Sequential Kernel Goodness-of-fit Testing

Zhengyu Zhou [1] **Weiwei Liu** [1]

Abstract

Goodness-of-fit testing, a classical statistical tool, has been extensively explored in the batch setting, where the sample size is predetermined. However, practitioners often prefer methods that adapt to the complexity of a problem rather than fixing the sample size beforehand. Classical batch tests are generally unsuitable for streaming data, as valid inference after data peeking requires multiple testing corrections, resulting in reduced statistical power. To address this issue, we delve into the design of consistent sequential goodness-of-fit tests. Following the principle of *testing by betting*, we reframe this task as selecting a sequence of payoff functions that maximize the wealth of a fictitious bettor, betting against the null in a repeated game. We conduct experiments to demonstrate the adaptability of our sequential test across varying difficulty levels of problems while maintaining control over type-I errors.

1. Introduction

Goodness-of-fit tests are fundamental tools in statistical analysis, dating back to the Kolmogorov–Smirnov test (Kolmogorov, 1933; Smirnov, 1948). Given observations Z sampled from the distribution q, we aim to test the null hypothesis that q matches the reference or target distribution p. Traditional goodness-of-fit measures, such as the Kolmogorov–Smirnov statistic (Kolmogorov, 1933; Smirnov, 1948) and Cramér–von Mises criterion, are primarily applicable to univariate random variables. Gorham and Mackey propose the Stein discrepancy (Gorham & Mackey, 2015), a measure of sample quality relative to a target. This measure is a maximum discrepancy between empirical sample expectations and target expectations over a large class of test functions. It is constructed to have zero expectation over the target distribution through a Stein operator, which solely depends on the derivative of the $\log p$. Kernel Stein Discrepancy (KSD) has gained substantial attention in non-parametric goodness-of-fit testing (Liu et al., 2016; Chwialkowski et al., 2016; Jitkrittum et al., 2017; Baum et al., 2023; Liu et al., 2023).

In the existing literature, significant attention has been devoted to *batch* testing, primarily when dealing with predetermined sample sizes. When the sample distribution deviates from the target, the necessary sample size for detecting such discrepancies is not known a priori. In cases where test results show promise but remain inconclusive, such as when a p-value slightly exceeds a chosen significance level, gathering additional data and repeating the study becomes necessary. Traditional batch tests are not designed to deal with these challenges. Therefore, this paper focuses on sequential tests that facilitate the scrutiny of observed data, enabling the decision to stop and reject the null hypothesis or continue with data collection.

Problem Setup. Suppose there is a continuous stream of data denoted as $\{Z_t\}_{t\geq 1} \subset \mathcal{Z}$, where $Z_t \overset{\text{iid}}{\sim} q$. Our objective is to investigate whether the distribution q aligns with a known reference or target distribution p. According to our formulation, the target distribution p is assumed to be known only up to a normalization constant. We also aim to design sequential tests for the following hypotheses:

$$H_0 : Z_t \overset{\text{iid}}{\sim} q,\ t \geq 1 \text{ and } q = p, \tag{1a}$$

$$H_1 : Z_t \overset{\text{iid}}{\sim} q,\ t \geq 1 \text{ and } q \neq p. \tag{1b}$$

Following the "test of power one" framework (Darling & Robbins, 1968), we define a level-α sequential test as a mapping $\Phi : \cup_{t=1}^{\infty} \mathcal{Z}^t \to \{0, 1\}$ that satisfies the formula:

$$\mathbb{P}_{H_0}(\exists\, t \geq 1 : \Phi(Z_1, \ldots, Z_t) = 1) \leq \alpha. \tag{2}$$

The output 0 stands for "do not reject the null yet and continue sampling," while 1 means "reject the null and stop." Additionally, defining the stopping time $\tau := \inf\{t \geq 1 : \Phi(Z_1, \ldots, Z_t) = 1\}$ as the first time the test outputs 1, a level-α sequential test must satisfy the formula:

$$\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha. \tag{3}$$

[1]School of Computer Science, National Engineering Research Center for Multimedia Software, Institute of Artificial Intelligence and Hubei Key Laboratory of Multimedia and Network Communication Engineering, Wuhan University, Wuhan, China. Correspondence to: Weiwei Liu <liuweiwei863@gmail.com>.

Proceedings of the 41st International Conference on Machine Learning, Vienna, Austria. PMLR 235, 2024.

We highlight the primary limitations of existing tests that our new method addresses.

Limitations of Corrected Batch Tests. Using batch tests without corrections for multiple testing results in an inflated false alarm rate under continuous monitoring (see Appendix A.1). Hence, naïve Bonferroni corrections restore type-I error control but generally result in low-power tests. As a result, directly designing sequential tests, without batch test correction, is necessary. Therefore, we perform valid sequential goodness-of-fit tests for target density function $p(x) \propto 1/(1+x^2)$ and $Z_t = \beta X_t + (1-\beta)Y_t$, where $X_t \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0,1)$ and $Y_t \overset{\text{i.i.d.}}{\sim} p$ are independent. We consider 11 β: $\beta \in \{0.5, 0.55, \dots, 1.0\}$ values, and for each β we repeat the simulation 100 times. We use Gaussian kernel $k(x,y) = \exp(-|x-y|^2/2)$ for all testing procedures. In this simulation, we compare two goodness-of-fit testing approaches:

1. **KSD-based SKGT** proposed in this work (Algorithm 2).
2. **Batch KSD Test** adapted for continuous monitoring via Bonferroni correction. We enable monitoring after processing every n, $n \in \{10, 100\}$, new points from q. In other words, the bootstrapped p-value (computed over 1000 wild bootstrapped samples) is compared with the rejection thresholds: $\alpha_i = \alpha/(i(i+1))$ and $i = 1, 2, \dots$, where i stands for the monitoring order.

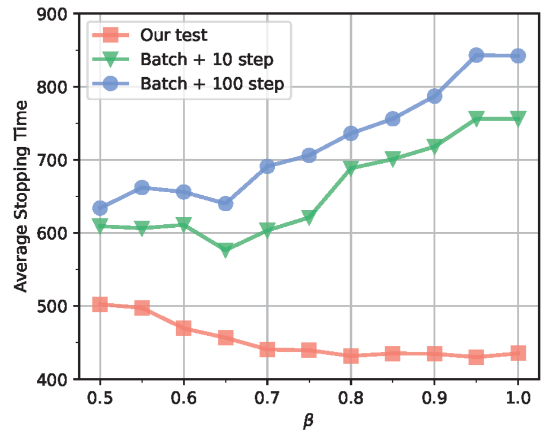

Figure 1. Average stopping time of different tests for continuous monitoring. Batch + n-step represents the batch KSD tests with Bonferroni correction applied every n steps.

As illustrated in Figure 1, our tests exhibit a lower average sample requirement compared to other methods. Additionally, as β approaches 1, indicating a less challenging task, the number of samples needed to reject the null hypothesis decreases. This property is not observed in other tests in this example.

Sequential tests serve as complementary tools to batch tests and are not designed to replace them. We consider two scenarios to highlight this point. When there are $2,000$ data points, recourse is limited if batch tests fail to reject the null hypothesis. However, if sequential tests fail to reject, analysts can collect more data and continue testing, retaining type-I error control. Conversely, when dealing with 2 million data points, the KSD may be time-consuming, due to the bootstrap procedure. Therefore, if the alternative hypothesis is true, and the signal is strong, sequential tests may reject the hypothesis within 200 samples and terminate. In essence, the capacity of sequential tests to continuously collect and analyze data proves advantageous, particularly in challenging situations.

1.1. Related Work

The principle of testing by betting can be traced back to Ville's 1939 doctoral thesis (Ville, 1939), which has recently been popularized by Shafer (2021). Shafer (2021) primarily focuses on parametric and simple hypotheses, which is distinct from our setting. The studies most closely related to the current paper include (Shekhar & Ramdas, 2023; Shaer et al., 2023; Podkopaev et al., 2023; Grünwald et al., 2023; Podkopaev & Ramdas, 2023), which also address non-parametric hypotheses. Notably, Shekhar and Ramdas utilize testing by betting to design sequential non-parametric two-sample tests (Shekhar & Ramdas, 2023), incorporating a state-of-the-art sequential kernel maximum mean discrepancy test. To the best of our knowledge, this paper is the first attempt to apply the principle of testing by betting to goodness-of-fit testing.

2. Sequential Kernel Goodness-of-fit Testing

This section summarizes the principle of testing by betting (Shafer, 2021; Shafer & Vovk, 2019). Given that a sequence of random variables $\{Z_t\}_{t\geq 1}$, where $Z_t \in \mathcal{Z}$, a player begins with initial wealth $\mathcal{K}_0 = 1$. At round t of the game, the player can select a *payoff* function $f_t : \mathcal{Z} \to [-1,1]$ that satisfies $\mathbb{E}_{Z\sim P_Z}[f_t(Z)|\mathcal{F}_{t-1}] = 0$ for all $P_Z \in H_0$, where $\mathcal{F}_{t-1} = \sigma(Z_1, \dots, Z_{t-1})$ is the σ-algebra generated by $\{Z_s : 1 \leq s \leq t-1\}$. Additionally, the player bets a fraction of the wealth $\lambda_t \mathcal{K}_{t-1}$ for an $\mathcal{F}_{t-1}$-measurable $\lambda_t \in [-1,1]$. Once Z_t is revealed, the player's *wealth* is updated as:

$$\mathcal{K}_t = \mathcal{K}_{t-1}(1 + \lambda_t f_t(Z_t)). \tag{4}$$

After that, a level-α sequential test is obtained using the following stopping rule: $\Phi(Z_1, \dots, Z_t) = \mathbb{1}\{\mathcal{K}_t \geq 1/\alpha\}$, i.e., the null hypothesis is rejected once the player's wealth exceeds $1/\alpha$. Under the null hypothesis, the imposed constraints on the payoff sequences $\{f_t\}_{t\geq 1}$ and betting fraction $\{\lambda_t\}_{t\geq 1}$ prevent the player from making a profit. Formally, the wealth process $\{\mathcal{K}_t\}_{t\geq 0}$ is a non-negative martingale. The validity of the resulting test is established through Ville's inequality (Ville, 1939).

To ensure that the resulting test has power under the alternative hypothesis, payoffs and betting fractions have to be selected carefully. Inspired by sequential two-sample tests proposed by Shekhar & Ramdas (2023), our approach relies on the KSD (Chwialkowski et al., 2016; Liu et al., 2016), which has a variational representation.

Kernel Stein Discrepancy. Let $\mathcal{G}$ be the reproducing kernel Hilbert space (RKHS) of real-valued functions on $\mathbb{R}^d$ with the reproducing kernel $k(\cdot,\cdot)$ and inner product $\langle\cdot,\cdot\rangle_{\mathcal{G}}$. Similarly, let $\mathcal{G}^d$ denote the product RKHS consisting of elements $f := (f_1,\dots,f_d)$ and $f_i \in \mathcal{G}$, with the standard inner product $\langle f, g\rangle_{\mathcal{G}^d} = \sum_{i=1}^d \langle f_i, g_i\rangle_{\mathcal{G}}$. Therefore, we define a Stein operator T_p acting on $f \in \mathcal{G}^d$ as:

$$(T_p f)(x) := \sum_{i=1}^d \left(\frac{\partial \log p(x)}{\partial x_i} f_i(x) + \frac{\partial f_i(x)}{\partial x_i} \right).$$

We observe that the operator can be expressed by defining a function that depends on the kernel and log-density's gradient, as shown below:

$$\xi_p(x,\cdot) := [\nabla \log p(x) k(x,\cdot) + \nabla k(x,\cdot)], \tag{5}$$

where the gradient is taken with respect to variable x. Thus, the expected inner product of $\xi_p(x,\cdot)$ with f corresponds to the expected value of the Stein operator,

$$\begin{aligned}\mathbb{E}_{Z\sim q} T_p f(Z) &= \langle f, \mathbb{E}_{Z\sim q}\xi_p(Z,\cdot)\rangle_{\mathcal{G}^d} \\ &= \sum_{i=1}^d \langle f_i, \mathbb{E}_{Z\sim q}\xi_{p,i}(Z,\cdot)\rangle_{\mathcal{G}},\end{aligned}$$

where $\xi_{p,i}(x,\cdot)$ is the i-th component of $\xi_p(x,\cdot)$. Furthermore, for samples from the target distribution, $\mathbb{E}_{X\sim p}(T_p f)(X) = 0$, which can be verified using integration by parts. We define the KSD as:

$$\begin{aligned} S_p(q) &:= \sup_{\|f\|_{\mathcal{G}^d}\leq 1} \mathbb{E}_{Z\sim q} T_p f(Z) - \mathbb{E}_{X\sim p}(T_p f)(X) \\ &= \sup_{\|f\|_{\mathcal{G}^d}\leq 1} \mathbb{E}_{Z\sim q} T_p f(Z) \\ &= \sup_{\|f\|_{\mathcal{G}^d}\leq 1} \langle f, \mathbb{E}_{Z\sim q}\xi_p(Z,\cdot)\rangle_{\mathcal{G}^d} \\ &= \|\mathbb{E}_q \xi_p(Z,\cdot)\|_{\mathcal{G}^d} \end{aligned} \tag{6}$$

To develop powerful goodness-of-fit tests, it is imperative that the function class $\mathcal{G}^d$ is sufficiently expressive to guarantee $S_p(q) > 0$ when $p \neq q$. The following theorem established by Chwialkowski et al. (2016) affirms that the KSD can be used to distinguish between two distributions. Before presenting the theorem, it is necessary to define the following:

$$\begin{aligned} h_p(x,y) := &\nabla \log p(x)^\top \nabla \log p(y) k(x,y) \\ &+ \nabla \log p(y)^\top \nabla_x k(x,y) \\ &+ \nabla \log p(x)^\top \nabla_y k(x,y) \\ &+ \langle \nabla_x k(x,\cdot), \nabla_y k(\cdot,y)\rangle_{\mathcal{G}^d}, \end{aligned} \tag{7}$$

where the last term can be written as $\sum_{i=1}^d \frac{\partial^2 k(x,y)}{\partial x_i \partial y_i}$.

