OPTIMAL TESTS FOLLOWING SEQUENTIAL EXPERIMENTS

KARUN ADUSUMILLI†

Abstract. We develop methods for inference following sequential experiments by studying the asymptotic properties of tests. We find that the large-sample power of any test can be matched by a test in a suitable limit-experiment involving Gaussian diffusions. This establishes that a fixed set of statistics are asymptotically sufficient for testing; these are the number of times each treatment has been sampled, and the final value of the score/efficient influence function process for each treatment. We also derive asymptotically optimal tests under various conditions and apply these findings to three types of sequential experiments: costly sampling, group sequential trials and bandit-experiments.

This version: September 17, 2024

I would like to thank Kei Hirano and Jack Porter for insightful discussions that stimulated this research. Thanks also to seminar participants at multiple universities and conferences for helpful comments.

†Department of Economics, University of Pennsylvania.

1. Introduction

Recent years have seen tremendous advances in the theory and application of sequential/adaptive experiments. Such experiments are now used being in a wide variety of fields, ranging from online advertising (Russo et al., 2017), to dynamic pricing (Ferreira et al., 2018), drug discovery (Wassmer and Brannath, 2016), public health (Athey et al., 2021), and economic interventions (Kasy and Sautmann, 2019). Compared to traditional randomized trials, these experiments allow one to target and achieve a more efficient balance of welfare, ethical, and economic considerations. In fact, starting from the Critical Path Initiative in 2006, the FDA has actively promoted the use of sequential designs in clinical trials for reducing trial costs and risks for participants (US Food and Drug Admin., 2019). For instance, group-sequential designs, wherein researchers conduct interim analyses at predetermined stages of the experiment, are now routinely used in clinical trials: if the analysis suggests a significant positive or negative effect from the treatment, the trial may be stopped early. Other examples of sequential experiments include bandit experiments (Lattimore and Szepesvári, 2020), best-arm identification (Russo and Van Roy, 2016) and costly sampling (Adusumilli, 2022), among many others.

Although hypothesis testing is not always the primary goal of sequential experiments, one may still desire to conduct a hypothesis test after the experiment is completed. For example, a pharmaceutical company may conduct an adaptive trial for drug testing with the explicit goal of maximizing welfare or minimizing costs, but may nevertheless be required to test the null hypothesis of a zero average treatment effect for the drug after the trial. Despite the practical importance of such inferential methods, there are currently few results characterizing optimal tests, or even identifying which sample statistics to use when conducting tests after sequential experiments. This paper aims to fill this gap.

To this end, we follow the standard approach in econometrics and statistics (see, e.g., Van der Vaart, 2000, Chapter 14) of studying the properties of various candidate tests by characterizing their power against local alternatives, also known as Pitman alternatives. These are alternatives that converge to the null at the parametric, i.e., $1/\sqrt{n}$ rate, leading to non-trivial asymptotic power. Here, n is typically the sample size, although it can have other interpretations in experiments which are open-ended, see Section 2 for a discussion. The main finding of this paper is that the asymptotic power function of any test can be matched by that of a test in a limit experiment where one observes a Gaussian process for each treatment, and the aim is to conduct inference on the drifts of the Gaussian processes.

As a by-product of this equivalence, we show that the power function of any candidate test (which in general depends on the entire data collected) can be matched asymptotically by one that only depends on a finite set of sufficient statistics. In the most general scenario, the sufficient statistics are the number of times each treatment has been sampled by the end of the experiment, along with final value of the score (for parametric models) or efficient influence function (for non-parametric models) process for each treatment. However, even these statistics can be further reduced under additional assumptions on the sampling and stopping rules. Our results thus show that a substantial dimension reduction is possible, and only a few statistics are relevant for conducting tests.

Furthermore, we characterize the optimal tests in the limit experiment. We then show that finite sample analogues of these are asymptotically optimal under the original sequential experiment. Our results can also be used to compute the power envelope, i.e., an upper bound on the asymptotic power function of any test. Although a uniformly most powerful test in the limit experiment may not always exist, some positive results are obtained for testing linear combinations under unbiasedness, α-spending restrictions or conditional size constraints. Alternatively, one may impose less stringent criteria for optimality, like weighted average power, and we show how to compute optimal tests under such criteria as well.

We provide two new asymptotic representation theorems (ARTs) for formalizing the equivalence of tests between the original and limit experiments. The first applies to 'stopping-time experiments', where the sampling rule is fixed beforehand but the stopping rule (which describes when the experiment is to be terminated) is fully adaptive (i.e., it can be updated after every new observation). Our second ART allows for the sampling rule to be adaptive as well, but we require the sampling and stopping decision to be updated only a finite number of times, after observing the data in batches. While constraining attention to batched experiments is undoubtedly a limitation, practical considerations often necessitate conducting sequential experiments in batches anyway. Also, as shown in Adusumilli (2021), a fully adaptive experiment can often be approximated by a batched experiment with a sufficiently large number of batches. Our second ART builds on, and extends, the recent work of Hirano and Porter (2023) on asymptotic representations. We refer to Sections 1.1 and 4.1 for a detailed comparison. Importantly, and in contrast to Hirano and Porter (2023), our analysis covers both parametric and non-parametric settings.

We apply our results to three important examples of sequential experiments: costly sampling, group sequential trials and bandit experiments. We suggest new inferential procedures for these experiments that are asymptotically optimal under various scenarios such as unbiasedness, α-spending etc.

1.1. **Related literature.** Despite the vast amount of work on the development of sequential learning algorithms, the literature on inference following the use of such algorithms is relatively sparse. One approach gaining some popularity in computer-science is called 'any-time inference'. Here, one seeks to construct tests and confidence intervals that are correctly sized no matter how, or when, the experiment is stopped. We refer to Ramdas et al. (2022) for a survey and to Grünwald et al. (2020), Howard et al. (2021), Johari et al. (2022) for some recent contributions. The uniform-in-time size constraint is a stringent requirement, and this comes at the expense of lower power than could be achieved otherwise. By contrast, our focus in this paper is on classical notions of testing, where size control is only achieved when the experimental protocol, i.e., the specific sampling rule and stopping time, is followed exactly. In essence, this requires the decision maker to pre-register the experiment and fully commit to the protocol. We believe this is valid assumption in most applications; adaptive experiments are usually constructed with the explicit goal of welfare maximization, so there is little incentive to deviate from the protocol as long as the preferences of the experimenter and the end-user of the experiment are aligned (e.g., in the case of online marketplaces they would be the same entity). In other situations, pre-registration of the experimental design is usually mandatory, see, e.g., the FDA guidance on sequential designs (US Food and Drug Admin., 2019, Section III.C).

There are other recent papers which propose inferential methods under the 'classical' hypothesis-testing framework. Zhang et al. (2020) and Hadad et al. (2021) suggest asymptotically normal tests for some specific classes of sequential experiments. These tests are based on re-weighing the observations. There are also a number of methods for group sequential and linear boundary designs commonly used in clinical trials, see Hall (2013) for a review. However, neither of them are optimal even within their specific use cases.

Finally, in prior and closely related work, Hirano and Porter (2023) obtain an Asymptotic Representation Theorem (ART) for batched sequential experiments and apply this to testing. The ART of Hirano and Porter (2023) is a lot more general than ours, e.g., it can be used to determine optimal conditional tests given outcomes from previous stages. However, this generality comes at a price as the state variables increase linearly with the number of batches. Here, we build on and extend these results to show that only a fixed number of sufficient statistics are needed to match the *unconditional* asymptotic power of any test, irrespective of the number of batches (our results also apply to asymptotic power *conditional* on stopping times). We also derive a number of additional results that are new to this literature: First, our ART for stopping-time experiments applies to fully adaptive experiments (this result is not based on Hirano and Porter, 2023; rather, it makes use of a representation theorem for stopping times due to Le Cam, 1979). Second, our analysis covers non-parametric models, which is important for applications. Third, we characterize the properties of optimal tests in a number of different scenarios, e.g., for testing linear combinations of parameters, or under unbiasedness and α-spending requirements. This is useful as UMP tests do not generally exist otherwise.

As noted earlier, this paper employs the local asymptotic power criterion to rank tests. This criterion naturally leads to 'diffusion asymptotics', where the limit experiment consists of Gaussian diffusions. Diffusion asymptotics were first introduced by Wager and Xu (2021) and Fan and Glynn (2021) to study the properties of a class of sequential algorithms. In previous work (Adusumilli, 2021), this author demonstrated some asymptotic equivalence results for comparing the Bayes and minimax risk of bandit experiments. Here, we apply the techniques devised in these papers to study inference.

1.2. **Examples.** Before describing our procedures, it can be instructive to consider some examples of sequential experiments.

1.2.1. *Costly sampling.* Consider a sequential experiment in which sampling is costly, and the aim is to select the best of two possible treatments. Previous work by this author (Adusumilli, 2022) showed that the minimax optimal strategy in this setting involves a fixed sampling rule (the Neyman allocation) and stopping when the average difference in treatment outcomes multiplied by the number of observations exceeds a specific threshold. In fact, the stopping rule here has the same form as the Sequential Probability Ratio Test (SPRT) procedure of Wald (1947), even though the latter is motivated by very different considerations. SPRT is itself a special case of 'fully sequential linear boundary designs', as discussed, e.g., in Whitehead (1997). Typically these procedures recommend sampling the two treatments in equal proportions instead of the Neyman allocation. In Section 5, we show that for 'horizontal fully sequential boundary designs' with any fixed sampling rule (including, but not restricted to, the Neyman allocation), the most powerful unbiased test for treatment effects depends only on the stopping time and rejects when it is below a specific threshold.

1.2.2. *Group sequential trials.* In many applications, it is not feasible to employ continuous-time monitoring designs that update the decision rule after each observation. Instead, one may wish to stop the experiment only at a limited number of pre-specified times. Such designs are known as group-sequential trials, see Wassmer and Brannath (2016) for a textbook treatment. Recently, these experiments have become very popular for conducting clinical trials; they have been used, e.g., to test the efficacy of Coronavirus vaccines (Baden et al., 2021). While a number of methods have been proposed for inference following these experiments, as reviewed, e.g., in Hall (2013), it is not clear which, if any, are optimal. In Section 5, we derive optimal non-parametric tests and confidence intervals for such designs under α-spending and conditional size criteria (see, Section 2.7).

1.2.3. *Multi-armed bandit experiments.* In the previous two examples, the decision maker could choose when to end the experiment, but the sampling strategy was fixed beforehand. In many experiments however, the sampling rule can also be modified based on the information revealed from past data. Multi-armed bandit experiments are a canonical example of these. Previously, Hirano and Porter (2023) derived asymptotic power envelopes for any test following batched parametric bandit experiments. In this paper, we refine the results of Hirano and Porter (2023) further by showing that only a finite number of sufficient statistics are needed for testing, irrespective of the number of batches. Our results apply to non-parametric models as well.

2. Optimal tests in experiments involving stopping times

In this section we study the asymptotic properties of tests for parametric stopping-time experiments, i.e., sequential experiments that involve a pre-determined stopping time.

2.1. **An empirical illustration.** We begin by demonstrating our methodology through an empirical application involving one-armed bandits, which we also use as a recurring example throughout our discussion of the general framework. This setup is inspired by a

Google Analytics example that describes the algorithms employed for website optimization.[1] The scenario is as follows: you currently own a website with a known conversion rate, θ_0, generating a Bernoulli(θ_0) distribution of outcomes.[2] You wish to test a new version of the website with an unknown conversion rate θ. To identify the superior variant, you conduct an adaptive experiment utilizing Thompson Sampling (TS).[3]

The TS algorithm starts with a prior belief over the unknown θ. Since the outcome distribution is Bernoulli, it is standard practice to set a beta-prior over θ. As samples are collected from the unknown website variant, the prior is revised through Bayesian updating. The TS algorithm then distributes traffic between the websites based on the posterior probability that each one is the best. The trial ends after n rounds of experimentation, where n is pre-determined. Following the end of the experiment you are interested in testing whether is a significant difference between the two websites, i.e., you want to test $H_0 : \theta = \theta_0$ against $H_1 : \theta \neq \theta_0$.

For this illustration, we follow the Google Analytics example and set $\theta_0 = 0.05$. As to the choice of prior for TS, it would be reasonable to pick one that is centered around θ_0, so we choose Beta$(1, 1/\theta_0)$. With these values for θ_0 and the prior, we can simulate multiple runs of the Thompson sampling algorithm for this setting. Based on these simulations, Figure 2.1 plots the finite sample size of a naive t-test for sample sizes n ranging between 2500 and 10000 (for comparison, the Google Analytics example used $n = 6600$). As the t-test ignores the adaptive nature of the algorithm, it is incorrectly sized; in this particular setting, it is under-sized with a nominal asymptotic size of 4%.[4]

The same figure also shows the size of one of our proposed tests, defined as

$$\varphi^*(\hat{\tau}, \bar{x}_n) := \mathbb{I}\left(\int_{-1.5}^{1.5} \exp\left\{\frac{h\hat{\tau}}{\sigma_0}\sqrt{n}(\bar{x}_n - \theta_0) - \frac{h^2\hat{\tau}}{2\sigma_0}\right\} dh > \gamma\right),$$

where $\sigma_0 := \sqrt{\theta_0(1-\theta_0)}$, $\hat{\tau}$ represents the fraction of times the new website variant was sampled, and $\bar{x}_n$ is the sample average of outcomes under the new website variant. This test is designed to maximize the weighted average power (WAP) using a uniform weight over alternatives in the interval $\left[\theta_0 - 1.5n^{-1/2}, \theta_0 + 1.5n^{-1/2}\right]$. For the sample sizes considered, this approximately corresponds to a range of $[-3\%, 7\%]$ for θ. Deng et al. (2013) survey the practice of website optimization and suggest that, in practice,

[1]The webpage describing the simulation study can be accessed here.
[2]The conversion rate is defined as the percentage of users who have completed a desired action, e.g., clicking an ad.
[3]While we use TS as an illustration, our analysis applies to any bandit algorithm.
[4]In other sequential experiments, the naive t-test can be over-sized, see Section 5.1 for an example.

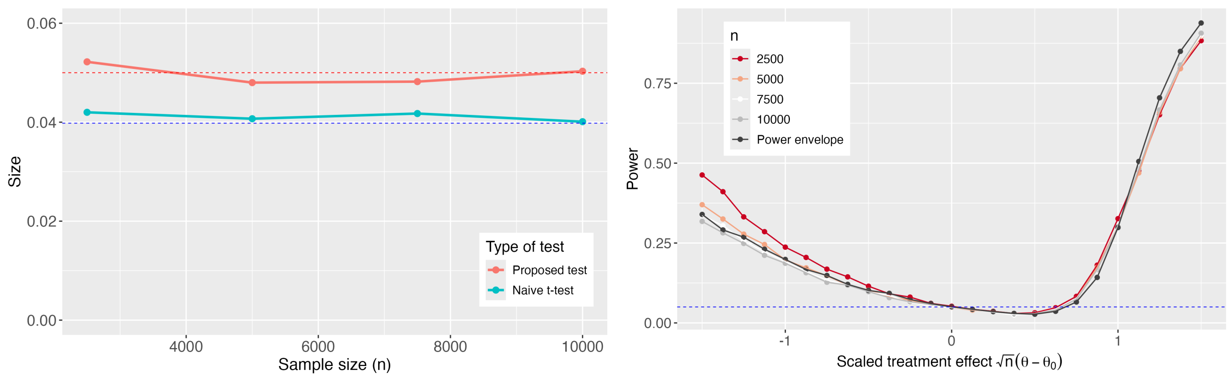

A: Size

B: Power function

Note: Panel A plots the size of our proposed weighted average power (WAP) test along with that of the naive t-test at the nominal 5% level (solid red line). The solid blue is the asymptotic size of the naive t-test. Panel B plots the finite sample power envelopes of the WAP test for different n, along with the associated asymptotic power envelope for the given weighting. The dashed blue line is the nominal size (5%).

FIGURE 2.1. Finite sample performance of the proposed WAP test

the difference $\theta - \theta_0$ between various variants of a website is typically less than 1%. Therefore, our chosen weighting aligns well with realistic values for θ.[5]

We determine the critical value γ by simulating the distribution of $\hat{\tau}, \bar{x}_n$ under the null. Specifically, we simulate the TS algorithm multiple times, drawing outcomes for both the old and new website variants from a Bernoulli(θ_0) distribution. This process ensures that our test maintains exact size control by construction.

Panel B of Figure 2.1 further depicts the power function of $\varphi^*(\cdot)$ across different sample sizes, showing that it closely approaches the power envelope for weighted average tests under the specified weighting. Notably, the power functions are not symmetric around θ_0. This asymmetry arises from the nature of the TS algorithm, which tends to over-sample the new website variant when $\theta > \theta_0$, resulting in higher power under those alternatives, but under-samples it when $\theta < \theta_0$.

2.2. **Setup and assumptions.** We now describe our general setup. Let Y denote the outcome variable(s) of interest. Before starting the experiment, a Decision-Maker (DM) commits to an experimental protocol with an associated stopping time, $\hat{\tau}$, that describes the eventual sample size in multiples of n observations (see below for the interpretation of n). The choice of $\hat{\tau}$ may involve balancing a number of considerations such as costs, ethics, welfare etc. We abstract away from these issues and take $\hat{\tau}$ as given. In the course of the experiment, the DM observes a sequence of outcomes $Y_1, Y_2, \ldots$ The stopping time $\hat{\tau}$ is assumed to be adapted to the filtration generated by the outcome observations. In

[5]Alternative weighting functions can also be employed, though we do not present those results here.

this section we employ a parametric model P_θ for the outcomes. Our interest is in testing $H_0 : \theta \in \Theta_0$ vs $H_1 : \theta \in \Theta_1$ where $\Theta_0 \cap \Theta_1 = \emptyset$.

There are two notions of asymptotics one could employ in this setting, and consequently, two different interpretations of n. In many examples, e.g., our empirical illustration, there is a limit on the maximum number of observations that can be collected; this limit is pre-specified and we take it to be n. Consequently, in these experiments, $\hat{\tau} \in [0, 1]$. Alternatively, we may have open-ended experiments where the stopping time is determined by balancing the benefit of experimentation with the cost for sampling each additional unit of observation. In this case, we employ small-cost asymptotics and n then indexes the rate at which the sampling costs go to 0 (alternatively, we can relate n to the population size in the implementation phase following the experiment, see Adusumilli, 2022). The results in this section apply to both asymptotic regimes.

In the context of the empirical illustration from Section 2.1, $Y \sim \text{Bernoulli}(\theta)$ is binary, and describes conversions under the new variant of the website. The observations from the original website are distributed as $\text{Bernoulli}(\theta_0)$, but as θ_0 is known, these observations are ancillary to the estimation of θ and can therefore be excluded from subsequent analysis. In this experiment, the stopping time is therefore $\hat{\tau} = \hat{Q}/n$, where $\hat{Q}$ denotes the total number of times the new website has been sampled over the course of the experiment.

Although we refer to Y_i as outcome variables for simplicity, our framework is more general and can accommodate covariates by incorporating them into Y_i. For example, in a one-armed contextual bandit, the "outcomes" z_i may be linked to covariates x_i through a parametric model, such as $z_i = x_i^\intercal \theta + \epsilon_i$, where $\epsilon_i \sim \mathcal{N}(0, \sigma^2)$. Incorporating this into our setup is straightforward: we simply define $Y_i := (z_i, x_i^\intercal)^\intercal$.

Returning to the general setup, let $\varphi_n \in [0, 1]$ denote a candidate test. It is required to be measurable with respect to $\sigma\{Y_1, \dots, Y_{\lfloor n\hat{\tau} \rfloor}\}$, i.e., it can only depend on information from the past outcomes. Now, it is fairly straightforward to construct tests that have power 1 against any fixed alternative as $n \to \infty$. Consequently, to obtain a more fine-grained characterization of tests, we consider their performance against local perturbations of the form $\{\theta_0 + h/\sqrt{n}; h \in \mathbb{R}^d\}$, where $\theta_0 \in \Theta_0$ denotes some reference parameter, chosen such that $\hat{\tau}$ has a non-trivial limit distribution. The choice of θ_0 is described in detail in Section 2.4, but in most of the applications it can be taken to be some parameter in the null set (as is the case with our empirical illustration).

The local perturbation analysis makes the testing problem harder even as the effective sample size $n \to \infty$. For instance, in the empirical illustration we tested $H_0 : \theta = \theta_0$, against $H_1 : \theta \neq \theta_0$. Deng et al. (2013) survey the practice of website optimization and document that the practical difference, $\theta - \theta_0$, between various variants of a website is quite small, less than 1% or so. Consequently, to distinguish between the two websites, it is important to develop tests that possess high power against close alternatives. This motivates our local perturbation analysis, where we aim to characterize power against local alternatives of the form $\theta_0 + h/\sqrt{n}$.

Let ν denote a dominating measure for the family $\{P_\theta : \theta \in \mathbb{R}^d\}$, and set $p_\theta := dP_\theta/d\nu$. We impose the following regularity conditions on the family P_θ:

Assumption 1. *The class $\{P_\theta : \theta \in \mathbb{R}^d\}$ is differentiable in quadratic mean around θ_0, i.e., there exists a score function $\psi(\cdot)$ such that for each $h \in \mathbb{R}^d$,*

$$\int \left[\sqrt{p_{\theta_0+h}} - \sqrt{p_{\theta_0}} - \frac{1}{2}h^\intercal \psi \sqrt{p_{\theta_0}}\right]^2 d\nu = o(|h|^2). \tag{2.1}$$

In the empirical illustration, Assumption 1 holds with $\psi(x) = x - \theta_0$. More broadly, this assumption is satisfied for a wide range of commonly used distributions, including the Normal, Cauchy, Exponential, and Poisson distributions.

Take $P_{nt,h}$ to be the joint probability measure over the iid sequence of outcomes $Y_1, \dots, Y_{nt}$ when each $Y_i \sim P_{\theta_0+h/\sqrt{n}}$. Let $\mathbb{E}_{nt,\boldsymbol{h}}[\cdot]$ represent the corresponding expectation under this measure. We define the (standardized) score process $x_n(t)$ as

$$x_n(t) = \frac{I^{-1/2}}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \psi(Y_i),$$

where $I := \mathbb{E}_{nT,0}[\psi(Y_i)\psi(Y_i)^\intercal]$ is the information matrix. It is well known, see e.g., Van der Vaart (2000, Chapter 7), that quadratic mean differentiability (Assumption 1) implies $\mathbb{E}_{nT,0}[\psi(Y_i)] = 0$ and ensures that I exists. Then, by a functional central limit theorem,

$$x_n(\cdot) \xrightarrow[P_{nT,0}]{d} x(\cdot); \; x(\cdot) \sim W(\cdot). \tag{2.2}$$

Here, and in what follows, $W(\cdot)$ denotes the standard d-dimensional Brownian motion. Assumption 1 also implies the important property of Sequential Local Asymptotic Normality (SLAN; Adusumilli, 2021): for any given $h \in \mathbb{R}^d$,

$$\sum_{i=1}^{\lfloor nt \rfloor} \ln \frac{dp_{\theta_0+h/\sqrt{n}}}{dp_{\theta_0}} = h^\intercal I^{1/2} x_n(t) - \frac{t}{2} h^\intercal I h + o_{P_{nT,0}}(1), \text{ uniformly over bounded } t. \tag{2.3}$$

The above states that the log-likelihood ratio admits a quadratic approximation uniformly over all bounded t.

Our next assumption concerns the properties of the stopping times $\hat{\tau}$:

Assumption 2. *There exists $T < \infty$ independent of n such that $\hat{\tau} \leq T$. Furthermore, $(\hat{\tau}, x_n(\hat{\tau}))$ has a weak limit under $P_{nT,0}$.*

Both requirements are fairly innocuous. We previously noted that $\tau \leq 1$ in many examples. By Prohorov's theorem, the first part of Assumption 2 along with (2.2) implies that $(\hat{\tau}, x_n(\hat{\tau}))$ is tight and therefore converges weakly along sub-sequences. The second part of Assumption 2 further disciplines the sequence by requiring it to converge weakly to the same limit under every subsequence. Since $\hat{\tau}$ indexes our experimental protocol, this essentially requires that the sequence of experimental protocols has a well defined asymptotic limit. This is a rather mild condition. In fact, it would be unusual for experimental protocols to have different weak limits across subsequences; if that were the case, they should be considered distinct protocols altogether.

2.3. **Asymptotic representation theorem.** We now characterize the limiting power of tests as $n \to \infty$ by associating each stopping-time experiment with a simpler experiment involving Gaussian diffusions.

Specifically, consider a limit experiment where one observes a Gaussian diffusion $x(t) := I^{1/2}ht + W(t)$, with an unknown h, along with a Uniform$[0,1]$ random variable U that is independent of the process $x(\cdot)$. Define $\mathcal{F}_t := \sigma\{x(s), U; s \leq t\}$ to be the filtration, i.e., the information set consisting of U and the knowledge of the stochastic process $x(\cdot)$ until time t. Let $\mathbb{P}_h$ denote the induced joint probability over the exogenous random variable U and the sample paths of $x(\cdot)$ given h, and take $\mathbb{E}_h[\cdot]$ to be its corresponding expectation. Suppose that we are interested in conducting inference on h using a test statistic φ that depends only on: (i) an $\mathcal{F}_t$-adapted stopping time τ that is the limiting version of $\hat{\tau}$ (in a sense made precise below); and (ii) the stopped process $x(\tau)$. The following theorem relates the original testing problem to the one in such a limit experiment:

Theorem 1. *Suppose Assumptions 1 and 2 hold. Let φ_n be some test function defined on the sample space $Y_1, \dots, Y_{n\hat{\tau}}$, and $\beta_n(h)$, its power against $P_{nT,h}$. Then:*
(i) (Le Cam, 1979) There exists an $\mathcal{F}_t$-adapted stopping time τ for which
$(\hat{\tau}, x_n(\hat{\tau})) \xrightarrow[P_{nT,0}]{d} (\tau, x(\tau))$.
(ii) Suppose that $\beta_n(h)$ converges point-wise for each $h \in \mathbb{R}^d$. Then, there exists a test

φ in the limit experiment depending only on $\tau, x(\tau)$ such that $\beta_n(h) \to \beta(h)$ for every $h \in \mathbb{R}^d$, where $\beta(h) := \mathbb{E}_h[\varphi(\tau, x(\tau))]$ is the power of φ in the limit experiment.

