}_t} = \sum_{\mathbf{k} \in \mathcal{G}_t} \mathbf{y}^t[\mathbf{k}] \left(\frac{k_j}{\widehat{\mu}_{tj}} - \sum_{j' \in [m] \setminus \{j\}} \frac{k_{j'}}{\widehat{\mu}_{tj'}} \right) = -(m-2)t,$$

where the last equality holds since $\sum_{\mathbf{k} \in \mathcal{G}_t} k_j \mathbf{y}^t[\mathbf{k}] = \sum_{i=1}^t y_{ij} = t \widehat{\mu}_{tj}$. Hence, the gradient

$$\nabla_{\mathbf{b}} W_t^{\mathbf{b}}(\mathbf{y}^t; \widehat{\boldsymbol{\mu}}_t^{-1})|_{\mathbf{b}=\widehat{\boldsymbol{\mu}}_t} = -(m-2)t \mathbf{1}_K$$

is perpendicular to the probability simplex Δ^{K-1} , and thus together with the log-concavity, we can conclude that $\mathbf{b} = \widehat{\boldsymbol{\mu}}_t$ is the global maximizer. \square

B. General Reduction for CS with WoR Sampling

Proposition 10 (Waudby-Smith & Ramdas (2020)). Suppose that a sequence $\{W_t(\widetilde{\mathbf{Z}}^t; \mathbf{m})\}_{t=0}^{\infty}$ with $W_0(\varphi; \mathbf{m}) = 1$ is a martingale for $\mathbf{m} = \widetilde{\boldsymbol{\mu}}$, if a stochastic process $(\widetilde{\mathbf{Z}})_{i=1}^{\infty}$ satisfies $E[\widetilde{\mathbf{Z}}_t | \mathcal{F}_{t-1}] = \widetilde{\boldsymbol{\mu}}$. If we define

$$\begin{aligned} \mathcal{C}_{\text{WoR}}(\mathbf{Z}^t; \delta) &:= \left\{ \mathbf{m} : W_t(\mathbf{Z}^t; \mathbf{m}_{N,t+1}^{\text{WoR}}(\mathbf{Z}^t)) < \frac{1}{\delta} \right\} \\ &= \left\{ \mathbf{m} : W_t\left(\mathbf{Z}^t; \frac{N\mathbf{m} - \sum_{i=1}^t \mathbf{Z}_i}{N-t}\right) < \frac{1}{\delta} \right\} \end{aligned}$$

for $t \in [N-1]$, then

$$P(\forall t \in [N-1], \boldsymbol{\mu} \in \mathcal{C}_{\text{WoR}}(\mathbf{Z}^t; \delta)) \geq 1 - \delta,$$

given that $(\mathbf{Z}_t)_{t=1}^N$ is sampled WoR from $\{\mathbf{z}_i\}_{i=1}^N$.

C. Posterior-Prior Ratio Interpretation of PPR CS

As alluded to earlier, the PPR sequence (11) can be derived as a posterior-prior ratio as follows. Note that we can write

$$R(\mathbf{Y}^t; \mathbf{m}) = \frac{\pi_0(N\mathbf{m})}{\pi_t(N\mathbf{m})},$$

if we define

$$\pi_t(\mathbf{n}) := \frac{B(\mathbf{n} + \boldsymbol{\alpha})}{B(\mathbf{K}_t + \boldsymbol{\alpha})} \binom{N-t}{\mathbf{n} - \mathbf{K}_t}$$

for $t \geq 0$ and $\mathbf{n} \in \mathcal{G}_{N,K}$. This quantity $\pi_t(\mathbf{n})$ is the Dirichlet multinomial distribution, which is the posterior distribution induced by the hypergeometric observation model, while $\pi_0(\mathbf{n})$ corresponds to the prior distribution. More precisely, if we assume $\mathbf{Y}_t | (\mathbf{Y}_1, \dots, \mathbf{Y}_{t-1}) \sim \text{MultiHyperGeo}(N - (t - 1), \mathbf{n} - \mathbf{s}_{t-1}, 1)$ and $\mathbf{n} \sim \text{DirMult}(N, \boldsymbol{\alpha})$. Then, by the conjugacy, one can show that $\mathbf{n} - \mathbf{K}_t | (\mathbf{Y}_1, \dots, \mathbf{Y}_t) \sim \text{DirMult}(N - t, \mathbf{K}_t + \boldsymbol{\alpha})$, which leads to the definition of $\pi_t(\mathbf{n})$; see (Waudby-Smith & Ramdas, 2020, Theorem C.1). It can be shown that any prior-to-posterior ratio sequence is a martingale if $\mathbf{m} = \boldsymbol{\mu}$, and thus can be used to construct a CS.