Theorem 2.1. *(Chwialkowski et al., 2016, Theorem 2.2) Let p, q be probability measures that $Z \sim q$. If the kernel k is C_0-universal (Carmeli et al., 2010, Definition 4.1), $\mathbb{E}_{Z\sim q} h_p(Z,Z) < \infty$, and $\mathbb{E}_{Z\sim q}\left\|\nabla\left(\log\frac{p(Z)}{q(Z)}\right)\right\|^2 < \infty$, then $S_p(q) = 0$ if and only if $q = p$.*

Examples of C_0-universal kernels on $\mathbb{R}^d$ include the Gaussian, Laplacian, inverse multiquadrics, and Matérn class, among others.

KSD-based Sequential Kernel Goodness-of-fit Testing. An element $g^* \in \mathcal{G}^d$ that achieves the supremum in (6), often referred to as the "witness function", can be regarded as the test function in $\mathcal{G}^d$ that best distinguishes q from p. Thus, we consider the payoffs $f(Z_t)$ of the following form:

$$s \cdot \langle g, \xi_p(Z_t,\cdot)\rangle_{\mathcal{G}^d} = s \cdot (\langle g(Z_t), \nabla \log p(Z_t)\rangle + \nabla \cdot g(Z_t)), \tag{8}$$

where $\nabla \cdot g = \sum_{i=1}^d \frac{\partial g}{\partial x_i}$ denotes the divergence of g, and the scaling factor $s > 0$ ensures that $f(z) \in [-1,1]$, for any $z \in \mathbb{R}^d$. Substituting g in (8) with the witness function g^*, we denote the resulting function as the "oracle payoff" f^*. Let the oracle wealth process $\{\mathcal{K}_t^*\}_{t\geq 0}$ be defined using f^* and the betting fraction, as shown below:

$$\lambda^* = \frac{\mathbb{E}_{Z\sim q}[f^*(Z)]}{\mathbb{E}_{Z\sim q}[f^*(Z)] + \mathbb{E}_{Z\sim q}[(f^*(Z))^2]}. \tag{9}$$

We have the following result regarding the oracle payoff function and betting fraction in (9).

Theorem 2.2. *Let $\mathcal{G}$ be the RKHS constructed from a C_0-universal kernel and $\mathcal{G}^d$ be the corresponding product RKHS:*

1. *Under H_0 in (1a), any payoff function of the form (8) satisfies $\mathbb{E}_{H_0}[f(Z)] = 0$.*
2. *Under H_1 in (1b), the oracle payoff function f^* based on the witness function g^* satisfies $\mathbb{E}_{H_1}[f^*(Z)] > 0$. Further, for λ^* defined in (9), it holds that $\mathbb{E}_{H_1}[\log(1+\lambda^* f^*(Z))] > 0$. Hence, $\mathcal{K}_t^* \overset{\text{a.s.}}{\to} +\infty$, which implies that the oracle test is consistent: $\mathbb{P}_{H_1}(\tau^* < \infty) = 1$, where $\tau^* = \inf\{t \geq 1 : \mathcal{K}_t^* \geq 1/\alpha\}$.*

Remark 2.1. While the betting fraction (9) suffices to guarantee the consistency of the corresponding test, the fastest growth rate of the wealth process is ensured by considering

$$\lambda^*_{\text{best}} \in \arg\max_{\lambda \in [-1,1]} \mathbb{E}_{Z\sim q}[\log(1+\lambda f^*(Z))].$$

However, casually selecting the betting fraction may result in the wealth tending to zero almost surely, as exemplified in (Podkopaev et al., 2023, Example 2).

Therefore, to construct a practical test, we must replace the oracle f^* and λ^* with predictable estimates $\{f_t\}_{t\geq 1}$ and $\{\lambda_t\}$. This indicates that they are computed using data observed prior to a given round of the game.

Assumption 2.3. Suppose the kernel k is C_0-universal, non-negative and satisfies: $\sup_{x\in\mathbb{R}^d} k(x,x) \leq 1$ and $\sup_{x\in\mathbb{R}^d} \sum_{i=1}^d \frac{\partial^2 k(x,y)}{\partial x_i \partial y_i}\Big|_{y=x} \leq 1$.

Assumption 2.4. Suppose that the target distribution p satisfies: $\sup_{z\in\mathbb{R}^d} \|\nabla \log p(z)\|_2 \leq 1$, where $\|v\|_2 := (\sum_{i=1}^d v_i^2)^{1/2}$ denotes the ℓ_2-norm.

The C_0-universality of kernel k in Assumption 2.3 ensures that $S_p(q) > 0$, when $p \neq q$. The boundedness requirement in Assumption 2.3 is fulfilled by several commonly used kernels, including the Gaussian and inverse multiquadrics with appropriate scaling. Assumption 2.3 together with Assumption 2.4 guarantees the boundedness of the payoff function.

Payoff Function f_t. Considering the KSD's variational formulation, the witness function adopts a closed form:

$$g^* = \frac{\mathbb{E}_{Z\sim q}\xi_p(Z,\cdot)}{\|\mathbb{E}_{Z\sim q}\xi_p(Z,\cdot)\|_{\mathcal{G}^d}}. \tag{10}$$

The oracle payoff $f^*(Z_t)$ based on KSD is given by:

$$\frac{1}{2}\langle g^*, \xi_p(Z_t,\cdot)\rangle_{\mathcal{G}^d} = \frac{1}{2}(\langle g^*(Z_t), \nabla \log p(Z_t)\rangle + \nabla\cdot g^*(Z_t)), \tag{11}$$

with the form (8) and $s = 1/2$. Therefore, to construct the test, we use estimators $\{f_t\}_{t\geq 1}$ of the oracle payoff function f^* obtained by replacing g^* in (11) with the plug-in estimator:

$$\hat{g}_t = \frac{\mathbb{E}_{Z\sim \hat{q}_{t-1}}\xi_p(Z,\cdot)}{\|\mathbb{E}_{Z\sim \hat{q}_{t-1}}\xi_p(Z,\cdot)\|_{\mathcal{G}^d}}, \tag{12}$$

where $\hat{q}_{t-1}$ is the empirical distribution of $\{Z_s\}_{s=1}^{t-1}$ and

$$\mathbb{E}_{Z\sim \hat{q}_{t-1}}\xi_p(Z,\cdot) = \frac{1}{t-1}\left(\sum_{i=1}^{t-1} \nabla \log p(Z_i) k(Z_i,\cdot) + \nabla k(Z_i,\cdot)\right). \tag{13}$$

Algorithm 1 Online Newton Step (ONS) strategy for selecting betting fractions

Input: sequence of payoffs $\{f_t(Z_t)\}_{t\geq 1}$, $\lambda_1^{\text{ONS}} = 0$, $a_0 = 1$.
for $t = 1, 2, \ldots$ **do**
  Observe $f_t(Z_t)$;
  Set $z_t = f_t(Z_t)/(1 - \lambda_t^{\text{ONS}} f_t(Z_t))$;
  Set $a_t = a_{t-1} + z_t^2$;
  Set $\lambda_{t+1}^{\text{ONS}} := \frac{1}{2} \wedge \left(-\frac{1}{2} \vee \left(\frac{2}{2-\log 3}\cdot\frac{z_t}{a_t} - \lambda_t^{\text{ONS}}\right)\right)$;
end for

Algorithm 2 KSD-based SKGT

Input: significance level $\alpha \in (0,1)$, data stream $\{Z_t\}_{t\geq 1}$, where $Z_t \sim q$, $\lambda_1^{\text{ONS}} = 0$.
for t=1,2,... **do**
  Use $Z_1, \ldots, Z_{t-1}$ to compute $\hat{g}_t$ as in (12);
  Compute KSD payoff $f_t(Z_t)$;
  Update the wealth process $\mathcal{K}_t$ as in (4);
  if $\mathcal{K}_t \geq 1/\alpha$ **then**
    Reject H_0 and stop;
  else
    compute $\lambda_{t+1}^{\text{ONS}}$ (Algorithm 1);
  end if
end for

Notably, in (12), the witness function's plug-in estimate is defined as an operator.

Betting Fraction λ_t. To select betting fractions in an online fashion, we employ the approach proposed by Cutkosky & Orabona (2018). They expressed the problem of choosing the optimal betting fraction for coin betting as an online optimization problem with exp-concave losses and proposed a strategy based on the Online Newton Step (ONS) (Hazan et al., 2007), which is summarized in Algorithm 1.

We conclude this section with formal guarantees regarding the time-uniform type-I error control and consistency of KSD-based SKGT. In addition, we show that the wealth process grows exponentially and characterize the wealth growth rate in terms of the true KSD. The proof is detailed in Appendix B.1.

Theorem 2.5. *Under Assumptions 2.3 and 2.4, the following claims hold for KSD-based SKGT (Algorithm 2):*

1. *Under H_0 in (1a), SKGT stops with probability at most α: $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.*

2. *Under H_1 in (1b), then it holds that $\mathcal{K} \xrightarrow{\text{a.s.}} +\infty$ and thus the SKGT is consistent: $\mathbb{P}_{H_1}(\tau < \infty) = 1$. Furthermore, the wealth grows exponentially, and the cor-*

responding growth rate satisfies the following formula:

$$\liminf_{t\to\infty} \frac{\log \mathcal{K}_t}{t} \geq \frac{\mathbb{E}_{H_1}[f^*(Z)]}{4}\left(\frac{\mathbb{E}_{H_1}[f^*(Z)]}{\mathbb{E}_{H_1}[(f^*(Z))^2]} \wedge 1\right) \tag{14}$$

almost surely, where f^* is the oracle payoff defined in (11).

