

Sequential testing for elicitable functionals via supermartingales

PHILIPPE CASGRAIN^{1,a}, MARTIN LARSSON^{2,b} and JOHANNA ZIEGEL^{3,c}

¹*ETH Zürich, Zürich, Switzerland, philippe.casgrain@math.ethz.ch*

²*Carnegie Mellon University, Pittsburgh, USA, larsson@cmu.edu*

³*University of Bern, Bern, Switzerland, johanna.ziegel@unibe.ch*

We design sequential tests for a large class of nonparametric null hypotheses based on elicitable and identifiable functionals. Such functionals are defined in terms of scoring functions and identification functions, which are ideal building blocks for constructing nonnegative supermartingales under the null. This in turn yields sequential tests via Ville's inequality. Using regret bounds from Online Convex Optimization, we obtain rigorous guarantees on the asymptotic power of the tests for a wide range of alternative hypotheses. Our results allow for bounded and unbounded data distributions, assuming that a sub- ψ tail bound is satisfied.

Keywords: Anytime valid testing; elicitable functionals; identifiable functionals; online optimization; sequential statistics

1. Introduction

We design sequential tests and confidence sequences for a large class of nonparametric null hypotheses based on elicitable and identifiable functionals. Such functionals include moments, quantiles, expectiles, and many other examples, all of which can be tested using the approach developed here. The null hypotheses that we cover are highly composite and nonparametric; for instance, the null could consist of *all* distributions whose median, say, is a given value. Our tests are *sequential*, or *anytime valid*, in the sense that data is observed sequentially through time, and at each point in time the decision to stop or continue may depend on all available data without compromising Type-I error guarantees. We also obtain guarantees on the power of our tests with respect to large composite nonparametric alternative hypotheses.

The basic mechanism we use to construct sequential tests rests on the notion of *test (super)martingales* due to Shafer et al. (2011). The idea is simple but powerful: a test statistic that is a nonnegative supermartingale if the null hypothesis is true can only reach large values with small probability. This can be quantified using Ville's inequality. Thus if one rejects the null only when a sufficiently large value of the test statistic has been observed, Type-I error control is ensured.

Our contribution rests on the observation that elicitable and identifiable functionals in the sense of Fissler and Ziegel (2016), Gneiting (2011), Lambert, Pennock and Shoham (2008) are ideal for constructing test supermartingales. Combined with a construction known as predictable mixing, one immediately obtains large families of test supermartingales which can be used as possible test statistics for the null hypothesis defined by a particular elicitable or identifiable functional.

The predictable mixing construction can be interpreted in terms of betting or trading. Finding a useful predictable mixture corresponds to determining a profitable trading strategy. The supermartingale condition under the null ensures that profits are limited if the null is true. But if the null is false, it may be possible to “bet against the null” in a way that leads to large profits and, hence, reject the null. Doing so requires two things. First, in order to bet against the null, one must specify a suitable distribution to bet on. Second, given this distribution, one must find a strategy that is likely to be profitable.

We address these two points at once by making use of ideas from online convex optimization (OCO). We demonstrate how off-the-shelf algorithms can be used to produce strong trading strategies. This leads to powerful test supermartingales given as predictable mixtures of the basic set of test supermartingales constructed from the elicitable or identifiable functional used to specify the null hypothesis. A major advantage of this approach is that these algorithms come with performance guarantees in the form of regret bounds. These regret bounds translate into rigorous guarantees on the power of the resulting sequential test under a wide variety of alternative hypotheses.

Sequential testing goes back to [Wald \(1945\)](#). A large body of literature on the subject exists, and martingale techniques have played an important role from the beginning. This is notable in the work of Darling, Lai, Robbins, and Siegmund going back to the 1960s, although the majority of their work was in a parametric context. We refer to Appendix F in [Waudby-Smith and Ramdas \(2020\)](#) for a historical overview and numerous references. The concept of a test martingale was introduced in [Shafer et al. \(2011\)](#), and there has recently been a number of papers related to this circle of ideas, for instance [Howard et al. \(2020, 2021\)](#) which derive time-uniform confidence sequences and concentration bounds. In particular, the closely related notion of e-variables and e-processes have received significant attention; see e.g. [Grünwald, de Heide and Koolen \(2019\)](#), [Ramdas et al. \(2020, 2022a\)](#), [Vovk and Wang \(2021\)](#), [Xu, Wang and Ramdas \(2021\)](#) as well as Remark 2.3 below. Most closely related to our paper is the work of [Waudby-Smith and Ramdas \(2020\)](#), which develops confidence sequences for the mean of a sequence of bounded random variables. That paper makes use of the same betting perspective, which enables the authors to obtain powerful confidence sequences. It also discusses various related strands of literature and the history of the subject. However, the authors do not consider other functionals beyond the mean, and they rely on the boundedness of the data in an essential way. Our work generalizes both of these points. Moreover, they do not make use of OCO to obtain regret bounds which then translate into statements about power. Our use of regret bounds is reminiscent of [Ramdas et al. \(2022a\)](#), where regret bounds are used to derive power guarantees for the particular problem of testing exchangeability of binary sequences. OCO and regret bounds are also ubiquitous in the online learning literature, and we point in particular to the papers [Jun and Orabona \(2019\)](#), [Orabona and Jun \(2021\)](#), [Shekhar and Ramdas \(2021\)](#) for further applications and developments based on betting ideas and martingales, as well as additional pointers to the literature. However, these papers are not concerned with elicitable and identifiable functionals. Another paper related to ours is [Henzi and Ziegel \(2022\)](#), which treats the problem of probability forecasting. The authors rely on the betting analogy to construct sequential tests for the statistical significance of score differences of competing forecasts. The paper [Choe and Ramdas \(2021\)](#) studies a sequential forecasting problem using martingales based on scoring functions.

The concepts of elicibility and identifiability go back to the PhD thesis of [Osband \(1985\)](#). However, the term elicibility was coined later ([Lambert, Pennock and Shoham, 2008](#)), and it was popularized by [Fissler and Ziegel \(2016\)](#), [Frongillo and Kash \(2021\)](#), [Gneiting \(2011\)](#), [Steinwart et al. \(2014\)](#). Prior to our work, no systematic approach to the sequential testing problem for elicitable and identifiable functionals has been developed. Thus, our paper demonstrates an essential link between test supermartingales and the concepts of elicibility and identifiability.

The paper is organized as follows. In Section 2, we review the definition of test supermartingales and discuss how they can be used to construct powerful sequential tests via predictable mixing. In Section 3, we discuss elicibility and identifiability, both as a way of specifying nonparametric null hypotheses and as the basis for constructing test supermartingales. Importantly, we show how sub- ψ tail bounds can be leveraged to handle unbounded data. Section 4 discusses how regret bounds from online convex optimization lead to statements about asymptotic power of the tests. In Section 5, we briefly discuss the related issue of confidence sequences. Section 6 contains a simulation study illustrating the techniques developed in this paper. The proofs of all results are collected in Appendix B.

2. Sequential testing via supermartingales

We consider a sequential testing environment in which a discrete-time stochastic process $(X_t)_{t \in \mathbb{N}}$, taking values in some measurable space X , is observed sequentially through time. The process $(X_t)_{t \in \mathbb{N}}$ is called the *data generating process*. Concrete examples of data generating processes include patient data collected from clinical trials or daily profit and loss values of a trading strategy. The filtration generated by the data generating process is denoted $\mathcal{F} := (\mathcal{F}_t)_{t \in \mathbb{N}}$ where $\mathcal{F}_t := \sigma(X_1, \dots, X_t)$ is the information set generated by the data collected until time t . We let \mathcal{F}_0 denote the trivial σ -algebra.

A statistical hypothesis is a collection $\mathcal{H} \subseteq \mathcal{M}_1(X^{\mathbb{N}})$, where $\mathcal{M}_1(X^{\mathbb{N}})$ is the set of all possible probability distributions of the data generating process. Thus an element $P \in \mathcal{M}_1(X^{\mathbb{N}})$ is a distribution of the entire sequence $(X_t)_{t \in \mathbb{N}}$. The hypothesis \mathcal{H} encodes the belief that the realized data was governed by one of the distributions $P \in \mathcal{H}$. A (*sequential*) *test* for a given null hypothesis \mathcal{H}_0 is defined as an \mathcal{F} -stopping time τ that specifies the time at which \mathcal{H}_0 is rejected. The requirement that τ be a stopping time means that the decision whether to stop and reject \mathcal{H}_0 , or to continue and observe more data, is only based on data available at the time the decision is made. Stopping times are allowed to take the value infinity, and this corresponds to the possibility that the test never rejects the null.

2.1. Anytime validity and test supermartingales

In contrast to traditional hypothesis testing with a fixed and known total sample size, the total sample size that will be produced by the X_t before stopping is not known in advance. As a consequence, repeatedly evaluating a test designed for fixed finite sample sizes will generate an inflated Type-I error (or size) of the test, see for example [Albers \(2019\)](#), [O’Neill and Wetherill \(1971\)](#). There are techniques, such as multiple comparison p-value adjustments, to correctly modify tests a-posteriori ([Hsu, 1996](#)). However, these methods have quickly decaying power as the number of repeated tests grows large and require that the number of tests to be performed be known in advance. In order to avoid these issues, we work with the concept of a *sequential*, or *anytime valid*, test. This allows for data-dependent testing policies that need not necessarily be specified in advance.

Definition 2.1 (sequential test). Let a null hypothesis \mathcal{H}_0 be given, and $\alpha \in (0, 1)$. A (valid) *level- α sequential test* for \mathcal{H}_0 is a stopping time τ_0 such that $P(\tau_0 < \infty) \leq \alpha$ for all $P \in \mathcal{H}_0$.

Validity thus refers to Type-I error control. Alternatively, a sequential test τ_0 can be understood through the rejection indicator process $\chi_t = \mathbf{1}_{\{\tau_0 \leq t\}}$. This process is zero until there is enough evidence to reject the null, and then becomes one. The validity property of τ_0 can be shown to be equivalent to

$$P(\chi_\tau = 1) \leq \alpha \text{ for all finite stopping times } \tau \text{ and all } P \in \mathcal{H}_0.$$

This formulation emphasizes the idea of *anytime validity*: regardless of when you inspect the level of evidence, even if you do so at a random and potentially data dependent time, validity is preserved. For further discussion of anytime valid testing and inference, see e.g. [Ramdas et al. \(2020, 2022b\)](#), [Waudby-Smith and Ramdas \(2020\)](#).

A classical method for constructing sequential tests is based on nonnegative supermartingales. The following definition goes back to [Shafer et al. \(2011\)](#).

Definition 2.2 (test supermartingale). Let a null hypothesis \mathcal{H}_0 be given. A *test supermartingale* (for \mathcal{H}_0) is a nonnegative adapted process $W := (W_t)_{t \in \mathbb{N}}$ with initial value $W_0 \leq 1$ that is a P -supermartingale for all $P \in \mathcal{H}_0$.

We recall that a random process $(Z_t)_{t \in \mathbb{N}}$ adapted to a filtration $(\mathcal{G}_t)_{t \in \mathbb{N}}$ such that $\mathbb{E}_P[|Z_t|] < \infty$ is a P -supermartingale (P -submartingale) if for all $t \in \mathbb{N}$ we have $\mathbb{E}_P[Z_{t+1} \mid \mathcal{G}_t] \leq Z_t$ ($\mathbb{E}_P[Z_{t+1} \mid \mathcal{G}_t] \geq Z_t$). A process that is both a P -supermartingale and a P -submartingale is called a P -martingale, and satisfies the equality $\mathbb{E}_P[Z_{t+1} \mid \mathcal{G}_t] = Z_t$ for all $t \in \mathbb{N}$. Here, \mathbb{E}_P means that the expectation is taken with respect to the probability measure $P \in \mathcal{M}_1(X^{\mathbb{N}})$.

Test supermartingales can be used to construct sequential tests. The basic tool for showing validity is Ville’s inequality (Ville, 1939), which states that any nonnegative P -supermartingale W with $W_0 \leq 1$ satisfies $P(\sup_{t \in \mathbb{N}} W_t > 1/\alpha) \leq \alpha$ for all $\alpha \in (0, 1)$. Thus, if W is a test supermartingale and $\alpha \in (0, 1)$ is fixed, then the test which rejects the null as soon as W reaches a value above $1/\alpha$,

$$\tau_0 = \inf \left\{ t \in \mathbb{N} : W_t > \frac{1}{\alpha} \right\}, \tag{1}$$

satisfies

$$P(\tau_0 < \infty) = P \left(\sup_{t \in \mathbb{N}} W_t > \frac{1}{\alpha} \right) \leq \alpha,$$

and is therefore valid at level α .

Remark 2.3. A notion closely related to test supermartingales is that of an e -process, which is a nonnegative adapted process W such that $\mathbb{E}_P[W_\tau] \leq 1$ for all $P \in \mathcal{H}_0$ and all stopping times τ (Ramdas et al., 2020, 2022a, Vovk and Wang, 2021, Xu, Wang and Ramdas, 2021). The stopping theorem implies that every test supermartingale is an e -process, but the converse is not true (Ramdas et al., 2020, 2022a). The ‘static’ or non-sequential analog of an e -process is known as an e -variable, which is a nonnegative random variable E such that $\mathbb{E}_P[E] \leq 1$ for all $P \in \mathcal{H}_0$. These notions have recently been studied extensively as a tool for safe inference (Grünwald, de Heide and Koolen, 2019).

2.2. Power and growth

In addition to validity, we are interested in power against suitable alternative hypotheses \mathcal{H}_1 disjoint from \mathcal{H}_0 . Loosely speaking, good power means that if the true data generating distribution belongs to \mathcal{H}_1 , the test should reject \mathcal{H}_0 quickly with high probability. For tests arising from test supermartingales W via (1), good power is achieved by designing W to grow quickly with high probability under distributions in \mathcal{H}_1 . Quick growth of W can be achieved by optimizing the Growth Rate Optimal (GRO) criterion, which has recently received significant attention in the context of e -values and e -process (Grünwald, de Heide and Koolen, 2019). In our setting, the GRO criterion is as follows. At each time T one seeks to maximize the expected logarithmic increment conditionally on data observed so far across all test supermartingale increments. More formally one aims to solve

$$\text{maximize } \mathbb{E}_Q \left[\log \frac{W_{T+1}}{W_T} \mid \mathcal{F}_T \right] \quad \text{subject to} \quad \sup_{P \in \mathcal{H}_0} \mathbb{E}_P \left[\frac{W_{T+1}}{W_T} \mid \mathcal{F}_T \right] \leq 1 \tag{2}$$

given the observed data X_1, \dots, X_T , where Q is a suitable distribution. As we will see, Q need not itself be the only element of \mathcal{H}_1 , or even belong to \mathcal{H}_1 at all. It is a purely computational device used to guide the choice of W . In implementing this idea, faces three key issues:

1. The problem (2) optimizes over the set of all test supermartingales. Solving it requires a description of this set, or of a sufficiently rich subset.