The requirement that $\beta_n(h)$ converges point-wise for each $h \in \mathbb{R}^d$ is rather mild. Classical representation theorems for tests, e.g., Van der Vaart (2000, Theorem 15.1), also require this. In any event, this assumption is not needed for deriving asymptotic upper bounds on tests. Suppose $\beta^*(h)$ denotes an upper bound on the power of tests in the limit experiment. Then, even if $\beta_n(\cdot)$ were not a convergent sequence, we can still employ Theorem 1 on subsequences (since $\beta_n(h) \in [0,1]$ is tight) to show that $\limsup_n \beta_n(h) \le \beta^*(h)$. In Section 2.8 we describe tests that attain the upper bound, showing that the bound is indeed tight.

2.3.1. *Discussion.* Theorem 1 states that irrespective of the algorithm used to conduct the adaptive experiment, the only statistics that matter are the stopping time and stopped (i.e., final) value of the score process. Different algorithms induce different joint distributions over these statistics under the null, and this joint distribution determines the critical value of the test.

The first part of Theorem 1 is essentially due to Le Cam (1979). It states that any experimental protocol (as associated with $\hat{\tau}$) can be matched with a suitable protocol in the limit experiment, in the sense of replicating the joint distributions of $\hat{\tau}, x_n(\hat{\tau})$. Note, however, that this does not by itself imply $(\hat{\tau}, x_n(\cdot)) \xrightarrow[P_{nT,0}]{d} (\tau, x(\cdot))$ as a process over t. Also, the result makes no claims about the distributions of statistics other than $\hat{\tau}, x_n(\hat{\tau})$.

The second part of Theorem 1 is new. The proof builds on the quadratic approximation to the log-likelihood, (2.3), previously derived in Adusumilli (2021). Combining with the first part of Theorem 1, this approximation implies

$$\ln \frac{dP_{n\hat{\tau}.,h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}) \xrightarrow[P_{nT,0}]{d} h^\intercal I^{1/2} x(\tau) - \frac{\tau}{2} h^\intercal I h. \tag{2.4}$$

By the Girsanov theorem, the right hand side of (2.4) corresponds to the log-likelihood ratio $\ln \varphi(\tau; h) := \ln \frac{d\mathbb{P}_h}{d\mathbb{P}_0}(\tau)$ in the limit experiment with the protocol τ. In classical, i.e., non-sequential settings, it is well established that weak convergence of likelihood ratios implies the equivalence of experiments and that one can achieve asymptotic dimension reduction by identifying sufficient statistics in the limit experiment. However, the sequential setting warrants a bit of caution. Although the likelihood ratio $\varphi(\tau; h)$ depends solely on $\tau, x(\tau)$, these quantities do not constitute sufficient statistics in the traditional sense. Specifically, the conditional distribution of the data $\{x(t), U : 0 \le t \le \tau\}$ given

$(\tau, x(\tau))$ generally remains dependent on h.[6] Consequently, we cannot hope to represent the distributions of all possible statistics from the sequential experiment in terms of just $\tau, x(\tau)$.

Nevertheless, for hypothesis testing, we demonstrate that it is indeed feasible to condition solely on $\tau, x(\tau)$. To build intuition for this result, note that the power of the test depends on h solely through the stopped value of the log-likelihood process $\ln \frac{dP_{nt,\cdot,h}}{dP_{nt,0}}(\mathbf{y}_{nt})$, i.e.,

$$\beta_n(h) = \mathbb{E}_{nT,h}[\varphi_n] = \mathbb{E}_{nT,0}\left[\varphi_n \exp\left\{\ln \frac{dP_{n\hat{\tau},\cdot,h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}})\right\}\right].$$

Specifically, as the log-likelihood ratio process is a martingale and also Markovian, its value at $\hat{\tau}$ summarizes all relevant information about the distribution of φ_n under h. This observation, combined with (2.4), suggests that $\tau, x(\tau)$ should be sufficient for testing. The formal proof relies on a novel change-of-measure argument that extends Le Cam's third lemma to the sequential setting. This derivation represents the main theoretical contribution of this paper.

2.3.2. *Power enhancement by conditioning on $\tau, x(\tau)$.* The proof of Theorem 1 reveals that the power of any test in the limit experiment can be enhanced by conditioning on $\tau, x(\tau)$ under $\mathbb{P}_0$. Specifically, if $\tilde{\varphi}(\cdot)$ is a valid test in the limit experiment, then the test defined by $\varphi(\tau, x(\tau)) := \mathbb{E}_0[\tilde{\varphi} \mid \tau, x(\tau)]$ is also valid. Moreover, the power of $\varphi(\cdot)$ under any local alternative indexed by h is at least as great as that of $\tilde{\varphi}(\cdot)$. This approach is analogous to Rao-Blackwellization in classical statistics, although it is important to recall that $\tau, x(\tau)$ do not constitute sufficient statistics.

2.4. **Drifting nulls and the choice of the reference parameter.** The reference parameter θ_0 needs to be chosen such that the limiting stopping time does not degenerate to a point mass at 0. To illustrate the pitfalls from a non-judicious choice of the reference parameter, consider the empirical illustration with the one-armed bandit and suppose we are interested in testing $H_0 : \theta = \bar{\theta}$ where $\bar{\theta} < \theta_0$. Any 'regret-consistent' bandit algorithm - this includes Thompson Sampling - would only sample the new website variant a vanishingly small fraction of the time if the 'reward gap' $\bar{\theta} - \theta_0$ is strictly negative. Hence, if $\bar{\theta}$ were chosen as the reference parameter, the stopping $\hat{\tau}$ would converge in

[6] Incidentally, this implies that Fisher's factorization theorem fails in sequential settings. If τ were fixed to a value T (say), as in classical experiments, the well-known properties of the Brownian bridge would imply that, given $x(T)$, the distribution of the sample paths $\{x(t) : 0 \leq t \leq T\}$ is independent of h and $x(T)$ would therefore be a sufficient statistic. However, this property no longer holds if τ is endogenously determined on the basis of the past values of the score process.

probability-$P_{\bar{\theta}}$ (i.e., under the null) to 0. Intuitively, because the bandit algorithm is designed to detect small shifts with respect to the baseline of θ_0, it is not particularly well suited for testing $\theta = \bar{\theta}$ when $\bar{\theta}$ is substantially smaller than θ_0.

Consequently, to provide inference in such settings, we employ a drifting null. Specifically, we set the reference parameter θ_0 to be the one that leads to a non-trivial limit for $\hat{\tau}$, but take the null to be $H_0 : h = h_0/\sqrt{n}$, where h_0 is fixed over n and calibrated as $\sqrt{n}(\bar{\theta} - \theta_0)$. The null, H_0, thus changes with n, but for the observed sample size we are still testing $\theta = \bar{\theta}$. It is then straightforward to show that Theorems 1 and 2 continue to apply in this setting: asymptotically, the inference problem is equivalent to testing that the drift of $x(\cdot)$ is $I^{1/2}h_0$ in the limit experiment. The asymptotic approximation will be more accurate the closer $\bar{\theta}$ is to θ_0.

Using drifting nulls, we can also obtain valid confidence intervals for θ by constructing tests for each local null value $h_0 \in \mathbb{R}^d$ and then employing test inversion. These confidence intervals shrink at $n^{-1/2}$ rates, as would be expected in parametric settings.

Previously, Romano (2005) employed the idea of a drifting null hypothesis parameter space for testing equivalence hypotheses. While the setup is different, the motivation behind the use of the drifting null is similar: it ensures the problem is asymptotically non-degenerate.

2.5. **Simulating the distribution of $(\tau, x(\tau))$.** In the next section, we characterize the form of optimal tests in the limit experiment. However, to determine their critical values, one would need to know, or be able to simulate from the joint distribution of $\tau, x(\tau)$ under θ_0 (i.e., when $h = 0$). Unfortunately, the first part of Theorem 1 does not explicitly characterize τ; it merely establishes that such a stopping time must exist. The second part of Theorem 1 then takes this τ as given, analogous to how $\hat{\tau}$ is treated as given in the original experiment.

Fortunately, many stopping times used in practice can be expressed as a function $\hat{\tau} = g(x_n(\cdot), U)$, where $g(\cdot)$ depends on the sample paths of the score process $x_n(\cdot)$ over the interval $[0, T]$, along with some exogenous randomization $U \sim \text{Uniform}[0, 1]$. Indeed, previous research (e.g., Adusumilli, 2022) has demonstrated that any asymptotically Bayes-optimal stopping time depends solely on these two components. The analysis is then straightforward if $g(\cdot)$ is known in closed form. For example, this is true for the optimal stopping time under costly sampling, where $\hat{\tau} = \inf\{t : |x_n(t)| \geq \gamma\}$, which implies $g(z(\cdot)) = \inf\{t : |z(t)| \geq \gamma\}$. In this case, it follows from the extended continuous mapping theorem that $\tau = g(x(\cdot))$. Given that $x(\cdot)$ follows a standard Brownian motion when $h = 0$, determining the joint distribution of $(\tau, x(\tau))$ becomes straightforward through simulating Brownian motion.

However, in many cases, the relationship between $x_n(\cdot)$ and $\hat{\tau}$ is more complicated. For example, in the context of Thompson sampling, the joint distribution of $(\tau, x(\tau))$ is obtained by solving a set of coupled stochastic differential equations, as demonstrated by Fan and Glynn (2021) and Wager and Xu (2021). Fortunately, there is a simpler approach: since θ_0 is generally known in advance, we can simulate experimental data by randomly drawing new outcome values from P_{θ_0} and following the given experimental protocol. This enables us to determine the exact distribution of $(\hat{\tau}, x_n(\hat{\tau}))$ under P_{θ_0}. The first part of Theorem 1 implies the distribution of $(\hat{\tau}, x_n(\hat{\tau}))$ converges weakly to that of $(\tau, x(\tau))$ under P_{θ_0}, so under mild conditions, the critical values derived from $(\hat{\tau}, x_n(\hat{\tau}))$ will converge to those obtained from $(\tau, x(\tau))$. In fact, even when it is feasible to simulate $(\tau, x(\tau))$ directly, it is much preferable to use the critical values based on $(\hat{\tau}, x_n(\hat{\tau}))$ since this approach ensures that the resulting test maintains exact size control in finite samples.

2.6. **Characterization of optimal tests in the limit experiment.** We now construct optimal tests in the limit experiment under various scenarios. The utility of this analysis is two-fold: First, it enables us to provide an asymptotic upper bound on tests. Second, we can construct asymptotically optimal tests as the sample counterparts of optimal tests in the limit experiment.

2.6.1. *Testing a parameter vector.* The simplest hypothesis testing problem in the limit experiment concerns testing $H_0 : h = 0$ vs $H_1 : h = h_1$. This is asymptotically equivalent to testing $H_0 : \theta = \theta_0$ vs $H_1 : \theta = \theta_0 + h/\sqrt{n}$ in the original experiment. By the Neyman-Pearson lemma, the uniformly most powerful (UMP) test is

$$\varphi_{h_1}^* = \mathbb{I}\left\{h_1^\intercal I^{1/2} x(\tau) - \frac{\tau}{2} h_1^\intercal I h_1 > \gamma_{h_1}\right\},$$

where the critical value $\gamma_{h_1} \in \mathbb{R}$ is chosen by the size requirement, and can be computed using the procedures described in Section 2.5.

Let $\beta^*(h_1)$ denote the power function of $\varphi_{h_1}^*$. Then, by Theorem 1, $\beta^*(\cdot)$ is an upper bound on the asymptotic power function of any test of $H_0 : \theta = \theta_0$.

2.6.2. *Testing linear combinations.* When θ is a vector, we are often interested in inference on a scalar sub-component, or perhaps more generally, a linear transformation, $a^\intercal\theta$, of it. This translates to testing linear combinations of h, i.e., $H_0 : a^\intercal h = 0$, in the limit experiment. In a standard (i.e., non-sequential) setting, it is known that testing linear combinations is associated with a further dimension reduction of the data since we can condition on an equivalent linear transformation of the score function. We show that a similar dimension reduction is possible in the sequential setting as well, but it requires the stopping time to also depend only on a reduced set of statistics.

Define $\sigma^2 := a^\intercal I^{-1} a$, $\tilde{x}(t) := \sigma^{-1} a^\intercal I^{-1/2} x(t)$, let U_1 denote a Uniform$[0,1]$ random variable independent of $\tilde{x}(\cdot)$, and take $\tilde{\mathcal{F}}_t$ to be the filtration generated by $\sigma\{U_1, \tilde{x}(s) : s \le t\}$. Note that $\tilde{x}(\cdot) \sim W(\cdot)$ under the null and is therefore a pivotal statistic.

Proposition 1. *Suppose that the stopping time τ in Theorem 1 is $\tilde{\mathcal{F}}_t$-adapted. Then, the UMP test of $H_0 : a^\intercal h = 0$ vs $H_1 : a^\intercal h = c$ in the limit experiment is*

$$\varphi_c^*(\tau, \tilde{x}(\tau)) = \mathbb{I}\left\{ c\tilde{x}(\tau) - \frac{c^2}{2\sigma}\tau > \gamma_c \right\}.$$

In addition, suppose Assumptions 1 and 2 hold, let $\beta^*(c)$ denote the power of φ_c^* for a given c, and $\beta_n(h)$ the power of some test, φ_n, of $H_0 : a^\intercal \theta = 0$ in the original experiment against local alternatives $\theta \equiv \theta_0 + h/\sqrt{n}$. Then, for each $h \in \mathbb{R}^d$, $\limsup_{n\to\infty} \beta_n(h) \le \beta^*(a^\intercal h)$.

The above result suggests that $\tilde{x}(\tau)$ and τ are sufficient statistics for the optimal test. An important caveat, however, is that the class of stopping times are further constrained to only depend on $\tilde{x}(t)$ in the limit. In practice, this would happen if the stopping time $\hat{\tau}$ in the original experiment is a function only of $\hat{\tilde{x}}_n(\cdot) := \sigma^{-1} a^\intercal I^{-1/2} x_n(\cdot)$. Fortunately, this is the case in a number of examples.

It is straightforward to show that the same power envelope, $\beta^*(\cdot)$, also applies to tests of the composite hypothesis $H_0 : a^\intercal\theta \le 0$.

2.6.3. *Unbiased tests.* A test is said to be unbiased if its power is greater than size under all alternatives. This is a desirable property and one may wish to enforce this. The following result describes a useful property of unbiased tests in the limit experiment:

Proposition 2. *Any unbiased test of $H_0 : h = 0$ vs $H_1 : h \neq 0$ in the limit experiment must satisfy $\mathbb{E}_0[x(\tau)\varphi(\tau, x(\tau))] = 0$.*

We can obtain the best unbiased test φ_b^* (if it exists) by solving the optimization problem

$$\arg\max_{\varphi(\cdot)} \mathbb{E}_h [\varphi] \text{ s.t. } \mathbb{E}_0 [\varphi] \leq \alpha \text{ and } \mathbb{E}_0[x(\tau)\varphi(\tau, x(\tau))] = 0. \tag{2.5}$$

Note that the trivial test $\varphi = \alpha$ satisfies the constraints in (2.5) since $\mathbb{E}_0[x(\tau)] = 0$ by the martingale property. Consequently, it follows by Proposition 2 that if (2.5) admits a solution φ_b^* which is independent of h, then it is best unbiased. See Section 5.1 for an application of this result.

2.6.4. *Weighted average power.* Except in some specific contexts, it is not usually possible characterize a uniquely most powerful test even if we impose unbiasedness. In such cases, it is common to use a weighted average power criterion. We specify a weight function, $w(\cdot)$, over the local alternatives $h \neq 0$, and aim to maximize weighted average power $\int \mathbb{E}_h [\varphi] \, dw(h)$. For instance, one could take $w(\cdot)$ to be normal or even uniform (within a compact set), as in the empirical application. The test of $H_0 : h = 0$ in the limit experiment that maximizes weighted average power is given by

$$\varphi_w^*(\tau, x(\tau)) = \mathbb{I}\left\{ \int e^{h^\intercal I^{1/2} x(\tau) - \frac{\tau}{2} h^\intercal I h} dw(h) > \gamma \right\}. \tag{2.6}$$

The critical value γ will need to be computed following the procedures described in Section 2.5.

2.7. **Alpha-spending and conditional size criteria.** In this section, we study inference under a stronger version of the size constraint, inspired by the $\boldsymbol{\alpha}$-spending approach in group sequential trials (Gordon Lan and DeMets, 1983). Suppose that the stopping time is discrete, taking only the values $t = 1, 2, \ldots, T$. Then, instead of an overall size constraint of the form $\mathbb{E}_{nT,\mathbf{0}}[\varphi_n] \leq \alpha$, we may specify a 'spending-vector' $\boldsymbol{\alpha} := (\alpha_1, \ldots, \alpha_T)$ satisfying $\sum_{t=1}^{T} \alpha_t = \alpha$, and require

$$\mathbb{E}_{nT,\mathbf{0}}[\mathbb{I}\{\hat{\tau} = t\}\varphi_n] \leq \alpha_t \ \forall \ t. \tag{2.7}$$

In what follows, we call a test, φ_n, satisfying (2.7) a level-$\boldsymbol{\alpha}$ test (with a boldface $\boldsymbol{\alpha}$). Intuitively, if each t corresponds to a different stage of the experiment, the $\boldsymbol{\alpha}$-spending constraint prescribes the maximum amount of Type-I error that may be expended at stage t. As a practical matter, it enables us to characterize a UMP or UMP unbiased test in settings where such tests do not otherwise exist. We also envision the criterion as a useful conceptual device: even if we are ultimately interested in a standard level-α test, we can obtain this by optimizing a chosen power criterion (average power, etc.) over the spending vectors $\boldsymbol{\alpha} := (\alpha_1, \ldots, \alpha_K)$ satisfying $\sum_k \alpha_k \leq \alpha$.

A particularly interesting example of an $\boldsymbol{\alpha}$-spending vector is $(\alpha P_{nT,0}(\hat{\tau} = 1), \ldots, \alpha P_{nT,0}(\hat{\tau} = k))$; this corresponds to the requirement that $\mathbb{E}_{nT,\mathbf{0}}\left[\varphi_n \middle| \hat{\tau} = t\right] \leq \alpha$ for all t, i.e., the test be conditionally level-α given any realization of the stopping time. Note, however, that $\hat{\tau}$ is not ancillary and tests based on this stronger notion of size constraint would have lower power since the criterion disregards information provided by the stopping time for discriminating between the hypotheses.

Under the $\boldsymbol{\alpha}$-spending constraint, a test that maximizes expected power also maximizes expected power conditional on each realization of stopping time. This is a simple consequence of the law of iterated expectations. Consequently, we focus on conditional power in this section. Our main result here is a generalization of Theorem 1 to $\boldsymbol{\alpha}$-spending restrictions. The limit experiment is the same as in Section 2.3.

Theorem 2. *Suppose Assumptions 1, 2 hold, and the stopping times are discrete, taking only the values $1, 2, \ldots, T$. Let φ_n be some level-$\boldsymbol{\alpha}$ test defined on the sample space $Y_1, \ldots, Y_{n\hat{\tau}}$, and $\beta_n(h|t)$, its conditional power against $P_{nT,h}$ given $\hat{\tau} = t$. Suppose that $\beta_n(h|t)$ converges point-wise for each h, t. Then, there exists a level-$\boldsymbol{\alpha}$ test, $\varphi(\cdot)$, in the limit experiment depending only on $\tau, x(\tau)$ such that, for every $h \in \mathbb{R}^d$ and $t \in \{1, 2, \ldots, T\}$ for which $\mathbb{P}_0(\tau = t) \neq 0$, $\beta_n(h|t)$ converges to $\beta(h|t)$, where $\beta(h|t) := \mathbb{E}_h[\varphi(\tau, x(\tau))|\tau = t]$ is the conditional power of $\varphi(\cdot)$ in the limit experiment.*

2.7.1. *Power envelope.* By the Neyman-Pearson lemma, the uniformly most powerful level-$\boldsymbol{\alpha}$ (UMP-$\boldsymbol{\alpha}$) test of $H_0 : h = 0$ vs $H_1 : h = h_1$ in the limit experiment is given by

$$\varphi^*_{h_1}(t, x(t)) = \begin{cases} 1 & \text{if } \mathbb{P}_0(\tau = t) \leq \alpha_t \\ \mathbb{I}\left\{h_1^\intercal I^{1/2} x(t) > \gamma(t)\right\} & \text{if } \mathbb{P}_0(\tau = t) > \alpha_t \end{cases}.$$

Here, $\gamma(t) \in \mathbb{R}$ is chosen by the $\boldsymbol{\alpha}$-spending requirement that $\mathbb{E}_0[\varphi^*_{h_1}(\tau, x(\tau))|\tau = t] \leq \alpha_t/\mathbb{P}_0(\tau = t)$ for each t. If we take $\beta^*(h_1|t)$ to be the power function of $\varphi^*_{h_1}(\cdot)$, Theorem 2 implies $\beta^*(\cdot|t)$ is an upper bound on the limiting conditional power function of any level-$\boldsymbol{\alpha}$ test of $H_0 : \theta = \theta_0$.

2.7.2. *Testing linear combinations.* A stronger result is possible for tests of linear combinations of θ. Recall the definitions of $\tilde{x}(t)$ and $\tilde{\mathcal{F}}_t$ from Section 2.6.2. If the limiting stopping time is $\tilde{\mathcal{F}}_t$ -adapted, we have, as in Proposition 1 that the sufficient statistics are only $\tilde{x}(\tau), \tau$. Furthermore, the UMP-$\boldsymbol{\alpha}$ test of $H_0 : a^\intercal h = 0$ vs $H_1 : a^\intercal h = c\ (> 0)$ in the limit experiment is

$$\breve{\varphi}^*(t, \tilde{x}(t)) = \begin{cases} 1 & \text{if } \mathbb{P}_0(\tau = t) \leq \alpha_t \\ \mathbb{I}\left\{c\tilde{x}(t) > \gamma_c(t)\right\} \equiv \mathbb{I}\left\{\tilde{x}(t) > \tilde{\gamma}(t)\right\} & \text{if } \mathbb{P}_0(\tau = t) > \alpha_t \end{cases}.$$

Here, $\tilde{\gamma}(t)$ is chosen such that $\mathbb{E}_0[\breve{\varphi}^*(\tau, \tilde{x}(\tau))|\tau = t] = \alpha_t/\mathbb{P}_0(\tau = t)$. Clearly, $\tilde{\gamma}(t)$ it is independent of c for $c > 0$. Since $\breve{\varphi}^*(\cdot)$ is thereby also independent of c for $c > 0$, we conclude that it is UMP-$\boldsymbol{\alpha}$ for testing the composite one-sided alternative $H_0 : a^\intercal h = 0$ vs $H_1 : a^\intercal h > 0$. Thus, there exists a uniquely most powerful one sided test in the $\boldsymbol{\alpha}$-spending setting. The test has the same form as a classical one-sided test, but the critical value $\tilde{\gamma}(t)$ is non-standard and needs to be determined by simulating the distribution of $\tilde{x}(t)$ given t.

We also note that by Theorem 2, the conditional power function, $\breve{\beta}^*(c|t)$, of $\breve{\varphi}^*(\cdot)$ is an asymptotic upper bound on the conditional power of any level-$\boldsymbol{\alpha}$ test, φ_n, of $H_0 : a^\intercal \theta = 0$ vs $H_1 : a^\intercal \theta > 0$ in the original experiment against local alternatives $\theta \equiv \theta_0 + h/\sqrt{n}$ satisfying $a^\intercal \theta = c/\sqrt{n}$.

2.7.3. *Conditionally unbiased tests.* We call a test conditionally unbiased if it is unbiased conditional on any possible realization of the stopping time. In analogy with Proposition 2, a necessary condition for $\varphi(\cdot)$ being conditionally unbiased in the limit experiment is that

$$\mathbb{E}_0\left[x(\tau)\left(\varphi(\tau, x(\tau)) - \alpha\right)|\tau = t\right] = 0\ \forall\ t. \tag{2.8}$$

Then, by a similar argument as in Lehmann and Romano (2005, Section 4.2), the UMP conditionally unbiased (level-$\boldsymbol{\alpha}$) test of $H_0 : a^\intercal h = 0$ vs $H_1 : a^\intercal h \neq 0$ in the limit experiment can be shown to be

$$\bar{\varphi}^*(t, \tilde{x}(t)) = \begin{cases} 1 & \text{if } \mathbb{P}_0(\tau = t) \leq \alpha_t \\ \mathbb{I}\left\{\tilde{x}(t) \notin [\gamma_L(t), \gamma_U(t)]\right\} & \text{if } \mathbb{P}_0(\tau = t) > \alpha_t \end{cases},$$

where $\gamma_L(t), \gamma_U(t)$ are chosen to satisfy both (2.7) and (2.8). In practice, this requires simulating the distribution of $\tilde{x}(\tau)$ given $\tau = t$. Also, $\gamma_L(\cdot) = -\gamma_U(\cdot)$ if the distribution of $\tilde{x}(\tau)$ given $\tau = t$ is symmetric around 0 under the null. As before, the test has the same form as a classical two-sided test, but it employs non-standard critical values.

2.8. **Attaining the power bounds.** So far we have described upper bounds on the asymptotic power functions of tests. Under Assumption 2, given a UMP test in the limit experiment, we can construct a finite sample version of this by replacing $\tau, x(\tau)$ with $\hat{\tau}, x_n(\hat{\tau})$. The resulting test statistic would then be applicable in finite samples and also asymptotically optimal, in in the sense of attaining the power envelope.

To outline the procedure, let $\varphi^*(\tau, x(\tau); \gamma^*)$ represent a UMP test in the limit experiment, indexed by its critical value γ^*. As detailed in Section 2.5, γ^* can be estimated either by simulating the exact distribution of $(\tau, x(\tau))$ or by using a finite-sample version based on the distribution of $(\hat{\tau}, x_n(\hat{\tau}))$. Let $\hat{\gamma}$ denote the resulting estimate, which we assume to be consistent under the reference probability distribution; that is, $\hat{\gamma} = \gamma^* + o_{P_{nT,0}}(1)$. A sufficient condition for this consistency is that the function $\alpha(\gamma) := \mathbb{E}_{P_{\theta_0}}[\varphi^*(\tau, x(\tau); \gamma)]$ is strictly monotone over γ.

The finite-sample version of φ^* is then given by $\varphi_n^* := \varphi^*(\hat{\tau}, x_n(\hat{\tau}); \hat{\gamma})$. Since $x_n(\hat{\tau})$ depends on the information matrix I, it is necessary to either calibrate it to $I(\theta_0)$ or replace it with a consistent estimator. We discuss variance estimators in Appendix B.1.