Proof sketch. Time-uniform type-I error control directly follows from Ville's inequality. To establish the consistency of KSD-based SKGT, we leverage (Cutkosky & Orabona, 2018, Theorem 1), showing that

$$\frac{\log \mathcal{K}_t}{t} \geq \frac{\log \mathcal{K}(\lambda_0)}{t} - C \cdot \frac{\log t}{t},$$

for any betting fraction λ_0. Hence, after selecting a specific betting fraction and employing fundamental inequalities, we further obtain

$$\frac{\log \mathcal{K}_t}{t} \geq \left(\frac{\frac{1}{t}\sum_{i=1}^t f_i}{4} \vee 0\right) \cdot \left(\frac{\frac{1}{t}\sum_{i=1}^t f_i}{\frac{1}{t}\sum_{i=1}^t f_i^2} \wedge 1\right) - \frac{C\log t}{t}.$$

Then, the asymptotic properties of the right-hand side are investigated in Lemma B.5. The proof's core lies in the convergence of V-statistics, analyzed using the non-asymptotic convergence based on Hoeffding's study (Hoeffding, 1963). Asymptotic convergence follows from combining the non-asymptotic results with the Borel–Cantelli lemma. □

Since $\mathbb{E}_{H_1}[f^*(Z)] = \frac{1}{2}S_p(q)$ and $f^*(z) \in [-1, 1]$, as established in the proof of Theorem 2.5, Theorem 2.5 implies that:

$$\liminf_{t\to\infty} \frac{\log \mathcal{K}_t}{t} \overset{\text{a.s.}}{\geq} \frac{1}{8}S_p(q).$$

Amongst the betting fractions that are constrained to lie in $[-1/2, 1/2]$, such as the ONS betting strategy, the optimal growth rate is ensured by using:

$$\lambda^* = \operatorname*{arg\,max}_{\lambda\in[-1/2,1/2]} \mathbb{E}[\log(1 + \lambda f^*(Z)]. \tag{15}$$

Consequently, we obtain the following results regarding the oracle test's growth rate:

Proposition 2.6. *The optimal log-wealth $S^* := \mathbb{E}[\log(1 + \lambda^* f^*(Z))]$ that can be achieved by an oracle betting scheme* (15)*, which knows f^* from* (11) *and the underlying distribution, satisfies the formula:*

$$S^* \leq \frac{\mathbb{E}[f^*(Z)]}{2}\left(\frac{8\mathbb{E}[f^*(Z)]}{3\mathbb{E}[(f^*(Z))^2]} \wedge 1\right). \tag{16}$$

Proof. The fact that $S^* \leq \mathbb{E}[f^*(Z)/2]$ trivially follows from $\mathbb{E}[\log(1 + \lambda f^*(Z))] \leq \lambda\mathbb{E}[f^*(Z)] \leq \mathbb{E}[f^*(Z)]/2$. Since for any $x \in [-1/2, 1/2]$, it holds that: $\log(1 + x) \leq x - 3x^2/8$, we know that:

$$S^* \leq \max_{\lambda\in[-1/2,1/2]} \left(\lambda\mathbb{E}[f^*(Z)] - \frac{3}{8}\lambda^2\mathbb{E}[(f^*(Z))^2]\right), \tag{17}$$

and by assuming the maximization problem, we obtains the upper bound:

$$S^* \leq \frac{2\left(\mathbb{E}[f^*(Z)]\right)^2}{3\mathbb{E}[(f^*(Z))^2]} \tag{18}$$

assuming $\frac{\mathbb{E}[f^*(Z)]}{\mathbb{E}[(f^*(Z))^2]} \leq 3/8$. On the other hand, it always holds that: $S^* \leq \mathbb{E}[f^*(Z)/2]$. To obtain the claimed bound, we multiply the RHS of (18) by two, which completes the proof of (16). □

3. Alternative Stein Discrepancies

Practically, data is sometimes encountered in discrete spaces or bounded domains. However, the KSD is specifically designed for smooth density functions on $\mathbb{R}^d$. This section introduces sequential goodness-of-fit tests based on Kernel Discrete Stein Discrepancy (KDSD) (Yang et al., 2018) and Bound-domain Kernel Stein Discrepancy (bd-KSD) (Xu, 2022).

3.1. KDSD-based Sequential Kernel Goodness-of-fit Testing

We recall the KDSD defined on $\mathcal{X} = \{0, \dots, L-1\}^d$ with $L > 1$. In place of derivative, we specify Δ_k as the cyclic forward difference with respect to k-th coordinate as follows:

$$\Delta_k f(\mathbf{x}) = f(x^1, \dots, \tilde{x}^k, \dots, x^d) - f(x^1, \dots, x^k, \dots, x^d),$$

where $\tilde{x}^k = x^k + 1 \bmod L$, with the corresponding vector-valued operator $\Delta = (\Delta_1, \dots, \Delta_d)^\top$. The inverse operator Δ_k^{-1} is given by the backward difference:

$$\Delta_k^{-1} f(\mathbf{x}) = f(x^1, \dots, x^k, \dots, x^d) - f(x^1, \dots, \bar{x}^k, \dots, x^d),$$

where $\bar{x}^k = x^k - 1 \bmod L$, and $\Delta^{-1} = (\Delta_1^{-1}, \dots, \Delta_d^{-1})^\top$. Then, the score is $\mathbf{s}_p(\mathbf{x}) := p(\mathbf{x})^{-1}\Delta p(\mathbf{x})$, where it is assumed that the probability mass function is strictly positive. For a vector-valued function $\mathbf{g} : \mathcal{X} \to \mathbb{R}^d$, the difference Stein operator is then defined as:

$$\begin{aligned}\mathcal{A}_p\mathbf{g}(\mathbf{x}) &:= \operatorname{tr}[\mathbf{g}(\mathbf{x})\mathbf{s}_p^\top(\mathbf{x}) + \Delta^{-1}\mathbf{g}(\mathbf{x})] \\ &= \sum_{i=1}^d \frac{\Delta_i p(\mathbf{x})}{p(\mathbf{x})} g_i(\mathbf{x}) + \Delta_i^{-1} g_i(\mathbf{x}).\end{aligned}$$

It can be shown that $\mathbb{E}_{\mathbf{x}\sim p}[\mathcal{A}_p\mathbf{g}(\mathbf{x})] = 0$ (Yang et al., 2018, Theorem 2). Given an RKHS $\mathcal{H}^d$ of vector-valued functions $\mathbf{g} : \mathcal{X} \to \mathbb{R}^d$, we obtain the KDSD as:

$$\operatorname{KDSD}(q\|p) := \sup_{\mathbf{g}\in\mathcal{H}^d, \|\mathbf{g}\|_{\mathcal{H}^d}\leq 1} \mathbb{E}_{\mathbf{x}\sim q}[\mathcal{A}_p\mathbf{g}(\mathbf{x})]. \tag{19}$$

Algorithm 3 KDSD-based SKGT

Input: significance level $\alpha \in (0,1)$, data stream $\{\mathbf{x}_t\}_{t\geq 1}$, where $\mathbf{x}_t \sim q$, $\lambda_1^{\text{ONS}} = 0$.
for t=1,2,... **do**
  Use $\mathbf{x}_1, \ldots, \mathbf{x}_{t-1}$ to compute $\hat{\mathbf{g}}_t$ as in (23);
  Compute KDSD payoff $f_t(\mathbf{x}_t) = s \langle \hat{\mathbf{g}}_t, \eta_p(\mathbf{x}_t, \cdot) \rangle_{\mathcal{H}^d}$;
  Update the wealth process $\mathcal{K}_t$ as in (4);
  if $\mathcal{K}_t \geq 1/\alpha$ **then**
    Reject H_0 and stop;
  else
    compute $\lambda_{t+1}^{\text{ONS}}$ (Algorithm 1);
  end if
end for

Witness Function for KDSD. Suppose that k is the reproducing kernel of $\mathcal{H}$, we define the kernel embedding as $\eta_p(\mathbf{x}, \cdot) := \mathbf{s}_p(\mathbf{x})k(\mathbf{x}, \cdot) + \Delta_{\mathbf{x}}^{-1} k(\mathbf{x}, \cdot)$, where $\Delta_{\mathbf{x}}^{-1}$ indicates that the operator Δ^{-1} is applied with respect to $\mathbf{x}$. Based on the kernel embedding, we can express KDSD as:

$$\text{KDSD}(q\|p) = \sup_{\mathbf{g}\in\mathcal{H}^d, \|\mathbf{g}\|_{\mathcal{H}^d}\leq 1} \langle \mathbf{g}, \mathbb{E}_{\mathbf{x}\sim q} \eta_p(\mathbf{x}, \cdot) \rangle_{\mathcal{H}^d}. \quad (20)$$

Based on the variational formulation, the witness function of KDSD has the closed form:

$$\mathbf{g}^* = \frac{\mathbb{E}_{\mathbf{x}\sim q} \eta_p(\mathbf{x}, \cdot)}{\|\mathbb{E}_{\mathbf{x}\sim q} \eta_p(\mathbf{x}, \cdot)\|_{\mathcal{H}^d}}. \quad (21)$$

The oracle payoff function f^* based on KDSD is given by the equation:

$$f^*(\mathbf{x}) = s \langle \mathbf{g}^*, \eta_p(\mathbf{x}, \cdot) \rangle_{\mathcal{H}^d}, \quad (22)$$

where the scaling factor $s > 0$ ensures $f^*(\mathbf{x}_t) \in [-1, 1]$. Given a sequence of samples $\{\mathbf{x}_t\}_{t\geq 1}$ from q, we use estimates $\{f_t\}_{t\geq 1}$ of the oracle payoff function obtained by replacing $\mathbf{g}^*$ in (22) with the plug-in estimator:

$$\hat{\mathbf{g}}_t = \frac{\mathbb{E}_{\mathbf{x}\sim \hat{q}_{t-1}} \eta_p(\mathbf{x}, \cdot)}{\left\|\mathbb{E}_{\mathbf{x}\sim \hat{q}_{t-1}} \eta_p(\mathbf{x}, \cdot)\right\|_{\mathcal{H}^d}}, \quad (23)$$

where $\hat{q}_{t-1}$ is the empirical distribution of $\{\mathbf{x}_1, \ldots, \mathbf{x}_{t-1}\}$.

Assumption 3.1. Suppose that a constant B_p for the target distribution p exists, such that $\sup_{\mathbf{x}\in\mathcal{X}} \|\Delta p(\mathbf{x})/p(\mathbf{x})\|_2 \leq B_p$.

Assumption 3.2. Suppose that the kernel k satisfies: $\sup_{\mathbf{x}\in\mathcal{X}} k(\mathbf{x}, \mathbf{x}) \leq B_{k,0}$, $\sup_{\mathbf{x}\in\mathcal{X}} \left\|\Delta_{\mathbf{x}}^{-1} k(\mathbf{x}, \mathbf{x})\right\|_2 \leq B_{k,1}$, and $\sup_{\mathbf{x}\in\mathcal{X}} \text{tr}[\Delta_{\mathbf{x}'}^{-1} \Delta_{\mathbf{x}}^{-1} k(\mathbf{x}, \mathbf{x}')]\big|_{\mathbf{x}'=\mathbf{x}} \leq B_{k,2}$.

Considering that we are working with distributions in a discrete space, the aforementioned assumption is not too restrictive. A common choice for the kernel in discrete space is the *exponential Hamming kernel*: $k(\mathbf{x}, \mathbf{x}') = \exp(-H(\mathbf{x}, \mathbf{x}'))$, where $H(\mathbf{x}, \mathbf{x}') := \frac{1}{d} \sum_{i=1}^d \mathbb{1}\{x_i \neq x_i'\}$ is the normalized Hamming distance. We have the following guarantees regarding our KDSD-based SKGT's time-uniform type-I error control, as detailed in Appendix B.2.

Theorem 3.3. *Assuming that Assumptions 3.1 and 3.2 are satisfied, and setting* $s = \frac{1}{\sqrt{B_{k,0}B_p^2 + 2B_{k,1}B_p + B_{k,2}}}$, *then, under the null hypothesis* H_0 *in (1a), the KDSD-based SKGT (Algorithm 3) ensures:* $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.

3.2. bd-KSD-based Sequential Kernel Goodness-of-fit Testing

Unlike densities on unbounded domains commonly assumed to vanish at infinity, densities on compact domains may not necessarily vanish at the boundary. Hence, a direct application of a Stein operator on $\mathbb{R}^d$ may require the knowledge of normalized density at the boundary, which defeats the purpose of KSD testing for unnormalized models.