2. A suitable distribution Q has to be specified.
3. One has to actually solve (2), at least numerically, and ideally derive performance guarantees with respect to the set \mathcal{H}_1 of alternatives.

In this paper, we consider null hypotheses \mathcal{H}_0 based on elicitable functionals and identifiable functionals, which admit large families of explicit test supermartingales. This addresses 1. In order to address 2 we focus on distributions Q which are not fixed in advance but rather learned in an online fashion as more and more data is observed. This idea has recently also been explored by [Waudby-Smith and Ramdas \(2020\)](#). Having dealt with 1 and 2, the GRO criterion becomes a concrete optimization problem which we solve using methods from Online Convex Optimization (OCO). A key feature of this approach is that OCO methods come with asymptotic performance guarantees in the form of regret bounds. We employ these bounds to show that the resulting tests have asymptotic power one under a large composite nonparametric alternative hypothesis \mathcal{H}_1 , in the sense that we obtain a test supermartingale which tends to infinity with probability one under every distribution in \mathcal{H}_1 . Consequently, if the true data distribution is some element of \mathcal{H}_1 then any test τ_0 of the form (1) is guaranteed to eventually reject the null: $P(\tau_0 < \infty) = 1$ for all $P \in \mathcal{H}_1$. This power guarantee addresses 3.

2.3. Test supermartingales via mixing

The null hypotheses considered in this paper will be constructed directly in terms of explicit families of test supermartingales $L^\theta = (L_t^\theta)_{t \in \mathbb{N}}$ indexed by a parameter $\theta \in \Theta$, where Θ is an (arbitrary) index set. Whenever such a family $\{L^\theta\}_{\theta \in \Theta}$ is available, it is possible to construct new test supermartingales by combining its members. For instance, it is clear that any convex combination of test supermartingales is again a test supermartingale. More generally, one can use predictably mixed test supermartingales as shown in the following lemma; see also [Waudby-Smith and Ramdas \(2020\)](#). In this way, one can assemble weak test supermartingales into more powerful ones.

Lemma 2.4 (Predictably mixed supermartingale). *Let $\{L^\theta\}_{\theta \in \Theta}$ be a family of test supermartingales and $(\pi_t)_{t \in \mathbb{N}}$ a predictable sequence of probability measures on Θ . Then the process $W = (W_t)_{t \in \mathbb{N}}$ defined by $W_0 = 1$ and*

$$W_t = \prod_{i=1}^t \int_{\Theta} \frac{L_i^\theta}{L_{i-1}^\theta} \pi_i(d\theta), \quad \text{for all } t \in \mathbb{N}, \tag{3}$$

is also a test supermartingale.

To be precise, we assume here that Θ is a measurable space, and that $(\omega, \theta) \mapsto L_t^\theta(\omega)$ is measurable for each t . The condition on $(\pi_t)_{t \in \mathbb{N}}$ means that for each $t \in \mathbb{N}$, $\pi_t = \pi_t(X_1, \dots, X_{t-1}; d\theta)$ is a probability measure on Θ that may depend in a measurable way on the preceding data points.

Example 1 (i.i.d. Gaussians). Consider the hypothesis $\mathcal{H}_0 = \{P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : X_t \sim \text{i.i.d. } \mathcal{N}(0, 1)\}$. Then for every $\lambda \in \mathbb{R}$ the process L^λ defined by $L_0^\lambda = 1$ and for $t \in \mathbb{N}$ by $L_t^\lambda = \exp\left(\lambda \sum_{i=1}^t X_i - \frac{1}{2} \lambda^2 t\right)$ is a test (super)martingale. The predictably mixed process W in (3) then takes the form

$$W_t = \prod_{i=1}^t \int_{\mathbb{R}} \exp\left(\lambda X_i - \frac{1}{2} \lambda^2 t\right) \pi_i(d\lambda).$$

Remark 2.5. The mixing construction (3) admits a useful interpretation in terms of trading a portfolio of financial assets. Treating the collection of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$ as a collection of tradable assets indexed by $\theta \in \Theta$, we may think of W_t in (3) as the value of a portfolio trading these assets. Indeed, regard $\pi_t(d\theta)$ as the portfolio weights specifying the proportion of capital allocated to each asset θ at time $t - 1$; this is observable at time $t - 1$ because $(\pi_t)_{t \in \mathbb{N}}$ is a predictable sequence. The portfolio return from time $t - 1$ to t is the weighted average of the individual asset returns,

$$\int_{\Theta} \frac{L_t^\theta - L_{t-1}^\theta}{L_{t-1}^\theta} \pi_t(d\theta).$$

Rearranging (3) one sees that this is equal to the overall portfolio return $(W_t - W_{t-1})/W_{t-1}$. We can think of selecting a strong allocation strategy $(\pi_t)_{t \in \mathbb{N}}$ as choosing bets against \mathcal{H}_0 in order to make W_t grow quickly, eventually exceeding the threshold α^{-1} to reject \mathcal{H}_0 . Indeed, if \mathcal{H}_0 is false, there may exist assets L^θ which are not supermartingales, enabling one to ‘bet against the null’ by selecting $(\pi_t)_{t \in \mathbb{N}}$ with weights on these processes such that the wealth process grows on average. In contrast, if \mathcal{H}_0 is true, Ville’s inequality shows that, regardless of the trading strategy employed, it is unlikely (with probability bounded by α) that our wealth ever exceeds the threshold α^{-1} .

Remark 2.6. By choosing $\pi_t(d\theta) = L_{t-1}^\theta F(d\theta) / \int_{\Theta} L_{t-1}^{\theta'} F(d\theta')$ in Lemma 2.4 for some fixed probability distribution $F(d\theta)$ on Θ , we obtain the wealth process $W_t = \int_{\Theta} L_t^\theta F(d\theta)$. This is the Robbins–Siegmund mixture (super-) martingale.

In practice it is not feasible to work with general predictable sequences $(\pi_t)_{t \in \mathbb{N}}$. Instead we consider parsimonious specifications that tend to work well in experiments. A key example is the Dirac specification, $\pi_t = \delta_{\theta_t}$, where $(\theta_t)_{t \in \mathbb{N}}$ is a Θ -valued predictable process. This is the simplest possible specification. In terms of the trading interpretation in Remark 2.5, a strategy of this kind chooses in each period one single asset where all capital is invested. The test supermartingale (3) simplifies to

$$W_t = \prod_{i=1}^t \frac{L_i^{\theta_i}}{L_{i-1}^{\theta_i}}. \tag{4}$$

The form 4 is similar in form to the adaptive likelihood ratio statistics of Robbins and Siegmund (1972), although our family of supermartingales L_t^θ will usually not be obtained as likelihood ratios.

3. Specifying the null hypothesis

We consider null hypotheses involving the value of certain statistical functionals of the (conditional) distributions of the data. For example, for a given value λ_0 we may want to test the hypothesis

$$\mathcal{H}_0 = \{P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : \lambda_0 \text{ is a median of the conditional distribution } P(X_t \in \cdot \mid \mathcal{F}_{t-1}) \text{ for all } t \in \mathbb{N}\}.$$

The hypotheses considered below generalize this example beyond medians to a large class of *elicitable functionals* and *identifiable functionals*. These concepts are reviewed below; they include quantiles, moments, expectiles, and many other examples. The key common feature of these hypotheses is that they can be expressed in the form

$$\mathcal{H}_0 = \{P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : L^\theta \text{ is a } P\text{-supermartingale for all } \theta \in \Theta\} \tag{5}$$

for some explicit family of nonnegative processes $L^\theta = (L_t^\theta)_{t \in \mathbb{N}}$ starting at $L_0^\theta = 1$, indexed by a parameter $\theta \in \Theta$ where Θ is an index set. In our applications Θ will be a subset of a finite-dimensional space. Thus by construction, $\{L^\theta\}_{\theta \in \Theta}$ constitutes a family of ‘base’ test supermartingales for \mathcal{H}_0 which can be used to form other test supermartingales through predictable mixing as explained in Subsection 2.3. This leads to general procedures for designing sequential tests for elicitable and identifiable functionals.

3.1. Definition of elicibility and identifiability

We review the definitions as given in Fissler and Ziegel (2016). Fix $k \in \mathbb{N}$ and a subset $\Lambda \subseteq \mathbb{R}^k$. A *scoring function* is simply a measurable map $s: \Lambda \times \mathcal{X} \rightarrow \mathbb{R}$. Let \mathcal{M} be a class of probability distributions on \mathcal{X} . If for each distribution $\mu \in \mathcal{M}$ the map

$$\lambda \mapsto \mathbb{E}_\mu[s(\lambda, X)] \tag{6}$$

is well-defined and finite, let $T(\mu)$ denote the set of its minimizers. The induced map T is called an *elicitable functional (with respect to \mathcal{M})* and $s(\lambda, x)$ a *strictly consistent scoring function* for T . Here, X denotes the canonical random variable on \mathcal{X} . If $T(\mu)$ only consists of one element, that is the minimizer in (6) is unique, we abuse notation and also use the notation $T(\mu)$ for the minimizer. Any given elicitable functional can have many different strictly consistent scoring functions.

Similarly, an *identification function* is a measurable map $m: \Lambda \times \mathcal{X} \rightarrow \mathbb{R}^k$. If for each distribution $\mu \in \mathcal{M}$ the map

$$\lambda \mapsto \mathbb{E}_\mu[m(\lambda, X)] \tag{7}$$

is well-defined and finite, let $T(\mu)$ denote the set of its zeros. Then T is called an *identifiable functional (with respect to \mathcal{M})* and $m(\lambda, x)$ a *strict identification function* for T . A zero of the expected identification function in (7) is understood to hold component-wise, since $m(\lambda, X) \in \mathbb{R}^k$. We remark that the identifiability of a functional is conceptually different from the identifiability of parameters in statistical models. However, on a certain level there are connections between the two concepts, see Dimitriadis, Fissler and Ziegel (2023, Supplementary Material). When applying an elicitable or identifiable functional T to a distribution μ in the following, we always implicitly assume that the functional is well-defined for this distribution.

Table 1 contains some examples of commonly used functionals that happen to be both elicitable and identifiable. The presented scoring functions are standard but not strictly consistent on the maximal possible domain \mathcal{M} of definition of the respective functionals. Different choices of strictly consistent scoring functions allow to show elicibility of these functionals on their natural domains of definition; see Gneiting (2011) for details. Some important functionals are neither identifiable nor elicitable: Examples are the variance and the expected shortfall, where the latter is a prominent risk measure in finance and insurance.

If T is an elicitable functional with a scoring function $s(\lambda, x)$ that is convex in λ , then, under suitable conditions, T is also an identifiable functional with identification function $m(\lambda, x) \in \partial_\lambda s(\lambda, x)$, an element of the subgradient of the scoring function with respect to λ . For the converse, linking an identification function $m(\lambda, x)$ to a unique convex scoring function $s(\lambda, x)$, more subtle conditions are needed, and we point interested readers to Rockafellar and Wets (2009, Theorem 12.25). In the absence of convexity, scoring and identification functions are still linked through gradients under sufficient differentiability assumptions which are formalized as Osband’s principle in Fissler and Ziegel (2016).

	$T(\mu)$	$s(\lambda, x)$	$m(\lambda, x)$
Mean	$\mathbb{E}_\mu[X]$	$\frac{1}{2}(x - \lambda)^2$	$x - \lambda$
α-Quantile	$\inf\{\lambda : \alpha \geq \mu(X < \lambda)\}$	$ x - \lambda \left(\alpha \mathbb{1}_{\{x < \lambda\}} + (1 - \alpha) \mathbb{1}_{\{x > \lambda\}} \right)$	$\mathbb{1}_{\{x > \lambda\}} - \alpha$
Regression	$\operatorname{argmin}_{\lambda \in \mathbb{R}^k} E_\mu[(\lambda'Y - Z)^2]$	$\frac{1}{2} \ \lambda'y - z\ ^2$	$\lambda yy' - zy'$

Table 1. Examples of statistical quantities that can be expressed as elicitable and identifiable functionals. In the elicitable case, for the mean, \mathcal{M} is the class of all distributions with finite second moment; for the quantiles, it is the class of all continuous distributions with finite first moment; for regression, $X = (Y, Z), x = (y, z) \in \mathbb{R}^{k+1}$, and \mathcal{M} contains all distributions on \mathbb{R}^{k+1} with finite expected squared norm. In the identifiable case, for the mean, \mathcal{M} is the class of all distributions with finite first moment; for the quantiles, it is the class of all distributions with continuous distribution function at the α -quantile; for regression, $X = (Y, Z), x = (y, z) \in \mathbb{R}^{k+1}$, and \mathcal{M} contains all distributions on \mathbb{R}^{k+1} with finite mean.

3.2. Elicitable and identifiable hypotheses

We now use the concepts of elicibility and identifiability to construct null hypotheses for the sequential testing problem. Let T be either an elicitable functional or an identifiable functional with scoring function $s(\lambda, x)$ (in the elicitable case) or identification function $m(\lambda, x)$ (in the identifiable case). Given a fixed value $\lambda_0 \in \Lambda$ we consider the null hypothesis

$$\mathcal{H}_0 = \left\{ P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : \lambda_0 \in T(P(X_t \in \cdot \mid \mathcal{F}_{t-1})) \text{ for all } t \in \mathbb{N}, P\text{-a.s.} \right\}. \tag{8}$$

Thus the hypothesis is that T returns a set containing λ_0 whenever it is applied to the conditional distribution of an observation X_t given all earlier observations. For instance, if T is the median functional, we recover the example at the beginning of this section.

Using the definition of elicibility or identifiability, we obtain a simpler representation of \mathcal{H}_0 in terms of supermartingales or martingales and the scoring or identification function. Specifically, we have

$$\mathcal{H}_0 = \left\{ P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : \sum_{i=1}^t (s(\lambda_0, X_i) - s(\lambda, X_i)) \text{ is a } P\text{-supermartingale for all } \lambda \in \Lambda \right\} \tag{9}$$

in the elicitable case, and

$$\mathcal{H}_0 = \left\{ P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : \sum_{i=1}^t m(\lambda_0, X_i) \text{ is a } P\text{-martingale} \right\} \tag{10}$$

in the identifiable case. Note that in the identifiable case, $\sum_{i=1}^t m(\lambda_0, x_i)$ is a *vector valued* martingale, that is, a vector valued process such that each component is a martingale. For convenience, if \mathcal{H}_0 is of the form (9) we call it an *elicitable hypothesis*, and if it is of the form (10) we call it an *identifiable hypothesis*. The focus of our paper is to derive sequential tests for hypotheses of the form (9) or (10). An interesting related question is to characterize the properties of the stochastic processes with distributions in \mathcal{H}_0 . This question has been studied by [Frongillo and Nobel \(2020\)](#).

