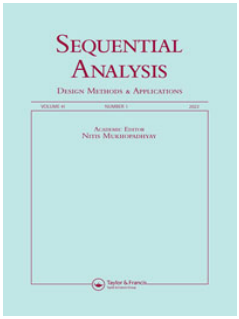

Sequential Analysis
Design Methods and Applications

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/lsqa20

Adversarially robust sequential hypothesis testing

Shuchen Cao, Ruizhi Zhang & Shaofeng Zou

To cite this article: Shuchen Cao, Ruizhi Zhang & Shaofeng Zou (2022) Adversarially robust sequential hypothesis testing, Sequential Analysis, 41:1, 81-103, DOI: 10.1080/07474946.2022.2043050

To link to this article: https://doi.org/10.1080/07474946.2022.2043050

Published online: 17 May 2022.

SEQUENTIAL ANALYSIS
2022, VOL. 41, NO. 1, 81–103
https://doi.org/10.1080/07474946.2022.2043050

Adversarially robust sequential hypothesis testing

Shuchen Cao[a], Ruizhi Zhang[a], and Shaofeng Zou[b]

[a]Department of Statistics, University of Nebraska–Lincoln, Lincoln, Nebraska, USA; [b]Department of Electrical Engineering, University at Buffalo, the State University of New York, Buffalo, New York, USA

ABSTRACT

The problem of sequential hypothesis testing is studied, where samples are taken sequentially, and the goal is to distinguish between the null hypothesis where the samples are generated according to a distribution p and the alternative hypothesis where the samples are generated according to a distribution q. The defender (decision maker) aims to distinguish the two hypotheses using as few samples as possible subject to false alarm constraints. The problem is studied under the adversarial setting, where the data generating distributions under the two hypotheses are manipulated by an adversary, whose goal is to deteriorate the performance of the defender—for example, increasing the probability of error and expected sample sizes—with minimal cost. Specifically, under the null hypothesis, the adversary picks a distribution $p \in \mathcal{P}$ with cost $c_0(p)$, and under the alternative hypothesis, the adversary picks a distribution $q \in \mathcal{Q}$ with cost $c_1(q)$. This problem is formulated as a non-zero-sum game between the defender and the adversary. A pair of strategies (the adversary's strategy and the defender's strategy) is proposed and proved to be a Nash equilibrium pair for the non-zero-sum game between the adversary and the defender asymptotically. The defender's strategy is a sequential probability ratio test and thus is computationally efficient for practical implementation.

ARTICLE HISTORY
Received 29 April 2021
Revised 2 December 2021
Accepted 5 December 2021

KEYWORDS
Nash equilibrium; non-zero-sum game; SPRT

SUBJECT CLASSIFICATIONS
62L10; 62L15; 60G40

1. INTRODUCTION

With recent advancements in wireless communication and sensing technology, rich and complex sequential high-dimensional data from large-scale distributed sensor networks and cyber-physical systems are made available for a wide range of statistical inference applications. However, the reliance on wireless communication and the sparse spatial distribution of these networked systems make them extremely vulnerable to adversarial attacks, such as measurement manipulation, communication blocks, and false data injection. Those attacks may result in substantial damage to critical infrastructures, the economy, the ecosystem, and even public safety and thus need to be detected in an efficient and reliable manner. For example, an adversary would like to inject false measurements into a smart grid. The adversary may choose different false data injection strategies and

CONTACT Ruizhi Zhang rzhang35@unl.edu Department of Statistics, University of Nebraska–Lincoln, Lincoln, Nebraska 68583, USA.
A preliminary version of this work was presented in part at IEEE International Symposium on Information Theory (ISIT), June 2020, Los Angeles, California, USA (Zhang and Zou 2020).

will receive different rewards. The utility of the defender can be defined according to the damage to the power system caused by the attack.

An adversary may adapt its strategy according to the detection scheme to increase the error probability and expected sample size and, therefore, classical hypothesis testing approaches (Kay 1993; Cover and Thomas 2012; Poor 2013; Moulin and Veeravalli 2018) may fail. This motivates the study of hypothesis testing in adversarial environments in this article, which is of great importance especially in security applications. The adversary is often caused by human factors (Marano, Matta, and Tong 2008; Mo, Hespanha, and Sinopoli 2014; Jin and Lai 2021). Thus, any attempts to improve the performance of test schemes may result in a dual effort to devise more powerful counterthreat detection techniques that leave less evidence. Although this is an unavoidable and possibly virtuous loop, which may finally lead to more powerful threat detection and counterthreat detection tools, it is desired to investigate the ultimate limits of such a procedure.

In this article, we propose a game-theoretic approach to study the limit of such a procedure; that is, Nash equilibrium (Reny 2008). The defender receives samples sequentially (see Tartakovsky, Nikiforov, and Basseville [2014]; Poor and Hadjiliadis [2009]; Moulin and Veeravalli [2018]; Siegmund [1985] for sequential hypothesis testing). Specifically, under the null hypothesis $\mathcal{H}_0$, the samples are generated independently by a distribution $p \in \mathcal{P}$ that is picked by the adversary with cost $c_0(p)$; and under the alternative hypothesis $\mathcal{H}_1$, the samples are generated independently by another distribution $q \in \mathcal{Q}$ that is picked by the adversary with cost $c_1(q)$. The defender aims to use as few samples as possible to accurately decide which hypothesis is correct, whereas the adversary aims to fool the defender with a minimal cost. To model the interaction between the adversary and the defender and to take into consideration the costs $c_0(p)$ and $c_1(q)$ of using strategies p and q by the adversary, we formulate the problem as a non-zero-sum game between the defender and the adversary. The goal of this article is to find a Nash equilibrium strategy pair for this non-zero-sum game and analyze the performance of the defender and the adversary in the Nash equilibrium.

Our problem in this article is related to the problem of robust hypothesis testing (Huber 1965, 1981; Veeravalli, Basar, and Poor 1994; Pandit, Meyn, and Veeravalli 2004; Levy 2009; Wilcox 2011; Gül and Zoubir 2017; Molloy and Ford 2017; Gao et al. 2018; Qin and Priebe 2017). For robust hypothesis testing, the goal is to design algorithms that are robust to model uncertainty. Specifically, the problem of robust hypothesis testing is usually formulated into a minimax problem, which is to first find the least favorable distributions and then design likelihood ratio test between the least favorable distributions. However, such an approach does not take into consideration the interaction between the adversary and the defender. In adversarial environments, where security is crucial, the adversary (defender) may change its strategy according to the defender's (adversary's) strategy and vice versa. Therefore, it is of more practical importance to formulate the problem into a game and study its Nash equilibrium. More important, in this article, we design different utility functions for the defender and the adversary that take into consideration the cost of picking particular strategies, which were not investigated in these robust hypothesis testing studies.

The problem of hypothesis testing using a game-theoretic approach was studied in Yasodharan and Loiseau (2019), Vamvoudakis et al. (2014), and Barni and Tondi (2013).

We note that those studies are nonsequential; that is, one decision is made using a fixed number of samples. However, as observed in sequential hypothesis testing (Tartakovsky, Nikiforov, and Basseville 2014), sequential tests can achieve the same accuracy using fewer samples on average than fixed-sample-size tests and are more suitable for online applications with time series data. In this article, we focus on the sequential setting, where samples arrive sequentially, one at each time instant. Specifically, in Vamvoudakis et al. (2014) and Barni and Tondi (2013), the problem was formulated as a zero-sum game, where the probability of error is used as the utility function. Moreover, Vamvoudakis et al. (2014) focused on the particular case where the distributions are Bernoulli. In Yasodharan and Loiseau (2019), the non-zero-sum setting was investigated, where the adversary picks a distribution $q \in \mathcal{Q}$ with cost $c(q)$. The utility functions used in Yasodharan and Loiseau (2019) are linear combinations of the error probabilities and the cost function $c(q)$. For a large number of samples, the error probabilities converge to zero; however, $c(q)$ does not scale with the sample size. Therefore, such utility functions fail to consider the error probabilities asymptotically because when comparing to the cost $c(q)$, the vanishing error probabilities are negligible. This article proposes designing utility functions using error exponents instead of error probabilities to overcome this challenge.

Our article is also related to the studies of sequential detection and quickest change detection in sensor networks with Byzantine sensors (Bayraktar and Lai 2015; Fellouris, Bayraktar, and Lai 2018; Li, Mo, and Hao 2019; Y.-C. Huang, Lin, and Huang 2019b; Y.-J. Huang, Lin, and Huang 2019). In Li, Mo, and Hao (2019), a game-theoretic approach was constructed to solve sequential hypothesis testing with Byzantine sensors. In Bayraktar and Lai (2015), Y.-J. Huang, Lin, and Huang (2019), Fellouris, Bayraktar, and Lai (2018), and Y.-C. Huang, Lin, and Huang (2019), the problem of quickest change detection with Byzantine sensors was studied but not under a game-theoretic framework. The main difference lies in that in Li, Mo, and Hao (2019), Bayraktar and Lai (2015), Fellouris, Bayraktar, and Lai (2018), Y.-J. Huang, Lin, and Huang (2019), and Y.-C. Huang, Lin, and Huang (2019), it was assumed that the adversary can modify samples from compromised sensors arbitrarily. However, only a limited number of sensors can be compromised; in this article, we only have one single data stream, and the adversary can only choose its strategy p and q from two fixed sets $\mathcal{P}$ and $\mathcal{Q}$ with costs $c_0(p)$ and $c_1(q)$.