We recall the bd-KSD defined on a bounded domain $V \subset \mathbb{R}^d$. We consider a bounded smooth function $\mathbf{h} : \mathbb{R}^d \to \mathbb{R}^d$ such that $h_i(\partial V) = 0, \forall i = 1, \ldots, d$. For the unnormalized p on V, the bounded-domain Stein operator is defined as: $T_{p,\mathbf{h}}\mathbf{g}(\mathbf{x}) = \frac{1}{p(\mathbf{x})} \sum_{i=1}^d \frac{\partial}{\partial x_i}(p(\mathbf{x}) g_i(\mathbf{x}) h_i(\mathbf{x}))$. Given an RKHS $\mathcal{H}^d$ of vector-valued functions $\mathbf{g} : \mathbb{R}^d \to \mathbb{R}^d$, the bd-KSD is defined as:

$$\text{bd-KSD}(q\|p) := \sup_{\mathbf{g}\in\mathcal{H}^d, \|\mathbf{g}\|_{\mathcal{H}^d}\leq 1} \mathbb{E}_{\mathbf{x}\sim q}[T_{p,\mathbf{h}}\mathbf{g}(\mathbf{x})]. \quad (24)$$

Witness Function for bd-KSD. Assuming that k is the reproducing kernel of $\mathcal{H}$, we define the kernel embedding $\zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot) := (\nabla \log p(\mathbf{x}) \odot \mathbf{h}(\mathbf{x}) + \nabla \odot \mathbf{h}(\mathbf{x}))k(\mathbf{x}, \cdot) + \mathbf{h}(\mathbf{x}) \odot \nabla k(\mathbf{x}, \cdot)$, where we denote $\mathbf{u} \odot \mathbf{v} := (u_1 v_1, \cdots, u_d v_d)^\top$ as the element-wise multiplication. Leveraging the kernel embedding, we can express bd-KSD as:

$$\text{bd-KSD}(q\|p) = \sup_{\mathbf{g}\in\mathcal{H}^d, \|\mathbf{g}\|_{\mathcal{H}^d}\leq 1} \langle \mathbf{g}, \mathbb{E}_{\mathbf{x}\sim q} \zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot) \rangle_{\mathcal{H}^d}. \quad (25)$$

Moreover, building upon the variational formulation, the witness function of bd-KSD adopts the closed form:

$$\mathbf{g}^* = \frac{\mathbb{E}_{\mathbf{x}\sim q} \zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot)}{\|\mathbb{E}_{\mathbf{x}\sim q} \zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot)\|_{\mathcal{H}^d}}. \quad (26)$$

Similar to the payoff function defined in the previous section, we use the plug-in estimator:

$$\hat{\mathbf{g}}_t = \frac{\mathbb{E}_{\mathbf{x}\sim \hat{q}_{t-1}} \zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot)}{\left\|\mathbb{E}_{\mathbf{x}\sim \hat{q}_{t-1}} \zeta_{p,\mathbf{h}}(\mathbf{x}, \cdot)\right\|_{\mathcal{H}^d}}, \quad (27)$$

where $\hat{q}_{t-1}$ is the empirical distribution of a sequence of samples $\{\mathbf{x}_1, \ldots, \mathbf{x}_{t-1}\}$. Finally, leveraging the plug-in

Algorithm 4 bd-KSD-based SKGT

Input: significance level $\alpha \in (0,1)$, data stream $\{\mathbf{x}_t\}_{t\geq 1}$, where $\mathbf{x}_t \sim q$, $\lambda_1^{\text{ONS}} = 0$.
for t=1,2,... **do**
  Use $\mathbf{x}_1, \ldots, \mathbf{x}_{t-1}$ to compute $\hat{\mathbf{g}}_t$ as in (27);
  Compute bd-KSD payoff $f_t(\mathbf{x}_t) = s\,\langle \hat{\mathbf{g}}_t, \zeta_p(\mathbf{x}_t, \cdot)\rangle_{\mathcal{H}^d}$;
  Update the wealth process $\mathcal{K}_t$ as in (4);
  if $\mathcal{K}_t \geq 1/\alpha$ **then**
    Reject H_0 and stop;
  else
    compute $\lambda_{t+1}^{\text{ONS}}$ (Algorithm 1);
  end if
end for

estimator, we construct the payoff function as

$$f_t(\mathbf{x}_t) = s\,\langle \hat{\mathbf{g}}_t, \zeta_{p,\mathbf{h}}(\mathbf{x}_t, \cdot)\rangle_{\mathcal{H}^d}\,. \tag{28}$$

Assumption 3.4. Suppose that a constant C_p for the target distribution p exists, such that $\sup_{\mathbf{x}\in V} \|\nabla \log p(\mathbf{x})\|_2 \leq C_p$.

Assumption 3.5. Suppose that the kernel k satisfies: $\sup_{\mathbf{x}\in V} k(\mathbf{x},\mathbf{x}) \leq C_{k,0}$, $\sup_{\mathbf{x}\in V} \|\nabla k(\mathbf{x},\mathbf{x})\|_2 \leq C_{k,1}$, and $\sup_{\mathbf{x}\in V} \mathrm{tr}[\nabla_{\mathbf{x}'}\nabla_{\mathbf{x}} k(\mathbf{x},\mathbf{x}')]\big|_{\mathbf{x}'=\mathbf{x}} \leq C_{k,2}$.

Assumption 3.6. Suppose that constants $C_{h,0}$ and $C_{h,1}$ for the function $\mathbf{h}$ exist, such that $\sup_{\mathbf{x}\in V} \|\mathbf{h}(\mathbf{x})\|_2 \leq C_{h,0}$ and $\sup_{\mathbf{x}\in V} \|\nabla \odot \mathbf{h}(\mathbf{x})\|_2 \leq C_{h,1}$.

Notably, as we are currently dealing with a bounded domain, the above assumption is not overly restrictive. The following result guarantees the time-uniform type-I error control of our bd-KSD-based SKGT. The proof is deferred to Appendix B.2.

Theorem 3.7. *Assuming that Assumptions 3.4, 3.5, and 3.6 hold, and setting* $s = \frac{1}{\sqrt{2(C_p^2 C_{h,0}^2 + C_{h,1}^2)C_{k,0} + 2(C_p C_{h,0} + C_{h,1})C_{k,1} + C_{h,0}^2 C_{k,2}}}$, *then, under the null hypothesis* H_0 *in* (1a)*, the bd-KSD-based SKGT (Algorithm 4) satisfies:* $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.

4. Numerical Simulations

This section describes the experiments that demonstrate our tests' capacity to adapt to a problem's unknown difficulty while maintaining type-I error control.

4.1. Student's t versus Normal

In our first experiment, we consider testing the null hypothesis that the observed samples come from a Student's t distribution with 1 degree of freedom. Student's t distribution has the probability density function given by the formula:

$$f(t) \propto \left(1 + \frac{t^2}{\nu}\right)^{-(\nu+1)/2},$$

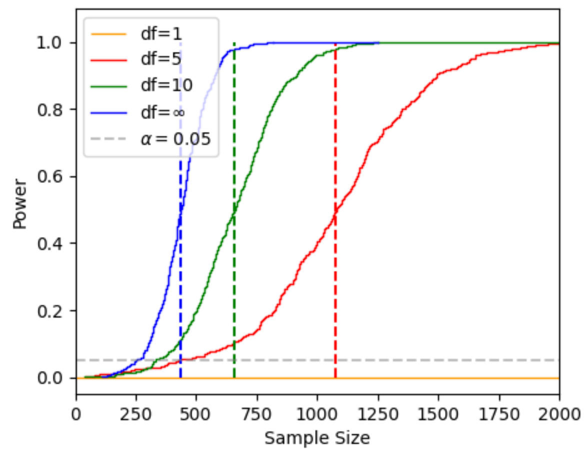

Figure 2. Ability of our sequential test to adapt to the problem's unknown difficulty while controlling the type-I error. Each solid curve is obtained by running 500 trials of the KSD-based SKGT (Algorithm 2) using samples from the Student's t distribution with the degrees of freedom at 1, 5, 10, and ∞. The dashed vertical line shows our SKGT's average stopping time. "df" denotes "degree of freedom."

where ν is the number of degrees of freedom. We consider samples from distributions with 1, 5, 10, and ∞, where ∞ is equivalent to sampling from a standard normal distribution. We employ our KSD-based SKGT (Algorithm 2) to test against the null hypothesis, utilizing the Gaussian kernel $k(x,y) = \exp(-|x-y|^2/2)$.

As displayed in Figure 2, we observed that under the null hypothesis, the type-I error is rigorously controlled. Furthermore, as the degrees of freedom approach 1, indicating a more challenging problem, the expected stopping time increases, thereby corroborating our theoretical result presented in Theorem 2.5.

4.2. Ising Model

In the second experiment, we consider testing the null hypothesis that the observed samples come from an Ising model. The Ising model (Ising, 1924) is a canonical example of a Markov random field. Consider an undirected graph $G = (V, E)$, where each vertex $i \in V$ is associated with a binary spin. The collection of spins form a random vector $\mathbf{x} = (x_1, \ldots, x_d) \in \{-1, +1\}^d$, whose components x_i and x_j directly interacts only if $(i,j) \in E$. The probability mass function is defined as:

$$p_\Theta(\mathbf{x}) = \frac{1}{Z(\Theta)} \exp\left(\sum_{(i,j)\in E} \theta_{ij} x_i x_j\right),$$

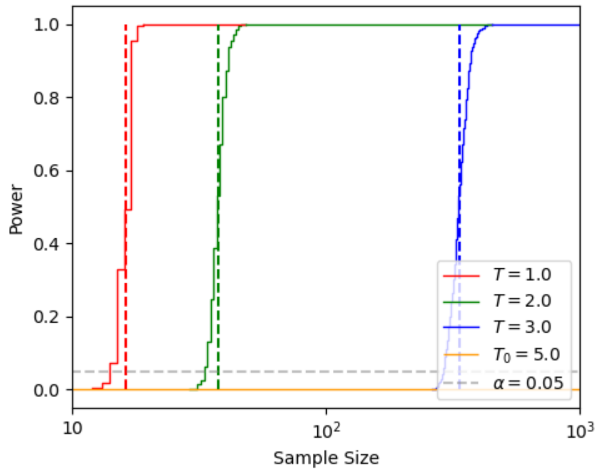

Figure 3. Our sequential test's capacity to adapt to a problem's unknown complexity while controlling the type-I error. Each solid curve is obtained by running 500 trials of the KDSD-based SKGT (Algorithm 3) using samples from the Ising models under temperatures 1, 2, 3, and 5. The dashed vertical lines show our SKGT's corresponding average stopping time.

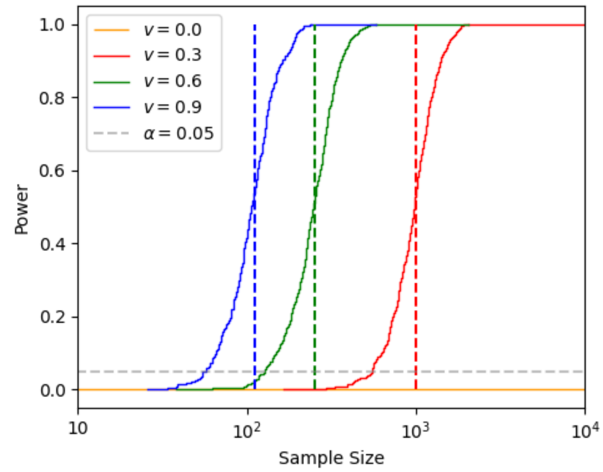

Figure 4. Our sequential test's ability to adapt to a problem's unknown difficulty while controlling the type-I error. Each solid curve is obtained by running 500 trials of the bd-KSD-based SKGT (Algorithm 4) using samples from the truncated Gaussians with shifts 0, 0.3, 0.6, and 0.9. The dashed vertical lines show our SKGT's corresponding average stopping time.

where θ_{ij} are the edge potentials and $Z(\Theta)$ is the partition function which is prohibitive to compute when d is high.

We consider a periodic 10-by-10 lattice, with $d = 100$ random variables. We focus on the ferromagnetic setting and set $\theta_{ij} = 1/T$, where T is the temperature of the system. For $T_0 = 5$ and various values of T, we test the hypotheses $H_0 : T = T_0$ vs. $H_1 : T \neq T_0$ using data samples drawn from the model under temperature T. To draw samples from the Ising model, we apply the Metropolis algorithm. We employ our KDSD-based SKGT (Algorithm 3) to test against the null hypothesis, using the exponential hamming kernel $k(\mathbf{x}, \mathbf{x}') = \exp(-H(\mathbf{x}, \mathbf{x}'))$, where $H(\mathbf{x}, \mathbf{x}') := \frac{1}{d}\sum_{i=1}^{d} \mathbb{1}\{x_i \neq x_i'\}$ is the normalized Hamming distance.

As illustrated in Figure 3, under the null hypothesis, our KDSD-based SKGT does not reject the null within 1000 samples. Conversely, under the alternative hypotheses, as the temperature approaches T_0, indicative of a more challenging task, the average stopping time increases.