The test φ_n^* is asymptotically optimal, in the sense of attaining the power envelope, under mild assumptions. In particular, we only require that $\varphi^*(\cdot, \cdot; \cdot)$ satisfy the conditions for an extended continuous mapping theorem. Together with (2.3) and the first part of Theorem 1, this implies

$$\begin{pmatrix} \varphi^*(\hat{\tau}, x_n(\hat{\tau}); \hat{\gamma}) \\ \sum_{i=1}^{\lfloor n\hat{\tau} \rfloor} \ln \frac{dp_{\theta_0 + h/\sqrt{n}}}{dp_{\theta_0}}(Y_i) \end{pmatrix} \xrightarrow[P_{nT,0}]{d} \begin{pmatrix} \varphi^*(\tau, x(\tau); \gamma^*) \\ h^\intercal I^{1/2} x(\tau) - \frac{\tau}{2} h^\intercal I h \end{pmatrix},$$

for any $h \in \mathbb{R}^d$. Then, a similar argument as in the proof of Theorem 1 shows that the local power of φ_n^* converges to that of φ^* in the limit experiment.

3. Testing in non-parametric settings

The previous section concentrated on parametric models, where the distribution of outcomes is predefined. However, in many real-world applications, the exact distribution is often unknown, leading to a non-parametric setting. In such cases, our objective might be to perform inference on a regular functional, denoted by $\mu := \mu(P)$, of the unknown data distribution P. Common examples of regular functionals include the mean, median, and quantiles. For simplicity, we assume that μ is a scalar.

The gist of our non-parametric results is that, under certain conditions on the stopping time, the sample analogs of the tests described in Section (2.6) - which were optimal in the limiting experiment - remain optimal in the non-parametric setting.

3.1. **Formal results.** Our formal analysis of the non-parametric regime follows Van der Vaart (2000). Let $\mathcal{P}$ denote the class of probability distributions with bounded variance and dominated by some measure ν. We then fix a reference $P_0 \in \mathcal{P}$, and surround it with various smooth one-dimensional parametric sub-models, $\{P_{s,h} : s \leq \eta\}$ for some $\eta > 0$, whose score function is $h(\cdot)$ and that pass through P_0 at $s = 0$ (i.e., $P_{0,h} = P_0$). Formally, these sub-models satisfy

$$\int \left[\frac{dP_{s,h}^{1/2} - dP_0^{1/2}}{s} - \frac{1}{2} h dP_0^{1/2} \right]^2 d\nu \to 0 \text{ as } s \to 0. \tag{3.1}$$

As in Section 2.4, the reference P_0 should be chosen such that the stopping time $\hat{\tau}$ has a non-trivial limit distribution under it.

By Van der Vaart (2000), (3.1) implies $\int h dP_0 = 0$ and $\int h^2 dP_0 < \infty$. The set of all such candidate h is termed the tangent space $T(P_0)$. This is a subset of the Hilbert space $L^2(P_0)$, endowed with the inner product $\langle f, g \rangle = \mathbb{E}_{P_0}[fg]$ and norm $\|f\| = \mathbb{E}_{P_0}[f^2]^{1/2}$. For any $h \in T(P_0)$, let $P_{nT,h}$ denote the joint probability measure over $Y_1, \dots, Y_{nT}$, when each Y_i is an iid draw from $P_{1/\sqrt{n},h}$. Also, take $\mathbb{E}_{nT,h}[\cdot]$ to be its corresponding expectation. An important implication of (3.1) is the SLAN property that for all $h \in T(P_0)$,

$$\sum_{i=1}^{\lfloor nt \rfloor} \ln \frac{dP_{1/\sqrt{n},h}}{dP_0}(Y_i) = \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} h(Y_i) - \frac{t}{2} \|h\|^2 + o_{P_{nT,0}}(1), \quad \text{uniformly over } t. \tag{3.2}$$

See Adusumilli (2021, Lemma 2) for the proof.

Let $\psi \in T(P_0)$ denote the efficient influence function corresponding to μ, in the sense that for any $h \in T(P_0)$,

$$\frac{\mu(P_{s,h}) - \mu(P_0)}{s} - \langle \psi, h \rangle = o(s). \tag{3.3}$$

Denote $\sigma^2 = \mathbb{E}_{P_0}[\psi^2]$. The analogue of the score process in the non-parametric setting is the efficient influence function process

$$x_n(t) := \frac{\sigma^{-1}}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \psi(Y_i).$$

At a high level, the theory for inference in non-parametric settings is closely related to that for testing linear combinations in parametric models (see, Section 2.6). It is not entirely surprising, then, that the assumptions described below are similar to those used in Proposition 1:

Assumption 3. *(i) The sub-models* $\{P_{s,h}; h \in T(P_0)\}$ *satisfy (3.1). Furthermore, they admit an efficient influence function,* $\psi(\cdot)$*, for* $\mu(P)$ *such that (3.3) holds.*

(ii) The stopping time $\hat{\tau}$ *is a continuous function of* $x_n(\cdot)$ *in the sense that* $\hat{\tau} = \tau(x_n(\cdot))$*, where* $\tau(\cdot)$ *satisfies the conditions for an extended continuous mapping theorem (Van Der Vaart and Wellner, 1996, Theorem 1.11.1).*

Assumption 3(i) is a mild regularity condition that is common in non-parametric analysis. Assumption 3(ii), which is substantive, states that the stopping time depends only on the efficient influence function process. This is indeed the case for the examples considered in Section 5. More generally, however, it may be that $\hat{\tau}$ depends on other statistics beyond $x_n(\cdot)$. In such situations, the set of asymptotically sufficient statistics should be expanded to include these additional ones. An extension of our results to these situations is straightforward, albeit case specific, see Section 4.4 for an illustration.

The definition of asymptotic size in the non-parametric regime requires a bit of care. We follow Choi et al. (1996) and call a test, φ_n, of $H_0 : \mu = 0$ asymptotically level-α if

$$\sup_{\{h \in T(P_0): \langle \psi, h \rangle = 0\}} \limsup_n \int \varphi_n dP_{nT,h} \leq \alpha.$$

Intuitively, this requires the test statistic to be correctly sized under all possible parametric sub-models.

3.1.1. *Power envelope.* Our first result in this section is a power envelope for asymptotically level-α tests. Consider a limit experiment where one observes a stopping time τ, which is the weak limit of $\hat{\tau}$, and a Gaussian process $x(\cdot) \sim \sigma^{-1}\mu \cdot + W(\cdot)$, where $W(\cdot)$ denotes 1-dimensional Brownian motion. By Assumption 3(ii) and the functional central limit theorem applied on $x_n(\cdot)$, we have $\tau = \tau(x(\cdot))$ and τ is therefore adapted to the filtration generated by the sample paths of $x(\cdot)$. For any $\mu \in \mathbb{R}$, let $\mathbb{E}_\mu[\cdot]$ denote the induced distribution over the sample paths of $x(\cdot)$ between $[0, T]$. Also, define

$$\varphi^*_\mu(\tau, x(\tau)) := \mathbb{I}\left\{\frac{\mu}{\sigma}x(\tau) - \frac{\mu^2}{2\sigma^2}\tau > \gamma\right\}, \tag{3.4}$$

with the critical value γ being determined by the requirement $\mathbb{E}_0[\varphi^*_\mu] = \alpha$, and set $\beta^*(\mu) := \mathbb{E}_\mu[\varphi^*_\mu]$.

Proposition 3. *Suppose Assumption 3 holds. Let $\beta_n(h)$ denote the power of some asymptotically level-α test, φ_n, of $H_0 : \mu = 0$ against local alternatives $P_{\delta/\sqrt{n},h}$. Then, for every $h \in T(P_0)$ and $\mu := \delta \langle \psi, h \rangle$, $\limsup_{n\to\infty} \beta_n(h) \leq \beta^* (\mu)$.*

Proposition 3 states that $\beta^*(\cdot)$ is an asymptotic upper bound on the power of any test φ_n. The power envelope depends solely on the functional $\mu(P)$ of the local alternative distribution P.

3.1.2. *Unbiased tests.* A similar result holds for unbiased tests. Following Choi et al. (1996), we say that a test φ_n of $H_0 : \mu = 0$ vs $H_1 : \mu \neq 0$ is asymptotically unbiased if

$$\sup_{\{h\in T(P_0):\langle\psi,h\rangle=0\}} \limsup_n \int \varphi_n dP_{nT,h} \leq \alpha, \text{ and}$$
$$\inf_{\{h\in T(P_0):\langle\psi,h\rangle\neq0\}} \liminf_n \int \varphi_n dP_{nT,h} \geq \alpha.$$

The following result indicates that the local power of such a test is bounded by that of the optimal unbiased test in the limiting experiment, assuming one exists. Given the nature of the limit experiment described above, the optimal unbiased tests have the same characterization (2.5) as in the parametric setting (with μ now replacing h in that characterization).

Proposition 4. *Suppose Assumption 3 holds and there exists a best unbiased test, $\tilde{\varphi}^*$, in the limit experiment with power function $\bar{\beta}^*(\mu)$. Let $\beta_n(h)$ denote the power of some asymptotically unbiased test, φ_n, of $H_0 : \mu = 0$ vs $H_1 : \mu \neq 0$ over local alternatives $P_{\delta/\sqrt{n},h}$. Then, for every $h \in T(P_0)$ and $\mu := \delta \langle \psi, h \rangle$, $\limsup_{n\to\infty} \beta_n(h) \leq \tilde{\beta}^* (\mu)$.*

The proof is analogous to that of Proposition 3, and is therefore omitted. Also, both Propositions 3 and 4 can be extended to $\boldsymbol{\alpha}$-spending constraints but we omit formal statements for brevity.

3.1.3. *Weighted average power.* In the non-parametric setting, a weight function corresponds to a distribution over the tangent space $T(P_0)$. Any element $h \in T(P_0)$ can be expressed as $h = \langle \psi/\sigma, h \rangle \psi/\sigma + \tilde{h}$, where $\tilde{h}$ is orthogonal to ψ (i.e., $\left\langle \psi, \tilde{h} \right\rangle = 0$). We focus on a specific class of weight functions, $m(\cdot)$ that are multiplicatively separable with respect to the weights they assign to $\langle \psi, h \rangle$ and $\tilde{h}$. Specifically, each $m(\cdot)$ takes the form $w \times \rho$ where w is a distribution over $\mu := \langle \psi, h \rangle$ and ρ is a distribution over $\tilde{h} \in T(P_0)$.[7]

[7]Since $T(P_0)$ is a Hilbert space, we can alternatively think of ρ as a distribution over the space of square integrable sequences.

Multiplicative separability can be motivated by the invariance requirement that inference over μ should be unaffected by the value of the 'nuisance parameter' $\tilde{h}$.

In analogy with (2.6), the test that maximizes weighted average power in the limit experiment from Section 3.1.1 is given by

$$\varphi_w^*(\tau, x(\tau)) = \mathbb{I}\left\{\int \exp\left(\frac{\mu}{\sigma}x(\tau) - \frac{\mu^2}{2\sigma^2}\tau\right) dw(\mu) > \gamma\right\}.$$

The following proposition demonstrates that, for any weight function of the form $m(\cdot)$ as described above, the finite sample weighted average power of any test is bounded above by that of $\varphi_w^*(\cdot)$ in the limit experiment:

Proposition 5. *Suppose Assumption 3 holds. Let $\beta(h)$ denote the power of some asymptotically level-α test, φ_n, of $H_0 : \mu = 0$. Then, for any weight function $m(\cdot)$ that is of the multiplicatively separable form $w \times \rho$ described above, we have $\limsup_{n\to\infty} \int \beta_n(h)dm(h) \leq \int \beta^*(\mu)\, dw(\mu)$, where $\beta^*(\mu)$ is the power function of $\varphi_w^*(\cdot)$ in the limit experiment.*

3.2. **Computation of critical values.** In contrast to the parametric framework, the reference distribution P_0 in our non-parametric setup primarily serves as a theoretical construct. It is more precise to consider P_0 as a member of an equivalence class $\mathcal{P}_0$ of distributions, all of which yield the same asymptotic joint distribution for $(\hat{\tau}, x_n(\hat{\tau}))$. As a result, we cannot apply the strategy used in Section 2.5 of determining critical values by simulating experimental data under P_0.

Instead, we leverage the fact that for any $P_0 \in \mathcal{P}_0$, the functional central limit theorem ensures $x_n(\cdot) \xrightarrow[P_0]{d} x(\cdot)$, where $x(\cdot)$ represents a standard d-dimensional Brownian motion. Following Assumption 3(ii) and the continuous mapping theorem, it then follows that the limit distribution of $(\hat{\tau}, x_n(\hat{\tau}))$ under any $P_0 \in \mathcal{P}_0$ is given by $(\tau, x(\tau))$, where $\tau := \tau(x(\cdot))$. Notably, the distribution of $(\tau, x(\tau))$ remains invariant with respect to the choice of P_0. This allows us to select P_0 from any convenient distributional family, provided it satisfies $\mathbb{E}_{P_0}[\psi(Y_i)] = 0$ and $\mathbb{E}_{P_0}[\psi(Y_i)^2] = \sigma^2$. For example, when $\mu(\cdot)$ represents the mean, $\psi(Y_i) = Y_i$, and P_0 can be chosen as the normal distribution $\mathcal{N}(0, \hat{\sigma}^2)$, where $\hat{\sigma}^2$ is any consistent estimator of σ^2. The joint distribution of $(\tau, x(\tau))$ can then be approximated by simulating $(\hat{\tau}, x_n(\hat{\tau}))$ under this chosen P_0. Although the critical values estimated in this manner, unlike those in Section 2.5, do not yield exactly sized tests, they do converge in probability to the asymptotic critical values.

Alternatively, as discussed in Section 2.5, the distribution of $(\tau, x(\tau))$ can also be determined analytically in several applications, offering a complementary approach to simulation-based methods.

3.3. **Attaining the power bounds.** It follows by similar reasoning as in Section 2.8 (but now using parametric sub-models) that we can attain the power bounds $\beta^*(\cdot), \tilde{\beta}^*(\cdot)$ by employing plug-in versions of the corresponding UMP tests in the limit experiment. This process simply involves replacing $\tau, x(\tau)$ with $\hat{\tau}, x_n(\hat{\tau})$, and, if necessary, substituting the asymptotic critical values with the estimates described in Section 3.2. As $x_n(\hat{\tau})$ is dependent on the variance σ, we must replace σ with a consistent estimate, as detailed in Appendix B.1.

3.4. **Two-sample tests.** In many sequential experiments it is common to test two treatments simultaneously. For instance, both the costly sampling and group sequential trial examples from Section 1.2 involve two treatments that are sampled in fixed proportions (but the stopping time is history dependent). We may then be interested in conducting inference on the difference between some regular functionals of the two treatments. A salient example of this is inference on the expected treatment effect.

To make matters precise, let $a \in \{0, 1\}$ index the two treatments, and $P^{(1)}, P^{(0)}$ denote the corresponding outcome distributions. Suppose that at each period, the DM samples treatment 1 at some fixed proportion π. It is without loss of generality to suppose that the outcomes from the two treatments are independent as we can only ever observe the effect of a single treatment. We are interested in conducting inference on the difference, $\mu(P^{(1)}) - \mu(P^{(0)})$, where $\mu(\cdot)$ is some regular scalar functional of the data distribution, e.g., its mean.

Let $P_0^{(1)}, P_0^{(0)}$ denote some reference probability distributions on the boundary of the null hypothesis so that $\mu(P_0^{(1)}) - \mu(P_0^{(0)}) = 0$. Following Van der Vaart (2000, Section 25.6), we analyze the power of tests against smooth one-dimensional sub-models of the form $\left\{ \left(P_{s,h_1}^{(1)}, P_{s,h_0}^{(0)} \right) : s \leq \eta \right\}$ for some $\eta > 0$, where $h_a(\cdot)$ is a measurable function satisfying

$$\int \left[\frac{\sqrt{dP_{s,h_a}^{(a)}} - \sqrt{dP_0^{(a)}}}{s} - \frac{1}{2} h_a \sqrt{dP_0^{(a)}} \right]^2 d\nu \to 0 \text{ as } s \to 0. \tag{3.5}$$

As before, the set of all possible h_a satisfying $\int h_a dP_0^{(a)} = 0$ and $\int h_a^2 dP_0^{(a)} < \infty$ forms a tangent space $T(P_0^{(a)})$. This is a subset of the Hilbert space $L^2(P_0^{(a)})$, endowed with the inner product $\langle f, g\rangle_a = \mathbb{E}_{P_0^{(a)}}[fg]$ and norm $\|f\|_a = \mathbb{E}_{P_0^{(a)}}[f^2]^{1/2}$. Let $\psi_a \in T(P_0^{(a)})$ denote the efficient influence function satisfying

$$\frac{\mu(P_{s,h_a}^{(a)}) - \mu(P_0^{(a)})}{s} - \langle \psi_a, h_a\rangle_a = o(s) \tag{3.6}$$

for any $h_a \in T(P_0^{(a)})$. Denote $\sigma_a^2 = \mathbb{E}_{P_0^{(a)}}[\psi_a^2]$. The sufficient statistic here is the differenced efficient influence function process

$$x_n(t) := \frac{1}{\sigma}\left(\frac{1}{\pi\sqrt{n}}\sum_{i=1}^{\lfloor n\pi t\rfloor}\psi_1(Y_i^{(1)}) - \frac{1}{(1-\pi)\sqrt{n}}\sum_{i=1}^{\lfloor n(1-\pi)t\rfloor}\psi_0(Y_i^{(0)})\right), \tag{3.7}$$

where $\sigma^2 := \left(\frac{\sigma_1^2}{\pi} + \frac{\sigma_0^2}{1-\pi}\right)$. Note that the number of observations from each treatment at time t is $\lfloor n\pi t\rfloor, \lfloor n(1-\pi)t\rfloor$. The assumptions below are analogous to Assumption 3:

Assumption 4. *(i) The sub-models $\{P_{s,h_a}^{(a)}; h_a \in T(P_0^{(a)})\}$ satisfy (3.5). Furthermore, they admit an efficient influence function, $\psi_a(\cdot)$, such that (3.6) holds.*

(ii) The stopping time $\hat{\tau}$ is a continuous function of $x_n(\cdot)$ in the sense that $\hat{\tau} = \tau(x_n(\cdot))$, where $\tau(\cdot)$ satisfies the conditions for an extended continuous mapping theorem (Van Der Vaart and Wellner, 1996, Theorem 1.11.1).

As discussed in Section 5, Assumption 4(ii) is satisfied in the context of both costly sampling and group sequential trials.

Set $\mu_a := \mu(P^{(a)})$. In analogy with Section 3.1, we call a test, φ_n, of $H_0 : \mu_1 - \mu_0 = 0$ asymptotically level-α if

$$\sup_{\left\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1 - \langle\psi_0,h_0\rangle_0 = 0\right\}} \limsup_n \int \varphi_n dP_{nT,\boldsymbol{h}} \le \alpha. \tag{3.8}$$

Similarly, a test, φ_n, of $H_0 : \mu_1 - \mu_0 = 0$ vs $H_1 : \mu_1 - \mu_0 \neq 0$ is asymptotically unbiased if

$$\sup_{\left\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1 - \langle\psi_0,h_0\rangle_0 = 0\right\}} \limsup_n \int \varphi_n dP_{nT,\boldsymbol{h}} \le \alpha, \text{ and}$$
$$\inf_{\left\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1 - \langle\psi_0,h_0\rangle_0 \neq 0\right\}} \liminf_n \int \varphi_n dP_{nT,\boldsymbol{h}} \ge \alpha. \tag{3.9}$$

Consider the limit experiment where one observes $x(\cdot) \sim \sigma^{-1}(\mu_1 - \mu_0)\cdot + W(\cdot)$ and a $\mathcal{F}_t \equiv \sigma\{x(s); s \le t\}$ adapted stopping time τ that is the weak limit of $\hat{\tau}$. Then, setting $\mu := \mu_1 - \mu_0$, define the power functions $\beta^*(\cdot)$ as in the previous section. The following result provides upper bounds on asymptotically level-α tests.

Proposition 6. *Suppose Assumption 4 holds. Let* $\beta_n(\boldsymbol{h})$ *the power of some asymptotically level-*α *test,* φ_n*, of* $H_0 : \mu_1 - \mu_0 = 0$ *against local alternatives* $P^{(1)}_{\delta_1/\sqrt{n},h_1} \times P^{(0)}_{\delta_0/\sqrt{n},h_0}$*. Then, for every* $\boldsymbol{h} \in T(P^{(1)}_0) \times T(P^{(0)}_0)$ *and* $\mu := \delta_1 \langle \psi_1, h_1 \rangle_1 - \delta_0 \langle \psi_0, h_0 \rangle_0$*,* $\limsup_{n\to\infty} \beta_n(\boldsymbol{h}) \leq \beta^*(\mu)$*.*

We prove Proposition 6 in Appendix A. As in Section 3.3, we can attain the power bound $\beta^*(\cdot)$ by employing plug-in versions of the corresponding UMP tests in the limit experiment. The sole difference is that $x_n(\hat{\tau})$ is now defined as (3.7).

The extension to unbiased and weighted average power tests is similar. We omit the formal statements for brevity.

4. Optimal tests in batched experiments

We now analyze sequential experiments with multiple treatments and where the sampling rule, i.e., the number of units allocated to each treatment, also changes over the course of the experiment. Since our results here draw on Hirano and Porter (2023), we restrict attention to batched experiments, where the sampling strategy is only allowed to be changed at some fixed, discrete set of times. However, as discussed in Section 4.1.1, we conjecture that our findings could extend to fully adaptive settings without modification.

Suppose there are K treatments under consideration. We take $K = 2$ to simplify the notation, but all our results extend to any fixed K. The outcomes, $Y^{(a)}$, under treatment $a \in \{0, 1\}$ are distributed according to some parametric model $\{P^{(a)}_{\theta^{(a)}}\}$. Here $\theta^{(a)} \in \mathbb{R}^d$ is some unknown parameter vector; we assume for simplicity that the dimension of $\theta^{(1)}, \theta^{(0)}$ is the same, but none of our results actually require this. It is without loss of generality to suppose that the outcomes from each treatment are independent conditional on $\theta^{(1)}, \theta^{(0)}$, as we only ever observe one of the two potential outcomes for any given observation. In the batch setting, the DM divides the observations into batches of size n, and registers a sampling rule $\{\hat{\pi}^{(a)}_j\}_j$ that prescribes the fraction of observations allocated to treatment a in batch j based on information from the previous batches $1, \ldots, j-1$. The experiment ends after J batches. It is possible to set $\pi^{(a)}_j = 0$ for some or all treatments (e.g., the experiment may be stopped early); we only require $\sum_a \hat{\pi}^{(a)}_j \leq 1$ for each j. We develop asymptotic representation theorems for tests of $H_0 : \theta = \Theta_0$ vs $H_1 : \theta \in \Theta_1$, where $\theta := (\theta^{(1)}, \theta^{(0)})$. Let $(\theta^{(1)}_0, \theta^{(0)}_0) \in \Theta_0$ denote some reference parameter in the null set.

Take $\hat{q}^{(a)}_j$ to be the proportion of observations allocated to treatment a up-to batch j, as a fraction of n. Also, let $Y^{(a)}_j$ denote the j-th observation of treatment a in the experiment. Clearly, any candidate test, $\delta(\cdot)$, is required to be

$$\sigma\left\{\left(Y_1^{(0)},\dots,Y_{nq_J^{(0)}}^{(0)}\right),\left(Y_1^{(1)},\dots,Y_{nq_J^{(1)}}^{(1)}\right)\right\}$$

measurable. As in the previous sections, we measure the performance of tests against local perturbations of the form $\{\theta_0^{(a)}+h_a/\sqrt{n}; h_a\in\mathbb{R}^d\}$. Let ν denote a dominating measure for $\{P_\theta^{(a)}:\theta\in\mathbb{R}^d, a\in\{0,1\}\}$, and set $p_\theta^{(a)}:=dP_\theta^{(a)}/d\nu$. We require $\{P_\theta^{(a)}\}$ to be quadratically mean differentiable (qmd):

Assumption 5. *The class* $\{P_\theta^{(a)}:\theta\in\mathbb{R}^d\}$ *is qmd around* $\theta_0^{(a)}$ *for each* $a\in\{0,1\}$*, i.e., there exists a score function* $\psi_a(\cdot)$ *such that for each* $h_a\in\mathbb{R}^d$,

$$\int\left[\sqrt{p^{(a)}_{\theta_0^{(a)}+h_a}}-\sqrt{p^{(a)}_{\theta_0^{(a)}}}-\frac{1}{2}h_a^\intercal\psi_a\sqrt{p_{\theta_0^{(a)}}}\right]^2 d\nu=o(|h_a|^2).$$

Furthermore, the information matrix $I_a:=\mathbb{E}_0[\psi_a\psi_a^\intercal]$ *is invertible for* $a\in\{0,1\}$.

Define $z_{j,n}^{(a)}(\hat{\pi}_j)$ as the standardized score process from each batch, where

$$z_{j,n}^{(a)}(t):=\frac{I_a^{-1/2}}{\sqrt{n}}\sum_{i=1}^{\lfloor nt\rfloor}\psi_a(Y_{i,j}^{(a)})$$

for each $t\in[0,1]$. Let $Y_{i,j}^{(a)}$ denote the i-th outcome observation from arm a in batch j. At each batch j, one can imagine that there is a potential set of outcomes, $\{\mathbf{y}_j^{(1)},\mathbf{y}_j^{(0)}\}$ with $\mathbf{y}_j^{(a)}:=\{Y_{i,j}^{(a)}\}_{i=1}^n$, that could be sampled from both arms, but only a sub-collection, $\{Y_{i,j}^{(a)}; i=1,\dots,n\hat{\pi}_j^{(a)}\}$, of these are actually sampled. Let $\boldsymbol{h}:=(h_1,h_0)$, take $P_{n,\boldsymbol{h}}$ to be the joint probability measure over

$$\{\mathbf{y}_1^{(1)},\mathbf{y}_1^{(0)},\dots,\mathbf{y}_J^{(1)},\mathbf{y}_J^{(0)}\}$$

when each $Y_{i,j}^{(a)}\sim P_{\theta_0^{(a)}+h_a/\sqrt{n}}$, and take $\mathbb{E}_{n,\boldsymbol{h}}[\cdot]$ to be its corresponding expectation. Then, by a standard functional central limit theorem,

$$z_{j,n}^{(a)}(t)\xrightarrow[P_{n,0}]{d}z(t);\ z(\cdot)\sim W_j^{(a)}(\cdot), \tag{4.1}$$

where $\{W_j^{(a)}\}_{j,a}$ are independent d-dimensional Brownian motions.

