

Anytime-valid off-policy inference for contextual bandits

IAN WAUDBY-SMITH, Dept. of Statistics & Data Science, Carnegie Mellon University, USA

LILI WU, Microsoft, USA

AADITYA RAMDAS, Dept. of Statistics & Data Science, Carnegie Mellon University, USA

NIKOS KARAMPATZIAKIS, Microsoft, USA

PAUL MINEIRO, Microsoft, USA

Contextual bandit algorithms are ubiquitous tools for active sequential experimentation in healthcare and the tech industry. They involve online learning algorithms that adaptively learn policies over time to map observed contexts X_t to actions A_t in an attempt to maximize stochastic rewards R_t . This adaptivity raises interesting but hard statistical inference questions, especially counterfactual ones: for example, it is often of interest to estimate the properties of a hypothetical policy that is different from the logging policy that was used to collect the data — a problem known as “off-policy evaluation” (OPE). Using modern martingale techniques, we present a comprehensive framework for OPE inference that relax unnecessary conditions made in some past works (such as performing inference at prespecified sample sizes, uniformly bounded importance weights, constant logging policies, constant policy values, among others), significantly improving on them both theoretically and empirically. Importantly, our methods can be employed while the original experiment is still running (that is, not necessarily post-hoc), when the logging policy may be itself changing (due to learning), and even if the context distributions are a highly dependent time-series (such as if they are drifting over time). More concretely, we derive confidence sequences for various functionals of interest in OPE. These include doubly robust ones for time-varying off-policy mean reward values, but also confidence bands for the entire cumulative distribution function of the off-policy reward distribution. All of our methods (a) are valid at arbitrary stopping times (b) only make nonparametric assumptions, (c) do not require importance weights to be uniformly bounded and if they are, we do not need to know these bounds, and (d) adapt to the empirical variance of our estimators. In summary, our methods enable anytime-valid off-policy inference using adaptively collected contextual bandit data.

1 INTRODUCTION

The so-called “contextual bandit” problem is an abstraction that can be used to describe several problem setups in statistics and machine learning [35, 36]. For example, it generalizes the multi-armed bandit problem by allowing for “contextual” side information, and it can be used to describe many adaptive sequential experiments. The general contextual bandit problem can be described informally as follows: an agent (such as a medical scientist in a clinical trial) views contextual information $X_t \in \mathcal{X}$ for subject t (such as the clinical patient’s demographics, medical history, etc.), takes an action $A_t \in \mathcal{A}$ (such as whether to administer a placebo, a low dose, or a high dose), and observes some reward R_t (such as whether their adverse symptoms have subsided). This description is made formal in the protocol for the generation of contextual bandit data in Algorithm 1. In the present paper, no restrictions are placed on the dimensionality or structure of the context and action spaces \mathcal{X} and \mathcal{A} beyond them being measurable, but it is often helpful to think about \mathcal{X} as a d -dimensional Euclidean space, and \mathcal{A} as $\{0, 1\}$ for binary treatments, or \mathbb{R} for different dosages, and so on. Indeed, while high-dimensional settings often pose certain challenges in contextual bandits (such as computational ones, or inflated variances), none of these issues will affect the *validity* of our statistical inference methods. Throughout, we will require that the rewards are real-valued and bounded in $[0, 1]$ — a common assumption in contextual bandits [28, 50] — except for Section 4 where we relax the boundedness constraint.

There are two main objectives that one can study in the contextual bandit setup: (1) policy optimization, and (2) off-policy evaluation (OPE) [11, 12, 35, 36]. Here, a “policy” $\pi(a | x)$ is simply

a conditional distribution over actions, such as the probability that patient t should receive various treatments given their context X_t . Policy *optimization* is concerned with finding policies that achieve high cumulative rewards (typically measured through regret), while off-policy *evaluation* is concerned with asking the counterfactual question: “how would we have done if we used some policy π instead of the policy that is currently collecting data?”. In this paper, we study the latter with a particular focus on statistical inference in adaptive, sequential environments under nonparametric assumptions.

1.1 Off-policy inference, confidence intervals, and confidence sequences

By far the most common parameter of interest in the OPE problem is the expected reward $v := \mathbb{E}_\pi(R)$ that would result from taking an action from the policy π . This expectation v is called the “value” of the policy π . While several estimators for v have been derived and refined over the years, many practical problems call for more than just a point estimate: we may also wish to quantify the uncertainty surrounding our estimates via statistical inference tools such as confidence intervals (CI). However, a major drawback of CIs is the fact that they are only valid at *fixed and prespecified* sample sizes, while contextual bandit data are collected in a sequential and adaptive fashion over time.

We lay out the assumed protocol for the generation of contextual bandit data in Algorithm 1, and in particular, all of our results will assume access to the output of this algorithm, namely the (potentially infinite) sequence of tuples $(X_t, A_t, R_t)_{t=1}^T$ for $T \in \mathbb{N} \cup \{\infty\}$. As is standard in OPE, we will always assume that the policy π is (almost surely) absolutely continuous with respect to h_t so that π/h_t is almost surely finite (without which, estimation and inference are not possible in general). Indeed, this permits many bandit techniques and in principle allows for Thompson sampling since it always assigns positive probability to an action (note that it may not always be easy to compute the probability of taking that action via Thompson sampling, but if those probabilities can be computed, they can be used directly within our framework). However, $(h_t)_{t=1}^\infty$ cannot be the result of Upper Confidence Bound (UCB)-style algorithms since they take conditionally deterministic actions given the past, violating the absolute continuity of π with respect to h_t .

In Algorithm 1, the term “exogenously time-varying” simply means that the context and reward distributions at time t can only depend on the past through $X_1^{t-1} \equiv (X_1, \dots, X_{t-1})$, and not on the actions taken (or rewards received). Formally, we allow for any joint distribution over $(X_t, A_t, R_t)_{t=1}^\infty$ as long as

$$p_{R_t}(r \mid x, a, \mathcal{H}_{t-1}) = p_{R_t}(r \mid x, a, X_1^{t-1}) \text{ and } p_{X_t}(x \mid \mathcal{H}_{t-1}) = p_{X_t}(x \mid X_1^{t-1}), \quad (1)$$

where \mathcal{H}_t is all of the history $\sigma((X_i, A_i, R_i)_{i=1}^t)$ up until time t . This conditional independence requirement (1) includes as a special case more classical setups where X_t is independent of all else given A_t , such as those considered in Bibaut et al. [4] or iid scenarios [28], but is strictly more general, since, for example, $(X_t)_{t=1}^\infty$ can be a highly dependent time-series. However, we do not go as far as to consider the adversarial setting that is sometimes studied in the context of regret minimization. We impose this conditional independence requirement since otherwise, the interpretation of $\mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1})$ changes depending on which sequence of actions were played by the logging policy. Making matters more concrete, the conditional off-policy value $\mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1})$ at time t is given by

$$v_t := \mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1}) \equiv \int_{X \times \mathcal{A} \times \mathbb{R}} r \cdot p_{R_t}(r \mid a, x, \mathcal{H}_{t-1}) \pi(a \mid x) p_{X_t}(x \mid \mathcal{H}_{t-1}) dx da dr \quad (2)$$

$$= \int_{X \times \mathcal{A} \times \mathbb{R}} r \cdot p_{R_t}(r \mid a, x, X_1^{t-1}) \pi(a \mid x) p_{X_t}(x \mid X_1^{t-1}) dx da dr, \quad (3)$$

and the equality (3) follows from (1). Notice that (2) could in principle depend on the logging policies and actions played, but (3) does not. Despite imposing the conditional independence assumption (1), the integral (2) is still a perfectly well-defined functional, and if (1) is not satisfied, then our CSs will still cover a quantity in terms of this functional. However, its interpretation would no longer be counterfactual with respect to the entire sequence of actions (only conditional on the past).

While most prior work on OPE in contextual bandits is not written *causally* in terms of potential outcomes (e.g. [4, 5, 25, 28, 50]), it is nevertheless possible to write down a causal target v_t^* (i.e. a functional of the potential outcome distribution) and show that it is equal to v_t under certain causal identification assumptions. These assumptions resemble the familiar *consistency*, *exchangeability*, and *positivity* conditions that are ubiquitous in the treatment effect estimation literature. Moreover, there is a close relationship between OPE and the estimation of so-called *stochastic interventions* in causal inference; indeed, they can essentially be seen as equivalent but with slightly different emphases and setups. However, given that neither the potential outcomes view nor the stochastic intervention interpretation of OPE are typically emphasized in the contextual bandit literature (with the exception of Zhan et al. [61], who use potential outcomes throughout), we leave this discussion to Appendix B.

Algorithm 1 Protocol for the generation of contextual bandit data

```

// Here,  $T \in \mathbb{N} \cup \{\infty\}$ .
for  $t = 1, 2, \dots, T$  do
  // The agent selects a policy  $h_t$  based on the history  $\mathcal{H}_{t-1} \equiv \sigma((X_i, A_i, R_i)_{i=1}^{t-1})$ .
   $h_t \in \mathcal{H}_{t-1}$ .

  // The environment draws a context from an (exogenously time-varying) distribution.
   $X_t \sim p_{X_t}(\cdot)$ 

  // The agent plays a random action drawn from the selected policy.
   $A_t \sim h_t(\cdot \mid X_t)$ .

  // The environment draws a reward from an (exogenously time-varying) distribution based on the
  // action and context.
   $R_t \sim p_{R_t}(\cdot \mid A_t, X_t)$ .
end for

// Return a (potentially infinite) sequence of contextual bandit data.
return  $(X_t, A_t, R_t)_{t=1}^T$ 

```

To illustrate the shortcomings of CIs for OPE, suppose we run a contextual bandit algorithm and want to see whether π is better than the current state-of-the-art policy h – e.g. whether $\mathbb{E}_\pi(R) > \mathbb{E}_h(R)$. (Here, we are implicitly assuming that $\mathbb{E}_{\pi'}(R) = \mathbb{E}_{\pi'}(R_t \mid \mathcal{H}_{t-1})$ for any policy π' for the sake of illustration.) Suppose we compute a CI for the value of π based on n samples (for some prespecified n), and while π seems promising, the CI turns out to be inconclusive (the CI for $\mathbb{E}_\pi(R)$ includes $\mathbb{E}_h(R)$ if the latter is known, or the two CIs overlap if the latter is unknown). It is tempting to collect more data, for a total of n' points, to see if the result is now conclusive; however the resulting sample size n' is now a *data-dependent* quantity, rendering the CI invalid. (This could happen more than once.)

Fortunately, there exist statistical tools that permit adaptive stopping in these types of sequential data collection problems: *confidence sequences* (CSs [9, 34]). A CS is a sequence of confidence intervals, valid at all sample sizes uniformly (and hence at arbitrary stopping times). Importantly for the aforementioned OPE setup, CSs allow practitioners to collect additional data and continuously monitor it, so that the resulting CI is indeed valid at the data-dependent stopped sample size n' . More formally, we say that a sequence of intervals $[L_t, U_t]_{t=1}^\infty$ is a CS for the parameter $\theta \in \mathbb{R}$ if

$$\mathbb{P}(\forall t \in \mathbb{N}, \theta \in [L_t, U_t]) \geq 1 - \alpha, \text{ or equivalently, } \mathbb{P}(\exists t \in \mathbb{N} : \theta \notin [L_t, U_t]) \leq \alpha. \quad (4)$$

Contrast (4) above with the definition of a CI which states that $\forall n \in \mathbb{N}, \mathbb{P}(\theta \in [L_n, U_n]) \geq 1 - \alpha$, so that the “ $\forall n$ ” is outside the probability $\mathbb{P}(\cdot)$ rather than inside. A powerful consequence of (4) is that $[L_\tau, U_\tau]$ is a valid CI for *any stopping time* τ . In fact, $[L_\tau, U_\tau]$ being a valid CI is not just an implication of (4) but the two statements are actually equivalent; see Howard et al. [24, Lemma 3].

The consequence for the OPE practitioner is that they can continuously update and monitor a CS for the value of π while the contextual bandit algorithm is running, and deploy π as soon as they are confident that it is better than the current state-of-the-art h . Karampatziakis et al. [28] refer to this adaptive policy switching as “gated deployment”, and we will return to this motivating example through the paper. Let us now lay out five desiderata that we want all of our CSs to satisfy.

1.2 Desiderata for anytime-valid off-policy inference

Throughout this paper, we will derive methods for off-policy evaluation and inference in a variety of settings — including fixed policies (Section 2), time-varying policies (Section 3), and for entire cumulative distribution functions (Section 4). However, what all of these approaches will have in common is that they will satisfy five desirable properties which we enumerate here.

- (1) **Nonasymptotic:** Our confidence sets will satisfy *exact* coverage guarantees for *any* sample size, unlike procedures based on the central limit theorem which only satisfy *approximate* guarantees for large samples.¹
- (2) **Nonparametric:** We will not make any parametric assumptions on the distribution of the contexts, policies, or rewards.
- (3) **Time-uniform / anytime-valid:** Our confidence sets will be *uniformly* valid for all sample sizes, and permit off-policy inference at arbitrary data-dependent stopping times.
- (4) **Adaptive data collection (via online learning):** All of our off-policy inference methods will allow for the sequence of logging policies $(h_t)_{t=1}^\infty$ to be predictable (i.e. h_t can depend on \mathcal{H}_{t-1}). In particular $(h_t)_{t=1}^\infty$ can be the result of an online learning algorithm.
- (5) **Unknown and unbounded w_{\max} :** In all of our algorithms, the maximal importance weight

$$w_{\max} := \operatorname{ess\,sup}_{t \in \mathbb{N}, a \in \mathcal{A}, x \in \mathcal{X}} \frac{\pi(a | x)}{h_t(a | x)}$$

can be unknown, and need not be uniformly bounded (i.e. w_{\max} can be infinite). Note that we do require that importance weights $\pi(a | x)/h_t(a | x)$ themselves are finite for each (t, a, x) , but their essential *supremum* need not be. Perhaps surprisingly, even if w_{\max} is infinite, it is still possible for our CSs to shrink to zero-width since they depend only on *empirical variances*. As an illustrative example, see Proposition 3 for a closed-form CS whose width can shrink at a rate of $\sqrt{\log \log t/t}$ as long as the importance-weighted rewards are well-behaved (e.g. in the iid setting, if they have finite second moments).

¹Note that nonasymptotic procedures may be more conservative than asymptotic ones as they satisfy more rigorous coverage (similarly, type-I error) guarantees. Whether one should sacrifice stronger guarantees for tightness comes down to philosophical preference. For the purposes of this paper, we focus solely on nonasymptotics.

In addition to the above, we will design procedures so that they have strong empirical performance and are straightforward to compute. While some of these desiderata are quite intuitive and common in statistical inference (such as nonasymptotic, nonparametric, and time-uniform validity, so as to avoid relying on large sample theory, unrealistic parametric assumptions, or prespecified sample sizes), desiderata 4 and 5 are more specific to OPE and have not been satisfied in several prior works as we outline in Sections 2, 3, and 4. Given their central importance to our work, let us elaborate on them here.

Why allow for logging policies to be predictable? For the purpose of policy optimization, contextual bandit algorithms are tasked with balancing exploration and exploitation: simultaneously figuring out which policies will yield high rewards (at the expense of trying out suboptimal policies) and playing actions from the policies that have proven effective so far. On the other hand, in adaptive sequential trials, an experimenter might aim to balance context distributions between treatment arms (such as via Efron’s biased coin design [14]) or to adaptively find the treatment policy that yields the most efficient treatment effect estimators [29]. In both cases, the logging policies $(h_t)_{t=1}^\infty$ are not only changing over time, but adaptively so based on previous observations. We strive to design procedures that permit inference in precisely these types of adaptive data collection regimes, despite most prior works on off-policy inference for contextual bandits having assumed that there is a fixed, prespecified logging policy [5, 25, 28, 50]. Of course, if a CS or CI is valid under adaptive data collection, they are also valid when fixed logging policies are used instead.

Why not rely on knowledge of w_{\max} ? Related to the previous discussion, it may not be known *a priori* how the range of the predictable logging policies will evolve over time. Moreover, one can imagine a situation where $\sup_{a,x} \pi(a | x)/h_t(a | x) \rightarrow \infty$, even if every individual importance weight π/h_t is finite. In such cases, having CSs be agnostic to the value of w_{\max} is essential. However, even if w_{\max} is known, it may be preferable to design CSs that do not depend on this worst-case value. Suppose for the sake of illustration that a logging policy h assigns a novel treatment (denoted by $a = 1$) with probability 1/5 and a placebo (denoted by $a = 0$) with probability 4/5 for most subjects, except for a small but high-risk subpopulation, who receive the novel treatment with probability 1/1000. To estimate the expected reward of the novel treatment, note that the importance weight for subject t will take on the value $w_t := 1/h(A_t | X_t)$ for treatment $A_t \in \{0, 1\}$ and context $X_t \in \mathcal{X}$. Despite the fact that most of the importance weights are only 5, and hence most importance-weighted pseudo-outcomes $w_t R_t$ will take values in $[0, 5]$, the worst-case w_{\max} is much larger at 1000. Consequently, we should expect a CS that scales with $w_{\max} = 1000$ to be much wider than one that only scales with a quantity like an empirical variance. For these reasons we prefer procedures that depend on an empirical variance term (defined later) rather than the worst-case importance weight w_{\max} .

1.3 Outline and contributions

Our fundamental contribution is in the derivation of CSs for various off-policy parameters, including fixed policy values, time-varying policy values, and quantiles of the off-policy reward distribution. We begin in Section 2 with the most common formulation of the OPE problem: estimating the value v of a target policy π . Theorem 1 presents time-uniform CSs for v , a result that generalizes and improves upon the current state-of-the-art CSs for v by Karampatziakis et al. [28]. In Section 3, we consider the more challenging problem of estimating a time-varying average policy value v_t , where the distribution of off-policy rewards can change over time in an arbitrary and unknown fashion. In Section 4, we derive CSs for quantiles of the off-policy reward distribution, and in particular, Theorem 3 presents a confidence band for the entire cumulative distribution function (CDF) that is both uniformly valid in time *and* in the quantiles. For the results of Sections 3 and 4, no other

solutions to this problem exist in the literature, to the best of our knowledge. Finally, in Section 5, we summarize our results and describe some natural extensions and implications of them, namely false discovery rate control under arbitrary dependence when evaluating several policies, and differentially private off-policy inference.

1.4 Related work

Throughout the paper, we will draw detailed comparisons to work that is most closely related to ours, i.e. papers that are broadly concerned with estimating policy values and/or the off-policy CDF from contextual bandit data in a model-free setting — here we are using the term “model-free” to mean that no restrictions are placed on the functional form between the rewards R and the actions A nor on the covariates X . Specifically, Table 1 and the preceding text provides a (selective) property-by-property comparison to the directly related works of Bibaut et al. [4], Karampatziakis et al. [28], Zhan et al. [61], and Howard et al. [24], and Table 2 provides a similar comparison to the works of Howard and Ramdas [22], Chandak et al. [5], and Huang et al. [25]. However, there are several other works that focus on *fixed- n* (i.e. not time-uniform) and *asymptotic* statistical inference from adaptively collected data (e.g. in the form of multi-armed bandits, contextual bandits, or more general reinforcement learning). For example, Ramprasad et al. [41] develop a bootstrap procedure for estimating policy values under Markov noise with temporal difference learning algorithms, Dimakopoulou et al. [10] perform adaptive inference in the multi-armed bandit setting, Hadad et al. [18] provide asymptotic confidence intervals for treatment effects in adaptive experiments, and Zhang et al. [63] provide distribution-uniform asymptotic procedures for M-estimation from contextual bandit data. Other works that consider more model-based approaches include Zhang et al. [62], Khamaru et al. [33], Shen et al. [48], and Chen et al. [6].

1.5 Notation: supermartingales, filtrations, and stopping times

Since all of our results will rely on the analysis of nonnegative (super)martingales, predictable processes, stopping times, and so on, it is worth defining some of these terms before proceeding. Consider a universe of distributions Π on a filtered probability space (Ω, \mathcal{F}) . A single draw from any distribution $P \in \Pi$ results in a sequence Z_1, Z_2, \dots of potentially dependent observations. (In the context of this paper, Z_t may represent (X_t, A_t, R_t) , for example, and the distribution P may be induced by the policy, and not specified in advance.) If Z_1, Z_2, \dots are independent and identically distributed (iid), we will explicitly say so, but in general we eschew iid assumptions in this paper.