In this article, we construct a pair of strategies and prove that they are in Nash equilibrium for the non-zero-sum game asymptotically. More specifically, our main contributions can be summarized as follows: (i) we propose a non-zero-sum game to model sequential detection in adversarial environments, which provides a general framework for this type of problem; (ii) we develop a novel scheme with utility functions defined by the error exponents of the sequential tests and the cost functions c_0 and c_1; (iii) we construct a pair of strategies, prove that they are in Nash equilibrium of the game asymptotically, and characterize their utilities at the Nash equilibrium; and (iv) provide some examples and extensive numerical results to validate our theoretical assertions.

1.1. Article Organization

The article is organized as follows. In Section 2, we introduce the problem formulation. In Section 3, we provide preliminaries and background information on Nash equilibrium

and the sequential probability ratio test (SPRT). In Section 4, we present our main results. Specifically, we provide general sufficient conditions to construct a pair of strategies in the Nash equilibrium asymptotically. We also characterize the utilities of both the defender and the adversary at the equilibrium. In Section 5, we use several examples to illustrate how to construct strategies in Nash equilibrium and present extensive numerical experiments to illustrate our theoretical results. In Section 6, we conclude our article with a few remarks.

2. PROBLEM FORMULATION

Suppose we are observing a data stream $\{X_t\}_{t=1}^{\infty}$ sequentially. There are two players: the adversary and the defender. Under the null hypothesis, X_ts are independent and identically distributed (i.i.d.) according to a distribution p picked by the adversary from a set of distributions $\mathcal{P}$. Under the alternative hypothesis, X_ts are i.i.d. according to another distribution $q \in \mathcal{Q}$ picked by the adversary. It is assumed that $\mathcal{P} \cap \mathcal{Q} = \emptyset$. The adversary's strategy is defined by $\kappa = (p, q)$, where $p \in \mathcal{P}$ and $q \in \mathcal{Q}$. The problem can be formulated as a sequential hypothesis testing problem where the null hypothesis is

$$\mathcal{H}_0 : X_t\text{'s} \sim p \text{ i.i.d.}, \quad (\text{i.e., adversary picks } p \in \mathcal{P}),$$

and the alternative hypothesis is

$$\mathcal{H}_1 : X_t\text{'s} \sim q \text{ i.i.d.}, \quad (\text{i.e., adversary picks } q \in \mathcal{Q}).$$

The adversary's goal is to pick a strategy $\kappa = (p, q)$ so that the defender needs more samples to accurately make a conclusion about the two hypotheses. We further define two cost functions $c_0(p) : \mathcal{P} \mapsto R^+$ and $c_1(q) : \mathcal{Q} \mapsto R^+$ to model the cost of the adversary using strategy κ; for example, for $\kappa = (p, q)$, $c_0(p)$ is the power consumption of choosing $p \in \mathcal{P}$ and $c_1(q)$ is the power consumption of choosing $q \in \mathcal{Q}$.

The defender's strategy is defined by a sequential decision rule $\delta = (T, d)$, where $T \in \{1, 2, \ldots\}$ is a stopping time and $d \in \{0, 1\}$ is the decision rule. Here $d = 0$ is to accept $\mathcal{H}_0$, and $d = 1$ is to reject $\mathcal{H}_0$. The defender's goal is to minimize the number of samples needed for correct decisions. There are two types of expected number of samples. One is the expected number of samples needed to make a decision under the null hypothesis, denoted by $\mathbf{E}_p[T]$. The other one is the expected number of samples needed to make a decision under the alternative hypothesis, denoted by $\mathbf{E}_q[T]$. We further define the type I error $\alpha = \mathbf{P}_p(d = 1)$ and the type II error $\beta = \mathbf{P}_q(d = 0)$. The defender considers stopping rules in the following set:

$$C(\alpha, \beta) := \{\delta = (T, d) : \mathbf{P}_p(d = 1) \le \alpha, \mathbf{P}_q(d = 0) \le \beta, \mathbf{E}_p[T] < \infty, \mathbf{E}_q[T] < \infty\}. \quad (2.1)$$

If the defender's strategy is $\delta \in C(\alpha, \beta)$, and the adversary's strategy is κ, the defender's utility function is defined as

$$u_D(\delta, \kappa) = \gamma \frac{|\log \beta|}{\mathbf{E}_p[T]} + \frac{|\log \alpha|}{\mathbf{E}_q[T]}, \quad (2.2)$$

where $\gamma > 0$ is a positive number to capture the weights for $\mathbf{E}_p[T]$ and $\mathbf{E}_q[T]$. Such a definition of the defender's utility function captures the error exponent of the two types

of errors, and the defender would like to maximize it so that the detection is more sample efficient.

The adversary's goal is to fool the defender as much as possible; that is, to maximize the expected number of samples needed. We then define the adversary's utility when it plays the strategy $\kappa = (p, q)$ and the defender plays the strategy $\delta \in C(\alpha, \beta)$ as follows:

$$u_A(\delta, \kappa) = -c_0(p) - c_1(q) - \gamma \frac{|\log \beta|}{\mathbf{E}_p[T]} - \frac{|\log \alpha|}{\mathbf{E}_q[T]}. \tag{2.3}$$

By comparing (2.2) and (2.3), it can be observed that if there is no cost for the adversary to use strategy κ—that is, $c_0(p) = c_1(q) = 0$—then the game is a zero-sum game. The goal of this article is to construct a Nash equilibrium pair of strategies for this game.

3. PRELIMINARIES

In this section, we provide some preliminaries for Nash equilibrium, which is the central concept of rational behavior in noncooperative game theory, and the SPRT, which is optimal for the simple sequential hypothesis testing problem.

3.1. Nash Equilibrium Strategy

Based on the discussion in the previous section, the problem of sequential hypothesis testing in an adversarial environment can be modeled as a two-player game, where the defender's strategy and utility function are $(\delta, u_D(\delta, \kappa))$ and the adversary's strategy and utility function are $(\kappa, u_A(\delta, \kappa))$. Both players intend to maximize their corresponding utility functions.

Under this game-theoretic framework, we are ready to introduce the concept of Nash equilibrium (Reny 2008).

Definition 3.1. *For the strategies* $\delta^* \in C(\alpha, \beta), \kappa^* = (p^*, q^*),\ p^* \in \mathcal{P}, q^* \in \mathcal{Q},\ (\delta^*, \kappa^*)$ *is a pure strategy Nash equilibrium if*

$$u_A(\delta^*, \kappa^*) \geq u_A(\delta^*, \kappa) \quad \text{for all } \kappa,$$
$$u_D(\delta^*, \kappa^*) \geq u_D(\delta, \kappa^*) \quad \text{for all } \delta \in C(\alpha, \beta).$$

In particular, for the zero-sum game where $u_A(\delta, \kappa) + u_D(\delta, \kappa) = 0$, (δ^*, κ^*) is a pure strategy Nash equilibrium if for all κ and $\delta \in C(\alpha, \beta)$, $u_A(\delta^*, \kappa) \leq u_A(\delta^*, \kappa^*) \leq u_A(\delta, \kappa^*)$.

Based on the definition, a strategy (δ^*, κ^*) is in Nash equilibrium if no player can do better by unilaterally changing its own strategy. It is interesting to see the difference between the Nash equilibrium strategy and the minimax optimal strategy robust sequential test. Specifically, the minimax optimal strategy $\tilde{\delta}$ is usually obtained by solving the following problem:

$$\min_{\delta \in C(\alpha, \beta)} \max_{\kappa} (u_A(\delta, \kappa)), \tag{3.1}$$

the solution to which is not necessarily in Nash equilibrium.

3.2. The Sequential Probability Ratio Test

For the simple sequential hypothesis testing problem, the SPRT, which was developed by Wald (1945), is widely used. Consider the following sequential hypothesis testing problem:

$$\mathcal{H}_0 : X_t\text{'s} \sim p \text{ i.i.d.}, \tag{3.2}$$

$$\mathcal{H}_1 : X_t\text{'s} \sim q \text{ i.i.d.}, \tag{3.3}$$

where $\{X_t\}_{t=1}^{\infty}$ is the observing data stream.

Let $S_n = \sum_{i=1}^{n} \log \frac{q(X_i)}{p(X_i)}$. Then the SPRT testing p against q is a pair of a stopping time T and a decision rule d, where

$$T(a,b) = \inf\{n : S_n \notin (a,b)\}, \tag{3.4}$$

$$d = \begin{cases} 0, & \text{if } S_T \le a, \\ 1, & \text{if } S_T \ge b. \end{cases} \tag{3.5}$$

Here a, b are two stopping thresholds that control the type I and type II errors of the SPRT. The SPRT has nice optimality that of all tests with the same power, on average, it requires the fewest observations (Wald and Wolfowitz 1948).

Denote the Kullback-Leibler divergences between p and q by $KL(p||q) = \int p(x) \log \frac{p(x)}{q(x)} dx$, and assume that $KL(p||q)$ and $KL(q||p)$ are well defined and finite for any $p \in \mathcal{P}$ and any $q \in \mathcal{Q}$. Then, as $a \to -\infty, b \to +\infty$, it can be shown that

$$\mathbf{P}_p(S_T \ge b) = (1 + o(1))e^{-b}, \tag{3.6}$$

$$\mathbf{P}_q(S_T \le a) = (1 + o(1))e^{a}. \tag{3.7}$$

Thus, the choices of $a = \log \beta$, and $b = \log \frac{1}{\alpha}$, guarantee that $\mathbf{P}_p(d = 1) = (1 + o(1))\alpha$ and $\mathbf{P}_q(d = 0) = (1 + o(1))\beta$. Additionally, the expected sample sizes of the SPRT can be characterized as follows:

$$\mathbf{E}_p[T(a,b)] = (1 + o(1)) \frac{|a|}{KL(p||q)}, \tag{3.8}$$

$$\mathbf{E}_q[T(a,b)] = (1 + o(1)) \frac{b}{KL(q||p)}. \tag{3.9}$$

See Tartakovsky, Nikiforov, and Basseville (2014), Moulin and Veeravalli (2018), and Siegmund (1985) for more details of the SPRT.

