

## MINIMAX OPTIMAL SEQUENTIAL HYPOTHESIS TESTS FOR MARKOV PROCESSES

BY MICHAEL FAUSS<sup>1,\*</sup>, ABDELHAK M. ZOUBIR<sup>1,\*\*</sup> AND H. VINCENT POOR<sup>2</sup>

<sup>1</sup>Signal Processing Group, TU Darmstadt, \*[mfauss@princeton.edu](mailto:mfauss@princeton.edu); \*\*[zoubir@spg.tu-darmstadt.de](mailto:zoubir@spg.tu-darmstadt.de)

<sup>2</sup>Department of Electrical Engineering, Princeton University, [poor@princeton.edu](mailto:poor@princeton.edu)

Under mild Markov assumptions, sufficient conditions for strict minimax optimality of sequential tests for multiple hypotheses under distributional uncertainty are derived. First, the design of optimal sequential tests for simple hypotheses is revisited, and it is shown that the partial derivatives of the corresponding cost function are closely related to the performance metrics of the underlying sequential test. Second, an implicit characterization of the least favorable distributions for a given testing policy is stated. By combining the results on optimal sequential tests and least favorable distributions, sufficient conditions for a sequential test to be minimax optimal under general distributional uncertainties are obtained. The cost function of the minimax optimal test is further identified as a generalized  $f$ -dissimilarity and the least favorable distributions as those that are most similar with respect to this dissimilarity. Numerical examples for minimax optimal sequential tests under different uncertainties illustrate the theoretical results.

**1. Introduction.** Sequential hypothesis tests are well known for being highly efficient in terms of the number of required samples and, as a consequence, for minimizing the average decision delay in time-critical applications. In his seminal book [62], Wald showed that, compared to fixed sample size tests, sequential tests can reduce the average number of samples by a factor of approximately two. In general, the ability to allow the overall number of samples to depend on the current history makes sequential procedures more flexible and adaptable than procedures whose sample size is chosen a priori. Comprehensive overviews of sequential hypothesis testing and related topics can be found in [22, 47, 53, 55, 62], to name just a few.

A well-established drawback of sequential hypothesis tests is that their higher efficiency depends critically on the assumption that the process generating the observations indeed follows the assumed model. If this is not the case, that is, if a model mismatch occurs, the number of samples can increase significantly; compare, for example, [55], Figure 3.4, which illustrates the influence of a model mismatch on the expected run length of a sequential test for the mean of an autoregressive process. This observation lead Kiefer and Weiss to propose a sequential test that, in addition to meeting the targeted error probabilities under the hypotheses, minimizes the *maximum* expected run length over all feasible distributions [11, 35]. Different variations of the corresponding optimization problem are known as Kiefer–Weiss problem or modified Kiefer–Weiss problem and have received considerable attention in the literature [10, 37, 45, 66]. However, to the present day, exact solutions to the Kiefer–Weiss problem have only been shown for special cases of binary hypothesis tests.

A natural generalization of the Kiefer–Weiss problem is to include the error probabilities in the minimax criterion, that is, to design a test whose maximum error probabilities are minimal over the set of feasible distributions. For fixed sample sizes this minimax approach to

---

Received September 2018; revised May 2019.

*MSC2020 subject classifications.* Primary 62L10; secondary 62C20.

*Key words and phrases.* Sequential analysis, minimax procedures, distributional uncertainty, robust hypothesis testing, multiple hypothesis testing.

the design of statistical tests was pioneered by Huber [29] and is known as *robust hypothesis testing* or *robust detection*. In general, robust hypothesis tests sacrifice some efficiency under ideal conditions in order to be less sensitive to deviations from the ideal case [30]. In this sense, robust hypothesis tests, and robust statistics in general, form a middle ground between parametric and nonparametric approaches. For an overview of existing results, recent advances and applications of robust statistics, see, for example, [1, 33, 38, 67, 68].

The idea underlying this paper is to leverage both sequential and robust hypothesis testing. Ideally, a robust sequential test is *fast and reliable*, that is, it requires fewer observations, on average, than a fixed sample size test and at the same time works reliably under model mismatch. In what follows, a coherent framework for the design of minimax optimal sequential tests is developed, including the Kiefer–Weiss problem as a special case. The main contributions are summarized as follows:

- Sufficient conditions for strict minimax optimality of sequential tests for multiple hypotheses under distributional uncertainty are derived. The results hold for stochastic processes with a time-homogeneous Markovian representation and under mild assumptions on the type of uncertainty.
- The tests satisfying the minimax optimality conditions are identified as optimal tests for least favorable distributions. The latter are shown to be data dependent such that the underlying stochastic process becomes Markovian, even if the increments are independent under the nominal model.
- It is shown that the least favorable distributions are those that are maximally similar with respect to a weighted statistical similarity measure of the  $f$ -dissimilarity type with the weights being data dependent. Both the similarity measure and the way in which the weights depend on the data are characterized explicitly.
- A method for the design of minimax optimal sequential tests in practice is proposed, and two numerical examples are provided to illustrate the theoretical findings.

The existing literature on minimax hypothesis testing can roughly be divided into two groups. The first group consists of results on strictly minimax optimal tests that are based on two alternating Choquet capacities [8, 31]. If the existence of capacity achieving distributions can be shown, then these distributions are least favorable and maximally similar with respect to *all* convex divergences; compare [31], Theorem 6.1. The derivation of strictly minimax fixed sample size tests in [15, 29, 32] as well as the strictly minimax change detection procedure in [57] are based on this result.

If no capacity-achieving distributions can be identified, least favorable distributions are usually defined as the minimizers of a given statistical divergence, playing the role of a surrogate objective function. That is, instead of solving the robust detection problem exactly, it is shown that choosing distributions that are maximally similar with respect to a suitable divergence are *almost* least favorable and the corresponding tests *almost* minimax. Typical examples are [3, 21, 44, 46, 65]. However, many asymptotic results can also be grouped under this approach in the sense that the solutions are obtained by solving an asymptotic surrogate problem which typically induces the Kullback–Leibler divergence as the corresponding similarity measure. Examples are abundant in the literature: [6, 9, 10, 12, 19, 52, 61, 66]. Comparable asymptotic result can also be found for the closely related problem of minimax quickest change detection [2, 5, 18]; the connection between the two problems is discussed in more detail in Section 7.

The few works on minimax sequential testing that do not fall under the two categories of asymptotic results or capacity based results are usually limited in scope. For example, the results in [60] are highly application specific, the results in [34, 39] are based on strong assumptions on the distributions and the results in [23] only hold for the single sample case.

In summary, what sets this paper apart from the existing literature is that *strictly minimax optimal* sequential tests are derived under *mild distributional assumptions* and that the proof of optimality is given *without invoking Choquet capacities*. Finally, note that the paper extends and generalizes [13] where minimax optimal sequential tests for two hypotheses are studied under slightly stricter assumptions.

The paper is organized as follows. The minimax sequential testing problem is stated formally in Section 2. Optimal sequential tests for multiple simple hypotheses are revisited in Section 3. Some useful results that relate properties of an optimal test to the partial derivatives of its cost function are shown in Section 4. Least favorable distributions and their properties are discussed in Section 5. The main result, sufficient conditions for minimax optimal sequential tests, is stated in Section 6 and followed by a discussion in Section 7. Two illustrative numerical examples are shown in Section 8.

**2. Notation and problem formulation.** In this section the minimax optimal sequential testing problem is defined, and some common notation are introduced. Notation not covered here are defined when they occur in the text.

2.1. *Notation.* Random variables are denoted by uppercase letters and their realizations by lowercase letters. Analogously, probability distributions are denoted by uppercase letters and their densities by the corresponding lowercase letters. Blackboard bold is used to indicate product measures. Measurable sets are denoted by tuples  $(\Omega, \mathcal{F})$ . Boldface lowercase letters are used to indicate vectors; no distinction is made between row and column vectors. The inner product of two vectors,  $\mathbf{x}$  and  $\mathbf{y}$ , is denoted by  $\langle \mathbf{x}, \mathbf{y} \rangle$  and the elementwise product by  $\mathbf{x}\mathbf{y}$ . All comparisons between vectors are defined elementwise. The indicator function of a set  $\mathcal{A}$  is denoted by  $\mathcal{I}(\mathcal{A})$ . All comparisons between functions are defined pointwise.

The notation  $\partial_{y_k} f(\mathbf{y})$  is used for the subdifferential [50], Section 23, of a convex function  $f: \mathcal{Y} \subset \mathbb{R}^K \rightarrow \mathbb{R}$  with respect to  $y_k$  evaluated at  $\mathbf{y}$ , that is,

$$(2.1) \quad \partial_{y_k} f(\mathbf{y}) := \{c \in \mathbb{R} : f(\mathbf{y}') - f(\mathbf{y}) \geq c(y'_k - y_k) \forall \mathbf{y}' \in \mathcal{Y}_k(\mathbf{y})\},$$

with  $\mathcal{Y}_k(\mathbf{y}) := \{\mathbf{y}' \in \mathcal{Y} : y'_j = y_j \forall j \in \{1, \dots, K\} \setminus \{k\}\}$ . The superdifferential of a concave function is defined analogously. Both are referred to as generalized differentials in what follows. The length of the interval corresponding to  $\partial_{y_k} f(\mathbf{y})$  is denoted by

$$(2.2) \quad |\partial_{y_k} f(\mathbf{y})| := \sup_{a, b \in \partial_{y_k} f(\mathbf{y})} |a - b|.$$

If a function  $f_{y_k}$  exists such that  $f_{y_k}(\mathbf{y}) \in \partial_{y_k} f(\mathbf{y}) \forall \mathbf{y} \in \mathcal{Y}$ , then  $f_{y_k}$  is called a partial generalized derivative of  $f$  with respect to  $y_k$ . The set of all partial generalized derivatives,  $f_{y_k}$ , is denoted by  $\partial_{y_k} f$ .

2.2. *Underlying stochastic process.* Let  $(X_n)_{n \geq 1}$  be a discrete-time stochastic process with values in  $(\Omega_X, \mathcal{F}_X)$ . The joint distribution of  $(X_n)_{n \geq 1}$  on the cylinder set

$$(2.3) \quad (\Omega_X^{\mathbb{N}}, \mathcal{F}_X^{\mathbb{N}}) := \left( \prod_{n \geq 1} \Omega_X, \prod_{n \geq 1} \mathcal{F}_X \right)$$

is denoted by  $\mathbb{P}$ , the conditional or marginal distributions of an individual random variable  $X$  on  $(\Omega_X, \mathcal{F}_X)$  by  $P$  and the natural filtration [7], Definition 2.32, of the process  $(X_n)_{n \geq 1}$  by  $(\mathcal{F}_X^n)_{n \geq 1}$ . In order to balance generality and tractability, the analysis in this paper is limited to stochastic processes that satisfy the following three assumptions.

ASSUMPTION 1. The process  $(X_n)_{n \geq 1}$  admits a time-homogeneous Markovian representation. That is, there exists a  $(\Omega_\Theta, \mathcal{F}_\Theta)$ -valued stochastic process  $(\Theta_n)_{n \geq 0}$  adapted to  $(\mathcal{F}_X^n)_{n \geq 1}$  such that

$$(2.4) \quad \begin{aligned} \mathbb{P}(X_{n+1} \in \mathcal{E} \mid X_1 = x_1, \dots, X_n = x_n) &= \mathbb{P}(X_{n+1} \in \mathcal{E} \mid \Theta_n = \theta_n) \\ &=: P_{\theta_n}(\mathcal{E}) \end{aligned}$$

for all  $n \geq 1$  and all  $\mathcal{E} \in \mathcal{F}_X$ . The distribution of  $X_1$  is denoted by  $P_{\theta_0}$ , where  $\theta_0$  is assumed to be deterministic and known a priori. An extension to randomly initialized  $\theta_0$  should not be hard but will not be entered.