Our next assumption is a weak convergence requirement that is also employed by Hirano and Porter (2023).

Assumption 6. *The sequence*

$$\xi_n:=\left(\left(\hat{\pi}_1^{(1)},\hat{\pi}_1^{(0)},z_{1,n}^{(1)}(\hat{\pi}_1^{(1)}),z_{1,n}^{(0)}(\hat{\pi}_1^{(0)})\right),\dots,\left(\hat{\pi}_J^{(1)},\hat{\pi}_J^{(0)},z_{J,n}^{(1)}(\hat{\pi}_J^{(1)}),z_{J,n}^{(0)}(\hat{\pi}_J^{(0)})\right)\right)$$

converges weakly under $P_{n,0}$.

By Prohorov's theorem, ξ_n already converges on sub-sequences. Assumption 6 tightens this further and ensures that the sequence of experimental protocols indexed by $\left\{(\hat{\pi}_1^{(1)}, \hat{\pi}_1^{(0)}), \dots, (\hat{\pi}_J^{(1)}, \hat{\pi}_J^{(0)})\right\}$ has a well defined asymptotic limit. As with Assumption 2, this is a rather mild assumption: if the sub-sequences converged to different quantities, they should really then be identified as indexing different protocols.

4.1. **Asymptotic representation theorem.** Consider a limit experiment where $\boldsymbol{h} := (h_1, h_0)$ is unknown, and for each batch j, one observes the stopped process $z_j^{(a)}(\pi_j^{(a)})$, where

$$z_j^{(a)}(t) := I_a^{1/2} h_a t + W_j^{(a)}(t), \tag{4.2}$$

and $\{W_j^{(a)}; j = 1, \dots, J; a = 0, 1\}$ are independent Brownian motions. Each $\pi_j^{(a)}$ is required to satisfy $\sum_a \pi_j^{(a)} \leq 1$ and also to be

$$\sigma\left\{(z_1^{(1)}, z_1^{(0)}, U_1), \dots, (z_{j-1}^{(1)}, z_{j-1}^{(0)}, U_{j-1})\right\}$$

measurable, where $U_j \sim \text{Uniform}[0,1]$ is exogenous to all the past values $\left\{z_{j'}^{(a)}, U_{j'} : j' < j\right\}$. Let φ denote a test statistic for $H_0 : h = 0$ that depends only on: (i) $q_a = \sum_j \pi_j^{(a)}$, i.e., the fraction of times each arm was pulled; and (ii) $x_a = \sum_j z_j^{(a)}(\pi_j^{(a)})$, i.e., the cumulative score process for each arm. Also, let $\mathbb{P}_{\boldsymbol{h}}$ denote the joint probability measure over $\{z_j^{(a)}(\cdot); a \in \{0, 1\}, j \in \{1, \dots, J\}\}$ when each $z_j^{(a)}(\cdot)$ is distributed as in (4.2), and take $\mathbb{E}_{\boldsymbol{h}}[\cdot]$ to be its corresponding expectation.

The following theorem shows that the power function of any test φ_n in the original testing problem can be matched by one such test, φ, in the limit experiment.

Theorem 3. *Suppose Assumptions 5 and 6 hold. Let φ_n be some test function in the original batched experiment, and $\beta_n(\boldsymbol{h})$, its power against $P_{n,\boldsymbol{h}}$. Then:*
(i) (Hirano and Porter, 2023) There exists a batched policy function $\pi = \{\pi_j^{(a)}\}_j$ and processes $\{z_j^{(a)}(\cdot)\}_{j,a}$ defined on the limit experiment for which

$$\left(\left(\hat{\pi}_1^{(1)}, \hat{\pi}_1^{(0)}, z_{1,n}^{(1)}(\hat{\pi}_1^{(1)}), z_{1,n}^{(0)}(\hat{\pi}_1^{(0)})\right), \dots, \left(\hat{\pi}_J^{(1)}, \hat{\pi}_J^{(0)}, z_{J,n}^{(1)}(\hat{\pi}_J^{(1)}), z_{J,n}^{(0)}(\hat{\pi}_J^{(0)})\right)\right)$$
$$\xrightarrow[P_{n,0}]{d} \left(\left(\pi_1^{(1)}, \pi_1^{(0)}, z_1^{(1)}(\pi_1^{(1)}), z_1^{(0)}(\pi_1^{(0)})\right), \dots, \left(\pi_J^{(1)}, \pi_J^{(0)}, z_J^{(1)}(\pi_J^{(1)}), z_J^{(0)}(\pi_J^{(0)})\right)\right).$$

(ii) Suppose that $\beta_n(\boldsymbol{h})$ converges point-wise for each $\boldsymbol{h}$. Then, there exists a test φ in the limit experiment depending only on q_1, q_0, x_1, x_0 such that $\beta_n(\boldsymbol{h}) \to \beta(\boldsymbol{h})$ for every $\boldsymbol{h} \in \mathbb{R}^d \times \mathbb{R}^d$, where $\beta(\boldsymbol{h}) := \mathbb{E}_{\boldsymbol{h}}[\varphi]$ is the power of φ in the limit experiment.

The first part of Theorem 3 is due to Hirano and Porter (2023); we only modify the terminology slightly. We require this result to be able to define the quantities q_1, q_0, x_1, x_0 in the limit experiment. The results of Le Cam (1979), used previously in Theorem 1, are not directly applicable here: while they provide a representation theorem for x_a, q_a from each arm a, this is not enough to capture the joint distribution of (x_1, q_1) and (x_0, q_0).

Hirano and Porter (2023) already show that any sequence of tests φ_n can be asymptotically matched by a test φ in the limit experiment that is measurable with respect to $\left\{(z_1^{(1)}, z_1^{(0)}, U_1), \ldots, (z_J^{(1)}, z_J^{(0)}, U_J)\right\}$. The novel contribution here lies in the second part of Theorem 3, which demonstrates that further dimension reduction is possible. A straightforward application of Hirano and Porter (2023) would require sufficient statistics that increase linearly with the number of batches, resulting in a vector of dimension $2dJ+1$ (the uniform random variables $U_1, \ldots, U_J$ can be subsumed into a single $U \sim \text{Uniform}[0,1]$). In contrast, we show that for testing it is sufficient to condition only on q_1, q_0, x_1, x_0, which have a fixed dimension $2d+2$ (or $2d+1$ if we impose $q^{(1)} + q^{(0)} = J$). This represents a significant reduction in dimensionality. Notably, this result does not directly follow from the first part of the theorem, as, similar to the discussion in Section 2.3.1, (q_1, q_0, x_1, x_0), are not sufficient statistics in the conventional sense. Consequently, to prove the second part of Theorem 3, we leverage the representation theorem of Hirano and Porter (2023), combine it with the SLAN property (2.3), and apply a change-of-measure argument that extends Le Cam's third lemma.

4.1.1. *An alternative representation of the limit experiment.* From the distribution of $z_j^{(a)}(\cdot)$ given in (4.2), it is easy to verify that

$$z_j^{(a)}(\pi_j^{(a)}) \sim I_a^{1/2} h_a \pi_j^{(a)} + W_j^{(a)}(\pi_j^{(a)}).$$

Combined with the definition $q_a = \sum_j \pi_j^{(a)}$ and the fact $\{W_j^{(a)}; j = 1, \ldots, J; a = 0, 1\}$ are independent Brownian motions, we obtain

$$x_a = \sum_j z_j^{(a)}(\pi_j^{(a)}) \sim I_a^{1/2} h_a q_a + W_a(q_a), \tag{4.3}$$

where $W_1(\cdot).W_0(\cdot)$ are standard d-dimensional Brownian motions that are again independent of each other. In view of the above, we can alternatively think of the limit experiment as observing $\{q_a\}_a$ along with $\{x_a\}_a$, with the latter distributed as in (4.3). The benefit of this formulation is that it does not depend on the number of batches.

This reformulation suggests that the second part of Theorem 3 may still hold in a continuous experimentation setting. However, the first part of Theorem 3 cannot be directly extended to fully adaptive experimentation. Specifically, it is possible to construct a sequence of fully adaptive sampling rules that do not converge to any well-defined (i.e., Lebesgue measurable) policy rule in continuous time.[8] In contrast, x_a, q_a can always be defined independently of a limiting policy rule, as they represent the limiting versions of the stopped likelihood ratio process for treatment a and the fraction of time that treatment was sampled. We therefore conjecture that an asymptotic representation theorem can be established for $\{(x_a, q_a)\}_a$ independently of the first part of Theorem 3, and that the second part of Theorem 3 can be extended to the fully adaptive setting without modification. The details of this extension are, however, left for future work.

4.2. **Characterization of optimal tests in the limit experiment.** It is generally unrealistic in batched sequential experiments for the sampling rule to depend on fewer statistics than q_1, q_0, x_1, x_0. Consequently, we do not have sharp results for testing linear combinations as in Proposition 1. We do, however, have analogues to the other results in Section 2.6.

4.2.1. *Power envelope.* Consider testing $H_0 : \boldsymbol{h} = 0$ vs $H_1 : \boldsymbol{h} = \boldsymbol{h}_1$ in the limit experiment. By the Neyman-Pearson lemma, and the Girsanov theorem applied on (4.3), the optimal test is given by

$$\varphi^*_{\boldsymbol{h}_1} = \mathbb{I}\left\{ \sum_{a\in\{0,1\}} \left(h_a^\intercal I_a^{1/2} x_a - \frac{q_a}{2} h_a^\intercal I_a h_a \right) > \gamma_{h_1} \right\}, \tag{4.4}$$

where $\gamma_{\boldsymbol{h}_1}$ is chosen such that $\mathbb{E}_0[\varphi^*_{\boldsymbol{h}_1}] = \alpha$. Take $\beta^*(\boldsymbol{h}_1)$ to be the power function of $\varphi^*_{\boldsymbol{h}_1}$ against $H_1 : \boldsymbol{h} = \boldsymbol{h}_1$. Theorem 3 shows that $\beta^*(\cdot)$ is an asymptotic power envelope for any test of $H_0 : \theta = \theta_0$ in the original experiment.

4.2.2. *Unbiased tests.* Suppose $\varphi(q_1, q_0, x_1, x_0)$ is an unbiased test of $H_0 : \boldsymbol{h} = 0$ vs $H_1 : \boldsymbol{h} \neq 0$ in the limit experiment. Then, in analogy with Proposition 2, it needs to satisfy the following property:

Proposition 7. *Any unbiased test of $H_0 : \boldsymbol{h} = 0$ vs $H_1 : \boldsymbol{h} \neq 0$ in the limit experiment must satisfy $\mathbb{E}_0[x_a \varphi(q_1, q_0, x_1, x_0)] = 0$ for all a, where $x_a \sim W_a(q_a)$ under $\mathbb{P}_0$.*

[8]This is related to the well known failure of the measurable selection theorem in stochastic optimal control, see, e.g., Bertsekas (2012, Appendix A) for a discussion of this in the discrete time setting.

4.2.3. *Weighted average power.* Let $w(\cdot)$ denote a weight function over alternatives $\boldsymbol{h} \neq 0$. Then, the uniquely optimal test of $H_0 : \boldsymbol{h} = 0$ that maximizes weighted average power over $w(\cdot)$ is given by

$$\varphi_w^* = \mathbb{I}\left\{\int \exp\left\{\sum_{a\in\{0,1\}}\left(h_a^\intercal I_a^{1/2}x_a - \frac{q_a}{2}h_a^\intercal I_a h_a\right)\right\}dw(\boldsymbol{h}) > \gamma\right\}.$$

The value of γ is chosen to satisfy $\mathbb{E}_0[\varphi_w^*] = \alpha$. In practice, it can be computed by simulation.

4.3. **Asymptotically optimal tests.** For each $a \in \{0,1\}$, let $\hat{q}_a = \sum_j \hat{\pi}_j^{(a)}$ and $\hat{x}_a = \sum_j z_{j,n}^{(a)}(\hat{\pi}_j^{(a)})$ denote the finite sample counterparts of q_a, x_a. As before, we can construct finite sample versions of the various optimal tests described in Section 4.2 by replacing q_1, q_0, x_1, x_0 with $\hat{q}_1, \hat{q}_0, \hat{x}_1, \hat{x}_0$. The critical values can be obtained by simulating the joint distribution of $\hat{q}_1, \hat{q}_0, \hat{x}_1, \hat{x}_0$ under the null (i.e., by sampling the outcomes for each treatment a from the corresponding null distribution $P^{(a)}_{\theta_0^{(a)}}$).

4.4. **Non-parametric tests.** For the non-parametric setting, we make use of the same notation as in Section 3.4. We are interested in conducting inference on some regular vector of functionals, $\left(\mu(P^{(1)}), \mu(P^{(0)})\right)$, of the outcome distributions $P^{(1)}, P^{(0)}$ for the two treatments. To simplify matters, we take $\mu_a := \mu(P^{(a)})$ to be scalar. The definition of asymptotically level-α and unbiased tests is unchanged from (3.8) and (3.9).

Let ψ_a, σ_a be defined as in Section 3.4. Set

$$z_{j,n}^{(a)} := \frac{1}{\sigma_a\sqrt{n}}\sum_{i=1}^{\lfloor nt\rfloor}\psi_a(Y_{i,j}^{(a)}),$$

and for $k = 1, \dots, K$, take $s_n(k) = \{\hat{q}_1(k), \hat{q}_0(k), \hat{x}_1(k), \hat{x}_0(k)\}$ to be a vector of state variables containing

$$\hat{q}_a(k) := \sum_{j=1}^{k}\hat{\pi}_j^{(a)}, \text{ and } \hat{x}_a(k) := \sum_{j=1}^{k} z_{n,j}^{(a)}(\hat{\pi}_j^{(a)}).$$

Assumption 7. *(i) The sub-models $\{P^{(a)}_{s,h_a}; h_a \in T(P^{(a)}_0)\}$ satisfy (3.5). Furthermore, they admit an efficient influence function, ψ_a, such that (3.6) holds.*

(ii) The sampling rule $\hat{\pi}_{j+1}$ in batch j is a continuous function of $s_n(j)$ in the sense that $\hat{\pi}_{j+1} = \pi_{j+1}(s_n(j))$, where $\pi_{j+1}(\cdot)$ satisfies the conditions for an extended continuous mapping theorem (Van Der Vaart and Wellner, 1996, Theorem 1.11.1) for each $j = 0, \dots, K-1$.

Assumption 7(i) is standard. Assumption 7(ii) implies that the sampling rule depends on a vector of four state variables. This is in contrast to the single sufficient statistic used in Section 3.4. We impose Assumption 7(ii) as it is more realistic; many commonly used algorithms, e.g., Thompson sampling, depend on all four statistics. The assumption still imposes a dimension reduction as it requires the sampling rule to be independent of the data conditional on knowing $s_n(\cdot)$. In practice, any Bayes or minimax optimal algorithm would only depend on $s_n(\cdot)$ anyway, as noted in Adusumilli (2021). In fact, we are not aware of any commonly used algorithm that requires more statistics beyond these four.

The reliance of the sampling rule on the vector $s_n(\cdot)$ implies that the optimal test should also depend on the full vector, and cannot be reduced further. The relevant limit experiment is the one described in Section 4.1.1, with μ_a replacing h_a. Also, let

$$\varphi_{\bar{\mu}_1,\bar{\mu}_0} = \mathbb{I}\left\{ \sum_{a\in\{0,1\}} \left(\frac{\bar{\mu}_a}{\sigma_a} x_a - \frac{q_a}{2\sigma_a^2}\bar{\mu}_a^2 \right) \geq \gamma_{\bar{\mu}_1,\bar{\mu}_0} \right\}$$

denote the Neyman-Pearson test of $H_0 : (\mu_1, \mu_0) = (0, 0)$ vs $H_1 : (\mu_1, \mu_0) = (\bar{\mu}_1, \bar{\mu}_0)$ in the limit experiment, with $\gamma_{\bar{\mu}_1,\bar{\mu}_0}$ determined by the size requirement. Take $\beta^*(\bar{\mu}_1, \bar{\mu}_0)$ to be its corresponding power.

Proposition 8. *Suppose Assumption 7 holds. Let $\beta_n(\boldsymbol{h})$ the power of some asymptotically level-α test, φ_n, of $H_0 : (\mu_1, \mu_0) = (0, 0)$ against local alternatives $P^{(1)}_{\delta_1/\sqrt{n},h_1} \times P^{(0)}_{\delta_0/\sqrt{n},h_0}$. Then, for every $\boldsymbol{h} \in T(P_0^{(1)}) \times T(P_0^{(0)})$ and $\mu_a := \delta_a \langle \psi_a, h_a \rangle_a$ for $a \in \{0, 1\}$, $\limsup_{n\to\infty} \beta_n(\boldsymbol{h}) \leq \beta^* (\mu_1, \mu_0)$.*

Proposition 8 describes the power envelope for testing that the parameter vector (μ_1, μ_0) takes on a given value. Suppose, however, that one is only interested in providing inference for single component of that vector, say μ_1. Then μ_0 is a nuisance parameter under the null, and one would need to employ the usual strategies for getting rid of the dependence on μ_0, e.g., through conditional inference or minimax tests. We leave the discussion of these possibilities for future research.

5. Applications

In this section, we apply the methods developed in this paper to the various examples of sequential experiments described in Section 1.2.

5.1. **Horizontal boundary designs.** As a first illustration of our methods, consider the class of horizontal boundary designs with a fixed sampling rule, π, and the stopping time $\hat{\tau} = \inf\{t : |x_n(t)| \geq \gamma\}$, where $x_n(t)$ is defined as in (3.7). As a concrete example, suppose μ_1, μ_0 denote the mean values of outcomes from each treatment, with σ_1, σ_0 their corresponding standard deviations. If the goal of the experiment is to determine the treatment with the largest mean while minimizing the number of samples, which are costly, then, as shown in Adusumilli (2022), the minimax optimal sampling strategy is the Neyman allocation $\pi_1^* = \sigma_1/(\sigma_1+\sigma_0)$, and optimal stopping rule is $\hat{\tau} = \inf\{t : |x_n(t)| \geq \gamma\}$ with the efficient influence functions $\psi_1(Y) = \psi_0(Y) = Y$.

We aim to test the null hypothesis of no treatment effect, $H_0 : \mu_1 - \mu_0 = 0$ against the alternative hypothesis $H_1 : \mu_1 - \mu_0 \neq 0$. Let $F_\mu(\cdot)$ denote the distribution of τ in the limit experiment where $x(t) \sim \sigma^{-1}\mu t + W(t)$ and $\tau = \inf\{t : |x(t)| \geq \gamma\}$. Utilizing our characterization of unbiased tests (see, 2.5), we demonstrate in Appendix B.2 that the optimal unbiased test in the limit experiment depends solely on τ, and is given by $\varphi^* = \mathbb{I}\{\tau \leq F_0^{-1}(\alpha)\}$. Consequently, the test $\hat{\varphi} := \mathbb{I}\{\hat{\tau} \leq F_0^{-1}(\alpha)\}$ constitutes the uniformly most powerful (UMP) asymptotically unbiased test in this setting. We summarize this result below:

Lemma 1. *Consider the sequential experiment described above with a fixed sampling rule π and stopping time $\hat{\tau} = \inf\{t : |x_n(t)| \geq \gamma\}$. The test, $\hat{\varphi} = \mathbb{I}\{\hat{\tau} \leq F_0^{-1}(\alpha)\}$, is the UMP asymptotically unbiased test (in the sense that it attains the upper bound in Proposition 3) of $H_0 : \mu_1 = \mu_0$ vs $H_1 : \mu_1 \neq \mu_0$ in this experiment.*

5.1.1. *Numerical Illustration.* To assess the finite sample performance of $\hat{\varphi}$, we conducted Monte Carlo simulations with the following setup: $Y_i^{(1)} = \delta + \epsilon_i^{(1)}$ and $Y_i^{(0)} = \epsilon_i^{(0)}$, where where $\epsilon_i^{(1)}, \epsilon_i^{(0)} \sim \sqrt{3} \times \text{Uniform}[-1, 1]$. We set the threshold γ to 0.536, corresponding to a sampling cost of $c = 1$ per observation in the costly sampling framework, and treatments were assigned in equal proportions ($\pi = 1/2$). It is important to note that the choice of error distribution is specific to the simulation and does not imply that the test assumes knowledge of this distribution. The test operates under a non-parametric framework.

Figure 5.1, Panel A illustrates the test size for various sample sizes n at the nominal 5% significance level. We used the simulation-based method outlined in Section 3.2 to estimate critical values. The results show that even for relatively small sample sizes, the test size approximates the nominal level. For comparison, we also include the size of the naive two-sample test, which, due to its adaptive stopping rule, is invalid and shows an actual size close to 9%.

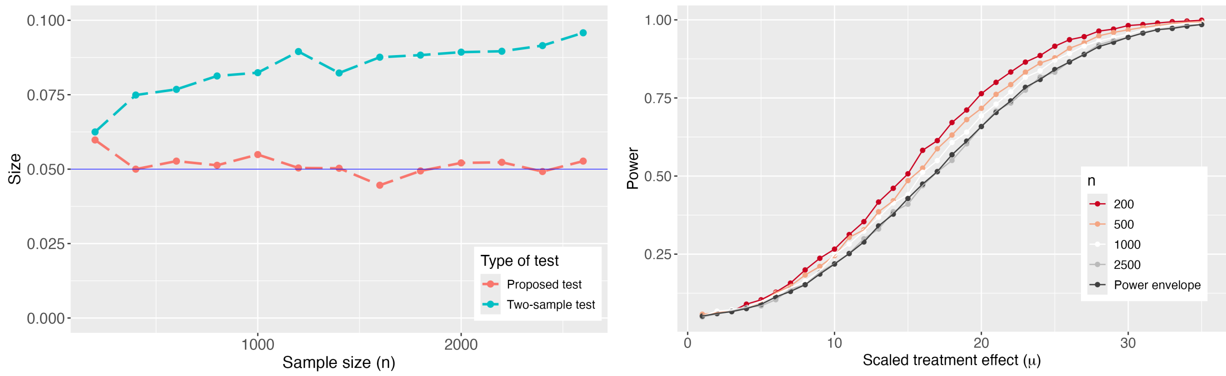

A: Size B: Power function

Note: Panel A plots the size of $\hat{\varphi}$ along with that of the standard two-sample test at the nominal 5% level (solid blue line) when the errors are drawn from a $\sqrt{3} \times \text{Uniform}[-1, 1]$ distribution for each treatment. Panel B plots the finite sample power envelopes of $\hat{\varphi}$ under different n, along with asymptotic power envelope for unbiased tests. The scaled treatment effect is defined as $\mu = \sqrt{n}|\delta|$.

FIGURE 5.1. Finite sample performance of $\hat{\varphi}$ under horizontal boundary designs

Panel B of the same figure plots the finite sample power functions for $\hat{\varphi}$ under different n. The power is computed against local alternatives; the reward gap in the figure is the scaled one, $\mu = \sqrt{n}|\delta|$. But for any given n, the actual difference in mean outcomes is $\mu/\sqrt{n}$. The same plot also displays the asymptotic power envelope for unbiased tests, obtained as the power function of the best unbiased test, $\varphi^* = \mathbb{I}\{\tau \le F_0^{-1}(\alpha)\}$, in the limit experiment. Even for small samples, the power function of $\hat{\varphi}$ is close to the asymptotic upper bound.

5.2. **Group sequential experiments.** In this application, we suggest methods for inference on treatment effects following group sequential experiments. To simplify matters, suppose that the researchers assign the two treatments with equal probability in each stage. Let μ_1, μ_0 denote the expectation of outcomes from the two treatments. Also, take $x_n(\cdot)$ to be the scaled difference in sample means, i.e., it is the quantity defined in (3.7) with $\psi_1(Y) = \psi_0(Y) = Y$. While there are a number of different group sequential designs, see, e.g., Wassmer and Brannath (2016) for a textbook overview, the general construction is that the experiment is terminated at the end of stage t if $x_n(t)$ is outside some interval $\mathcal{I}_t$. The stopping time $\hat{\tau}$ thus satisfies $\{\hat{\tau} > t-1\} \equiv \cap_{l=1}^{t-1} \{x_n(l) \in \mathcal{I}_l\}$. The intervals $\{\mathcal{I}_t\}_{t=1}^T$ are pre-determined and chosen by balancing various ethical, cost and power criteria. We take them as given.

We are interested in testing the drifting hypotheses $H_0 : \mu_1 - \mu_0 = \bar{\mu}/\sqrt{n}$ vs $H_1 : \mu_1 - \mu_0 > \bar{\mu}/\sqrt{n}$ at some spending level $\boldsymbol{\alpha}$ that is chosen by the experimenter.[9] We can then invert these tests to obtain one-sided confidence intervals for the treatment effect $\mu_1 - \mu_0$. The limit experiment in this setting consists of observing $x(t) \sim \sigma^{-1}\mu t + W(t)$, where $\mu := \mu_1 - \mu_0$, along with a discrete stopping time $\tau \in \{1, \dots, T\}$ such that $\{\tau > t - 1\}$ if and only if $x(l) \in \mathcal{I}_l$ for all $l = 1, \dots, t-1$. Let $\mathbb{P}_\mu(\cdot)$ denote the induced probability measure over the sample paths of $x(\cdot)$ between 0 and T, and $\mathbb{E}_\mu[\cdot]$ its corresponding expectation. In view of the results in Section 2.7, the optimal level-$\boldsymbol{\alpha}$ test $\varphi^*(\cdot)$ of $H_0 : \mu = \bar{\mu}$ vs $H_1 : \mu > \bar{\mu}$ in the limit experiment is given by

$$\varphi^*(\tau, x(\tau)) = \begin{cases} 1 & \text{if } \mathbb{P}_{\bar{\mu}}(\tau = t) \le \alpha_t \\ \mathbb{I}\{x(t) \ge \gamma(t)\} & \text{if } \mathbb{P}_{\bar{\mu}}(\tau = t) > \alpha_t, \end{cases} \tag{5.1}$$

where $\gamma(t)$ is chosen such that $\mathbb{E}_{\bar{\mu}}[\varphi^*(\tau, x(\tau))|\tau = t] = \alpha_t / \mathbb{P}_{\bar{\mu}}(\tau = t)$.