4. ASYMPTOTIC NASH EQUILIBRIUM STRATEGY

In this section, we construct a pair of strategies $(\delta^*, (p^*, q^*))$, yjsy are in Nash equilibrium asymptotically. Specifically, if we can find a distribution $p^* \in \mathcal{P}$ and a $q^* \in \mathcal{Q}$ satisfying

$$\begin{aligned} (p^*, q^*) = \arg\max_{p \in \mathcal{P}, q \in \mathcal{Q}} (&-c_0(p) - c_1(q) \\ &- \gamma \lambda_2(q, p^*, q^*) \mathbf{E}_p \log \frac{p^*(x)}{q^*(x)} - \lambda_1(p, p^*, q^*) \mathbf{E}_q \log \frac{q^*(x)}{p^*(x)}), \end{aligned} \tag{4.1}$$

$$\inf_{p\in\mathcal{P}}\left[KL(p||q^*)-KL(p||p^*)\right]>0, \tag{4.2}$$

and

$$\inf_{q\in\mathcal{Q}}\left[KL(q||p^*-KL(q||q^*)\right]>0, \tag{4.3}$$

and a sequential test $\delta^*\in C(\alpha,\beta)$, which is the SPRT testing q^* against p^* with $\mathbf{P}_{p^*}(d=1)=\alpha$ and $\mathbf{P}_{q^*}(d=0)=\beta$, then we have the following theorem of asymptotic Nash equilibrium. Note that $\lambda_1(p,p^*,q^*)$ is a positive number such that

$$\mathbf{E}_p\left[\left(\frac{q^*(x)}{p^*(x)}\right)^{\lambda_1}\right]=\int p(x)\left(\frac{q^*(x)}{p^*(x)}\right)^{\lambda_1}dx=1, \tag{4.4}$$

and $\lambda_2(q,p^*,q^*)$ is a positive number such that

$$\mathbf{E}_q\left[\left(\frac{p^*(x)}{q^*(x)}\right)^{\lambda_2}\right]=\int q(x)\left(\frac{p^*(x)}{q^*(x)}\right)^{\lambda_2}dx=1. \tag{4.5}$$

For simplification, we use $\lambda_1(p)$ and $\lambda_2(q)$ to denote them respectively.

Theorem 4.1. *Assume that there exists a distribution $p^*\in\mathcal{P}$ and a distribution $q^*\in\mathcal{Q}$ satisfying* (4.1), (4.2), *and* (4.3) *and* $\mathbf{E}_p\left[\left(\log\frac{q*(X_i)}{p*(X_i)}\right)^2\right]$ *and* $\mathbf{E}_q\left[\left(\log\frac{q*(X_i)}{p*(X_i)}\right)^2\right]$ *are finite for any $p\in\mathcal{P}$ and $q\in\mathcal{Q}$. Then as $\alpha,\beta\to 0$, $\beta\log\alpha\to 0$, $\alpha\log\beta\to 0$, $(\delta_{\kappa^*},\kappa^*)$ is in Nash equilibrium asymptotically for the utility functions $u_A(\delta,\kappa)$ and $u_D(\delta,\kappa)$; that is,*

$$\lim_{\alpha,\beta\to 0}u_D(\delta_{\kappa*},\kappa^*)\geq\lim_{\alpha,\beta\to 0}u_D(\delta,\kappa^*), \tag{4.6}$$

$$\lim_{\alpha,\beta\to 0}u_A(\delta_{\kappa*},\kappa^*)\geq\lim_{\alpha,\beta\to 0}u_A(\delta_{\kappa*},\kappa), \tag{4.7}$$

where $\delta=(T,d)$ is any sequential test in $C(\alpha,\beta)$, $\kappa^*=(p^*,q^*)$, $\delta_{\kappa*}$ is the SPRT for testing p^* against q^*.

We note that (4.6) also holds without the limit. Before providing the detailed proof of Theorem 4.1, we first present the following useful lemma.

Lemma 4.1. *Suppose that Y is a continuous random variable that takes both positive and negative values with nonzero probabilities, and assume that $\varphi(\lambda)=\mathbf{E}[e^{-\lambda Y}]$ is well defined over $-\infty<\lambda<\infty$. Then there exists a constant $\lambda^*>0$ satisfying $\varphi(\lambda^*)=1$ if and only if $\mathbf{E}[Y]>0$.*

The proof of Lemma 4.1 can be found in Zhang and Zou (2020). Note that in Lemma 4.1, we assume that Y is a continuous random variable for simplicity. In general, this lemma still holds with mild assumptions on Y. See appendix A2 of Wald (2013) for more details. With Lemma 4.1, we are ready to prove Theorem 4.1.

Proof. This is the proof of Theorem 4.1. To simplify the notation, in this proof, $\mathbf{P}_p$ and $\mathbf{E}_p$ denote the probability measure and expectation when the data have the probability distribution $p(x)$.

For fixed $\kappa \in \mathcal{P} \times \mathcal{Q}$, let $\delta_\kappa = (T_\kappa, d_\kappa) \in C(\alpha, \beta)$ be the SPRT testing q against p with $\mathbf{P}_p(d_\kappa = 1) = \alpha$ and $\mathbf{P}_q(d_\kappa = 0) = \beta$.

First, we show that $u_D(\delta_\kappa, \kappa) \geq u_D(\delta, \kappa)$ for any test $\delta = (T, d) \in C(\alpha, \beta)$. In fact, this follows easily from the optimality of the SPRT (Wald and Wolfowitz 1948). Specifically, for any test $\delta = (T, d) \in C(\alpha, \beta)$, we have that $\mathbf{E}_p(T_\kappa) \leq \mathbf{E}_p(T)$ and $\mathbf{E}_q(T_\kappa) \leq \mathbf{E}_q(T)$. Thus, $u_D(\delta_\kappa, \kappa) \geq u_D(\delta, \kappa)$. In particular, we have that $u_D(\delta_{\kappa*}, \kappa^*) \geq u_D(\delta, \kappa^*)$.

Next, we will show that $u_A(\delta_{\kappa*}, \kappa^*) \geq u_A(\delta_{\kappa*}, \kappa)$ for any $\kappa \in \mathcal{P} \times \mathcal{Q}$. For any fixed $\kappa \in \mathcal{P} \times \mathcal{Q}$, we first find the thresholds of SPRT δ_κ to guarantee the type I and type II error constraints. Specifically, consider the stopping time

$$T^*(a, b) = \inf\left\{ n : \sum_{i=1}^{n} \log \frac{q^*(X_i)}{p^*(X_i)} \notin (a, b) \right\}. \tag{4.8}$$

We need to determine a and b such that

$$\mathbf{P}_p\left(\sum_{i=1}^{T^*} \log \frac{q^*(X_i)}{p^*(X_i)} \geq b \right) = \alpha, \tag{4.9}$$

$$\mathbf{P}_q\left(\sum_{i=1}^{T^*} \log \frac{q^*(X_i)}{p^*(X_i)} \leq a \right) = \beta. \tag{4.10}$$

First we compute the following probability

$$\mathbf{P}_p\left(\sum_{i=1}^{T^*} \log \frac{q^*(X_i)}{p^*(X_i)} \geq b \right). \tag{4.11}$$

By the condition in (4.2), we have that $\int p(x) \log \frac{p^*(x)}{q^*(x)} dx > 0$. Thus, if $Y_1 = \log \frac{p^*(X)}{q^*(X)}$ is a continuous random variable and can take both positive and negative values, by Lemma 4.1, there exists a positive number $\lambda_1(p)$ (which depends on p, p^*, q^*) such that

$$\mathbf{E}_p[e^{-\lambda_1(p) \cdot Y_1}] = \int p(x) \left(\frac{q^*(x)}{p^*(x)} \right)^{\lambda_1(p)} dx = 1. \tag{4.12}$$

Then, if we define $h_1(x) = p(x)(\frac{q^*(x)}{p^*(x)})^{\lambda_1(p)}$, $h_1(x)$ is just a probability density function. Moreover, we have that $\log \frac{h_1(x)}{p(x)} = \lambda_1(p) \log \frac{q^*(x)}{p^*(x)}$. Therefore, the stopping time $T^*(a, b)$ would be the same as $T_1' = \inf\left\{ n : \sum_{i=1}^{n} \log \frac{h_1(X_i)}{p(X_i)} \notin (a\lambda_1(p), b\lambda_1(p)) \right\}$, which is the SPRT for h_1 against p. By the property of the SPRT in (3.6), it then follows that

$$\mathbf{P}_p\left(\sum_{i=1}^{T_1'} \log \frac{h_1(X_i)}{p(X_i)} \geq b\lambda_1(p) \right) = (1 + o(1))e^{-b\lambda_1(p)}. \tag{4.13}$$

Note that

$$\mathbf{P}_p\left(\sum_{i=1}^{T^*} \log \frac{q^*(X_i)}{p^*(X_i)} \geq b \right) = \mathbf{P}_p\left(\sum_{i=1}^{T_1'} \log \frac{h_1(X_i)}{p(X_i)} \geq b\lambda_1(p) \right),$$

which yields that

$$\mathbf{P}_p\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\geq b\right)=(1+o(1))e^{-b\lambda_1(p)}. \tag{4.14}$$

Next, we need to compute the following probability:

$$\mathbf{P}_q\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\leq a\right). \tag{4.15}$$