Let us spell out why, in the elicitable case, (8) essentially coincides with (9); the identifiable case is similar. Due to the definition of elicibility, $\lambda_0 \in T(P(X_t \in \cdot \mid \mathcal{F}_{t-1}))$ is equivalent to having $\mathbb{E}_P[s(\lambda_0, X_t) \mid \mathcal{F}_{t-1}] \leq \mathbb{E}_P[s(\lambda, X_t) \mid \mathcal{F}_{t-1}]$ for all $\lambda \in \Lambda$. This holds for all $t \in \mathbb{N}$ if and only if the process $\sum_{i=1}^t (s(\lambda_0, X_i) - s(\lambda, X_i))$ is a P -supermartingale for all $\lambda \in \Lambda$. Thus the right-hand sides of (8)

and (9) are essentially the same. There is one subtlety that we have neglected in this argument which is usually irrelevant in applications. Since strictly consistent scoring functions are not unique, it may happen that the elicitable functional is defined for a larger class of distributions than the one where the chosen strictly consistent scoring function in (9) has finite expectation. This means that the moment conditions on the conditional distributions in (9) may be slightly stronger than in (8). However, for many examples including the ones in Table 1, this problem does not arise since we work with score differences.

The scoring function which elicits a functional T is usually not unique. For example, the class of consistent scoring functions for the mean consists of all Bregman loss functions (Frongillo and Kash, 2015, Savage, 1971). Although there are no general guidelines for how one should select a scoring function, it is often natural to give preference to scoring functions that satisfy certain additional desirable properties, a relevant example in our setting being convexity in the first argument. For the mean, ratios of expectations and quantiles, convex strictly consistent scoring functions are essentially unique, see Fissler (2017, Corollary 4.2.17), Caponnetto (2005) and Steinwart et al. (2014, Corollary 10). In the context of estimation in semi-parametric models for a quantile or the mean, Dimitriadis, Fissler and Ziegel (2020), Komunjer and Vuong (2010a,b) show that there exist unique choices of scoring functions which maximize the asymptotic efficiency of the estimators, but these are different from the convex choices described above.

The form of \mathcal{H}_0 in (9) and (10) are suggestive of how one could construct families $\{L^\theta\}_{\theta \in \Theta}$ of base test supermartingales. If $s(\lambda, x)$ or $m(\lambda, x)$ is uniformly bounded, this is straightforward, see Subsection 3.3. The unbounded case requires to include additional moment bounds, which we achieve by imposing a sub- ψ condition, see Subsection 3.4.

3.3. Uniformly bounded scoring and identification functions

We construct parametric families $\{L^\theta\}_{\theta \in \Theta}$ of test martingales when the score difference, $(\lambda, x) \mapsto s(\lambda_0, x) - s(\lambda, x)$, or the norm of the identification function, $x \mapsto m(\lambda_0, x)$, are uniformly bounded.

Lemma 3.1 (Test martingales for elicitable hypotheses). *Consider an elicitable hypothesis \mathcal{H}_0 of the form (9), and assume that $\inf_{\lambda \in \Lambda, x \in \mathcal{X}} \{s(\lambda_0, x) - s(\lambda, x)\} > -1$. For each $\lambda \in \Lambda$, define the processes $L^\lambda = (L_t^\lambda)_{t \in \mathbb{N}}$ by*

$$L_t^\lambda = \prod_{i=1}^t (1 + s(\lambda_0, X_i) - s(\lambda, X_i)) .$$

Then the collection of processes $\{L^\lambda\}_{\lambda \in \Lambda}$ forms a family of \mathcal{H}_0 test supermartingales.

The lower bound of -1 appearing in the assumption that $\inf_{\lambda \in \Lambda, x \in \mathcal{X}} \{s(\lambda_0, x) - s(\lambda, x)\} > -1$ is without loss of generality since scoring functions may be rescaled by positive constants leaving the elicitable functional itself unchanged. In a similar fashion, we may construct test martingales for identifiable hypotheses as follows.

Lemma 3.2 (Test martingales for identifiable hypotheses). *Consider an identifiable hypothesis \mathcal{H}_0 of the form (10), assume that $\sup_{x \in \mathcal{X}} \|m(\lambda_0, x)\| < \infty$ and define*

$$A_{m, \lambda_0}^{\text{bd}} := \{\eta \in \mathbb{R}^d : \inf_{x \in \mathcal{X}} \langle \eta, m(\lambda_0, x) \rangle > -1\} .$$

For each $\eta \in A_{m,\lambda_0}^{\text{bd}}$, define the process $L^\eta = (L_t^\eta)_{t \in \mathbb{N}}$ by

$$L_t^\eta = \prod_{i=1}^t (1 + \langle \eta, m(\lambda_0, X_i) \rangle) .$$

Then $A_{m,\lambda_0}^{\text{bd}}$ is convex with non-empty interior, and the set $\{L^\eta\}_{\eta \in A_{m,\lambda_0}^{\text{bd}}}$ forms a family of \mathcal{H}_0 test martingales.

Remark 3.3. Whenever the scoring function $s(\lambda, x)$ is convex in λ and satisfies the conditions of Lemma 3.1, there is a direct connection between the two martingale constructions presented above. Indeed, if we let $m(\lambda_0, x) \in \partial_\lambda s(\lambda_0, x)$, then

$$-1 < s(\lambda_0, x) - s(\lambda, x) \leq \langle \lambda_0 - \lambda, m(\lambda_0, x) \rangle \tag{11}$$

for all x and λ . Hence, each increment of the test martingale construction of Lemma 3.2 can be thought of as a linearization of the increments of the processes defined in Lemma 3.1. Moreover, equation (11) implies that for any elicitable hypotheses with convex scoring function, identifiable test martingales generated by Lemma 3.2 will dominate the elicitable test martingale Lemma 3.1 whenever $\eta = \lambda - \lambda_0$. Indeed, it is easy to verify that

$$0 < \prod_{i=1}^t (1 + s(\lambda_0, x_i) - s(\lambda, x_i)) \leq \prod_{i=1}^t (1 + \langle \lambda_0 - \lambda, m(\lambda_0, x_i) \rangle) ,$$

and that the right-hand side produces a valid test supermartingale for all $\lambda, \lambda_0 \in \Lambda$. This observation suggests that whenever an elicitable functional T admits a bounded and convex scoring function, the test generated by its subgradient using Lemma 3.2 will always be more powerful than the one generated by Lemma 3.1.

Fissler (2017, Proposition 3.2.1) shows that, under suitable conditions, identification functions are unique up to multiplication with a matrix valued function in λ . Therefore, Remark 3.3 does not only apply to a subgradient of a convex scoring function but to any identification function, as long as a convex scoring function for the respective functional exists, and with a suitable modification of the relation $\eta = \lambda_0 - \lambda$.

In the setting of analyzing the asymptotic efficiency of semi-parametric estimators of elicitable and identifiable functionals, a similar relation is observed, in which an estimator generated by the identification function will always be asymptotically more efficient than its elicitable counterpart (Dimitriadis, Fissler and Ziegel, 2020).

The uniform boundedness assumptions in Lemmas 3.1 and 3.2 may appear to be restrictive. However, they cover a number of cases of interest including the mean whenever the data generating process $\{X_i\}_{i \in \mathbb{N}}$ is bounded, see also Waudby-Smith and Ramdas (2020). A second relevant example are (vectors of) quantiles, where the uniform boundedness assumption for the identification function is met, regardless of whether the data generating process is bounded or not. Indeed, it is easy to see from the rightmost column of Table 1 that $\|m(\lambda, x)\| \leq \max\{\alpha, 1 - \alpha\}$ is uniformly bounded. Hence the family of test martingales of Lemma 3.2 is always valid in the case of testing quantiles.

3.4. Test supermartingales for sub- ψ hypotheses

In the more general case that the scoring function or identification function is unbounded, we construct families of test martingales under the assumption of a tail bound on the scoring or identification function which involves bounding the cumulant generating function. We introduce the definition of a sub- ψ process below, a notion related to those introduced in [de la Peña, Klass and Lai \(2004\)](#), [Freedman \(1975\)](#) but most closely related to [Howard et al. \(2020, Definition 1\)](#).

Definition 3.4 (Sub- ψ Process). Let $Y = (Y_t)_{t \in \mathbb{N}}$ and $V = (V_t)_{t \in \mathbb{N}}$ be \mathcal{F} -adapted processes, where the variance process V_t is assumed to be non-negative and $Y_0 = V_0 = 0$. We say that (Y, V) is sub- ψ if there is a $u_{\max} > 0$ and a nonnegative convex function $\psi : [0, u_{\max}) \rightarrow [0, \infty)$ satisfying $\psi(0) = \psi'(0) = 0$, where $\psi'(0)$ is its right derivative, and for each $u \in [0, u_{\max})$,

$$(e^{u(Y_t - \bar{Y}_t) - V_t \psi(u)})_{t \in \mathbb{N}}$$

is a supermartingale, where $\bar{Y}_t = \sum_{i=1}^t \mathbb{E}[Y_i - Y_{i-1} \mid \mathcal{F}_{i-1}]$.

Definition 3.4 is similar to [Howard et al. \(2020, Definition 1\)](#), but there are some noteworthy differences. In particular, [Howard et al. \(2020, Definition 1\)](#) is a weaker condition in that it allows $e^{u(Y_t - \bar{Y}_t) - V_t \psi(u)}$ to only be upper-bounded by a supermartingale, rather than be a supermartingale itself. We make the choice of requiring the supermartingale condition in order to be able work with the supermartingale predictable mixing introduced in Section 2.3, which would break down without this assumption. Further discussion of the sub- ψ condition and its applications in time-uniform confidence bounds can be found in [Howard et al. \(2020, 2021\)](#). In particular, we point the reader to [Howard et al. \(2020, Section 3.1, Table 3\)](#) for a collection of commonly used ψ functions and variance processes V_t which are valid under a wide variety of assumptions.

Typically, $V_t = t$ is the simplest possible choice of variance process, and we use it for all concrete examples in this paper. We have chosen to state our theoretical results for the more general Definition 3.4 in order to be consistent with the existing literature. When $V_t = t$, the sub- ψ condition specializes to (conditional, one-sided versions of) sub-Gaussian, sub-Gamma, sub-Exponential, sub-Bernoulli and related conditions on the increments $\Delta Y_t = Y_{t+1} - Y_t$ of Y , obtained by choosing ψ to be the corresponding cumulant generating function. Specifically, the condition in Definition 3.4 is then equivalent to

$$\mathbb{E} \left[e^{u(\Delta Y_t - \mathbb{E}[\Delta Y_t \mid \mathcal{F}_t])} \mid \mathcal{F}_{t-1} \right] \leq e^{\psi(u)} \text{ for all } u \in [0, u_{\max}), t \in \mathbb{N}.$$

This condition implies a bound on the right tail probabilities of the increments of a sub- ψ process Y . Indeed, Chernoff's inequality (see e.g. [Hagerup and Rüb, 1990](#)) states that

$$\log P(\Delta Y_t > c + \mathbb{E}[\Delta Y_t \mid \mathcal{F}_t] \mid \mathcal{F}_t) \leq -\psi^*(c) \text{ for all } t \in \mathbb{N},$$

where $\psi^*(c) = \sup\{uc - \psi(u) : u \in [0, u_{\max})\}$ is the convex conjugate of ψ .

The following lemma shows that, for a given sub- ψ process $Y = (Y_t)_{t \in \mathbb{N}}$, the supermartingale property on Y is equivalent to the existence of a non-negative supermartingale.

Lemma 3.5. *Suppose that $(Y, V) = (Y_t, V_t)_{t \in \mathbb{N}}$ is an \mathcal{F} -adapted sub- ψ process. Then Y is a supermartingale if and only if $(e^{uY_t - V_t \psi(u)})_{t \in \mathbb{N}}$ is a supermartingale for all $u \in [0, u_{\max})$.*

We say that a family $\{Y^\theta\}_{\theta \in \Theta}$ of integrable processes indexed by Θ is sub- ψ if for each $\theta \in \Theta$ there is a function $\psi_\theta : [0, u_{\max}) \rightarrow [0, \infty)$ and a process V^θ such that for each $\theta \in \Theta$, Y^θ is sub- ψ_θ . We note

that although $\theta \mapsto (\psi_\theta(\cdot), V^\theta)$ varies with θ , the interval $[0, u_{\max})$ is assumed to be the same for all $\theta \in \Theta$. Using Lemma 3.5, we construct families of test supermartingales under the assumption that the scoring or identification functions satisfy a sub- ψ condition, allowing us to extend the sequential testing methodology to unbounded data.

Lemma 3.6 (Test supermartingales for sub- ψ elicitable hypotheses). *Let T be an elicitable functional with scoring function $s(\lambda, x)$ and let \mathcal{H}_0 be an elicitable hypothesis of the form (9). For every $\lambda \in \Lambda$ define $Y_t^\lambda = \sum_{i=1}^t s(\lambda_0, X_i) - s(\lambda, X_i)$ and $V_t^\lambda = \sum_{i=1}^t v_i^\lambda$ for some nonnegative \mathcal{F} -adapted process $(v_i^\lambda)_{i \in \mathbb{N}}$. If the family $\{(Y^\lambda, V^\lambda)\}_{\lambda \in \Lambda}$ is sub- ψ under every measure in \mathcal{H}_0 , then for each $\lambda \in \Lambda$ and $u \in [0, u_{\max})$, the process $L^{\lambda, u} = (L_t^{\lambda, u})_{t \in \mathbb{N}}$ defined by*

$$L_t^{\lambda, u} = e^{uY_t^\lambda - V_t^\lambda \psi_\lambda(u)} = \prod_{i=1}^t e^{u(s(\lambda_0, X_i) - s(\lambda, X_i)) - v_i^\lambda \psi_\lambda(u)}$$

is an \mathcal{H}_0 test supermartingale.

Lemma 3.7 (Test martingales for sub- ψ identifiable hypotheses). *Let T be an identifiable functional with identification function $(\lambda, x) \mapsto m(\lambda, x) \in \mathbb{R}^k$, and let \mathcal{H}_0 be an identifiable hypothesis of the form (10). Let $A_{m, \lambda_0}^\psi \subseteq \mathbb{R}^k$, and define for each $\eta \in A_{m, \lambda_0}^\psi$ the processes $Y_t^\eta = \sum_{i=1}^t \langle \eta, m(\lambda_0, X_i) \rangle$ and $V_t^\eta = \sum_{i=1}^t v_i^\eta$ for some nonnegative \mathcal{F} -adapted process $(v_i^\eta)_{i \in \mathbb{N}}$. If the family of processes $\{(Y^\eta, V^\eta)\}_{\eta \in A_{m, \lambda_0}^\psi}$ is sub- ψ under every measure in \mathcal{H}_0 , then for each $\eta \in A_{m, \lambda_0}^\psi$ and $u \in [0, u_{\max})$ the process $L^{\eta, u} = (L_t^{\eta, u})_{t \in \mathbb{N}}$ defined by*

$$L_t^{\eta, u} = \prod_{i=1}^t e^{u \langle \eta, m(\lambda_0, X_i) \rangle - v_i^\eta \psi_\eta(u)}$$

is an \mathcal{H}_0 test supermartingale.