A finite sample version, $\hat{\varphi}$, of this test can be constructed by replacing $\tau, x(\tau)$ in φ^* with $\hat{\tau}, x_n(\hat{\tau})$. The resulting test would be asymptotically optimal under a suitable non-parametric version of the $\boldsymbol{\alpha}$-spending requirement; we refer to Appendix B.3 for the details and also for the proof that $\hat{\varphi}$ is asymptotically optimal, in the sense that it attains the power of φ^* in the limit experiment. A two-sided test for $H_0 : \mu_1 - \mu_0 = \bar{\mu}/\sqrt{n}$ vs $H_1 : \mu_1 - \mu_0 \neq \bar{\mu}/\sqrt{n}$ can be similarly constructed by imposing a conditional unbiasedness restriction as in Section 2.7.3.

5.2.1. *Numerical Illustration.* To illustrate the methodology, consider a group sequential trial based on the widely-used design of O'Brien and Fleming (1979), with $T = 2$ stages. This corresponds to setting $\mathcal{I}_1 = [-2.797, 2.797]$. We would like to test $H_0 : \mu_1 - \mu_0 = \bar{\mu}/\sqrt{n}$ vs $H_1 : \mu_1 - \mu_0 > \bar{\mu}/\sqrt{n}$ at the spending level $(\alpha/\mathbb{P}_{\bar{\mu}}(\tau = 1), \alpha/\mathbb{P}_{\bar{\mu}}(\tau = 2))$, equivalent to a conditional size constraint, $\mathbb{P}_{\bar{\mu}}(\varphi = 1|\tau = t) = \alpha \ \forall\ t$. Figure 5.2 Panel A plots the asymptotic critical values, $(\gamma(1), \gamma(2))$, for this test under $\alpha = 0.05$ and $\sigma_1 = \sigma_0 = 1$. Unsurprisingly, the thresholds are increasing in $\bar{\mu}$, but it is interesting to observe that they cross at some $\bar{\mu}$.

To describe the finite sample performance of this test, we ran Monte-Carlo simulations with $Y_i^{(1)} = \bar{\mu}/\sqrt{n} + \epsilon_i^{(1)}$ and $Y_i^{(0)} = \epsilon_i^{(0)}$ where $\epsilon_i^{(1)}, \epsilon_i^{(0)} \sim \sqrt{3} \times \text{Uniform}[-1, 1]$. The

[9]In most examples of group sequential designs, the intervals $\mathcal{I}_t$ are themselves chosen to maximize power under some $\bar{\boldsymbol{\alpha}}$-spending criterion, given the null of $\mu_1 = \mu_0$. In general, our $\boldsymbol{\alpha}$ here may be different from $\bar{\boldsymbol{\alpha}}$. Furthermore, we are interested in conducting inference on general null hypotheses of the form $H_0 : \mu_1 - \mu_0 = \bar{\mu}/\sqrt{n}$; these are different from the null hypothesis of no average treatment effect used to motivate the group sequential design.

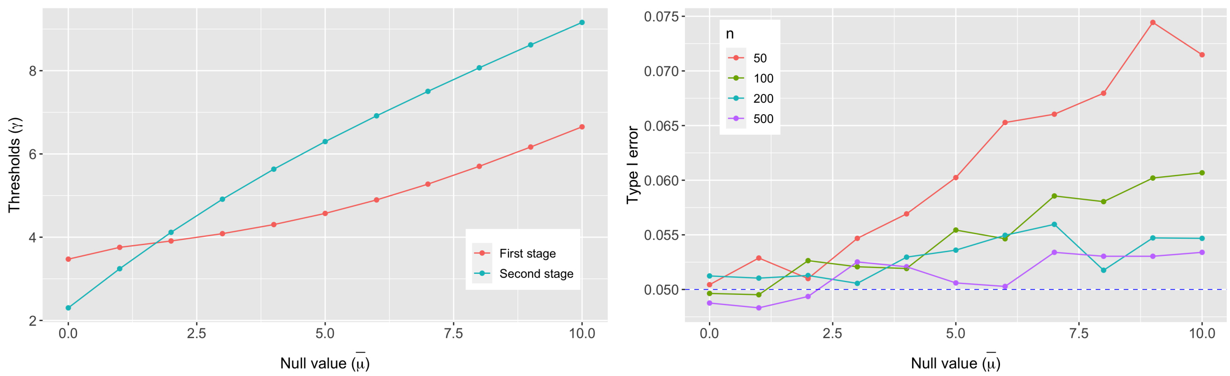

A: Critical values B: Finite sample size

Note: Panel A plots the threshold values in each stage for the optimal, one-sided, level-$\boldsymbol{\alpha}$ test, (5.1), at the $(0.05/\mathbb{P}_{\bar{\mu}}(\tau = 1), 0.05/\mathbb{P}_{\bar{\mu}}(\tau = 2))$ spending level. Panel B plots the overall type-I error in finite samples for different values of n and null values, $\bar{\mu}$, when the errors are drawn from a $\sqrt{3} \times \text{Uniform}[-1, 1]$ distribution for each treatment.

FIGURE 5.2. Testing in group sequential experiments

treatments were sampled in equal proportions ($\pi = 1/2$). Since σ_1, σ_0 are unknown in practice, we estimate them using data from the first stage. Figure 5.2, Panel B plots the overall size of the test (which is the sum of the α-spending values at each stage) for different values of n and $\bar{\mu}$ under the nominal $\boldsymbol{\alpha}$-spending level of $(0.05/\mathbb{P}_{\bar{\mu}}(\tau = 1), 0.05/\mathbb{P}_{\bar{\mu}}(\tau = 2))$. We see that the asymptotic approximation worsens for larger values of $\bar{\mu}$, but overall, the size is close to nominal even for relatively small values of n.

5.3. **Bandit experiments.** In Section 2.1, we described an example with a one-armed bandit. Here, we describe inferential procedures for the batched multi-armed bandit problem, focusing again on the Thompson-Sampling algorithm. For illustration, we employ $K = 2$ treatments and $J = 10$ batches. Let $(\bar{\mu}_1, \bar{\mu}_0)$ and (σ_1^2, σ_0^2) denote the population means and variances for each treatment. For simplicity, we take $\sigma_1^2 = \sigma_0^2 = 1$. The limit experiment can be described as follows: Suppose the decision maker (DM) employs the sampling rule $\pi_j^{(a)}$ in batch j. The DM then observes $Z_j^{(a)} \sim \mathcal{N}(\bar{\mu}_a \pi_a, \pi_a \sigma_a^2)$ for $a \in \{0, 1\}$ and updates the state variables x_a, q_a (which are initially set to 0) as

$$x_a \leftarrow x_a + Z_j^{(a)}, \quad q_a \leftarrow q_a + \pi_a.$$

Under an under-smoothed prior, recommended by Wager and Xu (2021), the Thompson sampling rule in batch $j+1$ is

$$\pi_{j+1}^{(1)} = \Phi\left(\frac{q_1^{-1} x_1 - q_0^{-1} x_0}{\sqrt{j/q_1 q_0}}\right).$$

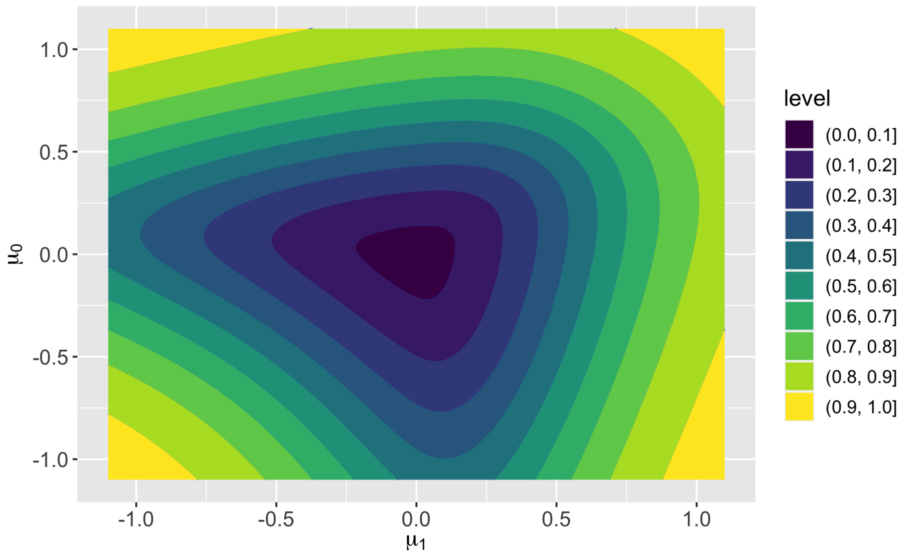

Note: The figure plots the asymptotic power envelope for any test of $H_0 : (\mu, \mu) = (0, 0)$ against different values (μ_1, μ_0) under the alternative.

FIGURE 5.3. Power envelope for Thompson-sampling with 10 batches

We set $\pi_1^{(a)} = 1/2$ for first batch. In what follows, we let $\mu_a := J\bar{\mu}_a$. We are interested in testing $H_0 : (\mu_1, \mu_0) = (0, 0)$.

Figure 5.3, Panel A plots the asymptotic power envelope for testing $H_0 : (\mu_1, \mu_2) = (0, 0)$. Clearly, the envelope is not symmetric; distinguishing $(a, 0)$ from $(0, 0)$ is easier than distinguishing $(-a, 0)$ from $(0, 0)$ for any $a > 0$. This is because of the asymmetry in treatment allocation under Thompson sampling; under $(-a, 0)$, treatment 1 is sampled more often than treatment 0 but the data from treatment 1 is uninformative for distinguishing $(-a, 0)$ from $(0, 0)$.

5.3.1. *Numerical illustration.* To determine the accuracy of our asymptotic approximations, we ran Monte-Carlo simulations with $Y_i^{(a)} = \mu_a + \epsilon_i^{(a)}$ where $\epsilon_i^{(1)}, \epsilon_i^{(0)} \sim \sqrt{3} \times$ Uniform$[-1, 1]$. Figure 5.4, Panel A plots the finite sample performance of the Neyman-Pearson tests in the limit experiment for testing $H_0 : (\mu_1, \mu_0) = (0, 0)$ vs $H_1 : (\mu_1, \mu_0) = (\mu, \mu)$ under various values of μ (due to symmetry, we only report the results for positive μ). Panel B repeats the same calculation, but against alternatives of the form $H_1 : (\mu, 0)$. As noted earlier, power is higher for $\mu > 0$ as opposed to $\mu < 0$. Both plots show that the asymptotic approximation is quite accurate even for n as small as 20 (note that the number of batches is 10, so this corresponds to 200 observations overall). The approximation is somewhat worse for testing $\mu < 0$; this is because Thompson-sampling allocates

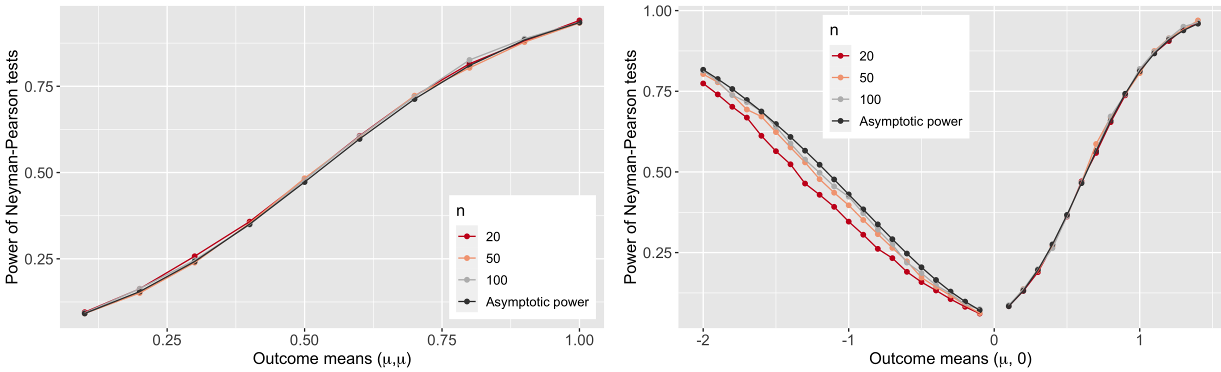

A: Power against $H_1 : (\mu, \mu)$ B: Power against $H_1 : (\mu, 0)$

Note: Panel A plots the finite sample power of Neyman-Pearson tests at the nominal 5% level (solid blue line) for testing $H_0 : (\mu_1, \mu_0) = (0, 0)$ against $H_1 : (\mu_1, \mu_0) = (\mu, \mu)$ when the errors are drawn from a $\sqrt{3} \times \text{Uniform}[-1, 1]$ distribution for each treatment. Panel B repeats the same calculation for alternatives of the form $H_1 : (\mu_1, \mu_0) = (\mu, 0)$. Both panels also display the asymptotic power envelope.

FIGURE 5.4. Finite sample performance of Neyman-Pearson tests in bandit experiments

much fewer units to treatment 0 in this instance, even though it is only data from this treatment that is informative for distinguishing the two hypotheses.

6. Conclusion

Conducting inference after sequential experiments is a challenging task. However, significant progress can be made by analyzing the optimal inference problem under an appropriate limit experiment. We showed that the data from any sequential experiment can be condensed into a finite number of sufficient statistics, while still maintaining the power of tests. Furthermore, we were able to establish uniquely optimal tests under reasonable constraints such as unbiasedness, $\boldsymbol{\alpha}$-spending and conditional power, in both parametric and non-parametric regimes. Taken together, these findings offer a comprehensive framework for conducting optimal inference following sequential experiments.

Despite these results, there are still several avenues for future research. While we believe that our results for experiments with adaptive sampling rules apply without batching, this needs be formally verified. Our characterization of uniquely optimal tests is also limited in this context, as $\boldsymbol{\alpha}$-spending restrictions are not feasible. Therefore, exploring other types of testing considerations such as invariance or conditional inference may be worthwhile. We believe that the techniques developed in this paper will prove useful for analyzing these other classes of tests.

References

K. Adusumilli, "Risk and optimal policies in bandit experiments," *arXiv preprint arXiv:2112.06363*, 2021.

——, "How to sample and when to stop sampling: The generalized wald problem and minimax policies," *arXiv preprint arXiv:2210.15841*, 2022.

S. Athey, K. Bergstrom, V. Hadad, J. C. Jamison, B. Özler, L. Parisotto, and J. D. Sama, "Shared decision-making," *Development Research*, 2021.

L. R. Baden *et al.*, "Efficacy and safety of the mrna-1273 sars-cov-2 vaccine," *New England journal of medicine*, vol. 384, no. 5, pp. 403–416, 2021.

D. Bertsekas, *Dynamic programming and optimal control: Volume II.* Athena scientific, 2012, vol. 1.

S. Choi, W. J. Hall, and A. Schick, "Asymptotically uniformly most powerful tests in parametric and semiparametric models," *The Annals of Statistics*, vol. 24, no. 2, pp. 841–861, 1996.

A. Deng, Y. Xu, R. Kohavi, and T. Walker, "Improving the sensitivity of online controlled experiments by utilizing pre-experiment data," in *Proceedings of the sixth ACM international conference on Web search and data mining*, 2013, pp. 123–132.

L. Fan and P. W. Glynn, "Diffusion approximations for thompson sampling," *arXiv preprint arXiv:2105.09232*, 2021.

K. J. Ferreira, D. Simchi-Levi, and H. Wang, "Online network revenue management using thompson sampling," *Operations research*, vol. 66, no. 6, pp. 1586–1602, 2018.

D. Fudenberg, P. Strack, and T. Strzalecki, "Speed, accuracy, and the optimal timing of choices," *American Economic Review*, vol. 108, no. 12, pp. 3651–84, 2018.

K. Gordon Lan and D. L. DeMets, "Discrete sequential boundaries for clinical trials," *Biometrika*, vol. 70, no. 3, pp. 659–663, 1983.

P. Grünwald, R. de Heide, and W. M. Koolen, "Safe testing," in *2020 Information Theory and Applications Workshop (ITA).* IEEE, 2020, pp. 1–54.

V. Hadad, D. A. Hirshberg, R. Zhan, S. Wager, and S. Athey, "Confidence intervals for policy evaluation in adaptive experiments," *Proceedings of the national academy of sciences*, vol. 118, no. 15, p. e2014602118, 2021.

W. J. Hall, "Analysis of sequential clinical trials," *Modern Clinical Trial Analysis*, pp. 81–125, 2013.

K. Hirano and J. R. Porter, “Asymptotic representations for sequential decisions, adaptive experiments, and batched bandits,” *arXiv preprint arXiv:2302.03117*, 2023.

S. R. Howard, A. Ramdas, J. McAuliffe, and J. Sekhon, “Time-uniform, nonparametric, nonasymptotic confidence sequences,” *The Annals of Statistics*, vol. 49, no. 2, 2021.

R. Johari, P. Koomen, L. Pekelis, and D. Walsh, “Always valid inference: Continuous monitoring of a/b tests,” *Operations Research*, vol. 70, no. 3, pp. 1806–1821, 2022.

M. Kasy and A. Sautmann, “Adaptive treatment assignment in experiments for policy choice,” 2019.

T. Lattimore and C. Szepesvári, *Bandit algorithms*. Cambridge University Press, 2020.

L. Le Cam, “A Reduction Theorem for Certain Sequential Experiments. II,” *The Annals of Statistics*, vol. 7, no. 4, pp. 847 – 859, 1979.

E. L. Lehmann and J. P. Romano, *Testing statistical hypotheses*. Springer, 2005, vol. 3.

P. C. O’Brien and T. R. Fleming, “A multiple testing procedure for clinical trials,” *Biometrics*, pp. 549–556, 1979.

A. Ramdas, P. Grünwald, V. Vovk, and G. Shafer, “Game-theoretic statistics and safe anytime-valid inference,” *arXiv preprint arXiv:2210.01948*, 2022.

J. P. Romano, “Optimal testing of equivalence hypotheses,” *The Annals of Statistics*, vol. 33, no. 3, pp. 1036–1047, 2005.

D. Russo and B. Van Roy, “An information-theoretic analysis of thompson sampling,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 2442–2471, 2016.

D. Russo, B. Van Roy, A. Kazerouni, I. Osband, and Z. Wen, “A tutorial on thompson sampling,” *arXiv preprint arXiv:1707.02038*, 2017.

US Food and Drug Admin., “Guidance document: Adaptive design clinical trials for drugs and biologics,” 2019. [Online]. Available: "https://www.fda.gov/regulatory-information/search-fda-guidance-documents/adaptive-design-clinical-trials-drugs-and-biologics-guidance-industry"

A. W. Van der Vaart, *Asymptotic statistics*. Cambridge university press, 2000.

A. W. Van Der Vaart and J. Wellner, *Weak convergence and empirical processes: with applications to statistics*. Springer Science & Business Media, 1996.

S. Wager and K. Xu, “Diffusion asymptotics for sequential experiments,” *arXiv preprint arXiv:2101.09855*, 2021.

A. Wald, “Sequential analysis,” Tech. Rep., 1947.

G. Wassmer and W. Brannath, *Group sequential and confirmatory adaptive designs in clinical trials*. Springer, 2016, vol. 301.

J. Whitehead, *The design and analysis of sequential clinical trials*. John Wiley & Sons, 1997.

K. Zhang, L. Janson, and S. Murphy, "Inference for batched bandits," *Advances in neural information processing systems*, vol. 33, pp. 9818–9829, 2020.

Appendix A. Proofs

A.1. **Proof of Theorem 1.** As noted earlier, the first part of the theorem follows from Le Cam (1979, Theorem 1). We therefore focus on proving the second claim. To this end, denote $\mathbf{y}_{nt} = (Y_1, \dots, Y_{nt})$. Defining

$$\ln \frac{dP_{nt,h}}{dP_{nt,0}}(\mathbf{y}_{nt}) = \sum_{i=1}^{\lfloor nt \rfloor} \ln \frac{dp_{\theta_0 + h/\sqrt{n}}}{dp_{\theta_0}}(Y_i),$$

we have by the SLAN property, (2.3), and Assumption 2 that

$$\ln \frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}) = h^{\intercal} I^{1/2} x_n(\hat{\tau}) - \frac{\hat{\tau}}{2} h^{\intercal} I h + o_{P_{nT,0}}(1).$$

Combining the above with the first part of the theorem gives

$$\ln \frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}) \xrightarrow[P_{nT,0}]{d} h^{\intercal} I^{1/2} x(\tau) - \frac{\tau}{2} h^{\intercal} I h, \tag{A.1}$$

where $x(\cdot)$ has the same distribution as d-dimensional Brownian motion.

Now, φ_n is tight since $\varphi_n \in [0, 1]$. Together with (A.1), this implies the joint

$$\left(\varphi_n, \ln \frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}})\right)$$

is also tight. Hence, by Prohorov's theorem, given any sequence $\{n_j\}$, there exists a further sub-sequence $\{n_{j_m}\}$ - represented as $\{n\}$ for simplicity - such that

$$\begin{pmatrix} \varphi_n \\ \frac{dP_{n\hat{\tau},.,h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}) \end{pmatrix} \xrightarrow[P_{nT,0}]{d} \begin{pmatrix} \bar{\varphi} \\ V \end{pmatrix}; \quad V \sim \exp\left\{h^{\intercal} I^{1/2} x(\tau) - \frac{\tau}{2} h^{\intercal} I h\right\}, \tag{A.2}$$

where $\bar{\varphi} \in [0, 1]$. It is a well known property of Brownian motion that

$$M(t) := \exp\left\{h^{\intercal} I^{1/2} x(t) - \frac{t}{2} h^{\intercal} I h\right\}$$

is a martingale with respect to the filtration $\mathcal{F}_t$. Since τ is an $\mathcal{F}_t$-adapted stopping time, the optional stopping theorem then implies $E[V] \equiv E[M(\tau)] = E[M(0)] = 1$.

We now claim that

$$\varphi_n \xrightarrow[P_{nT,h}]{d} L; \text{ where } L(B) := E[\mathbb{I}\{\bar{\varphi} \in B\} V] \ \forall \ B \in \mathcal{B}(\mathbb{R}). \tag{A.3}$$

It is clear from $V \geq 0$ and $E[V] = 1$ that $L(\cdot)$ is a probability measure, and that for every measurable function $f : \mathbb{R} \to \mathbb{R}$, $\int f dL = E[f(\bar{\varphi})V]$. Furthermore, for any $f(\cdot)$

continuous and non-negative,

$$\liminf \mathbb{E}_{nT,h}[f(\varphi_n)] \geq \liminf \mathbb{E}_{nT,0}\left[f(\varphi_n)\frac{dP_{nT,h}}{dP_{nT,0}}\right]$$
$$= \liminf \mathbb{E}_{nT,0}\left[f(\varphi_n)\frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}\right] \geq E[f(\bar{\varphi})V],$$

where the equality follows from the law of iterated expectations since φ_n is a function only of $\mathbf{y}_{n\hat{\tau}}$ and $dP_{nt,h}/dP_{nt,0}$ is a martingale under $P_{nT,0}$; and the last inequality follows from applying the portmanteau lemma on (A.2). Finally, applying the portmanteau lemma again, in the converse direction, gives (A.3).

Since φ_n is bounded, (A.3) implies

$$\lim_{n\to\infty} \beta_n(h) := \lim_{n\to\infty} \mathbb{E}_{nT,h}\left[\varphi_n\right] = E\left[\bar{\varphi}e^{h^\intercal I^{1/2}x(\tau)-\frac{\tau}{2}h^\intercal Ih}\right]. \tag{A.4}$$

Define $\varphi(\tau, x(\tau)) := E[\bar{\varphi}|\tau, x(\tau)]$; this is a test statistic since $\varphi \in [0,1]$. The right hand side of (A.4) then becomes

$$E\left[\varphi(\tau, x(\tau))e^{h^\intercal I^{1/2}x(\tau)-\frac{\tau}{2}h^\intercal Ih}\right].$$

But by the Girsanov theorem, this is just the expectation, $\mathbb{E}_h[\varphi(\tau, x(\tau))]$, of $\varphi(\tau, x(\tau))$ when $x(t)$ is distributed as a Gaussian process with drift $I^{1/2}h$, i.e., when $x(t) \sim I^{1/2}ht + W(t)$. We have thus shown that $\beta_n(h)$ converges to $\mathbb{E}_h[\varphi(\tau, x(\tau))] := \beta(h)$ on sub-sequences (since (A.2) only holds true on subsequences). However, $\beta_n(h)$ is a convergent sequence by assumption, so we can remove the sub-sequence qualification.

A.2. **Proof of Proposition 1.** We start by proving the first claim. Denote $H_0 \equiv \{h : a^\intercal h = 0\}$ and $H_1 \equiv \{h : a^\intercal h = c\}$. Let $\mathbb{P}_h$ denote the induced probability measure over the sample paths generated by $x(t) \sim I^{1/2}ht + W(t)$ between $t \in [0, T]$. As before, $\mathcal{F}_t$ denotes the filtration generated by $\{U, x(s) : s \leq t\}$. Given any $h_1 \in H_1$, define $h_0 = h_1 - (a^\intercal h_1/a^\intercal I^{-1}a)I^{-1}a$. Note that $a^\intercal h_1 = c$ and $h_0 \in H_0$. Let $\ln \frac{d\mathbb{P}_{h_1}}{d\mathbb{P}_{h_0}}(\mathcal{F}_t)$ denote the likelihood ratio between the probabilities induced by the parameters h_1, h_0 over the filtration $\mathcal{F}_t$. By the Girsanov theorem,

$$\ln \frac{d\mathbb{P}_{h_1}}{d\mathbb{P}_{h_0}}(\mathcal{F}_\tau) = \left(h_1^\intercal I^{1/2}x(\tau) - \frac{\tau}{2}h_1^\intercal Ih_1\right) - \left(h_0^\intercal I^{1/2}x(\tau) - \frac{\tau}{2}h_0^\intercal Ih_0\right)$$
$$= \frac{1}{\sigma}c\tilde{x}(\tau) - \frac{c^2}{2\sigma^2}\tau,$$

where $\tilde{x}(t) := \sigma^{-1}a^{\intercal}I^{-1/2}x(t)$. Hence, an application of the Neyman-Pearson lemma shows that the UMP test of $H_0' : h = h_0$ vs $H_1' : h = h_1$ is given by

$$\varphi_c^* = \mathbb{I}\left\{c\tilde{x}(\tau) - \frac{c^2}{2\sigma}\tau \geq \gamma\right\},$$

where γ is chosen by the size requirement. Now, for any $h_0 \in H_0$,

$$\tilde{x}(t) \equiv \sigma^{-1}a^{\intercal}I^{-1/2}x(t) \sim W(t).$$

Hence, the distribution of the sample paths of $\tilde{x}(\cdot)$ is independent of h_0 under the null. Combined with the assumption that τ is $\tilde{\mathcal{F}}_t$-adapted, this implies φ_c^* does not depend on h_0 and, by extension, h_1, except through c. Since $h_1 \in H_1$ was arbitrary, we are led to conclude φ_c^* is UMP more generally for testing $H_0 : a^{\intercal}h = 0$ vs $H_1 : a^{\intercal}h = c$.