By the condition in (4.3), we have that $\int q(x)\log\frac{q^*(x)}{p^*(x)}dx>0$. Thus, if $Y_2=\log\frac{q^*(X)}{p^*(X)}$ is a continuous random variable and can take both positive and negative values, by Lemma 4.1, there exists a positive number $\lambda_2(q)$ (which depends on q,p^*,q^*) such that

$$\mathbf{E}_q\left[e^{-\lambda_2(q)\cdot Y_2}\right]=\int q(x)\left(\frac{p^*(x)}{q^*(x)}\right)^{\lambda_2(q)}dx=1. \tag{4.16}$$

Then, if we define $h_2(x)=q(x)(\frac{p^*(x)}{q^*(x)})^{\lambda_2(q)}$, $h_2(x)$ is a probability density function. Moreover, we have that $\log\frac{q(x)}{h_2(x)}=\lambda_2(q)\log\frac{q^*(x)}{p^*(x)}$. Therefore, the stopping time $T^*(a,b)$ would be the same as $T_2'=\inf\left\{n:\sum_{i=1}^n\log\frac{q(X_i)}{h_2(X_i)}\notin(a\lambda_2(q),b\lambda_2(q))\right\}$, which is the SPRT for q against h_2. By the property of the SPRT in (3.6), it then follows that

$$\mathbf{P}_q\left(\sum_{i=1}^{T_2'}\log\frac{q(X_i)}{h_2(X_i)}\leq a\lambda_2(q)\right)=(1+o(1))e^{a\lambda_2(q)}. \tag{4.17}$$

Note that

$$\mathbf{P}_q\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\leq a\right)=\mathbf{P}_q\left(\sum_{i=1}^{T_2'}\log\frac{q(X_i)}{h_2(X_i)}\leq a\lambda_2(q)\right),$$

which yields that

$$\mathbf{P}_q\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\leq a\right)=(1+o(1))e^{a\lambda_2(q)}. \tag{4.18}$$

Combining (4.14) and (4.18), as $\alpha,\beta\to 0$, we can choose

$$a=(\log\beta)/\lambda_2(q),\ \text{ and }\quad b=-(\log\alpha)/\lambda_1(p). \tag{4.19}$$

to achieve the type I and type II error constraints asymptotically.

By Wald's identity, we have that

$$\mathbf{E}_p\left[\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\right]=\mathbf{E}_p[T^*]\mathbf{E}_p\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right], \tag{4.20}$$

$$\mathbf{E}_q\left[\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\right]=\mathbf{E}_q[T^*]\mathbf{E}_q\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right]. \tag{4.21}$$

Note that as $\alpha,\beta\to 0$, $\alpha\log\beta\to 0$ and $\beta\log\alpha\to 0$, and $\mathbf{E}_p\left[(\log\frac{q^*(X_i)}{p^*(X_i)})^2\right]$ and $\mathbf{E}_q\left[(\log\frac{q^*(X_i)}{p^*(X_i)})^2\right]$ are finite, then the overshoot can be ignored (Siegmund 1985). Thus, we have that

$$\begin{aligned}\mathbf{E}_p\left[\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\right] &= a\mathbf{P}_p\left(\sum_{i=1}^{T^*}\log\frac{q^*X_i)}{p^*(X_i)}\le a\right)+b\mathbf{P}_p\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\ge b\right)+O(1)\\ &= a(1-\alpha)+b\alpha+O(1)\\ &= \frac{\log\beta}{\lambda_2(q)}-\frac{\alpha\log\beta}{\lambda_2(q)}-\frac{\alpha\log\alpha}{\lambda_1(p)}+O(1)\\ &= \frac{\log\beta}{\lambda_2(q)}+O(1),\end{aligned}$$

and

$$\begin{aligned}\mathbf{E}_q\left[\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\right] &= a\mathbf{P}_q\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\le a\right)+b\mathbf{P}_q\left(\sum_{i=1}^{T^*}\log\frac{q^*(X_i)}{p^*(X_i)}\ge b\right)+O(1)\\ &= a\beta+b(1-\beta)+O(1)\\ &= -\frac{\log\alpha}{\lambda_1(p)}+\frac{\beta\log\alpha}{\lambda_1(p)}+\frac{\beta\log\beta}{\lambda_2(q)}+O(1)\\ &= -\frac{\log\alpha}{\lambda_1(p)}+O(1).\end{aligned}$$

Therefore, we have that for any $p\in\mathcal{P}$ and $q\in\mathcal{Q}$,

$$\begin{aligned}\mathbf{E}_p(T^*) &= (1+o(1))\frac{\log\beta}{\lambda_2(q)\mathbf{E}_p\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right]},\\ \mathbf{E}_q(T^*) &= (1+o(1))\frac{-\log\alpha}{\lambda_1(p)\mathbf{E}_q\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right]},\end{aligned}$$

and the utility function $\lim_{\alpha,\beta\to0}u_A(\delta_{\kappa*},\kappa)=-c_0(p)-c_1(q)-\gamma\lambda_2(q)\mathbf{E}_p[\log\frac{p^*(X_i)}{q^*(X_i)}]-\lambda_1(p)\mathbf{E}_q[\log\frac{q^*(X_i)}{p^*(X_i)}]$. It then follows from the definition of κ^* in (4.1) that

$$\begin{aligned}\lim_{\alpha,\beta\to0}u_A(\delta_{\kappa*},\kappa^*) &= -c_0(p^*)-c_1(q^*)-\gamma\lambda_2(q^*)\mathbf{E}_{p^*}\left[\log\frac{p^*(X_i)}{q^*(X_i)}\right]-\lambda_1(p^*)\mathbf{E}_{q^*}\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right]\\ &\ge -c_0(p)-c_1(q)-\gamma\lambda_2(q)\mathbf{E}_p\left[\log\frac{p^*(X_i)}{q^*(X_i)}\right]-\lambda_1(p)\mathbf{E}_q\left[\log\frac{q^*(X_i)}{p^*(X_i)}\right].\end{aligned}$$

We then conclude that $\lim_{\alpha,\beta\to0}u_A(\delta_{\kappa*},\kappa^*)\ge\lim_{\alpha,\beta\to0}u_A(\delta_{\kappa*},\kappa)$. Therefore, $(\delta_{\kappa*},\kappa^*)$ is a pair of strategies in Nash equilibrium asymptotically as $\alpha,\beta\to0$, which completes the proof.

Based on Theorem 4.1, we can get a pair of strategies in Nash equilibrium asymptotically once we can find a pair of distributions $p^*\in\mathcal{P},q^*\in\mathcal{Q}$ that satisfy conditions (4.1), (4.2), and (4.3). For a special case when the cost functions $c_0(p)=c_1(q)=0$, we can get the least favorable pair distributions that satisfy these conditions. In the next section, we will consider several examples to construct strategies in Nash equilibrium when the cost functions are nonzero.

Furthermore, based on Theorem 4.1, we can get the following corollary, which summarizes the asymptotic utilities for the strategies in Nash equilibrium.

Corollary 4.1. *If there exist distributions* $\kappa^* = (p^* \in \mathcal{P}, q^* \in \mathcal{Q})$ *satisfying* (4.1), (4.2), *and* (4.3) *and* $(\delta_{\kappa^*}, \kappa^*)$ *are in Nash equilibrium asymptotically, then the utilities of adversary and defender are*

$$\lim_{\alpha,\beta\to 0} u_A(\delta_{\kappa^*}, \kappa^*) = -c_0(p^*) - c_1(q^*) - \gamma KL(p^*||q^*) - KL(q^*||p^*),$$
$$\lim_{\alpha,\beta\to 0} u_D(\delta_{\kappa^*}, \kappa^*) = \gamma KL(p^*||q^*) + KL(q^*||p^*).$$

5. EXAMPLES AND NUMERICAL RESULTS

In this section, we consider four examples to illustrate our theoretical result.

5.1. Sparse Attack

In this subsection, we consider a special case when the attack is sparse. Specifically, suppose under the null hypothesis when there is no attack, X_ts are i.i.d with probability density function $p(x)$. Under the alternative hypothesis, the adversary will attack each data with probability $\epsilon > 0$. The data that will be attacked have probability density function $q(x)$. In the other word, we assume the set $\mathcal{P} = \{p\}$ has only one distribution. The set $\mathcal{Q} = \{(1-\epsilon)p + \epsilon q | 0 < \epsilon < 1\}$, where q is a known density. In this case, the adversary's strategy is determined by the attack probability ϵ and its goal is to choose an ϵ that can maximize its utility.

By Theorem 4.1, we can get the following corollary, which provides a sufficient condition to a pair of strategies in Nash equilibrium.

Corollary 5.1. *Suppose that* $\mathcal{P} = \{p\}$ *and* $\mathcal{Q} = \{(1-\epsilon)p + \epsilon q | 0 < \epsilon < 1\}$, *where* p *and* q *are two fixed probability density functions. Let* $c(\epsilon)$ *be the cost function by choosing* ϵ. *If there exists* ϵ^* *minimizing*

$$-c(\epsilon) - \int \left[(1-\epsilon)p(x) + \epsilon q(x)\right] \log \frac{(1-\epsilon^*)p(x) + \epsilon^* q(x)}{p(x)}\, dx \tag{5.1}$$

and

$$\int \left[(1-\epsilon)p(x) + \epsilon q(x)\right] \log \frac{(1-\epsilon^*)p(x) + \epsilon^* q(x)}{p(x)}\, dx > 0 \tag{5.2}$$

for all $\epsilon \in (0,1)$, *then* $(\delta_{\epsilon^*}, \epsilon^*)$ *is in Nash equilibrium asymptotically.*

Note the ϵ^* depends on distributions p and q. In general, there is no closed form to get ϵ^* from (5.1) explicitly. Fortunately, numeric methods can be used to search for ϵ^*.