4.3. Truncated Gaussians in Unit Ball

In the third experiment, we consider testing the null hypothesis that samples come from the standard Gaussian truncated in $\boldsymbol{B}_1(\mathbb{R}^3) = \{\mathbf{x} \in \mathbb{R}^3 : \|\mathbf{x}\|_2 \leq 1\}$. Specifically, our testing procedure is applied to samples generated from truncated Gaussians with various mean shifts, formulated as

$$q(\mathbf{x}) \propto \mathcal{N}\left(\boldsymbol{\mu}, \boldsymbol{I}_3\right), \mathbf{x} \in \boldsymbol{B}_1(\mathbb{R}^3),$$

where $\boldsymbol{\mu} = (v, v, v)^\top$ and $\boldsymbol{I}_3$ is the identity matrix in $\mathbb{R}^{3\times 3}$. We consider different values of $v \in \{0, 0.3, 0.6, 0.9\}$. We employ our bd-KSD-based KDSD, equipped with the Gaussian kernel $k(\mathbf{x}, \mathbf{y}) = \exp\left(-\|\mathbf{x} - \mathbf{y}\|_2^2/2\right)$. To satisfy boundary conditions and consider the rotational invariance of the unit ball, we set the auxiliary functions as $h_i(\mathbf{x}) = 1 - \|\mathbf{x}\|_2^2, \ i = 1, 2, 3$.

As depicted in Figure 4, under the null hypothesis, our bd-KSDs-based SKGT does not reject the null within 10^4 samples. Conversely, under the alternative hypothesis, as the mean shift decreases, indicating a more challenging task, the average stopping time increases.

5. Conclusion

In this paper, we introduce the SKGT based on the principle of testing by betting. The SKGT allows for continuous monitoring of data and adaptation to a problem's unknown complexity. After that, we provide formal guarantees regarding time-uniform type-I error control and the consistency of the KSD-based SKGT. Specifically, we demonstrate that the wealth process exhibits exponential growth and characterize the rate of wealth growth in terms of the true KSD. To handle data in discrete spaces and bounded domains, we propose SKGTs based on KDSD and bd-KSD. Our experiments demonstrate the adaptability of our tests to a problem's unknown complexity while maintaining type-I error control.

Impact Statement

To the best of our knowledge, this work has no negative social impact. This work mainly provides sequential kernel goodness-of-fit tests. Hence, our work may promote the development of the related fields.

Acknowledgements

This work is supported by the National Key R&D Program of China under Grant 2023YFC3604702, and the Fundamental Research Fund Program of LIESMARS.

References

Baum, J., Kanagawa, H., and Gretton, A. A kernel stein test of goodness of fit for sequential models. In *ICML*, 2023.

Carmeli, C., De Vito, E., Toigo, A., and Umanitá, V. Vector valued reproducing kernel hilbert spaces and universality. *Analysis and Applications*, 2010.

Chwialkowski, K., Strathmann, H., and Gretton, A. A kernel test of goodness of fit. In *ICML*, 2016.

Cutkosky, A. and Orabona, F. Black-box reductions for parameter-free online learning in banach spaces. In *COLT*, 2018.

Darling, D. A. and Robbins, H. Some nonparametric sequential tests with power one. *Proceedings of the National Academy of Sciences*, 1968.

Fan, X., Grama, I., and Liu, Q. Exponential inequalities for martingales with applications. *Electronic Journal of Probability*, 2015.

Gorham, J. and Mackey, L. Measuring sample quality with stein's method. In *NeurIPS*, 2015.

Grünwald, P., Henzi, A., and Lardy, T. Anytime-valid tests of conditional independence under model-x. *Journal of the American Statistical Association*, 2023.

Hazan, E., Agarwal, A., and Kale, S. Logarithmic regret algorithms for online convex optimization. *Machine Learning*, 2007.

Hoeffding, W. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 1963.

Ising, E. *Beitrag zur theorie des ferro-und paramagnetismus*. PhD thesis, Grefe & Tiedemann Hamburg, Germany, 1924.

Jitkrittum, W., Xu, W., Szabó, Z., Fukumizu, K., and Gretton, A. A linear-time kernel goodness-of-fit test. In *NeurIPS*, 2017.

Kolmogorov, A. Sulla determinazione empirica di una legge didistribuzione. *Giorn Dell'inst Ital Degli Att*, 1933.

Liu, Q., Lee, J., and Jordan, M. A kernelized stein discrepancy for goodness-of-fit tests. In *ICML*, 2016.

Liu, X., Duncan, A. B., and Gandy, A. Using perturbation to improve goodness-of-fit tests based on kernelized stein discrepancy. In *ICML*, 2023.

Podkopaev, A. and Ramdas, A. Sequential predictive two-sample and independence testing. In *NeurIPS*, 2023.

Podkopaev, A., Blöbaum, P., Kasiviswanathan, S., and Ramdas, A. Sequential kernelized independence testing. In *ICML*, 2023.

Shaer, S., Maman, G., and Romano, Y. Model-x sequential testing for conditional independence via testing by betting. In *AISTATS*, 2023.

Shafer, G. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society Series A: Statistics in Society*, 2021.

Shafer, G. and Vovk, V. *Game-theoretic foundations for probability and finance*. 2019.

Shekhar, S. and Ramdas, A. Nonparametric two-sample testing by betting. *IEEE Transactions on Information Theory*, 2023.

Smirnov, N. Table for estimating the goodness of fit of empirical distributions. *The annals of mathematical statistics*, 1948.

Ville, J. *Étude critique de la notion de collectif*. Thèses de l'entre-deux-guerres. Gauthier-Villars, 1939.

Xu, W. Standardisation-function kernel stein discrepancy: A unifying view on kernel stein discrepancy tests for goodness-of-fit. In *AISTATS*, 2022.

Yang, J., Liu, Q., Rao, V., and Neville, J. Goodness-of-fit testing for discrete distributions via stein discrepancy. In *ICML*, 2018.

A. Goodness-of-fit Testing for Streaming Data

A.1. Failure of Batch KSD under Continuous Monitoring

It is straightforward to estimate the squared Stein discrepancy from samples $\{Z_i\}_{i=1}^n$: a quadratic time estimator is a V-statistic, and takes the form:

$$V_n = \frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n h_p(Z_i, Z_j).$$

For kernel k, we choose the Gaussian kernel $k(x,x') = \exp(\|x-x'\|_2^2/2)$. To conduct goodness-of-fit testing using batch KSD, we use the Wild Bootstrap Testing (Chwialkowski et al., 2016). A simple Markov chain taking values in $\{-1,1\}$, starting from $W_{1,n}=1$,

$$W_{t,n} = \mathbb{1}\{U_t > a_n\}W_{t-1,n} - \mathbb{1}\{U_t < a_n\}W_{t-1,n},$$

where the U_t are uniform $(0,1)$ i.i.d. random variables and a_n is the probability of $W_{t,n}$ changing sign. For i.i.d. data, we set $a_n = 0.5$. This leads to a bootstrapped V-statistic

$$B_n = \frac{1}{n^2}\sum_{i=1}^n\sum_{j=1}^n W_{i,n}W_{j,n}h_p(Z_i, Z_j).$$

Having obtained wild bootstrap samples $\{B_{n,i}\}_{i=1}^D$, we use the p-value: $P = \frac{1}{D}\sum_{i=1}^D \mathbb{1}\{B_{n,i} > V_n\}$, with $D = 1000$ wild bootstrap samples. Next, we study batch KSD under (a) *fixed-time* and (b) *continuous* monitoring. We consider a simple case when Z is from the student-t distribution with degree of freedom 1 and the target distribution is also the student-t distribution with density $\frac{1}{\pi}\left(1+t^2\right)^{-1}$. We conduct a test at 12 different samples sizes: $t \in \{50, 100, \ldots, 600\}$:

(a) Under fixed-time monitoring, for each value of t, we sample a sequence $Z_1, \ldots, Z_t$ (150 times for each t) and conduct batch-KSD test. The goal is to confirm that batch-KSD controls type I error by tracking the standard miscoverage rate.

(b) Under continuous monitoring, we sample new datapoints and re-conduct the test. We illustrate inflated type I error by tracking the *cumulative miscoverage rate*, that is, the fraction of times, the test falsely rejects the null.

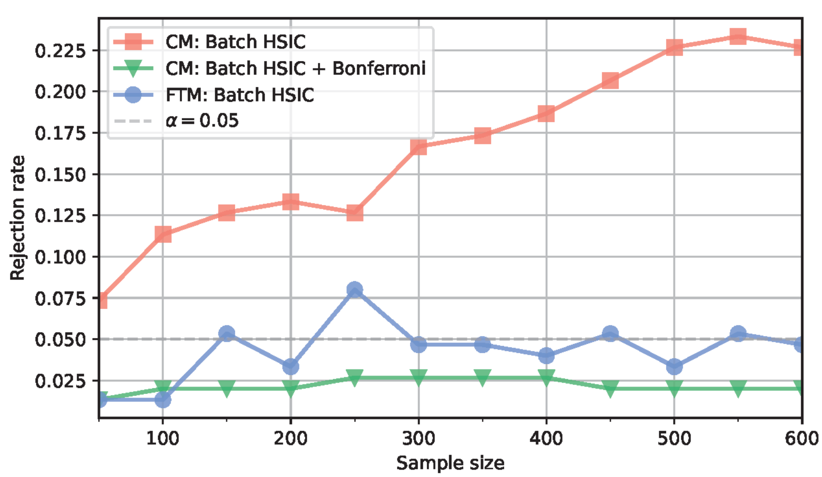

Figure 5. Inflated false discovery rate of batch KSD under continuous monitoring (CM, red line with squares) for the case when $p = q$. Bonferroni correction (CM, green line with triangles) restore type I error control. Type I error is controlled at a specified level under fixed-time monitoring (FTM, blue line with circles).

The results are presented in Figure 5. For Bonferroni correction, we compute p-values every 50 samples and decompose the error budget as: $\alpha = \sum_{i=1}^\infty \frac{\alpha}{i(i+1)}$, that is, for i-th test we use threshold $\alpha_i = \frac{\alpha}{i(i+1)}$.

B. Proofs

B.1. Proofs for Section 2

Theorem 2.2. *Let $\mathcal{G}$ be the RKHS constructed from a C_0-universal kernel and $\mathcal{G}^d$ be the corresponding product RKHS:*

1. *Under H_0 in (1a), any payoff function of the form (8) satisfies $\mathbb{E}_{H_0}[f(Z)] = 0$.*
2. *Under H_1 in (1b), the oracle payoff function f^* based on the witness function g^* satisfies $\mathbb{E}_{H_1}[f^*(Z)] > 0$. Further, for λ^* defined in (9), it holds that $\mathbb{E}_{H_1}[\log(1+\lambda^* f^*(Z))] > 0$. Hence, $\mathcal{K}_t^* \overset{\text{a.s.}}{\to} +\infty$, which implies that the oracle test is consistent: $\mathbb{P}_{H_1}(\tau^* < \infty) = 1$, where $\tau^* = \inf\{t \geq 1 : \mathcal{K}_t^* \geq 1/\alpha\}$.*

Proof. 1. Under H_0 in (1a), we have that:

$$\begin{aligned}
\mathbb{E}_{Z\sim p}[f^*(Z)] &= \mathbb{E}_{Z\sim p}[s \cdot (\langle g^*(Z_t), \nabla \log p(Z_t)\rangle + \nabla \cdot g^*(Z_t))] \\
&= s \int_{\mathbb{R}^d} (\langle g^*(x), \nabla \log p(x)\rangle + \nabla \cdot g^*(x))\, p(x) \mathrm{d}x \\
&= s \int_{\mathbb{R}^d} \sum_{i=1}^d g_i^*(x) \frac{\partial p(x)}{\partial x_i} \mathrm{d}x + s \int_{\mathbb{R}^d} (\nabla \cdot g^*(x)) p(x) \mathrm{d}x \\
&\overset{\text{(i)}}{=} -s \int_{\mathbb{R}^d} \sum_{i=1}^d p(x) \frac{\partial g_i^*(x)}{\partial x_i} \mathrm{d}x + s \int_{\mathbb{R}^d} (\nabla \cdot g^*(x)) p(x) \mathrm{d}x \\
&= 0,
\end{aligned} \tag{29}$$

where equality (i) follows from integration by parts.