As is common in the statistics literature, we will use upper-case letters like Z to refer to random variables and lower-case letters z to refer to non-stochastic values in the same space that Z takes values. Let Z_1^t denote the tuple (Z_1, \dots, Z_t) and let $\mathcal{H} \equiv (\mathcal{H}_t)_{t=1}^\infty$ by default represent the data (or “canonical”) filtration, meaning that $\mathcal{H}_t = \sigma(Z_1^t)$.

A sequence of random variables $Y \equiv (Y_t)_{t=1}^\infty$ is called a *process* if it is adapted to \mathcal{H} , that is if Y_t is measurable with respect to \mathcal{H}_t for every t . A process Y is *predictable* if Y_t is measurable with respect to \mathcal{H}_{t-1} — informally “ Y_t only depends on the past”. A process M is a *martingale* for P with respect to \mathcal{H} if

$$\mathbb{E}_P[M_t \mid \mathcal{H}_{t-1}] = M_{t-1} \quad (5)$$

for all $t \geq 1$. M is a *supermartingale* for P if it satisfies (5) with “=” relaxed to “ \leq ”. A (super)martingale is called a *test (super)martingale* if it is nonnegative and $M_0 = 1$. A process M is called a test (super)martingale for $\mathcal{P} \subset \Pi$ if it is a test (super)martingale for every $P \in \mathcal{P}$.

Throughout, if an expectation \mathbb{E} operator is used without a subscript P , or if a **boldface** \mathbb{P} is used to denote a probability, these are always referring to the distribution of $(X_t, A_t, R_t)_{t \geq 1}$ induced by Algorithm 1 and the logging policies $(h_t)_{t=1}^\infty$.

An \mathcal{H} -stopping time τ is a $\mathbb{N} \cup \{\infty\}$ -valued random variable such that $\{\tau \leq t\} \in \mathcal{H}_t$ for each $t \geq 0$. Informally and in the context of this paper, a stopping time can be thought of as a sample size that was chosen based on all of the information \mathcal{H}_t up until time t .

2 WARMUP: OFF-POLICY INFERENCE FOR CONSTANT POLICY VALUES

This section deals with the case where v_t from (2) does not depend on t , meaning that it is constant as a function of time. We handle the time-varying case in the next section.

We begin by extending a result of Karampatziakis et al. [28, Section 5.2] which applied in the iid setting, meaning that the logging policy h is fixed and the contexts and rewards are assumed to be iid. Their paper derives several CSs for the value $v := \mathbb{E}_{A \sim \pi}(R)$ of the policy π for $[0, 1]$ -bounded rewards R , but some of their CSs require knowledge of w_{\max} , which we would like to avoid as per our desiderata in Section 1.2. However, their so-called ‘‘scalar betting’’ approach in [28, Section 5.2] makes use of importance-weighted random variables and does not depend on knowing w_{\max} . To elaborate, let w_t be the importance weight for the target policy π versus the logging policy h given by

$$w_t := \frac{\pi(A_t | X_t)}{h(A_t | X_t)}, \quad (6)$$

and let $\phi_t^{(IW-\ell)} := w_t R_t$ and $\phi_t^{(IW-u)} := w_t(1 - R_t)$ be importance-weighted rewards that will be used to construct lower and upper bounds respectively. Note that $\phi_t^{(IW-\ell)}$ is ubiquitous in the bandit and causal inference literatures, and the authors were not concerned with deriving *new estimators*, but rather *new confidence sequences* using existing estimators. While $w_t \equiv w_t(X_t, A_t)$ does depend on both A_t and X_t , we leave the dependence on them implicit going forward to reduce notational clutter.

Proposition 1 (Scalar betting off-policy CS [28]). *Suppose $(X_t, A_t, R_t)_{t=1}^\infty$ are iid with $[0, 1]$ -valued rewards $(R_t)_{t=1}^\infty$, and the logging policy h is fixed. For each $v' \in [0, 1]$, let $(\lambda_t^L(v'))_{t=1}^\infty$ be any $[0, 1/v']$ -valued predictable sequence. Then,*

$$L_t^{IW} := \inf \left\{ v' \in [0, 1] : \prod_{i=1}^t \left(1 + \lambda_i^L(v') \cdot (\phi_i^{(IW-\ell)} - v') \right) < \frac{1}{\alpha} \right\} \quad (7)$$

forms a lower $(1 - \alpha)$ -CS for v , meaning $\mathbb{P}(\forall t \in \mathbb{N}, v \geq L_t^{IW}) \geq 1 - \alpha$. Similarly, for any $[0, 1/(1 - v')]$ -valued predictable sequence $(\lambda_t^U(v'))_{t=1}^\infty$.

$$U_t^{IW} := 1 - \inf \left\{ 1 - v' \in [0, 1] : \prod_{i=1}^t \left[1 + \lambda_i^U(v') \cdot (\phi_i^{(IW-u)} - (1 - v')) \right] < \frac{1}{\alpha} \right\} \quad (8)$$

forms an upper $(1 - \alpha)$ -CS for v , meaning $\mathbb{P}(\forall t \in \mathbb{N}, v \leq U_t^{IW}) \geq 1 - \alpha$. A two-sided CS can be formed using $[L_t^{IW}, U_t^{IW}]_{t=1}^\infty$ combined with a union bound.

The above CS $[L_t^{IW}, U_t^{IW}]_{t=1}^\infty$ due to Karampatziakis et al. [28] has a number of desirable properties. Namely, it satisfies the first four of five desiderata in Section 1.2, meaning it is a nonasymptotic, nonparametric, time-uniform confidence sequence that does not require knowledge of w_{\max} . Note that while infima appear in the definitions of L_t^{IW} and U_t^{IW} they are straightforward to compute (e.g. via line or grid search) when the product is quasiconvex in $v' \in [0, 1]$ which is often the case as we discuss in Section 2.2. The idea behind Proposition 1 is to show that the product inside the above infima are nonnegative martingales when $v' = v$ and then apply Ville’s inequality to it [53]. Our main results in the coming sections use similar techniques albeit with very different (super)martingales tailored to different problem settings.

We also wish to highlight that indeed, $[L_t^{IW}, U_t^{IW}]_{t=1}^\infty$ forms a valid $(1-\alpha)$ -CS *regardless* of how the sequences $(\lambda_t^L(v'))_t$ and $(\lambda_t^U(v'))_t$ are chosen. Such phenomena are common in martingale-based statistical procedures such as in Waudby-Smith and Ramdas [58] (see also the review paper of Ramdas et al. [40]) and will be seen in several of the results to follow. We will discuss some guiding principles for how to choose these sequences in Remark 1.

2.1 Tighter confidence sequences via doubly robust pseudo-outcomes

Here, we generalize and improve upon Proposition 1 in three ways. First, we show that the logging policy h can be replaced by a *sequence* of predictable logging policies $(h_t)_{t=1}^\infty$ so that h_t can be built from the entire history \mathcal{H}_{t-1} up until time $t-1$ (and in particular, $(h_t)_{t=1}^\infty$ can be the result of an online learning algorithm), so that the importance weight w_t at time t is given by

$$w_t := \frac{\pi(A_t | X_t)}{h_t(A_t | X_t)}. \quad (9)$$

Second, we show how the importance-weighted pseudo-outcomes $(\phi_t)_{t=1}^\infty$ can be made doubly robust in the sense of Dudík et al. [11, 12]. Indeed, define the lower and upper doubly robust pseudo-outcomes $(\phi_t^{(\text{DR-}\ell)})_{t=1}^\infty$ and $(\phi_t^{(\text{DR-u})})_{t=1}^\infty$ given by

$$\phi_t^{(\text{DR-}\ell)} := w_t \cdot \left(R_t - \left[\widehat{r}_t(X_t; A_t) \wedge \frac{k_t}{w_t} \right] \right) + \mathbb{E}_{a \sim \pi(\cdot | X_t)} \left(\widehat{r}_t(X_t; a) \wedge \frac{k_t}{w_t} \right), \quad (10)$$

$$\phi_t^{(\text{DR-u})} := w_t \cdot \left(1 - R_t - \left[(1 - \widehat{r}_t(X_t; A_t)) \wedge \frac{k_t}{w_t} \right] \right) + \mathbb{E}_{a \sim \pi(\cdot | X_t)} \left([1 - \widehat{r}_t(X_t; a)] \wedge \frac{k_t}{w_t} \right), \quad (11)$$

where $\widehat{r}_t(X_t; A_t)$ is any $[0, 1]$ -valued predictor of R_t built from \mathcal{H}_{t-1} and k_t is a $\mathbb{R}_{\geq 0} \cup \{\infty\}$ -valued tuning parameter built from \mathcal{H}_{t-1} that determines how “doubly robust” ϕ_t^{DR} should be. Note that $\phi_t^{(\text{DR-}\ell)}$ and $\phi_t^{(\text{DR-u})}$ are both at least $-k_t$ by construction, and have conditional means of v and $1-v$ respectively. Note that the phrase “doubly robust” is sometimes used to refer to *properties* of estimators (e.g. that their bias is second order and depends only on products of nuisance errors in observational studies where importance weights are unknown [32]) and sometimes to refer to *types* of estimators that enjoy variance-reduction without compromising validity in experiments where importance weights are known. We are using this phrase in the second sense following the conventions of Dudík et al. [11, 12].

Similar to the discussion surrounding Proposition 1, the doubly robust pseudo-outcomes in (10) are ubiquitous in the causal inference and bandit literatures [11, 12, 44, 51] with the minor tweak that we are truncating the reward predictor. Note that we are *not* doing this for the purposes of deriving better estimators — instead, we are doing so for the purposes of sharp concentration of measure in the pursuit of tighter CSs.

Setting $k_1 = k_2 = \dots = 0$ recovers the IW outcomes exactly, while setting $k_1 = k_2 = \dots = \infty$ recovers the classic doubly robust outcomes [11, 12] (this could also be achieved by setting $k_1 = k_2 = \dots = w_{\max}$, provided w_{\max} is finite and known). We discuss the need to truncate \widehat{r}_t in Remark 2, but the motivation to include a reward predictor at all is to reduce the variance of $\phi_t^{(\text{DR-}\ell)}$ and $\phi_t^{(\text{DR-u})}$ if R_t can be well-predicted by \widehat{r}_t , a well-known phenomenon in doubly robust estimation [7, 44, 52]. Indeed, we find that the resulting CSs are able to adapt to this reduced variance accordingly for large t (see Figure 1 for an illustration).

Third and finally, we relax the iid assumption, and only require that $\mathbb{E}_\pi(R_t | \mathcal{H}_{t-1}) = v \equiv \mathbb{E}_\pi(R_t)$ and $R_t \in [0, 1]$ almost surely. This relaxation of assumptions can be obtained for free, without any change to the resulting CSs whatsoever, and with only a slight modification to the proof. We summarize our extensions in the following theorem.

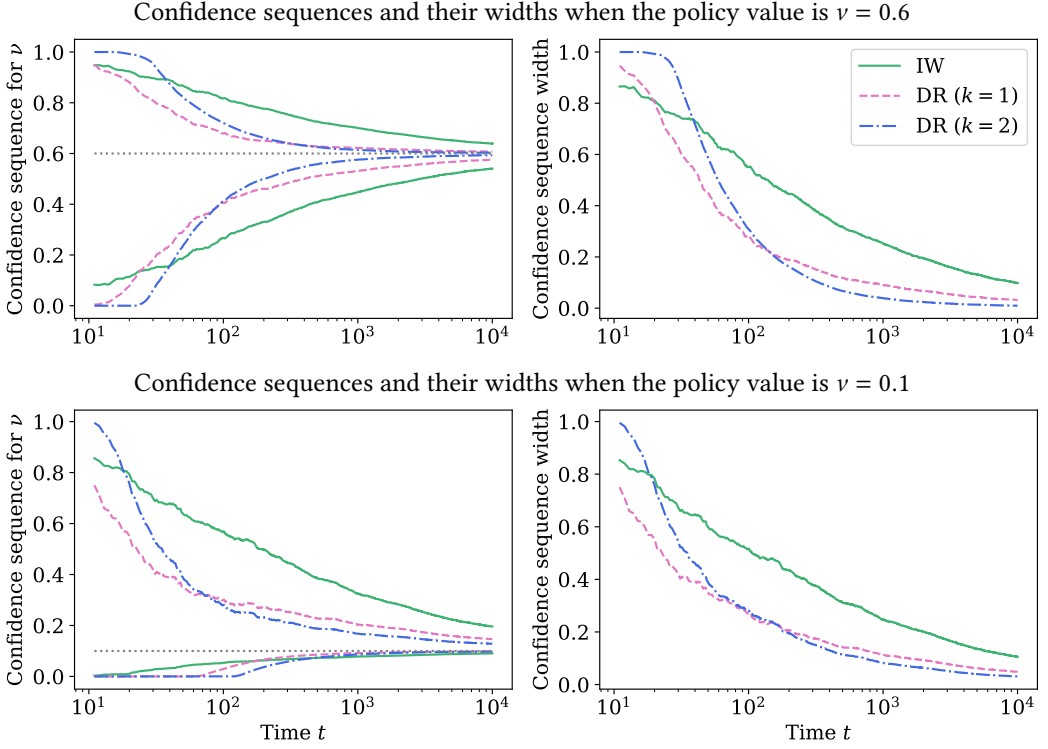


Fig. 1. Three confidence sequences for a policies with values $\nu = 0.6$ and $\nu = 0.1$. The first CS is built from importance-weighted pseudo-outcomes (“IW”), and the other two are built from doubly robust pseudo-outcomes (“DR”) with k taking values 1 and 2, respectively. In these examples, the reward R_t can be predicted easily, a property that only the doubly robust CSs can exploit. Notice that a larger value of k allows the doubly robust CS to become narrower for large t , but it pays for this adaptivity with wider bounds at small t . Nevertheless, all three CSs are time-uniform, and nonasymptotically valid in both simulation scenarios.

THEOREM 1 (DOUBLY ROBUST BETTING OFF-POLICY CS). *Suppose $(X_t, A_t, R_t)_{t=1}^\infty$ is an infinite sequence of contextual bandit data generated by the predictable policies $(h_t)_{t=1}^\infty$ whose $[0, 1]$ -valued reward R_t at time t has conditional mean $\mathbb{E}_\pi(R_t | \mathcal{H}_{t-1}) = \nu \equiv \mathbb{E}_\pi(R_t)$. For any predictable sequence $(\lambda_t^L(\nu'))_{t=1}^\infty$ such that $\lambda_t^L(\nu') \in [0, (\nu' + k_t)^{-1}]$, we have*

$$L_t^{\text{DR}} := \inf \left\{ \nu' \in [0, 1] : \prod_{i=1}^t \left[1 + \lambda_i^L(\nu') \cdot (\phi_i^{(\text{DR-}l)} - \nu') \right] < \frac{1}{\alpha} \right\} \quad (12)$$

forms a lower $(1 - \alpha)$ -CS for ν . Similarly, if $\lambda_t^U(\nu') \in [0, (1 - \nu' + k_t)^{-1}]$ is predictable, then

$$U_t^{\text{DR}} := 1 - \inf \left\{ 1 - \nu' \in [0, 1] : \prod_{i=1}^t \left[1 + \lambda_i^U(\nu') \cdot (\phi_i^{(\text{DR-u})} - (1 - \nu')) \right] < \frac{1}{\alpha} \right\} \quad (13)$$

forms an upper $(1 - \alpha)$ -CS for ν .

The proof of Theorem 1 can be found in Appendix A.2 and relies on applying Ville’s inequality [53] to the products in (12) and (13). Note that the dimensionality of \mathcal{X} does *not* in any way affect the validity of Theorem 1. Moreover, $\phi_t^{(\text{DR-}l)}$ and $\phi_t^{(\text{DR-u})}$ are always unbiased and yield valid CSs

regardless of how \widehat{r}_t is chosen. The reason to introduce these doubly robust pseudo-outcomes is to obtain lower-variance CSs (as illustrated in Figure 1) since doubly robust estimators can be semiparametric efficient thereby attaining the optimal asymptotic mean squared error in a local minimax sense. These details are outside the scope of the present paper, but we direct the interested reader to Kennedy [32] and Uehara et al. [51] for modern reviews discussing this subject.

Notice that Theorem 1 is a generalization of Proposition 1. Indeed, if the logging policies do not change (i.e. $h_1 = h_2 = \dots = h$), and if the observations $(X_t, A_t, R_t)_{t=1}^\infty$ are iid, and if $k_t = 0$ for each t , then Theorem 1 recovers Proposition 1 exactly. For this reason, we do not elaborate on empirical comparisons between Proposition 1 and Theorem 1 – any CS that can be derived using the former is a special case of the latter. Moreover, in the *on-policy* setting with all importance weights set to 1 and without a reward predictor, Theorem 1 recovers the betting-style CSs of Waudby-Smith and Ramdas [58, Theorem 3]. As alluded to in the discussion following Proposition 1, the infima above are straightforward to compute for many choices of predictable sequences $(\lambda_t^L(v'))_{t=1}^\infty$ and $(\lambda_t^U(v'))_{t=1}^\infty$ including all of those discussed in the following section.

2.2 Tuning, truncating, and mirroring

We make three remarks below, that are important on both theoretical and practical fronts.

Remark 1 (Tuning $(\lambda_t^L)_{t=1}^\infty$ and $(k_t)_{t=1}^\infty$). As stated, Theorem 1 yields a valid lower CS for v using *any* predictable sequence of $[0, (v' + k_t)^{-1}]$ -valued tuning parameters $(\lambda_t^L(v'))$ – referred to as “bets” by Karampatziakis et al. [28] and Waudby-Smith and Ramdas [58], but how should these bets be chosen? Waudby-Smith and Ramdas [58, Appendix B] discuss several possible options, but in practice none of them uniformly dominate the others. (This should not be surprising, since there is a certain formal sense in which different nontrivial nonnegative martingales cannot uniformly dominate each other for a given composite sequential testing problem; see Ramdas et al. [39] for a precise statement.) For a simple-to-implement and empirically compelling option, we suggest scaling ϕ_t^{DR} as $\xi_t := \phi_t^{\text{DR}} / (k_t + 1)$ and setting $\lambda_t^L(v')$ as

$$\lambda_t^L(v') := \sqrt{\frac{2 \log(1/\alpha)}{\widehat{\sigma}_{t-1}^2 t \log(1+t)}} \wedge \frac{c}{k_t + v'}, \text{ where} \quad (14)$$

$$\widehat{\sigma}_t^2 := \frac{\sigma_0^2 + \sum_{i=1}^t (\xi_i - \bar{\xi}_i)^2}{t+1}, \text{ and } \bar{\xi}_t := \left(\frac{1}{t} \sum_{i=1}^t \xi_i \right) \wedge \frac{1}{k_t + 1}, \quad (15)$$

with a similar definition for $\lambda_t^U(v')$ but with $c/(k_t + v')$ replaced by $c/(k_t + (1 - v'))$. Here, $c \in (0, 1)$ is some truncation scale, reasonable values of which may lie between 1/4 and 3/4, but it is of relatively minor practical importance, and for sufficiently large t , the choice of c will be inconsequential. A justification for why $\lambda_t^L(v')$ is a sensible choice can be found in Waudby-Smith and Ramdas [58, Section B.1], but the practitioner is nevertheless free to use any other sequence of bets, as long as they are predictable and satisfy the aforementioned boundedness constraints. Furthermore, as in Waudby-Smith and Ramdas [58], when $\lambda_t^L(v')$ is chosen as above, the product in (12) is quasiconvex in $v' \in [0, 1]$ and hence the infima in Theorem 1 (and Proposition 1) can be computed straightforwardly via line or grid search (see Waudby-Smith and Ramdas [58, Section A.5]).

The sequence of nonnegative $(k_t)_{t=1}^\infty$ that truncate the reward predictors can also be chosen in any way as long as they are predictable. There are several heuristics that one might use, with increasing levels of complexity. One option is to have a prior guess for w_{\max} (or some value w'_{\max} that the practitioner believes will upper-bound most importance weights) and set $k_t = w'_{\max} / C$ for some $C \geq 1$, e.g. $C = 2$. For a more adaptive option, one could set $k_t := \text{median}(w_1, \dots, w_{t-1})$, or even try out a grid of values $\{k^{(1)}, \dots, k^{(J)}\}$ and choose the $k^{(j)}$ that would have yielded the

tightest CSs in hindsight. Nevertheless, all three of these options yield nonasymptotically valid $(1 - \alpha)$ -CSs for v .