For example, let $p(x) = \text{Beta}(2,1)$, $q(x) = \text{Beta}(3,1)$, and $c(\epsilon) = (1-\epsilon)^2$. Define function $f(\epsilon) = -(1-\epsilon)^2 - \int_0^1 (1 - \epsilon + 2\epsilon x) \log[1 + (2x-1)\epsilon^*] dx$. By Corollary 5.1, if ϵ^* maximizes $f(\epsilon)$, then $(\delta_{\epsilon^*}, \epsilon^*)$ are in Nash equilibrium. Note that $f(\epsilon)$ is a concave function of ϵ. If ϵ^* is the maximizer of $f(\epsilon)$, we should have $\frac{\partial f}{\partial \epsilon}(\epsilon^*) = 0$, which yields

$$\epsilon^* = 1 - \frac{1}{2} \int (3x^2 - 2x) \log\left[1 + \left(\frac{3}{2}x - 1\right)\epsilon^*\right] dx. \tag{5.3}$$

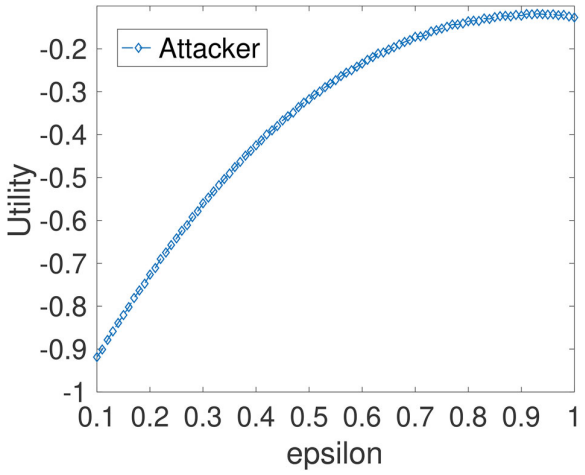

Figure 1. Utilities of the attacker in beta distribution.

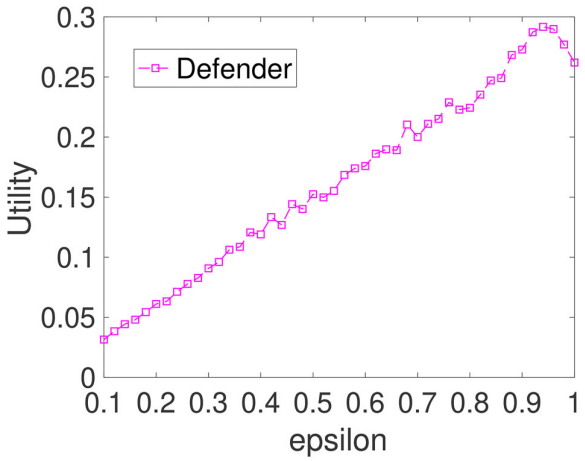

Figure 2. Utilities of the defender in beta distribution.

Then, we can use numerical method such as bisection to search for the solution of (5.3) and get $\epsilon^* = 0.93$.

The simulation result below shows a similar result. We first simulate the utility of the attacker when the defender chooses the strategy δ_{ϵ^*} with $\epsilon^* = 0.93$. We can see from Figure 1 that the utility of the attacker is maximized at $\epsilon = 0.93$. Similarly, we simulate the defender's utility when the attacker fixes $\epsilon^* = 0.93$. The resulting defender's utility is shown in Figure 2. We can see it is maximized when ϵ is closed to 0.93.

5.2. Mean Attack to Gaussian Distributions

Let $\mathcal{P}$ be a set of normal distributions $\mathcal{N}(\theta_0, 1)$, where $\theta_0 \in [-1, 1]$, and $\mathcal{Q}$ be another set of normal distributions $\mathcal{N}(\theta_1, 1)$, where $\theta_1 \in [2, 4]$. By Theorem 4.1, if there exist θ_0^*

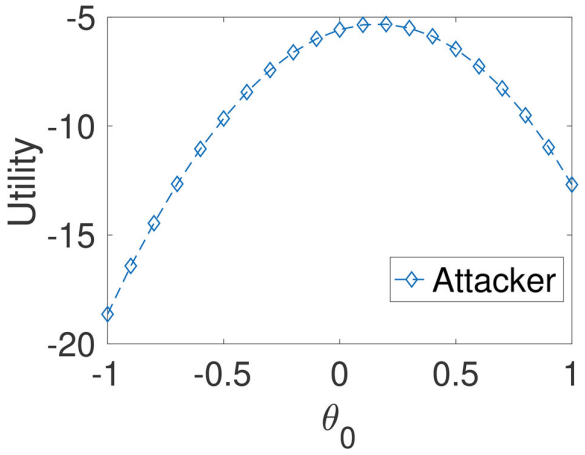

Figure 3. Utilities of the attacker at $\theta_1 = 2.0625$ of mean attack in Gaussian distribution.

and θ_1^* satisfying (4.1), (4.2), and (4.3); that is,

$$(\theta_0^*, \theta_1^*) = \underset{\theta_0 \in [-1,1], \theta_1 \in [2,4]}{\arg\max} \Big\{ -c_0(\theta_0) - c_1(\theta_1) + (\gamma+1)\Big[(\theta_0^* + \theta_1^*)^2/2 - (\theta_0 + \theta_1)(\theta_0^* + \theta_1^*) + 2\theta_0\theta_1\Big]\Big\} \tag{5.4}$$

,

$$\inf_{\theta_0 \in [-1,1]} \left[-(\theta_0^{*2} - \theta_1^{*2})/2 + (\theta_0^* - \theta_1^*)\theta_0\right] > 0, \tag{5.5}$$

and

$$\inf_{\theta_1 \in [2,4]} \left[-(\theta_1^{*2} - \theta_0^{*2})/2 + (\theta_1^* - \theta_0^*)\theta_1\right] > 0, \tag{5.6}$$

then let $\kappa^* = (\theta_0^*, \theta_1^*)$ and $(\delta_{\kappa^*}, \kappa^*)$ is in Nash equilibrium asymptotically. Specifically, if we set $c_0(\theta_0) = 10\theta_0^2$, $c_1(\theta_1) = 2(\theta_1 - 3)^2$, and $\gamma = 1$, we have $\theta_0^* = 0.1875$ and $\theta_1^* = 2.0625$, which correspond to $\mathcal{N}(0.1875, 1)$ in $\mathcal{P}$ and $\mathcal{N}(2.0625, 1)$ in $\mathcal{Q}$. If we set $c_0(\theta_0) = \theta_0^2$, $c_1(\theta_1) = 5(\theta_1 - 2.2)^2$, and $\gamma = 1$, we have $\theta_0^* = 1$ and $\theta_1^* = 2$, which correspond to $\mathcal{N}(1, 1)$ in $\mathcal{P}$ and $\mathcal{N}(2, 1)$ in $\mathcal{Q}$.

We then simulate the adversary's utility and the defender's utility by 10^5 Monte Carlo simulations to validate our theoretical results. We first simulate the adversary's utilities when the defender fixes its strategy to be the SPRT testing p^* to q^*. Figure 3 shows the simulated adversary's utility when the attacker chooses strategy $(\theta_0, \theta_1^* = 2.0625)$ with varies $\theta_0 \in [-1, 1]$. From Figure 3, we can see that the attacker gets its utility maximized when θ_0 is close to $\theta_0^* = 0.1875$. Similarly, in Figure 4, we plot the adversary's utilities when the attacker chooses its strategy $(\theta_0^* = 0.1875, \theta_1)$ with varies $\theta_1 \in [2, 4]$. From Figure 4, we can see that the attacker gets its utility maximized when θ_1 is close to $\theta_1^* = 2.0625$.

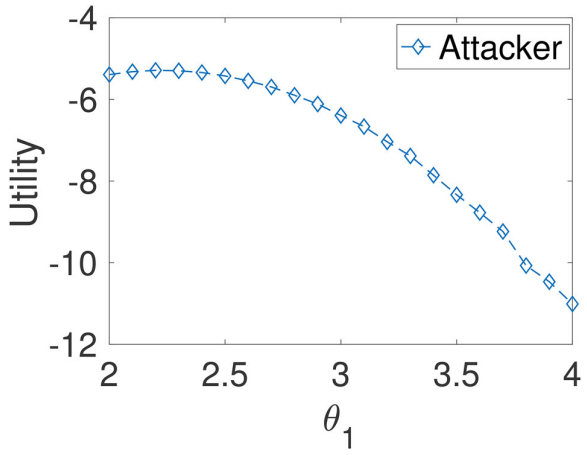

Figure 4. Utilities of the attacker at $\theta_0 = 0.1875$ of mean attack in Gaussian distribution.

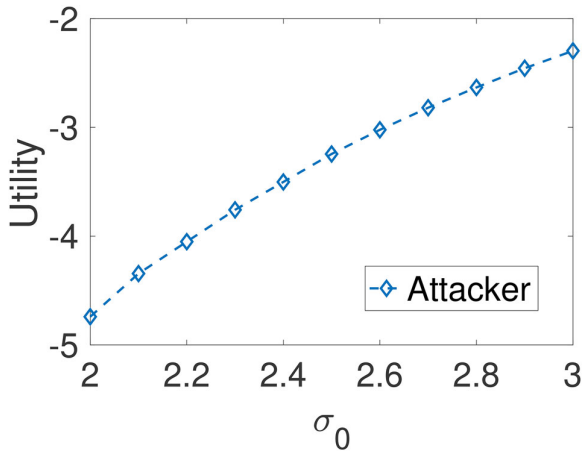

Figure 5. Utilities of the attacker at $\sigma_1 = 9$ of variance attack in Gaussian distribution.