2. Under H_1 in (1b) and the i.i.d. setting, we have

$$\mathbb{E}[f^*(Z_t) | \mathcal{F}_{t-1}] = \mathbb{E}[f^*(Z)] = s \cdot S_p(q) > 0.$$

Let $W := f^*(Z)$, and consider the quantity $\mathbb{E}_{H_1}[\log(1+\lambda W)]$. We use the following inequality (Fan et al., 2015, Equation (4.12)): for any $y \geq -1$ and $\lambda \in [0, 1)$, we have

$$\log(1+\lambda y) \geq \lambda y + y^2(\log(1-\lambda) + \lambda).$$

Hence, we get the following:

$$\mathbb{E}_{H_1}[\log(1+\lambda W)]\lambda \geq \lambda \mathbb{E}_{H_1}[W] + \mathbb{E}_{H_1}[W^2](\log(1-\lambda) + \lambda).$$

Finally, by choosing $\lambda^* = \mathbb{E}_{H_1}[W]/(\mathbb{E}_{H_1}[W] + \mathbb{E}_{H_1}[W^2]) \in (0, 1)$ and the fact that $\log(1-x) + x \geq -\frac{x^2}{2(1-x)}$ for $x \in [0, 1)$, we arrive at

$$\mathbb{E}_{H_1}[\log(1+\lambda^* W)] \geq \frac{1}{2} \cdot \frac{(\mathbb{E}_{H_1}[W])^2}{\mathbb{E}_{H_1}[W] + \mathbb{E}_{H_1}[W^2]} > 0.$$

The wealth process corresponding to the oracle test satisfies:

$$\mathcal{K}_t^* = \prod_{i=1}^t (1 + \lambda^* f^*(Z_t)) = \exp\left(t \cdot \frac{1}{t} \sum_{i=1}^t \log(1 + \lambda^* f^*(Z_i))\right)$$

By the Strong Law of Large Numbers (SLLN), we have:

$$\frac{1}{t} \sum_{i=1}^t \log(1+\lambda^* f^*(Z_i)) \xrightarrow{\text{a.s.}} \mathbb{E}_{H_1}[\log(1+\lambda^* f^*(Z))] > 0.$$

The above result implies that $\mathcal{K}_t^* \xrightarrow{\text{a.s.}} +\infty$, thereby ensuring the consistency of the oracle test. □

Theorem 2.5. *Under Assumptions 2.3 and 2.4, the following claims hold for KSD-based SKGT (Algorithm 2):*

1. *Under H_0 in (1a), SKGT stops with probability at most α: $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.*
2. *Under H_1 in (1b), then it holds that $\mathcal{K} \xrightarrow{\text{a.s.}} +\infty$ and thus the SKGT is consistent: $\mathbb{P}_{H_1}(\tau < \infty) = 1$. Furthermore, the wealth grows exponentially, and the corresponding growth rate satisfies the following formula:*

$$\liminf_{t \to \infty} \frac{\log \mathcal{K}_t}{t} \geq \frac{\mathbb{E}_{H_1}[f^*(Z)]}{4} \left(\frac{\mathbb{E}_{H_1}[f^*(Z)]}{\mathbb{E}_{H_1}[(f^*(Z))^2]} \wedge 1 \right) \tag{14}$$

almost surely, where f^* is the oracle payoff defined in (11).

Proof. 1. First, let us show that the predictable estimates of the oracle payoff function are bounded when the scaling factors $s = 1/2$ is used. Recall that:

$$f_t(x) = \frac{1}{2} \langle \hat{g}_t, \nabla \log p(x) k(x, \cdot) + \nabla k(x, \cdot) \rangle_{\mathcal{G}^d} \tag{30}$$

The Cauchy-Schwartz inequality implies that

$$\begin{aligned} |f_t(x)| &\leq \frac{1}{2} \|\hat{g}_t\|_{\mathcal{G}^d} \|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d} + \frac{1}{2} \|\hat{g}_t\|_{\mathcal{G}^d} \|\nabla k(x, \cdot)\|_{\mathcal{G}^d} \\ &= \frac{1}{2} \|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d} + \frac{1}{2} \|\nabla k(x, \cdot)\|_{\mathcal{G}^d}, \end{aligned}$$

where the second equality follows from the constraint that $\|\hat{g}_t\|_{\mathcal{G}^d} = 1$. The first term $\|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d}$ can be computed as:

$$\begin{aligned} \|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d} &= \sqrt{\sum_{i=1}^{d} \langle \partial_{x_i} \log p(x) k(x, \cdot), \partial_{x_i} \log p(x) k(x, \cdot) \rangle} \\ &= \|\nabla \log p(x)\|_2 \sqrt{k(x, x)} \leq 1, \end{aligned}$$

where the last inequality follows from Assumption 2.3 and 2.4. The second term $\|\nabla k(x, \cdot)\|_{\mathcal{G}^d}$ is calculated as follows:

$$\begin{aligned} \|\nabla k(x, \cdot)\|_{\mathcal{G}^d} &= \sqrt{\sum_{i=1}^{d} \langle \partial_{x_i} k(x, \cdot), \partial_{x_i} k(x, \cdot) \rangle_{\mathcal{G}}} \\ &= \sqrt{\sum_{i=1}^{d} \partial_{x_i} \langle \partial_{x_i} k(x, \cdot), k(x, \cdot) \rangle_{\mathcal{G}}} \\ &= \sqrt{\sum_{i=1}^{d} \frac{\partial^2 k(x, y)}{\partial x_i \partial y_i} \Big|_{y=x}} \leq 1, \end{aligned}$$

where the last inequality follows from Assumption 2.3. Hence $f_t(x) \in [-1, 1]$. Next, we show that the constructed payoff function yields a fair bet. Note that f_t is constructed from $\{Z_s\}_{s=1}^{t-1}$; thus f_t is $\mathcal{F}_{t-1}$-measurable. Based on this fact, we have:

$$\mathbb{E}[f_t(Z_t) | \mathcal{F}_{t-1}] = \frac{1}{2} \mathbb{E}_{Z_t \sim q}[\langle \hat{g}_t(Z_t), \nabla \log p(Z_t) \rangle + \nabla \cdot \hat{g}_t(Z_t)],$$

and in particular, the above implies that $\mathbb{E}_{H_0}[f_t(Z_t) | \mathcal{F}_{t-1}] = 0$ for H_0 in (1a). In the following, we show that the resulting wealth process is a non-negative martingale for all strategies for selecting betting fractions that are considered in this work. Since the betting fractions are predictable, i.e. λ_t is $\mathcal{F}_{t-1}$-measurable, we have:

$$\begin{aligned} \mathbb{E}_{H_0}[\mathcal{K}_t | \mathcal{F}_{t-1}] &= \mathbb{E}_{H_0}[\mathcal{K}_{t-1}(1 + \lambda_t f_t(Z_t)) | \mathcal{F}_{t-1}] \\ &= \mathcal{K}_{t-1} + \lambda_t \mathbb{E}_{H_0}[f_t(Z_t) | \mathcal{F}_{t-1}] \\ &= \mathcal{K}_{t-1}. \end{aligned}$$

The assertion of the theorem then follows directly from Ville's inequality (Theorem B.6) when $a = 1/\alpha$.

2. In the following, we establish the consistency of KSD-based SKGT with the ONS betting strategy. Under the ONS betting strategy, for any sequence of outcomes $\{f_t\}_{t\geq 1}$, $f_t \in [-,1,1], i \geq 1$, it holds that (Cutkosky & Orabona, 2018, Proof of Theorem 1):

$$\log \mathcal{K}_t(\lambda_0) - \log \mathcal{K}_t = O\left(\log\left(\sum_{i=1}^t f_i^2\right)\right), \tag{31}$$

where $\mathcal{K}_t(\lambda_0)$ is the wealth process of any constant betting strategy $\lambda_0 \in [-1/2, 1/2]$ and $\mathcal{K}_t$ is the wealth process corresponding to the ONS strategy. It follows that the wealth process corresponding to the ONS strategy satisfies

$$\frac{\log \mathcal{K}_t}{t} \geq \frac{\log \mathcal{K}_t(\lambda_0)}{t} - C \cdot \frac{\log t}{t}, \tag{32}$$

for some absolute constant $C > 0$. Next, let us consider:

$$\lambda_0 = \frac{1}{2}\left(\left(\frac{\sum_{i=1}^t f_i}{\sum_{i=1}^t f_i^2} \wedge 1\right) \vee 0\right).$$

We obtain:

$$\begin{aligned}\frac{\log \mathcal{K}_t(\lambda_0)}{t} &= \frac{1}{t}\sum_{i=1}^t \log(1+\lambda_0 f_i)\\ &\overset{(a)}{\geq} \frac{1}{t}\sum_{i=1}^t \left(\lambda_0 f_i - \lambda_0^2 f_i^2\right)\\ &= \left(\frac{\frac{1}{t}\sum_{i=1}^t f_i}{4} \vee 0\right)\cdot\left(\frac{\frac{1}{t}\sum_{i=1}^t f_i}{\frac{1}{t}\sum_{i=1}^t f_i^2} \wedge 1\right),\end{aligned} \tag{33}$$

where in (a) we use the inequality: $\log(1+x) \geq x - x^2$ for $x \in [-1/2, 1/2]$. From Lemma B.5, it follows for $f_i = f_i(Z_i)$ that:

$$\frac{\frac{1}{t}\sum_{i=1}^t f_i(Z_i)}{4} \cdot \left(\frac{\frac{1}{t}\sum_{i=1}^t f_i(Z_i)}{\frac{1}{t}\sum_{i=1}^t (f_i(Z_i))^2} \wedge 1\right) \xrightarrow{\text{a.s.}} \frac{\mathbb{E}[f^*(Z)]}{4}\cdot\left(\frac{\mathbb{E}[f^*(Z)]}{\mathbb{E}[(f^*(Z))^2]} \wedge 1\right). \tag{34}$$

Using (32), we conclude that the growth rate of the ONS wealth process satisfies

$$\liminf_{t\to\infty} \frac{\log \mathcal{K}_t}{t} \geq \frac{\mathbb{E}[f^*(Z)]}{4}\cdot\left(\frac{\mathbb{E}[f^*(Z)]}{\mathbb{E}[(f^*(Z))^2]} \wedge 1\right). \tag{35}$$

We conclude that the test is consistent, that is, if H_1 is true, then $\mathbb{P}(\tau < \infty) = 1$.

□

B.2. Proofs for Section 3

Theorem 3.3. *Assuming that Assumptions 3.1 and 3.2 are satisfied, and setting* $s = \frac{1}{\sqrt{B_{k,0}B_p^2 + 2B_{k,1}B_p + B_{k,2}}}$, *then, under the null hypothesis* H_0 *in (1a), the KDSD-based SKGT (Algorithm 3) ensures:* $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.