Remark 2 (Why truncate the reward predictor \widehat{r}_t ?). Readers familiar with doubly robust estimation in causal inference or contextual bandits will notice that if $k_1 = k_2 = \dots = \infty$, then $\phi_t^{(\text{DR-}\ell)}$ takes the form of a classical doubly robust estimator of the policy value, and that such estimators often vastly outperform those based on importance weighting alone (and in many cases, are provably more efficient, at least in an asymptotic sense), so why would we want to truncate \widehat{r}_t at all?

The reason has to do with the fact that for nonasymptotic inference, we exploit the lower-boundedness of $\phi_t^{(\text{DR-}\ell)}$ in order to show that the product in (12) is a nonnegative martingale. However, if we introduce a non-truncated reward predictor, we can only say that $\phi_t^{(\text{DR-}\ell)}$ is lower-bounded by $-w_{\max}$, which we do not want to assume knowledge of (or that it is finite at all). Truncation of \widehat{r}_t allows us to occupy a middle ground, so that many of the efficiency gains from doubly robust estimation can be realized, without entirely losing the lower-boundedness structure of $\phi_t^{(\text{DR-}\ell)}$.

This same line of thought helps to illustrate why including a reward predictor tightens our CSs for large t , potentially at the expense of tightness for smaller t . Notice that truncating the reward predictor at k_t/w_t simply restricts the tuning parameter $\lambda_t^L(v')$ to lie in $[0, (v' + k_t)^{-1}]$ instead of $[0, v'^{-1}]$ for importance weighting. Without dwelling on the details too much, it is known that smaller values of λ_t correspond to CSs and CIs being tighter for larger t – e.g. the role that λ plays in Hoeffding’s CIs looks like $\sqrt{\log(2/\alpha)/2n}$ [21] – but we refer the reader to papers on CSs for more in-depth discussions [24, 58]. The important takeaway for our purposes here, is that larger k_t corresponds to more variance adaptivity via double robustness, but does more to restrict the CSs tightness at small t . Nevertheless, this tradeoff is clearly worth it in some cases (see Figure 1).

We note that the idea to truncate \widehat{r}_t based on k_t/w_t was inspired by the so-called “reduced-variance” estimators of Zimmert and Lattimore [64], Zimmert and Seldin [65]. However, their reduced-variance estimators are slightly different since they multiply by an indicator $\mathbb{1}(w_i \leq \eta)$ for some $\eta \geq 0$, which sends the reward predictor to zero for large importance weights, whereas ours only truncates the reward predictor.

Remark 3 (Mirroring trick for upper CSs). Notice that in Proposition 1 and Theorem 1, upper CSs for v were obtained by importance weighting $1 - R_t$ rather than R_t to obtain a *lower* CS for $1 - v$, which was then translated into an *upper* CS for v . This “mirroring trick” – first used in the OPE setting by Thomas et al. [50] to the best of our knowledge – applies to all of the results that follow, but for the sake of succinctness, we will only explicitly write the lower CSs.

2.3 Closed-form confidence sequences

In Theorem 1, we derived CSs for the policy value that generalize and improve on prior work [28]. These bounds are empirically tight and can be computed efficiently, but are not closed-form, which may be desirable in practice. In this section, we derive a simple, closed-form, variance-adaptive CS for the fixed policy value $v := \mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1})$. Let

$$\xi_t := \frac{\phi_t^{(\text{DR-}\ell)}}{k_t + 1}, \text{ and } \widehat{\xi}_{t-1} := \left(\frac{1}{t-1} \sum_{i=1}^{t-1} \xi_i \right) \wedge \frac{1}{k_t + 1}. \quad (16)$$

With the above notation in mind, we are ready to state the main result of this section.

Proposition 2 (Closed-form predictable plug-in CS for v). *Given contextual bandit data $(X_t, A_t, R_t)_{t=1}^\infty$ with $[0, 1]$ -valued rewards, choose nonnegative and predictable tuning parameters $(k_t)_{t=1}^\infty$, and define*

$(\lambda_t)_{t=1}^\infty$ as

$$\lambda_t := \sqrt{\frac{2 \log(1/\alpha)}{\widehat{\sigma}_{t-1}^2 t \log(1+t)}} \wedge c, \quad \widehat{\sigma}_t^2 := \frac{\sigma_0^2 + \sum_{i=1}^t (\xi_i - \bar{\xi}_i)^2}{t+1}, \quad \bar{\xi}_t := \left(\frac{1}{t} \sum_{i=1}^t \xi_i \right) \wedge \frac{1}{k_t+1}, \quad (17)$$

where $c \in (0, 1)$ is some truncation parameter (reasonable values of which may include 1/2 or 3/4); and $\xi_0 \in (0, 1)$ and $\sigma_0^2 > 0$ are some user-chosen parameters that can be thought of as prior guesses for the mean and variance of ξ , respectively. Then,

$$L_t^{\text{PrPI}} := \left(\frac{\sum_{i=1}^t \lambda_i \xi_i}{\sum_{i=1}^t \lambda_i / (k_i + 1)} - \frac{\log(1/\alpha) + \sum_{i=1}^t (\xi_i - \widehat{\xi}_{i-1})^2 \psi_E(\lambda_i)}{\sum_{i=1}^t \lambda_i / (k_i + 1)} \right) \vee 0 \quad (18)$$

forms a lower $(1 - \alpha)$ -CS for v . An analogous upper CS $(U_t^{\text{PrPI}})_{t=1}^\infty$ follows by mirroring (Remark 3).




Fig. 2. Betting-based (Theorem 1) and predictable plug-in (PrPI) (Proposition 2) CSs for v with both importance-weighted (IW) and doubly robust (DR) variants. Notice that for both IW and DR CSs, the betting-based approach of Theorem 1 outperforms the PrPI CSs. Nevertheless, the closed-form PrPI CSs are simpler to implement, and can still benefit from doubly robust variance adaptation.

The proof can be found in Appendix A.4 and relies on Ville’s inequality [53] applied to a predictable plug-in empirical Bernstein supermartingale similar to that of Waudby-Smith and Ramdas [58, Section 3.2] but with a variant of Fan’s inequality [15] for lower-bounded random variables with upper-bounded means. As seen in Figure 2, the betting CSs of Theorem 1 still have better empirical performance than the closed-form predictable plug-in CSs of Proposition 2, but the latter are more computationally and analytically convenient, and are valid under the same set of assumptions. A similar phenomenon was observed by Waudby-Smith and Ramdas [58] for bounded random variables (outside the context of OPE). In the on-policy setting with all importance weights set to 1 and without a reward predictor, Proposition 2 recovers Waudby-Smith and Ramdas [58, Theorem 2].

It is important not to confuse $\widehat{\xi}_t$ with $\bar{\xi}_t$. The difference between them may seem minor since the former simply has access to one less data point than the latter, but they play two very different roles in Proposition 2: $\widehat{\xi}_{t-1}$ is a *predictable* sample mean that shows up in the width of L_t^{PrPI} explicitly, and its predictability is fundamental to the proof technique, while $\bar{\xi}_t$ is just used as a tool to obtain better estimates of $\widehat{\sigma}_t^2$ so that they can be plugged in to the tuning parameters λ_t . Consequently, $\widehat{\xi}_t$

can be found in CSs that rely on a similar proof technique (such as Theorem 2), while $\bar{\xi}_t$ can be found in other CSs that make use of predictable tuning parameters (such as Theorem 1).

2.4 Fixed-time confidence intervals

While this paper is focused on time-uniform CSs for OPE, our methods also naturally give rise to fixed-time CIs that are *not* anytime-valid but can still benefit from our general techniques. In this section, we will briefly discuss what minor modifications are needed to derive sharp fixed-time instantiations of our otherwise time-uniform bounds. We will also compare our fixed-time CIs to the CIs of Thomas et al. [50], but this comparison is by no means comprehensive. Indeed, our goal is not to show that our methods are “better” than prior work, even if some simulations may suggest this — instead, we aim to provide the reader with some context as to how our fixed-time instantiations fit within the broader literature on CIs for OPE.

Confidence intervals for policy values. We begin by deriving fixed-time analogues of the CSs for v presented in Theorem 1 and Proposition 2 — the former being a “betting-style” CS that is very tight in practice, and the latter being a closed-form predictable plug-in (PrPl) CS that is slightly more analytically and computationally convenient. In both cases, our suggested modification is essentially the same: choose a predictable sequence $(\lambda_t)_{t=1}^n$ that is tuned for the desired sample size n — to be elaborated on shortly — and take the intersection of the implicit CS $(C_t)_{t=1}^n$ that is formed from times 1 through n . Concretely, define the predictable sequence $(\lambda_t)_{t=1}^n$ given by

$$\hat{\lambda}_{t,n} := \sqrt{\frac{2 \log(1/\alpha)}{n \hat{\sigma}_{t-1}^2}}, \quad (19)$$

where $\hat{\sigma}_t^2$ is given as in (15). Then, we have the following corollary for betting-style CIs for v .

Corollary 1. *Let $(X_t, A_t, R_t)_{t=1}^n$ be a finite sequence of contextual bandit data with $[0, 1]$ -valued rewards and define $(\lambda_t^L(v'))_{t=1}^n$ and $(\lambda_t^U(v'))_{t=1}^n$ by*

$$\lambda_t^L(v') := \hat{\lambda}_{t,n} \wedge \frac{c}{k_t + v'}, \text{ and } \lambda_t^U(v') := \hat{\lambda}_{t,n} \wedge \frac{c}{k_t + (1 - v')}, \quad (20)$$

where $c \in (0, 1)$ is some truncation scale as in Theorem 1. Let L_t^{DR} and U_t^{DR} be as in (12) and (13). Then,

$$\hat{L}_n := \max_{1 \leq t \leq n} L_t^{\text{DR}} \text{ and } \hat{U}_n := \min_{1 \leq t \leq n} U_t^{\text{DR}} \quad (21)$$

form lower and upper $(1 - \alpha)$ -CIs for v , respectively, meaning $\mathbb{P}(v \in [\hat{L}_n, \hat{U}_n]) \geq 1 - \alpha$.

Corollary 1 is an immediate consequence of Theorem 1 where the sequence of tuning parameters was chosen to tighten the CI for the sample size n . This particular choice of $(\lambda_t^L(v'))_{t=1}^n$ is inspired by the fact that the product in (12) resembles an exponential supermartingale whose resulting CI can be tightened using tuning parameters that are well-estimated by (20). For more details, we refer the reader to Waudby-Smith and Ramdas [58, Section 3]. The fact that the maximum and minimum can be taken in (21) follows from the fact that $(L_t^{\text{DR}})_{t=1}^n$ satisfies $\mathbb{P}(\forall t \in \{1, \dots, n\}, v \geq L_t^{\text{DR}}) \geq 1 - \alpha$, and similarly for the upper CI \hat{U}_n . Figure 3 demonstrates what these CIs may look like in practice.

Similar to how Corollary 1 is a fixed-time instantiation of Theorem 1, the following corollary is a fixed-time instantiation of the closed-form PrPl CSs of Proposition 2.

Corollary 2. *Let $(X_t, A_t, R_t)_{t=1}^n$ be a finite sequence of contextual bandit data with $[0, 1]$ -valued rewards and define $(\lambda_t)_{t=1}^n$ by*

$$\lambda_t := \hat{\lambda}_{t,n} \wedge c, \quad (22)$$




Fig. 3. Fixed-time 90% confidence intervals for ν using three different methods: a betting-based CI (Corollary 1), a predictable plug-in (PrPI) CI (Corollary 2), and those presented in a paper entitled “High-confidence off-policy evaluation” (HCOPE15) by Thomas et al. [50]. Notice that the betting-based CI outperforms the closed-form PrPI CI, which itself significantly outperforms the bounds in Thomas et al. [50].

with c chosen as in Proposition 2. Then, with L_t^{PrPI} and U_t^{PrPI} defined in Proposition 2, we have that

$$\hat{L}_n^{\text{PrPI}} := \max_{1 \leq t \leq n} L_t^{\text{PrPI}} \text{ and } \hat{U}_n^{\text{PrPI}} := \min_{1 \leq t \leq n} U_t^{\text{PrPI}} \quad (23)$$

form lower and upper $(1 - \alpha)$ -CIs for ν , respectively, meaning $\mathbb{P}(\nu \in [\hat{L}_n^{\text{PrPI}}, \hat{U}_n^{\text{PrPI}}]) \geq 1 - \alpha$.

Corollary 2 is an immediate consequence of Proposition 2 instantiated for a different choice of λ_t and with an intersection being taken over the implicit CS from times 1 through n .

While the methods of this section improve on past work both theoretically and empirically, each of our results thus far have assumed that $\nu \equiv \mathbb{E}_\pi(R_t)$ is fixed and does not change over time, an assumption that we may not always wish to make in practice (e.g. if the environment is nonstationary). Fortunately, it is still possible to design CSs that capture an interpretable parameter: the *time-varying average policy value thus far*. However, we will need completely different supermartingales to achieve this, which we outline in the following section.

3 INFERENCE FOR TIME-VARYING POLICY VALUES

Let us now consider the more challenging task of performing anytime-valid off-policy inference for a time-varying average policy value. Concretely, suppose that the value of the $[0, 1]$ -bounded reward R_t under policy π is given by $v_t := \mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1}) \in [0, 1]$, and hence $(v_t)_{t=1}^\infty$ is now a *sequence* of conditional policy values. Our goal is to derive CSs for $(\tilde{\nu}_t)_{t=1}^\infty$ where $\tilde{\nu}_t := \frac{1}{t} \sum_{i=1}^t v_i$ is the *average conditional policy value so far*. In addition to satisfying desiderata 1–5 in Section 1.2 our CSs will impose no restrictions on how $(v_t)_{t=1}^\infty$ changes over time. Unfortunately, the techniques of Proposition 1 and Theorem 1 will not work here because their underlying test statistics cannot be written explicitly as functions of a candidate average policy value $\tilde{\nu}_t := \frac{1}{t} \sum_{i=1}^t v_i$, but only of the entire candidate tuple (v_1, \dots, v_t) . To remedy this, we will rely on test statistics that are functions of a candidate average $\tilde{\nu}_t$ to derive CSs for $\tilde{\nu}_t$ in Theorems 2 and 3.

We will present two CSs for $\tilde{\nu}_t$: (1) the “empirical Bernstein” CS in Theorem 2 whose underlying supermartingale is constructed using Robbins’ method of mixtures [42], and (2) the “iterated logarithm” CS in Proposition 3 which uses the stitching technique. Both yield time-uniform, nonasymptotically valid bounds, and are easy to compute. However, the former tends to have better

empirical performance in finite samples, while the latter achieves the (optimal) rate of convergence, matching the law of the iterated logarithm. Nevertheless, both boundaries shrink to zero-width at a rate of $\widetilde{O}(\sqrt{V_t}/t)$ [24]. Here, $\widetilde{O}(\cdot)$ means $O(\cdot)$ up to logarithmic factors.

In order to write down the empirical Bernstein CS, we first define the scaled doubly robust pseudo-outcomes $\xi_t := \phi_t/(1+k)$ where $\phi_t \equiv \phi_t^{(\text{DR-}\ell)}$ is given in (10) and with all k_t equal to a fixed k . Then, define the corresponding centered sum process $(S_t(\widetilde{v}_t))_{t=1}^\infty$ and variance process $(V_t)_{t=1}^\infty$ given by

$$S_t(\widetilde{v}_t) := \sum_{i=1}^t \xi_i - \frac{t\widetilde{v}_t}{1+k}, \text{ and} \quad (24)$$

$$V_t := \sum_{i=1}^t (\xi_i - \widehat{\xi}_{i-1})^2, \text{ where } \widehat{\xi}_t := \left(\frac{1}{t} \sum_{i=1}^t \xi_i \right) \wedge \frac{1}{1+k}, \quad (25)$$

and $\xi_0 \in [0, (1+k)^{-1}]$ is chosen by the user. With this setup and notation in mind, we are ready to state the empirical Bernstein CS for \widetilde{v}_t .

THEOREM 2 (EMPIRICAL BERNSTEIN CONFIDENCE SEQUENCE FOR \widetilde{v}_t). *Let $(X_t, A_t, R_t)_{t=1}^\infty$ be an infinite sequence of contextual bandit data with $[0, 1]$ -valued rewards generated by the sequence of policies $(h_t)_{t=1}^\infty$, and let $(S_t(\widetilde{v}_t))_{t=1}^\infty$ and $(V_t)_{t=1}^\infty$ be the centered sum and variance processes as in (24) and (25). Then for any $\rho > 0$,*

$$M_t^{\text{EB}}(\widetilde{v}_t) := \left(\frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \right) \left(\frac{1}{V_t + \rho} \right) {}_1F_1(1, V_t + \rho + 1, S_t(\widetilde{v}_t) + V_t + \rho) \quad (26)$$

forms a nonnegative supermartingale starting at one, where ${}_1F_1(\cdot, \cdot, \cdot)$ is Kummer's confluent hypergeometric function and $\Gamma(\cdot, \cdot)$ is the upper incomplete gamma function. Consequently,

$$L_t^{\text{EB}} := \inf \{ \widetilde{v}_t \in [0, 1] : M_t^{\text{EB}}(\widetilde{v}_t) < 1/\alpha \} \quad (27)$$

forms a lower $(1 - \alpha)$ -CS for \widetilde{v}_t , meaning $\mathbb{P}(\forall t \in \mathbb{N}, \widetilde{v}_t \geq L_t^{\text{EB}}) \geq 1 - \alpha$. Similarly, an upper CS can be derived by using $\phi_t^{(\text{DR-u})}$ defined in (11) and employing the mirroring trick described in Remark 3.

The proof of Theorem 2 can be found in Appendix A.3 and relies on an inequality due to Fan et al. [15] along with a mixture supermartingale analogous to that of Howard et al. [24, Proposition 9]. In the on-policy setting with all importance weights set to 1 and with no reward predictor (and hence $k = 0$), we have that Theorem 2 recovers the gamma-exponential mixture CS of Howard et al. [24, Proposition 9].

The tuning parameter $\rho > 0$ effectively dictates the neighborhood of intrinsic time — i.e. the value of V_t — at which L_t^{EB} is tightest, and it is rather straightforward to choose $\rho > 0$ given this interpretation. Following Howard et al. [24], $\rho > 0$ can be chosen to (approximately) tighten L_t^{EB} at $V_t = V^*$ by setting

$$\rho(V^*) := \sqrt{\frac{-2 \log \alpha + \log(-2 \log \alpha + 1)}{V^*}}. \quad (28)$$

Nevertheless, L_t^{EB} forms a valid lower $(1 - \alpha)$ -CS for \widetilde{v}_t regardless of how $\rho > 0$ is chosen, as long as this is done data-independently.

Readers familiar with gamma-exponential mixture supermartingales such as those in Howard et al. [23, 24] and Choe and Ramdas [8] may have expected to see a lower incomplete gamma function $\gamma(a, b)$ instead of ${}_1F_1(1, a, b)$. Indeed, ${}_1F_1(1, a, b)$ reduces to $\gamma(a, b)$ when $b \equiv S_t(\widetilde{v}_t) + V_t + \rho$ is nonnegative, but unlike the lower incomplete gamma, ${}_1F_1(1, a, b)$ is well-defined when $b < 0$. Writing (26) in terms of a lower incomplete gamma would have required lower-bounding this term

by a piece-wise function when $b \equiv S_t + V_t + \rho < 0$ as in Choe and Ramdas [8, Appendix C]; see Remark 5 for details.

While L_t^{EB} is not a closed-form bound, it can be computed efficiently using root-finding algorithms, and it can also be shown to achieve an asymptotic width of $O(\sqrt{V_t \log V_t}/t)$ (the justifications in Howard et al. [24] carry over to this scenario). While this rate is sufficient for deriving CSs with strong empirical performance that shrink to zero-width, one can derive CSs that achieve an improved rate of $O(\sqrt{V_t \log \log V_t}/t)$ using a different technique known as “stitching”.




Fig. 4. Various CSs for the time-varying policy value \tilde{v}_t . The left-hand side plot illustrates that while the betting-style CS of Theorem 1 is tight when \tilde{v}_t remains fixed, it fails to cover when \tilde{v}_t changes (in this case, there is an abrupt change at $t = 1000$). The right-hand side plot illustrates how Theorem 2 and Proposition 3 compare, both using their importance-weighted (IW) and doubly robust (DR) variants. Notice that while LIL-IW and LIL-DR attain optimal rates of convergence, the empirical Bernstein CSs (EB-IW and EB-DR) are much tighter in practice. In both cases, the DR variant outperforms the IW variant due to the reward being easy to predict in this particular example.