ASSUMPTION 2. There exists a function  $\xi : \Omega_\Theta \times \Omega_X \rightarrow \Omega_\Theta$  that is measurable with respect to  $\mathcal{F}_\Theta \otimes \mathcal{F}_X$  and that satisfies

$$(2.5) \quad \Theta_{n+1} = \xi(\Theta_n, X_{n+1}) =: \xi_{\Theta_n}(X_{n+1})$$

for all  $n \geq 0$ .

ASSUMPTION 3. For all  $\theta \in \Omega_\Theta$ , the probability measure  $P_\theta$ , defined in Assumption 1, admits a density  $p_\theta$  with respect to some  $\sigma$ -finite reference measure  $\mu$ .

The set of distributions  $\mathbb{P}$  on  $(\Omega_X^{\mathbb{N}}, \mathcal{F}_X^{\mathbb{N}})$  that satisfy these three assumptions is denoted by  $\mathbb{M}$ . The set of distributions  $P$  on  $(\Omega_X, \mathcal{F}_X)$  that admit densities with respect to  $\mu$  is denoted by  $\mathcal{M}_\mu$ .

The above assumptions are rather mild and are introduced primarily to simplify the presentation of the results. In general, the sufficient statistic  $\Theta$  can be chosen as a sliding window of past samples, that is,  $\Theta_n = (X_{n-m}, \dots, X_n)$ , where  $m$  is a finite positive integer. Hence, the presented results apply to every discrete-time Markov process of finite order. However, in order to implement the test in practice,  $\Omega_\Theta$  should be sufficiently low-dimensional (compare the examples in Section 8). As long as the existence of the corresponding densities is guaranteed, the reference measure  $\mu$  in Assumption 3 can be chosen arbitrarily. This aspect can be exploited to simplify the numerical design of minimax sequential tests and is discussed in more detail in Sections 7 and 8.

2.3. *Uncertainty model and hypotheses.* For general Markov processes the question of how to model distributional uncertainty is nontrivial and has far-reaching implications on the definition of minimax robustness. In the most general case the joint distribution  $\mathbb{P}$  is subject to uncertainty. However, defining meaningful uncertainty models for  $\mathbb{P}$  is an intricate task and is usually neither feasible nor desirable. An approach that is more tractable and more useful in practice is to assume that at any given time instant,  $n \geq 1$ , the *marginal* or *conditional* distribution of  $X_n$  is subject to uncertainty.

In this paper it is assumed that the conditional distributions  $P_\theta$ , as defined in (2.4), are subject to uncertainty. More precisely, for each  $\theta \in \Omega_\Theta$  the conditional distribution  $P_\theta$  is replaced by an *uncertainty set* of feasible distributions  $\mathcal{P}_\theta \subset \mathcal{M}_\mu$ . This model induces an uncertainty set for  $\mathbb{P}$  which is given by

$$(2.6) \quad \mathcal{P} := \left\{ \mathbb{P} \in \mathbb{M} : \mathbb{P} = \prod_{n \geq 0} P_{\theta_n}, P_{\theta_n} \in \mathcal{P}_{\theta_n} \right\}$$

and is completely specified by the corresponding family of uncertainty sets for the conditional distributions  $\{\mathcal{P}_\theta : \theta \in \Omega_\Theta\}$ .

The goal of this paper is to characterize and to design minimax optimal sequential tests for multiple hypotheses under the assumption that under each hypothesis the distribution is subject to the type of uncertainty introduced above. That is, each hypothesis is given by

$$(2.7) \quad \mathcal{H}_k: \mathbb{P} \in \mathcal{P}_k, \quad k = 1, \dots, K,$$

where all  $\mathcal{P}_k$  are of the form (2.6) and are defined by a corresponding family of conditional uncertainty sets  $\{\mathcal{P}_\theta^{(k)} : \theta \in \Omega_\Theta\}$ . Note that the parameter  $\theta$ , which corresponds to the sufficient statistic in Assumption 2, does not depend on  $k$ , that is, the statistic needs to be chosen such that it is sufficient under all hypotheses. Finally, the set  $\mathcal{P}_0$  defines the uncertainty in the distribution under which the expected run length is supposed to be minimum. It is assumed to be of the form (2.6) as well. In practice, one is often interested in minimizing the worst case expected run length under one hypothesis or under any hypotheses, that is,  $\mathcal{P}_0 = \mathcal{P}_k$  for some  $k \in \{1, \dots, K\}$  or  $\mathcal{P}_0 = \bigcup_{k=1}^K \mathcal{P}_k$ , respectively. In principle, however,  $\mathcal{P}_0$  can be chosen freely by the test designer.

Before proceeding, it is useful to illustrate the assumptions on the underlying stochastic process and the proposed uncertainty model with an example. Consider an exponentially weighted moving-average process, that is,

$$(2.8) \quad X_{n+1} = \sum_{l=1}^{\infty} a^l X_{n+1-l} + W_{n+1},$$

where  $a \in (-1, 1)$  is a known scalar and  $(W_n)_{n \geq 1}$  is a sequence of independent random variables that are identically distributed according to  $P_W$ . This process can equivalently be written as

$$(2.9) \quad X_{n+1} = a\Theta_n + W_{n+1},$$

where the sufficient statistic  $\Theta_n$  can be updated recursively via

$$(2.10) \quad \Theta_{n+1} = \xi_{\Theta_n}(X_{n+1}) = \Theta_n + X_{n+1}.$$

In order to introduce uncertainty, it is assumed that with probability  $\varepsilon$  the increment  $W_n$  is replaced by an arbitrarily distributed outlier. This model yields the following family of conditional uncertainty sets

$$(2.11) \quad \mathcal{P}_\theta = \{P \in \mathcal{M}_\mu : P(\mathcal{E}) = (1 - \varepsilon)P_W(\mathcal{E} - a\theta) + \varepsilon H(\mathcal{E}), H \in \mathcal{M}_\mu\},$$

where  $\mathcal{E} \in \mathcal{F}_X$  and  $\mathcal{E} - a\theta$  is shorthand for  $\{x \in \Omega_X : x + a\theta \in \mathcal{E}\}$ . In Section 8 a variant of this example is used to illustrate the design of a minimax optimal test with dependencies in the underlying stochastic process.

*2.4. Testing policies and test statistics.* A sequential test is specified via two sequences of randomized decision rules,  $(\psi_n)_{n \geq 1}$  and  $(\delta_n)_{n \geq 1}$ , that are adapted to the filtration  $(\mathcal{F}_X^n)_{n \geq 1}$ . Each  $\psi_n : \Omega_X^n \rightarrow [0, 1]$  denotes the probability of stopping at time instant  $n$ . Each  $\delta_n : \Omega_X^n \rightarrow \Delta^K$  is a  $K$ -dimensional vector,  $\delta_n = (\delta_{1,n}, \dots, \delta_{K,n})$ , whose  $k$ th element denotes the probability of deciding for  $\mathcal{H}_k$ , given that the test has stopped at time instant  $n$ . The randomization is assumed to be performed by independently drawing from a Bernoulli distribution with success probability  $\psi_n$  and a discrete distribution on  $\{1, \dots, K\}$  with associated probabilities  $(\delta_{1,n}, \dots, \delta_{K,n})$ , respectively. The set of randomized  $K$ -dimensional decision rules defined on  $(\Omega_X^n, \mathcal{F}_X^n)$  is denoted by  $\Delta_n^K$ . The stopping time of the test is denoted by  $\tau = \tau(\psi)$ .

For the sake of a more concise notation, let  $\pi = (\pi_n)_{n \geq 1}$  with  $\pi_n = (\psi_n, \delta_n) \in \Delta_n^1 \times \Delta_n^K$ ; denote a sequence of tuples of stopping and decision rules. In what follows,  $\pi$  is referred to as a *testing policy*, and the set of all feasible policies is denoted by  $\Pi := \times_{n \geq 1} (\Delta_n^1 \times \Delta_n^K)$ .

A test statistic is a stochastic process,  $(T_n)_{n \geq 0}$ , that is adapted to the filtration  $(\mathcal{F}_X^n)_{n \geq 1}$  and allows the stopping and decision rules to be defined as functions mapping from the codomain of  $T_n$  to the unit interval. Of particular importance for this paper is the case where the sequence of test statistics,  $(T_n)_{n \geq 0}$ , is itself a time-homogeneous Markov process and the stopping and decision rules are independent of the time index  $n$ . The corresponding testing policies are in the following referred to as time homogeneous. This property is formalized in the following definition:

DEFINITION 1. A policy  $\pi \in \Pi$  is referred to as *time homogeneous* if there exists a  $(\Omega_T, \mathcal{F}_T)$ -valued stochastic process  $(T_n)_{n \geq 0}$  that is adapted to the filtration  $(\mathcal{F}_X^n)_{n \geq 1}$ , and it holds that

$$(2.12) \quad \psi_n = \psi(T_n) \quad \text{and} \quad \delta_n = \delta(T_n),$$

where the functions  $\psi : \Omega_T \rightarrow [0, 1]$  and  $\delta : \Omega_T \rightarrow [0, 1]^K$  are independent of the index  $n$ .

Focusing on time-homogeneous policies significantly simplifies both the derivation and the presentation of the main results. Moreover, it will become clear in the course of the paper that optimal tests for time-homogeneous Markov processes can always be realized with time-homogeneous policies so that this restriction can be made without sacrificing optimality.

2.5. *Performance metrics and problem formulation.* The performance metrics considered in this paper are the probability of erroneously rejecting the  $k$ th hypothesis,  $\alpha_k$ , and the expected run length of the sequential test,  $\gamma$ . Both are defined as functions of the testing policy and the true distribution:

$$(2.13) \quad \gamma(\pi, \mathbb{P}) := E_{\pi, \mathbb{P}}[\tau(\psi)],$$

$$(2.14) \quad \alpha_k(\pi, \mathbb{P}) := E_{\pi, \mathbb{P}}[1 - \delta_{k, \tau}],$$

with  $k = 1, \dots, K$ . Here,  $E_{\pi, \mathbb{P}}$  denotes the expected value taken jointly with respect to the distribution of  $(X_n)_{n \geq 1}$  and the randomized policy  $\pi$ . A generalization to performance metrics that are defined in terms of the pairwise error probabilities is possible but would considerably complicate notation while adding little conceptual insight.

It is important to note that for the design of robust sequential tests the error probabilities and the expected run length need to be treated as equally important performance metrics. On the one hand, reducing the sample size is typically the reason for using sequential tests in the first place. On the other hand, a test whose error probabilities remain bounded over a given uncertainty set but whose expected run length can increase arbitrarily, cannot be considered robust. In other words, a robust test should not be allowed to delay a decision indefinitely in order to avoid making a wrong decision.

The first optimality criterion considered in this paper is the weighted sum cost, that is,

$$(2.15) \quad L_\lambda(\pi, \mathbb{P}) = \gamma(\pi, \mathbb{P}_0) + \sum_{k=1}^K \lambda_k \alpha_k(\pi, \mathbb{P}_k),$$

where  $\mathbb{P} = (\mathbb{P}_0, \dots, \mathbb{P}_K)$  denotes a  $K + 1$  dimensional vector of distributions and  $\lambda = (\lambda_1, \dots, \lambda_K)$  denotes a  $K$  dimensional vector of nonnegative cost coefficients. The minimax problem corresponding to the cost function in (2.15) reads as

$$(2.16) \quad \inf_{\pi \in \Pi} \sup_{\mathbb{P} \in \mathcal{P}} L_\lambda(\pi, \mathbb{P}),$$

where  $\mathbb{P} \in \mathcal{P}$  is used as a compact notation for  $\mathbb{P}_k \in \mathcal{P}_k, k = 0, \dots, K$ .