Proof. It suffices to show that the proposed payoff functions are bounded. Note that:

$$\begin{aligned}\left|\langle \hat{\mathbf{g}}_t, \eta_p(\mathbf{x}_t,\cdot)\rangle_{\mathcal{H}^d}\right| &\leq \|\hat{\mathbf{g}}_t\|_{\mathcal{H}^d}\|\eta_p(\mathbf{x}_t,\cdot)\|_{\mathcal{H}^d}\\ &= \|\eta_p(\mathbf{x}_t,\cdot)\|_{\mathcal{H}^d}\\ &= \sqrt{k(\mathbf{x}_t,\mathbf{x}_t)\|\mathbf{s}_p(\mathbf{x}_t)\|_2^2 + 2\mathbf{s}_p(\mathbf{x}_t)^\top \Delta_{\mathbf{x}}^{-1} k(\mathbf{x}_t,\mathbf{x}_t) + \operatorname{tr}[\Delta_{\mathbf{x}'}^{-1}\Delta_{\mathbf{x}}^{-1} k(\mathbf{x},\mathbf{x}')]\big|_{\mathbf{x}'=\mathbf{x}}}\\ &\leq \sqrt{B_{k,0}B_p^2 + 2B_{k,1}B_p + B_{k,2}},\end{aligned}$$

where the first inequality follows from the Cauchy-Schwartz inequality and the second inequality is due to Assumption 3.1 and Assumption 3.2. We conclude that any predictive estimate of the oracle payoff function for KDSD satisfies

$$|f_t(\mathbf{x}_t)| \leq 1.$$

Next, we show that the constructed payoff function yields a fair bet. Note that f_t is constructed from $\{\mathbf{x}_s\}_{s=1}^{t-1}$; thus f_t is $\mathcal{F}_{t-1}$-measurable. Based on this fact, we have:

$$\mathbb{E}[f_t(\mathbf{x}_t)|\mathcal{F}_{t-1}] = s\mathbb{E}_{\mathbf{x}_t\sim q}[\langle \hat{\mathbf{g}}_t(\mathbf{x}_t), \mathbf{s}_p(\mathbf{x}_t)\rangle + \Delta^{-1}\cdot \hat{\mathbf{g}}_t(\mathbf{x}_t)],$$

and in particular, the above implies that $\mathbb{E}_{H_0}[f_t(\mathbf{x}_t)|\mathcal{F}_{t-1}] = 0$ for H_0 in (1a). In the following, we show that the resulting wealth process is a non-negative martingale for all strategies for selecting betting fractions that are considered in this work. Since the betting fractions are predictable, i.e. λ_t is $\mathcal{F}_{t-1}$-measurable, we have:

$$\begin{aligned}\mathbb{E}_{H_0}[\mathcal{K}_t|\mathcal{F}_{t-1}] &= \mathbb{E}_{H_0}[\mathcal{K}_{t-1}(1+\lambda_t f_t(\mathbf{x}_t))|\mathcal{F}_{t-1}] \\ &= \mathcal{K}_{t-1} + \lambda_t \mathbb{E}_{H_0}[f_t(\mathbf{x}_t)|\mathcal{F}_{t-1}] \\ &= \mathcal{K}_{t-1}.\end{aligned}$$

The assertion of the theorem then follows directly from Ville's inequality (Theorem B.6) when $a = 1/\alpha$. ☐

Theorem 3.7. *Assuming that Assumptions 3.4, 3.5, and 3.6 hold, and setting* $s = \frac{1}{\sqrt{2(C_p^2C_{h,0}^2+C_{h,1}^2)C_{k,0}+2(C_pC_{h,0}+C_{h,1})C_{k,1}+C_{h,0}^2C_{k,2}}}$*, then, under the null hypothesis* H_0 *in (1a), the bd-KSD-based SKGT (Algorithm 4) satisfies:* $\mathbb{P}_{H_0}(\tau < \infty) \leq \alpha$.

Proof. It suffices to show that the proposed payoff functions are bounded. Note that:

$$\begin{aligned}&\left|\langle \hat{\mathbf{g}}_t, \zeta_{p,\mathbf{h}}(\mathbf{x}_t,\cdot)\rangle_{\mathcal{H}^d}\right|^2 \\ &\leq \|\hat{\mathbf{g}}_t\|_{\mathcal{H}^d}^2 \|\zeta_{p,\mathbf{h}}(\mathbf{x}_t,\cdot)\|_{\mathcal{H}^d}^2 \\ &= \underbrace{\|\nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t) + \nabla\odot\mathbf{h}(\mathbf{x}_t)\|_2^2\, k(\mathbf{x}_t,\mathbf{x}_t)}_{\text{(I)}} \\ &\quad + 2\underbrace{\langle \nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t) + \nabla\odot\mathbf{h}(\mathbf{x}_t), \nabla k(\mathbf{x}_t,\mathbf{x}_t)\rangle}_{\text{(II)}} \\ &\quad + \underbrace{\sum_{i=1}^d h_i^2(\mathbf{x}_t)\frac{\partial^2 k(\mathbf{x},\mathbf{x}')}{\partial x_i \partial x_i'}\Big|_{\mathbf{x}=\mathbf{x}'=\mathbf{x}_t}}_{\text{III}},\end{aligned}$$

where the first inequality is due to Cauchy-Schwartz inequality. For the term (I), we have

$$\begin{aligned}\|\nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t) + \nabla\odot\mathbf{h}(\mathbf{x}_t)\|_2^2\, k(\mathbf{x}_t,\mathbf{x}_t) &\leq 2(\|\nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t)\|_2^2 + \|\nabla\odot\mathbf{h}(\mathbf{x}_t)\|_2^2)k(\mathbf{x}_t,\mathbf{x}_t) \\ &\leq 2(\|\nabla \log p(\mathbf{x}_t)\|_2^2\,\|\mathbf{h}(\mathbf{x}_t)\|_2^2 + \|\nabla\odot\mathbf{h}(\mathbf{x}_t)\|_2^2)k(\mathbf{x}_t,\mathbf{x}_t) \\ &\leq 2(C_{h,0}^2C_p^2 + C_{h,1}^2)C_{k,0},\end{aligned} \tag{36}$$

where the first inequality follows from the triangle inequality and $(x+y)^2 \leq 2x^2+2y^2,\ \forall x,y\in\mathbb{R}$, the second inequality is due to the fact that $\|\mathbf{u}\odot\mathbf{v}\|_2^2 \leq \|\mathbf{u}\|_2^2\,\|\mathbf{v}\|_2^2$, and the third inequality follows from Assumptions 3.4, 3.5 and 3.6. For the term (II), we have

$$\begin{aligned}\langle \nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t) + \nabla\odot\mathbf{h}(\mathbf{x}_t), \nabla k(\mathbf{x}_t,\mathbf{x}_t)\rangle &\leq 2(\|\nabla \log p(\mathbf{x}_t)\odot \mathbf{h}(\mathbf{x}_t)\|_2 + \|\nabla\odot\mathbf{h}(\mathbf{x}_t)\|_2)\,\|\nabla k(\mathbf{x}_t,\mathbf{x}_t)\|_2 \\ &\leq 2(C_pC_{h,0}+C_{h,1})C_{k,1},\end{aligned} \tag{37}$$

where the first inequality follows from the Cauchy-Schwartz inequality and the second inequality is due to our assumptions and the fact that $\|\mathbf{u} \odot \mathbf{v}\|_2^2 \leq \|\mathbf{u}\|_2^2 \|\mathbf{v}\|_2^2$. For the term (III), we have

$$\begin{aligned}\sum_{i=1}^{d} h_i^2(\mathbf{x}_t) \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x_i'}\Big|_{\mathbf{x}=\mathbf{x}'=\mathbf{x}_t} &\leq \sum_{i=1}^{d} \|\mathbf{h}(\mathbf{x}_t)\|_2^2 \frac{\partial^2 k(\mathbf{x}, \mathbf{x}')}{\partial x_i \partial x_i'}\Big|_{\mathbf{x}=\mathbf{x}'=\mathbf{x}_t} \\ &\leq C_{h,0}^2 C_{k,2},\end{aligned}$$

where the first inequality is because $u_i^2 \leq \|\mathbf{u}\|_2^2$, for all $i = 1, \ldots, d$, and the second inequality is due to our assumptions. Combining these terms, we obtain $\left|\langle \hat{\mathbf{g}}_t, \zeta_{p,\mathbf{h}}(\mathbf{x}_t, \cdot)\rangle_{\mathcal{H}^d}\right| \leq \sqrt{2(C_p^2 C_{h,0}^2 + C_{h,1}^2)C_{k,0} + 2(C_p C_{h,0} + C_{h,1})C_{k,1} + C_{h,0}^2 C_{k,2}}$. By setting $s = \frac{1}{\sqrt{2(C_p^2 C_{h,0}^2 + C_{h,1}^2)C_{k,0} + 2(C_p C_{h,0} + C_{h,1})C_{k,1} + C_{h,0}^2 C_{k,2}}}$, we conclude that $|f_t(\mathbf{x}_t)| \leq 1$. Following the same procedure as in the proof of Theorem 3.3, we can show that the resulting wealth process is a martingale. The assertion of the theorem directly follows from Ville's inequality. □

B.3. Supporting Lemmas

To begin, let's review the definition of $h_p(x, y)$ before presenting the first result.

$$\begin{aligned}h_p(x, y) := &\nabla \log p(x)^\top \nabla \log p(y) k(x, y) + \nabla \log p(y)^\top \nabla_x k(x, y) \\ &+ \nabla \log p(x)^\top \nabla_y k(x, y) + \langle \nabla_x k(x, y), \nabla_y k(x, y)\rangle_{\mathcal{G}^d},\end{aligned}$$

Lemma B.1 (Convergence of V-statistic). *Suppose that Assumption 2.3 and 2.4 hold. Define* $V_n = \frac{1}{n^2}\sum_{i=1}^{n}\sum_{i=1}^{n} h_p(Z_i, Z_j)$ *and* $S_p^2(q) = \mathbb{E}_{Z,Z'}[h_p(Z, Z')]$, *where* $\{Z_i\}_{i=1}^n$ *is a sample from distribution* q *and* Z' *is independent of* Z *with identical distribution* q. *Then for* $n > 1$ *and any* $\delta \in (0, 1)$, *it holds with probability of at least* $1 - \delta$ *that:*

$$\left|V_n - S_p^2(q)\right| \leq 8\sqrt{\frac{\log(2/\delta)}{n}} + \frac{8}{n}.$$

Proof. First, note that V_n is a biased estimator of $S_p^2(q)$:

$$\mathbb{E}[V_n] = \frac{n-1}{n}\mathbb{E}_{Z,Z'}[h_p(Z, Z')] + \frac{1}{n}\mathbb{E}_Z[h_p(Z, Z)].$$

The triangle inequality implies:

$$|V_n - S_p^2(q)| \leq |V_n - \mathbb{E}[V_n]| + \frac{1}{n}\left|S_p^2(q) - \mathbb{E}_Z[h_p(Z, Z)]\right|$$

Then the problem is decomposed into upper bounding the first and second term separately.

(i) Upper bounding $\left|S_p^2(q) - \mathbb{E}_Z[h_p(Z, Z)]\right|$.

We start with providing an upper bound for $h_p(x, y)$. $h_p(x, y)$ can be written as:

$$h_p(x, y) = \langle \nabla \log p(x) k(x, \cdot) + \nabla k(x, \cdot), \nabla \log p(y) k(y, \cdot) + \nabla k(y, \cdot)\rangle_{\mathcal{G}^d}$$

Using Cauchy-Schwartz inequality, we obtain:

$$|h_p(x, y)| \leq \|\nabla \log p(x) k(x, \cdot) + \nabla k(x, \cdot)\|_{\mathcal{G}^d} \cdot \|\nabla \log p(y) k(y, \cdot) + \nabla k(y, \cdot)\|_{\mathcal{G}^d}$$

By the triangle inequality of norm $\|\cdot\|_{\mathcal{G}^d}$, we have:

$$\|\nabla \log p(x) k(x, \cdot) + \nabla k(x, \cdot)\|_{\mathcal{G}^d} \leq \|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d} + \|\nabla k(x, \cdot)\|_{\mathcal{G}^d}$$

The first term $\|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d}$ can be computed as:

$$\begin{aligned}\|\nabla \log p(x) k(x, \cdot)\|_{\mathcal{G}^d} &= \sqrt{\sum_{i=1}^{d} \langle \partial_{x_i} \log p(x) k(x, \cdot), \partial_{x_i} \log p(x) k(x, \cdot)\rangle} \\ &= \|\nabla \log p(x)\|_2 \sqrt{k(x, x)} \leq 1,\end{aligned}$$

where the last inequality follows from Assumption 2.3 and 2.4. The second term $\|\nabla k(x,\cdot)\|_{\mathcal{G}^d}$ is calculated as follows:

$$\begin{aligned}\|\nabla k(x,\cdot)\|_{\mathcal{G}^d} &= \sqrt{\sum_{i=1}^d \langle \partial_{x_i} k(x,\cdot), \partial_{x_i} k(x,\cdot)\rangle_{\mathcal{G}}} \\ &= \sqrt{\sum_{i=1}^d \partial_{x_i} \langle \partial_{x_i} k(x,\cdot), k(x,\cdot)\rangle_{\mathcal{G}}} \\ &= \sqrt{\sum_{i=1}^d \frac{\partial^2 k(x,y)}{\partial x_i \partial y_i}\Big|_{y=x}} \le 1,\end{aligned}$$

where the last inequality follows from Assumption 2.3. Hence $\|\nabla \log p(x) k(x,\cdot) + \nabla k(x,\cdot)\|_{\mathcal{G}^d} \le 2$. The same argument also holds for $\|\nabla \log p(y) k(y,\cdot) + \nabla k(y,\cdot)\|_{\mathcal{G}^d}$, therefore, we obtain $|h_p(x,y)| \le 4$ for any $x, y \in \mathbb{R}^d$. It implies that

$$\left|S_p^2(q) - \mathbb{E}_Z[h_p(Z,Z)]\right| \le 8. \tag{38}$$

(ii) Upper bounding $|V_n - \mathbb{E}[V_n]|$.