Proposition 3 (Variance-adaptive iterated logarithm confidence sequence for \tilde{v}_t). Let $\bar{V}_t := V_t \vee 1$ where V_t is given as in (25). Define the function $\ell_t(\alpha)$

$$\ell_t(\alpha) := 2 \log(\log \bar{V}_t + 1) + \log\left(\frac{1.65}{\alpha}\right). \quad (29)$$

Then we have that under the same conditions as Theorem 2,

$$L_t^{\text{LIL}} := (k+1) \left(\frac{1}{t} \sum_{i=1}^t \xi_i - \frac{\sqrt{2.13\ell_t(\alpha)\bar{V}_t + 1.76\ell_t(\alpha)^2}}{t} - \frac{1.33\ell_t(\alpha)^2}{t} \right) \vee 0 \quad (30)$$

forms a lower $(1 - \alpha)$ -CS for $\tilde{v}_t := \frac{1}{t} \sum_{i=1}^t v_i$, meaning $\mathbb{P}(\forall t, \tilde{v}_t \geq L_t^{\text{LIL}}) \geq 1 - \alpha$. An analogous upper CS $(U_t^{\text{LIL}})_{t=1}^\infty$ can be derived using the mirroring trick of Remark 3.

The proof in Appendix A.5 uses the “stitching” (sometimes called “peeling”) technique that is common in the derivation of LIL-type bounds [9, 27, 30] applied to linear sub-exponential boundaries. Importantly, $L_t^{\text{LIL}} \asymp O\left(t^{-1}\sqrt{V_t \log \log V_t}\right)$ which matches the unimprovable rate implied by the law of the iterated logarithm. We take a maximum with 0 in L_t^{LIL} (and hence an implicit minimum

with 1 in U_t^{LLL}) since $\tilde{v}_t \in [0, 1]$ for every $t \in \mathbb{N}$ by assumption. Of course, if one is in a stationary environment so that $v_1 = v_2 = \dots = v$, then $\tilde{v}_t = v$, and hence $(L_t^{\text{LLL}}, U_t^{\text{LLL}})$ forms a $(1 - \alpha)$ -CS for v .

Comparison of Theorems 1 and 2 with prior work. To the best of our knowledge, Thomas et al. [50] were the first to derive nonasymptotic confidence intervals for policy values in contextual bandits, and they did so without knowledge of w_{\max} . However, their bounds are not time-uniform, and the authors do not consider time-varying policy values nor data-dependent logging policies $(h_t)_{t=1}^{\infty}$. Four other prior works stand out as being related to the results of this section, namely Karampatziakis et al. [28], Howard et al. [24, Section 4.2], Bibaut et al. [4], and Zhan et al. [61] and we discuss each of them in some detail below. Note that in the last row labeled “Doubly robust” of Table 1, we are referring to the property of confidence sets to be potentially sharpened in the presence of regression estimators without compromising validity (as discussed in the paragraphs surrounding Theorem 1).

- **KMR21:** The off-policy CSs of Karampatziakis et al. [28] — reviewed in Proposition 1 — can in several ways be seen as an improvement of Thomas et al. [50] since they are time-uniform, in addition to being empirically tight. As discussed in Section 2, their importance-weighted off-policy CSs are strictly generalized and extended in our Theorem 1, but their result nevertheless yields the current state-of-the-art CSs for v .
- **BDKCvdL21 & ZHHA21:** Bibaut et al. [4] and Zhan et al. [61] both study the off-policy inference problem in contextual bandits from an asymptotic point of view. Their off-policy estimators take the form of sample averages of influence functions — what Bibaut et al. [4] refer to as the canonical gradient — to which martingale central limit theorems may be applied to obtain asymptotically valid inference. In contrast to our work, the confidence intervals of Bibaut et al. [4] and Zhan et al. [61] (a) are asymptotic, and hence do not have finite-sample guarantees, (b) are not time-uniform, and hence cannot be used at stopping times, and (c) do not track time-varying policy values.
- **HRMS21:** Howard et al. [24, Section 4.2] derive time-uniform, nonasymptotic CSs for the average treatment effect (ATE) in randomized experiments. The main difference between our results and those of [24, Section 4.2] is that they do not study the contextual bandit off-policy evaluation problem. However, since estimating the ATE in randomized experiments can be seen as a special case of the contextual bandit problem, it is natural to wonder how our approach differs in this special case. The main difference here is that Howard et al. [24] require knowledge of w_{\max} — or equivalently in their setup, the maximal and minimal propensity scores — while ours do not. Moreover, our results allow w_{\max} to be infinite and nevertheless enjoy variance-adaptivity. See Section 3.2 for a more detailed discussion of the implications of our bounds for ATE estimation in randomized experiments.

Table 1. Comparison of various CSs and CIs for mean off-policy values.

| | KMR21 | BDKCvdL21 & ZHHA21 | HRMS21 | Thm. 1 | Thm. 2 |
|------------------------------------|-------|--------------------|--------|--------|--------|
| Contextual bandits | ✓ | ✓ | | ✓ | ✓ |
| Time-varying rewards | | | ✓ | | ✓ |
| Nonasymptotic | ✓ | | ✓ | ✓ | ✓ |
| Time-uniform | ✓ | | ✓ | ✓ | ✓ |
| Predictable $(h_t)_{t=1}^{\infty}$ | | ✓ | ✓ | ✓ | ✓ |
| w_{\max} -free | ✓ | ✓ | | ✓ | ✓ |
| Doubly robust | | ✓ | ✓ | ✓ | ✓ |

3.1 A remark on policy value differences

The results in this paper have taken the form of CSs for policy values, e.g. a sequence of sets $[L_t, U_t]_{t=1}^\infty$ such that $\mathbb{P}(\forall t \in \mathbb{N}, v \in [L_t, U_t]) \geq 1 - \alpha$, but it may be of interest to directly estimate policy value *differences* — e.g. $\delta \equiv v_1 - v_2 \equiv v(\pi_1) - v(\pi_2)$, where $v(\pi)$ is the value of some policy π , and π_1 and π_2 are two policies we would like to compare. In many cases including the gated deployment problem studied by Karampatziakis et al. [28], π_1 is some target policy of interest and $\pi_2 = h_1 = h_2 = \dots = h$ is the logging policy so that $v(\pi_1) - v(h)$ can be interpreted as the additional value (or “lift”) in the target policy π_1 over the logging policy. However, our setup allows for π_1 and π_2 to be any two policies that are absolutely continuous with respect to the logging policies.

Of course, one can always solve this problem by union bounding: construct $(1 - \alpha/2)$ -CSs for v_1 and v_2 separately to yield a $(1 - \alpha)$ -CS for their difference. However, it is possible to remove this small amount of slack introduced by union bounding and instead derive a CS for the difference directly.

The idea is simple: rather than only leverage lower-boundedness of importance-weighted rewards $w_t R_t$, we construct a new random variable $\theta_t := w_t^{(1)} R_t - [1 - w_t^{(2)}(1 - R_t)]$ and leverage its lower-boundedness directly — here, $w_t^{(1)} = \pi_1/h_t$ and $w_t^{(2)} = \pi_2/h_t$ are the importance weights for policies π_1 and π_2 , respectively. In particular, notice that

$$\mathbb{E}[\theta_t] = \delta, \text{ and } \theta_t \geq -1, \text{ and hence} \quad (31)$$

$$\frac{1}{2}[\theta_t - \delta] \geq -1 \text{ almost surely.} \quad (32)$$

Consequently, (32) can be used in the proofs of our theorems to derive a CS for δ directly, since those proofs fundamentally rely on the centered (i.e. with their mean subtracted) random variables being almost surely lower-bounded by -1 . For instance, we have that for any $(0, 1)$ -valued predictable sequence $(\lambda_t(\delta))_{t=1}^\infty$,

$$M_t(\delta) := \prod_{i=1}^t (1 + \lambda_i(\delta) \cdot (\theta_i - \delta)/2) \quad (33)$$

forms a test martingale and hence $L_t := \inf \{\delta' \in [-1, 1] : M_t(\delta') < 1/\alpha\}$ forms a lower $(1 - \alpha)$ -CS for δ . As usual, the mirroring trick can be used to obtain an upper CS for this policy value difference. Moreover, the above discussion can be extended to time-varying policy value differences and doubly robust pseudo-outcomes (rather than just their importance-weighted counterparts), as well as sequences of policies — i.e. analyzing the sequences $(\pi_1^{(t)})_{t=1}^\infty$ and $(\pi_2^{(t)})_{t=1}^\infty$ — but we omit these derivations for the sake of brevity.

3.2 Time-varying treatment effects in adaptive experiments

While this paper is focused on anytime-valid *contextual bandit* inference — i.e. inference for policy values or their CDFs from contextual bandit data — one can nevertheless view off-policy evaluation as a generalization of treatment effect estimation from adaptive experiments. Consequently, every single result in this paper also has powerful implications for nonasymptotic inference for treatment effects from such experiments. In this section, we will focus on adaptive experiments with binary treatments for simplicity, but the analogy extends to more general settings.

From contextual bandits to adaptive experiments with binary treatments. The contextual bandit problem can be seen as a generalization of adaptive experiments since the latter has three key notational differences.

- (1) The “context” X_t is typically referred to as a “covariate” or a “feature”, and may be used to represent baseline demographics and medical history in a clinical trial, for example.

- (2) The “action” A_t (which is binary in this case) is referred to as a “treatment”, and the policy h_t is called the “propensity score”, and is simply the probability of a subject with covariates X_t receiving treatment $A_t = 1$ at time t .
- (3) The “reward” R_t is often referred to as the “outcome” for subject t .

There are many reasons why one may wish to run an adaptive sequential experiment rather than a simple Bernoulli(h) experiment with a constant pre-specified propensity score h . Two simple examples include: (a) balancing designs such as Efron’s biased coin [14] which vary the propensity scores $(h_t)_{t=1}^\infty$ over time to ensure that treatment groups are “balanced” within certain levels of the covariates, and (b) the experimental designs of Kato et al. [29] which adaptively choose propensity scores to minimize the variance of the resulting doubly robust and inverse propensity-weighted (IPW) estimators, yielding sharper confidence sets. (In the language of contextual bandits and off-policy evaluation, IPW and importance weighting are equivalent.) Both (a) and (b) – or any other design that varies propensity scores adaptively over time – can be paired with the CSs of the current paper.

Implications for causal inference in adaptive experiments. From the perspective of treatment effect estimation, the current paper provides nonasymptotic, nonparametric, time-uniform inference for treatment effects, all without knowledge of the minimal propensity score $h_{\min} := \text{ess inf}_{t,a,x} h_t(a | x)$ and this essential infimum can even be 0 (as long as it is not attained, meaning each $h_t(a | x)$ is itself positive). Contrast this with prior work on nonasymptotic, nonparametric, time-uniform inference for treatment effects such as Howard et al. [24, Section 4.2], which require *a priori* knowledge of $h_{\min} > 0$. Their bounds necessarily scale with an implied upper-bound on the variance of $h_t^{-1}R_t$ implied by h_{\min}^{-1} while ours only scale with the *empirical variance* of $(h_t^{-1}R_t)_{t=1}^\infty$ – the latter always being smaller.

Concretely, Theorem 1 can be used to derive CSs for the average treatment effect from adaptively collected data in experiments with binary treatments and bounded outcomes. Theorem 2 goes further, enabling the construction of CSs for *time-varying average treatment effects* similar to Howard et al. [24, Section 4.2]. Finally, Theorem 3 – to be presented in Section 4 – allows for the construction of time-uniform confidence bands for the CDF of the outcome distribution under a given treatment. Moreover, all of this is possible in a nonparametric, nonasymptotic framework, without knowledge (or strict positivity) of h_{\min} . To the best of our knowledge, all three of these implied results are new in the literature for treatment effect estimation.

3.3 Sequential testing and anytime p -values for off-policy inference

While we have thus far taken an *estimation* perspective (i.e. deriving CSs and CIs rather than p -values), all of our results have hypothesis testing analogues. In particular, the CSs and CIs developed in this paper have all been built by first deriving implicit *e-processes*. Formally, given a set of distributions \mathcal{P}_0 (referred to as “the null hypothesis”), an *e-process* $E \equiv (E_t)_{t=1}^\infty$ for \mathcal{P}_0 is a nonnegative process such that $\mathbb{E}_P[E_\tau] \leq 1$ for any $P \in \mathcal{P}_0$ and any stopping time τ . (In particular, all test supermartingales for \mathcal{P}_0 are *e-processes* by the optional stopping theorem, but not vice versa.)

While *e-processes* can serve as tools to derive CSs, they can also be used as interpretable testing tools in their own right, or as a way to derive anytime p -values – p -values that are uniformly valid over time in the same sense as CSs. Formally, an anytime p -value for \mathcal{P}_0 is an \mathcal{H} -adapted process $(p_t)_{t=1}^\infty$ such that

$$\sup_{P \in \mathcal{P}_0} P(\exists t \in \mathbb{N} : p_t \leq \alpha) \leq \alpha. \quad (34)$$

Compare (34) with a traditional fixed-time p -value p_n that satisfies $\forall n \in \mathbb{N}$, $\sup_{P \in \mathcal{P}_0} P(p_n \leq \alpha) \leq \alpha$. On the other hand, an e -process $(E_t)_{t=1}^\infty$ for \mathcal{P}_0 also satisfies Ville's inequality:

$$\sup_{P \in \mathcal{P}_0} P(\exists t \in \mathbb{N} : E_t \geq 1/\alpha) \leq \alpha. \quad (35)$$

As a direct consequence of (35), notice that e -processes yield anytime p -values via the transformation $p_t := (1/E_t) \wedge 1$.

Remark 4 (Which should you choose: e or p ?). There are several philosophical and practical reasons why one may wish to use e -processes over anytime p -values, despite the fact that they can both be used for sequential hypothesis testing. Philosophically, e -processes (and hence test supermartingales) have game-theoretic interpretations and connections to Bayesian statistics [17, 47, 57], and have been argued to serve as a better foundation for statistical communication [46]. Practically, stopped e -processes form e -values — nonnegative random variables with expectation at most one [54] — which have several attractive properties over p -values, including the fact that they are very straightforward to combine for the sake of testing a global null [54], to perform meta-analyses [49], or to control the false discovery rate under arbitrary dependence [55]. In this section, we remain agnostic as to which of the two one should use: we will simply derive e -processes and note that their philosophical and practical properties can be enjoyed within OPE, and that if anytime p -values are preferred, they are always available via the transformation $p_t := (1/E_t) \wedge 1$.

Following Section 3.1, let us now derive sequential tests for whether a policy π_1 has a higher average value than some other policy π_2 . Technically, one could also replace π_1 or π_2 with a *sequence* of predictable policies, but for simplicity we will only discuss fixed policies. Concretely, let Δ_t denote the difference in the values of policies π_1 and π_2 at time t ,

$$\delta_t := v_t(\pi_1) - v_t(\pi_2) \equiv \mathbb{E}_{A_t \sim \pi_1}(R_t) - \mathbb{E}_{A_t \sim \pi_2}(R_t), \quad (36)$$

and let $\Delta_t := \frac{1}{t} \sum_{i=1}^t \delta_i$ denote the running average difference. We are interested in testing the *weak* null hypothesis H_0 ,

$$H_0 : \forall t, \Delta_t \leq 0, \text{ vs } H_1 : \exists t : \Delta_t > 0. \quad (37)$$

In words, H_0 says that “ π_1 is no better than π_2 on average thus far” and was used to compare sequential forecasters in Choe and Ramdas [8]. This “weak null” should be contrasted with the “strong null” that would posit $H_0^* : \forall t, \delta_t \leq 0$ — clearly, the latter implies the former, and hence any e -process (or anytime p -value) for H_0 can also be used for H_0^* . Mathematically, H_0 is a composite superset of the point null H_0^* . From a practical perspective, H_0 may be a favorable null to test since it allows for $v_t(\pi_1) > v_t(\pi_2)$ at various t as long as $\frac{1}{t} \sum_{i=1}^t v_i(\pi_1) \leq \frac{1}{t} \sum_{i=1}^t v_i(\pi_2)$ for all t , whereas H_0^* requires π_1 to be uniformly dominated by π_2 .

Let us now derive an explicit e -process for H_0 . Using the techniques of Section 3.1, define

$$\theta_t := w_t^{(1)} R_t - (1 - w_t^{(2)})(1 - R_t), \quad (38)$$

where $w_t^{(1)} := \pi_1(A_t | X_t)/h_t(A_t | X_t)$ and $w_t^{(2)} := \pi_2(A_t | X_t)/h_t(A_t | X_t)$ are the importance weights for policies π_1 and π_2 . As before, we note that $\theta_t \geq -1$ and $\mathbb{E}(\theta_t | \mathcal{H}_{t-1}) = \delta_t$, and hence

$$\mathbb{E} \left(\frac{1}{t} \sum_{i=1}^t \theta_i \mid \mathcal{H}_{t-1} \right) = \Delta_t. \quad (39)$$

Given the above setup, we are ready to derive an e -process (and hence an anytime p -value) for the weak null H_0 .

Proposition 4. Given contextual bandit data $(X_t, A_t, R_t)_{t=1}^\infty$ and two target policies π_1 and π_2 that we would like to compare, define $S_t(\Delta'_t)$ and V_t by

$$S_t(\Delta'_t) := \frac{1}{2} \left(\sum_{i=1}^t \theta_i - t\Delta'_t \right) \text{ and} \quad (40)$$

$$V_t := \frac{1}{2} \sum_{i=1}^t (\theta_i - \widehat{\theta}_{i-1})^2, \text{ where } \widehat{\theta}_t := \frac{1}{2} \left[\left(\frac{1}{t} \sum_{i=1}^t \theta_i \right) \wedge 1 \right] \quad (41)$$

Then, given any $\rho > 0$, we have that

$$M_t^{\text{EB}}(0) := \left(\frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \right) \left(\frac{1}{V_t + \rho} \right) {}_1F_1(1, V_t + \rho + 1, S_t(0) + V_t + \rho) \quad (42)$$

forms an e -process for H_0 . Consequently, $p_t := (1/M_t^{\text{EB}}) \wedge 1$ forms an anytime p -value for the weak null $H_0 : \forall t, \Delta_t \leq 0$, meaning $\sup_{P \in H_0} P(\exists t : p_t \leq \alpha) \leq \alpha$.




Fig. 5. An illustration of how the anytime p -value derived in Proposition 4 can be used to test the weak null $H_0 : \forall t, \Delta_t \leq 0$. In the left-hand side plot, notice that δ_t ventures above 0 at several points prior to $t = 2037$, but the average policy value difference is positive for the first time at $t = 2037$. In the right-hand side plot, we see that the anytime p -value dips below α shortly after $\Delta_t > 0$, at which point the weak null can be safely rejected, with no penalties for the p -value having been continuously monitored.

The fact that $M_t^{\text{EB}}(0)$ forms an e -process is easy to see: under H_0 , we have that $\Delta_t \leq 0$ and notice that M_t^{EB} with $S_t(0)$ replaced with $S_t(\Delta_t)$ forms a test supermartingale using the same techniques as Section 3. Since $S_t(\cdot)$ is nonincreasing and ${}_1F_1$ is nondecreasing in its third argument, we have that M_t^{EB} is upper-bounded by the aforementioned test supermartingale whenever $\Delta_t \leq 0$. The claimed e -process property is then an immediate consequence of the optional stopping theorem applied to the above test supermartingale.

4 TIME-UNIFORM INFERENCE FOR THE OFF-POLICY CDF

Thus far we have focused on off-policy inference for *mean policy values*, i.e. functionals of the form $v := \mathbb{E}_\pi(R)$. In some cases, however, it may be of interest to study quantiles (e.g. median or 75th percentile) or perhaps the entire cumulative distribution function (CDF) of the reward distribution under policy π . In this section, we focus on the latter, deriving confidence bands for the CDF $\mathbb{P}_\pi(R \leq r)$ of the reward R under policy π . Our confidence bands will be uniform in two senses: in time, and in the quantiles. Concretely, if $Q(p)$ and $Q^-(p)$ are the right (standard) and left quantiles,

respectively — meaning $Q(p) := \sup \{x \in \mathbb{R} : \mathbb{P}_\pi(R \leq x)\}$ and $Q^-(p) := \sup \{x \in \mathbb{R} : \mathbb{P}_\pi(R < x)\}$ — then we will derive a sequence of confidence bands $[L_t(p), U_t(p)]_{t \in \mathbb{N}}$ such that

$$\mathbb{P}(\forall t \in \mathbb{N}, p \in (0, 1), L_t(p) < Q^-(p) \text{ and } Q(p) < U_t(p)) \geq 1 - \alpha. \quad (43)$$

Such a guarantee enables anytime-valid inference at arbitrary stopping times for *all quantiles* simultaneously, (as well as any functional thereof). In addition, all our confidence bands will satisfy all five desiderata laid out in Section 1.2, and they will consistently shrink to the true quantile $Q(p)$ for all p .