The second optimality criterion is the expected run length under constraints on the error probabilities. The corresponding minimax problem reads as

$$(2.17) \quad \inf_{\pi \in \Pi} \sup_{\mathbb{P}_0 \in \mathcal{P}_0} \gamma(\pi, \mathbb{P}_0) \quad \text{s.t.} \quad \sup_{\mathbb{P}_k \in \mathcal{P}_k} \alpha_k(\pi, \mathbb{P}_k) \leq \bar{\alpha}_k,$$

where the constraint holds for all  $k = 1, \dots, K$  and  $\bar{\alpha}_k$  denotes an upper bound on the probability of erroneously deciding against  $\mathcal{H}_k$ . The notation for the minimax optimal policies is fixed below and concludes the section.

DEFINITION 2. The set of time-homogeneous policies that are optimal in the sense of (2.16) is denoted by  $\Pi_\lambda^*(\mathcal{P})$ . The set of time-homogeneous policies that are optimal in the sense of (2.17) is denoted by  $\Pi_{\bar{\alpha}}^*(\mathcal{P})$ .

**3. Optimal tests.** Assume that the distributions  $\mathbb{P}_0, \dots, \mathbb{P}_K \in \mathbb{M}$  are given and fixed. In this case, problems (2.16) and (2.17) reduce to the design of an optimal test for  $K$  simple hypotheses, that is,

$$(3.1) \quad \inf_{\pi \in \Pi} L_\lambda(\pi, \mathbb{P})$$

and

$$(3.2) \quad \inf_{\pi \in \Pi} \gamma(\pi, \mathbb{P}_0) \quad \text{s.t.} \quad \alpha_k(\pi, \mathbb{P}_k) \leq \bar{\alpha}_k.$$

The notation for the corresponding optimal policies is fixed in the next definition.

DEFINITION 3. The set of time-homogeneous policies that are optimal in the sense of (3.1) is denoted by  $\Pi_\lambda^*(\mathbb{P})$ . The set of time-homogeneous policies that are optimal in the sense of (3.2) is denoted by  $\Pi_{\bar{\alpha}}^*(\mathbb{P})$ .

The solutions to both the unconstrained problem (3.1) and the constrained problem (3.2) can be found in the literature. The binary case ( $K = 2$ ) was treated in [14] under the same assumptions as stated in Section 2.2. In [41], the general solution for an arbitrary number of hypotheses and arbitrary underlying stochastic processes is derived.

For easier reference the solution of (3.1) is restated in this section. To this end, the functions  $G_\lambda^{(k)}, g_\lambda: \mathbb{R}_{\geq 0}^{K+1} \rightarrow \mathbb{R}_{\geq 0}$  are introduced. Let

$$(3.3) \quad G_\lambda^{(k)}(\mathbf{z}) := \sum_{i=1, i \neq k}^K \lambda_i z_i = \left( \sum_{i=1}^K \lambda_i z_i \right) - \lambda_k z_k$$

and

$$(3.4) \quad g_\lambda(\mathbf{z}) := \min_{k=1, \dots, K} G_\lambda^{(k)}(\mathbf{z}) = \left( \sum_{i=1}^K \lambda_i z_i \right) - \max_{k=1, \dots, K} \lambda_k z_k,$$

where  $\mathbf{z} = (z_0, \dots, z_K) \in \mathbb{R}_{\geq 0}^{K+1}$  and  $\lambda \in \mathbb{R}_{\geq 0}^K$  is the vector of cost coefficients introduced in Section 2.5. Note that both  $G_\lambda^{(k)}$  and  $g_\lambda$  are independent of  $z_0$ ; defining them as functions of the  $K + 1$  dimensional vector  $\mathbf{z}$  unifies the notation in what follows.

The cost function that characterizes the optimal test is stated in the following theorem. It extends Theorem 2.1 in [14] to multiple hypotheses.

THEOREM 1. Let  $\lambda \geq 0$ , let  $\mathbb{P} \in \mathbb{M}^{K+1}$ , and let  $\rho_\lambda : \Omega_\rho \rightarrow \mathbb{R}_{\geq 0}$ , where

$$(3.5) \quad \Omega_\rho := \mathbb{R}_{\geq 0}^{K+1} \times \Omega_\Theta.$$

The integral equation

$$(3.6) \quad \rho_\lambda(\mathbf{z}, \theta) = \min \left\{ g_\lambda(\mathbf{z}), z_0 + \int \rho_\lambda(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) \mu(dx) \right\},$$

with  $g_\lambda$  defined in (3.4), has a unique solution and it holds that

$$(3.7) \quad \inf_{\pi \in \Pi} L_\lambda(\pi, \mathbb{P}) = \rho_\lambda(\mathbf{1}, \theta_0).$$

Theorem 1 follows directly from Theorem 5 and Lemma 6 in [41] and the Markov property of the stochastic process  $(X_n)_{n \geq 1}$ . Also, compare Theorem 5 in [13]. The optimal test statistic and testing policies are obtained by comparing the cost for stopping with the expected cost for continuing under the optimal policy.

COROLLARY 1. The optimal test statistic of a test solving (3.1) is given by

$$(3.8) \quad T_n(x_1, \dots, x_n) = (z_n, \theta_n) \in \Omega_\rho,$$

where  $\theta_n$  is a sufficient statistic for  $(x_1, \dots, x_n)$  in the sense of (2.4) and  $\mathbf{z}_n = (z_{0,n}, \dots, z_{K,n})$  is the vector of likelihood ratios (Radon–Nikodym derivatives)

$$(3.9) \quad z_{k,n} = \prod_{i=1}^n \frac{dP_{\theta_{i-1}}^{(k)}}{d\mu}(x_i) = \prod_{i=1}^n p_{\theta_{i-1}}^{(k)}(x_i).$$

By Assumption 2 in Section 2.2, the test statistic can be calculated recursively via

$$(3.10) \quad \theta_{n+1} = \xi_{\theta_n}(x_{n+1}), \quad \theta_0: \text{given a priori},$$

$$(3.11) \quad \mathbf{z}_{n+1} = \mathbf{z}_n \mathbf{p}_{\theta_n}(x_{n+1}), \quad \mathbf{z}_0 = \mathbf{1}$$

with  $\xi_\theta$  is defined in (2.5).

COROLLARY 2. Let  $\lambda \geq 0$ , let  $\mathbb{P} \in \mathbb{M}^{K+1}$ , and let  $\rho_\lambda$  be as defined in Theorem 1. A policy  $\pi$  is time homogeneous in the sense of Definition 1 and optimal in the sense of (3.1), that is,  $\pi \in \Pi_\lambda^*(\mathbb{P})$ , if and only if its stopping and decision rules are of the form

$$(3.12) \quad \psi_n = \psi(\mathbf{z}_n, \theta_n) \quad \text{and} \quad \delta_{k,n} = \delta_k(\mathbf{z}_n),$$

where

$$(3.13) \quad \mathcal{I}(\{g_\lambda(\mathbf{z}) < \rho_\lambda(\mathbf{z}, \theta)\}) \leq \psi(\mathbf{z}, \theta) \leq \mathcal{I}(\{g_\lambda(\mathbf{z}) \leq \rho_\lambda(\mathbf{z}, \theta)\}),$$

$$(3.14) \quad \delta_k(\mathbf{z}) \leq \mathcal{I}(\{G_\lambda^{(k)}(\mathbf{z}) = g_\lambda(\mathbf{z})\})$$

and  $(\mathbf{z}_n, \theta_n)$  is defined in Corollary 1.

Corollary 1 and Corollary 2 follow immediately from Theorem 6 in [41] and characterize the set of optimal time-homogeneous policies under Markov assumptions. The two components of the optimal test statistic  $(\mathbf{z}_n, \theta_n)$  correspond to the two types of information that are necessary to apply the optimal stopping rule. The likelihood ratios  $\mathbf{z}$  are needed to evaluate the cost for stopping, while the state of the Markov process  $\theta$  is needed to evaluate the conditional expectation that determines the cost for continuing.

The policies defined in Corollary 2 are optimal in the sense of the unconstrained problem (3.1). The solution of the constrained problem is closely related, but its statement is deferred



to the next section since it relies on properties of the optimal cost function  $\rho_\lambda$  that need to be established first. Before turning to the latter, a more compact notation and a useful way of characterizing the performance of time-homogeneous tests of the form (3.8) are introduced.

For policies of the form given in Corollaries 1 and 2, it is convenient to define the expected run length and the error probabilities of the underlying test as functions of the initial state of the test statistic, that is,

$$(3.15) \quad \gamma_{\pi, \mathbb{P}}(\mathbf{z}, \theta) := E_{\pi, \mathbb{P}}[\tau(\psi) \mid \mathbf{Z}_0 = \mathbf{z}, \Theta_0 = \theta],$$

$$(3.16) \quad \alpha_{\pi, \mathbb{P}}^{(k)}(\mathbf{z}, \theta) := E_{\pi, \mathbb{P}}[1 - \delta_{\tau, k} \mid \mathbf{Z}_0 = \mathbf{z}, \Theta_0 = \theta],$$

where  $\gamma_{\pi, \mathbb{P}}: \Omega_\rho \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$  and  $\alpha_{\pi, \mathbb{P}}: \Omega_\rho \rightarrow [0, 1]$ . Since  $(\mathbf{Z}_n, \Theta_n)_{n \geq 0}$  is a time-homogeneous Markov process,  $\gamma_{\pi, \mathbb{P}}$  and  $\alpha_{\pi, \mathbb{P}}^{(k)}$  solve the Chapman–Kolmogorov equations [28]

$$(3.17) \quad \gamma_{\pi, \mathbb{P}}(\mathbf{z}, \theta) = (1 - \psi(\mathbf{z}, \theta)) \left( 1 + \int \gamma_{\pi, \mathbb{P}}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) P_\theta(dx) \right)$$

and

$$(3.18) \quad \begin{aligned} \alpha_{\pi, \mathbb{P}}^{(k)}(\mathbf{z}, \theta) &= \psi(\mathbf{z}, \theta)(1 - \delta_k(\mathbf{z}, \theta)) \\ &+ (1 - \psi(\mathbf{z}, \theta)) \int \alpha_{\pi, \mathbb{P}}^{(k)}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) P_\theta(dx). \end{aligned}$$

Conditioning on the true initial states of the statistics reduces the conditional performance metrics in (3.15) and (3.16) to the unconditional ones in (2.13) and (2.14), that is,

$$(3.19) \quad \gamma_{\pi, \mathbb{P}}(\mathbf{1}, \theta_0) = \gamma(\pi, \mathbb{P}) \quad \text{and} \quad \alpha_{\pi, \mathbb{P}}^{(k)}(\mathbf{1}, \theta_0) = \alpha_k(\pi, \mathbb{P}).$$

Being able to characterize the performance of a test in terms of solutions of (3.17) and (3.18) is of critical importance for the proofs given in later sections.