The following examples illustrate two situations where test supermartingales can be constructed for unbounded data using sub- ψ assumptions.

Example 2 (Sub- ψ mean). Recall from Table 1 that the mean is identifiable with identification function $m(\lambda, x) = x - \lambda$. Now suppose that a real-valued data-generating process $(X_t)_{t \in \mathbb{N}}$ has conditionally sub-Gaussian increments so that $\log \mathbb{E}[e^{u(X_t - \mathbb{E}[X_t | \mathcal{F}_{t-1}])} | \mathcal{F}_{t-1}] \leq \tilde{\psi}(u)$ for $\tilde{\psi}(u) = \sigma^2 u^2 / 2$ and some $\sigma > 0$. Under this assumption, we have that $\sum_{i=1}^t \langle z, X_i - \lambda_0 \rangle$ is sub- ψ with $\psi_z(u) = z^2 \tilde{\psi}(u)$ and $V_t = t$, allowing us to apply Lemma 3.7.

Example 3 (Sub- ψ regression). Consider the hypothesis that the data follows an AR(k) linear time series model, $X_t = \sum_{i=1}^k \beta_i X_{t-i} + \epsilon_t$, where $X_t, \epsilon_t \in \mathbb{R}$, $(\epsilon_t)_{t \in \mathbb{N}}$ is a martingale difference sequence where $(\pm \sum_{i=1}^t \epsilon_i)_{t \in \mathbb{N}}$ is sub- ψ with variance process $V_t = t$ and $\beta \in \mathbb{R}^k$ is unknown. We wish to test whether $\beta = \beta_0 \in \mathbb{R}^k$. In each time step, β is the value of the identifiable functional $T(P) = \operatorname{argmin}_\beta \mathbb{E}_P[\|\sum_{i=1}^k \beta_i X_{t-i} - X_t\|^2 | \mathcal{F}_{t-1}]$. In view of the regression example of Table 1, this functional has identification function

$$m(\beta, (y, \mathbf{x})) = \left(\left(\sum_{j=1}^k \beta_j x_j - y \right) \frac{x_i}{\|\mathbf{x}\|} \right)_{i=1}^k,$$

where $\mathbf{x} = (x_i)_{i=1}^k$ and where the re-scaling by $1/\|\mathbf{x}\|$ is possible because $\mathbf{X}_{(t-1):(t-k)} = (X_{t-i})_{i=1}^k$ is \mathcal{F}_{t-1} -measurable. In order to apply the testing methodology of Lemma 3.7, we show that the processes Y_t^z with $Y_t^z - Y_{t-1}^z = \langle z, m(\beta_0, (X_t, \mathbf{X}_{(t-1):(t-k)})) \rangle$ are sub- ψ_z with $V_t^z = t$ for all $z \in \mathbb{R}^d$ with $\|z\| \leq 1$. Computing

$$\langle z, m(\beta_0, (X_t, \mathbf{X}_{(t-1):(t-k)})) \rangle - \mathbb{E} \left[\langle z, m(\beta_0, (X_t, \mathbf{X}_{(t-1):(t-k)})) \rangle \middle| \mathcal{F}_{t-1} \right] = c(z) \epsilon_t,$$

where $c(z) = \langle z, \mathbf{X}_{(t-1):(t-k)} \rangle / \|\mathbf{X}_{(t-1):(t-k)}\|$, we see that if $\pm \sum_{i=1}^t \epsilon_i$ are sub- ψ with $V_t = t$, then Y_t^z will also be sub- ψ since $\log \mathbb{E}_P \left[\exp(c(z)u \epsilon_t) \right] \leq \psi(|c(z)|u) \leq \psi(u)$, which follows due to the fact that $|c(z)| \leq 1$. Hence, a sufficient condition for this identifiable functional to satisfy the necessary sub- ψ condition is simply that the signed residual processes $(\pm \sum_{i=1}^t \epsilon_i)_{t \in \mathbb{N}}$ are sub- ψ .

There is a relationship between tests for identifiable and elicitable hypotheses with a convex scoring function in the sub- ψ case, analogous to the one pointed out in Remark 3.3. Indeed, let $T(\mu)$ be an elicitable functional with a convex scoring function $(\lambda, x) \mapsto s(\lambda, x)$. As pointed out in Remark 3.3, $T(\mu)$ is also identifiable with identification function $m(\lambda, x) \in \partial_\lambda s(\lambda, x)$. If we assume in addition that for a fixed $\lambda_0 \in \Lambda$, the family of processes $\{(Y_t^\lambda)_{t \geq 0}\}_{\lambda \in \Lambda}$ with increments

$$Y_t^\lambda - Y_{t-1}^\lambda = \langle \lambda_0 - \lambda, m(\lambda_0, X_t) \rangle$$

is sub- ψ with $V_t = t$, then for any $\lambda \in \Lambda$ and $u \in [0, u_{\max})$, the process

$$L_t^{\lambda, u} = \prod_{i=1}^t e^{u \langle \lambda_0 - \lambda, m(\lambda_0, X_i) \rangle - \psi_\lambda(u)},$$

is a valid test supermartingale according to Lemma 3.7. However, since $s(\lambda, x)$ is convex and the $\langle \lambda_0 - \lambda, m(\lambda_0, X_i) \rangle$ are increments of the sub- ψ process Y_t^λ , we have that under \mathcal{H}_0 ,

$$\mathbb{E} \left[e^{u(s(\lambda_0, X_i) - s(\lambda, X_i))} \middle| \mathcal{F}_{i-1} \right] \leq \mathbb{E} \left[e^{u \langle \lambda_0 - \lambda, m(\lambda_0, X_i) \rangle} \middle| \mathcal{F}_{i-1} \right] \leq e^{\psi(u)}.$$

Hence, the processes

$$\tilde{L}_t^{\lambda, u} = \prod_{i=1}^t e^{u(s(\lambda_0, X_i) - s(\lambda, X_i)) - \psi_\lambda(u)}$$

are valid \mathcal{H}_0 test supermartingales which match the form of the test supermartingales presented in Lemma 3.6. Due to (3.4), however, we note that $\tilde{L}_t^{\lambda, u} \leq L_t^{\lambda, u}$. Hence, whenever $\lambda \mapsto s(\lambda, x)$ is convex and the families of processes $\{(s(\lambda_0, X_i) - s(\lambda, X_i))_{i \in \mathbb{N}}\}_{\lambda \in \Lambda}$ and $\{(\langle \lambda_0 - \lambda, m(\lambda_0, X_i) \rangle)_{i \in \mathbb{N}}\}_{\lambda \in \Lambda}$ are both increments of a sub- ψ process, we find that the tests generated by the identification function m according to Lemma 3.7 will always be more powerful than a test generated by the scoring function s according to Lemma 3.6, yielding a conclusion analogous to that in Remark 3.3.

Remark 3.8 (Bridging the sub- ψ and bounded test supermartingales). Although the sub- ψ and uniformly bounded hypothesis testing methodologies may appear disjoint, there are in fact some connections which are worth highlighting.

The first and arguably most important remark is that all processes with bounded increments are sub-Gaussian, and hence sub- ψ . Indeed, whenever a process $(Y_t)_{t \in \mathbb{N}}$ satisfies $\Delta Y_t = Y_{t+1} - Y_t \in [a, b]$, it follows by Hoeffding’s lemma (see e.g. Hertz (2020), Hoeffding (1994)) that

$$\log \mathbb{E} \left[e^{u(\Delta Y_t - \mathbb{E}[\Delta Y_t | \mathcal{F}_t])} \middle| \mathcal{F}_t \right] \leq \frac{1}{8} u^2 (b - a)^2, \tag{12}$$

where the second expression in the above inequality is the cumulant generating functions of a Gaussian random variable. Hence, we have that processes with bounded increments are sub- ψ , where ψ is given by the right-hand side in (12).

There exists a deeper connection between the two as follows. Let $Z = (Z_t)_{t \in \mathbb{N}}$ be a supermartingale difference process and let $Y_t = \sum_{i=1}^t Z_i$. Define for each $n \in \mathbb{N}$ and $u \geq 0$ the process $M_t^{u,n} = \prod_{i=1}^t (1 + uZ_i/n)^n e^{-\mu_n(u)}$, where the collection of compensators $\{\mu_n\}_{n \in \mathbb{N}}$ with $\mu_n : [0, \infty) \rightarrow [0, \infty]$ satisfy $\log \mathbb{E}[(1 + uZ_i/n)^n | \mathcal{F}_{i-1}] \leq \mu_n(u)$ for all $i \in \mathbb{N}$ and $u > 0$. Since the function $z \mapsto (1 + z/n)^n$ is monotone, nonnegative and convex function on $z \in [-n, \infty)$, we find that whenever $uZ_i \geq -n$ for all $i \in \mathbb{N}$, the process $L^{u,n}$ is a nonnegative supermartingale with initial value $L_0^{u,n} = 1$.

Whenever $n = 1$, we recover a set of processes with multiplicative increments that are linear in the Z_i , $M_t^{u,1} = \prod_{i=1}^t (1 + uZ_i) e^{-\mu_1(u)}$. Noting that if $uZ_i > -1$ and $\mathbb{E}[Z_i | \mathcal{F}_{i-1}] \leq 0$, we find that $\mu_1 \equiv 0$ is a valid compensator and $M^{u,1}$, defined in this manner, matches the structure of the uniformly bounded test supermartingales presented in Section 3.3. On the other hand, we note that as $n \rightarrow \infty$, $(1 + z/n)^n \rightarrow e^z$ pointwise, which yields the process $M_t^{u,\infty} = \prod_{i=1}^t e^{uZ_i - \mu_\infty(u)}$, revealing a similar structure to the test supermartingales discussed in Section 3.4. Indeed, whenever the process Y_i is sub- ψ and $u \geq 0$, we find that $L^{u,\infty}$ is a nonnegative supermartingale.

The families of test supermartingales presented in Lemmas 3.2 and 3.7 can be thought of as generalizations of the processes presented in [Waudby-Smith and Ramdas \(2020\)](#) for the purpose of building confidence sequences for means of bounded random variables. In particular, let us consider the functional $T(\mu) = \mathbb{E}_\mu[X]$ with identification function $m(\lambda, x) = \lambda - x$, where the random process $\{X_i\}_{i \in \mathbb{N}}$ is constrained to the interval $[0, 1]$. Applying Lemma 3.2 to this particular setting, we recover the capital process in [Waudby-Smith and Ramdas \(2020, §4\)](#). Similarly, as noted in Remark 3.8, the process $\sum_{i=1}^t \langle z, m(X_i, \lambda) \rangle$ is sub- ψ for all $\lambda \in [0, 1]$ with $\psi_\lambda(u) = \frac{1}{8}u^2$ and $V_t = t$, allowing us to recover via Lemma 3.7 the so-called Chernoff and predictably-mixed Chernoff martingales ([Waudby-Smith and Ramdas, 2020, §2.3 & §3.1](#)).

4. Power via online convex optimization

Consider a null hypothesis \mathcal{H}_0 of the form (5) given in terms of a family of base test supermartingales $\{L^\theta\}_{\theta \in \Theta}$. We assume that each L^θ is of product form

$$L_t^\theta = \prod_{i=1}^t f_\theta(X_i) \tag{13}$$

for some nonnegative function $f_\theta(x)$. This is the case for all the hypotheses considered in Section 3. We now address the problem of designing a powerful test supermartingale W . As a starting point we take the GRO criterion in (2) and search among processes obtained by predictable mixing as in (3). We must choose a distribution $\hat{Q} \in \mathcal{M}(X^{\mathbb{N}})$ to bet on. This distribution will not be fully specified ahead of time, but rather learnt adaptively. For $t \in \mathbb{N}$ we set $\hat{Q}(X_{t+1} \in \cdot | \mathcal{F}_t) = (1/t) \sum_{i=1}^t \delta_{X_i}$, the empirical measure of the data observed so far. Together with $\hat{Q}(X_1 \in \cdot)$, which we choose arbitrarily, this uniquely specifies the distribution \hat{Q} . The GRO problem (2) at time T , restricted to predictable mixtures W as in (3), now takes the form

$$\underset{\pi_{T+1}}{\text{maximize}} \sum_{i=1}^T \log \int_{\Theta} f_\theta(X_i) \pi_{T+1}(d\theta). \tag{14}$$

Let us give some further details for how to get from (2) to (14). First, the supermartingale constraint in (2) is automatically satisfied since W is a predictable mixture. Next, the form of W and L^θ imply that $W_{T+1}/W_T = \int_{\Theta} (L_{T+1}^\theta/L_T^\theta)\pi_{T+1}(d\theta) = \int_{\Theta} f_\theta(X_{T+1})\pi_{T+1}(d\theta)$. Finally, because $\hat{Q}(X_{T+1} \in \cdot | \mathcal{F}_T)$ is defined as the empirical measure of X_1, \dots, X_T it follows that the objective function in (2) is

$$\mathbb{E}_{\hat{Q}} \left[\log \frac{W_{T+1}}{W_T} \mid \mathcal{F}_T \right] = \mathbb{E}_{\hat{Q}} \left[\log \int_{\Theta} f_\theta(X_{T+1})\pi_{T+1}(d\theta) \mid \mathcal{F}_T \right] = \frac{1}{T} \sum_{i=1}^T \log \int_{\Theta} f_\theta(X_i)\pi_{T+1}(d\theta).$$

Since the factor $1/T$ does not affect the optimization, we arrive at (14).

Although the maximization problem (14) is concave in π_{T+1} , for practical reasons we wish to avoid optimizing over this potentially infinite dimensional quantity. Instead, as discussed in Subsection 2.3, we restrict the optimization to the much smaller set of ‘one-asset strategies’ of the form $\pi_{T+1} = \delta_{\theta_{T+1}}$ for some $\theta_{T+1} \in \Theta$ that may depend on X_1, \dots, X_T . Doing so simplifies (14) further to

$$\text{maximize}_{\theta_{T+1}} \sum_{i=1}^T \log f_{\theta_{T+1}}(X_i). \tag{15}$$

Due to the product form of L^θ the objective function is actually of the even simpler form $\log L_T^{\theta_{T+1}}$. Here, the dependence on the optimization variable θ_{T+1} is no longer concave in general. Nonetheless, the following lemma shows that in a wide range of examples concavity does, in fact, hold. This will allow us to apply results from Online Convex Optimization, either directly to (15) itself, or to regularized versions of it.