In order to state our main result, we first need to define a few terms. Define $W_t, \bar{W}_t, \bar{q}_t(p)$, and $\ell_t(p; \alpha)$ given by

$$W_t := \sum_{i=1}^t w_i^2, \quad \bar{W}_t := W_t \vee 1, \quad (44)$$

$$\bar{q}_t(p) := \text{logit}^{-1} \left(\text{logit}(p) + 4 \sqrt{\frac{e}{\bar{W}_t}} \right), \quad (45)$$

$$\ell_t(p; \alpha) := 2 \log \left(\log \bar{W}_t + 1 \right) + 2 \log \left(\left\| \frac{\sqrt{\bar{W}_t} \text{logit}(p)}{4} \right\| \vee 1 \right) + \log \left(\frac{7.06}{\alpha} \right), \quad (46)$$

$$\text{and } \mathfrak{B}_t(p; \alpha) := \frac{\sqrt{2.13 \ell_t(p; \alpha) \bar{W}_t + 1.76 \bar{q}_t(p)^2 \ell_t(p; \alpha)^2}}{t} \quad (47)$$

$$+ \frac{1.33 \bar{q}_t(p) \ell_t(p; \alpha) + t(\bar{q}_t(p) - p)}{t}. \quad (48)$$

While some of the above may seem complicated, they arise naturally from the proof technique discussed below, and it is straightforward to implement them into code. Given the above setup, we are ready to state the main result of this section.

THEOREM 3 (TIME-UNIFORM CONFIDENCE BAND FOR THE OFF-POLICY CDF). *Consider a sequence of contextual bandit data $(X_t, A_t, R_t)_{t=1}^\infty$ with real-valued (i.e. not necessarily $[0, 1]$ -bounded) rewards. Let $\widehat{F}_t^\pi(x) := \frac{1}{t} \sum_{i=1}^t w_i \mathbb{1}(R_i \leq x)$ be the importance-weighted empirical CDF, and let $\widehat{Q}_t(p)$ and $\widehat{Q}_t^-(p)$ be the upper and lower empirical quantiles, meaning*

$$\widehat{Q}_t(p) := \sup \left\{ x \in \mathbb{R} : \widehat{F}_t^\pi(x) \leq p \right\}, \quad (49)$$

and similarly for $\widehat{Q}_t^-(p)$ with \leq in the above supremum replaced by a strict inequality $<$. Then,

$$\mathbb{P} \left(\forall t \in \mathbb{N}, p \in (0, 1), Q(p) < \widehat{Q}_t^-(p + \mathfrak{B}_t(p; \alpha)) \wedge 1 \right) \geq 1 - \alpha. \quad (50)$$

Similarly, after applying the mirroring trick of Remark 3, we have that

$$\mathbb{P} \left(\forall t \in \mathbb{N}, p \in (0, 1), \widehat{Q}_t \left(\left[p + \frac{1}{t} \sum_{i=1}^t w_i - 1 - \mathfrak{B}_t(1-p; \alpha) \right] \vee 0 \right) < Q^-(p) \right) \geq 1 - \alpha. \quad (51)$$

The proof in Appendix A.6 modifies the “double stitching” technique of Howard and Ramdas [22, Theorem 5] to handle importance-weighted observations, and relies on a sub-exponential concentration inequality rather than a sub-Bernoulli one. Notice that (50) and (51) could be written without \widehat{Q}_t^- and Q^- — i.e. replacing \widehat{Q}_t^- and Q^- with \widehat{Q}_t and Q , respectively — but this would never result in a tighter bound.




Fig. 6. Time- and quantile-uniform 90% confidence band for the off-policy CDF $\mathbb{P}_\pi(R \leq Q)$ in a Bernoulli(1/2) experiment with the target policy set to $\pi(a | x) = \mathbb{1}(a = 1)$ – i.e. “always play action 1”. Here, the off-policy distribution of R is a beta distribution with $\alpha = \beta = 10$. These confidence bands are simultaneously valid for all $Q \in \mathbb{R}$ and all $t \in \mathbb{N}$ (though we only display them at $t \in \{10^3, 10^4, 10^5\}$ above). In particular, notice that as t gets larger, the confidence bands shrink towards the true CDF (and will continue to do so in the limit).

Many of the CSs throughout this paper have recovered prior CSs in the literature when specialized to the on-policy regime (that is, when all importance weights are set to 1 and reward predictors are set to 0). Examples include Theorem 1 recovering Waudby-Smith and Ramdas [58, Theorem 3] or Theorem 2 recovering Howard et al. [24, Proposition 9]. However, Theorem 3 does not recover the on-policy bound it most resembles (Howard and Ramdas [22, Theorem 5]). The reason for this is subtle, and has to do with the fact that in the on-policy regime, $\mathbb{1}(R_t \leq Q(p)) - p$ is a Bernoulli(p) random variable, hence their partial sums $\frac{1}{t} \sum_{i=1}^t [\mathbb{1}(R_i \leq Q(p)) - p]$ form sub-Bernoulli processes [22–24] with variance process $tp(1-p)$ regardless of the value of $Q(p)$. On the other hand, in the off-policy setting, we use importance-weighted indicators $w_t \mathbb{1}(R_t \leq Q(p)) - p$ whose partial sums are not sub-Bernoulli, but are instead sub-exponential with a variance process that depends on $Q(p)$. This fundamental difference changes the test supermartingales that we have access to, and consequently alters the downstream CSs.

Comparison with prior work. There are three prior works that are related to the results of this section, namely those of Howard and Ramdas [22], Chandak et al. [5], and Huang et al. [25], but it is important to note that none of them solve the problem that we are studying – time-uniform confidence bands for off-policy CDFs – and hence we focus on a theoretical comparison with these prior works, rather than an empirical one. We discuss each of them in detail below and summarize how they compare with the present paper in Table 2.

- **HR22:** Howard and Ramdas [22] derive time- and quantile-uniform confidence bands for the CDF of iid random variables in the *on-policy* setting, and in particular, Theorem 5 in their paper satisfies a guarantee of the form (43). However, Howard and Ramdas [22] do not consider the off-policy inference problem that we do here, and hence the “Predictable $(h_t)_{t=1}^\infty$ ” and “ w_{\max} -free” rows are not applicable (N/A). In addition, our setup can be seen as a generalization of theirs if all importance weights are set to 1.

- UnO21:** In the paper entitled “Universal off-policy evaluation” (UnO), Chandak et al. [5] derive fixed-time quantile-uniform confidence bands for the off-policy CDF. Cleverly exploiting monotonicity of CDFs, they reduce the problem to computing finitely many confidence intervals for means of importance-weighted bounded random variables, and taking a union bound over them. Notably, their bounds do not require knowledge of w_{\max} . The main difference between our bounds and those of Chandak et al. [5] is that ours are both time- and quantile-uniform, while theirs are only quantile-uniform. Note that Chandak et al. [5] do consider the more general setup of reinforcement learning in Markov Decision Processes (MDP) – an area to which we intend to extend all of our CSs in the future – and MDPs include the contextual bandit setting as a special case. When focusing on contextual bandits specifically, however, and even when ignoring time-uniformity, Theorem 3 improves on the fixed-time results of Chandak et al. [5] in the two following ways. First, the confidence bands of Chandak et al. [5] are not guaranteed to shrink to the true CDF as the sample size grows to infinity while ours are (and at an explicit rate of $O(\sqrt{\log t/t})$), which we refer to as “consistency” in Table 2. Second, their bounds assume that the logging policy h is fixed, while ours can take the form of a sequence of data-dependent logging policies $(h_t)_{t=1}^\infty$, such as those that result from online learning algorithms. However, since Chandak et al. [5, Theorem 2] simply requires taking a union bound over several CIs for importance-weighted bounded random variables, their Theorem 2 can presumably be extended to handle predictable logging policy sequences by employing the CIs provided in Corollaries 1 or 2.
- HLLA21:** Huang et al. [25] derive impressive quantile-uniform confidence bands for the off-policy CDF. Their bounds are elegant and simple to state, resembling the famous Dvoretzky-Kiefer-Wolfowitz (DKW) inequalities and their sharpened forms [13, 37]. Notably, their bounds are consistent for the true CDF at $O(1/\sqrt{n})$ rates. Similar to Chandak et al. [5], however, their results are not time-uniform, and hence do not permit valid inference at stopping times, unlike ours. Moreover, all of their bounds require knowledge of w_{\max} (and for this value to be finite), while our bounds do not. Finally, similar to Chandak et al. [5], their bounds assume that the logging policy h is fixed.

Table 2. Comparison of various uniform confidence bands for the CDF

| | HR22 | UnO21 | HLLA21 | Thm. 3 |
|----------------------------------|------|-------|--------|--------|
| Off-policy | | ✓ | ✓ | ✓ |
| Time-uniform | ✓ | | | ✓ |
| Consistency | ✓ | | ✓ | ✓ |
| Predictable $(h_t)_{t=1}^\infty$ | N/A | | | ✓ |
| w_{\max} -free | N/A | ✓ | | ✓ |

5 SUMMARY & EXTENSIONS

This paper derived time-uniform confidence sequences for various parameters in off-policy evaluation which remain valid even in contextual bandit setups where data are collected adaptively and sequentially over time. We began in Section 1.2 by laying out our desiderata for off-policy inference: we sought methods that (1) are exact and nonasymptotically valid, (2) only make nonparametric assumptions such as boundedness, (3) are time-uniform, and hence valid at arbitrary stopping times, (4) do not require knowledge of extreme values of importance weights, and (5) allow data to be collected by data-dependent logging policies.

In Section 2, we began by studying the most classical off-policy parameter — a fixed policy value v — and we derived CSs that strictly generalize prior state-of-the-art CSs by weakening the required assumptions and allowing for variance reduction via double robustness. In the same section, we also develop the first closed-form confidence sequences for policy values, as well as some tight fixed-time confidence intervals that are instantiations of our time-uniform bounds. Section 3 then developed CSs for a more general parameter: the time-varying average policy value $\bar{v}_t := \frac{1}{t} \sum_{i=1}^t v_i$, and we discussed what implications these bounds have for adaptive sequential experiments, such as online A/B tests.

Finally, in Section 4, we derived simultaneously valid CSs for every quantile of the off-policy reward distribution. Said differently, these bounds form time-uniform confidence bands for the CDF of the off-policy reward distribution.

There are a few other works that consider CSs and test supermartingales without reference to OPE or contextual bandits, but that nevertheless now have interesting consequences in the OPE problem once paired with the present paper. In particular, we want to highlight the implications that Wang and Ramdas [55] and Xu et al. [60] have on false discovery rate control in OPE, and how Waudby-Smith et al. [59] immediately yields algorithms for differentially private OPE. We briefly discuss these implications below, but omit their full derivations since these extensions are rather simple and not central to the current paper.

False discovery/coverage rate control under arbitrary dependence. Suppose that rather than estimate a single policy value v , we are interested in a collection (v_1, \dots, v_J) containing the values of the policies (π_1, \dots, π_J) . When testing several hypothesis or constructing several CIs, etc., it is often of interest to control some multiple testing metric, such as the false discovery rate (FDR), or the false coverage rate (FCR), respectively [2, 3]. Rather surprisingly, Wang and Ramdas [55] and Xu et al. [60] show that for tests and CIs built from e -values — nonnegative test statistics with expectation at most one — the FDR and FCR can be controlled under *arbitrary* dependence with virtually no modification to the famous Benjamini-Hochberg [2] and Benjamini-Yekutieli [3] procedures, while this fact is not true for generic tests and CIs. Relevant to the current paper, *all* of our CSs are fundamentally built from test supermartingales which form e -values at arbitrary stopping times. As a concrete consequence, we can take a collection of stopped CSs $C_\tau^{(1)}, \dots, C_\tau^{(J)}$ for $(\pi_j)_{j \in [J]}$, adjust them via the e -BY procedure of Xu et al. [60] to produce $\tilde{C}_\tau^{(1)}, \dots, \tilde{C}_\tau^{(J)}$, so that the FCR is controlled at some desired level $\delta \in (0, 1)$. The ability control the FCR under arbitrary dependence is crucial for our setting since the CSs $(C_\tau^{(j)})_{j \in [J]}$ are highly dependent and constructed from the same data, but with different importance weights. Similar implications hold for sequential tests and control of the FDR via the e -BH procedure of Wang and Ramdas [55].

Locally differentially private off-policy evaluation in contextual bandits. Waudby-Smith et al. [59] developed nonparametric CSs and CIs for means of bounded random variables under privacy constraints. The authors developed a so-called “Nonparametric randomized response” (NPRR) mechanism that serves as a nonparametric generalization of Warner’s randomized response [56], mapping a $[0, 1]$ -bounded random variable Y_t to a new random variable Z_t so that each Z_t is an ε -locally differentially private view of Y_t with mean $r\mathbb{E}(Y_t) + (1-r)/2$, where r is a known quantity that depends on ε (and hence it is possible to work out what $\mathbb{E}(Y_t)$ is). While Waudby-Smith et al. [59] did not explicitly consider the contextual bandit setup, they did develop CSs for time-varying treatment effects in sequential experiments, similar to the discussion in Section 3.2. However, like other prior work, their CSs require *a priori* knowledge of the minimal propensity score (in the language of this paper: they require knowledge of w_{\max}). Nevertheless, it is possible to derive locally private CSs for (time-varying) policy values without knowledge of w_{\max} using the techniques of

the current paper. Moreover, several policies can be evaluated from a *single* application of NPRR, thereby avoiding inflation of the privacy parameter ϵ from evaluating multiple policies. That is, given $[0, 1]$ -bounded rewards $(R_t)_{t=1}^{\infty}$, we can use NPRR to generate private views $(Z_t)_{t=1}^{\infty}$ of these rewards, and notice that $\mathbb{E}(w_t Z_t) = \mathbb{E}_{A_t \sim \pi}(Z_t) = r\mathbb{E}_{A_t \sim \pi}(R_t) + (1 - r)/2$, and hence a CS for $\mathbb{E}(w_t Z_t)$ can be translated into a CS for $\mathbb{E}_{A_t \sim \pi}(R_t)$ even though we only see a privatized version of R_t . In particular, practitioners can derive locally private CSs for time-varying policy values using Theorem 2 for several policies π_1, \dots, π_J , with only a single application of NPRR.

We believe that this paper presents a comprehensive treatment of OPE inference, yielding procedures that are theoretically valid under more general settings and yet deliver state-of-the-art practical performance. A challenging open problem is to extend these techniques to the off-policy MDP (Markov Decision Process) setting, where the actions at each step affect subsequent covariate and reward distributions, as captured by state variables. Another important open problem is to design practical OPE inference methods not just for one policy, but uniformly over an entire family of policies.

Acknowledgements

IW-S thanks Martin Larsson, Alec McClean, Steve Howard, Ruohan Zhan, and Edward H. Kennedy for helpful discussions. The authors acknowledge support from NSF grants IIS-2229881 and DMS-2310718. AR acknowledges funding from NSF Grant DMS2053804 and from ARL IoBT REIGN. Research reported in this paper was sponsored in part by the DEVCOM Army Research Laboratory under Cooperative Agreement W911NF-17-2-0196 (ARL IoBTCRA). The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation herein.

REFERENCES

- [1] Alberto Abadie, Susan Athey, Guido W Imbens, and Jeffrey M Wooldridge. Sampling-based versus design-based uncertainty in regression analysis. *Econometrica*, 88(1):265–296, 2020.
- [2] Yoav Benjamini and Yoel Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300, 1995.
- [3] Yoav Benjamini and Daniel Yekutieli. False discovery rate–adjusted multiple confidence intervals for selected parameters. *Journal of the American Statistical Association*, 100(469):71–81, 2005.
- [4] Aurélien Bibaut, Maria Dimakopoulou, Nathan Kallus, Antoine Chambaz, and Mark van der Laan. Post-contextual-bandit inference. *Advances in Neural Information Processing Systems*, 34:28548–28559, 2021.
- [5] Yash Chandak, Scott Niekum, Bruno da Silva, Erik Learned-Miller, Emma Brunskill, and Philip S Thomas. Universal off-policy evaluation. *Advances in Neural Information Processing Systems*, 34, 2021.
- [6] Haoyu Chen, Wenbin Lu, and Rui Song. Statistical inference for online decision making: In a contextual bandit setting. *Journal of the American Statistical Association*, 116(533):240–255, 2021.
- [7] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *The Econometrics Journal*, 21(1): C1–C68, 2018.
- [8] Yo Joong Choe and Aaditya Ramdas. Comparing sequential forecasters. *Operations Research*, 2023.
- [9] DA Darling and Herbert Robbins. Confidence sequences for mean, variance, and median. *Proceedings of the National Academy of Sciences of the United States of America*, 58(1):66, 1967.
- [10] Maria Dimakopoulou, Zhimei Ren, and Zhengyuan Zhou. Online multi-armed bandits with adaptive inference. *Advances in Neural Information Processing Systems*, 34:1939–1951, 2021.
- [11] Miroslav Dudík, John Langford, and Lihong Li. Doubly robust policy evaluation and learning. In *Proceedings of the 28th International Conference on Machine Learning*, pages 1097–1104, 2011.
- [12] Miroslav Dudík, Dumitru Erhan, John Langford, and Lihong Li. Doubly robust policy evaluation and optimization. *Statistical Science*, 29(4):485–511, 2014.

- [13] Aryeh Dvoretzky, Jack Kiefer, and Jacob Wolfowitz. Asymptotic minimax character of the sample distribution function and of the classical multinomial estimator. *The Annals of Mathematical Statistics*, pages 642–669, 1956.
- [14] Bradley Efron. Forcing a sequential experiment to be balanced. *Biometrika*, 58(3):403–417, 1971.
- [15] Xiequan Fan, Ion Grama, and Quansheng Liu. Exponential inequalities for martingales with applications. *Electronic Journal of Probability*, 2015.
- [16] Ronald Aylmer Fisher. Design of experiments. *British Medical Journal*, 1(3923):554, 1936.
- [17] Peter Grünwald, Rianne de Heide, and Wouter M Koolen. Safe testing. *arXiv preprint arXiv:1906.07801*, 2019.
- [18] Vitor Hadad, David A Hirshberg, Ruohan Zhan, Stefan Wager, and Susan Athey. Confidence intervals for policy evaluation in adaptive experiments. *Proceedings of the national academy of sciences*, 118(15):e2014602118, 2021.
- [19] Dae Woong Ham, Iavor Bojinov, Michael Lindon, and Martin Tingley. Design-based confidence sequences for anytime-valid causal inference. *arXiv preprint arXiv:2210.08639*, 2022.
- [20] Sebastian Haneuse and Andrea Rotnitzky. Estimation of the effect of interventions that modify the received treatment. *Statistics in medicine*, 32(30):5260–5277, 2013.
- [21] Wassily Hoeffding. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.
- [22] Steven R Howard and Aaditya Ramdas. Sequential estimation of quantiles with applications to A/B testing and best-arm identification. *Bernoulli*, 28(3):1704–1728, 2022.
- [23] Steven R. Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform Chernoff bounds via nonnegative supermartingales. *Probability Surveys*, 17:257–317, 2020.
- [24] Steven R Howard, Aaditya Ramdas, Jon McAuliffe, and Jasjeet Sekhon. Time-uniform, nonparametric, nonasymptotic confidence sequences. *The Annals of Statistics*, 2021.
- [25] Audrey Huang, Liu Leqi, Zachary Lipton, and Kamyar Azizzadenesheli. Off-policy risk assessment in contextual bandits. *Advances in Neural Information Processing Systems*, 34:23714–23726, 2021.
- [26] Guido W Imbens and Donald B Rubin. *Causal inference in statistics, social, and biomedical sciences*. Cambridge University Press, 2015.
- [27] Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil’ UCB: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pages 423–439. PMLR, 2014.
- [28] Nikos Karampatziakis, Paul Mineiro, and Aaditya Ramdas. Off-policy confidence sequences. *International Conference on Machine Learning*, 2021.
- [29] Masahiro Kato, Takuya Ishihara, Junya Honda, and Yusuke Narita. Adaptive experimental design for efficient treatment effect estimation. *arXiv preprint arXiv:2002.05308*, 2020.
- [30] Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *The Journal of Machine Learning Research*, 17(1):1–42, 2016.
- [31] Edward H Kennedy. Nonparametric causal effects based on incremental propensity score interventions. *Journal of the American Statistical Association*, 114(526):645–656, 2019.
- [32] Edward H Kennedy. Semiparametric doubly robust targeted double machine learning: a review. *arXiv preprint arXiv:2203.06469*, 2022.
- [33] Koulik Khamaru, Yash Deshpande, Lester Mackey, and Martin J Wainwright. Near-optimal inference in adaptive linear regression. *arXiv preprint arXiv:2107.02266*, 2021.
- [34] Tze Leung Lai. On confidence sequences. *The Annals of Statistics*, 4(2):265–280, 1976.
- [35] John Langford and Tong Zhang. The epoch-greedy algorithm for contextual multi-armed bandits. *Advances in neural information processing systems*, 20(1):96–1, 2007.
- [36] Lihong Li, Wei Chu, John Langford, and Robert E Schapire. A contextual-bandit approach to personalized news article recommendation. In *Proceedings of the 19th International Conference on World Wide Web*, pages 661–670, 2010.
- [37] Pascal Massart. The tight constant in the Dvoretzky-Kiefer-Wolfowitz inequality. *The Annals of Probability*, pages 1269–1283, 1990.
- [38] Jerzy Neyman. On the application of probability theory to agricultural experiments, essay on principles, section 9. *Statistical Science*, pages 465–472, 1923/1990.
- [39] Aaditya Ramdas, Johannes Ruf, Martin Larsson, and Wouter Koolen. Admissible anytime-valid sequential inference must rely on nonnegative martingales. *arXiv preprint arXiv:2009.03167*, 2020.
- [40] Aaditya Ramdas, Peter Grünwald, Vladimir Vovk, and Glenn Shafer. Game-theoretic statistics and safe anytime-valid inference. *Statistical Science*, 2023.
- [41] Pratik Ramprasad, Yuantong Li, Zhuoran Yang, Zhaoran Wang, Will Wei Sun, and Guang Cheng. Online bootstrap inference for policy evaluation in reinforcement learning. *Journal of the American Statistical Association*, pages 1–14, 2022.
- [42] Herbert Robbins. Statistical methods related to the law of the iterated logarithm. *The Annals of Mathematical Statistics*, 41(5):1397–1409, 1970.