The second claim is an easy consequence of the first claim and Theorem 1.

A.3. **Proof of Proposition 2.** By the Girsanov theorem,

$$\beta(h) := \mathbb{E}_h[\varphi] = \mathbb{E}_0\left[\varphi(\tau, x(\tau))e^{h^{\intercal}I^{1/2}x(\tau)-\frac{\tau}{2}h^{\intercal}Ih}\right].$$

It can be verified from the above that $\beta(h)$ is differentiable around $h = 0$. But unbiasedness requires $\mathbb{E}_h[\varphi] \geq \alpha$ for all h and $\mathbb{E}_0[\varphi] = \alpha$. This is only possible if $\beta'(0) = 0$, i.e., $\mathbb{E}_0[x(\tau)\varphi(\tau, x(\tau))] = 0$.

A.4. **Proof of Theorem 2.** Since $\hat{\tau}$ is bounded, it follows by similar arguments as in the proof of Theorem 1 that $\left(\varphi_n, \hat{\tau}, \ln \frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}})\right)$ is tight. Consequently, by Prohorov's theorem, given any sequence $\{n_j\}$, there exists a further sub-sequence $\{n_{j_m}\}$ - represented as $\{n\}$ for simplicity - such that

$$\begin{pmatrix} \varphi_n \\ \hat{\tau} \\ \frac{dP_{n\hat{\tau},h}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}) \end{pmatrix} \xrightarrow[P_{nT,0}]{d} \begin{pmatrix} \bar{\varphi} \\ \tau \\ V \end{pmatrix}; \quad V \sim \exp\left\{h^{\intercal}I^{1/2}x(\tau) - \frac{\tau}{2}h^{\intercal}Ih\right\}. \tag{A.5}$$

It then follows as in the proof of Theorem 1 that

$$\begin{pmatrix} \varphi_n \\ \hat{\tau} \end{pmatrix} \xrightarrow[P_{nT,h}]{d} L; \text{ where } L(B) := E[\mathbb{I}\{(\bar{\varphi}, \tau) \in B\}V] \ \forall\ B \in \mathcal{B}(\mathbb{R}^2). \tag{A.6}$$

The above in turn implies

$$\lim_{n\to\infty} \mathbb{E}_{nT,h}\left[\varphi_n \mathbb{I}\{\hat{\tau}=t\}\right] = E\left[\bar{\varphi}\mathbb{I}\{\tau=t\}e^{h^\intercal I^{1/2}x(\tau)-\frac{\tau}{2}h^\intercal I h}\right], \text{ and} \tag{A.7}$$

$$\lim_{n\to\infty} \mathbb{E}_{nT,h}\left[\mathbb{I}\{\hat{\tau}=t\}\right] = E\left[\mathbb{I}\{\tau=t\}e^{h^\intercal I^{1/2}x(\tau)-\frac{\tau}{2}h^\intercal I h}\right]. \tag{A.8}$$

for every $t \in \{1, 2, \ldots, T\}$.

Denote $\varphi(\tau, x(\tau)) = E[\bar{\varphi}|\tau, x(\tau)]$; this is a level-$\boldsymbol{\alpha}$ test, as can be verified by setting $h = 0$ in (A.7). The right hand side of (A.7) then becomes

$$E\left[\varphi(\tau, x(\tau))\mathbb{I}\{\tau=t\}e^{h^\intercal I^{1/2}x(\tau)-\frac{\tau}{2}h^\intercal I h}\right].$$

An application of the Girsanov theorem then shows that the right hand sides of (A.7) and (A.8) are just the expectations $\mathbb{E}_h[\varphi(\tau, x(\tau))\mathbb{I}\{\tau = t\}]$ and $\mathbb{E}_h[\mathbb{I}\{\tau = t\}]$ when $x(t) \sim I^{1/2}ht + W(t)$. What is more, the measures $\mathbb{P}_0(\cdot), \mathbb{P}_h(\cdot)$ are absolutely continuous, so $\mathbb{P}_0(\tau = t) = 0$ if and only if $\mathbb{P}_h(\tau = t) = 0$ for any $h \in \mathbb{R}^d$. We are thus led to conclude that

$$\lim_{n\to\infty} \beta_n(h|t) := \lim_{n\to\infty} \frac{\mathbb{E}_{nT,h}\left[\varphi_n \mathbb{I}\{\hat{\tau}=t\}\right]}{\mathbb{E}_{nT,h}\left[\mathbb{I}\{\hat{\tau}=t\}\right]} = \frac{\mathbb{E}_h\left[\varphi_n \mathbb{I}\{\hat{\tau}=t\}\right]}{\mathbb{E}_h\left[\mathbb{I}\{\hat{\tau}=t\}\right]} := \beta(h|t) \tag{A.9}$$

for every $h \in \mathbb{R}^d$, and $t \in \{1, 2, \ldots, T\}$ satisfying $\mathbb{P}_0(\tau = t) \neq 0$.

While we have only demonstrated (A.9) for sub-sequences, the assumption that $\beta_n(h|t)$ is a convergent sequence implies this result holds more generally for the entire sequence.

A.5. **Proof of Proposition 3.** Fix some arbitrary $g_1 \in T(P_0)$. To simplify matters, we set $\delta = 1$. The case of general δ can be handled by simply replacing g_1 with g_1/δ. By standard results for Hilbert spaces, we can write $g_1 = \sigma^{-1}\langle\psi, g\rangle(\psi/\sigma) + \tilde{g}_1$, where $\tilde{g}_1 \perp (\psi/\sigma)$. Define $\boldsymbol{g} := (\psi/\sigma, \tilde{g}_1/\|\tilde{g}_1\|)^\intercal$, and consider sub-models of the form $P_{1/\sqrt{n},\boldsymbol{h}^\intercal\boldsymbol{g}}$ for $\boldsymbol{h} \in \mathbb{R}^2$. By (3.2),

$$\sum_{i=1}^{\lfloor nt\rfloor} \ln \frac{dP_{1/\sqrt{n},\boldsymbol{h}^\intercal\boldsymbol{g}}}{dP_0}(Y_i) = \frac{\boldsymbol{h}^\intercal}{\sqrt{n}}\sum_{i=1}^{\lfloor nt\rfloor}\boldsymbol{g}(Y_i) - \frac{t}{2}\boldsymbol{h}^\intercal\boldsymbol{h} + o_{P_{nT,0}}(1), \quad \text{uniformly over } t. \tag{A.10}$$

Comparing with (2.3), we observe that $\left\{P_{1/\sqrt{n},\boldsymbol{h}^\intercal\boldsymbol{g}} : \boldsymbol{h} \in \mathbb{R}^2\right\}$ is equivalent to a parametric model with score $\boldsymbol{g}(\cdot)$ and local parameter $\boldsymbol{h}$ (note that $\mathbb{E}_{P_0}[\boldsymbol{g}\boldsymbol{g}^\intercal] = I$). Let $G_n(t) := n^{-1/2}\sum_{i=1}^n \boldsymbol{g}(Y_i)$ denote the score process. By the functional central limit theorem, $G_n(t) \xrightarrow[P_{nT,0}]{d} G(t) \equiv (x(t), \tilde{G}(t))$, where $x(\cdot), \tilde{G}(\cdot)$ are independent one-dimensional Brownian motions. Take $\mathcal{G}_t := \sigma\{G(s) : s \le t\}$, $\mathcal{F}_t := \sigma\{x(s) : s \le t\}$ to be the filtrations generated by $G(\cdot)$ and $x(\cdot)$ respectively until time t. Since the first component of $G_n(\cdot)$ is

$x_n(\cdot)$ and $\hat{\tau} = \tau(x_n(\cdot))$ by Assumption 3(ii), the extended continuous mapping theorem implies

$$(\hat{\tau}, G_n(\hat{\tau})) \xrightarrow[P_{nT,0}]{d} (\tau, G(\tau)), \tag{A.11}$$

where τ is a $\mathcal{F}_t$-adapted stopping time, and therefore, $\mathcal{G}_t$-adapted by extension.

Let φ_n denote any asymptotically level-α test. Consider the limit experiment where one observes a $\mathcal{G}_t$-adapted stopping time τ along with a diffusion process $G(t) := \boldsymbol{h}t + W(t)$, where $W(\cdot)$ is 2-dimensional Brownian motion. Using (A.10) and (A.11), we can argue as in the proof of Theorem 1 that the power function, $\beta_n(\boldsymbol{h}^\intercal \boldsymbol{g}) := \int \varphi_n dP_{nT,\boldsymbol{h}^\intercal \boldsymbol{g}}$, of φ_n in the parametric model $\left\{P_{1/\sqrt{n},\boldsymbol{h}^\intercal \boldsymbol{g}} : \boldsymbol{h} \in \mathbb{R}^2\right\}$ can be matched along sub-sequences (since $\beta_n(\cdot) \in [0,1]$ is tight) by the power function $\beta(\boldsymbol{h})$ of some test $\varphi(\tau, G(\tau))$ in the limit experiment (the choice of $\varphi(\cdot)$ is allowed to depend on the sub-sequence). Note that by our definitions, $\langle \psi, \boldsymbol{h}^\intercal \boldsymbol{g}\rangle$ is simply the first component of $\boldsymbol{h}$ divided by σ. This in turn implies, as a consequence of the definition of asymptotically level-α tests, that $\varphi(\cdot)$ is level-α for testing $H_0 : (1,0)^\intercal \boldsymbol{h} = 0$ in the limit experiment.

Now, by a similar argument as in the proof of Proposition 1, along with the fact $(1,0)^\intercal G(t) = x(t)$, the optimal level-$\alpha$ test of $H_0 : (1,0)^\intercal \boldsymbol{h} = 0$ vs $H_1 : (1,0)^\intercal \boldsymbol{h} = \mu_1/\sigma$ in the limit experiment is given by

$$\varphi^*_{\mu_1}(\tau, x(\tau)) := \mathbb{I}\left\{\mu_1 x(\tau) - \frac{\mu_1^2}{2\sigma}\tau \geq \gamma\right\}.$$

Since $G(t) := \boldsymbol{h}t + W(t)$, for all $\boldsymbol{h} \in H_1$ (i.e., all $\boldsymbol{h}$ in the alternative set),

$$x(t) = (1,0)^\intercal G(t) \sim \sigma^{-1}\mu_1 t + \tilde{W}(t),$$

where $\tilde{W}(\cdot)$ is 1-dimensional Brownian motion. As τ is $\mathcal{F}_t$-adapted, the joint distribution of $(\tau, x(\tau))$ therefore depends only on μ_1 for $\boldsymbol{h} \in H_1$. Consequently, the power, $\mathbb{E}_{\boldsymbol{h}}[\varphi^*_{\mu_1}(\tau, x(\tau))]$, of $\varphi^*_{\mu_1}(\cdot)$ against such alternatives depends only on μ_1, and is denoted by $\beta^*(\mu_1)$. Since $\varphi^*_{\mu_1}(\cdot)$ is the optimal test and $\mu_1 = \langle \psi, \boldsymbol{h}^\intercal \boldsymbol{g}\rangle$, we conclude $\beta(\boldsymbol{h}) \leq \beta^*(\langle \psi, \boldsymbol{h}^\intercal \boldsymbol{g}\rangle)$.

Since the above holds for any $\beta(\boldsymbol{h})$ that is the limit of a sub-sequence of $\beta_n(\boldsymbol{h}^\intercal g)$, we conclude that $\limsup_n \beta_n(\boldsymbol{h}^\intercal \boldsymbol{g}) \leq \beta^*(\langle \psi, \boldsymbol{h}^\intercal \boldsymbol{g}\rangle)$ for any $\boldsymbol{h} \in \mathbb{R}^2$. Setting $\boldsymbol{h} = (\langle \psi, g_1\rangle/\sigma, \|\tilde{g}_1\|)^\intercal$ then gives $\limsup_n \beta_n(g_1) \leq \beta^*(\langle \psi, g_1\rangle)$. Since $g_1 \in T(P_0)$ was arbitrary, the claim follows.

A.6. **Proof of Proposition 5.** Recall that we can represent any $g_1 \in T(P_0)$ as $g_1 = \sigma^{-1}\langle \psi, g\rangle(\psi/\sigma) + \tilde{g}_1$, where $\tilde{g}_1 \perp (\psi/\sigma)$. Let φ_n denote any asymptotically level-α test.

By the form of the weight function $m(\cdot)$ and Fubini's theorem, we can write the weighted average power of φ_n as

$$\begin{aligned}\int \beta_n(g_1)dm(g_1) &= \int \left(\int \beta_n\left(\frac{\mu}{\sigma^2}\psi + \tilde{g}_1\right) dw(\mu)\right) d\rho(\tilde{g}_1)\\ &:= \int \left(\int \beta_n(\mu; \tilde{g}_1)dw(\mu)\right) d\rho(\tilde{g}_1).\end{aligned}$$

Since $\int \beta_n(\mu; \tilde{g}_1)dw(\mu) \leq 1$, an application of Fatou's lemma gives

$$\limsup_{n\to\infty} \int \beta_n(g_1)dm(g_1) \leq \int \limsup_{n\to\infty} \left(\int \beta_n(\mu; \tilde{g}_1)dw(\mu)\right) d\rho(\tilde{g}_1).$$

The claim thus follows if we show that

$$\limsup_{n\to\infty} \int \beta_n(\mu; \tilde{g}_1)dw(\mu) \leq \int \beta^*(\mu)dw(\mu) \text{ for each } \tilde{g}_1. \tag{A.12}$$

To this end, define $\boldsymbol{g} = (\psi/\sigma, \tilde{g}_1/\left\|\tilde{g}_1\right\|)^\intercal$, and consider sub-models of the form $P_{1/\sqrt{n},\boldsymbol{h}_\mu^\intercal \boldsymbol{g}}$ where $\boldsymbol{h}_\mu := (\mu/\sigma, \|\tilde{g}_1\|)$. By (3.2),

$$\sum_{i=1}^{\lfloor nt\rfloor} \ln \frac{dP_{1/\sqrt{n},\boldsymbol{h}_\mu^\intercal \boldsymbol{g}}}{dP_0}(Y_i) = \frac{\boldsymbol{h}_\mu^\intercal}{\sqrt{n}} \sum_{i=1}^{\lfloor nt\rfloor} \boldsymbol{g}(Y_i) - \frac{t}{2}\boldsymbol{h}_\mu^\intercal \boldsymbol{h}_\mu + o_{P_{nT,0}}(1), \quad \text{uniformly over } t. \tag{A.13}$$

Let $\beta_n\left(\boldsymbol{h}_\mu^\intercal \boldsymbol{g}\right) := \int \varphi_n dP_{nT,\boldsymbol{h}_\mu^\intercal \boldsymbol{g}}$ denote the power function of φ_n under the parametric model $\left\{P_{1/\sqrt{n},\boldsymbol{h}_\mu^\intercal \boldsymbol{g}} : \mu \in \mathbb{R}\right\}$. Note that by construction $\beta_n(\mu; \tilde{g}_1) = \beta_n\left(\boldsymbol{h}_\mu^\intercal \boldsymbol{g}\right)$.

We now employ similar arguments as in the proof of Proposition 3. In that proof we defined a diffusion process $G(t) := \boldsymbol{h}_\mu t + W(t)$, where $W(\cdot)$ is 2-dimensional Brownian motion. Also, we set

$$x(t) := (1, 0)^\intercal G(t) \sim \sigma^{-1}\mu t + \tilde{W}(t),$$

where $\tilde{W}(\cdot)$ denotes 1-dimensional Brownian motion, and took $\mathcal{F}_t$ to be the filtration generated by $x(t)$ until time t. Now consider a limit experiment where one observes a $\mathcal{F}_t$-adapted stopping time τ along with the diffusion process $G(t)$. By similar arguments as in the proof of Proposition 3, $\beta_n\left(\boldsymbol{h}_\mu^\intercal \boldsymbol{g}\right)$ can be matched along sub-sequences by the power function $\beta(\boldsymbol{h}_\mu)$ of some test $\varphi(\tau, G(\tau))$ in the limit experiment. As $\beta_n(\cdot) \in [0, 1]$, the dominated convergence theorem then implies

$$\lim_{n\to\infty} \int \beta_n\left(\boldsymbol{h}_\mu^\intercal \boldsymbol{g}\right) dw(\mu) = \int \beta(\boldsymbol{h}_\mu)dw(\mu)$$

on subsequences. Note that the limit $\beta(\boldsymbol{h}_\mu)$ is allowed to be potentially different across sub-sequences.

Since $\left\langle \psi, \boldsymbol{h}_{\mu}^{\intercal}\boldsymbol{g}\right\rangle = \mu/\sigma$, the definition of asymptotically level-α tests, together with the fact $\beta_n\left(\boldsymbol{h}_{\mu}^{\intercal}\boldsymbol{g}\right)$ converges to $\beta(\boldsymbol{h}_{\mu})$ on sub-sequences, implies $\varphi(\cdot)$ is level-α for testing $H_0 : (1,0)^{\intercal}\boldsymbol{h}_{\mu} = 0$ in the limit experiment. Now, given a weight function $w(\cdot)$, an application of Girsanov's theorem shows that the optimal weighted average power test of $H_0 : (1,0)^{\intercal}\boldsymbol{h}_{\mu} = 0$ vs $H_1 : (1,0)^{\intercal}\boldsymbol{h}_{\mu} \neq 0$ in the limit experiment is given by

$$\varphi_w^*(\tau, x(\tau)) = \mathbb{I}\left\{\int \exp\left(\frac{\mu}{\sigma}x(\tau) - \frac{\mu^2}{2\sigma^2}\tau\right) dw(\mu) > \gamma\right\}.$$

We therefore conclude $\int \beta(\boldsymbol{h}_{\mu})dw(\mu) \leq \int \beta^*(\mu)dw(\mu)$, where $\beta^*(\mu)$ is the power function of $\varphi_w^*(\cdot)$. Since this inequality holds for any $\beta(\boldsymbol{h}_{\mu})$ that is the limit of a sub-sequence of $\beta_n(\boldsymbol{h}_{\mu}^{\intercal}g)$, we have

$$\limsup_{n\to\infty} \int \beta_n\left(\boldsymbol{h}_{\mu}^{\intercal}\boldsymbol{g}\right) dw(\mu) \leq \int \beta^*(\mu)dw(\mu).$$

The above proves (A.12) and the claim thus follows.

A.7. **Proof of Proposition 6.** Fix some arbitrary $\boldsymbol{g} = (g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$. To simplify matters, we set $\delta_1 = \delta_0 = 1$. The case of general δ can be handled by simply replacing g_a with g_a/δ_a. In what follows, let $\pi_1 = \pi$ and $\pi_0 = 1 - \pi$. The vectors $\mathbf{y}_{nt}^{(1)} = (Y_1^{(1)}, \dots, Y_{n\pi_1 t}^{(1)})$ and $\mathbf{y}_{nt}^{(0)} = (Y_1^{(0)}, \dots, Y_{n\pi_0 t}^{(0)})$ denote the collection of outcomes from treatments 1 and 0 until time t, and we set $\mathbf{y}_{nt} = (\mathbf{y}_{nt}^{(1)}, \mathbf{y}_{nt}^{(0)})$. Define $P_{nt,\boldsymbol{g}}$ as the joint probability measure over $\mathbf{y}_{nt}$ when each $Y_i^{(a)}$ is an iid draw from $P_{1/\sqrt{n},g_a}^{(a)}$.

As in the proof of Proposition 3, we can write $g_a = \sigma_a^{-1}\left\langle \psi_a, g_a\right\rangle_a (\psi_a/\sigma_a) + \tilde{g}_a$, where $\tilde{g}_a \perp (\psi_a/\sigma_a)$. Define $\boldsymbol{g}_a := (\psi_a/\sigma_a, \tilde{g}_a/\left\|\tilde{g}_a\right\|_a)^{\intercal}$, and consider sub-models of the form $P_{1/\sqrt{n},\boldsymbol{h}_1^{\intercal}\boldsymbol{g}_1} \times P_{1/\sqrt{n},\boldsymbol{h}_0^{\intercal}\boldsymbol{g}_0}$ for $\boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2$. By the SLAN property, (3.2), and the fact that $\mathbf{y}_{nt}^{(1)}, \mathbf{y}_{nt}^{(0)}$ are independent,

$$\begin{aligned}\ln \frac{dP_{nt,(\boldsymbol{h}_1^{\intercal}\boldsymbol{g}_1,\boldsymbol{h}_0^{\intercal}\boldsymbol{g}_0)}}{dP_{nt,\boldsymbol{0}}}(\mathbf{y}_{nt}) &= \frac{\boldsymbol{h}_1^{\intercal}}{\sqrt{n}}\sum_{i=1}^{\lfloor n\pi_1 t\rfloor} \boldsymbol{g}_1(Y_i^{(1)}) - \frac{\pi_1 t}{2}\boldsymbol{h}_1^{\intercal}\boldsymbol{h}_1 + \dots \\ &\dots + \frac{\boldsymbol{h}_0^{\intercal}}{\sqrt{n}}\sum_{i=1}^{\lfloor n\pi_0 t\rfloor} \boldsymbol{g}_0(Y_i^{(0)}) - \frac{\pi_0 t}{2}\boldsymbol{h}_0^{\intercal}\boldsymbol{h}_0 + o_{P_{nT,0}}(1), \quad \text{uniformly over } t. \end{aligned} \tag{A.14}$$

Let $G_{a,n}(t) := n^{-1/2}\sum_{i=1}^{\lfloor n\pi_a t\rfloor} \boldsymbol{g}_a(Y_i^{(a)})$ for $a \in \{0,1\}$. By a standard functional central limit theorem,

$$G_{a,n}(t) \xrightarrow[P_{nT,\boldsymbol{0}}]{d} G_a(t) \equiv (z_a(t), \tilde{G}_a(t)),$$

where $z_a(\cdot)/\sqrt{\pi_a}, \tilde{G}_a(\cdot)/\sqrt{\pi_a}$ are independent 1-dimensional Brownian motions. Furthermore, since $\mathbf{y}_{nt}^{(1)}, \mathbf{y}_{nt}^{(0)}$ are independent of each other, $G_1(\cdot), G_0(\cdot)$ are independent Gaussian processes. Define $\sigma^2 := \left(\frac{\sigma_1^2}{\pi_1} + \frac{\sigma_0^2}{\pi_0}\right)$,

$$x(t) := \frac{1}{\sigma}\left(\frac{\sigma_1}{\pi_1} z_1(t) - \frac{\sigma_0}{\pi_0} z_0(t)\right)$$

and take $\mathcal{G}_t := \sigma\{(G_1(s), G_0(s)) : s \leq t\}$, $\mathcal{F}_t := \sigma\{x(s) : s \leq t\}$ to be the filtrations generated by $\boldsymbol{G}(\cdot) := (G_1(\cdot), G_0(\cdot))$ and $x(\cdot)$ respectively until time t. Using Assumption 4(ii), the extended continuous mapping theorem implies

$$(\hat{\tau}, G_{1,n}(\hat{\tau}), G_{0,n}(\hat{\tau})) \xrightarrow[P_{nT,0}]{d} (\tau, G_1(\tau), G_0(\tau)), \tag{A.15}$$

where τ is a $\mathcal{F}_t$-adapted stopping time, and thereby $\mathcal{G}_t$-adapted, by extension.

Let φ_n denote any asymptotically unbiased test. Consider the limit experiment where one observes a $\mathcal{G}_t$-adapted stopping time τ along with diffusion processes $G_a(t) := \pi_a \boldsymbol{h}_a t + \sqrt{\pi_a} W_a(t),\ a \in \{0,1\}$, where $W_1(\cdot), W_0(\cdot)$ are independent 2-dimensional Brownian motions. By Lemma 2 in Appendix B, the power function

$$\beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0) := \int \varphi_n dP_{nT,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}$$

of φ_n under the parametric model $\left\{P_{1/\sqrt{n}, \boldsymbol{h}_1^\intercal \boldsymbol{g}_1} \times P_{1/\sqrt{n}, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0} : \boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2\right\}$ can be matched, along sub-sequences, by the power function, $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0)$, of some test $\varphi(\tau, \boldsymbol{G}(\tau))$ in the limit experiment (the choice of $\varphi(\cdot)$ is allowed to depend on the sub-sequence). Note that by our definitions, the first component of $\boldsymbol{h}_a$ is $\langle \psi_a, \boldsymbol{h}_a^\intercal \boldsymbol{g}_a \rangle_a / \sigma_a$. This in turn implies, as a consequence of the definition of asymptotically level-α tests, that $\varphi(\cdot)$ is level-α for testing $H_0 : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = 0$ in the limit experiment.

Now, by Lemma 3 in Appendix B, the optimal level-α test of $H_0 : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = 0$ vs $H_1 : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = \mu$ in the limit experiment is

$$\varphi_\mu^*(\tau, x(\tau)) := \mathbb{I}\left\{\mu x(\tau) - \frac{\mu^2}{2\sigma}\tau \geq \gamma\right\}.$$

For all $\boldsymbol{h} \in H_1 \equiv \{\boldsymbol{h} : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = \mu\}$ in the alternative set,

$$\begin{aligned} x(t) &\sim \sigma^{-1}\mu t + \frac{1}{\sigma}\left(\sqrt{\frac{\sigma_1^2}{\pi_1}}(1,0)^\intercal W_1(t) - \sqrt{\frac{\sigma_0^2}{\pi_0}}(1,0)^\intercal W_0(t)\right) \\ &\sim \sigma^{-1}\mu t + \tilde{W}(t), \end{aligned}$$

where $\tilde{W}(\cdot)$ is standard 1-dimensional Brownian motion. As τ is $\mathcal{F}_t$-adapted, it follows that the joint distribution of $(\tau, x(\tau))$ depends only on μ for $\boldsymbol{h} \in H_1$. Consequently, the power, $\mathbb{E}_{\boldsymbol{h}}[\varphi_\mu^*(\tau, x(\tau))]$, of φ_μ^* against the values in the alternative hypothesis H_1 depends only on μ, and is denoted by $\beta^*(\mu)$. Since $\varphi^*_\mu(\cdot)$ is the optimal test, $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0) \le \beta^*(\mu)$ for all $\boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2$.