Lemma 4.1. *Let $\{L^\theta\}_{\theta \in \Theta}$ satisfy one of the following conditions.*

1. *There exists a convex set $\Theta \subseteq \Lambda$ such that $\{L^\theta\}_{\theta \in \Theta}$ are test supermartingales for a bounded elicitable hypothesis defined according to Lemma 3.1, where the scoring function $\Theta \ni \lambda \mapsto s(\lambda, x)$ is convex for all $x \in \mathcal{X}$.*
2. *There exists a convex set $\Theta \subseteq A_{m, \lambda_0}^{\text{bd}}$ such that $\{L^\theta\}_{\theta \in \Theta}$ are test martingales for a bounded identifiable hypothesis defined according to Lemma 3.2.*
3. *There exists a finite dimensional convex set Θ and a map*

$$\Theta \ni \theta \mapsto (\lambda(\theta), u(\theta)) \subseteq \Lambda \times [0, u_{\max}]$$

such that the collection $\{L^\theta\}_{\theta \in \Theta} = \{L^{\lambda(\theta), u(\theta)}\}_{\theta \in \Theta}$ are test supermartingales for a sub- ψ elicitable hypothesis defined according to Lemma 3.6 and

$$\Theta \ni \theta \mapsto u(\theta) (s(\lambda_0, x_i) - s(\lambda(\theta), x_i)) - v_i^{\lambda(\theta)} \psi_{\lambda(\theta)}(u(\theta))$$

is almost surely concave.

4. *There exists a finite-dimensional convex set Θ and a map*

$$\Theta \ni \theta \mapsto (\eta(\theta), u(\theta)) \subseteq A_{m, \lambda_0}^\psi \times [0, u_{\max}]$$

such that $\{L^\theta\}_{\theta \in \Theta} = \{L^{\eta(\theta), u(\theta)}\}_{\theta \in \Theta}$ are test supermartingales for a sub- ψ identifiable hypothesis defined according to Lemma 3.7, and

$$\Theta \mapsto u(\theta) \langle \eta(\theta), m(\lambda_0, x_i) \rangle - v_i^{\lambda(\theta)} \psi_{\eta(\theta)}(u(\theta))$$

is almost surely concave.

Then for each $t \in \mathbb{N}$, the map $\theta \mapsto \log L_t^\theta$ is concave a.s.

The proof of the lemma uses the fact that the composition of an increasing concave function with a concave function is concave.

4.1. Regret and asymptotic power

By repeatedly solving (15) (or a regularized version of it) in each time period, we obtain a predictable sequence $(\theta_t)_{t \in \mathbb{N}}$ which produces the test supermartingale

$$W_t = \prod_{i=1}^t f_{\theta_i}(X_i). \tag{16}$$

A standard way of measuring the quality of the sequence $(\theta_t)_{t \in \mathbb{N}}$ is the *regret*, defined at each time t by

$$\text{Regret}_t := \max_{\theta \in \Theta} \{ \log L_t^\theta - \log W_t \}. \tag{17}$$

The regret represents the difference between the log-value of the best retrospectively chosen single-asset portfolio, $\max_{\theta \in \Theta} \log L_t^\theta$, and the given log-wealth $\log W_t$. Various well-known algorithms for solving either (15) or regularized versions of it, yield regret that grows sublinearly,

$$\text{Regret}_T = o(T),$$

where we emphasize that this holds almost surely, that is, $\lim_{T \rightarrow \infty} \text{Regret}_T/T = 0$ a.s. We review some of these algorithms in Subsection 4.2. The following theorem shows that if regret grows sublinearly, then the test (1) constructed from the test supermartingale W has asymptotic power one.

Theorem 4.2 (Sublinear regret implies asymptotic power). *Let $W = (W_t)_{t \in \mathbb{N}}$ be a predictably mixed test supermartingale process defined according to (16), generated by a family of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$ and a sequence $(\theta_t)_{t \in \mathbb{N}}$ achieving $\text{Regret}_T = o(T)$. Let $\alpha \in (0, 1]$, and let τ_0 be the test induced by W , that is, $\tau_0 = \inf\{t \in \mathbb{N} : W_t > \alpha^{-1}\}$. Consider a probability measure $Q \notin \mathcal{H}_0$ such that there exists $\theta' \in \Theta$ for which*

$$Q \left(\liminf_{T \rightarrow \infty} \frac{\log L_T^{\theta'}}{T} > 0 \right) = 1. \tag{18}$$

Then

$$Q(\tau_0 \text{ rejects } \mathcal{H}_0) = Q(\tau_0 < \infty) = Q \left(\lim_{T \rightarrow \infty} \max_{0 \leq t \leq T} W_t > \alpha^{-1} \right) = 1,$$

that is, τ_0 eventually rejects \mathcal{H}_0 with probability one.

Each of the algorithms presented in Subsection 4.2 achieve sublinear regret, as required in the statement of Theorem 4.2. The second assumption in Theorem 4.2 that must be met to achieve asymptotic power one, is the existence of a θ' that satisfies (18). In the next lemma, we show that there are easily verifiable sufficient conditions to guarantee this condition when data is generated by a stationary ergodic process.

Proposition 4.3. *Suppose that for $Q \notin \mathcal{H}_0$, the data-generating process $(X_t)_{t \in \mathbb{N}}$ is stationary and ergodic with invariant measure Q_∞ . Let $\{L_t^\theta\}_{\theta \in \Theta}$ be a collection of \mathcal{H}_0 test supermartingales where one of the following conditions holds.*

1. \mathcal{H}_0 is a uniformly bounded elicitable hypothesis with $\{L_t^\theta\}_{\theta \in \Theta}$ defined according to Lemma 3.1 where $\Theta \subseteq \Lambda$ and there exists $\lambda' \in \Theta$ such that $\mathbb{E}_{Q_\infty}[\log(1 + s(\lambda_0, X_\infty) - s(\lambda', X_\infty))] > 0$.
2. \mathcal{H}_0 is a uniformly bounded identifiable hypothesis with $\{L_t^\theta\}_{\theta \in \Theta}$ defined according to Lemma 3.2 where $\Theta \subseteq A_{m, \lambda_0}^{\text{bd}}$ and there exists $\eta' \in \Theta$ such that $\mathbb{E}_{Q_\infty}[\log(1 + \langle \eta', m(\lambda_0, X_\infty) \rangle)] > 0$.
3. \mathcal{H}_0 is a sub- ψ elicitable hypothesis with $\{L_t^\theta\}_{\theta \in \Theta}$ defined according to Lemma 3.6 where $\Theta = \Theta' \times [0, \epsilon)$ in which $0 < \epsilon \leq u_{\max}$, $\Theta' \subseteq \Lambda$, $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t v_i^\lambda < \infty$ a.s. for all $\lambda \in \Lambda$, and there exists $\lambda' \in \Theta'$ such that $\mathbb{E}_{Q_\infty}[s(\lambda_0, X_\infty) - s(\lambda', X_\infty)] > 0$.
4. \mathcal{H}_0 is a sub- ψ identifiable hypothesis with $\{L_t^\theta\}_{\theta \in \Theta}$ defined according to Lemma 3.7 where $\Theta = \Theta' \times [0, \epsilon)$ in which $0 < \epsilon \leq u_{\max}$, $\Theta' \subseteq A_{m, \lambda_0}^\psi$, $\lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t v_i^\eta < \infty$ a.s. for all $\eta \in A_{m, \lambda_0}^\psi$, and there exists $\eta \in \Theta'$ such that $\mathbb{E}_{Q_\infty}[\langle \eta, m(\lambda_0, X_\infty) \rangle] > 0$.

Then there exists $\theta' \in \Theta$ for which (18) holds.

We remark here that two common cases in which Proposition 4.3 may apply are when $(X_t)_{t \in \mathbb{N}}$ is either i.i.d. or generated by a stationary, irreducible and aperiodic Markov chain. Moreover, conditions (iii) and (iv) of Proposition 4.3 hold for any $Q \notin \mathcal{H}_0$, which can be seen by inspecting (9) or (10).

To conclude, Theorem 4.2 can be interpreted as a high-level result stating that whenever there are alternative hypotheses which offer ‘sufficient evidence against the null’, regret-optimal strategies will asymptotically reject the null with probability one. Proposition 4.3 quantifies the notion of ‘sufficient evidence against the null’ in the case of a stationary ergodic data generating process.

Regret bounds have been used by various authors outside the context of elicitable and identifiable functionals as a way to ensure the power of statistical inference procedures based on betting. See for example Jun and Orabona (2019), Orabona and Jun (2021), Ramdas et al. (2022a), Shekhar and Ramdas (2021). The latter reference in particular contains a literature overview as well as links to online learning. Further historical remarks can be found in Appendix F of Waudby-Smith and Ramdas (2020).

4.2. Online convex optimization algorithms

We summarize a few simple but effective OCO algorithms that can be applied to optimizing the predictably mixed test supermartingale $V = (V_t)_{t \in \mathbb{N}}$ by selecting a sequence of single-asset portfolios given by $(\theta_t)_{t \in \mathbb{N}}$ which maximize the growth rate of (16). We point interested readers to Hazan (2016), Shalev-Shwartz (2011) for a comprehensive introduction to OCO.

For the remainder of the section, we assume that for all $t \in \mathbb{N}$, the maps of forward differences defined on a convex set Θ ,

$$\theta \mapsto \Delta \log L_{t-1}^\theta = \log L_t^\theta - \log L_{t-1}^\theta$$

are concave. We say that a convex function $f : \Theta \rightarrow \mathbb{R}$ is strongly convex with parameter $\mu > 0$ whenever $f(\theta_1) - f(\theta_0) - \langle v_0, \theta_1 - \theta_0 \rangle \geq \mu \|\theta_1 - \theta_0\|^2$ for all $\theta_0, \theta_1 \in \Theta$, where $v_0 \in \partial_\theta f(\theta_0)$ is an element the subgradient of f at θ_0 . Similarly, f is strongly concave whenever $-f$ is strongly convex. We define the norm of the subgradient of a convex function $f : \Theta \rightarrow \mathbb{R}$ as $\|\partial_\theta f(\theta)\| = \sup_{v \in \partial_\theta f(\theta)} \|v\|$.

Follow The Leader. We define the Follow The Leader (FTL) algorithm as choosing at each iteration $t + 1$

$$\theta_{t+1} \in \operatorname{argmin}_{\theta \in \Theta} -\log L_t^\theta = \operatorname{argmin}_{\theta \in \Theta} - \sum_{1 \leq i \leq t} \Delta \log L_{i-1}^\theta. \tag{FTL}$$

Hence, at each iteration, (FTL) picks $\theta \in \Theta$ such that $\log L_t^\theta$ has the largest average growth rate in hindsight. Implicitly, this algorithm assumes that computing the argmax at each round can be done

relatively easily. Under the additional assumptions that (i) $\sup_{\theta \in \Theta} \|\partial_{\theta}(-\Delta \log L_{t-1}^{\theta})\| \leq G < \infty$ and (ii) the $\Delta \log L_t^{\theta}$ are strongly concave, we have that $\text{Regret}_T \leq \frac{G^2}{2}(1 + \log T)$. We point the reader to McMahan (2017, §3.6) for a derivation of this bound. This algorithm is not recommended for the case when $\Delta \log L_t^{\theta}$ are not strongly concave. Indeed, it is possible to create counterexamples, such as those presented in Shalev-Shwartz (2011, Example 2.2) or Hazan (2016, Page 65), of non strongly-concave functions where the algorithm induces super-linear growth in the regret. The bounded gradient assumption (i) will be satisfied by, for example, Lipschitz continuous functions.

Follow The Regularized Leader. A simple fix for this problem leads to the second algorithm, which involves the inclusion of regularization, increasing the algorithm’s stability. The most obvious implementation of this concept is the Follow The Regularized Leader (FTRL) algorithm, in which we introduce regularization terms to the optimization problem in (FTL). Precisely, at each iteration $t + 1$, FTRL selects

$$\theta_{t+1} \in \underset{\theta \in \Theta}{\operatorname{argmin}} \left\{ -\log L_t^{\theta} + \sum_{1 \leq i \leq t} r_{i-1}(\theta) \right\}, \tag{FTRL}$$

where the $\{r_t\}_{t \in \mathbb{N}}$ is a sequence of strongly convex functions for which either (a) $x_0 = \operatorname{argmin}_{\theta \in \Theta} r_t(\theta)$ for all t , in which the algorithm is the *FTRL-Centered* variant or (b) for which we assume that $x_t = \operatorname{argmin}_{\theta \in \Theta} r_t(\theta)$ where this variant is named *FTRL-Proximal*. Under the additional assumptions that (i) $\sup_{\theta \in \Theta} \|\partial_{\theta}(-\Delta \log L_{t-1}^{\theta})\| \leq G < \infty$ and that (ii) $\operatorname{diam}(\Theta) = \sup_{x,y \in \Theta} \|x - y\| < \infty$, one can devise a sequence of centered or proximal regularizers $\{r_t\}_{t \in \mathbb{N}}$ such that the algorithms enjoy the regret bound $\text{Regret}_T = O(\sqrt{T})$. We point the reader to McMahan (2017, Section 3) for a summary of the exact conditions and specific bounds.

Online Gradient Descent. The argmin expressions in (FTL) and (FTRL) may be difficult to compute directly introducing problems in the implementation of FTL or FTRL. Instead, using gradients collected in each step, one can follow the direction of steepest ascent of $\Delta \log L_{t-1}^{\theta}$ at each iteration. This produces the Online Gradient Descent (OGD) algorithm, where we choose at each iteration

$$\theta_{t+1} = \Pi_{\Theta} \{ \theta_t - \eta_t v_t \} \text{ where } v_t \in \partial_{\theta} \left(-\Delta \log L_{t-1}^{\theta} \right), \tag{OGD}$$

where $\{\eta_t\}_{t \in \mathbb{N}}$ is a sequence of positive learning rates and $\Pi_{\Theta}(x) = \operatorname{argmin}_{y \in \Theta} \|y - x\|$ is the projection operator. This algorithm has the advantage that it is extremely simple to compute, provided that gradients are available at each step. For an appropriate choice of $\{\eta_t\}_{t \in \mathbb{N}}$, and under the additional assumption that (i) $\sup_{\theta \in \Theta} \|\partial_{\theta}(-\Delta \log (L_{t-1}^{\theta}))\| \leq G < \infty$ and (ii) $\operatorname{diam}(\Theta) < \infty$, the algorithm has a regret bound of $\text{Regret}_T = O(\sqrt{T})$. We point the reader to Zinkevich (2003) and McMahan (2017, §3.1) for the specific bounds and conditions.