- [43] James Robins. A new approach to causal inference in mortality studies with a sustained exposure period—application to control of the healthy worker survivor effect. *Mathematical modelling*, 7(9-12):1393–1512, 1986.
- [44] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American Statistical Association*, 89(427):846–866, 1994.
- [45] Donald B Rubin. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of educational Psychology*, 66(5):688, 1974.
- [46] Glenn Shafer. Testing by betting: A strategy for statistical and scientific communication. *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, 184(2):407–431, 2021.
- [47] Glenn Shafer, Alexander Shen, Nikolai Vereshchagin, and Vladimir Vovk. Test martingales, Bayes factors and p-values. *Statistical Science*, 26(1):84–101, 2011.
- [48] Ye Shen, Hengrui Cai, and Rui Song. Doubly robust interval estimation for optimal policy evaluation in online learning. *arXiv preprint arXiv:2110.15501*, 2021.
- [49] Judith ter Schure and Peter Grünwald. All-in meta-analysis: breathing life into living systematic reviews. *arXiv preprint arXiv:2109.12141*, 2021.
- [50] Philip Thomas, Georgios Theodorou, and Mohammad Ghavamzadeh. High-confidence off-policy evaluation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 29, 2015.
- [51] Masatoshi Uehara, Chengchun Shi, and Nathan Kallus. A review of off-policy evaluation in reinforcement learning. *arXiv preprint arXiv:2212.06355*, 2022.
- [52] Mark J van der Laan and Sherri Rose. *Targeted learning: causal inference for observational and experimental data*. Springer Science & Business Media, 2011.
- [53] Jean Ville. Etude critique de la notion de collectif. *Bull. Amer. Math. Soc*, 45(11):824, 1939.
- [54] Vladimir Vovk and Ruodu Wang. E-values: Calibration, combination and applications. *The Annals of Statistics*, 49(3): 1736–1754, 2021.
- [55] Ruodu Wang and Aaditya Ramdas. False discovery rate control with e-values. *Journal of the Royal Statistical Society: Series B (Methodological)*, 2022.
- [56] Stanley L Warner. Randomized response: A survey technique for eliminating evasive answer bias. *Journal of the American Statistical Association*, 60(309):63–69, 1965.
- [57] Ian Waudby-Smith and Aaditya Ramdas. Confidence sequences for sampling without replacement. *Advances in Neural Information Processing Systems*, 33, 2020.
- [58] Ian Waudby-Smith and Aaditya Ramdas. Estimating means of bounded random variables by betting. *Journal of the Royal Statistical Society, Series B (to appear with discussion)*, 2023.
- [59] Ian Waudby-Smith, Zhiwei Steven Wu, and Aaditya Ramdas. Locally private nonparametric confidence intervals and sequences. *International Conference on Machine Learning*, 2023.
- [60] Ziyu Xu, Ruodu Wang, and Aaditya Ramdas. Post-selection inference for e-value based confidence intervals. *arXiv preprint arXiv:2203.12572*, 2022.
- [61] Ruohan Zhan, Vitor Hadad, David A Hirshberg, and Susan Athey. Off-policy evaluation via adaptive weighting with data from contextual bandits. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, pages 2125–2135, 2021.
- [62] Kelly Zhang, Lucas Janson, and Susan Murphy. Inference for batched bandits. *Advances in neural information processing systems*, 33:9818–9829, 2020.
- [63] Kelly Zhang, Lucas Janson, and Susan Murphy. Statistical inference with m-estimators on adaptively collected data. *Advances in neural information processing systems*, 34:7460–7471, 2021.
- [64] Julian Zimmert and Tor Lattimore. Connections between mirror descent, Thompson sampling and the information ratio. *Advances in Neural Information Processing Systems*, 32, 2019.
- [65] Julian Zimmert and Yevgeny Seldin. Tsallis-INF: An optimal algorithm for stochastic and adversarial bandits. *Journal of Machine Learning Research*, 22(28):1–49, 2021.

A PROOFS OF THE MAIN RESULTS

A.1 A technical lemma

Lemma 1. *Let Z and \widehat{Z} be \mathcal{H} -adapted processes such that $Z_t - \widehat{Z}_{t-1} \geq -1$ almost surely for all t . Denoting $\mu_t := \mathbb{E}(Z_t \mid \mathcal{H}_{t-1})$, we have that for any $(0, 1)$ -valued predictable process $(\lambda_t)_{t=1}^\infty$,*

$$M_t := \exp \left(\sum_{i=1}^t \lambda_i (Z_i - \mu_i) - \sum_{i=1}^t \psi_E(\lambda_i) (Z_i - \widehat{Z}_{i-1})^2 \right), \quad (52)$$

forms a test supermartingale, where $\psi_E(\lambda) := -\log(1 - \lambda) - \lambda$.

Above, \widehat{Z}_{t-1} is to be interpreted as an estimator of Z_t using the first $t - 1$ samples. Closely related lemmas have appeared in [15, 23, 24, 58], but those papers assumed $Z_t - \mu_t \geq -1$, which does not suffice for our purposes. What is somewhat surprising above is that we do not require a particular lower bound on Z_t or an upper bound on μ_t , as long as $Z_t - \widehat{Z}_{t-1} \geq -1$.

PROOF. First, note that $M_0 \equiv 1$ by construction, and M_t is always positive. It remains to show that M_t forms a supermartingale. Writing out the conditional expectation of M_t given \mathcal{H}_{t-1} , we have that

$$\mathbb{E}(M_t \mid \mathcal{H}_{t-1}) = M_{t-1} \underbrace{\mathbb{E} \left(\exp \left\{ \lambda_t (Z_t - \mu_t) - \psi_E(\lambda_t) (Z_t - \widehat{Z}_{t-1})^2 \right\} \mid \mathcal{H}_{t-1} \right)}_{(\dagger)}, \quad (53)$$

and hence it suffices to prove that $(\dagger) \leq 1$. Denote for the sake of succinctness,

$$Y_t := Z_t - \mu_t \quad \text{and} \quad \delta_t := \widehat{Z}_{t-1} - \mu_t,$$

and note that $\mathbb{E}(Y_t \mid \mathcal{H}_{t-1}) = 0$. Using the proof of Fan et al. [15, Proposition 4.1], we have that $\exp\{b\lambda - b^2\psi_E(\lambda)\} \leq 1 + b\lambda$ for any $\lambda \in [0, 1)$ and $b \geq -1$. Setting $b := Y_t - \delta_t = Z_t - \widehat{Z}_{t-1}$,

$$\begin{aligned} & \mathbb{E} \left[\exp \left\{ \lambda_t Y_t - (Y_t - \delta_t)^2 \psi_E(\lambda_t) \right\} \mid \mathcal{H}_{t-1} \right] \\ &= \mathbb{E} \left[\exp \left\{ \lambda_t (Y_t - \delta_t) - (Y_t - \delta_t)^2 \psi_E(\lambda_t) \right\} \mid \mathcal{H}_{t-1} \right] \exp(\lambda_t \delta_t) \\ &\leq \mathbb{E} [1 + (Y_t - \delta_t) \lambda_t \mid \mathcal{H}_{t-1}] \exp(\lambda_t \delta_t) \\ &= \mathbb{E} [1 - \delta_t \lambda_t \mid \mathcal{H}_{t-1}] \exp(\lambda_t \delta_t) \leq 1, \end{aligned}$$

where the last line follows from the fact that Y_t is conditionally mean zero and the inequality $1 - x \leq \exp(-x)$ for all $x \in \mathbb{R}$. This completes the proof. \square

A.2 Proof of Theorem 1

We will only derive the lower CS for ν , since the upper CS follows analogously. Consider the process $(M_t(\nu))_{t=1}^\infty$ given by

$$M_t(\nu) := \prod_{i=1}^t \left[1 + \lambda_i^L(\nu) \cdot (\phi_i^{(\text{DR-}\ell)} - \nu) \right]. \quad (54)$$

The proof proceeds in three steps, following the strategy of [24, 58] and [28]. In Step 1, we show that the pseudo-outcomes have conditional mean ν , i.e. $\mathbb{E}(\phi_t^{(\text{DR-}\ell)} \mid \mathcal{H}_{t-1}) = \mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1}) = \nu$. In Step 2, we use Step 1 to show that $M_t(\nu)$ forms a test martingale and apply Ville's inequality to it. In Step 3, we "invert" this test martingale to obtain the lower CS found in Theorem 1.

Step 1: Computing the conditional mean of the doubly robust pseudo-outcomes. Writing out the conditional expectation of $\phi_t^{(\text{DR-}\ell)}$, we have

$$\begin{aligned}
& \mathbb{E}[\phi_t^{(\text{DR-}\ell)} \mid \mathcal{H}_{t-1}] \\
&= \mathbb{E}(w_t R_t \mid \mathcal{H}_{t-1}) - \mathbb{E} \left\{ w_t \cdot \left(\widehat{r}_t(X_t; A_t) \wedge \frac{k_t}{w_t} \right) - \mathbb{E}_{a \sim \pi(\cdot \mid X_t)} \left(\widehat{r}_t(X_t; a) \wedge \frac{k_t}{w_t} \right) \mid \mathcal{H}_{t-1} \right\} \\
&= \mathbb{E}(w_t R_t \mid \mathcal{H}_{t-1}) \\
&= \int_{(x,a,r)} \frac{\pi(a \mid x)}{h_t(a \mid x)} r \cdot p_{R_t}(r \mid a, x, \mathcal{H}_{t-1}) h_t(a \mid x) p_{X_t}(x \mid \mathcal{H}_{t-1}) dx da dr \\
&= \int_{(x,a,r)} r \cdot p_{R_t}(r \mid a, x, \mathcal{H}_{t-1}) \pi(a \mid x) p_{X_t}(x \mid \mathcal{H}_{t-1}) dx da dr \\
&= \mathbb{E}_\pi(R_t \mid \mathcal{H}_{t-1}) = v.
\end{aligned}$$

Step 2: Showing that $M_t(v)$ forms a test martingale. First, note that $M_0 \equiv 1$ by construction. To show that M_t is nonnegative, notice that since $R_t, \widehat{r}_t \in [0, 1]$ almost surely, we have that $\phi_t^{(\text{DR-}\ell)} \geq -k_t$. Therefore, for any $v \in [0, 1]$,

$$\begin{aligned}
1 + \lambda_t^L(v) \cdot (\phi_t^{(\text{DR-}\ell)} - v) &\geq 1 + \lambda_t^L(v) \cdot (-k_t - v) \\
&> 1 + \frac{-k_t - v}{k_t + v} \quad (\text{since } \lambda_t^L(v) \in [0, (v + k_t)^{-1}]) \\
&= 0.
\end{aligned}$$

Lastly, it remains to show that $\mathbb{E}[M_t(v) \mid \mathcal{H}_{t-1}] = M_{t-1}(v)$. Writing out the conditional expectation of $M_t(v)$, we have

$$\begin{aligned}
\mathbb{E}[M_t(v) \mid \mathcal{H}_{t-1}] &= \mathbb{E} \left[M_{t-1}(v) \left\{ 1 + \lambda_t^L(v) \cdot (\phi_t^{(\text{DR-}\ell)} - v) \right\} \mid \mathcal{H}_{t-1} \right] \\
&= M_{t-1}(v) \cdot \mathbb{E} \left[\left\{ 1 + \lambda_t^L(v) \cdot (\phi_t^{(\text{DR-}\ell)} - v) \right\} \mid \mathcal{H}_{t-1} \right] \\
&= M_{t-1}(v) \cdot (1 + \lambda_t^L(v) \cdot 0) = M_{t-1}(v),
\end{aligned}$$

where the second line follows from the fact that M_{t-1}, λ_t^L are predictable, and the third line follows from Step 1. Therefore, by Ville's inequality for nonnegative supermartingales [53], we have

$$\mathbb{P} \left(\exists t \in \mathbb{N}, M_t(v) \geq \frac{1}{\alpha} \right) \leq \alpha. \tag{55}$$

Step 3: Inverting Ville's inequality to obtain a lower CS. Recall the lower boundary given by (12),

$$L_t^{\text{DR}} := \inf \left\{ v' \in [0, 1] : \prod_{i=1}^t \left[1 + \lambda_i^L(v') \cdot (\phi_i^{(\text{DR-}\ell)} - v') \right] < \frac{1}{\alpha} \right\}$$

and notice that if $v < L_t^{\text{DR}}$, then $M_t(v) \geq 1/\alpha$ by definition of L_t^{DR} . Consequently,

$$\mathbb{P}(\exists t \in \mathbb{N}, v < L_t^{\text{DR}}) \leq \mathbb{P} \left(\exists t \in \mathbb{N}, M_t(v) \geq \frac{1}{\alpha} \right) \leq \alpha.$$

Therefore, we have $\mathbb{P}(\forall t \in \mathbb{N}, v \geq L_t^{\text{DR}}) \geq 1 - \alpha$, so L_t^{DR} forms a lower $(1 - \alpha)$ -CS for v , which completes the proof. \square

A.3 Proof of Theorem 2

PROOF OF THEOREM 2. The proof proceeds in three steps, following the high level outline of the conjugate mixture method in [24]. First, we invoke Lemma 1 to derive a test supermartingale for each $\lambda \in (0, 1)$. Second, we mix over $\lambda \in (0, 1)$ using the truncated gamma density to obtain (26). Third and finally, we invert this test supermartingale to obtain a lower CS for \tilde{v}_t .

Step 1: Deriving a test supermartingale indexed by $\lambda \in (0, 1)$. Let $Z_t := \xi_t$ and $\widehat{Z}_{t-1} := \widehat{\xi}_{t-1}$ as in the setup of Theorem 2. First, notice that $\mathbb{E}(\xi_t | \mathcal{H}_{t-1}) = v_t$:

$$\mathbb{E}(\xi_t | \mathcal{H}_{t-1}) = \mathbb{E}(w_t R_t | \mathcal{H}_{t-1}) \quad (56)$$

$$= \int_{x,a,r} \frac{\pi(a | x)}{h_t(a | x)} r p_{R_t}(r | a, x, \mathcal{H}_{t-1}) h_t(a | x) p_{X_t}(x | \mathcal{H}_{t-1}) dx da dr. \quad (57)$$

Notice that $\xi_t - \widehat{\xi}_{t-1} \geq -1$, and hence by Lemma 1, we have that for any $\lambda \in (0, 1)$

$$M_t(\tilde{v}_t; \lambda) := \exp \{ \lambda S_t(\tilde{v}_t) - V_t \psi_E(\lambda) \}$$

forms a test supermartingale.

Step 2: Mixing over λ using the truncated gamma density. For any distribution F on $(0, 1)$,

$$M_t^{\text{EB}}(\tilde{v}_t) := \int_{\lambda \in (0,1)} M_t(\tilde{v}_t; \lambda) dF(\lambda) \quad (58)$$

forms a test supermartingale by Fubini's theorem. In particular, we will use the truncated gamma density $f(\lambda)$ given by

$$f(\lambda) = \frac{\rho^\rho e^{-\rho(1-\lambda)} (1-\lambda)^{\rho-1}}{\Gamma(\rho) - \Gamma(\rho, \rho)}, \quad (59)$$

as the mixing density. Writing out $M_t(v)$ using $dF(\lambda) := f(\lambda)d\lambda$, we have

$$\begin{aligned} M_t^{\text{EB}}(\tilde{v}_t) &:= \int_0^1 \exp \{ \lambda S_t(\tilde{v}_t) - V_t \psi_E(\lambda) \} f(\lambda) d\lambda \\ &= \int_0^1 \exp \{ \lambda S_t(\tilde{v}_t) - V_t \psi_E(\lambda) \} \frac{\rho^\rho e^{-\rho(1-\lambda)} (1-\lambda)^{\rho-1}}{\Gamma(\rho) - \Gamma(\rho, \rho)} d\lambda \\ &= \frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \int_0^1 \exp \{ \lambda (\rho + S_t + V_t) \} (1-\lambda)^{V_t + \rho - 1} d\lambda \\ &= \left(\frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \right) \left(\frac{1}{V_t + \rho} \right) \left(\frac{\Gamma(b)}{\Gamma(a)\Gamma(b-a)} \int_0^1 e^{zu} u^{a-1} (1-u)^{b-a-1} du \right) \Bigg|_{\substack{a=1 \\ b=V_t+\rho+1 \\ z=S_t+V_t+\rho}} \\ &= \left(\frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \right) \left(\frac{1}{V_t + \rho} \right) {}_1F_1(1, V_t + \rho + 1, S_t + V_t + \rho), \end{aligned}$$

which completes this step.

Step 3: Inverting the mixture test supermartingale to obtain (26). Similar to Step 3 of the proof of Theorem 1, we have that $\tilde{v}_t < L_t^{\text{EB}}$ if and only if $M_t(\tilde{v}_t) \geq 1/\alpha$, and hence by Ville's inequality for nonnegative supermartingales, we have that

$$\mathbb{P}(\exists t : \tilde{v}_t < L_t^{\text{EB}}) = \mathbb{P}(\exists t : M_t^{\text{EB}}(\tilde{v}_t) \geq 1/\alpha) \leq \alpha,$$

and hence L_t^{EB} forms a lower $(1-\alpha)$ -CS for \tilde{v}_t . This completes the proof. \square

Remark 5 (Writing (26) in terms of the lower incomplete gamma function). For readers familiar with Howard et al. [24, Proposition 9], we can rewrite (26) in terms of the lower incomplete gamma function via the identity ${}_1F_1(1, b, z) = (b-1)e^z z^{1-b}(\Gamma(b-1) - \Gamma(b-1, z))$, resulting in

$$\begin{aligned} & \left(\frac{\rho^\rho e^{-\rho}}{\Gamma(\rho) - \Gamma(\rho, \rho)} \right) \left(\frac{1}{v + \rho} \right) {}_1F_1(1, v + \rho + 1, s + v + \rho) \\ &= \left(\frac{\rho^\rho}{\Gamma(\rho)\gamma(\rho, \rho)} \right) \frac{\Gamma(v + \rho)\gamma(v + \rho, s + v + \rho)}{(s + v + \rho)^{v+\rho}} \exp(s + v), \end{aligned}$$

where $\gamma(\cdot, \cdot)$ is the lower regularized incomplete gamma function and $v = V_t$ and $s = S_t(\tilde{v}_t)$. This matches Howard et al. [24, Eq. (66)] when setting $c = 1$. The final representation above is real-valued after some complex terms are cancelled (in the case where $(s + v + \rho)$ is negative), but the representation in terms of ${}_1F_1(1, \cdot, \cdot)$ sidesteps this subtlety altogether, which is why we prefer to use it in Theorem 2.