Since the above holds for any $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0)$ that is the limit of a sub-sequence of $\beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)$, we conclude that $\limsup_n \beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0) \le \beta^*(\mu)$ for any $\mu \in \mathbb{R}$ and $\boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2$ satisfying $\langle \psi_1, \boldsymbol{h}_1^\intercal \boldsymbol{g}_1 \rangle_1 - \langle \psi_0, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0 \rangle_0 = \mu$. Setting $\boldsymbol{h}_a = (\sigma_a^{-1} \langle \psi_a, g_a \rangle_a, \|\tilde{g}_a\|_a)^\intercal$ for $a \in \{0, 1\}$ then gives $\limsup_n \int \varphi_n dP_{nT,(g_1,g_0)} \le \beta^*(\mu)$. Since $(g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$ was arbitrary, the claim follows.

A.8. **Proof of Theorem 3.** As noted previously, the first claim is shown in Hirano and Porter (2023). Consequently, we only focus on proving the second claim. Let $\mathbf{y}^{(a)}_{j,nq}$ denote the first nq observations from treatment a in batch j. Define

$$\ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\mathbf{y}^{(a)}_{j,nq}) = \sum_{i=1}^{\lfloor nq \rfloor} \ln \frac{dp_{\theta_0^{(a)}+h_a/\sqrt{n}}}{dp_{\theta_0}}(Y^{(a)}_{i,j}).$$

By the SLAN property, which is a consequence of Assumption 5,

$$\ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\mathbf{y}^{(a)}_{j,n\hat{\pi}_j^{(a)}}) = h_a^\intercal I_a^{1/2} z^{(a)}_{j,n}(\hat{\pi}_j^{(a)}) - \frac{\hat{\pi}_j^{(a)}}{2} h_a^\intercal I_a h_a + o_{P_{n,0}}(1). \tag{A.16}$$

The above is true for all j, a.

Denote the observed set of outcomes by $\bar{\mathbf{y}} = \left(\mathbf{y}^{(1)}_{1,n\hat{\pi}_1^{(1)}}, \mathbf{y}^{(0)}_{1,n\hat{\pi}_1^{(0)}}, \dots, \mathbf{y}^{(1)}_{J,n\hat{\pi}_J^{(1)}}, \mathbf{y}^{(0)}_{J,n\hat{\pi}_J^{(0)}}\right)$. The likelihood ratio of the observations satisfies

$$\begin{aligned} \ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\bar{\mathbf{y}}) &= \sum_j \sum_{a \in \{0,1\}} \ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\mathbf{y}^{(a)}_{j,nq}) \\ &= \sum_j \sum_{a \in \{0,1\}} \left\{ h_a^\intercal I_a^{1/2} z^{(a)}_{j,n}(\hat{\pi}_j^{(a)}) - \frac{\hat{\pi}_j^{(a)}}{2} h_a^\intercal I_a h_a \right\}, \end{aligned} \tag{A.17}$$

where the second equality follows from (A.16). Combining the above with the first part of the theorem, we find

$$\ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\bar{\mathbf{y}}) \xrightarrow[P_{n,0}]{d} \sum_j \sum_{a \in \{0,1\}} \left\{ h_a^\intercal I_a^{1/2} z^{(a)}_j(\pi_j^{(a)}) - \frac{\pi_j^{(a)}}{2} h_a^\intercal I_a h_a \right\}, \tag{A.18}$$

where $z^{(a)}_j(t)$ is distributed as d-dimensional Brownian motion.

Note that φ_n is required to be measurable with respect to $\bar{\mathbf{y}}$. Furthermore, φ_n is tight since $\varphi_n \in [0, 1]$. Together with (A.18), this implies the joint $\left(\varphi_n, \ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\bar{\mathbf{y}})\right)$ is also tight. Hence, by Prohorov's theorem, given any sequence $\{n_j\}$, there exists a further sub-sequence $\{n_{j_m}\}$ - represented as $\{n\}$ without loss of generality - such that

$$\begin{pmatrix} \varphi_n \\ \ln \frac{dP_{n,\boldsymbol{h}}}{dP_{n,0}}(\bar{\mathbf{y}}) \end{pmatrix} \xrightarrow[P_{n,0}]{d} \begin{pmatrix} \bar{\varphi} \\ V \end{pmatrix}; \quad V \sim \prod_{j=1,\dots,J} \prod_{a\in\{0,1\}} \exp\left\{ h_a^{\intercal} I_a^{1/2} z_j^{(a)}(\pi_j^{(a)}) - \frac{\pi_j^{(a)}}{2} h_a^{\intercal} I_a h_a \right\}, \tag{A.19}$$

where $\bar{\varphi} \in [0,1]$. Define

$$V_j^{(a)} := \exp\left\{ h_a^{\intercal} I_a^{1/2} z_j^{(a)}(\pi_j^{(a)}) - \frac{\pi_j^{(a)}}{2} h_a^{\intercal} I_a h_a \right\},$$

so that $V = \prod_{j=1,\dots,J}\prod_{a\in\{0,1\}} V_j^{(a)}$. By the definition of $z_j^{(a)}(\cdot)$ and $\pi_j^{(a)}$ in the limit experiment, we have that the process $z_j^{(a)}(\cdot)$ is independent of data from the all past batches, and consequently, is also independent of $\pi_j^{(a)}$. Hence, by the martingale property of $M_j^{(a)}(t) := \exp\left\{ h_a^{\intercal} I_a^{1/2} z_j^{(a)}(t) - \frac{t}{2} h_a^{\intercal} I_a h_a \right\}$,

$$E[V_j^{(a)} | z_1^{(1)}, z_1^{(0)}, \pi_1^{(1)}, \pi_1^{(0)} \dots, z_{j-1}^{(1)}, z_{j-1}^{(0)}, \pi_{j-1}^{(1)}, \pi_{j-1}^{(0)}] = 1$$

for all j and $a \in \{0,1\}$. This implies, by an iterative argument, that $E[V] = 1$. Consequently, we can employ similar arguments as in the proof of Theorem 1 to show that

$$\begin{aligned} \lim_{n\to\infty} \beta_n(\boldsymbol{h}) &:= \lim_{n\to\infty} \mathbb{E}_{n,\boldsymbol{h}}[\varphi_n] \\ &= E\left[\bar{\varphi} \prod_{j=1,\dots,J} \prod_{a\in\{0,1\}} e^{h_a^{\intercal} I_a^{1/2} z_j^{(a)}(\pi_j^{(a)}) - \frac{\pi_j^{(a)}}{2} h_a^{\intercal} I_a h_a} \right] \\ &= E\left[\bar{\varphi} \prod_{a\in\{0,1\}} e^{h_a^{\intercal} I_a^{1/2} x_a - \frac{q_a}{2} h_a^{\intercal} I_a h_a} \right], \end{aligned} \tag{A.20}$$

where the last equality follows from the definition of x_a, q_a. Define

$$\varphi(q_1, q_0, x_1, x_0) := E[\bar{\varphi} | q_1, q_0, x_1, x_0].$$

Then, the right hand side of (A.20) becomes

$$E\left[\varphi(q_1, q_0, x_1, x_0) \prod_{a\in\{0,1\}} e^{h_a^{\intercal} I_a^{1/2} x_a - \frac{q_a}{2} h_a^{\intercal} I_a h_a} \right].$$

But by a repeated application of the Girsanov theorem, this is just the expectation, $\mathbb{E}_{\boldsymbol{h}}[\varphi]$, of φ when each $z_j^{(a)}(t)$ is distributed as a Gaussian process with drift $I_a^{1/2} h_a$, i.e., when $z_j^{(a)}(t) \sim I_a^{1/2} h_a t + W_j^{(a)}(t)$, and $\{W_j^{(a)}(\cdot)\}_{j,a}$ are independent Brownian motions.

While we have only demonstrated (A.20) for sub-sequences (since (A.19) applies only on sub-sequences), the assumption that $\beta_n(h|t)$ is a convergent sequence implies this result holds more generally for the entire sequence.

A.9. **Proof of Proposition 8.** Denote the observed set of outcomes by

$$\bar{\mathbf{y}} = \left(\mathbf{y}^{(1)}_{1,n\hat{\pi}^{(1)}_1}, \mathbf{y}^{(0)}_{1,n\hat{\pi}^{(0)}_1}, \dots, \mathbf{y}^{(1)}_{J,n\hat{\pi}^{(1)}_J}, \mathbf{y}^{(0)}_{J,n\hat{\pi}^{(0)}_J}\right).$$

Fix some arbitrary $\boldsymbol{g} = (g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$. As in the proof of Proposition 6, we can write $g_a = \sigma_a^{-1} \langle \psi_a, g_a \rangle_a (\psi_a/\sigma_a) + \tilde{g}_a$, where $\tilde{g}_a \perp (\psi_a/\sigma_a)$. Define $\boldsymbol{g}_a := (\psi_a/\sigma_a, \tilde{g}_a/\left\|\tilde{g}_a\right\|_a)^\intercal$, and consider sub-models of the form $P_{1/\sqrt{n},\boldsymbol{h}_1^\intercal \boldsymbol{g}_1} \times P_{1/\sqrt{n},\boldsymbol{h}_0^\intercal \boldsymbol{g}_0}$ for $\boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2$. Following similar simplifications as in the proofs of Propositions 3 and 6, we set $\delta_1 = \delta_0 = 1$ without loss of generality.

Let $P_{n,\boldsymbol{h}}$ and $P_{n,0}$ be defined as in Section 4.1, and set

$$\boldsymbol{Z}^{(a)}_{j,n}(t) := \frac{1}{\sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \boldsymbol{g}_a(Y^{(a)}_{i,j}), \text{ and } z^{(a)}_{j,n}(t) := \frac{1}{\sigma_a \sqrt{n}} \sum_{i=1}^{\lfloor nt \rfloor} \psi_a(Y^{(a)}_{i,j}).$$

By similar arguments as that leading to (A.17), the likelihood ratio,

$$\ln \frac{dP_{n,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}}{dP_{n,0}}(\bar{\mathbf{y}}),$$

of all observations, $\bar{\boldsymbol{y}}$, under the sub-model $P_{1/\sqrt{n},\boldsymbol{h}_1^\intercal \boldsymbol{g}_1} \times P_{1/\sqrt{n},\boldsymbol{h}_0^\intercal \boldsymbol{g}_0}$ satisfies

$$\ln \frac{dP_{n,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}}{dP_{n,0}}(\bar{\mathbf{y}}) = \sum_a \sum_j \left\{ \frac{\boldsymbol{h}_a^\intercal}{\sqrt{n}} \boldsymbol{Z}^{(a)}_{j,n}(\hat{\pi}^{(a)}_j) - \frac{\hat{\pi}^{(a)}_j}{2} \boldsymbol{h}_a^\intercal \boldsymbol{h}_a \right\} + o_{P_{nT,0}}(1). \tag{A.21}$$

Now, by iterative use of the functional central limit theorem and the extended continuous mapping theorem (using Assumption 7),

$$\begin{pmatrix} \hat{\pi}^{(a)}_j \\ \boldsymbol{Z}^{(a)}_{j,n}(\hat{\pi}^{(a)}_j) \end{pmatrix} \xrightarrow[P_{nT,\mathbf{0}}]{d} \begin{pmatrix} \pi^{(a)}_j \\ \boldsymbol{Z}^{(a)}_j(\pi^{(a)}_j) \end{pmatrix}, \quad \boldsymbol{Z}^{(a)}_j(\cdot) \sim W_{a,j}(\cdot), \tag{A.22}$$

where $\{W_{a,j}\}_{a,j}$ are independent 2-dimensional Brownian motions, and $\pi^{(a)}_j$ is measurable with respect to $\sigma\left\{z^{(a)}_l(\cdot); l \le j-1\right\}$ since $\hat{\pi}^{(a)}_j$ is measurable with respect to $\sigma\left\{z^{(a)}_{l,n}(\cdot); l \le j-1\right\}$.

Let φ_n denote any asymptotically level-α test. Consider the limit experiment where one observes $q_a = \sum_j \pi^{(a)}_j$ and $x_a := \sum_j z^{(a)}_j(\pi^{(a)}_j)$, where

$$z^{(a)}_j(t) := \mu_a t + W^{(a)}_j(t), \tag{A.23}$$

and π_j is measurable with respect to $\sigma\left\{z_l^{(a)}(\cdot); l \leq j-1\right\}$. Using (A.21), (A.22) and employing similar arguments as in Theorem 3, we find that the power function

$$\beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0) := \int \varphi_n dP_{nT,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}$$

of φ_n in the parametric model $\left\{P_{1/\sqrt{n},\boldsymbol{h}_1^\intercal \boldsymbol{g}_1} \times P_{1/\sqrt{n},\boldsymbol{h}_0^\intercal \boldsymbol{g}_0} : \boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2\right\}$ can be matched along sub-sequences by the power function $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0)$ of some test $\varphi(q_1, q_0, \boldsymbol{G}_1, \boldsymbol{G}_0)$ that depends only on $\boldsymbol{G}_1, \boldsymbol{G}_0, q_1, q_0$ in the limit experiment (the choice of $\varphi(\cdot)$ is allowed to depend on the sub-sequence). Note that by our definitions, the first component of $\boldsymbol{h}_a$ is $\langle \psi_a, \boldsymbol{h}_a^\intercal \boldsymbol{g}_a \rangle_a / \sigma_a$. This in turn implies, as a consequence of the definition of asymptotically level-α tests, that $\varphi(\cdot)$ is level-α for testing

$$H_0 : ((\sigma_1, 0)^\intercal \boldsymbol{h}_1, (\sigma_0, 0)^\intercal \boldsymbol{h}_0) = (0, 0)$$

in the limit experiment.

Now, by Lemma 4 in Appendix B, the optimal level-α test of the null H_0 vs H_1 : $((\sigma_1, 0)^\intercal \boldsymbol{h}_1, (\sigma_0, 0)^\intercal \boldsymbol{h}_0) = (\mu_1, \mu_0)$ in the limit experiment is

$$\varphi^*_{\mu_1,\mu_0} = \mathbb{I}\left\{\sum_{a\in\{0,1\}} \left(\frac{\mu_a}{\sigma_a} x_a - \frac{q_a}{2\sigma_a^2}\mu_a^2\right) \geq \gamma_{\mu_1,\mu_0}\right\}.$$

Using (A.23) and the fact π_j depends only on the past values of $z_j^{(a)}(\cdot)$, it follows that the joint distribution of (q_1, q_0, x_1, x_0) depends only on μ_1, μ_0 for $\boldsymbol{h} \in H_1$. Consequently, the power, $\mathbb{E}_{\boldsymbol{h}}\left[\varphi^*_{\mu_1,\mu_0}\right]$, of $\varphi^*_{\mu_1,\mu_0}$ against the values in the alternative hypothesis H_1 depends only on (μ_1, μ_0), and is denoted by $\beta^*(\mu_1, \mu_0)$. Since $\varphi^*_{\mu_1,\mu_0}$ is the optimal test, $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0) \leq \beta^*(\mu_1, \mu_0)$.

Since the above holds for any $\beta(\boldsymbol{h}_1, \boldsymbol{h}_0)$ that is the limit of a sub-sequence of $\beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)$, we conclude that $\limsup_n \beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0) \leq \beta^*(\mu_1, \mu_0)$ for any $(\mu_1, \mu_0) \in \mathbb{R}$ and $\boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2$ such that $\langle \psi_a, \boldsymbol{h}_a^\intercal \boldsymbol{g}_a \rangle_a = \mu_a$. Setting $\boldsymbol{h}_a = (\sigma_a^{-1} \langle \psi_a, g_a \rangle_a, \|\tilde{g}_a\|_a)^\intercal$ for $a \in \{0, 1\}$ then gives $\limsup_n \int \varphi_n dP_{nT,(g_1,g_0)} \leq \beta^*(\mu_1, \mu_0)$. Since $(g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$ was arbitrary, the claim follows.

Appendix B. Additional results

B.1. **Variance estimators.** The score/efficient influence function process $x_n(\cdot)$ depends on the information matrix I (in the case of parametric models) or on the variance σ (in the case of non-parametric models). For parametric models, the reference parameter, θ_0, is generally known, and we could simply set $I = I(\theta_0)$. In non-parametric settings, however, this would be unknown, and we would need to replace I and σ with consistent estimators. Here, we discuss various proposals for variance estimation (note that I can be thought of as variance since $E_0[\psi\psi^\intercal] = I$).

Batched experiments. If the experiment is conducted in batches, we can simply use the data from the first batch to construct consistent estimators of the variances. This of course has the drawback of not using all the data, but it is unbiased and $\sqrt{n}$-consistent under very weak assumptions (i.e., existence of second moments).

Running-estimator of variance. For an estimator that is more generally valid and uses all the data, we recommend the running-variance estimate

$$\hat{\Sigma}_{a,t} = \frac{1}{nt}\sum_{i=1}^{\lfloor nt \rfloor} \psi_a(Y_i^{(a)})\psi_a(Y_i^{(a)})^\intercal - \left(\frac{1}{nt}\sum_{i=1}^{\lfloor nt \rfloor} \psi_a(Y_i^{(a)})\right)\left(\frac{1}{nt}\sum_{i=1}^{\lfloor nt \rfloor} \psi_a(Y_i^{(a)})\right)^\intercal, \tag{B.1}$$

for each treatment a. The final estimate of the variance would then be $\hat{\Sigma}_{a,\hat{\tau}}$ for stopping-times experiments, and $\hat{\Sigma}_{a,q_a}$ for batched experiments. Let $\Sigma_a := E_{0,a}[\psi_a\psi_a^\intercal]$ and suppose that $\psi_a\psi_a^\intercal$ is λ-sub-Gaussian for some $\lambda > 0$. Then using standard concentration inequalities, see e.g., Lattimore and Szepesvári (2020, Corollary 5.5), we can show that

$$P_{nT,0}\left(\bigcup_{t=1}^{T}\left\{\left|\hat{\Sigma}_{a,t} - \Sigma_a\right| \geq C\sqrt{\frac{\ln(1/\delta)}{nt}}\right\}\right) \leq nT\delta \quad \forall \quad \delta \in [0,1],$$

where C is independent of n, t, δ (but does depend on λ). Setting $\delta = n^{-a}$ for some $a > 0$ then implies that $\hat{\Sigma}_{a,\hat{\tau}}$ and $\hat{\Sigma}_{a,q_a}$ are $\sqrt{n}$-consistent for Σ_a (up to log factors) as long as $\hat{\tau}, q_a > 0$ almost-surely under $P_{nT,0}$.

Bayes estimators. Yet a third alternative is to place a prior on Σ_a and continuously update its value using posterior means. As a default, we suggest employing an inverse-Wishart prior and computing the posterior by treating the outcomes as Gaussian (this is of course justified in the limit). Since posterior consistency holds under mild assumptions, we expect this estimator to perform similarly to (B.1).

B.2. **Supporting information for Section 5.1.** In this section, we provide a proof of Lemma 1. The proof proceeds in two steps: First, we characterize the best unbiased test in the limit experiment described in Section 5.1. Then, we show that the finite sample counterpart of this test attains the power envelope for asymptotically unbiased tests.

Step 1: Consider the problem of testing $H_0 : \mu = 0$ vs $H_1 : \mu \neq 0$ in the limit experiment. Let $\mathbb{P}_\mu(\cdot)$ denote the induced probability measure over the sample paths of $x(\cdot)$ in the limit experiment, and $\mathbb{E}_\mu[\cdot]$ its corresponding expectation. Due to the nature of the stopping time, $x(\tau)$ can only take on two values $\gamma, -\gamma$. Let δ denote the sign of $x(\tau)$. Then, by sufficiency, any test φ, in the limit experiment can be written as a function only of τ, δ. Furthermore, by Proposition 2, any unbiased test, $\varphi(\tau, \delta)$, must satisfy $\mathbb{E}_0[\delta\varphi(\tau, \delta)] = 0$.

Fix some alternative $\mu \neq 0$ and consider the functional optimization problem

$$\max_{\varphi(\cdot)} \mathbb{E}_\mu[\varphi(\tau, \delta)] \equiv \mathbb{E}_0\left[\varphi(\tau, \delta)e^{\frac{1}{\sigma}\mu\delta\gamma - \frac{\tau}{2\sigma^2}\mu^2}\right] \tag{B.2}$$
$$\text{s.t } \mathbb{E}_0[\varphi(\tau, \delta)] \leq \alpha \text{ and } \mathbb{E}_0[\delta\varphi(\tau, \delta)] = 0.$$

Here, and in what follows, it should implicitly understood that the candidate functions, $\varphi(\cdot)$, are tests, i.e., their range is $[0, 1]$. Let φ^* denote the optimal solution to (B.2). Note that φ^* is unbiased since $\varphi = \alpha$ also satisfies the constraints in (B.2); indeed, $\mathbb{E}_0[\delta] = 0$ by symmetry. Consequently, if φ^* is shown to be independent of μ, we can conclude that it is the best unbiased test.

Now, by Fudenberg et al. (2018), δ is independent of τ given μ. Furthermore, by symmetry, $\mathbb{P}_0(\delta = 1) = \mathbb{P}_0(\delta = -1) = 1/2$ for $\mu = 0$. Based on these results, we have

$$\mathbb{E}_0[\delta\varphi(\tau, \delta)] = \frac{1}{2}\int \{\varphi(\tau, 1) - \varphi(\tau, 0)\}\, dF_0(\tau),$$
$$\mathbb{E}_0[\varphi(\tau, \delta)] = \frac{1}{2}\int \{\varphi(\tau, 1) + \varphi(\tau, 0)\}\, dF_0(\tau), \text{ and}$$
$$\mathbb{E}_0\left[\varphi(\tau, \delta)e^{\frac{1}{\sigma}\mu\delta\gamma - \frac{\tau}{2\sigma^2}\mu^2}\right] = \frac{e^{\mu\gamma/\sigma}}{2}\int \varphi(\tau, 1)e^{-\frac{\tau}{2\sigma^2}\mu^2}dF_0(\tau) + \frac{e^{-\mu\gamma/\sigma}}{2}\int \varphi(\tau, 0)e^{-\frac{\tau}{2\sigma^2}\mu^2}dF_0(\tau).$$

The first two equations above imply $\mathbb{E}_0[\varphi(\tau, 1)] = \mathbb{E}_0[\varphi(\tau, 0)] = \mathbb{E}_0[\varphi(\tau, \delta)]$ when $\mathbb{E}_0[\delta\varphi(\tau, \delta)] = 0$. Hence, we can rewrite the optimization problem (B.2) as

$$\max_{\varphi(\cdot)}\left\{\frac{e^{\mu\gamma/\sigma}}{2}\int \varphi(\tau, 1)e^{-\frac{\tau}{2\sigma^2}\mu^2}dF_0(\tau) + \frac{e^{-\mu\gamma/\sigma}}{2}\int \varphi(\tau, 0)e^{-\frac{\tau}{2\sigma^2}\mu^2}dF_0(\tau)\right\} \tag{B.3}$$
$$\text{s.t.} \int \varphi(\tau, 1)dF_0(\tau) \leq \alpha, \int \varphi(\tau, 0)dF_0(\tau) \leq \alpha \text{ and}$$
$$\int \varphi(\tau, 1)dF_0(\tau) = \int \varphi(\tau, 0)dF_0(\tau).$$

Let us momentarily disregard the last constraint in (B.3). Then the optimization problem factorizes, and the optimal $\varphi(\cdot)$ can be determined by separately solving for $\varphi(\cdot,1), \varphi(\cdot,0)$ as the functions that optimize

$$\max_{\varphi(\cdot,a)} \int \varphi(\tau,a)e^{-\frac{\tau}{2\sigma^2}\mu^2}dF_0(\tau) \quad \text{s.t.} \quad \int \varphi(\tau,a)dF_0(\tau) \leq \alpha$$

for $a \in \{0,1\}$. Let $\varphi^*(\cdot,a)$ denote the optimal solution. It is immediate from the optimization problem above that $\varphi^*(\tau,1) = \varphi^*(\tau,0) := \varphi^*(\tau)$, i.e., the optimal φ^* is independent of δ. Hence, the last constraint in (B.3) is satisfied. Furthermore, by the Neyman-Pearson lemma,

$$\varphi^*(\tau) = \mathbb{I}\left\{e^{-\frac{\tau}{2\sigma^2}\mu^2} \geq \gamma\right\} \equiv \mathbb{I}\left\{\tau \leq c\right\},$$

where $c = F_0^{-1}(\alpha)$ due to the requirement that $\int \varphi(\tau,a)dF_0(\tau) \leq \alpha$. Consequently, the solution, $\varphi^*(\cdot)$, to (B.2) is given by $\mathbb{I}\left\{\tau \leq F_0^{-1}(\alpha)\right\}$. This is obviously independent of μ. We conclude that it is the best unbiased test in the limit experiment.

Step 2: The finite sample counterpart of $\varphi^*(\cdot)$ is given by $\hat{\varphi}(\hat{\tau}) := \mathbb{I}\left\{\hat{\tau} \leq F_0^{-1}(\alpha)\right\}$, where it may be recalled that $\hat{\tau} = \inf\{t : |x_n(t)| \geq \gamma\}$. Fix some arbitrary $\boldsymbol{g} := (g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$. Let $P_{nT,\boldsymbol{g}}$ be defined as in the proof of Proposition 6. By similar arguments as in the proofs of Adusumilli (2022, Theorems 3 and 5),

$$\hat{\tau} \xrightarrow[P_{nT,\boldsymbol{g}}]{d} \tau := \inf\{t : |x(t)| \geq \gamma\}$$

along sub-sequences, where $x(t) \sim \sigma^{-1}\mu t + \tilde{W}(t)$ and $\mu := \langle\psi_1, g_1\rangle_1 - \langle\psi_0, g_0\rangle_0$. Hence,

$$\lim_{n\to\infty} \hat{\beta}(g_1, g_0) := \lim_{n\to\infty} P_{nT,(g_1,g_0)}\left(\hat{\tau} \leq F_0^{-1}(\alpha)\right) = \mathbb{P}_\mu\left(\tau \leq F_0^{-1}(\alpha)\right),$$

where $\mathbb{P}_\mu(\cdot)$ is the probability measure defined in Step 1. But $\tilde{\beta}^*(\mu) := P_\mu\left(\tau \leq F_0^{-1}(\alpha)\right)$ is just the power function of the best unbiased test, φ^*, in limit experiment. Hence, $\hat{\varphi}(\cdot)$ is an asymptotically optimal unbiased test.