5.3. Variance Attack to Gaussian Distributions

Let $\mathcal{P}$ be a set of distributions $\mathcal{N}(0, \sigma_0^2)$, where $\sigma_0 \in [a, b]$, and $\mathcal{Q}$ be a set of distributions $\mathcal{N}(0, \sigma_1^2)$, where $\sigma_1 \in [c, d]$. Assume that $0 < a < b < c < d$. Based on Theorem 4.1, we get the following corollary, which provides a pair of strategies that are in Nash equilibrium.

Corollary 5.2. *Assume that $c_0(\sigma)$ is a decreasing function—that is, $c_0'(\sigma) \leq 0$ for all $\sigma \in [a, b]$,—and $c_1(\sigma)$ is an increasing function; that is, $c_1'(\sigma) \geq 0$ for all $\sigma \in [c, d]$. Let $\gamma = 1$, $\sigma_0^* = b$, $\sigma_1^* = c$ and $\kappa^* = (\sigma_0^*, \sigma_1^*)$. Then $(\delta_{\kappa^*}, \kappa^*)$ is in Nash equilibrium asymptotically.*

The proof of Corollary 5.2 is postponed to Appendix A.1.

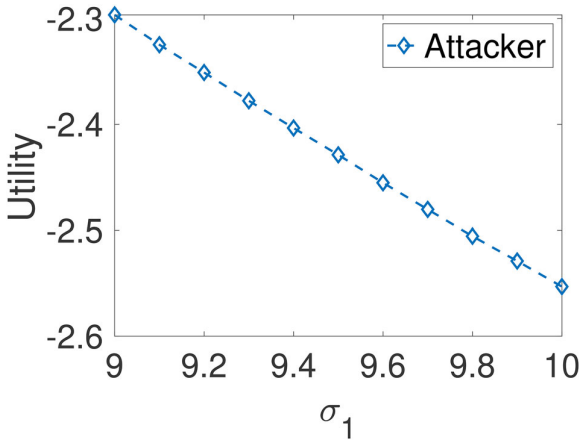

Figure 6. Utilities of the attacker at $\sigma_0 = 3$ of variance attack in Gaussian distribution.

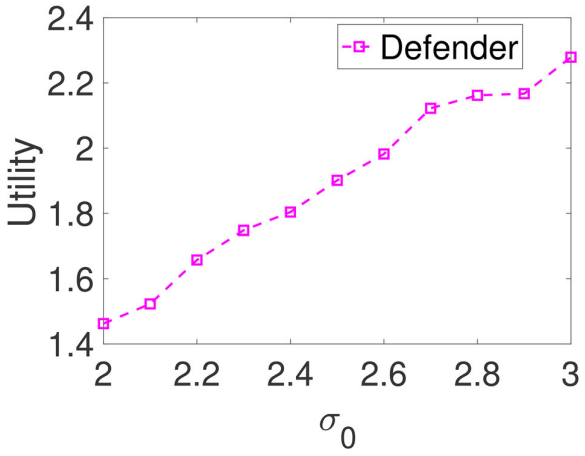

Figure 7. Utilities of the defender at $\sigma_1 = 9$ of variance attack in Gaussian distribution.

Figures 5 and 6 show the utility of the attacker when the defender fixes its strategy to be δ_{κ^*}, and Figures 7 and 8 show the utility of the defender when the attacker fixes its strategy to be κ^*. Specifically, we choose $\sigma_0 \in [2,3]$ and $\sigma_1 \in [9,10]$ with $c_0(\sigma_0) = c_1(\sigma_1) = 0$. Thus, by Corollary 5.2, $\sigma_0^* = 3$ and $\sigma_1^* = 9$ so that $p^* = \mathcal{N}(0,9)$ and $q^* = \mathcal{N}(0,81)$. Figure 5 shows the utility of the attacker at σ_1^* when the defender fixes its strategy to be $\delta_{\kappa^*} = (p^*, q^*)$, and it can be seen that when $\sigma_0 = \sigma_0^* = 3$, the utility of the attacker reaches its maximum. Figure 6 shows the utility of the attacker at σ_0^* when the defender chooses its strategy at $\delta_{\kappa^*} = (p^*, q^*)$, and it can be seen that the utility of the attacker reaches the maximum at $\sigma_1^* = 9$. Therefore, Figures 5 and 6 show that when the defender fixes its strategy to be δ_{κ^*}, the utility of the attacker reaches the maximum at $\sigma_0^* = 3$ and $\sigma_1^* = 9$. In Figure 7, we plot the utility of the defender at σ_1^* when the attacker chooses $\kappa^* = (\sigma_0^* = 3, \sigma_1^* = 9)$. We can see that the utility of the

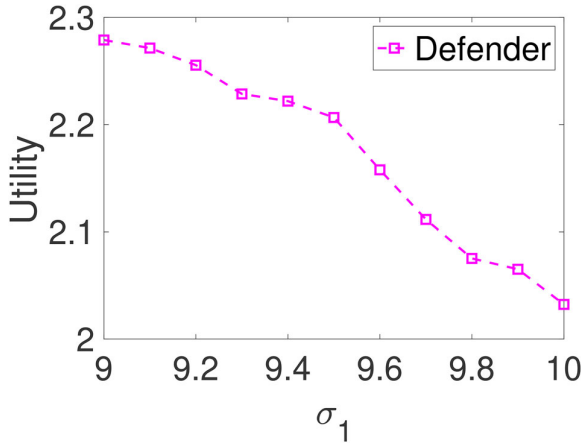

Figure 8. Utilities of the defender at $\sigma_0 = 3$ of variance attack in Gaussian distribution.

defender reaches the maximum at $\sigma_0 = 3$, which is σ_0^*. Figure 8 shows that the utility of the defender reaches the maximum at $\sigma_1 = 9$, which is σ_1^*, when the attacker chooses $\kappa^* = (\sigma_0^* = 3, \sigma_1^* = 9)$. Therefore, Figures 7 and 8 indicate that when the attacker fixes its strategy to be κ^*, the utility of the defender is maximized by $\sigma_0^* = 3$ and $\sigma_1^* = 9$. Hence, $(\kappa^*, \delta_{\kappa^*})$ is a pair of strategies in Nash equilibrium.

5.4. Mean Attack to Bernoulli Distributions

Let $\mathcal{P}$ be a set of distributions Bernoulli(θ_1), where $\theta_1 \in [a, b]$, and $\mathcal{Q}$ be a set of distributions Bernoulli(θ_2), where $\theta_2 \in [c, d]$, and assume that $0 < a < b < c < d$; that is, $\theta_1 < \theta_2$. Based on Theorem 4.1, we can find a pair of strategies in Nash equilibrium as stated in the following corollary.

Corollary 5.3. *Assume that $c(\theta_1)$ is a decreasing function—that is, $c'(\theta_1) \leq 0$ for all $\theta \in [a, b]$,—and $c(\theta_2)$ is an increasing function; that is, $c'(\theta_2) \geq 0$ for all $\theta \in [c, d]$. Let $\theta_1^* = b$, $\theta_2^* = c$, and $\kappa^* = (\theta_1^*, \theta_2^*)$. Then $(\delta_{\kappa^*}, \kappa^*)$ is in Nash equilibrium asymptotically.*

The proof of Corollary 5.3 is postponed to Appendix A.2.

In the experiments, we choose $\theta_1 \in [0.2, 0.3]$ and $\theta_2 \in [0.5, 0.6]$ with $c(\theta_1) = c(\theta_2) = 0$; then by Corollary 5.3, $\theta_1^* = 0.3$ and $\theta_2^* = 0.5$ so that $p^* =$ Bernoulli(0.3) and $q^* =$ Bernoulli(0.5). In Figures 9 and 10, we plot the utilities for the attacker when the attacker changes its strategy and the defender fixes its strategy to be $\delta_{\kappa^*} = (p^*, q^*)$. In Figure 9, the utility of the attacker is maximized by $\theta_1 = 0.3 = \theta_1^*$, and in Figure 10, the utility of the attacker is maximized by $\theta_2 = 0.5 = \theta_2^*$. Figures 11 and 12 show the utility for the defender when the defender changes its strategy and the attacker fixes its strategy to be $\kappa^* = (p^*, q^*)$. In Figure 11, the utility of defender is maximized by $\theta_1 = 0.3 = \theta_1^*$, whereas in Figure 12, the utility of the defender is maximized by $\theta_2 = 0.5 = \theta_2^*$. Therefore, $\kappa^*, \delta_{\kappa^*}$ provides a pair of strategies in Nash equilibrium.

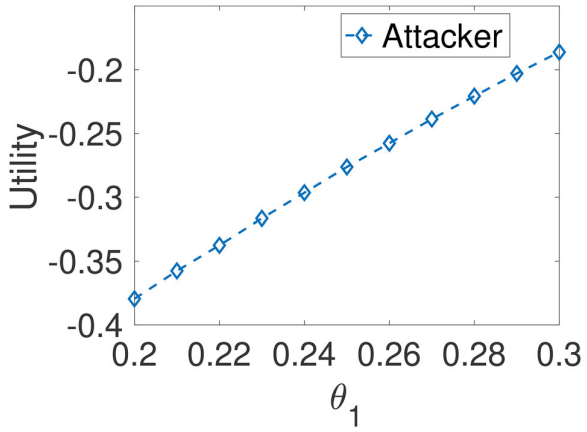

Figure 9. Utilities of the attacker in at $\theta_2 = 0.5$ of mean attack in Bernoulli distribution.