In order to simplify the notation of the central integral equations of this section, let  $\{\mu_{\mathbf{z}, \theta} : (\mathbf{z}, \theta) \in \Omega_\rho\}$  and  $\{P_{\mathbf{z}, \theta} : (\mathbf{z}, \theta) \in \Omega_\rho\}$  be two families of probability measures on  $\Omega_\rho$  that are defined via

$$(3.20) \quad \mu_{\mathbf{z}, \theta}(\mathcal{E}_z \times \mathcal{E}_\theta) := \mu(\{x \in \Omega_X : \mathbf{z} \mathbf{p}_\theta(x) \in \mathcal{E}_z, \xi_\theta(x) \in \mathcal{E}_\theta\}),$$

$$(3.21) \quad P_{\mathbf{z}, \theta}(\mathcal{E}_z \times \mathcal{E}_\theta) := P_\theta(\{x \in \Omega_X : \mathbf{z} \mathbf{p}_\theta(x) \in \mathcal{E}_z, \xi_\theta(x) \in \mathcal{E}_\theta\}),$$

where  $\mathcal{E}_z \times \mathcal{E}_\theta \in \mathcal{F}_\rho$  with  $\mathcal{F}_\rho$  denoting the natural  $\sigma$ -algebra on  $\Omega_\rho$ . The notation  $P_{\mathbf{z}, \theta}^{(k)}$  is used to refer to the probability measure in (3.21) with  $P_\theta$  chosen to be  $P_\theta^{(k)}$ . Using this notation, the integral equations in (3.6), (3.17) and (3.18) can be written more compactly as

$$(3.22) \quad \rho_\lambda = \min \left\{ g_\lambda, z_0 + \int \rho_\lambda d\mu_{\mathbf{z}, \theta} \right\}$$

and

$$(3.23) \quad \gamma_{\pi, \mathbb{P}} = (1 - \psi) \left( 1 + \int \gamma_{\pi, \mathbb{P}} dP_{\mathbf{z}, \theta} \right),$$

$$(3.24) \quad \alpha_{\pi, \mathbb{P}}^{(k)} = \psi(1 - \delta_k) + (1 - \psi) \int \alpha_{\pi, \mathbb{P}}^{(k)} dP_{\mathbf{z}, \theta}.$$

Both notations are used in what follows.

**4. Properties of the cost function  $\rho_\lambda$ .** While Theorem 1 is well known in the literature, the cost function  $\rho_\lambda$  in (3.6) has rarely been studied in detail. The connection between  $\rho_\lambda$  and the properties of sequential tests using policies  $\pi \in \Pi_\lambda^*(\mathbb{P})$  is the subject of this section. It extends and generalizes the results in Section 3 of [14].

**THEOREM 2.** *Let  $\rho_\lambda$  be as defined in Theorem 1, and let  $\Pi_\lambda^*(\mathbb{P})$  be as defined in Definition 3. For all  $\pi \in \Pi_\lambda^*(\mathbb{P})$ , it holds that*

$$(4.1) \quad \rho_\lambda(\mathbf{z}, \theta) = z_0 \gamma_{\pi, \mathbb{P}_0}(\mathbf{z}, \theta) + \sum_{k=1}^K \lambda_k z_k \alpha_{\pi, \mathbb{P}_k}^{(k)}(\mathbf{z}, \theta),$$

where  $\gamma_{\pi, \mathbb{P}}$  and  $\alpha_{\pi, \mathbb{P}}^{(k)}$  are defined in (3.15) and (3.16), respectively.

See Appendix A in the Supplementary Material [17] for a proof. Theorem 2 connects the performance metrics in Section 2.5 to the optimal cost function  $\rho_\lambda$  and is key to solving the unconstrained minimax problem (2.16). However, obtaining a solution to the constrained minimax problem (2.17) additionally requires control over the individual error probabilities. In what follows, a connection between the latter and the partial generalized derivatives of  $\rho_\lambda$  is established. First, some useful technical properties of  $\rho_\lambda$  are shown.

**LEMMA 1.** *For all  $\lambda \geq 0$  and all  $\mathbb{P} \in \mathbb{M}^{K+1}$ , the function  $\rho_\lambda$  that solves (3.6) is nondecreasing, concave and homogeneous of degree one in  $\mathbf{z}$ .*

A proof is detailed in Appendix B in the Supplementary Material [17]. Lemma 1 is significant for two reasons. First, it ensures that  $\rho_\lambda$  admits a generalized differential, a property that is used in the next theorem to establish a connection between  $\rho_\lambda$  and the error probabilities  $\alpha_k$ . Second, being concave and homogeneous qualifies  $\rho_\lambda$  as an  $f$ -dissimilarity, that is, a statistical measure for the joint similarity of  $P_\theta^{(0)}, \dots, P_\theta^{(K)}$ . This property of  $\rho_\lambda$  is not used explicitly in the derivations of the minimax optimal sequential test, but it puts the results into context and will later be shown to provide a unified interpretation of minimax optimal sequential and minimax optimal fixed sample size test. A more detailed discussion of this aspect is deferred to Section 7.

**THEOREM 3.** *Let  $\mathbb{P} \in \mathbb{M}^{K+1}$ , and let  $\Pi_\lambda^*(\mathbb{P})$  be as defined in Definition 3. For  $\rho_\lambda$ , as defined in Theorem 1, and  $\gamma_{\pi, \mathbb{P}}, \alpha_{\pi, \mathbb{P}}^{(k)}$ , as defined in (3.15) and (3.16), respectively, it holds that:*

1. For all  $\pi \in \Pi_\lambda^*(\mathbb{P})$  and all  $k = 1, \dots, K$

$$(4.2) \quad \gamma_{\pi, \mathbb{P}_0} \in \partial_{z_0} \rho_\lambda,$$

$$(4.3) \quad \lambda_k \alpha_{\pi, \mathbb{P}_k}^{(k)} \in \partial_{z_k} \rho_\lambda.$$

2. For all  $\pi \in \Pi_\lambda^*(\mathbb{P})$ , all  $(\mathbf{z}, \theta) \in \Omega_\rho$ , and all  $k = 1, \dots, K$

$$(4.4) \quad \{\gamma_{\pi, \mathbb{P}_0}(\mathbf{z}, \theta) : \pi \in \Pi_\lambda^*(\mathbb{P})\} = \partial_{z_0} \rho_\lambda(\mathbf{z}, \theta),$$

$$(4.5) \quad \{\lambda_k \alpha_{\pi, \mathbb{P}_k}^{(k)}(\mathbf{z}, \theta) : \pi \in \Pi_\lambda^*(\mathbb{P})\} = \partial_{z_k} \rho_\lambda(\mathbf{z}, \theta).$$

Theorem 3 is proven in Appendix C in the Supplementary Material [17]. Its two parts correspond to a *global* and a *local* statement about the generalized differentials of  $\rho_\lambda$ . The first part states that for all  $\pi \in \Pi_\lambda^*(\mathbb{P})$ , the functions  $\gamma_{\pi, \mathbb{P}_0}$  and  $\alpha_{\pi, \mathbb{P}_k}^{(k)}$  are valid generalized

differentials of  $\rho_\lambda$ . The second part states that at every point  $(z, \theta) \in \Omega_\rho$ , the generalized differential of  $\rho_\lambda$  coincides with the set of all error probabilities that can be realized by policies  $\pi \in \Pi_\lambda^*(\mathbb{P})$ . Note that since the optimal policy is deterministic on the interior of the decision regions, the remaining degrees of freedom in terms of the error probabilities are exclusively due to the randomization on the boundary of the stopping and decision regions. That is, the subdifferentials on the right-hand sides of (4.4) and (4.5) are spanned by policies with different randomizations on the left-hand side. If only deterministic policies are allowed, the result only holds for the end points of the interval, that is, the left and right derivative.

The local statement in Theorem 3 cannot be extended to a global statement since the integral equations (C.4) and (C.5) in the Supplementary Material [17] establish a coupling between the local differentials so that they cannot be chosen independently of each other. This coupling reflects the fact that changing the randomization on the boundaries of a decision region also affects the overall performance of the corresponding sequential test.

Based on Theorem 3, the following optimality result can be obtained for the constrained sequential testing problem in (3.2).

**THEOREM 4.** *Let  $\mathbb{P} \in \mathbb{M}^{K+1}$ , and let  $\pi^* \in \Pi_{\bar{\alpha}}^*(\mathbb{P})$ , that is,  $\pi^*$  solves (3.2). If*

$$(4.6) \quad \lambda^* \in \arg \max_{\lambda \in \mathbb{R}_{\geq 0}^K} \left\{ \rho_\lambda(\mathbf{1}, \theta_0) - \sum_{k=1}^K \lambda_k \bar{\alpha}_k \right\},$$

with  $\rho_\lambda$  defined in Theorem 1, then for all  $\pi \in \Pi_{\lambda^*}^*(\mathbb{P})$  and all  $k = 1, \dots, K$  it holds that

$$(4.7) \quad |\gamma(\pi, \mathbb{P}_0) - \gamma(\pi^*, \mathbb{P}_0)| \leq |\partial_{z_0} \rho_{\lambda^*}(\mathbf{1}, \theta_0)|,$$

$$(4.8) \quad \lambda_k^* |\alpha_k(\pi, \mathbb{P}_k) - \bar{\alpha}_k| \leq |\partial_{z_k} \rho_{\lambda^*}(\mathbf{1}, \theta_0)|.$$

Moreover, it holds that

$$(4.9) \quad \Pi_{\lambda^*}^*(\mathbb{P}) \cap \Pi_{\bar{\alpha}}^*(\mathbb{P}) \neq \emptyset,$$

that is, there exists at least one  $\pi \in \Pi_{\lambda^*}^*(\mathbb{P})$  that is optimal in the sense of (3.2).

A proof is detailed in Appendix D in the Supplementary Material [17]. The significance of Theorem 4 lies in the fact that it characterizes solutions of the constrained problem (3.2) in terms of the optimal cost function of the unconstrained problem (3.1). From an algorithmic point of view, Theorem 4 makes it possible to design constrained sequential tests via a systematic optimization of the cost coefficients  $\lambda$  instead of Monte Carlo simulations or resampling techniques [54, 56].

Theorems 2, 3 and 4 form the basis for the derivation of minimax optimal tests, which will be characterized as optimal tests for least favorable distributions. The latter are introduced and discussed in the next section.

**5. Least favorable distributions.** The counterpart of the optimal testing problem investigated in the previous sections is the problem of determining the least favorable distributions for a given testing policy  $\pi$ . In this case the unconstrained problem (3.1) reduces to

$$(5.1) \quad \sup_{\mathbb{P} \in \mathcal{P}} L_\lambda(\pi, \mathbb{P}) = \sup_{\mathbb{P} \in \mathcal{P}} \left( \gamma(\pi, \mathbb{P}_0) + \sum_{k=1}^K \lambda_k \alpha_k(\pi, \mathbb{P}_k) \right).$$

Since the expected run length and the error probabilities are coupled only via the policy, the joint problem in (5.1) decouples into  $K + 1$  individual maximization problems

$$(5.2) \quad \sup_{\mathbb{P}_0 \in \mathcal{P}_0} \gamma(\pi, \mathbb{P}_0) \quad \text{and} \quad \sup_{\mathbb{P}_k \in \mathcal{P}_k} \alpha_k(\pi, \mathbb{P}_k), \quad k = 1, \dots, K$$

which can be solved independently.

For arbitrary stopping and decision rules, solving the problems in (5.2) exactly is challenging and, in general, the least favorable distributions depend on the time instant  $n$  as well as on the history of the random process. However, for time-homogeneous policies of the form (3.12), a more elegant solution can be obtained.