A.4 Proof of Proposition 2

PROOF. Consider the process $M \equiv (M_t)_{t=1}^\infty$ given by

$$M_t := \exp \left\{ \sum_{i=1}^t \lambda_i \left(\xi_i - \frac{v}{k_i + 1} \right) - \sum_{i=1}^t \left(\xi_i - \tilde{\xi}_{i-1} \right)^2 \psi_E(\lambda_i) \right\}. \quad (60)$$

Then by Lemma 1, we have that M is a test supermartingale, and hence by Ville's inequality, $\mathbb{P}(\exists t : M_t \geq 1/\alpha) \leq \alpha$. Inverting this time-uniform concentration inequality, we have that with probability at least $(1 - \alpha)$ and for all $t \in \mathbb{N}$,

$$\begin{aligned} M_t < 1/\alpha &\iff \exp \left\{ \sum_{i=1}^t \lambda_i \left(\xi_i - \frac{v}{k_i + 1} \right) - \sum_{i=1}^t \left(\xi_i - \tilde{\xi}_{i-1} \right)^2 \psi_E(\lambda_i) \right\} < \frac{1}{\alpha} \\ &\iff \sum_{i=1}^t \lambda_i \left(\xi_i - \frac{v}{k_i + 1} \right) - \sum_{i=1}^t \left(\xi_i - \tilde{\xi}_{i-1} \right)^2 \psi_E(\lambda_i) < \log(1/\alpha) \\ &\iff \sum_{i=1}^t \lambda_i \xi_i - v \sum_{i=1}^t \frac{\lambda_i}{k_i + 1} - \sum_{i=1}^t \left(\xi_i - \tilde{\xi}_{i-1} \right)^2 \psi_E(\lambda_i) < \log(1/\alpha) \\ &\iff v > \frac{\sum_{i=1}^t \lambda_i \xi_i}{\sum_{i=1}^t \lambda_i / (k_i + 1)} - \frac{\log(1/\alpha) + \sum_{i=1}^t \left(\xi_i - \tilde{\xi}_{i-1} \right)^2 \psi_E(\lambda_i)}{\sum_{i=1}^t \lambda_i / (k_i + 1)}, \end{aligned}$$

which completes the proof. \square

A.5 Proof of Proposition 3

We will prove a more general result below for arbitrary $\eta, s > 1$, but the exact constants in Proposition 3 can be obtained by setting $\eta = e, s = 2$. By Lemma 1 combined with Howard et al. [23, Table 5, row 7], we have that $S_t(\tilde{v}_t)$ is a sub-gamma process with scale parameter $c = 1$, meaning for any $\lambda \in [0, 1)$

$$M_t^G(\lambda) := \exp \{ \lambda S_t(\tilde{v}_t) - V_t \psi_G(\lambda) \}, \quad (61)$$

where $\psi_G(\lambda) \equiv \psi_{G,1}(\lambda) = \frac{\lambda^2}{2(1-\lambda)}$. Define the following parameters:

$$\begin{aligned}\lambda_k &:= \psi^{-1}(\log(1/\alpha)/\eta^{k+1/2}), \text{ where } \psi_G^{-1}(a) := \frac{2}{1 + \sqrt{1 + 2/a}}, \\ \alpha_k &:= \frac{\alpha}{(k+1)^s \zeta(s)}, \text{ and} \\ b_{t,k} &:= \frac{V_t \psi_G(\lambda_k) + \log(1/\alpha_k)}{\lambda_k}.\end{aligned}$$

Taking a union bound over $k \in \mathbb{N}$, we have that

$$\mathbb{P}(\forall t \in \mathbb{N}, k \in \mathbb{N}, S_t(\bar{v}_t) \leq b_{t,k}) \geq 1 - \alpha.$$

It remains to find a deterministic upper bound on $b_{t,k}$ that does not depend on k . Indeed, similar to Howard et al. [24, Eq. (39)], we have that

$$b_{t,k} = A\left(\frac{\log(1/\alpha_k)}{\eta^{k+1/2}}\right) \underbrace{\left[\sqrt{\frac{\eta^{k+1/2}}{V_t}} + \sqrt{\frac{V_t}{\eta^{k+1/2}}} \right]}_{(\star)} \sqrt{\frac{\log(1/\alpha_k)V_t}{2}}, \quad (62)$$

where $A(a) := \sqrt{2a}/\psi_G^{-1}(a) = \sqrt{1+a/2} + \sqrt{a/2}$. Now, notice that (\star) is convex in V_t on $V_t \in [\eta^k, \eta^{k+1}]$, and hence (\star) is maximized at the endpoints η^k and η^{k+1} . Consequently, on the k^{th} epoch – i.e. when $\eta^k \leq V_t \leq \eta^{k+1}$ – we have that

$$\begin{aligned}b_{t,k} &\leq A\left(\frac{\log(1/\alpha_k)}{\eta^{k+1/2}}\right) \left[\eta^{-1/4} + \eta^{1/4} \right] \sqrt{\frac{\log(1/\alpha_k)V_t}{2}} \\ &\leq A\left(\frac{\sqrt{\eta} \log(1/\alpha_k)}{V_t}\right) \left[\eta^{-1/4} + \eta^{1/4} \right] \sqrt{\frac{\log(1/\alpha_k)V_t}{2}} \\ &= \left[\sqrt{1 + \frac{\sqrt{\eta} \log(1/\alpha_k)}{2V_t}} + \sqrt{\frac{\sqrt{\eta} \log(1/\alpha_k)}{2V_t}} \right] \cdot \left[\eta^{-1/4} + \eta^{1/4} \right] \cdot \sqrt{\frac{\log(1/\alpha_k)V_t}{2}}\end{aligned} \quad (63)$$

where the first inequality follows from our analysis of (\star) , the second follows from monotonicity of $A(\cdot)$ and the fact that $V_t \leq \eta^{k+1}$ on the k^{th} epoch, the third follows from the definition of $A(\cdot)$. Rewriting the final line (63) more succinctly, we have the following upper bound on $b_{t,k}$ for every $k \in \mathbb{N}$,

$$b_{t,k} \leq \sqrt{\gamma_1^2 \log(1/\alpha_k)V_t + \gamma_2^2 \log^2(1/\alpha_k) + \gamma_2 \log(1/\alpha_k)}, \quad (64)$$

$$\text{where } \gamma_1 := \frac{\eta^{1/4} + \eta^{-1/4}}{\sqrt{2}}, \text{ and } \gamma_2 := \frac{\sqrt{\eta} + 1}{2}. \quad (65)$$

Now, notice that the above bound only depends on k through $\log(1/\alpha_k)$. As such, we will upper bound $\log(1/\alpha_k)$ solely in terms of V_t and other constants. Indeed, on the k^{th} epoch, we have

$$\begin{aligned}\log(1/\alpha_k) &\equiv \log\left(\frac{(k+1)^s \zeta(s)}{\alpha}\right) = s \log(k+1) + \log\left(\frac{\zeta(s)}{\alpha}\right) \\ &\leq s \log(\log_\eta V_t + 1) + \log\left(\frac{\zeta(s)}{\alpha}\right) \equiv \ell_t\end{aligned} \quad (66)$$

where the final line used the upper bound $k \leq \overline{\log_\eta V_t}$ which follows because $\eta^k \leq V_t$ on the k^{th} epoch. Combining (64) and (66), we have that

$$b_{t,k} \leq \sqrt{\gamma_1^2 \ell_t V_t + \gamma_2^2 \ell_t^2} + \gamma_2 \ell_t, \text{ where } \ell_t := s \log \left(\log_\eta V_t + 1 \right) + \log \left(\frac{\zeta(s)}{\alpha} \right), \quad (67)$$

which no longer depends on k . Consequently, we have that

$$\begin{aligned} 1 - \alpha &\leq \mathbb{P} \left(\forall t \in \mathbb{N}, \sum_{i=1}^t \xi_i - \sum_{i=1}^t v_i \leq \sqrt{\gamma_1^2 \ell_t V_t + \gamma_2^2 \ell_t^2} + \gamma_2 \ell_t \right) \\ &= \mathbb{P} \left(\forall t \in \mathbb{N}, \tilde{v}_t \geq \underbrace{\frac{1}{t} \sum_{i=1}^t \xi_i - \frac{\sqrt{\gamma_1^2 \ell_t V_t + \gamma_2^2 \ell_t^2}}{t}}_{(\dagger)} - \frac{\gamma_2 \ell_t}{t} \right), \end{aligned}$$

and hence (\dagger) forms a lower $(1 - \alpha)$ -CS for \tilde{v}_t .

A.6 Proof of Theorem 3

We will prove a more general result below for arbitrary $\eta, s, \delta > 1$, but the exact constants in Proposition 3 can be obtained by setting $\eta = e$, and $s = \delta = 2$. The proof will proceed in five steps. First, we derive an exponential e -process – i.e. an adapted process upper-bounded by a test supermartingale – from $S_t(p) := \sum_{i=1}^t w_i \mathbb{1}(R_i \leq Q(p)) - tp$. Second, we apply Ville’s inequality to the aforementioned e -process to obtain a level- α linear boundary $b_t(p)$ on $S_t(p)$, meaning $\mathbb{P}(\exists t \in \mathbb{N} : S_t(p) \geq b_t(p)) \leq \alpha$. Third, we derive one level- $\alpha_{k,j}$ linear boundary for each $k \in \mathbb{N}, j \in \mathbb{Z}$ using the techniques of Step 2 so that $\sum_{k \in \mathbb{N}} \sum_{j \in \mathbb{Z}} \alpha_{k,j} \leq \alpha$ and take a union bound over all of them. Here, $k \in \mathbb{N}$ will index exponentially spaced epochs of time $t \in \mathbb{N}$, while $j \in \mathbb{Z}$ will index evenly-spaced log-odds of $p \in (0, 1)$. Fourth, we modify the boundaries derived in Step 3 to obtain a boundary that is uniform in both $t \in \mathbb{N}$ and in $p \in (0, 1)$. Fifth and finally, we obtain an analytic upper bound on the boundary derived in Step 4.

At several points throughout the proof, we will make use of various functions that depend on k and j . While we will define them as they are needed, we also list them here for reference.

$$W_t := \sum_{i=1}^t w_i^2, \quad (68a)$$

$$\alpha_{k,j} := \frac{\alpha}{(k+1)^s (|j| \vee 1)^s \zeta(s) (2\zeta(s) + 1)}, \quad (68b)$$

$$q(k, j) := \frac{1}{1 + \exp \{-2j\delta/\eta^{k/2}\}}, \quad (68c)$$

$$j(k, p) := \left\lceil \frac{\eta^{k/2} \text{logit}(p)}{2\delta} \right\rceil, \quad (68d)$$

$$\lambda(k, j) := \psi_{G,q(k,j)}^{-1}(\log(1/\alpha_{k,j})/\eta^{k+1/2}), \text{ where } \psi_{G,c}^{-1}(a) := \frac{2}{c + \sqrt{c^2 + 2/a}}, \text{ and} \quad (68e)$$

$$b_{t,k}(p) := \frac{W_t \psi_{G,p}(\lambda_{k,j}) + \log(1/\alpha_{k,j})}{\lambda_{k,j}}. \quad (68f)$$

Step 1: Deriving an e -process. Invoking Lemma 1 combined with Howard et al. [23, Table 5, row 7] we have that for any $p \in (0, 1)$, $S_t(p)$ is sub-gamma [23, 24] with variance process $V_t(p) := \sum_{i=1}^t (w_i \mathbb{1}\{R_i \leq Q(p)\})^2$ and scale $c = p$, meaning we have that for any $\lambda \in [0, 1/c)$,

$$M_t^G(\lambda; p) := \exp \{ \lambda S_t(p) - V_t(p) \psi_{G,p}(\lambda) \} \quad (69)$$

forms a test supermartingale. Now, since $V_t(p) \leq \sum_{i=1}^t w_i^2 \equiv W_t$ almost surely, we have that

$$E_t^G(\lambda; p) := \exp \{ \lambda S_t(p) - W_t \psi_{G,p}(\lambda) \} \leq M_t^G(\lambda; p) \quad (70)$$

forms an e -process — i.e. it is upper-bounded by a test supermartingale. This completes the first step of the proof.

Step 2: Applying Ville's inequality to $E_t^G(\lambda; p)$, yielding a time-uniform linear boundary. In Step 1, we showed that $E_t^G(\lambda; p)$ forms an e -process. By Ville's maximal inequality for nonnegative supermartingales [53], we have that

$$\mathbb{P}(\exists t \in \mathbb{N} : E_t^G(\lambda; p) \geq 1/\alpha) \leq \mathbb{P}(\exists t \in \mathbb{N} : M_t^G(\lambda; p) \geq 1/\alpha) \leq \alpha. \quad (71)$$

Now, we will rewrite the inequality $E_t^G(\lambda; p) \geq 1/\alpha$ slightly more conveniently so that we can derive a time-uniform concentration inequality for $S_t(p)$. Indeed,

$$\begin{aligned} E_t^G(\lambda; p) \geq 1/\alpha &\iff \lambda S_t(p) - W_t \psi_{G,p}(\lambda) \geq \log(1/\alpha) \\ &\iff S_t(p) \geq \underbrace{\frac{W_t \psi_{G,p}(\lambda) + \log(1/\alpha)}{\lambda}}_{b_t(p)}. \end{aligned}$$

In summary, we have the following time-uniform concentration inequality on $S_t(p)$ for any $p \in (0, 1)$, $\alpha \in (0, 1)$ and $\lambda \in [0, 1/p)$,

$$\mathbb{P}(\exists t \in \mathbb{N} : S_t(p) \geq b_t(p)) \leq \alpha, \text{ where } b_t(p) := \frac{W_t \psi_{G,p}(\lambda) + \log(1/\alpha)}{\lambda}, \quad (72)$$

which could also be written as a time-uniform high-probability upper bound on $S_t(p)$:

$$\mathbb{P}(\forall t \in \mathbb{N}, S_t(p) < b_t(p)) \geq 1 - \alpha. \quad (73)$$

Step 3: Union-bounding over infinitely many choices of λ , α , and p . In Step 2, we showed that $b_t(p)$ forms a time-uniform high-probability upper bound for $S_t(p)$. We will now take a union bound over a countably infinite two-dimensional grid of t and p . Concretely, for each $k \in \mathbb{N}$ and $j \in \mathbb{Z}$, recall $\alpha_{k,j}$, $q(k, j)$, and $\lambda(k, j)$ as in (68b), (68c), and (68e). The exact choices of $q(k, j)$ and $\lambda(k, j)$ will become relevant later. For now, note that by (72) from Step 2 combined with a union bound, we have that

$$\mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, j \in \mathbb{Z} : S_t(q(k, j)) \geq b_{t,k}(q(k, j))) \leq \sum_{k \in \mathbb{N}} \sum_{j \in \mathbb{Z}} \alpha_{k,j}, \quad (74)$$

$$\text{where } b_{t,k}(q(k, j)) := \frac{W_t \psi_{G,q(k,j)}(\lambda_{k,j}) + \log(1/\alpha_{k,j})}{\lambda_{k,j}}. \quad (75)$$

We will now show that $\sum_{k \in \mathbb{N}} \sum_{j \in \mathbb{Z}} \alpha_{k,j} = \alpha$ so that (74) holds with probability at most α . Indeed,

$$\begin{aligned} \sum_{k \in \mathbb{N}} \sum_{j \in \mathbb{Z}} \alpha_{k,j} &= \frac{\alpha}{\zeta(s)(2\zeta(s)+1)} \sum_{k \in \mathbb{N}} \frac{1}{(k+1)^s} \sum_{j \in \mathbb{Z}} \frac{1}{(|j| \vee 1)^s} \\ &= \frac{\alpha}{\zeta(s)(2\zeta(s)+1)} \underbrace{\sum_{k=0}^{\infty} \frac{1}{(k+1)^s}}_{=\zeta(s)} \underbrace{\left(1 + 2 \sum_{m=1}^{\infty} \frac{1}{m^s}\right)}_{=2\zeta(s)+1} \\ &= \alpha. \end{aligned}$$

Therefore, in summary, we have that

$$\mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, j \in \mathbb{Z} : S_t(q(k, j)) \geq b_{t,k}(q(k, j))) \leq \alpha. \quad (76)$$

Step 4: Removing dependence on $j \in \mathbb{Z}$ and obtaining p -uniformity. We will obtain a bound that is uniform in $p \in (0, 1)$ by replacing j with $j(k, p)$ – a function of both $k \in \mathbb{N}$ and $p \in (0, 1)$. For any $k \in \mathbb{N}$ and any $p \in (0, 1)$, define $j(k, p)$ as

$$j(k, p) := \left\lceil \frac{\eta^{k/2} \operatorname{logit}(p/(1-p))}{2\delta} \right\rceil. \quad (77)$$

Of course, $j(k, p) \in \mathbb{Z}$ is not unique. It is easy to check that $p \leq q(k, j(k, p))$, a fact that we will use shortly. Abusing notation slightly, let j_1, j_2, \dots denote the integers generated by $j(k, p)$ for every $k \in \mathbb{N}$ and $p \in (0, 1)$, and let $\mathcal{J} := \{j_1, j_2, \dots\} \subseteq \mathbb{Z}$ denote their image. Given this setup and applying (76) from Step 3, we have that

$$\begin{aligned} &\mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, p \in (0, 1) : S_t(q(k, j(k, p))) \geq b_{t,k}(q(k, j(k, p)))) \\ &= \mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, j \in \mathcal{J} : S_t(q(k, j)) \geq b_{t,k}(q(k, j))) \\ &\leq \mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, j \in \mathbb{Z} : S_t(q(k, j)) \geq b_{t,k}(q(k, j))) \\ &\leq \alpha, \end{aligned}$$

where the second line follows from the definition of \mathcal{J} , the third follows from the fact that $\mathcal{J} \subseteq \mathbb{Z}$, and the last follows from (76). In summary, we have the time- and p -uniform concentration inequality given by

$$\mathbb{P}(\exists t \in \mathbb{N}, k \in \mathbb{N}, p \in (0, 1) : S_t(q(k, j(k, p))) \geq b_{t,k}(q(k, j(k, p)))) \leq \alpha, \text{ or equivalently,} \quad (78)$$

$$\mathbb{P}(\forall t \in \mathbb{N}, k \in \mathbb{N}, p \in (0, 1), S_t(q(k, j(k, p))) < b_{t,k}(q(k, j(k, p)))) \geq 1 - \alpha \quad (79)$$

Step 5: Obtaining a time- and p -uniform upper bound on $S_t(p)$. While (79) is now written to be p -uniform, the quantity $b_{t,k}(q(k, j(k, p)))$ is only a high-probability upper bound on $S_t(q(k, j(k, p)))$, but what we need is a high-probability upper bound on $S_t(p)$. To this end, we use a similar technique to Howard and Ramdas [22] to bound the distance between $S_t(q(k, j(k, p)))$ and $S_t(p)$ for any $p \in (0, 1)$. Indeed, consider the following representation of $S_t(p)$ in terms of $S_t(q(k, j(k, p)))$:

$$\begin{aligned} S_t(p) &:= \sum_{i=1}^t w_i \mathbb{1}(R_i \leq Q(p)) - tp \\ &\leq \sum_{i=1}^t w_i \mathbb{1}(R_i \leq Q(q(k, j(k, p)))) - tp \\ &= S_t(q(k, j(k, p))) + t(q(k, j(k, p)) - p), \end{aligned}$$

where the first line follows by definition of $S_t(p)$, the second by monotonicity of $Q \mapsto \mathbb{1}(R_t \leq Q)$ and the fact that $p \leq p_{k,j(k,p)}$, and the third follows from the definition of $S_t(q(k, j(k, p)))$. Combining (79) with the above representation of $S_t(p)$, we have that

$$\mathbb{P} \left(\forall t \in \mathbb{N}, k \in \mathbb{N}, p \in (0, 1), S_t(p) < \underbrace{b_{t,k}(q(k, j(k, p)))}_{(i)} + \underbrace{t(q(k, j(k, p)) - p)}_{(ii)} \right) \geq 1 - \alpha, \quad (80)$$

where (i) $\equiv b_{t,k}(q(k, j(k, p)))$ is given by

$$b_{t,k}(q(k, j(k, p))) := \frac{W_t \psi_{G, q(k, j(k, p))}(\lambda_{k, j(k, p)}) + \log(1/\alpha_{k, j(k, p)})}{\lambda_{k, j(k, p)}}. \quad (81)$$