We summarize the bounds and assumptions for the OCO algorithms in Table 2, where we emphasize that the growth rate of the regret in the right-hand column holds almost surely regardless of the data-generating measure. Beyond the algorithms presented here, there exist a plethora of online optimization algorithms which may leverage the geometry of the index set Θ or past information about gradients in order to improve rates of convergence. For a broader survey of available OCO algorithms, we point the reader to Hazan (2016), McMahan (2017), Shalev-Shwartz (2011). We also point out that although the algorithms here offer guarantees for convex loss functions, they can in principle be used on non-convex optimization problems as well. However, while still yielding valid tests, the worst-case guarantees may no longer hold.

Assumptions			
Algorithm	$-\Delta \log W_t(\theta)$	$\text{diam}(\Theta)$	Regret_T
FTL	strongly-convex + bounded gradients	any	$O(\log T)$
FTRL	bounded gradients	$\leq R$	$O(\sqrt{T})$
OGD	bounded gradients	$\leq R$	$O(\sqrt{T})$

Table 2. Summary of assumptions and regret bounds for the OCO algorithms presented in Section 4.2.

The use of OCO methods for building more powerful predictably mixed test supermartingales is related to various online methods presented in [Waudby-Smith and Ramdas \(2020\)](#) for confidence sequence building for means of bounded random variables. In particular, Kelly betting ([Waudby-Smith and Ramdas, 2020, §5.2,5.3,5.6](#)) can be interpreted as a variation of the **FTL** algorithm to a family of processes of the form of Lemma 3.2 for bounded identifiable functionals. Similarly, the Online Newton Step Algorithm discussed in [Waudby-Smith and Ramdas \(2020, §5.5\)](#) can be thought of another method in the family of OCO algorithms. The authors of [Waudby-Smith and Ramdas \(2020\)](#) acknowledge and elaborate on the connection between their confidence sequence building methods with OCO and coin-betting algorithms in [Waudby-Smith and Ramdas \(2020, Appendix D\)](#) but they do not use regret bounds to derive asymptotic power guarantees.

Although largely ignored over the course of this section, another important aspect to consider when selecting an algorithm for online convex optimization is the computational complexity of the algorithm. For example, it is typically the case that (**FTL**), (**FTRL**) or other loss-minimization based algorithms will be run slower than gradient-based algorithms such as (**OGD**). On the other hand, although (**OGD**) and (**FTRL**) have similar asymptotic performance, we typically find that (**FTRL**) may perform better on average, which can be reflected in the constants associated with their regret bounds, see for example [McMahan \(2017\)](#).

Although the OCO algorithms presented in Section 4.2 offer regret guarantees which translate into asymptotic power, it is possible to generalize these methods in order to further improve performance in special cases. The algorithms **FTL**, **FTRL** and **OGD**, attempt to maximize some version of the data-generated objective function

$$\theta \mapsto \frac{1}{t} \log L_t^\theta = \frac{1}{t} \sum_{i=1}^t \Delta \log L^\theta(X_i) = \mathbb{E}_{\hat{P}_t} [\Delta \log L^\theta(X)] , \tag{19}$$

where we write $\Delta \log L_i^\theta = \Delta \log L^\theta(X_i)$ as an explicit function of the data point X_i , and where $\mathbb{E}_{\hat{P}_t}$ represents the expected value with respect to the empirical measure $\hat{P}_t = (1/t) \sum_{i=1}^t \delta_{X_i}$.

If additional distributional information is known about $\{X_t\}_{t \in \mathbb{N}}$, one could replace the empirical distribution \hat{P}_t with a predictive measure \tilde{P}_t which may better represent the true data distribution. Modifying (**FTL**) with this new measure, one obtains the algorithm

$$\theta_{t+1} \in \underset{\theta \in \Theta}{\text{argmin}} \mathbb{E}_{\tilde{P}_t} [\log \Delta L^\theta(X)] , \tag{FTLP}$$

and the same principle can be applied to modify both, (**FTRL**) and (**OGD**). The fact that we have complete freedom in choosing \tilde{P}_t at each step can be a major advantage of the betting approach to sequential testing, as has been discussed in detail elsewhere; see [Ramdas et al. \(2022a\)](#), [Waudby-Smith and Ramdas \(2020\)](#).

5. Confidence sequences and inverting tests

Our focus has been on methodology for testing elicitable and identifiable hypotheses in a sequential setting. However, the techniques we have developed allow for the construction of confidence sequences. Given a functional $T(P) : \mathcal{M}_1(X) \rightarrow \Lambda \subseteq \mathbb{R}^d$ define a *confidence sequence* $\{C_t\}_{t \in \mathbb{N}}$ at level $\alpha \in (0, 1)$ as a sequence $(C_t)_{t \in \mathbb{N}}$ of confidence sets $C_t \subseteq \Lambda$ for all $t \in \mathbb{N}$, satisfying the property that

$$\sup_{P \in \mathcal{H}_0(\lambda_0)} P(\forall t \in \mathbb{N} : \lambda_0 \notin C_t) \leq \alpha, \tag{20}$$

for any $\lambda_0 \in \Lambda$. Here, $\mathcal{H}_0(\lambda_0)$ denotes the null hypothesis at (8) where we now emphasize the dependence on λ_0 in the notation. Equation (20) can be interpreted as guaranteeing that, with high probability, the true value of the functional, $\lambda_0 \in T(P)$ is contained within, not only a single confidence set, but the entire sequence of confidence sets at once.

Using the sequential hypothesis testing methodology, we can construct such confidence sequences for the elicitable and identifiable martingale hypotheses of the form (9) and (10). Assume that for each $\lambda_0 \in \Lambda$ there exists an $\mathcal{H}_0(\lambda_0)$ test supermartingale denoted by $W^{\lambda_0} = (W_t^{\lambda_0})_{t \in \mathbb{N}}$. Each W^{λ_0} induces a sequential test for $\mathcal{H}_0(\lambda_0)$, so we may construct a confidence sequence $C = (C_t)_{t \in \mathbb{N}}$ by ‘inverting’ these tests as follows. For each $t \in \mathbb{N}$, we define

$$C_t = \left\{ \lambda \in \Lambda : \max_{0 \leq i \leq t} V_i^\lambda \leq \alpha^{-1} \right\}, \tag{21}$$

the set of λ whose associated null hypotheses have not yet been rejected by the tests induced by the associated W^{λ_0} . This construction produces a confidence sequence since

$$\sup_{P \in \mathcal{H}_0(\lambda_1)} P(\exists t \text{ s.t.} : \lambda_1 \notin C_t) = \sup_{P \in \mathcal{H}_0(\lambda_1)} P(\max_{t \in \mathbb{N}} W_t^{\lambda_1} > \alpha^{-1}) \leq \alpha.$$

We may apply the techniques from previous sections to construct confidence sequences for elicitable and identifiable functionals as follows. Depending on whether the functional T in question is elicitable or identifiable and on whether it is *uniformly bounded* or *sub-ψ*, for each $\lambda_0 \in \Lambda_0$ we construct a family of test supermartingales according to Lemmas 3.1, 3.2, 3.6 or 3.7, which we use to construct a mixture test supermartingale W^{λ_0} according to Lemma 2.4. In Appendix A, we show how this methodology can be extended for the purpose of testing one-sided and set-valued hypotheses.

6. Numerical examples

We begin by applying our tests to sequences of independent and identically distributed data. Let $\{X_t\}_{t \in \mathbb{N}} \sim$ i.i.d. Beta(α, β), where $\alpha = 2$ and $\beta = 5$ are the parameters of the beta distribution.

Our first experiment will be to test the mean and standard deviation of this data generating process, simultaneously. That is, we consider a functional $T : \mathcal{M}_1(X) \rightarrow \mathbb{R}^2$ where $T(Q) = (\mathbb{E}_Q[X_t], \sqrt{\mathbb{V}_Q[X_t]})'$. Under the assumed ground truth data generating measure P_{GT} , the value of this functional is approximately $T(P_{GT}) \approx (0.3, 0.16)'$. First, we are interested in testing the null hypothesis

$$\mathcal{H}_0 = \left\{ P \in \mathcal{M}_1(X^{\mathbb{N}}) : T(P(X_t \in \cdot | \mathcal{F}_{t-1})) = \begin{pmatrix} 0.4 \\ 0.4 \end{pmatrix} \text{ for all } t \in \mathbb{N}, P\text{-a.s.} \right\}. \tag{22}$$

Although this functional is both an elicitable and identifiable functional, we choose to test it as the latter, where we use the identification function $m((\lambda_\mu, \lambda_\sigma), x) = (\lambda_\mu - x, \lambda_\mu^2 + \lambda_\sigma^2 - x^2)'$ which satisfies

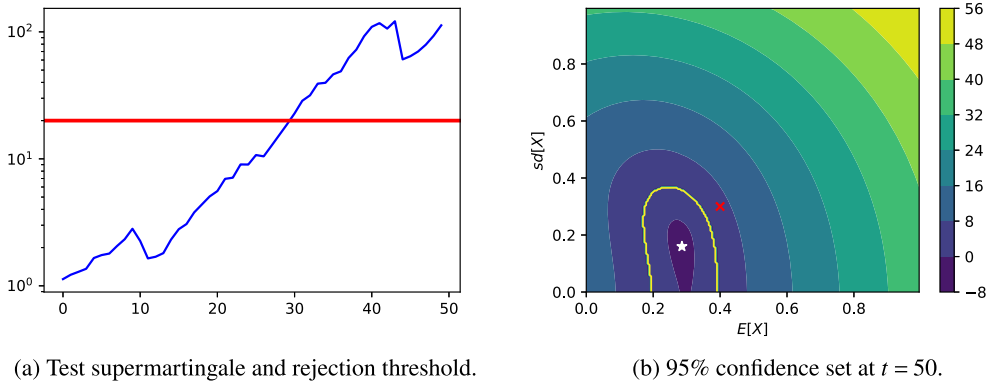


Figure 1. Hypothesis tests and confidence sets for the joint mean and standard deviation of the i.i.d. Beta data generating process. **(1a)** The path of the test supermartingale W_t constructed to test the hypothesis (22) (blue) with $\alpha = 0.05$ rejection threshold (red). **(1b)** The confidence set obtained at $t = 50$ (yellow line). The color gradient represents the value of $\log(W_{50})$ over the space of possible of null hypotheses. The white \star represents the true parameter values, and the red x represents the null hypothesis that is being tested in Figure (1a).

the uniform boundedness conditions required in order to generate a family of test supermartingales according to Lemma 3.2. Using this family of tests, we apply the FTL algorithm. The result of this test and of a confidence set on a single simulated path of the data generating process is displayed in Figure 1.

We conduct a second experiment with the same data generating distribution on the elicitable functional $T(Q) = (\text{VaR}_{0.05}(Q), \text{CVaR}_{0.05}(Q))'$, producing both the Value-at-Risk, that is, the 5%-quantile, and conditional Value-at-Risk or expected shortfall, that is $\mathbb{E}_Q[X_i | X_i < \text{VaR}_{0.05}(Q)]$, where this functional is identifiable with identification function $m((\lambda_v, \lambda_c), x) = (\mathbb{1}_{x \leq \lambda_v} - \alpha_0, x \mathbb{1}_{x \leq \lambda_v} - \alpha_0 \lambda_c)'$ where $\alpha_0 = 0.05$, see Fissler and Ziegel (2016, Theorem 5.2). We test the hypothesis


$$\mathcal{H}_0 = \left\{ P \in \mathcal{M}_1(X^{\mathbb{N}}) : T(P_{X_t \in \cdot | \mathcal{F}_{t-1}}) = \begin{pmatrix} 0.2 \\ 0.1 \end{pmatrix} \text{ for all } t \in \mathbb{N}, P\text{-a.s.} \right\}, \tag{23}$$

where we approximately have that $T(P_{\text{GT}}) \approx (0.06, 0.04)'$ under the true data generating measure P_{GT} . Since the underlying random variables are bounded, we construct a composite test supermartingale using the family of test supermartingales for uniformly bounded elicitable hypotheses given in Lemma 3.1. We optimize this composite martingale using the FTL algorithm. The results are shown in Figure 2.


Let us now consider the problem of estimating the linear coefficient of an AR(1) model. Specifically, assume that the data generating process is

$$X_{t+1} = \beta X_t + \xi_{t+1},$$

where the $\{\xi_t\}_{t \in \mathbb{N}}$ are a 1-sub-Gaussian martingale difference sequence. This assumption may be replaced with σ -sub-Gaussianity for any $\sigma > 0$ with straightforward minor modifications. We are interested in estimating and testing hypotheses regarding the value of $\beta \in \mathbb{R}$. For this purpose, we write this coefficient as the functional $\beta(P(\cdot | \mathcal{F}_t)) = \min_{\beta' \in \mathbb{R}} \mathbb{E}_P[(X_{t+1} - \beta' X_t)^2 | \mathcal{F}_t]$, which admits the identification function $m(\beta_0, (x_t, x_{t+1})) = (x_t | x_t)(\beta_0 x_t - x_{t+1})$ which will be 1-sub-Gaussian with



(a) Test supermartingale and rejection threshold.



(b) 95% confidence set at $t = 150$

Figure 2. Hypothesis tests and confidence sets for the joint quantile and expected shortfall of the i.i.d. Beta data generating process. **(2a)** The path of the test supermartingale W_t constructed to test the hypothesis (23) (blue) with $\alpha = 0.05$ rejection threshold (red). **(2b)** The confidence set obtained at $t = 50$ (white line). The color gradient represents the value of $\log(W_{50})$ over the space of possible of null hypotheses, lying in the space of values such that $CVaR_{0.05} \leq VaR_{0.05}$. The white \star represents the true parameter values, and the red x represents the null hypothesis that is being tested in Figure (2a).

$\psi(u) = (1/2)u^2$ as shown in Example 3. Hence, we may define the family of test supermartingales,

$$L_t^\eta = \prod_{i=1}^t \exp\left(\eta m(\beta_0, (x_i, x_{i+1})) - \frac{\eta^2}{2}\right),$$

indexed by $\eta \in \Theta = \mathbb{R}$. Using this family of test supermartingales and noting that $\eta \mapsto \log L_t^\eta$ is strongly concave, we construct a composite test supermartingale by applying the the FTL algorithm. This update rule has the advantage that it can be computed in closed form at each iteration. Indeed, letting $\bar{m}_t = (1/t) \sum_{i=1}^t m(\beta_0, (x_i, x_{i+1}))$, we have the update rule $\eta_{t+1} = \operatorname{argmin}_{\eta \in \Theta} \{-\log L_t^\eta\} = \operatorname{argmax}_{\eta \in \mathbb{R}} \{\eta \bar{m}_t - \frac{1}{2}\eta^2\} = \bar{m}_t$. Applying this OCO algorithm to randomly generated data from the AR(1) process, we test the hypothesis that $\mathcal{H}_0 = \{P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : \beta(P(\cdot | \mathcal{F}_t)) = 0.65 \text{ for all } t \in \mathbb{N}\}$ when data is generated from an AR(1) process with $\beta = 0.5$ and where $\{\epsilon_t\}_{t \in \mathbb{N}}$ are i.i.d. with $\epsilon_t \sim N(0, 0.8)$. The results of this hypothesis test are displayed in Figure 3a. Similarly, in Figure 3b, we display the running confidence sequence generated using the method described in Section 5.