**THEOREM 5.** *Let  $\mathcal{P}$  be an uncertainty set of the form (2.6), and let  $\gamma_{\pi, \mathcal{P}} : \Omega_\rho \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$  and  $\alpha_{\pi, \mathcal{P}}^{(k)} : \Omega_\rho \rightarrow \mathbb{R}_{\geq 0}$ . For all testing policies  $\pi$  of the form (3.12) and all  $k = 1, \dots, K$ , it holds that the integral equations*

$$(5.3) \quad \gamma_{\pi, \mathcal{P}} = (1 - \psi) \left( 1 + \sup_{H \in \mathcal{P}_\theta} \int \gamma_{\pi, \mathcal{P}}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) H(dx) \right),$$

$$(5.4) \quad \alpha_{\pi, \mathcal{P}}^{(k)} = \psi(1 - \delta_k) + (1 - \psi) \left( \sup_{H \in \mathcal{P}_\theta} \int \alpha_{\pi, \mathcal{P}}^{(k)}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) H(dx) \right)$$

have unique solutions.

**THEOREM 6.** *Let  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_K)$  be uncertainty sets of the form (2.6), and let  $\pi$  be of the form (3.12).*

- *If it holds that for every  $(\mathbf{z}, \theta) \in \Omega_\rho$*

$$(5.5) \quad \mathcal{Q}_{\mathbf{z}, \theta}^{(0)} := \arg \max_{H \in \mathcal{P}_\theta^{(0)}} \int \gamma_{\pi, \mathcal{P}_0}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) H(dx) \neq \emptyset$$

with  $\gamma_{\pi, \mathcal{P}}$  defined in (5.3), then every distribution

$$(5.6) \quad \mathbb{Q}_0 \in \mathcal{Q}_0 := \left\{ \mathbb{P} \in \mathbb{M} : \mathbb{P} = \prod_{n \geq 0} P_{\mathbf{z}_n, \theta_n}, P_{\mathbf{z}_n, \theta_n} \in \mathcal{Q}_{\mathbf{z}_n, \theta_n}^{(0)} \right\}$$

is least favorable with respect to the expected run length of the test, that is,

$$(5.7) \quad \gamma(\pi, \mathbb{Q}_0) = \sup_{\mathbb{P} \in \mathcal{P}_0} \gamma(\pi, \mathbb{P}).$$

- *If it holds that for every  $(\mathbf{z}, \theta) \in \Omega_\rho$*

$$(5.8) \quad \mathcal{Q}_{\mathbf{z}, \theta}^{(k)} := \arg \max_{H \in \mathcal{P}_\theta^{(k)}} \int \alpha_{\pi, \mathcal{P}_k}^{(k)}(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) H(dx) \neq \emptyset$$

with  $k = 1, \dots, K$  and  $\alpha_{\pi, \mathcal{P}}^{(k)}$  defined in (5.4), then every distribution

$$(5.9) \quad \mathbb{Q}_k \in \mathcal{Q}_k := \left\{ \mathbb{P} \in \mathbb{M} : \mathbb{P} = \prod_{n \geq 0} P_{\mathbf{z}_n, \theta_n}, P_{\mathbf{z}_n, \theta_n} \in \mathcal{Q}_{\mathbf{z}_n, \theta_n}^{(k)} \right\}$$

is least favorable with respect to the probability of erroneously deciding against  $\mathcal{H}_k$ , that is,

$$(5.10) \quad \alpha_k(\pi, \mathbb{Q}_k) = \sup_{\mathbb{P} \in \mathcal{P}_k} \alpha_k(\pi, \mathbb{P}).$$

Theorems 5 and 6 are proven in Appendices E and F in the Supplementary Material [17], respectively. Note that the distributions whose densities are used to update the test statistic  $\mathbf{z} \mathbf{p}_\theta$  are not the least favorable distribution but are determined by the policy  $\pi$ ; compare Corollary 1. From Theorem 6 it follows that, under the least favorable distributions, the process  $(X_n)_{n \geq 1}$  is a Markov process with sufficient statistic  $(\mathbf{z}_n, \theta_n)$ . That is, the least favorable distributions adapt to the policy of the test as well as to the history of the stochastic process.

It is worth highlighting that even in the case where the  $X_n$  are independent, that is, no additional statistic  $\theta_n$  is required, the least favorable distributions still generate a Markov process with sufficient statistic  $z_n$ . That is, the least favorable distributions depend on the history of the process via the likelihood ratios. This aspect distinguishes robust sequential hypothesis tests from existing robust tests for fixed sample sizes or for change points. It is discussed in more detail in Section 7.

Having characterized optimal tests and least favorable distributions, everything is in place for the derivation of minimax optimal sequential tests.

**6. Minimax optimal sequential tests.** In this section sufficient conditions for strict minimax optimality of sequential hypothesis tests are given. Following the procedure for the optimal sequential test without uncertainty, the solutions of the unconstrained problem (2.16) are derived first and are then shown to contain a solution of the constrained problem (2.17).

The following three theorems are stated in sequence and constitute the main contribution of the paper. A discussion and an interpretation of the results are deferred to the next section.

**THEOREM 7.** *Let  $\lambda \geq 0$ , let  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_K)$  be uncertainty sets of the form (2.6), and let  $\rho_\lambda, d_\lambda : \Omega_\rho \rightarrow \mathbb{R}_{\geq 0}$  and  $D_\lambda : \Omega_\rho \times \mathcal{M}_\mu \rightarrow \mathbb{R}_{\geq 0}$ . The equation system*

$$(6.1) \quad \rho_\lambda(z, \theta) = \min\{g_\lambda(z), z_0 + d_\lambda(z, \theta)\},$$

$$(6.2) \quad d_\lambda(z, \theta) = \sup_{\mathbf{P} \in \mathcal{P}_\theta} D_\lambda(z, \theta; \mathbf{P}),$$

$$(6.3) \quad D_\lambda(z, \theta; \mathbf{P}) = \int \rho_\lambda(\mathbf{z} \mathbf{p}_\theta(x), \xi_\theta(x)) \mu(dx)$$

with  $g_\lambda$  defined in (3.4) has a unique solution.

**THEOREM 8.** *Let  $\lambda \geq 0$ , let  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_K)$  be uncertainty sets of the form (2.6), and let  $\rho_\lambda, d_\lambda$  and  $D_\lambda$  be as defined in Theorem 7. If for all  $(z, \theta) \in \Omega_\rho$*

$$(6.4) \quad \mathcal{Q}_{z,\theta} := \arg \max_{\mathbf{P} \in \mathcal{P}_\theta} D_\lambda(z, \theta; \mathbf{P}) \neq \emptyset,$$

then every policy  $\pi \in \Pi_\lambda^*(\mathcal{Q})$  with

$$(6.5) \quad \mathcal{Q} \in \mathcal{Q} = \left\{ \mathbf{P} \in \mathbb{M}^{K+1} : \mathbb{P}_k = \prod_{n \geq 0} P_{z_n, \theta_n}^{(k)}, \mathbf{P}_{z_n, \theta_n} \in \mathcal{Q}_{z_n, \theta_n} \right\}$$

is minimax optimal in the sense of (2.16), that is,

$$(6.6) \quad \{\pi \in \Pi_\lambda^*(\mathcal{Q}) : \mathcal{Q} \in \mathcal{Q}\} \subset \Pi_\lambda^*(\mathcal{P}).$$

**THEOREM 9.** *Let  $\mathcal{P} = (\mathcal{P}_0, \dots, \mathcal{P}_K)$  be uncertainty sets of the form (2.6), and let  $\pi^* \in \Pi_{\bar{\alpha}}^*(\mathcal{P})$  such that  $(\pi^*, \mathbb{P}^*)$  solves (2.17). If*

$$(6.7) \quad \lambda^* \in \arg \max_{\lambda \in \mathbb{R}_{\geq 0}^K} \left\{ \rho_\lambda(\mathbf{1}, \theta_0) - \sum_{k=1}^K \lambda_k \bar{\alpha}_k \right\}$$

with  $\rho_\lambda$  defined in Theorem 7, then for all  $\pi \in \Pi_{\lambda^*}^*(\mathcal{Q})$  with  $\mathcal{Q} \in \mathcal{Q}$  defined in Theorem 8 and all  $k = 1, \dots, K$ , it holds that

$$(6.8) \quad |\gamma(\pi, \mathcal{Q}_0) - \gamma(\pi^*, \mathbb{P}_0^*)| \leq |\partial_{z_0} \rho_{\lambda^*}(\mathbf{1}, \theta_0)|,$$

$$(6.9) \quad \lambda_k^* |\alpha_k(\pi, \mathcal{Q}_k) - \bar{\alpha}_k| \leq |\partial_{z_k} \rho_{\lambda^*}(\mathbf{1}, \theta_0)|.$$

Moreover, it holds that

$$(6.10) \quad \{\pi \in \Pi_{\lambda^*}^*(\mathcal{Q}) : \mathcal{Q} \in \mathcal{Q}\} \cap \Pi_{\lambda^*}^*(\mathcal{P}) \neq \emptyset,$$

that is, there exists at least one pair  $(\pi, \mathcal{Q})$ , with  $\mathcal{Q} \in \mathcal{Q}$  and  $\pi \in \Pi_{\lambda^*}^*(\mathcal{Q})$ , that is optimal in the sense of (2.17).

Theorems 7, 8 and 9 are proven in Appendices G, H and I in the Supplementary Material [17], respectively.

**7. Discussion.** Theorems 8 and 9 in the previous section provide sufficient conditions for the characterization of minimax optimal tests in terms of optimal testing policies and least favorable distributions. In this section the question of existence is discussed, and an interpretation in terms of statistical similarity measures is given that provides additional insight and establishes a connection to minimax optimal fixed samples size tests and change detection procedures.

*7.1. Statistical similarity measures.* In order to obtain a better conceptual understanding of minimax optimal sequential tests, it is helpful to introduce a class of statistical similarity measures known as  $f$ -dissimilarities. They were first proposed by Györfi and Nemetz [24–26] as an extension of  $f$ -divergences to multiple distributions and play an important role in the theory of statistical decision making. In particular, the connection between  $f$ -dissimilarities and Bayesian risks has been a topic of high interest in statistics [40], signal processing [58] and machine learning [48].

In this section it is shown that the function  $\rho_{\lambda}$  in Theorem 7 induces an  $f$ -dissimilarity and that this  $f$ -dissimilarity provides a sufficient characterization of the minimax optimal test. For this purpose a variation on the concept of  $f$ -dissimilarities is useful which is defined as follows:

**DEFINITION 4.** Let  $P_1, \dots, P_K$  be probability measures on a measurable space  $(\Omega, \mathcal{F})$ , and let  $f : \mathbb{R}_{\geq 0}^K \times \Omega \rightarrow \mathbb{R}$ , with  $f = f(\mathbf{y}, \omega) = f(y_1, \dots, y_K, \omega)$ , be homogeneous of degree one and concave in  $(y_1, \dots, y_K)$ . The functional

$$(7.1) \quad I_f(P_1, \dots, P_K) = \int f(p_1(\omega), \dots, p_K(\omega), \omega) \mu(d\omega)$$

is called  $f$ -similarity of  $P_1, \dots, P_K$ .

$I_f$  in (7.1) is referred to as a *similarity* measure since for concave and homogeneous functions  $f$ , the functional  $-I_f = I_{-f}$ , is a *dissimilarity* measure in the sense of Györfi and Nemetz [26], Definition 1. Allowing  $f$  to depend on the integration variable directly is a minor generalization that, in the context of this paper, allows the similarity measure to depend on the history of the random process. Similar generalizations have been introduced in the literature before [4, 43, 49].

Using Definition 4, an intuitive characterization of the family of least favorable distributions can be given in terms of a corresponding family of  $f$ -similarities.