As stated in the seminar work of Hoeffding (1963), any V-statistic can be written as a U-statistic :

$$V_n = \frac{1}{n(n-1)} \sum_{i \ne j} h_p^*(Z_i, Z_j),$$

where $h_p^*(Z_i,Z_j) = \frac{n-1}{n} h_p(Z_i,Z_j) + \frac{1}{n} h_p(Z_i,Z_i)$ for $i \ne j$. We have shown that $h_p(x,y) \in [-4,4]$, then it also holds that $h_p^*(Z_i,Z_j) \in [-4,4]$. We are ready to use the result of Hoeffding (1963, Equation (5.7)):

$$\mathbb{P}(|V_n - \mathbb{E}[V_n]| \ge t) \le 2\exp(-nt^2/64).$$

Choosing $t = 8\sqrt{\frac{\log(2/\delta)}{n}}$, then we have, with probability of at least $1-\delta$:

$$|V_n - \mathbb{E}[V_n]| \le 8\sqrt{\frac{\log(2/\delta)}{n}}. \tag{39}$$

Combining (38) and (39), we have, with probability of at least $1-\delta$:

$$|V_n - S_p^2(q)| \le 8\sqrt{\frac{\log(2/\delta)}{n}} + \frac{8}{n}.$$

□

Now, we recall the definition of $\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]$:

$$\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)] := \frac{1}{t-1} \sum_{i=1}^{t-1} \xi_p(Z_i,\cdot) \tag{40}$$

Lemma B.2. *Suppose that Assumption 2.3 and 2.4 hold. For $\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]$ defined in (40), it holds that*

$$\left\|\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]\right\|_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} \|\mathbb{E}_q[\xi_p(Z,\cdot)]\|_{\mathcal{G}^d}. \tag{41}$$

Proof. We have

$$\begin{aligned}\|\mathbb{E}_q[\xi_p(Z,\cdot)]\|_{\mathcal{G}^d}^2 &= S_p^2(q) \\ \left\|\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]\right\|_{\mathcal{G}^d}^2 &= \frac{1}{(t-1)^2} \sum_{i=1}^{t-1} \sum_{i=1}^{t-1} h_p(Z_i,Z_j) =: V_{t-1},\end{aligned}$$

where V_{t-1} is defined in Lemma B.1. From Lemma B.1 and the Borel-Cantelli lemma, it follows that:

$$\left\|\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]\right\|^2_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} \|\mathbb{E}_q[\xi_p(Z,\cdot)]\|^2_{\mathcal{G}^d}.$$

The result then follows from the continuous mapping theorem. □

Lemma B.3. *Suppose that Assumption 2.3 and 2.4 hold. Under H_1 in (1b), for the oracle (10) and plug-in (12) witness function, it holds that:*

$$\langle \hat{g}_t, g^* \rangle_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} 1. \tag{42}$$

As a consequence, $\|\hat{g}_t - g^*\|_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} 0$.

Proof. Assuming that the alternative (1b) is true, it follows that:

$$\|\mathbb{E}_q[\xi_p(Z,\cdot)]\|_{\mathcal{G}^d} > 0.$$

We aim to show that:

$$\left\langle \frac{\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]}{\left\|\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]\right\|_{\mathcal{G}^d}}, \frac{\mathbb{E}_q[\xi_p(Z,\cdot)]}{\|\mathbb{E}_q[\xi_p(Z,\cdot)]\|_{\mathcal{G}^d}} \right\rangle_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} 1.$$

From Lemma B.2, we know that $\left\|\mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)]\right\|_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} \|\mathbb{E}_q[\xi_p(Z,\cdot)]\|_{\mathcal{G}^d}$. Hence, it suffices to show that

$$\left\langle \mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)], \mathbb{E}_q[\xi_p(Z,\cdot)] \right\rangle_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} \|\mathbb{E}_q[\xi_p(Z,\cdot)]\|^2_{\mathcal{G}^d}. \tag{43}$$

We rewrite $\left\langle \mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)], \mathbb{E}_q[\xi_p(Z,\cdot)] \right\rangle_{\mathcal{G}^d}$ as:

$$\begin{aligned} &\left\langle \mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)], \mathbb{E}_{Z\sim q}[\xi_p(Z,\cdot)] \right\rangle_{\mathcal{G}^d} \\ =&\frac{1}{t-1}\sum_{i=1}^{t-1} \langle \xi_p(Z_i,\cdot), \mathbb{E}_{Z\sim q}[\xi_p(Z,\cdot)]\rangle_{\mathcal{G}^d} \\ =&\frac{1}{t-1}\sum_{i=1}^{t-1} \mathbb{E}_{Z\sim q}[h_p(Z_i,Z)]. \end{aligned}$$

Since we have shown in the proof of (B.1) that $|h_p(x,y)| \le 4$, by the SLLN, we obtain:

$$\frac{1}{t-1}\sum_{i=1}^{t-1} \mathbb{E}_{Z\sim q}[h_p(Z_i,Z)] \xrightarrow{\text{a.s.}} \mathbb{E}_{Z,Z'}[h_p(Z,Z')] = S_p^2(q).$$

Therefore, we deduce that:

$$\left\langle \mathbb{E}_{\hat{q}_{t-1}}[\xi_p(Z,\cdot)], \mathbb{E}_{Z\sim q}[\xi_p(Z,\cdot)] \right\rangle_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} \|\mathbb{E}_q[\xi_p(Z,\cdot)]\|^2_{\mathcal{G}^d}.$$

$\|\hat{g}_{t-1} - g^*\|_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} 0$ simply follows from the fact that

$$\|\hat{g}_t - g^*\|_{\mathcal{G}^d} = \sqrt{2(1 - \langle \hat{g}_t, g^* \rangle_{\mathcal{G}^d})}.$$

□

Lemma B.4. *Suppose that $\{x_t\}_{t\ge1}$ is a sequence of numbers such that $\lim_{t\to\infty} x_t = x$. Then the corresponding sequence of partial averages also converges to x, that is, $\lim_{n\to\infty} \frac{1}{n}\sum_{t=1}^n x_t = x$. That also implies that if $\{X_t\}_{t\ge1}$ is a sequence of random variables such that $X_t \xrightarrow{\text{a.s.}} X$, then $\frac{1}{n}\sum_{t=1}^n X_t \xrightarrow{\text{a.s.}} X$.*

Proof. Fix any $\varepsilon > 0$. Since $\{x_t\}_{t\geq 1}$ is converging, then $\exists M > 0$, such that:

$$|x_t - x| \leq M, \quad \forall t \geq 1.$$

Now, let n_0 be such that $|x_t - x| \leq \varepsilon/2$ for all $n > n_0$. Further, choose any $n_1 > n_0$: $Mn_0/n_1 \leq \varepsilon/2$. Hence, for any $m > n_1$, it holds that:

$$\begin{aligned}\left|\frac{1}{m}\sum_{t=1}^{m}(x_t - x)\right| &\leq \left|\frac{1}{m}\sum_{t=1}^{n_0}(x_t - x)\right| + \left|\frac{1}{m}\sum_{t=n_0+1}^{m}(x_t - x)\right| \\ &\leq \frac{1}{m}\sum_{t=1}^{n_0}|x_t - x| + \frac{1}{m}\sum_{t=n_0+1}^{m}|x_t - x| \\ &\leq \frac{n_0}{m}M + \frac{m-n_0}{m}\frac{\varepsilon}{2} \leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2},\end{aligned}$$

which implies the result. □

Before, we state the next result, recall that KSD-based payoffs are based on the predictable estimates $\{g_i\}_{i\geq 1}$ of the oracle witness function g^* and have the following form:

$$\begin{aligned} f_i(Z_i) &= \frac{1}{2}(\langle \hat{g}_i(Z_i), \nabla \log p(Z_i)\rangle + \nabla \cdot \hat{g}_i(Z_i)), \quad i \geq 1, \\ f^*(Z_i) &= \frac{1}{2}(\langle g^*(Z_i), \nabla \log p(Z_i)\rangle + \nabla \cdot g^*(Z_i)). \end{aligned} \tag{44}$$

Lemma B.5. *Suppose that Assumption 2.3 and 2.4 hold. Under H_1 in (1b), it holds that:*

$$\frac{1}{t}\sum_{i=1}^{t} f_i(Z_i) \xrightarrow{\text{a.s.}} \mathbb{E}_q[f^*(Z)], \tag{45}$$

$$\frac{1}{t}\sum_{i=1}^{t}(f_i(Z_i))^2 \xrightarrow{\text{a.s.}} \mathbb{E}_q[(f^*(Z))^2]. \tag{46}$$

Proof. First, we prove that (45). Note that:

$$\left|\frac{1}{t}\sum_{i=1}^{t} f_i(Z_i) - \mathbb{E}[f^*(Z)]\right| \leq \left|\frac{1}{t}f_i(Z_i) - \frac{1}{t}\sum_{i=1}^{t} f^*(Z_i)\right| + \underbrace{\left|\frac{1}{t}\sum_{i=1}^{t} f^*(Z_i) - \mathbb{E}[f^*(Z)]\right|}_{\xrightarrow{\text{a.s.}} 0},$$

where the second term converges almost surely to 0 by the SLLN. For the first term, we have that:

$$\left|\frac{1}{t}\sum_{i=1}^{t} f^*(Z_i) - \mathbb{E}[f^*(Z)]\right| \leq \frac{1}{t}\sum_{i=1}^{t}|f_i(Z_i) - f^*(Z_i)|.$$

Finally, note that

$$\begin{aligned} |f_i(Z_i) - f^*(Z_i)| &= \frac{1}{2}|\langle \hat{g}_i - g^*, \xi_p(Z_i)\rangle| \\ &\leq \frac{1}{2}\|\hat{g}_i - g^*\|_{\mathcal{G}^d} \cdot \|\xi_p(Z_i, \cdot)\|_{\mathcal{G}^d} \\ &\leq \|\hat{g}_i - g^*\|_{\mathcal{G}^d} \xrightarrow{\text{a.s.}} 0. \end{aligned} \tag{47}$$

where the last inequality follows from the fact that $\|\xi_p(Z_i, \cdot)\|_{\mathcal{G}^d} \leq 2$ and the convergence result is due to Lemma B.3. The result (45) then follows after invoking Lemma B.4.

Next, we show that (46) holds. Note that:

$$\begin{aligned}\frac{1}{t}\sum_{i=1}^{t}\left(f_i(Z_i)\right)^2 &= \frac{1}{2}\sum_{i=1}^{t}\left(f_i(Z_i) - f^*(Z_i) + f^*(Z_i)\right)^2 \\ &= \underbrace{\frac{1}{t}\sum_{i=1}^{t}\left(f_i(Z_i) - f^*(Z_i)\right)^2}_{\xrightarrow{\text{a.s.}} 0} \\ &+ \frac{2}{t}\sum_{i=1}^{t} f^*(Z_i)\left(f_i(Z_i) - f^*(Z_i)\right) \\ &+ \underbrace{\frac{1}{t}\sum_{i=1}^{t}\left(f^*(Z_i)\right)^2}_{\xrightarrow{\text{a.s.}} \mathbb{E}[(f^*(Z))]},\end{aligned}$$

where the first convergence result is due to (47) and Lemma B.4 and the second converge result is due to the SLLN. Using (47), Lemma B.4, and the fact that f^* is bounded by 1, we deduce that:

$$\left|\frac{2}{t}\sum_{i=1}^{t} f^*(Z_i)\left(f_i(Z_i) - f^*(Z_i)\right)\right| \leq 2\cdot\frac{1}{t}\sum_{i=1}^{t}\left|f_i(Z_i) - f^*(Z_i)\right| \xrightarrow{\text{a.s.}} 0,$$

and hence we conclude that the convergence (46) holds. □

B.4. Auxiliary Results

Theorem B.6 (Ville's inequality (Ville, 1939))**.** *Suppose that $\{\mathcal{M}_t\}_{t\geq 0}$ is a non-negative supermartingale process adapted to a filtration $\{\mathcal{F}_t : t \geq 0\}$. Then, for any $a > 0$ it holds that:*

$$\mathbb{P}(\exists t \geq 1 : \mathcal{M}_t \geq a) \leq \frac{\mathbb{E}[\mathcal{M}_0]}{a}.$$