Appendix A: One-sided and set-valued hypotheses

The methodology considered in this paper may be extended towards testing one-sided and more generally, set-valued hypotheses on elicitable and identifiable functionals. Indeed, for a set $K \subseteq \Lambda$ let us consider the hypothesis

$$\mathcal{H}_0(K) = \{P \in \mathcal{M}_1(\mathcal{X}^{\mathbb{N}}) : T(P(X_t \in \cdot | \mathcal{F}_{t-1})) \in K \text{ for all } t \in \mathbb{N}, P\text{-a.s.}\}.$$

where K can be thought of as a set or interval in which we believe the functional T should lie when applied to the data-generating measure.

We can test the hypothesis $\mathcal{H}_0(K)$ by building an e-process out of test supermartingales for pointwise hypotheses. Indeed, for each $\lambda_0 \in K$, assume that $(W_t^{\lambda_0})_{t \in \mathbb{N}}$ is a test supermartingale for $\mathcal{H}_0(\{\lambda_0\}) =$

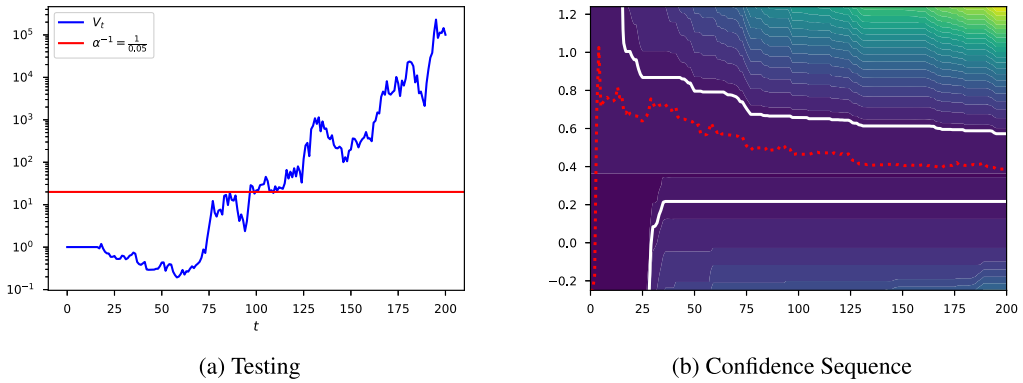


Figure 3. Hypothesis tests and confidence sets for the parameter β of an AR(1) data generating process. **(3a)** The path of the test supermartingale W_t (blue) constructed and the rejection threshold (red), constructed to test the hypothesis that $\beta(P(\cdot | \mathcal{F}_t)) = 0.65$ at the level $\alpha = 0.05$. **(3b)** The boundary of the confidence sequence (white) and the running estimate of the parameter β (red). The color gradient represents the value of $\log(W_t)$ over the space of possible of null hypotheses (y axis) at each point in time.

$\mathcal{H}_0(\lambda_0)$ given at (8). If we define the composite process $(M_t^K)_{t \in \mathbb{N}}$ as $M_t^K = \min_{\lambda_0 \in K} W_t^{\lambda_0}$, it is easy to see that for any \mathcal{F} -adapted stopping time τ , we have that

$$\mathbb{E} [M_\tau^K] = \mathbb{E} \left[\min_{\lambda_0 \in K} W_\tau^{\lambda_0} \right] \leq \min_{\lambda_0 \in K} \mathbb{E} [W_t^{\lambda_0}] \leq 1,$$

showing that $M_\tau^K > \alpha^{-1}$ can serve as a valid test for \mathcal{H}_0^K . In principle, any test derived over the course of this paper for pointwise hypotheses can be used to generate the set of test supermartingales $W_t^{\lambda_0}$ used to test the set-valued hypothesis.

Appendix B: Proofs

Proof of Lemma 2.4. Since W is nonnegative and $W_0 = 1$, we only need to check the P -supermartingale property for any $P \in \mathcal{H}_0$. For each $t \in \mathbb{N}$ we use Tonelli’s theorem, the fact that W_t and π_{t+1} are \mathcal{F}_t -measurable, the P -supermartingale property of L^θ , and the fact that π_{t+1} has total mass one, to obtain

$$\begin{aligned} \mathbb{E}_P [W_{t+1} | \mathcal{F}_t] &= W_t \mathbb{E}_P \left[\int_{\theta \in \Theta} \frac{L_{t+1}^\theta}{L_t^\theta} \pi_{t+1}(d\theta) | \mathcal{F}_t \right] \\ &= W_t \int_{\theta \in \Theta} \mathbb{E}_P \left[\frac{L_{t+1}^\theta}{L_t^\theta} | \mathcal{F}_t \right] \pi_{t+1}(d\theta) \leq W_t \pi_{t+1}(\Theta) = W_t. \end{aligned}$$

This shows that W is a P -supermartingale. □

Proof of Lemma 3.1. First, we note that $\inf_{\lambda \in \Lambda, x \in \mathcal{X}} \{s(\lambda_0, x) - s(\lambda, x)\} > -1$ implies that each L^λ is nonnegative. Next, for each $t \in \mathbb{N}$, $\lambda \in \Lambda$ and $P \in \mathcal{H}_0$, we may compute

$$\mathbb{E}_P [L_{t+1}^\lambda | \mathcal{F}_t] = L_t^\lambda (1 + \mathbb{E}_P [s(\lambda_0, X_{t+1}) - s(\lambda, X_{t+1}) | \mathcal{F}_t]).$$

By assumption, since $P \in \mathcal{H}_0$ with \mathcal{H}_0 defined according to equation (9), we have that

$$-1 < \mathbb{E}_P [s(\lambda_0, X_{t+1}) - s(\lambda, X_{t+1}) \mid \mathcal{F}_t] \leq 0,$$

which implies that

$$L_t^\lambda (1 + \mathbb{E}_P [s(\lambda_0, X_{t+1}) - s(\lambda, X_{t+1}) \mid \mathcal{F}_t]) \leq L_t^\lambda,$$

and hence $\mathbb{E}_P [L_{t+1}^\lambda \mid \mathcal{F}_t] \leq L_t^\lambda$. Lastly, $L_0^\lambda = 1$. □

Proof of Lemma 3.2. We begin with the proof that $A_{m,\lambda_0}^{\text{bd}}$ is convex with non-empty interior. By assumption, there exists $C > 0$ such that $\sup_{x \in \mathcal{X}} \|m(\lambda_0, x)\| \leq C < \infty$. Hence, for any $\|\eta\| < 1/C$, we have by the Cauchy-Schwarz inequality that

$$\langle \eta, m(\lambda_0, x) \rangle \geq -\|\eta\| \|m(\lambda_0, x)\| \geq -\|\eta\| \sup_{x \in \mathcal{X}} \|m(\lambda_0, x)\| > -1,$$

showing that $A_{m,\lambda_0}^{\text{bd}}$ has non-empty interior. Convexity follows from linearity of the scalar product in the first argument.

For $\eta \in A_{m,\lambda_0}^{\text{bd}}$, we have $1 + \langle \eta, m(\lambda_0, X_t) \rangle > 0$, and hence the L^η are all nonnegative. For each $t \in \mathbb{N}$, and $P \in \mathcal{H}_0$, it holds that

$$\mathbb{E}_P [L_{t+1}^\eta \mid \mathcal{F}_t] = L_t^\eta (1 + \langle \eta, \mathbb{E}_P [m(\lambda_0, X_{t+1}) \mid \mathcal{F}_t] \rangle).$$

By equation (10), we have that $\mathbb{E}_P [\langle \eta, m(\lambda_0, X_{t+1}) \rangle \mid \mathcal{F}_t] = 0$ and thus $\mathbb{E}_P [L_{t+1}^\eta \mid \mathcal{F}_t] = L_t^\eta$ for all t . Lastly, since $L_0^\eta = 1$ and $L_t^\eta > 0$ we conclude that L^η is an \mathcal{H}_0 test martingale. □

Proof of Lemma 3.5. Let Y be a supermartingale. By the sub- ψ property, and denoting the forward increments of processes as $\Delta Z_t = Z_{t+1} - Z_t$, we have that

$$\begin{aligned} e^{u\bar{Y}_t} \mathbb{E} \left[e^{u(Y_t - \bar{Y}_t) - V_t \psi(u)} \mid \mathcal{F}_{t-1} \right] &\leq e^{uY_{t-1} - V_{t-1} \psi(u)} e^{u\Delta \bar{Y}_{t-1}} \\ &= e^{uY_{t-1} - V_{t-1} \psi(u)} e^{u\mathbb{E}[\Delta Y_{t-1} \mid \mathcal{F}_{t-1}]} \leq e^{uY_{t-1} - V_{t-1} \psi(u)}, \end{aligned}$$

where the last inequality follows since the supermartingale property implies that $\mathbb{E}[Y_t - Y_{t-1} \mid \mathcal{F}_{t-1}] \leq 0$.

Now, assume that $(e^{uY_t - V_t \psi(u)})_{t \in \mathbb{N}}$ is a supermartingale for all $u \in [0, u_{\max})$. Since $\psi'(0) = 0$ we have

$$\begin{aligned} \mathbb{E}[\Delta Y_{t-1} \mid \mathcal{F}_{t-1}] &= \mathbb{E}[\Delta Y_{t-1} - \Delta V_{t-1} \psi'(0) \mid \mathcal{F}_{t-1}] \\ &= \mathbb{E} \left[\lim_{u \downarrow 0} \frac{e^{u\Delta Y_{t-1} - \Delta V_{t-1} \psi(u)} - 1}{u} \mid \mathcal{F}_{t-1} \right] = \lim_{u \downarrow 0} \frac{\mathbb{E} [e^{u\Delta Y_{t-1} - \Delta V_{t-1} \psi(u)} \mid \mathcal{F}_{t-1}] - 1}{u}, \end{aligned}$$

where in the third line we use the fact that since ψ is convex and differentiable at zero, there exists a closed neighborhood including zero in which $(e^{u\Delta Y_{t-1} - \Delta V_{t-1} \psi(u)} - 1)/u$ is continuous and hence has an integrable upper bound, allowing us to exchange the limit and the expectation by the dominated convergence theorem. Lastly, noting that $(\mathbb{E}[e^{u\Delta Y_{t-1} - \Delta V_{t-1} \psi(u)} \mid \mathcal{F}_{t-1}] - 1)/u \leq 0$ by the assumed supermartingale property, we conclude that $\mathbb{E}[\Delta Y_{t-1} \mid \mathcal{F}_{t-1}] \leq 0$, demonstrating that $(Y_t)_{t \in \mathbb{N}}$ is a supermartingale, as desired. □

Proof of Lemmas 3.6 and 3.7. Let us first consider the setting of Lemma 3.6. By construction, each of the processes $L^{\lambda,u}$ is nonnegative and satisfies $L_0^{\lambda,u} = 1$. Fixing (λ, u) , we have that for all $P \in \mathcal{H}_0$,

$$\mathbb{E}_P \left[L_{t+1}^{\lambda,u} \mid \mathcal{F}_t \right] = \mathbb{E}_P \left[e^{uY_{t+1}^\lambda - V_{t+1}^\lambda \psi_\lambda(u)} \mid \mathcal{F}_t \right] \leq e^{uY_t^\lambda - V_t^\lambda \psi_\lambda(u)} = L_t^{\lambda,u} \text{ a.s.,}$$

which follows by noting that Y^λ is a supermartingale according to the definition of \mathcal{H}_0 and by applying Lemma 3.5. Hence, each $L^{\lambda,u}$ is a valid test supermartingale.

For Lemma 3.7, we note that the proof works analogously to that of Lemma 3.6. □

Proof of Lemma 4.1. We prove this claim by proving it for the four separate classes of test supermartingales of Lemmas 3.1, 3.2, 3.6 and 3.7.

Lemma 3.1. In this case, the family of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$, where $\theta = \lambda \in \Lambda$, can be written as

$$\log L_t^\theta = \sum_{1 \leq i \leq t} \log(1 + s(\lambda_0, x_i) - s(\lambda, x_i)) . \tag{24}$$

Let us define $f_i(\lambda) = s(\lambda_0, x_i) - s(\lambda, x_i)$, which is concave since $s(\lambda, x_i)$ is convex. Using the concavity and monotonicity of log and the concavity of the f_i , we have that for any $\lambda, \lambda' \in \Lambda$ and $\rho \in (0, 1)$,

$$\begin{aligned} \log(1 + f_i(\rho\lambda + (1 - \rho)\lambda')) &\geq \log(\rho(1 + f_i(\lambda)) + (1 - \rho)f_i(\lambda')) \\ &\geq \rho \log(1 + f_i(\lambda)) + (1 - \rho) \log(1 + f_i(\lambda')), \end{aligned}$$

showing that $\log(1 + f_i(\lambda))$ is concave. Hence, $\log L_t^\theta$ is a sum of concave functions of θ , and hence is concave as well.

Lemma 3.2. In this case, the family of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$, where $\eta = \theta \in \Theta \subseteq A_{m, \lambda_0}^{bd}$, can be written as

$$\log L_t^\theta = \sum_{1 \leq i \leq t} \log(1 + \langle \eta, m(\lambda_0, x_i) \rangle) . \tag{25}$$

By the concavity of the log, each of the $\log(1 + \langle \eta, m(\lambda_0, x_i) \rangle)$ are concave in η . Hence, $\log L_t^\theta$ is a sum of concave functions of θ , and hence is concave as well.

Lemma 3.6. In this case, the family of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$, where $(\lambda, u) = \theta \in \Theta \subseteq [0, u_{\max}] \times \Lambda$, can be written as

$$\log L_t^\theta = \sum_{1 \leq i \leq t} \left\{ u(\theta) (s(\lambda_0, x_i) - s(\lambda(\theta), x_i)) - v_i^{\lambda(\theta)} \psi_{\lambda(\theta)}(u(\theta)) \right\} . \tag{26}$$

By assumption, each summand is concave in $\theta = (\lambda, u)$ and hence $\log L_t^\theta$ must also be concave in θ .