Step 5(i): Upper-bounding (i) without dependence on k . Applying Lemma 2 but with $j(k, p)$ in place of j , we have that for every $\eta^k \leq W_t \leq \eta^{k+1}$,

$$b_{t,k}(q(k, j(k, p))) \leq \sqrt{\gamma_1^2 \log(1/\alpha_{k, j(k, p)}) W_t + \gamma_2^2 q(k, j(k, p))^2 \log^2(1/\alpha_{k, j(k, p)})} + \gamma_2 q(k, j(k, p)) \log(1/\alpha_{k, j(k, p)}).$$

Now notice that the above upper bound depends on k solely through $q(k, j(k, p))$ and $\log(1/\alpha_{k, j(k, p)})$, each of which we will upper-bound independently of k . By Lemma 3, we have that

$$q(k, j(k, p)) \leq \bar{q}_t(p) \equiv \text{logit}^{-1} \left(\text{logit}(p) + 2\delta \sqrt{\frac{\eta}{W_t}} \right) \text{ for all } \eta^k \leq W_t \leq \eta^{k+1}, \quad (82)$$

so it remains to upper-bound $\log(1/\alpha_{k, j(k, p)})$. Recall the definition of $\alpha_{k, j}$ for any $k \in \mathbb{N}, j \in \mathbb{Z}$ given in (68b). Then we can write $\log(1/\alpha_{k, j(k, p)})$ as

$$\log(1/\alpha_{k, j(k, p)}) = \underbrace{s \log(k+1)}_{(\star k)} + \underbrace{2 \log(|j(k, p)| \vee 1)}_{(\star j)} + \log \zeta(s) + \log(2\zeta(s) + 1) + \log(1/\alpha), \quad (83)$$

and we observe that $(\star k)$ and $(\star j)$ are the only terms depending on k . Firstly, notice that $(\star j)$ can be upper bounded for every $\eta^k \leq W_t \leq \eta^{k+1}$ as

$$\begin{aligned} (\star j) &\equiv 2 \log(|j(k, p)| \vee 1) = 2 \log \left(\left\lceil \left\| \frac{\eta^{k/2} \text{logit}(p)}{2\delta} \right\| \vee 1 \right\rceil \right) \\ &\leq 2 \log \left(\left\lceil \left\| \frac{\sqrt{W_t} \text{logit}(p)}{2\delta} \right\| \vee 1 \right\rceil \right). \end{aligned}$$

Second, notice that we can easily upper-bound $(\star k)$ on epoch $\eta^k \leq W_t \leq \eta^{k+1}$ as

$$s \log(k+1) \leq s \log \left(\log_\eta W_t + 1 \right).$$

Therefore, we have the following upper-bound on $\log(1/\alpha_{k, j(k, p)})$:

$$\begin{aligned} \log(1/\alpha_{k, j}) &\leq s \log \left(\log_\eta W_t + 1 \right) + 2 \log \left(\left\lceil \left\| \frac{\sqrt{W_t} \text{logit}(p)}{2\delta} \right\| \vee 1 \right\rceil \right) + \log \zeta(s) + \log(2\zeta(s) + 1) + \log(1/\alpha), \\ &\equiv \ell_t(p), \end{aligned}$$

which no longer depends on k . In summary, we have that

$$b_{t,k}(q(k, j(k, p))) \leq \sqrt{\gamma_1^2 \ell_t(p) W_t + \gamma_2^2 \bar{q}_t(p)^2 \ell_t(p)^2} + \gamma_2 \bar{q}_t(p) \ell_t(p).$$

Step 5(ii): Upper-bounding (ii) without dependence on k . By Lemma 3, we have that $q(k, j(k, p)) \leq \bar{q}_t(p) \equiv \text{logit}^{-1} \left(\text{logit}(p) + 2\delta \sqrt{\frac{\eta}{W_t}} \right)$. Therefore, we can upper bound (ii) as

$$(ii) \equiv t(q(k, j(k, p)) - p) \leq t(\bar{q}_t(p) - p) \equiv t \left[\text{logit}^{-1} \left(\text{logit}(p) + 2\delta \sqrt{\frac{\eta}{W_t}} \right) - p \right], \quad (84)$$

where the final inequality no longer depends on k . In sum, with probability at least $1 - \alpha$,

$$\forall t \in \mathbb{N}, p \in (0, 1), S_t(p) < \sqrt{\gamma_1^2 \ell_t(p) W_t + \gamma_2^2 \bar{q}_t(p)^2 \ell_t(p)^2} + \gamma_2 \bar{q}_t(p) \ell_t(p) + t(\bar{q}_t(p) - p). \quad (85)$$

Lemma 2. For any $k \in \mathbb{N}$ and any $j \in \mathbb{Z}$, we have that for all $\eta^k \leq W_t \leq \eta^{k+1}$,

$$b_{t,k}(q(k, j)) \leq \sqrt{\gamma_1^2 \log(1/\alpha_{k,j}) W_t + \gamma_2^2 q(k, j)^2 \log^2(1/\alpha_{k,j})} + q(k, j) \gamma_2 \log(1/\alpha_{k,j}), \quad (86)$$

where γ_1, γ_2 are constants defined as

$$\gamma_1 := \frac{\eta^{1/4} + \eta^{-1/4}}{\sqrt{2}} \text{ and } \gamma_2 := \frac{\sqrt{\eta} + 1}{2}. \quad (87)$$

PROOF. Recall the chosen value of $\lambda_{k,j}$ given in (68e),

$$\lambda(k, j) := \psi_{G,q(k,j)}^{-1}(\log(1/\alpha_{k,j})/\eta^{k+1/2}), \text{ where } \psi_{G,c}^{-1}(a) := \frac{2}{c + \sqrt{c^2 + 2/a}} \quad (88)$$

Similar to Howard et al. [24, Eq. (39)], some algebra will reveal that for any $t, k \in \mathbb{N}, j \in \mathbb{Z}$, we have that

$$b_{t,k}(q(k, j)) = A_{q(k,j)} \left(\frac{\log(1/\alpha_{k,j})}{\eta^{k+1/2}} \right) \underbrace{\left[\sqrt{\frac{\eta^{k+1/2}}{W_t}} + \sqrt{\frac{W_t}{\eta^{k+1/2}}} \right]}_{(\star)} \sqrt{\frac{\log(1/\alpha_{k,j}) W_t}{2}},$$

where $A_c(a) := \sqrt{2a}/\psi_{G,c}^{-1}(a) = \sqrt{1 + c^2 a/2} + c\sqrt{a/2}$. Now, notice that the second derivative of (\star) with respect to W_t is positive on $W_t \in [\eta^k, \eta^{k+1}]$, and hence (\star) is convex in W_t . As such, for every $W_t \in [\eta^k, \eta^{k+1}]$ – i.e. the k^{th} epoch – we have that (\star) is maximized at the endpoints $W_t = \eta^k$ and $W_t = \eta^{k+1}$, and we thus have the following upper bound on $b_{t,k}(q(k, j))$ on the k^{th} epoch:

$$b_{t,k}(q(k, j)) \leq A_{q(k,j)} \left(\frac{\log(1/\alpha_{k,j})}{\eta^{k+1/2}} \right) \left[\eta^{1/4} + \eta^{-1/4} \right] \sqrt{\frac{\log(1/\alpha_{k,j}) W_t}{2}}. \quad (89)$$

Furthermore, since $W_t/\sqrt{\eta} \leq \eta^{k+1/2}$ on the k^{th} epoch, we also have that

$$A_{q(k,j)} \left(\frac{\log(1/\alpha_{k,j})}{\eta^{k+1/2}} \right) \leq A_{q(k,j)} \left(\frac{\sqrt{\eta} \log(1/\alpha_{k,j})}{W_t} \right) \text{ for all } \eta^k \leq W_t \leq \eta^{k+1}. \quad (90)$$

Putting (89) and (90) together, we have that for all $\eta^k \leq W_t \leq \eta^{k+1}$,

$$\begin{aligned} b_{t,k}(q(k, j)) &\leq \frac{\eta^{1/4} + \eta^{-1/4}}{\sqrt{2}} \left(\sqrt{\log(1/\alpha_{k,j}) W_t + \frac{\sqrt{\eta} q(k, j)^2 \log^2(1/\alpha_{k,j})}{2}} + q(k, j) \frac{\eta^{1/4} \log(1/\alpha_{k,j})}{\sqrt{2}} \right) \\ &= \sqrt{\gamma_1^2 \log(1/\alpha_{k,j}) W_t + \gamma_2^2 q(k, j)^2 \log^2(1/\alpha_{k,j})} + q(k, j) \gamma_2 \log(1/\alpha_{k,j}), \end{aligned} \quad (91)$$

where γ_1, γ_2 are constants defined in (87). This completes the proof of Lemma 2. \square

Lemma 3. Define $\bar{q}_t(p)$ as

$$\bar{q}_t(p) := \text{logit}^{-1} \left(\text{logit}(p) + 2\delta \sqrt{\frac{\eta}{W_t}} \right). \quad (92)$$

For all $\eta^k \leq W_t \leq \eta^{k+1}$, we have that $q(k, j(k, p)) \leq \bar{q}_t(p)$.

PROOF. The result follows by definition of $q(k, j(k, p))$. Indeed, we have that for all $\eta^k \leq W_t \leq \eta^{k+1}$,

$$\begin{aligned} q(k, j(k, p)) &:= \frac{1}{1 + \exp \left\{ -2j(k, p)\delta/\eta^{k/2} \right\}} \\ &= \left(1 + \exp \left\{ -2 \left[\frac{\eta^{k/2} \text{logit}(p)}{2\delta} \right] \delta/\eta^{k/2} \right\} \right)^{-1} \\ &\leq \left(1 + \exp \left\{ -2 \left(\frac{\eta^{k/2} \text{logit}(p)}{2\delta} + 1 \right) \delta/\eta^{k/2} \right\} \right)^{-1} \\ &= \left(1 + \exp \left\{ -(\text{logit}(p) - 2\delta/\eta^{k/2}) \right\} \right)^{-1} \\ &= \text{logit}^{-1}(\text{logit}(p) + 2\delta/\eta^{k/2}) \\ &\leq \text{logit}^{-1} \left(\text{logit}(p) + 2\delta \sqrt{\frac{\eta}{W_t}} \right), \end{aligned}$$

which completes the proof. \square

B A CAUSAL VIEW OF CONTEXTUAL BANDITS VIA POTENTIAL OUTCOMES

In Section 1, we discussed how the OPE problem can be interpreted as asking a *counterfactual* question, such as “how would the rewards have been, had we used a different policy π than the logging policy h that collected the data?”. While it is somewhat reasonable to think about the functional $v_t = \mathbb{E}_\pi(R_t | X_1^{t-1})$ in a counterfactual sense, the Neyman-Rubin potential outcomes framework was designed for the rigorous study of precisely these types of causal questions [38, 45]. In this section, we will define a target *causal* functional v_t^* in terms of potential outcomes, and outline the identification assumptions under which v_t^* is equal to v_t (and hence, the conditions under which our CSs can be interpreted as covering the causal quantity $\tilde{v}_t^* := \frac{1}{t} \sum_{i=1}^t v_i^*$). We emphasize that these identification assumptions are not required for the CSs to be useful or sensible – indeed, \tilde{v}_t^* is still an interpretable statistical quantity that we may wish to estimate – but they cannot otherwise be said to cover a causal functional defined in terms of potential outcomes.

Making our setup precise, we posit that for each time t , there is one potential outcome $R_t(a)$ for every action $a \in \mathcal{A}$. The functional we are ultimately interested in estimating is the *conditional mean potential outcome reward under the policy π* , i.e.

$$v_t^* \equiv \mathbb{E}_\pi(R_t(G) | X_1^{t-1}) := \mathbb{E} \left\{ \mathbb{E}_{G \sim \pi(\cdot | X_t)}(R_t(G) | X_t, X_1^{t-1}) | X_1^{t-1} \right\} \quad (93)$$

$$= \int_{\mathcal{A} \times \mathcal{X}} \mathbb{E}(R_t(g) | G = g, X_t = x, X_1^{t-1}) \pi(g | x) p_{X_t}(x | X_1^{t-1}) dg dx. \quad (94)$$

In words, v_t is the average of the potential outcomes $\{R_t(g)\}_{g \in \mathcal{A}}$ conditional on X_1^{t-1} with respect to the distribution $\pi(\cdot | X_t)$. We use g and G in place of a and A to avoid confusion between the actual (random) action A_t played according to the logging policy $h_t(\cdot | X_t)$ and the hypothetical (random) action G . Without further assumptions, however, the counterfactual quantity v_t^* is not necessarily identified, meaning it cannot necessarily be written as a functional of the distribution of the observed data $(X_t, A_t, R_t)_{t=1}^\infty$. This is simply due to the fact that the potential outcome $R_t(g)$ is not

directly observable from $(X_t, A_t, R_t)_{t=1}^\infty$. To remedy this, consider the following causal identification assumptions for every subject t ,

(IA1): **Consistency**: $A_t = a \implies R_t(a) = R_t$ for every $a \in \mathcal{A}$ with positive π -density,

(IA2): **Sequential exchangeability**: $A_t \perp\!\!\!\perp R_t(a) \mid X_1^t$, and

(IA3): **Positivity**: $\pi \ll h_t$, meaning $h_t(A_t \mid X_t) = 0 \implies \pi(A_t \mid X_t) = 0$ almost surely.

Notice that in the contextual bandit setup, IA2 and IA3 are known to hold *by design*, while IA1 is more subtle (e.g. IA1 may not hold even in a randomized experiment due to interference between subjects, such as in a vaccine trial). Nevertheless, with IA1, IA2, and IA3 in mind, we are ready to state the main identification result of this section.

Lemma 4. *Under causal assumptions IA1, IA2, and IA3, we have that*

$$v_t^* = v_t, \text{ and hence } \tilde{v}_t^* := \frac{1}{t} \sum_{i=1}^t v_i^* = \frac{1}{t} \sum_{i=1}^t v_i =: \tilde{v}_t. \quad (95)$$

In other words, the counterfactual conditional mean v_t^ can be represented as a function of the distribution of observed data $(X_t, A_t, R_t)_{t=1}^\infty$, and that representation is given by v_t .*

PROOF. The proof is an exercise in causal identification and essentially follows that of Kennedy [31, Theorem 1] and Robins' g -formula [43], but we nevertheless provide a derivation here for completeness.

Writing out the definition of v_t^* , we have

$$v_t^* := \int_{\mathcal{A} \times \mathcal{X}} \mathbb{E}(R_t(g) \mid G = g, X_t = x, X_1^{t-1}) \pi(g \mid x_t) p_{X_t}(x \mid X_1^{t-1}) dg dx \quad (96)$$

$$= \int_{\mathcal{A} \times \mathcal{X}} \mathbb{E}(R_t(g) \mid G = g, A_t = g, X_t = x, X_1^{t-1}) \pi(g \mid x) p_{X_t}(x \mid X_1^{t-1}) dg dx \quad (97)$$

$$= \int_{\mathcal{A} \times \mathcal{X}} \mathbb{E}(R_t(g) \mid A_t = g, X_t = x, X_1^{t-1}) \pi(g \mid x) p_{X_t}(x \mid X_1^{t-1}) dg dx \quad (98)$$

$$= \int_{\mathcal{A} \times \mathcal{X}} \mathbb{E}(R_t \mid A_t = g, X_t = x, X_1^{t-1}) \pi(g \mid x) p_{X_t}(x \mid X_1^{t-1}) dg dx \quad (99)$$

$$= \mathbb{E} \left\{ \mathbb{E}_{A_t \sim \pi(\cdot \mid X_t)}(R_t \mid X_t, X_1^{t-1}) \mid X_1^{t-1} \right\} \quad (100)$$

$$\equiv \mathbb{E}_\pi(R_t \mid X_1^{t-1}) =: v_t, \quad (101)$$

where (97) follows from IA2 (sequential exchangeability), (98) follows from the fact that $R_t(G) \perp\!\!\!\perp G \mid X_1^t$ (by definition), and (99) follows from IA1 (consistency). Throughout, we implicitly used IA3 (positivity) so that the outer integral is well-defined. That is, we conditioned on $A_t = g$ at several points, which implicitly leaves us with a factor of $\pi(g \mid x)/h_t(g \mid x)$ – positivity ensures that this quantity is well-defined with probability one. This completes the proof of Lemma 4. \square

Remark 6 (On the relationship between OPE and stochastic intervention effect estimation). It is no surprise that the proof of Lemma 4 follows Robins [43] and Kennedy [31] who study *stochastic intervention effects* in causal inference. Indeed, OPE and estimation of stochastic interventions are two different framings of essentially the same problem, and use the same importance-weighted and doubly robust estimators. The main differences between these fields lie in their emphases: the former is focused on adaptive experiments where logging policies are data-adaptive and *known*, whereas the latter typically places more emphasis on potential outcomes and causal identification, observational studies (i.e. where logging policies must be *estimated*), and more complex causal functionals, such as those of Haneuse and Rotnitzky [20] and Kennedy [31]. Of course, these are

incomplete characterizations made with broad strokes; for a more detailed summary of prior work in stochastic interventions, see Kennedy [31, Section 1].

Remark 7 (Implications for design-based confidence sequences). As an alternative to estimating treatment effects in superpopulations, one can opt to consider a so-called “design-based” approach to causal inference where the potential outcomes of all individuals are conditioned on, and confidence intervals are constructed for the *sample* average treatment effect (SATE) given by $\text{SATE}_t := \frac{1}{t} \sum_{i=1}^t (R_i(1) - R_i(0))$ where $R_i(a)$ is subject i 's potential outcome under treatment $a \in \{0, 1\}$. Here, the resulting confidence intervals cover the SATE with high probability, where the probability is taken with respect to the randomness in the treatment assignment mechanism only. The design-based approach goes back to Fisher and has a deep and extensive literature [16, 26, 38], and more recent work has constructed nonasymptotic CSs for the time-varying effect $(\text{SATE}_t)_{t=1}^\infty$ in Howard et al. [24, Section 4.2] and asymptotic ones in Ham et al. [19]. For a more comprehensive literature review, we direct readers to Abadie et al. [1] and Ham et al. [19] as well as the references therein.

We simply remark here that the results of Section 3 simultaneously apply to the design-based *and* superpopulation settings as immediate corollaries. Indeed, in the stochastic (non-design-based) setting for binary experiments and under the causal identification assumptions IA1–IA3, we have that Lemma 4 yields

$$\Delta_t^\star := \frac{1}{t} \sum_{i=1}^t \mathbb{E} [R_i(1) - R_i(0)] = \frac{1}{t} \sum_{i=1}^t [\mathbb{E}(R_i | A_i = 1) - \mathbb{E}(R_i | A_i = 0)] =: \Delta_t. \quad (102)$$

Now, to recover CSs for $(\text{SATE}_t)_{t=1}^\infty$ we simply condition on $(R_t(1), R_t(0), X_t)_{t=1}^\infty$ so that $(A_t)_{t=1}^\infty$ are the only non-degenerate random variables here. The techniques for time-varying treatment effects described in Section 3.1 and Section 3.2, yield a $(1 - \alpha)$ -CS $[L_t, U_t]_{t=1}^\infty$ for $(\Delta_t)_{t=1}^\infty \equiv (\Delta_t^\star)_{t=1}^\infty$ and hence for $(\text{SATE}_t)_{t=1}^\infty$. Going further, when instantiated for the design-based setting, our CSs substantially improve on Howard et al. [24, Section 4.2], both practically and theoretically. Indeed, as discussed in Ham et al. [19, Section 3.2], one of the drawbacks of existing nonasymptotic CSs in the literature is that the minimal propensity score – i.e. $p_{\min} := \text{ess inf}_{t,a,x} \mathbb{P}(A_t = a | \mathcal{H}_{t-1})$ – must be specified in advance, and the downstream CSs always scale with p_{\min}^{-1} . However, as we have emphasized throughout this paper, beginning with desideratum 5 in Section 1.2, *none* of our CSs suffer from this limitation.

Simultaneously, if we consider the superpopulation setting where $\mathbb{E}(R_t(1)) - \mathbb{E}(R_t(0)) = \delta$ for all $t \geq 1$ and for some $\delta \in [-1, 1]$, then under identification assumptions IA1–IA3, the same CS $[L_t, U_t]_{t=1}^\infty$ also covers δ by Lemma 4. In this way, our time-varying CSs simultaneously handle the stationary superpopulation setting where treatment effects do not change over time, as well as the design-based setting where all potential outcomes are conditioned on, since these are both special cases of the time-varying stochastic setting considered in Section 3.