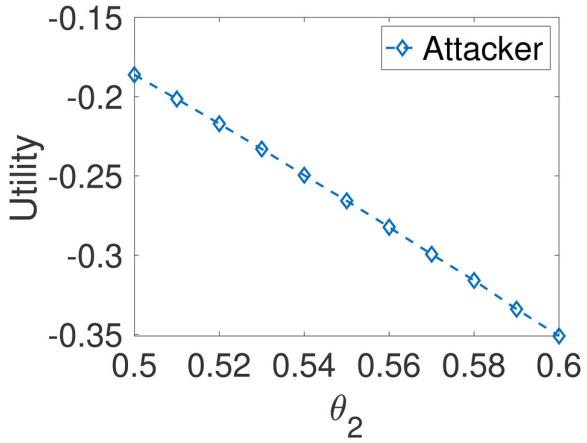

Figure 10. Utilities of the attacker at $\theta_1 = 0.3$ of mean attack in Bernoulli distribution.

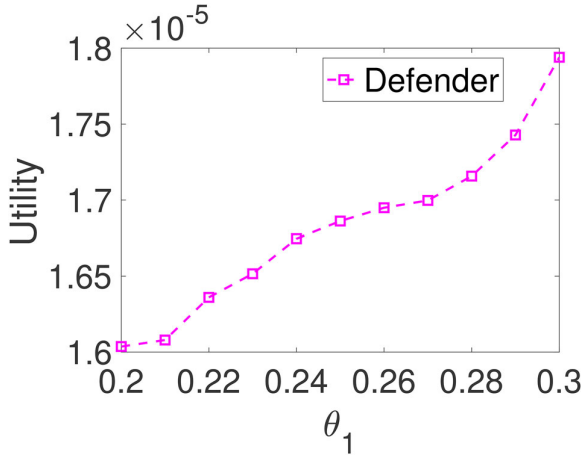

Figure 11. Utilities of the defender at $\theta_2 = 0.5$ of mean attack in Bernoulli distribution.

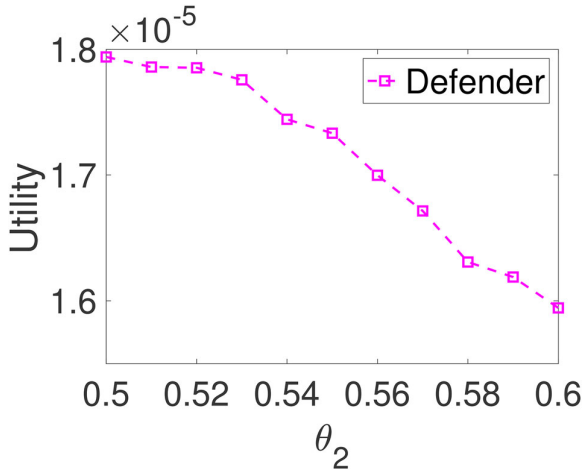

Figure 12. Utilities of the defender at $\theta_1 = 0.3$ of mean attack in Bernoulli distribution.

6. CONCLUSION

In this article, we proposed a game-theoretic approach for sequential detection in adversarial environments. We derived a pair of strategies and proved that they are in Nash equilibrium asymptotically. The analysis in this article, especially the analysis of the adversary's utility function, is in the asymptotic regime with vanishing type I and type II errors. It is of future interest to develop Nash equilibrium strategies in the nonasymptotic setting. In this article, we assume that the adversary chooses distributions $p \in \mathcal{P}$ and $q \in \mathcal{Q}$, and the samples are then generated i.i.d. by p and q under the two hypotheses, respectively; that is, the adversary is memoryless. For the adversary with memory, we can model the strategy of the adversarial as a finite-order Markov process and using the transition probability as the new observation; however, the analyses will be more challenging. We also note that in some cases, it may not be possible to find a pair of p^* and q^* satisfying (4.1) and (4.2) simultaneously.

APPENDIX

A.1. Proof of Corollary 5.2

Let $\sigma_0^* = b$ and $\sigma_1^* = c$. By (4.1) and Theorem 4.1, define $f(\sigma_0, \sigma_1) = -c_0(\sigma_0) - c_1(\sigma_1) - \lambda_2(\sigma_1)[\log \frac{\sigma_1^*}{\sigma_0^*} + \frac{\sigma_0^{*2} - \sigma_1^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}\sigma_0^2] - \lambda_1(\sigma_0)[\log \frac{\sigma_0^*}{\sigma_1^*} + \frac{\sigma_1^{*2} - \sigma_0^{*2}}{1\sigma_0^{*2}\sigma_1^{*2}}\sigma_1^2]$, where $\lambda_1(\sigma_0), \lambda_2(\sigma_1)$ are constants satisfying (4.12) and (4.16); that is,

$$\lambda_1(\sigma_0)\log\frac{\sigma_0^*}{\sigma_1^*} = \frac{1}{2}\log\left(\sigma_0^{*2}\sigma_1^{*2} - \lambda_1(\sigma_0)\sigma_0^2(\sigma_1^{*2} - \sigma_0^{*2})\right) - \log\left(\sigma_0^*\sigma_1^*\right), \tag{A.1}$$

$$\lambda_2(\sigma_1)\log\frac{\sigma_1^*}{\sigma_0^*} = \frac{1}{2}\log\left(\sigma_0^{*2}\sigma_1^{*2} - \lambda_2(\sigma_1)\sigma_1^2(\sigma_0^{*2} - \sigma_1^{*2})\right) - \log\left(\sigma_0^*\sigma_1^*\right). \tag{A.2}$$

In the proof below, we will show that $\sigma_0 = \sigma_0^* = b, \sigma_1 = \sigma_1^* = c$ maximize function $f(\sigma_0, \sigma_1)$. Thus, by Theorem 4.1, $(\delta_{\kappa^*}, \kappa^*)$ is in Nash equilibrium asymptotically.

Note that

$$\frac{\partial f}{\partial \sigma_0^2} = -c_0'(\sigma_0) - \lambda_2(\sigma_1)\frac{\sigma_0^{*2} - \sigma_1^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}} - \frac{\partial \lambda_1(\sigma_0)}{\partial \sigma_0^2}\left[\log\frac{\sigma_0^*}{\sigma_1^*} + \frac{\sigma_1^{*2} - \sigma_0^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}\sigma_1^2\right], \tag{A.3}$$

and

$$\frac{\partial f}{\partial \sigma_1^2} = -c_1'(\sigma_1) - \frac{\partial \lambda_2(\sigma_1)}{\partial \sigma_1^2}\left[\log\frac{\sigma_1^*}{\sigma_0^*} + \frac{\sigma_0^{*2} - \sigma_1^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}\sigma_0^2\right] - \lambda_1(\sigma_0)\frac{\sigma_1^{*2} - \sigma_0^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}. \tag{A.4}$$

First, we will show that $\frac{\partial f}{\partial \sigma_0^2} > 0$ so that f is an increasing function of σ_0^2. Note that $c_0'(\sigma_0) \le 0$ and $-\lambda_2(\sigma_1)\frac{\sigma_0^{*2} - \sigma_1^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}} > 0$ because $\sigma_1^{*2} > \sigma_0^{*2}$. Moreover, we have that the inequality

$$\log\frac{\sigma_0^*}{\sigma_1^*} + \frac{\sigma_1^{*2} - \sigma_0^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}\sigma_1^2 > 0 \tag{A.5}$$

holds for $\sigma_1 \in [c, d]$. Thus, the condition in (4.3) will always hold.

Next, we will show that $\frac{d\lambda_1(\sigma_0)}{d\sigma_0^2} < 0$. By (A.1) we have that

$$\sigma_0^2 = \frac{\sigma_0^{*2}\sigma_1^{*2} - \exp\left[2\lambda_1(\sigma_0)\log\frac{\sigma_0^*}{\sigma_1^*} + 2\log(\sigma_0^*\sigma_1^*)\right]}{\lambda_1(\sigma_0) \times (\sigma_1^{*2} - \sigma_0^{*2})}. \tag{A.6}$$

Taking the derivatives of σ_0^2 with respect to λ_1 yields

$$\frac{\partial \sigma_0^2}{\partial \lambda_1} = \frac{(xe^{-x} - 1 + e^{-x})(\sigma_0^*\sigma_1^*)^2(\sigma_1^{*2} - \sigma_0^{*2})}{\lambda_1^2(\sigma_0) \times (\sigma_1^{*2} - \sigma_0^{*2})^2}, \tag{A.7}$$

where $x = 2\lambda_1(\sigma_0)\log\frac{\sigma_1^*}{\sigma_0^*}$. Note for $x > 0$,

$$xe^{-x} - 1 + e^{-x} = (x+1)e^{-x} - 1 < e^x e^{-x} - 1 = 0,$$

and all of the other terms on the right-hand side of (A.7) are positive. Therefore, we have $\frac{\partial \sigma_0^2}{\partial \lambda_1} < 0$, which implies $\frac{\partial \lambda_1}{\partial \sigma_0^2} < 0$, and $\frac{\partial f}{\partial \sigma_0^2} > 0$.