**COROLLARY 3.** At every time instant  $n \geq 0$ , the least favorable distributions of  $X_{n+1}$ , conditioned on the state  $(\mathbf{Z}_n, \Theta_n) = (\mathbf{z}, \theta)$ , are the feasible distributions that are most similar with respect to the  $f$ -similarity defined by

$$(7.2) \quad f_{\mathbf{z}, \theta}(\mathbf{y}, x) = \rho_{\lambda}(\mathbf{z}\mathbf{y}, \xi_{\theta}(x))$$

with  $\rho_{\lambda}$  given in Theorem 7.

The family of  $f$ -divergences, defined by  $\{f_{z,\theta} : (z, \theta) \in \Omega_\rho\}$ , can be interpreted as follows: The equation system in Theorem 7 defines the optimal cost function  $\rho_\lambda$  which in turn defines the similarity measure in Corollary 3. The sequential aspect of the test is captured by the parameters  $(z, \theta)$ . The likelihood ratios  $z$  determine the weights of the individual densities. That is, the larger  $z_k$ , the larger the influence of the  $k$ th distribution on the similarity measure. In terms of the underlying hypothesis test, a high value of  $z_k$  implies that the test is likely to decide in favor of  $\mathcal{H}_k$ . Consequently, depending on whether  $\mathcal{H}_k$  is true or not, the least favorable distributions need to be as similar or dissimilar to  $\mathcal{H}_k$  as possible. On the other hand, a low value of  $z_k$  implies that a decision in favor of  $\mathcal{H}_k$  is highly unlikely so that the corresponding distribution contributes little to the overall similarity measure. The influence of the parameter  $\theta$ , that is, the history of the underlying stochastic process, does not affect the relative weighting of the individual distributions but rather the shape of  $\rho_\lambda$ . It strongly depends on the stochastic process itself, and it is less predictable than the effect of  $z$ .

The use of statistical similarity measures for the design of (sequential) tests for composite hypotheses has been suggested previously in the literature. However, as discussed in the [Introduction](#), in most works a suitable similarity measure is chosen *beforehand*—usually based on asymptotic results, bounds or approximations—and is used as a *surrogate* objective whose optimization is easier than solving the actual testing problem. Here, by contrast, it is shown that the minimax testing problem *induces* a similarity measure and that the distributions that maximize the latter solve the former.

It is also instructive to compare the minimax optimal tests in the previous section to tests whose least favorable distributions achieve Choquet capacities. In [15, 31], it is shown that in this case the pair of least favorable distributions jointly minimizes *all*  $f$ -divergences over the uncertainty sets. This implies that the least favorable distributions do not depend on the decision rule and can be calculated a priori; compare also [59]. For the minimax sequential test this is no longer the case. Instead, the least favorable distributions need to minimize a particular  $f$ -dissimilarity that depends on  $\rho_\lambda$  and the current state of the test statistic. Interestingly, the existence of capacity achieving distributions seems to be closely related to the existence of a single threshold test. Both fixed sample size detection and quickest change detection are problems for which optimal single threshold procedures exist and whose least favorable distributions have been shown to be of the capacity type [31, 57]. Sequential detection, on the other hand, is inherently a multithreshold procedure. This difference is subtle but turns out to be crucial. An intuitive explanation is as follows: For a single threshold test the least favorable distributions are those that maximize or minimize the drift of the test statistic towards the threshold. In the example of quickest change detection, maximizing the expected detection delay corresponds to minimizing the drift towards the threshold under the postchange distribution. On the other hand, in order to maximize the expected run length of a two-sided sequential test, one needs to minimize the drift of the test statistic towards *either* threshold. Depending on the data observed so far, this can translate to a drift towards the upper or towards the lower threshold; the least favorable distributions become data dependent. Similar considerations also apply to fixed sample size tests for multiple hypotheses which might explain why, as of now, no capacity based minimax solutions for the latter can be found in the literature.

In summary, the discussion shows that minimax optimal sequential tests are based the same principles that underpin existing minimax tests. Optimality is achieved by using a policy that leads to the *best separation* of the *most similar* distributions. The difference is that in the sequential case the similarity measure depends on the testing policy and on the observed data.

7.2. *Existence.* The results presented so far allow for some statements about the existence of minimax optimal tests. Stronger statements can be made for the unconstrained problem formulation (2.16) than for the constrained formulation (2.17). The former is considered first.

Since by Theorem 1 it holds that

$$\inf_{\pi \in \Pi} \sup_{\mathbb{P} \in \mathcal{P}} L_{\lambda}(\pi, \mathbb{P}) \leq g_{\lambda}(\mathbf{1}) \leq \sum_{k=1}^K \lambda_k,$$

the minimax optimal objective value in (2.16) is guaranteed to be finite for all cost coefficients  $\lambda \geq 0$  and all uncertainty sets  $\mathcal{P}$ . This includes scenarios where two uncertainty sets overlap or are identical. However, it does not imply that a testing policy exists that achieves this value. In order to guarantee (6.4), that is, that the supremum is attained, the uncertainty sets  $\mathcal{P}_1, \dots, \mathcal{P}_K$  need to be compact [51], Theorem 4.16. While this assumption seems restrictive in theory, it is automatically fulfilled when the problem is discretized in order to be solved numerically. Moreover, given that the least favorable distributions are maximizers of concave functionals, we conjecture that compactness is not necessary for the existence of minimax solutions. A more thorough discussion would need to go into the technical details of optimization on function spaces which is beyond the scope of this paper.

For the constrained problem (2.17) the situation is more involved, since, in order to be able to satisfy the constraints on the error probabilities, the uncertainty sets  $\mathcal{P}_1, \dots, \mathcal{P}_K$  need to be sufficiently *separated*, that is, the distance between the two sets needs to be large enough to statistically separate them using a finite number of samples. The appropriate way to measure this distance is via the  $f$ -similarity that is induced by the corresponding cost function  $\rho_{\lambda}$ . This leads to the following result:

**COROLLARY 4.** *The minimax optimal objective in (2.17) is finite, if and only if the right-hand side of (6.7) is bounded, that is, if*

$$(7.3) \quad \sup_{\lambda \in \mathbb{R}_{\geq 0}^K} \left\{ \rho_{\lambda}(\mathbf{1}, \theta_0) - \sum_{k=1}^K \lambda_k \bar{\alpha}_k \right\} < \infty$$

with  $\rho_{\lambda}$  defined in Theorem 7.

A proof of Corollary 4 is detailed in Appendix J in the Supplementary Material [17]. As in the unconstrained case, existence of a finite supremum does not imply that an optimal policy exists, unless the uncertainty sets are compact. Finally, note that no deterministic policy might exist that attains the bounds on the error probabilities with equality which implies that, in general, the optimal policy is randomized. In view of Theorem 3, this happens when the bounds are located in the interior of the subdifferentials of the optimal cost function.

Interestingly, the conditions for the existence of minimax optimal sequential tests are rather mild, especially for the unconstrained problem formulation. This is in contrast to the fixed sample size case for which much stricter sufficient conditions are given in the literature [15, 31]. This suggests the conjecture that these stricter conditions are only necessary to guarantee that the optimal policy and the least favorable distributions are *decoupled* and that minimax optimal fixed sample size tests for arbitrary uncertainty sets exist but that they require a joint design of the policy and the least favorable distributions. Unfortunately, applying the results presented in this paper to the fixed sample size case is not straightforward. First, the concept of ordered samples, which is essential to sequential hypothesis testing and induces state-dependent least favorable distributions, in general does not apply to fixed sample size tests. Second, owing to the deterministic stopping rule, the policies of fixed sample size tests are



not time homogeneous in the sense of Definition 1. However, the second numerical example given in the next section indicates that it might be possible to implement truncated tests with time-homogeneous policies; see Section 8.2 for more details.

Finally, it should be highlighted that the least favorable distributions do not depend on the reference measure  $\mu$ . More precisely, if  $\mu$  is absolutely continuous with respect to an alternative reference measure  $\tilde{\mu}$ , it follows from the homogeneity of  $\rho_\lambda$  that

$$\begin{aligned} D_\lambda(z, \theta; P) &= \int \rho_\lambda(z, \mathbf{p}, \xi_\theta) d\mu = \int \rho_\lambda(z, \mathbf{p}, \xi_\theta) \frac{d\mu}{d\tilde{\mu}} d\tilde{\mu} \\ &= \int \rho_\lambda\left(z, \mathbf{p} \frac{d\mu}{d\tilde{\mu}}, \xi_\theta\right) d\tilde{\mu} =: \int \rho_\lambda(z, \tilde{\mathbf{p}}, \xi_\theta) d\tilde{\mu}, \end{aligned}$$

where  $\tilde{\mathbf{p}} = (\tilde{p}_0, \dots, \tilde{p}_K)$  are probability densities with respect to  $\tilde{\mu}$ .

**8. Example and numerical results.** In order to illustrate the presented results, two examples for minimax optimal sequential tests are given in this section. First, a test for three hypotheses that is robust with respect to the error probabilities is designed under the assumption that  $(X_n)_{n \geq 1}$  is a sequence of independent random variables with identical uncertainty sets. Second, the case when the underlying processes admits dependencies is illustrated by solving the Kiefer–Weiss problem for a variant of the uncertainty model in Section 2.3. In both examples the targeted error probabilities are set to  $\bar{\alpha}_k = 0.01$  for all  $k = 1, \dots, K$ .

The test design is based on the following iterative procedure: First, all  $P_\theta^{(k)}$  are initialized to some feasible distribution. Keeping these distributions fixed, an optimal sequential test is designed by solving (4.6) for  $\lambda^*$  and  $\rho_{\lambda^*}$ . In the second step,  $\rho_{\lambda^*}$  is kept fixed, and the distributions are updated by solving the optimization problem in (6.4). Both steps are iterated until the changes in the function  $\rho_{\lambda^*}$  are small enough to assume convergence. For both examples convergence was assumed if the relative difference between two consecutive approximations of  $\rho_{\lambda^*}$  fell below  $10^{-3}$ . The question whether this procedure is guaranteed to converge in general is certainly worth investigating but beyond the scope of this paper.

It should be noted that this iterative procedure does *not* alternate between the design of optimal testing policies and least favorable distributions. Unless the procedure has converged, the distributions that solve (6.4) are not least favorable in the sense of Theorem 6. Moreover, the test statistic, which is part of the optimal policy, depends on the likelihood ratios and is hence affected by the update of the distributions in the second step.

In order to solve (4.6) and (6.4) numerically, both the state space  $\Omega_\rho$  and the sample space  $\Omega_X$  are discretized using a regularly spaced grid, and linear interpolation is used to evaluate functions between grid points. This straightforward approach works well for the examples presented here. However, if a larger number of hypotheses or more complex dependencies need to be considered, more sophisticated approximations need to be used [20, 36]. The linear programming algorithm detailed [14] was used to efficiently solve (4.6) jointly for  $\lambda^*$  and  $\rho_{\lambda^*}$ . However, in principle, any suitable convex optimization algorithm can be used to solve (4.6), and any method for solving nonlinear integral equations can be used to obtain  $\rho_{\lambda^*}$ .

In both examples the distributional uncertainty is of the density band type [32], that is,

$$(8.1) \quad \mathcal{P} = \{P \in \mathcal{M}_\mu : p' \leq p \leq p''\},$$

where  $0 \leq p' \leq p'' \leq \infty$ ,  $P'(\Omega_X) \leq 1$  and  $P''(\Omega_X) \geq 1$ . Here,  $P'$  and  $P''$  denote the measures corresponding to  $p'$  and  $p''$ , respectively. The reason for using this uncertainty model is twofold. First, it contains several popular uncertainty models as special cases, for example, the  $\varepsilon$ -contamination model [29], the bounded distribution function model [27, 42] and

the  $f$ -divergence ball model [16]. Second, an efficient iterative algorithm for the minimization of convex functionals of probability distribution under density band constraints exists that makes it possible to obtain accurate numerical solutions with moderate computational efforts. A more detailed discussion of the band model, its properties and how to obtain the least favorable distributions numerically can be found in [32] and [15].