B.3. **Supporting information for Section 5.2.**

B.3.1. *Nonparametric level-$\boldsymbol{\alpha}$ and conditionally unbiased tests.* First, we define nonparametric versions of the level-$\boldsymbol{\alpha}$ and conditionally unbiased requirements. We follow the same notation as in Section 3.4. A test, φ_n, of $H_0 : \mu_1 - \mu_0 = \mu/\sqrt{n}$ is said to asymptotically level-$\boldsymbol{\alpha}$ if

$$\sup_{\left\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1-\langle\psi_0,h_0\rangle_0=\mu\right\}} \limsup_n \int \mathbb{I}\{\hat{\tau}=k\}\varphi_n dP_{nT,\boldsymbol{h}} \leq \alpha_k \;\forall\; k. \tag{B.4}$$

Similarly, a test, φ_n, of $H_0 : \mu_1 - \mu_0 = \mu/\sqrt{n}$ vs $H_1 : \mu_1 - \mu_0 \neq \mu/\sqrt{n}$ is asymptotically conditionally unbiased if

$$\sup_{\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1-\langle\psi_0,h_0\rangle_0=\mu\}} \limsup_n \int \mathbb{I}\{\tau = k\}\varphi_n dP_{nT,\boldsymbol{h}}$$
$$\geq \inf_{\{\boldsymbol{h}:\langle\psi_1,h_1\rangle_1-\langle\psi_0,h_0\rangle_0\neq\mu\}} \liminf_n \int \varphi_n dP_{nT,\boldsymbol{h}}.$$

B.3.2. *Attaining the bound.* Recall the definition of $x_n(\cdot)$ in (3.7). While $x_n(\cdot)$ depends on the unknown quantities σ_1, σ_0, we can replace them with consistent estimates $\hat{\sigma}_1, \hat{\sigma}_0$ using data from the first batch without affecting the asymptotic results, so there is no loss of generality in taking them to be known. Let $\hat{\varphi} := \varphi^*(\hat{\tau}, x_n(\hat{\tau}))$ denote the finite sample counterpart of φ^*.

By an extension of Proposition 6 to α-spending tests, as in Theorem 2, the conditional power function, $\beta^*(\mu|k)$, of φ^* in the limit experiment is an upper bound on the asymptotic power function of any test in the original experiment. We now show that the local (conditional) power, $\hat{\beta}(g_1, g_0|k)$, of $\hat{\varphi}$ against sub-models $P_{1/\sqrt{n},g_1} \times P_{1/\sqrt{n},g_0}$ converges to $\beta^*(\mu|k)$. This implies that $\hat{\varphi}$ is an asymptotically optimal level-$\boldsymbol{\alpha}$ test in this experiment.

Fix some arbitrary $\boldsymbol{g} := (g_1, g_0) \in T(P_0^{(1)}) \times T(P_0^{(0)})$. Let $P_{nT,\boldsymbol{g}}$ be defined as in the proof of Proposition 6. By similar arguments as in the proofs of Adusumilli (2022, Theorems 3 and 5),

$$x_n(\cdot) \xrightarrow[P_{nT,\boldsymbol{g}}]{d} x(\cdot)$$

along sub-sequences, where $x(t) \sim \sigma^{-1}\mu t + \tilde{W}(t)$ and $\mu := \langle\psi_1, g_1\rangle_1 - \langle\psi_0, g_0\rangle_0$. Since $\hat{\tau}$ is a function of $x_n(\cdot)$, the above implies, by an application of the extended continuous mapping theorem (Van Der Vaart and Wellner, 1996, Theorem 1.11.1), that

$$\lim_{n\to\infty} \int \mathbb{I}\{\hat{\tau} = k\}\hat{\varphi} P_{nT,(g_1,g_0)} = \int \mathbb{I}\{\tau = k\}\varphi^* d\mathbb{P}_\mu, \text{ and}$$
$$\lim_{n\to\infty} \int \mathbb{I}\{\hat{\tau} = k\} P_{nT,(g_1,g_0)} = \int \mathbb{I}\{\tau = k\} d\mathbb{P}_\mu.$$

Hence, as long as $\mathbb{P}_0(\tau = k) \neq 0$, by the definition of conditional power, we obtain

$$\lim_{n\to\infty} \hat{\beta}(g_1, g_0|k) = \frac{\int \mathbb{I}\{\tau = k\}\varphi^* d\mathbb{P}_\mu}{\mathbb{I}\{\tau = k\} d\mathbb{P}_\mu} := \beta^*(\mu|k),$$

for any $\mu \in \mathbb{R}$. This implies that $\hat{\varphi}$ is asymptotically level-$\boldsymbol{\alpha}$ (as can be verified by setting $\mu = 0$ etc), and furthermore, its conditional power attains the upper bound $\beta^*(\cdot|k)$. Hence, $\hat{\varphi}$ is an asymptotically optimal level-$\boldsymbol{\alpha}$ test.

B.4. **Supporting results for the proof of Proposition 6.**

Lemma 2. *Consider the setup in the proof of Proposition 6. Let $P^{(a)}_{1/\sqrt{n},\boldsymbol{h}_a^\intercal \boldsymbol{g}_a}$ denote the probability sub-model for treatment a, and suppose that it satisfies the SLAN property*

$$\ln \frac{dP_{nt,\boldsymbol{h}_a^\intercal \boldsymbol{g}_a}}{dP_{nt,\mathbf{0}}}(\mathbf{y}_{nt}^{(a)}) = \frac{\boldsymbol{h}_a^\intercal}{\sqrt{n}} \sum_{i=1}^{\lfloor n\pi_a t\rfloor} \boldsymbol{g}_a(Y_i^{(a)}) - \frac{\pi_a t}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a + + o_{P_{nT,\mathbf{0}}}(1), \quad \textit{uniformly over } t.$$

Then, any test in the parametric model $\left\{P_{1/\sqrt{n},\boldsymbol{h}_1^\intercal \boldsymbol{g}_1} \times P_{1/\sqrt{n},\boldsymbol{h}_0^\intercal \boldsymbol{g}_0} : \boldsymbol{h}_1, \boldsymbol{h}_0 \in \mathbb{R}^2\right\}$ can be matched (along sub-sequences) by a test that depends only on $\boldsymbol{G}(\tau), \tau$ in the limit experiment.

Proof. Recall that $G_{a,n}(t) := n^{-1/2}\sum_{i=1}^{\lfloor n\pi_a t\rfloor} \boldsymbol{g}_a(Y_i^{(a)})$ for $a \in \{0,1\}$. Then, by the statement of the lemma, we have

$$\ln \frac{dP_{n\hat{\tau},\boldsymbol{h}_a^\intercal \boldsymbol{g}_a}}{dP_{n\hat{\tau},0}}(\mathbf{y}_{n\hat{\tau}}^{(a)}) = \boldsymbol{h}_a^\intercal G_{a,n}(\hat{\tau}) - \frac{\pi_a \hat{\tau}}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a + o_{P_{nT,\mathbf{0}}}(1), \tag{B.5}$$

for $a \in \{0,1\}$. In the proof of Proposition 6, we argued that

$$(\hat{\tau}, G_{1,n}(\hat{\tau}), G_{0,n}(\hat{\tau})) \xrightarrow[P_{nT,\mathbf{0}}]{d} (\tau, G_1(\tau), G_0(\tau)), \tag{B.6}$$

where $G_a(t) \sim \sqrt{\pi_a} W_a(t)$ with $W_1(\cdot), W(\cdot)$ being independent 2-dimensional Brownian motions; and τ is a $\mathcal{G}_t$-adapted stopping time. Equations (B.5) and (B.6) imply

$$\ln \frac{dP_{nt,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1,\boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}}{dP_{nt,\mathbf{0}}}(\mathbf{y}_{nt}) \xrightarrow[P_{nT,\mathbf{0}}]{d} \sum_{a\in\{0,1\}} \left\{\boldsymbol{h}_a^\intercal G_a(\tau) - \frac{\pi_a \tau}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a\right\}. \tag{B.7}$$

Now, any two-sample test, φ_n, is tight since $\varphi_n \in [0,1]$. Then, as in the proof of Theorem 1, we find that given any sequence $\{n_j\}$, there exists a further sub-sequence $\{n_{j_m}\}$ - represented as $\{n\}$ without loss of generality - such that

$$\begin{pmatrix} \varphi_n \\ \frac{dP_{nt,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1,\boldsymbol{h}_0^\intercal \boldsymbol{g}_0)}}{dP_{nt,\mathbf{0}}}(\mathbf{y}_{nt}) \end{pmatrix} \xrightarrow[P_{nT,\mathbf{0}}]{d} \begin{pmatrix} \bar{\varphi} \\ V \end{pmatrix}; \quad V \sim \exp \sum_a \left\{\boldsymbol{h}_a^\intercal G_a(\tau) - \frac{\pi_a \tau}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a\right\}, \tag{B.8}$$

where $\bar{\varphi} \in [0,1]$. Now, given that $G_a(t) \sim \sqrt{\pi_a} W_a(t)$,

$$V \sim \exp \sum_a \left\{\sqrt{\pi_a}\boldsymbol{h}_a^\intercal W_a(\tau) - \frac{\pi_a \tau}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a\right\}.$$

Clearly, V is the stochastic/Doléans-Dade exponential of $\sum_a \left\{\sqrt{\pi_a}\boldsymbol{h}_a^\intercal W_a(\tau)\right\}$. Since $W_1(\cdot), W_0(\cdot)$ are independent, the latter quantity is in turn distributed as $\left(\sum_a \pi_a \boldsymbol{h}_a^\intercal \boldsymbol{h}_a\right)^{1/2} \tilde{W}(t)$, where $\tilde{W}(\cdot)$ is standard 1-dimensional Brownian motion. Hence, by standard results on stochastic exponentials,

$$M(t) := \exp\sum_a \left\{ \boldsymbol{h}_a^\intercal G_a(t) - \frac{\pi_a t}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a \right\} \sim \exp\sum_a \left\{ \sqrt{\pi_a}\boldsymbol{h}_a^\intercal W_a(t) - \frac{\pi_a t}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a \right\}$$

is a martingale with respect to the filtration $\mathcal{G}_t$. Since τ is an $\mathcal{G}_t$-adapted stopping time, $E[V] \equiv E[M(\tau)] = E[M(0)] = 1$ using the optional stopping theorem.

The above then implies, as in the proof of Theorem 1, that

$$\lim_{n\to\infty} \beta_n(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0) := \lim_{n\to\infty} \int \varphi_n dP_{nT,(\boldsymbol{h}_1^\intercal \boldsymbol{g}_1, \boldsymbol{h}_0^\intercal \boldsymbol{g}_0)} = E\left[\bar{\varphi} e^{\sum_a \left\{ \boldsymbol{h}_a^\intercal G_a(\tau) - \frac{\pi_a \tau}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a \right\}} \right]. \quad \text{(B.9)}$$

Define $\varphi(\tau, \boldsymbol{G}(\tau)) := E[\bar{\varphi}|\tau, \boldsymbol{G}(\tau)]$; this is a test statistic since $\varphi \in [0,1]$. The right hand side of (B.9) then becomes

$$E\left[\varphi(\tau, \boldsymbol{G}(\tau)) e^{\sum_a \left\{ \boldsymbol{h}_a^\intercal G_a(\tau) - \frac{\pi_a \tau}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a \right\}} \right].$$

But by the Girsanov theorem, this is just the expectation, $\mathbb{E}_{\boldsymbol{h}}[\varphi(\tau, \boldsymbol{G}(\tau))]$, of $\varphi(\tau, \boldsymbol{G}(\tau))$ when $G_a(t) \sim \pi_a \boldsymbol{h}_a t + \sqrt{\pi_a} W_a(t)$. This proves the desired claim. □

Lemma 3. *Consider the limit experiment where one observes a stopping time τ and independent diffusion processes $G_1(\cdot), G_0(\cdot)$, where $G_a(t) := \pi_a \boldsymbol{h}_a t + \sqrt{\pi_a} W_a(t)$. Let σ, $x(\cdot)$ and $\mathcal{F}_t$ be as defined in the proof of Proposition 6, and suppose that τ is $\mathcal{F}_t$-adapted. Then, the optimal level-α test of $H_0 : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = 0$ vs $H_1 : (\sigma_1, 0)^\intercal \boldsymbol{h}_1 - (\sigma_0, 0)^\intercal \boldsymbol{h}_0 = \mu$ in the limit experiment is given by*

$$\varphi_\mu^*(\tau, x(\tau)) := \mathbb{I}\left\{ \mu x(\tau) - \frac{\mu^2}{2\sigma}\tau \geq \gamma \right\}.$$

Proof. For each a we employ a change of variables $\boldsymbol{h}_a \to \boldsymbol{\Delta}_a$ via $\boldsymbol{\Delta}_a = \Lambda_a \boldsymbol{h}_a$, where

$$\Lambda_a := \begin{bmatrix} \sigma_a & 0 \\ 0 & 1 \end{bmatrix}.$$

Set $\boldsymbol{\Delta} := (\boldsymbol{\Delta}_1, \boldsymbol{\Delta}_0)$. The null and alternative regions are then $H_0 \equiv \{\boldsymbol{\Delta} : (1,0)^\intercal \boldsymbol{\Delta}_1 - (1,0)^\intercal \boldsymbol{\Delta}_0 = 0\}$ and $H_1 \equiv \{\boldsymbol{\Delta} : (1,0)^\intercal \boldsymbol{\Delta}_1 - (1,0)^\intercal \boldsymbol{\Delta}_0 = \mu\}$. Let $\mathbb{P}_{\boldsymbol{\Delta}} \equiv \mathbb{P}_{\boldsymbol{h}}$ denote the induced probability measure over the sample paths generated by $G_1(\cdot), G_0(\cdot)$ between $t \in [0,T]$, when $G_a(t) \sim \pi_a \Lambda_a^{-1} \boldsymbol{\Delta}_a t + \sqrt{\pi_a} W_a(t)$. Also, recall that

$$x(t) := \frac{1}{\sigma}\left(\frac{\sigma_1}{\pi_1} z_1(t) - \frac{\sigma_0}{\pi_0} z_0(t) \right),$$

where $z_1(\cdot), z_2(\cdot)$ are the first components of $G_1(\cdot), G_0(\cdot)$.

Fix some $\bar{\boldsymbol{\Delta}} \equiv (\bar{\boldsymbol{\Delta}}_1, \bar{\boldsymbol{\Delta}}_0) \in H_1$. Let $\bar{\Delta}_{11}$ and $\bar{\Delta}_{01}$ denote the first components of $\bar{\boldsymbol{\Delta}}_1, \bar{\boldsymbol{\Delta}}_0$, and define γ, η so that

$$(\bar{\Delta}_{11}, \bar{\Delta}_{01}) = \left(\gamma + \frac{\sigma_1^2 \eta}{\pi_1}, \gamma - \frac{\sigma_0^2 \eta}{\pi_0}\right). \tag{B.10}$$

Clearly, $\eta = \mu/\sigma^2$ and $\gamma = \bar{\Delta}_{11} - \sigma_1^2\eta/\pi_1$. Now construct $\tilde{\boldsymbol{\Delta}} = (\tilde{\boldsymbol{\Delta}}_1, \tilde{\boldsymbol{\Delta}}_0)$ as follows: The second components of $\tilde{\boldsymbol{\Delta}}_1, \tilde{\boldsymbol{\Delta}}_0$ are the same as that of $\bar{\boldsymbol{\Delta}}_1, \bar{\boldsymbol{\Delta}}_0$. As for the first components, $\tilde{\Delta}_{11}, \tilde{\Delta}_{01}$ of $\tilde{\boldsymbol{\Delta}}_1, \tilde{\boldsymbol{\Delta}}_0$, take them to be

$$(\tilde{\Delta}_{11}, \tilde{\Delta}_{01}) = (\gamma, \gamma). \tag{B.11}$$

By construction, $(\tilde{\boldsymbol{\Delta}}_1, \tilde{\boldsymbol{\Delta}}_0) \in H_0$.

Consider testing $H_0' : \boldsymbol{\Delta} = \tilde{\boldsymbol{\Delta}}$ vs $H_1' : \boldsymbol{\Delta} = \bar{\boldsymbol{\Delta}}$. Let $\ln \frac{d\mathbb{P}_{\bar{\boldsymbol{\Delta}}}}{d\mathbb{P}_{\tilde{\boldsymbol{\Delta}}}}(\mathcal{G}_t)$ denote the likelihood ratio between the probabilities induced by the parameters $\tilde{\boldsymbol{\Delta}}, \bar{\boldsymbol{\Delta}}$ over the filtration $\mathcal{G}_t$. Since $G_1(\cdot), G_0(\cdot)$ are independent, the Girsanov theorem gives

$$\begin{aligned}
\ln \frac{d\mathbb{P}_{\bar{\boldsymbol{\Delta}}}}{d\mathbb{P}_{\tilde{\boldsymbol{\Delta}}}}(\mathcal{G}_t) &= \left(\bar{\boldsymbol{\Delta}}_1^\intercal \Lambda_1^{-1} G_1(\tau) - \frac{\pi_1 \tau}{2} \bar{\boldsymbol{\Delta}}_1^\intercal \Lambda_1^{-2} \bar{\boldsymbol{\Delta}}_1\right) - \left(\tilde{\boldsymbol{\Delta}}_1^\intercal \Lambda_1^{-1} G_1(\tau) - \frac{\pi_1 \tau}{2} \tilde{\boldsymbol{\Delta}}_1^\intercal \Lambda_1^{-2} \tilde{\boldsymbol{\Delta}}_1\right) \\
&\quad + \left(\bar{\boldsymbol{\Delta}}_0^\intercal \Lambda_0^{-1} G_0(\tau) - \frac{\pi_0 \tau}{2} \bar{\boldsymbol{\Delta}}_0^\intercal \Lambda_0^{-2} \bar{\boldsymbol{\Delta}}_0\right) - \left(\tilde{\boldsymbol{\Delta}}_0^\intercal \Lambda_0^{-1} G_0(\tau) - \frac{\pi_0 \tau}{2} \tilde{\boldsymbol{\Delta}}_0^\intercal \Lambda_0^{-2} \tilde{\boldsymbol{\Delta}}_0\right) \\
&= \sigma\eta x(\tau) - \frac{\eta^2\sigma^2}{2}\tau,
\end{aligned}$$

where the last step follows from some algebra after making use of (B.10) and (B.11). Based on the above, an application of the Neyman-Pearson lemma shows that the UMP test of $H_0' : \boldsymbol{\Delta} = \tilde{\boldsymbol{\Delta}}$ vs $H_1' : \boldsymbol{\Delta} = \bar{\boldsymbol{\Delta}}$ is given by

$$\varphi_\mu^* = \mathbb{I}\left\{\sigma\eta x(\tau) - \frac{\eta^2\sigma^2}{2}\tau \geq \tilde{\gamma}\right\} = \mathbb{I}\left\{\mu x(\tau) - \frac{\mu^2}{2\sigma}\tau \geq \gamma\right\}.$$

Here, γ is to be determined by the size requirement. Now, for any $\boldsymbol{\Delta} \in H_0$,

$$x(t) \equiv \frac{1}{\sigma}\left(\sqrt{\frac{\sigma_1^2}{\pi_1}}(1,0)^\intercal W_1(t) - \sqrt{\frac{\sigma_0^2}{\pi_0}}(1,0)^\intercal W_0(t)\right) \sim \tilde{W}(t),$$

where $\tilde{W}(\cdot)$ is standard 1-dimensional Brownian motion. Hence, the distribution of the sample paths of $x(\cdot)$ is independent of the value of $\boldsymbol{\Delta}$ under the null. Combined with the assumption that τ is $\mathcal{F}_t$-adapted, this implies φ_μ^* does not depend on $\tilde{\boldsymbol{\Delta}}$ and, by extension, $\bar{\boldsymbol{\Delta}}$, except through μ. Since $\bar{\boldsymbol{\Delta}} \in H_1$ was arbitrary, we are led to conclude φ_μ^* is UMP more generally for testing $H_0 \equiv \{\boldsymbol{\Delta} : (1,0)^\intercal\boldsymbol{\Delta}_1 - (1,0)^\intercal\boldsymbol{\Delta}_0 = 0\}$ vs $H_1 \equiv \{\boldsymbol{\Delta} : (1,0)^\intercal\boldsymbol{\Delta}_1 - (1,0)^\intercal\boldsymbol{\Delta}_0 = \mu\}$. ∎

B.5. **Supporting results for the proof of Proposition 8.**

Lemma 4. *Consider the limit experiment where one observes* $q_a = \sum_j \pi_j^{(a)}$ *and* $x_a := (1,0)^\intercal \sum_j \boldsymbol{Z}_j^{(a)}(\pi_j^{(a)})$, *where*

$$\boldsymbol{Z}_j^{(a)}(t) := \boldsymbol{h}_a t + W_j^{(a)}(t),$$

and π_j *is measurable with respect to*

$$\mathcal{F}_{j-1} \equiv \sigma\left\{(1,0)^\intercal \boldsymbol{Z}_l^{(a)}(\cdot); l \le j-1, a \in \{0,1\}\right\}.$$

Then, the optimal level-α test of $H_0 : ((1,0)^\intercal \boldsymbol{h}_1, (1,0)^\intercal \boldsymbol{h}_0) = (0,0)$ *vs* $H_1 : ((1,0)^\intercal \boldsymbol{h}_1, (1,0)^\intercal \boldsymbol{h}_0) = (\mu_1, \mu_0)$ *in the limit experiment is*

$$\varphi^*_{\mu_1,\mu_0} = \mathbb{I}\left\{\sum_{a\in\{0,1\}} \left(\mu_a x_a - \frac{q_a}{2}\mu_a^2\right) \ge \gamma_{\mu_1,\mu_0}\right\}.$$

Proof. Denote

$$H_0 \equiv \{\boldsymbol{h} : ((1,0)^\intercal \boldsymbol{h}_1, (1,0)^\intercal \boldsymbol{h}_0) = (0,0)\}, \quad \text{and}$$
$$H_1 \equiv \{\boldsymbol{h} : ((1,0)^\intercal \boldsymbol{h}_1, (1,0)^\intercal \boldsymbol{h}_0) = (\mu_1, \mu_0)\}.$$

Let $\mathbb{P}_h$ denote the induced probability measure over the sample paths generated by $\{z_j^{(a)}(t) : t \le \pi_j^{(a)}\}_{j,a}$.

Given any $(\boldsymbol{h}_1, \boldsymbol{h}_0) \in H_1$, define $\tilde{\boldsymbol{h}}_a = \boldsymbol{h}_a - (1,0)^\intercal \boldsymbol{h}_a (1,0)$ for $a \in \{0,1\}$. Note that $(\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_0) \in H_0$ and $(1,0)^\intercal \boldsymbol{h}_a = \mu_a$. Let

$$\ln \frac{d\mathbb{P}_{(\tilde{\boldsymbol{h}}_1,\tilde{\boldsymbol{h}}_0)}}{d\mathbb{P}_{(\boldsymbol{h}_1,\boldsymbol{h}_0)}}(\mathcal{G})$$

denote the likelihood ratio between the probabilities induced by the parameters $(\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_0)$, $(\boldsymbol{h}_1, \boldsymbol{h}_0)$ over the filtration

$$\mathcal{G} \equiv \sigma\left\{\boldsymbol{Z}_j^{(a)}(t) : t \le \pi_j^{(a)}; j = 1, \dots, J; a \in \{0,1\}\right\}.$$

By the Girsanov theorem, noting that $\{z_j^{(a)}(t) : t \le \pi_j^{(a)}\}_j$ are independent across a and defining $G_a := \sum_j \boldsymbol{Z}_j^{(a)}(\pi_j^{(a)})$, we obtain after some straightforward algebra that

$$\begin{aligned}
\ln \frac{d\mathbb{P}_{(\tilde{\boldsymbol{h}}_1,\tilde{\boldsymbol{h}}_0)}}{d\mathbb{P}_{(\boldsymbol{h}_1,\boldsymbol{h}_0)}}(\mathcal{F}) &= \sum_a \left\{\left(\tilde{\boldsymbol{h}}_a^\intercal G_a - \frac{q_a}{2}\tilde{\boldsymbol{h}}_a^\intercal \tilde{\boldsymbol{h}}_a\right) - \left(\boldsymbol{h}_a^\intercal G_a - \frac{q_a}{2}\boldsymbol{h}_a^\intercal \boldsymbol{h}_a\right)\right\} \\
&= \sum_a \left(\mu_a x_a(\tau) - \frac{\mu_a^2}{2} q_a\right),
\end{aligned}$$

where x_a is the first component of G_a. Hence, an application of the Neyman-Pearson lemma shows that the UMP test of $H_0' : \boldsymbol{h} = (\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_0)$ vs $H_1' : \boldsymbol{h} = (\boldsymbol{h}_1, \boldsymbol{h}_0)$ is given by

$$\varphi^*_{\mu_1,\mu_0} = \mathbb{I}\left\{\sum_a \left(\mu_a x_a(\tau) - \frac{\mu_a^2}{2} q_a\right) \geq \gamma\right\},$$

where γ is determined by the size requirement.

Now, for any $\boldsymbol{h} \in H_0$, both x_a and q_a measurable with respect to $\mathcal{F}$ by assumption. Since $(1,0)^\intercal \boldsymbol{Z}_j^{(a)}(\cdot)$ is independent of $\boldsymbol{h}_a$ given μ_a for all j, a, it follows that the distribution of x_a, q_a is independent of the value of $\boldsymbol{h} \in H_0$ under the null. This implies that $\varphi^*_{\mu_1,\mu_0}$ does not depend on $(\tilde{\boldsymbol{h}}_1, \tilde{\boldsymbol{h}}_0)$ and, by extension, $(\boldsymbol{h}_1, \boldsymbol{h}_0)$, except through $(\mu_1 \mu_0)$. Since $(\boldsymbol{h}_1, \boldsymbol{h}_0) \in H_1$ was arbitrary, we are led to conclude $\varphi^*_{\mu_1,\mu_0}$ is UMP more generally for testing the composite hypotheses H_0 vs H_1. □