Lemma 3.7. In this case, the family of test supermartingales $\{L^\theta\}_{\theta \in \Theta}$, where $(\eta, u) = \theta \in \Theta \subseteq A_{m, \lambda_0}^\psi \times [0, u_{\max}]$, can be written as

$$\log L_t^\theta = \sum_{1 \leq i \leq t} \left\{ \langle u(\theta), m(\lambda_0, x_i) \rangle - v_i^{\lambda(\theta)} \psi_{\eta(\theta)}(u(\theta)) \right\} . \tag{27}$$

By assumption, each summand is concave and hence $\log L_t^\theta$ is a sum of concave functions and hence is itself almost surely concave. □

Proof of Theorem 4.2. Using the definition of regret in equation (17), we have that for all $\theta \in \Theta$,

$$\frac{1}{T} \log W_T \geq \frac{1}{T} \log L_T^\theta - \frac{1}{T} B_T,$$

where $B_T = \text{Regret}_T$. Applying the assumptions that $B_T = o(T)$ and (18), we have that

$$\liminf_{T \rightarrow \infty} \frac{\log \max_{0 \leq t \leq T} W_t}{T} \geq \liminf_{T \rightarrow \infty} \frac{\log W_T}{T} \geq \liminf_{T \rightarrow \infty} \left\{ \frac{\log L_T^\theta}{T} - \frac{B_T}{T} \right\} = \liminf_{T \rightarrow \infty} \frac{\log L_T^\theta}{T} > 0,$$

where these inequalities hold Q -almost surely. This implies that, Q -almost surely, there exists a $T' \in \mathbb{N}$ and $\epsilon > 0$ such that $\max_{0 \leq t \leq T} W_t \geq e^{\epsilon T}$ for all $T > T'$. Hence, we may conclude that $\lim_{T \rightarrow \infty} \max_{0 \leq t \leq T} W_t \geq \lim_{T \rightarrow \infty} e^{\epsilon T} = \infty$ Q -almost surely, as desired. \square

Proof of Proposition 4.3. Cases (i) and (ii): By Lemma 3.1 or 3.2, we may write each test submartingale as $\log L_t^\theta = \sum_{i=1}^t \log(1 + f(\theta, X_i))$, where we have either $f(\theta, x) = s(\lambda_0, x) - s(\theta, x)$ or $f(\theta, x) = \langle \theta, m(\lambda_0, x) \rangle$, respectively. Since $(X_t)_{t \in \mathbb{N}}$ is assumed to be stationary and ergodic under Q_∞ , we may apply the Birkhoff-Khinchin ergodic theorem (Cornfeld, Fomin and Sinai, 2012, Theorem 1) to obtain

$$\lim_{t \rightarrow \infty} \frac{\log L_t^\theta}{t} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t \log(1 + f(\theta, X_i)) = \mathbb{E}_{Q_\infty} [\log(1 + f(\theta, X_\infty))] . \tag{28}$$

Cases (iii) and (iv): By Lemma 3.6 or 3.7, we may write each test submartingale as

$$\log L_t^\theta = \log L_t^{(\theta', u)} = \sum_{i=1}^t \left\{ u f(\theta', X_i) - v_i^{\theta'} \psi_{\theta'}(u) \right\},$$

where have $u \in [0, \epsilon]$ with $\epsilon > 0$, $v_i^{\theta'} > 0$ and, either $f(\theta', x) = s(\lambda_0, x) - s(\theta', x)$ in the elicitable case, or $f(\theta', x) = \langle \theta', m(\lambda_0, x) \rangle$ in the identifiable case. By assumption, in each case there exists a θ'_0 (where $\Theta' \subseteq \Lambda$ in the elicitable case and $\Theta' \subseteq A_{m, \lambda_0}^\psi$ in the identifiable case) such that $E_{Q_\infty} f(\theta'_0, X_\infty) = c > 0$. Since $(X_t)_{t \in \mathbb{N}}$ is assumed to be stationary and ergodic under Q , we may apply the Birkhoff-Khinchin ergodic theorem (Cornfeld, Fomin and Sinai, 2012, Theorem 1) to $L^{(\theta'_0, u)}$ to obtain

$$\begin{aligned} \lim_{t \rightarrow \infty} \frac{\log L_t^{(\theta'_0, u)}}{t} &= \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t \left\{ u f(\theta'_0, X_i) - v_i^{\theta'_0} \psi_{\theta'_0}(u) \right\} \\ &= u \mathbb{E}_{Q_\infty} [f(\theta'_0, X_\infty)] - \bar{v} \psi_{\theta'_0}(u), \end{aligned}$$

where $0 < \bar{v} = \lim_{t \rightarrow \infty} \frac{1}{t} \sum_{i=1}^t \bar{v}$ which is non-negative and finite by assumption. By Definition 3.4, we find that $\psi_{\theta'_0}(0) = 0$ and $\lim_{u \searrow 0} \psi_{\theta'_0}(0)/u = 0$. Hence, letting $g(u) = u c - \bar{v} \psi_{\theta'_0}(u)$, we find that, $g(0) = 0$ and $g'(0) = \lim_{u \searrow 0} g(u)/u = c > 0$. Hence, there exists $u'_0 \in (0, \epsilon)$ such that $g(u'_0) > 0$. By picking this u'_0 , we then have that $\lim_{t \rightarrow \infty} \log L_t^{(\theta'_0, u'_0)} / t = g(u'_0) = c > 0$, yielding the desired result. \square

References

Albers, C. (2019). The problem with unadjusted multiple and sequential statistical testing. *Nat. Commun.* **10** 1–4.

- Caponnetto, A. (2005). A note on the role of squared loss in regression. CBCL Paper, Massachusetts Institute of Technology, Cambridge, MA.
- Choe, Y.J. and Ramdas, A. (2021). Comparing sequential forecasters. Preprint. Available at [arXiv:2110.00115](https://arxiv.org/abs/2110.00115).
- Cornfeld, I.P., Fomin, S.V. and Sinai, Y.G. (2012). *Ergodic Theory*. New York: Springer.
- de la Peña, V.H., Klass, M.J. and Lai, T.L. (2004). Self-normalized processes: Exponential inequalities, moment bounds and iterated logarithm laws. *Ann. Probab.* **32** 1902–1933. [MR2073181 https://doi.org/10.1214/009117904000000397](https://doi.org/10.1214/009117904000000397)
- Dimitriadis, T., Fissler, T. and Ziegel, J.F. (2020). The efficiency gap. Preprint. Available at [arXiv:2010.14146](https://arxiv.org/abs/2010.14146).
- Dimitriadis, T., Fissler, T. and Ziegel, J. (2023). Characterizing M-estimators. *Biometrika* **asad026**. <https://doi.org/10.1093/biomet/asad026>
- Fissler, T. (2017). On higher order elicibility and some limit theorems on the Poisson and Wiener space. Ph.D. thesis, Univ. Bern.
- Fissler, T. and Ziegel, J.F. (2016). Higher order elicibility and Osband’s principle. *Ann. Statist.* **44** 1680–1707. [MR3519937 https://doi.org/10.1214/16-AOS1439](https://doi.org/10.1214/16-AOS1439)
- Freedman, D.A. (1975). On tail probabilities for martingales. *Ann. Probab.* **3** 100–118. [MR0380971 https://doi.org/10.1214/aop/1176996452](https://doi.org/10.1214/aop/1176996452)
- Frongillo, R. and Kash, I. (2015). Vector-valued property elicitation. In *Proceedings of the 28th Conference on Learning Theory*. PMLR **40** 710–727.
- Frongillo, R.M. and Kash, I.A. (2021). Elicitation complexity of statistical properties. *Biometrika* **108** 857–879. [MR4341356 https://doi.org/10.1093/biomet/asaa093](https://doi.org/10.1093/biomet/asaa093)
- Frongillo, R. and Nobel, A. (2020). Memoryless sequences for general losses. *J. Mach. Learn. Res.* **21** Paper No. 80, 28. [MR4119148](https://doi.org/10.48550/jmlr.2020.21.80)
- Gneiting, T. (2011). Making and evaluating point forecasts. *J. Amer. Statist. Assoc.* **106** 746–762. [MR2847988 https://doi.org/10.1198/jasa.2011.r10138](https://doi.org/10.1198/jasa.2011.r10138)
- Grünwald, P., de Heide, R. and Koolen, W. (2019). Safe testing. Preprint. Available at [arXiv:1906.07801](https://arxiv.org/abs/1906.07801).
- Hagerup, T. and Rüb, C. (1990). A guided tour of Chernoff bounds. *Inform. Process. Lett.* **33** 305–308. [MR1045520 https://doi.org/10.1016/0020-0190\(90\)90214-I](https://doi.org/10.1016/0020-0190(90)90214-I)
- Hazan, E. (2016). Introduction to online convex optimization. *Found. Trends Optim.* **2** 157–325.
- Henzi, A. and Ziegel, J.F. (2022). Valid sequential inference on probability forecast performance. *Biometrika* **109** 647–663. [MR4472840 https://doi.org/10.1093/biomet/asab047](https://doi.org/10.1093/biomet/asab047)
- Hertz, D. (2020). Improved Hoeffding’s lemma and Hoeffding’s tail bounds. Preprint. Available at [arXiv:2012.03535](https://arxiv.org/abs/2012.03535).
- Hoeffding, W. (1994). Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding* 409–426. New York, NY: Springer.
- Howard, S.R., Ramdas, A., McAuliffe, J. and Sekhon, J. (2020). Time-uniform Chernoff bounds via nonnegative supermartingales. *Probab. Surv.* **17** 257–317. [MR4100718 https://doi.org/10.1214/18-PS321](https://doi.org/10.1214/18-PS321)
- Howard, S.R., Ramdas, A., McAuliffe, J. and Sekhon, J. (2021). Time-uniform, nonparametric, nonasymptotic confidence sequences. *Ann. Statist.* **49** 1055–1080. [MR4255119 https://doi.org/10.1214/20-aos1991](https://doi.org/10.1214/20-aos1991)
- Hsu, J.C. (1996). *Multiple Comparisons: Theory and Methods*. London: CRC Press. [MR1629127 https://doi.org/10.1007/978-1-4899-7180-7](https://doi.org/10.1007/978-1-4899-7180-7)
- Jun, K.-S. and Orabona, F. (2019). Parameter-free online convex optimization with sub-exponential noise. In *Conference on Learning Theory* 1802–1823. PMLR.
- Komunjer, I. and Vuong, Q. (2010a). Semiparametric efficiency bound in time-series models for conditional quantiles. *Econometric Theory* **26** 383–405. [MR2600568 https://doi.org/10.1017/S0266466609100038](https://doi.org/10.1017/S0266466609100038)
- Komunjer, I. and Vuong, Q. (2010b). Efficient estimation in dynamic conditional quantile models. *J. Econometrics* **157** 272–285. [MR2661600 https://doi.org/10.1016/j.jeconom.2010.01.001](https://doi.org/10.1016/j.jeconom.2010.01.001)
- Lambert, N., Pennock, D.M. and Shoham, Y. (2008). Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce* 129–138. Chicago, IL, USA. Extended abstract.
- McMahan, H.B. (2017). A survey of algorithms and analysis for adaptive online learning. *J. Mach. Learn. Res.* **18** Paper No. 90, 50. [MR3714253](https://doi.org/10.48550/jmlr.2017.18.90)
- O’Neill, R. and Wetherill, G.B. (1971). The present state of multiple comparison methods. *J. Roy. Statist. Soc. Ser. B* **33** 218–241, 241–250. [MR0321252](https://doi.org/10.2307/2343252)

- Orabona, F. and Jun, K.-S. (2021). Tight concentrations and confidence sequences from the regret of universal portfolio. Preprint. Available at [arXiv:2110.14099](https://arxiv.org/abs/2110.14099).
- Osband, K.H. (1985). Providing incentives for better cost forecasting. Ph.D. thesis, Univ. California, Berkeley.
- Ramdas, A., Ruf, J., Larsson, M. and Koolen, W. (2020). Admissible anytime-valid sequential inference must rely on nonnegative martingales. Preprint. Available at [arXiv:2009.03167](https://arxiv.org/abs/2009.03167).
- Ramdas, A., Ruf, J., Larsson, M. and Koolen, W.M. (2022a). Testing exchangeability: Fork-convexity, supermartingales and e-processes. *Internat. J. Approx. Reason.* **141** 83–109. [MR4364897 https://doi.org/10.1016/j.ijar.2021.06.017](https://doi.org/10.1016/j.ijar.2021.06.017)
- Ramdas, A., Grünwald, P., Vovk, V. and Shafer, G. (2022b). Game-theoretic statistics and safe anytime-valid inference. Preprint. Available at [arXiv:2210.01948](https://arxiv.org/abs/2210.01948).
- Robbins, H. and Siegmund, D. (1972). A class of stopping rules for testing parametric hypotheses. In *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability (Univ. California, Berkeley, Calif., 1970/1971), Vol. IV: Biology and Health* 37–41. Berkeley, Calif.: Univ. California Press. [MR0403111](https://doi.org/10.1214/10-STS347)
- Rockafellar, R.T. and Wets, R.J.-B. (2009). *Variational Analysis*. Berlin, Heidelberg: Springer.
- Savage, L.J. (1971). Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* **66** 783–801. [MR0331571](https://doi.org/10.1080/01621459.1971.10473157)
- Shafer, G., Shen, A., Vereshchagin, N. and Vovk, V. (2011). Test martingales, Bayes factors and p -values. *Statist. Sci.* **26** 84–101. [MR2849911 https://doi.org/10.1214/10-STS347](https://doi.org/10.1214/10-STS347)
- Shalev-Shwartz, S. (2011). Online learning and online convex optimization. *Found. Trends Mach. Learn.* **4** 107–194.
- Shekhar, S. and Ramdas, A. (2021). Nonparametric two-sample testing by betting. Preprint. Available at [arXiv:2112.09162](https://arxiv.org/abs/2112.09162).
- Steinwart, I., Pasin, C., Williamson, R. and Zhang, S. (2014). Elicitation and identification of properties. In *Conference on Learning Theory* 482–526. PMLR.
- Ville, J. (1939). *Étude Critique de la Notion de Collectif*. NUMDAM [place of publication not identified]. [MR3533075](https://doi.org/10.1214/aoms/1177731118)
- Vovk, V. and Wang, R. (2021). E-values: Calibration, combination and applications. *Ann. Statist.* **49** 1736–1754. [MR4298879 https://doi.org/10.1214/20-aos2020](https://doi.org/10.1214/20-aos2020)
- Wald, A. (1945). Sequential tests of statistical hypotheses. *Ann. Math. Stat.* **16** 117–186. [MR0013275 https://doi.org/10.1214/aoms/1177731118](https://doi.org/10.1214/aoms/1177731118)
- Waudby-Smith, I. and Ramdas, A. (2023). Estimating means of bounded random variables by betting. *J. Roy. Statist. Soc. Ser. B* [qkad009. https://doi.org/10.1093/jrsssb/qkad009](https://doi.org/10.1093/jrsssb/qkad009)
- Xu, Z., Wang, R. and Ramdas, A. (2021). A unified framework for bandit multiple testing. Preprint. Available at [arXiv:2107.07322](https://arxiv.org/abs/2107.07322).
- Zinkevich, M. (2003). Online convex programming and generalized infinitesimal gradient ascent. In *Proceedings of the 20th International Conference on Machine Learning* 928–936.

Received April 2022 and revised May 2023