8.1. *IID process under three hypotheses.* For the first example all  $X_n, n \geq 1$ , are assumed to be independent and distributed on the interval  $\Omega_X = [-1, 1]$ . Let  $P_n$  denote the distribution of  $X_n$ . The task is to decide between the following three hypotheses:

$$(8.2) \quad \mathcal{H}_1: P_n \in \mathcal{P}_1, \quad \mathcal{H}_2: P_n \in \mathcal{P}_2, \quad \mathcal{H}_3: P_n = \mathcal{U}_{[-1,1]}$$

for all  $n \geq 1$ . Here,  $\mathcal{U}_{[a,b]}$  denotes the continuous uniform distribution on the interval  $[a, b]$ , and the uncertainty sets  $\mathcal{P}_1, \mathcal{P}_2$  are of the form (8.1) with

$$(8.3) \quad p'_1(x) = ae^{-2x} + 0.1, \quad p''_1(x) = ae^{-2x} + 0.3,$$

$$(8.4) \quad p'_2(x) = ae^{2x} + 0.1, \quad p''_2(x) = ae^{2x} + 0.3,$$

where  $a \approx 0.1907$  was chosen such that  $P'_1(\Omega_X) = P'_2(\Omega_X) = 0.9$  and  $P''_1(\Omega_X) = P''_2(\Omega_X) = 1.1$ . The expected run length was minimized under  $\mathcal{H}_3$ , that is,  $P_n = \mathcal{U}_{[-1,1]}$  for all  $n$ . Moreover, in order to keep the domain of the cost function two-dimensional, the reference measure  $\mu$  was set to  $\mu = \mathcal{U}_{[-1,1]}$  so that  $z_0 = z_3 = 1$  and  $\rho_\lambda$  becomes a function of  $(z_1, z_2)$  only.

The scenario in this example can arise, for instance, in monitoring applications, where  $\mathcal{H}_3$  corresponds to an “in control” state in which the distribution of the data is known almost exactly, while  $\mathcal{H}_1$  and  $\mathcal{H}_2$  correspond to two different “out of control” states with only partially known distributions. If it needs to be established that the system is “in control” before a certain procedure starts, it is reasonable to minimize the expected run length of the test under the “in control” distribution while still requiring it to be insensitive to distributional uncertainties in the “out of control” states.

In order to solve this example numerically, the likelihood ratio plane  $(z_1, z_2)$  was discretized on  $[-20, 10] \times [-20, 10]$  using  $301 \times 301$  uniformly spaced grid points, and the sample space  $\Omega_X = [-1, 1]$  was discretized using 101 uniformly spaced grid points. The design procedure detailed above converged after five iterations. The optimal weights were found to be  $\lambda^* \approx (133.41, 133.41, 45.41)$ . The resulting cost function  $\rho_{\lambda^*}$ , as well as the corresponding testing policy, are depicted in Figure 1. While the cost function as such provides little insight, the testing policy lends itself to an intuitive interpretation. In analogy to the regular sequential probability ratio test (SPRT), the minimax optimal test consists of two corridors that correspond to a binary test between  $\mathcal{H}_{\{1,2\}}$  and  $\mathcal{H}_3$ , respectively. Interestingly, there is a rather sharp intersection of the two corridors so that the test quickly reduces to a quasi-binary scenario.

The expected run length and the error probabilities as functions of the state of the test statistic can be obtained either via the partial derivatives of  $\rho_{\lambda^*}$  or by solving the integral equations (3.16) and (2.13) and are depicted in Figure 2. The “blocky” appearance of some of the functions is due to them having being downsampled to a coarser grid for plotting. Moreover, no smoothing was applied in order not to smear the hard transitions between the decision regions. Finally, note that the plots are oriented differently to provide a better visual representation of the respective function.

The stopping and decision rules in Figure 1 are depicted as functions of the log-likelihood ratios. The latter are in turn defined in terms of the least favorable distributions, that is, the distributions that solve the maximization in (6.4). Four examples of densities of least favorable distributions are depicted in Figure 3. As can be seen, the densities change significantly,




FIG. 1. Segment of the optimal cost function  $\rho_\lambda$  in logarithmic scale and the corresponding testing policy as a function of the log-likelihood ratios.

depending on the state of the test statistic. In the top left plot the test statistic is in its initial state, meaning that there is no preference for either hypothesis. Consequently, the least favorable densities are chosen such that all three distributions are equally similar to each other which in this case implies that they are symmetric around the y-axis and that  $q_z^{(1)}$  and  $q_z^{(2)}$  jointly mimic the uniform density  $p_z^{(3)}$ . Also, note that  $q_z^{(1)}$  and  $q_z^{(2)}$  overlap on an interval around  $x = 0$  so that observations in this interval are statistically indistinguishable under  $\mathcal{H}_1$  and  $\mathcal{H}_2$ . As the test statistic is updated, the least favorable distributions change. In the upper right and the lower left plot of Figure 3, two cases are depicted where the test has a strong preference for  $\mathcal{H}_1$  or  $\mathcal{H}_2$ , respectively; compare the decision regions in Figure 1. In both cases the least favorable densities are no longer symmetric, but their probability masses are shifted, their tail behavior is noticeably different and the interval of overlap can no longer be




FIG. 2. Performance metrics as functions of the log-likelihood ratios. Clockwise from the top left: expected run length, error probability of the first, second and third type.




FIG. 3. Examples of least favorable distributions for different states of the test statistic.

observed. Finally, in the lower right plot, there is a strong preference for  $\mathcal{H}_3$  which leads to  $q_z^{(1)}$  and  $q_z^{(2)}$  both shifting as much probability mass as possible to their tails in order to reduce the significance of the corresponding observations. It is interesting to observe the effect that an imminent decision for  $\mathcal{H}_3$  has on  $q_z^{(1)}$  and  $q_z^{(2)}$ , namely, that they become less similar to each other in order to increase the joint similarity to  $p_z^{(3)}$ . This is in contrast to the initial state depicted in the upper left plot, where  $q_z^{(1)}$  and  $q_z^{(2)}$  also try to approximate  $p_z^{(3)}$  but at the same time need to be similar to each other as well.

In order to verify the numerical results, 10,000 Monte Carlo simulations were performed using the testing policy depicted in Figure 1. The observations were drawn from the least favorable distributions which were calculated on the fly by solving (6.4) for the current weights  $(z_1, z_2)$ . The resulting confusion matrix and the average run length of the tests are shown in Table 1.

TABLE I  
Results of  $10^4$  Monte Carlo runs of the minimax optimal test using the policy depicted in Figure 1. The target detection probability is 99%, the theoretical expected run length under  $\mathcal{H}_3$  is  $E_{\mathbb{P}_3}[\tau(\psi)] \approx 36.31$  samples

True hypothesis	Decisions in %			Average run-length
	$\mathcal{H}_1$	$\mathcal{H}_2$	$\mathcal{H}_3$	
$\mathcal{H}_1$	99.22	0.00	0.78	25.46
$\mathcal{H}_2$	0.00	98.83	1.17	26.84
$\mathcal{H}_3$	0.61	0.35	99.04	36.04

8.2. *Binomial AR(1) process under two hypotheses.* For the second example a discrete version of the uncertainty model in Section 2.3 is considered, namely, a binomial AR(1) process [63] which has applications, for example, in finance and monitoring [64]. Due to space constraints, the example is not detailed here but can be found in [17].

**Acknowledgments.** The authors would like to sincerely thank the Associate Editor and the anonymous reviewers for their constructive feedback which greatly helped to improve the quality of the paper.

The first author was supported in part by the German Research Foundation (DFG) under Grant number 390542458, and in part by the U.S. National Science Foundation under Grant CCF-1908308.

## SUPPLEMENTARY MATERIAL

**Supplement to “Minimax optimal sequential hypothesis tests for Markov processes”** (DOI: [10.1214/19-AOS1899SUPP](https://doi.org/10.1214/19-AOS1899SUPP); .pdf). The supplement contains an additional numerical example and the technical proofs.

## REFERENCES

- [1] AVELLA MEDINA, M. and RONCHETTI, E. (2015). Robust statistics: A selective overview and new directions. *Wiley Interdiscip. Rev.: Comput. Stat.* **7** 372–393. MR3415420 <https://doi.org/10.1002/wics.1363>
- [2] BANERJEE, T. and VEERAVALLI, V. V. (2015). Data-efficient minimax quickest change detection with composite post-change distribution. *IEEE Trans. Inform. Theory* **61** 5172–5184. MR3386508 <https://doi.org/10.1109/TIT.2015.2458864>
- [3] BASU, A., HARRIS, I. R., HJORT, N. L. and JONES, M. C. (1998). Robust and efficient estimation by minimising a density power divergence. *Biometrika* **85** 549–559. MR1665873 <https://doi.org/10.1093/biomet/85.3.549>
- [4] BREUER, T. and CSISZÁR, I. (2016). Measuring distribution model risk. *Math. Finance* **26** 395–411. MR3481309 <https://doi.org/10.1111/mafi.12050>
- [5] BRODSKY, B. and DARKHOVSKY, B. (2008). Minimax methods for multihypothesis sequential testing and change-point detection problems. *Sequential Anal.* **27** 141–173. MR2439985 <https://doi.org/10.1080/07474940801989111>
- [6] BRODSKY, B. E. and DARKHOVSKY, B. S. (2008). Minimax sequential tests for many composite hypotheses. I. *Theory Probab. Appl.* **52** 565–579.
- [7] CAPASSO, V. and BAKSTEIN, D. (2015). *An Introduction to Continuous-Time Stochastic Processes: Theory, Models, and Applications to Finance, Biology, and Medicine*, 3rd ed. *Modeling and Simulation in Science, Engineering and Technology*. Birkhäuser/Springer, New York. MR3362672 <https://doi.org/10.1007/978-1-4939-2757-9>
- [8] CHOQUET, G. (1954). Theory of capacities. *Ann. Inst. Fourier (Grenoble)* **5** 131–295. MR0080760
- [9] DEGROOT, M. H. (1960). Minimax sequential tests of some composite hypotheses. *Ann. Math. Stat.* **31** 1193–1200. MR0120734 <https://doi.org/10.1214/aoms/1177705690>
- [10] DRAGALIN, V. P. and NOVIKOV, A. A. (1987). Asymptotic solution of the Kiefer–Weiss problem for processes with independent increments. *Theory Probab. Appl.* **32** 617–627.
- [11] DVORETZKY, A., KIEFER, J. and WOLFOWITZ, J. (1953). Sequential decision problems for processes with continuous time parameter. Testing hypotheses. *Ann. Math. Stat.* **24** 254–264. MR0054911 <https://doi.org/10.1214/aoms/1177729031>
- [12] EL-SAWY, A. H. and VANDELINDE, V. D. (1979). Robust sequential detection of signals in noise. *IEEE Trans. Inform. Theory* **25** 346–353. MR0528012 <https://doi.org/10.1109/TIT.1979.1056049>
- [13] FAUSS, M. (2016). Design and analysis of optimal and minimax robust sequential hypothesis tests. Ph.D. thesis, TU Darmstadt, Institute of Communications.
- [14] FAUSS, M. and ZOUBIR, A. M. (2015). A linear programming approach to sequential hypothesis testing. *Sequential Anal.* **34** 235–263. MR3348300 <https://doi.org/10.1080/07474946.2015.1030981>
- [15] FAUSS, M. and ZOUBIR, A. M. (2016). Old bands, new tracks—Revisiting the band model for robust hypothesis testing. *IEEE Trans. Signal Process.* **64** 5875–5886. MR3554590 <https://doi.org/10.1109/TSP.2016.2600505>

- [16] FAUSS, M., ZOUBIR, A. M. and POOR, V. H. (2018). On the equivalence of  $f$ -divergence balls and density bands in robust detection. In *Proc. of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*.
- [17] FAUSS, M., ZOUBIR, A. M. and POOR, H. V. (2020). Supplement to “Minimax optimal sequential hypothesis tests for Markov processes.” <https://doi.org/10.1214/19-AOS1899SUPP>.
- [18] FELLOURIS, G. and TARTAKOVSKY, A. G. (2012). Nearly minimax one-sided mixture-based sequential tests. *Sequential Anal.* **31** 297–325. MR2958638
- [19] FELLOURIS, G. and TARTAKOVSKY, A. G. (2013). Almost optimal sequential tests of discrete composite hypotheses. *Statist. Sinica* **23** 1717–1741. MR3222817
- [20] FERRARI, S. and STENGEL, R. F. (2005). Smooth function approximation using neural networks. *IEEE Trans. Neural Netw.* **16** 24–38.
- [21] GAO, R., XIE, L., XIE, Y. and XU, H. (2018). Robust hypothesis testing using Wasserstein uncertainty sets. *Adv. Neural Inf. Process. Syst.* 7913–7923.
- [22] GHOSH, B. K. and SEN, P. K., eds. (1991) *Handbook of Sequential Analysis. Statistics: Textbooks and Monographs* **118**. Dekker, New York. MR1174296
- [23] GÜL, G. and ZOUBIR, A. M. (2016). Robust hypothesis testing with  $\alpha$ -divergence. *IEEE Trans. Signal Process.* **64** 4737–4750. MR3538377 <https://doi.org/10.1109/TSP.2016.2569405>
- [24] GYÖRFI, L. and NEMETZ, T. (1975). On the dissimilarity of probability measures. Technical report, Mathematical Institute of the Hungarian Academy of Science.
- [25] GYÖRFI, L. and NEMETZ, T. (1977).  $f$ -dissimilarity: A general class of separation measures of several probability measures. In *Topics in Information Theory (Second Colloq., Keszthely, 1975)*. *Colloq. Math. Soc. János Bolyai* **16** 309–321. MR0459923
- [26] GYÖRFI, L. and NEMETZ, T. (1978).  $f$ -dissimilarity: A generalization of the affinity of several distributions. *Ann. Inst. Statist. Math.* **30** 105–113. MR0507086 <https://doi.org/10.1007/BF02480206>
- [27] HAFNER, R. (1993). Construction of minimax-tests for bounded families of probability-densities. *Metrika* **40** 1–23. MR1210451 <https://doi.org/10.1007/BF02613659>
- [28] HAZEWINDEL, M., ed. (1994). Kolmogorov–Chapman equation. In *Encyclopaedia of Mathematics* **5** 292. Springer, Dordrecht, Netherlands.
- [29] HUBER, P. J. (1965). A robust version of the probability ratio test. *Ann. Math. Stat.* **36** 1753–1758. MR0185747 <https://doi.org/10.1214/aoms/1177699803>
- [30] HUBER, P. J. (1981). *Robust Statistics*. Wiley, Hoboken, NJ.
- [31] HUBER, P. J. and STRASSEN, V. (1973). Minimax tests and the Neyman–Pearson lemma for capacities. *Ann. Statist.* **1** 251–263. MR0356306
- [32] KASSAM, S. A. (1981). Robust hypothesis testing for bounded classes of probability densities. *IEEE Trans. Inform. Theory* **27** 242–247. MR0633425 <https://doi.org/10.1109/TIT.1981.1056314>
- [33] KASSAM, S. A. and POOR, H. V. (1985). Robust techniques for signal processing: A survey. *Proc. IEEE* **73** 433–481.
- [34] KHARIN, A. (2002). On robustifying of the sequential probability ratio test for a discrete model under “contaminations”. *Aust. J. Stat.* **31** 267–277.
- [35] KIEFER, J. and WEISS, L. (1957). Some properties of generalized sequential probability ratio tests. *Ann. Math. Stat.* **28** 57–74. MR0087290 <https://doi.org/10.1214/aoms/1177707037>
- [36] KUNSCH, R. (2017). High-dimensional function approximation: Breaking the curse with Monte Carlo methods. Ph.D. thesis, Friedrich-Schiller-Universität Jena, Jena, Germany.
- [37] LORDEN, G. (1976). 2-SPRT’s and the modified Kiefer–Weiss problem of minimizing an expected sample size. *Ann. Statist.* **4** 281–291. MR0405750
- [38] MARONNA, R. A., MARTIN, R. D. and YOHAI, V. J. (2006). *Robust Statistics: Theory and Methods*. *Wiley Series in Probability and Statistics*. Wiley, Chichester. MR2238141 <https://doi.org/10.1002/0470010940>
- [39] MAURICE, R. J. (1957). A minimax procedure for choosing between two populations using sequential sampling. *J. Roy. Statist. Soc. Ser. B* **19** 255–261.
- [40] NGUYEN, X., WAINWRIGHT, M. J. and JORDAN, M. I. (2009). On surrogate loss functions and  $f$ -divergences. *Ann. Statist.* **37** 876–904. MR2502654 <https://doi.org/10.1214/08-AOS595>
- [41] NOVIKOV, A. (2009). Optimal sequential multiple hypothesis tests. *Kybernetika (Prague)* **45** 309–330. MR2518154
- [42] ÖSTERREICHER, F. (1978). On the construction of least favourable pairs of distributions. *Z. Wahrsch. Verw. Gebiete* **43** 49–55. MR0493481 <https://doi.org/10.1007/BF00535275>
- [43] PAPAIOORGIOU, N. S. (1997). Convex integral functionals. *Trans. Amer. Math. Soc.* **349** 1421–1436. MR1321585 <https://doi.org/10.1090/S0002-9947-97-01478-5>
- [44] PARDO, L. (2006). *Statistical Inference Based on Divergence Measures. Statistics: Textbooks and Monographs* **185**. CRC Press/CRC, Boca Raton, FL. MR2183173

- [45] PAVLOV, I. V. (1991). A sequential procedure for testing composite hypotheses with application to the Kiefer–Weiss problem. *Theory Probab. Appl.* **35** 280–292.
- [46] POOR, H. V. (1980). Robust decision design using a distance criterion. *IEEE Trans. Inform. Theory* **26** 575–587. MR0583942 <https://doi.org/10.1109/TIT.1980.1056249>
- [47] POOR, H. V. and HADJILIADIS, O. (2009). *Quickest Detection*. Cambridge Univ. Press, Cambridge. MR2482527
- [48] REID, M. D. and WILLIAMSON, R. C. (2011). Information, divergence and risk for binary experiments. *J. Mach. Learn. Res.* **12** 731–817. MR2786911
- [49] ROCKAFELLAR, R. T. (1968). Integrals which are convex functionals. *Pacific J. Math.* **24** 525–539. MR0236689
- [50] ROCKAFELLAR, R. T. (1970). *Convex Analysis. Princeton Mathematical Series* **28**. Princeton Univ. Press, Princeton, NJ. MR0274683
- [51] RUDIN, W. (1976). *Principles of Mathematical Analysis*, 3rd ed. *International Series in Pure and Applied Mathematics*. McGraw-Hill, New York–Auckland–Düsseldorf. MR0385023
- [52] SCHMITZ, N. (1987). Minimax sequential tests of composite hypotheses on the drift of a Wiener process. *Stat. Hefte* **28** 247–261. MR0975919
- [53] SIEGMUND, D. (1985). *Sequential Analysis: Tests and Confidence Intervals. Springer Series in Statistics*. Springer, New York. MR0799155 <https://doi.org/10.1007/978-1-4757-1862-1>
- [54] SOCHMAN, J. and MATAS, J. (2005). WaldBoost—Learning for time constrained sequential detection. In *Proc. of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (C. Schmid, S. Soatto and C. Tomasi, eds.) 150–156.
- [55] TARTAKOVSKY, A., NIKIFOROV, I. and BASSEVILLE, M. (2015). *Sequential Analysis: Hypothesis Testing and Changepoint Detection. Monographs on Statistics and Applied Probability* **136**. CRC Press, Boca Raton, FL. MR3241619
- [56] TARTAKOVSKY, A. G., LI, X. R. and YARALOV, G. (2003). Sequential detection of targets in multichannel systems. *IEEE Trans. Inform. Theory* **49** 425–445. MR1966790 <https://doi.org/10.1109/TIT.2002.807288>
- [57] UNNIKRISHNAN, J., VEERAVALLI, V. V. and MEYN, S. P. (2011). Minimax robust quickest change detection. *IEEE Trans. Inform. Theory* **57** 1604–1614. MR2815837 <https://doi.org/10.1109/TIT.2011.2104993>
- [58] VARSHNEY, K. R. (2011). Bayes risk error is a Bregman divergence. *IEEE Trans. Signal Process.* **59** 4470–4472. MR2865998 <https://doi.org/10.1109/TSP.2011.2159500>
- [59] VERDÚ, S. and POOR, H. V. (1984). On minimax robustness: A general approach and applications. *IEEE Trans. Inform. Theory* **30** 328–340. MR0754862 <https://doi.org/10.1109/TIT.1984.1056876>
- [60] VOS, H. J. (2001). A minimax procedure in the context of sequential testing problems in psychodiagnostics. *Br. J. Math. Stat. Psychol.* **54** 139–159.
- [61] VOUDOURI, E. and KURZ, L. (1988). A robust approach to sequential detection. *IEEE Trans. Acoust. Speech Signal Process.* **36** 1200–1210.
- [62] WALD, A. (1947). *Sequential Analysis*. Wiley, Hoboken, NJ. MR0020764
- [63] WEISS, C. H. and KIM, H.-Y. (2013). Binomial AR(1) processes: Moments, cumulants, and estimation. *Statistics* **47** 494–510. MR3060925 <https://doi.org/10.1080/02331888.2011.605893>
- [64] WEISS, C. H. and KIM, H.-Y. (2013). Parameter estimation for binomial AR(1) models with applications in finance and industry. *Statist. Papers* **54** 563–590. MR3072889 <https://doi.org/10.1007/s00362-012-0449-y>
- [65] YILMAZ, E., GEMMEKE, J. F. and HAMME, H. V. (2016). Noise robust exemplar matching with alpha–beta divergence. *Speech Commun.* **76** 127–142.
- [66] ZHITLUKHIN, M. V., MURAVLEV, A. A. and SHIRYAEV, A. N. (2013). The optimal decision rule in the Kiefer–Weiss problem for a Brownian motion. *Russian Math. Surveys* **68** 389–391.
- [67] ZOUBIR, A. M., KOIVUNEN, V., CHAKHCHOUKH, Y. and MUMA, M. (2012). Robust estimation in signal processing: A tutorial-style treatment of fundamental concepts. *IEEE Signal Process. Mag.* **29** 61–80.
- [68] ZOUBIR, A. M., KOIVUNEN, V., OLLILA, E. and MUMA, M. (2018). *Robust Statistics for Signal Processing*. Cambridge Univ. Press, Cambridge, UK.