Similarly, we can show that $\frac{\partial f}{\partial \sigma_1^2} < 0$ so that f is a decreasing function of σ_1^2. Note that $c_1'(\sigma_1) \ge 0$ and $-\lambda_1(\sigma_0)\frac{\sigma_1^{*2} - \sigma_0^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}} < 0$ because $\sigma_1^{*2} > \sigma_0^{*2}$. Moreover, we have that

$$\log\frac{\sigma_1^*}{\sigma_0^*} + \frac{\sigma_0^{*2} - \sigma_1^{*2}}{2\sigma_0^{*2}\sigma_1^{*2}}\sigma_0^2 > 0 \tag{A.8}$$

holds for any $\sigma_0 \in [a, b]$. Thus, condition (4.2) will always hold. Next, we will show that $\frac{\partial \lambda_2}{\partial \sigma_1^2} > 0$. By (A.2) we have

$$\sigma_1^2 = \frac{\sigma_0^{*2}\sigma_1^{*2} - \exp\left[2\lambda_2(\sigma_1)\log\frac{\sigma_1^*}{\sigma_0^*} + 2\log(\sigma_0^*\sigma_1^*)\right]}{\lambda_2(\sigma_1) \times (\sigma_0^{*2} - \sigma_1^{*2})}. \tag{A.9}$$

Taking the derivatives of σ_1^2 with respect to $\lambda_2(\sigma_1)$ and letting $x = 2\lambda_2(\sigma_1)\log\frac{\sigma_1^*}{\sigma_0^*}$ yields

$$\frac{\partial \sigma_1^2}{\partial \lambda_2} = \frac{(xe^x - 1 + e^x)(\sigma_0^*\sigma_1^*)^2(\sigma_1^{*2} - \sigma_0^{*2})}{[\lambda_2(\sigma_1)]^2(\sigma_1^{*2} - \sigma_0^{*2})^2}. \tag{A.10}$$

Note for $x > 0$, $xe^x - 1 + e^x > 0$, and all of the other terms on the right-hand side of (A.10) are positive, so $\frac{\partial \sigma_1^2}{\partial \lambda_2} > 0$. Therefore, $\frac{\partial f}{\partial \sigma_1^2} < 0$.

Because $\frac{\partial f}{\partial \sigma_0^2} > 0$ and $\frac{\partial f}{\partial \sigma_1^2} < 0$, $f(\sigma_0, \sigma_1)$ is maximized when $\sigma_0 = b$ and $\sigma_1 = c$, which completes the proof.

A.2. Proof of Corollary 5.3

Let $\theta_1^* = b$ and $\theta_2^* = c$. By (4.1) and Theorem 4.1, define

$$\begin{aligned} f(\theta_1,\theta_2) = &-c(\theta_1) - c(\theta_2) - \gamma\lambda_2(\theta_2)\left\{\log\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right) + \theta_1 \log\left[\frac{\theta_1^*(1-\theta_2^*)}{\theta_2^*(1-\theta_1^*)}\right]\right\} \\ &- \lambda_1(\theta_1)\left\{\log\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right) + \theta_2 \log\left[\frac{\theta_2^*(1-\theta_1^*)}{\theta_1^*(1-\theta_2^*)}\right]\right\}, \end{aligned} \tag{A.11}$$

where $\lambda_1(\theta_1), \lambda_2(\theta_2)$ are constants satisfying (4.12) and (4.16); that is,

$$\theta_1\left(\frac{\theta_2^*}{\theta_1^*}\right)^{\lambda_1} + (1-\theta_1)\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right)^{\lambda_1} = 1, \tag{A.12}$$

$$\theta_2\left(\frac{\theta_1^*}{\theta_2^*}\right)^{\lambda_2} + (1-\theta_2)\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right)^{\lambda_2} = 1. \tag{A.13}$$

In the proof below, we will show that $\theta_1 = \theta_1^* = b, \theta_2 = \theta_2^* = c$ maximize function $f(\theta_1,\theta_2)$. Let $\kappa^* = (\theta_1^*, \theta_2^*)$. Then, by Theorem 4.1, $(\delta_{\kappa^*}, \kappa^*)$ is in Nash equilibrium asymptotically.

Note that

$$\frac{\partial f}{\partial \theta_1} = -c'(\theta_1) - \gamma\lambda_2(\theta_2)\log\left[\frac{\theta_1^*(1-\theta_2^*)}{\theta_2^*(1-\theta_1^*)}\right] - \frac{\partial\lambda_1(\theta_1)}{\partial\theta_1}\left\{\log\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right) + \theta_2\log\left[\frac{\theta_2^*(1-\theta_1^*)}{\theta_1^*(1-\theta_2^*)}\right]\right\}, \tag{A.14}$$

and

$$\frac{\partial f}{\partial \theta_2} = -c'(\theta_2) - \lambda_1(\theta_1)\log\left[\frac{\theta_2^*(1-\theta_1^*)}{\theta_1^*(1-\theta_2^*)}\right] - \frac{\partial\lambda_2(\theta_2)}{\partial\theta_2}\gamma\left\{\log\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right) + \theta_1\log\left[\frac{\theta_1^*(1-\theta_2^*)}{\theta_2^*(1-\theta_1^*)}\right]\right\}. \tag{A.15}$$

First, we show that $\frac{\partial f}{\partial\theta_1} > 0$, and thus f is an increasing function of θ_1. Note that $c'(\theta_1) < 0$ and $-\gamma\lambda_2(\theta_2)\log\left[\frac{\theta_1^*(1-\theta_2^*)}{\theta_2^*(1-\theta_1^*)}\right] > 0$. Moreover, we have that

$$\log\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right) + \theta_2\log\left[\frac{\theta_2^*(1-\theta_1^*)}{\theta_1^*(1-\theta_2^*)}\right] > 0 \tag{A.16}$$

holds for any $\theta_2 \in [c,d]$. Therefore, the condition in (4.3) always hold. Next, we will show that $\frac{\partial\lambda_1}{\partial\theta_1} < 0$ so that the right-hand side of (A.14) is positive. By (A.12) we have that

$$\theta_1 = \frac{\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right)^{\lambda_1} - 1}{\left(\frac{1-\theta_2^*}{1-\theta_1^*}\right)^{\lambda_1} - \left(\frac{\theta_2^*}{\theta_1^*}\right)^{\lambda_1}}. \tag{A.17}$$

Let $x = \log\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right)$ and $y = \log\left(\frac{\theta_2^*}{\theta_1^*}\right)$. Clearly, $x, y > 0$. Rewriting (A.17), we can get

$$\theta_1 = \frac{e^{-x\lambda_1} - 1}{e^{-x\lambda_1} - e^{y\lambda_1}}. \tag{A.18}$$

Taking the derivatives of θ_1 with respect to $\lambda_1(\theta_1)$ yields

$$\begin{aligned} \frac{\partial\theta_1}{\partial\lambda_1} &= \frac{-xe^{-\lambda_1 x}(1-e^{\lambda_1 y}) + ye^{\lambda_1 y}(e^{-\lambda_1 x} - 1)}{(e^{\lambda_1 y} - e^{-\lambda_1 x})^2} \\ &= \frac{xe^{-\lambda_1 x}e^{\lambda_1 y}(1-e^{-\lambda_1 y}) + ye^{\lambda_1 y}e^{-\lambda_1 x}(1-e^{\lambda_1 x})}{(e^{\lambda_1 y} - e^{-\lambda_1 x})^2}, \end{aligned} \tag{A.19}$$

which is always negative because

$$xe^{-\lambda_1 x}e^{\lambda_1 y}(1-e^{-\lambda_1 y})+ye^{\lambda_1 y}e^{-\lambda_1 x}(1-e^{\lambda_1 x})<xe^{-\lambda_1 x}e^{\lambda_1 y}\lambda_1 y-ye^{\lambda_1 y}e^{-\lambda_1 x}\lambda_1 x=0. \tag{A.20}$$

Hence, $\frac{\partial f}{\partial \theta_1}>0$.

Next, we show that $\frac{\partial f}{\partial \theta_1}>0.$, and thus f is a decreasing function of θ_2. Note that $c'(\theta_2)>0$ and $-\lambda_1(\theta_1)\log\left[\frac{\theta_2^*(1-\theta_1^*)}{\theta_1^*(1-\theta_2^*)}\right]<0$ because $\theta_1^*<\theta_2^*$ by the assumption. Moreover, we have

$$\log\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right)+\theta_1\log\left[\frac{\theta_1^*(1-\theta_2^*)}{\theta_2^*(1-\theta_1^*)}\right]>0, \tag{A.21}$$

for any $\theta_1\in[a,b]$. Therefore, condition (4.2) always holds. Now we need to show that $\frac{\partial\lambda_2}{\partial\theta_2}>0$ so that the right-hand side of (A.15) is negative. By (A.13) we can get

$$\theta_2=\frac{\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right)^{\lambda_2}-1}{\left(\frac{1-\theta_1^*}{1-\theta_2^*}\right)^{\lambda_2}-\left(\frac{\theta_1^*}{\theta_2^*}\right)^{\lambda_2}}. \tag{A.22}$$

Recalling $x=\log\frac{1-\theta_1^*}{1-\theta_2^*}$ and $y=\log\frac{\theta_2^*}{\theta_1^*}, x>0, y>0,$ we can get

$$\theta_2=\frac{e^{x\lambda_2}-1}{e^{x\lambda_2}-e^{-y\lambda_2}}. \tag{A.23}$$

Taking the derivative of θ_2 with respect to λ_2, we have that

$$\begin{aligned}\frac{\partial\theta_2}{\partial\lambda_2}&=\frac{xe^{\lambda_2 x}(1-e^{\lambda_2 x})-ye^{-\lambda_2 y}(e^{\lambda_2 x}-1)}{\left(e^{\lambda_2 x}-e^{-\lambda_2 y}\right)^2}\\&=\frac{xe^{\lambda_2 x}e^{-\lambda_2 y}(e^{\lambda_2 y}-1)+ye^{-\lambda_2 y}e^{\lambda_2 x}(e^{-\lambda_2 x}-1)}{\left(e^{\lambda_2 x}-e^{-\lambda_2 y}\right)^2},\end{aligned} \tag{A.24}$$

which is always positive because

$$xe^{\lambda_2 x}e^{-\lambda_2 y}(e^{\lambda_2 y}-1)+ye^{-\lambda_2 y}e^{\lambda_2 x}(e^{-\lambda_2 x}-1)>xe^{\lambda_2 x}e^{-\lambda_2 y}\lambda_2 y-ye^{-\lambda_2 y}e^{\lambda_2 x}\lambda_2 x=0. \tag{A.25}$$

Hence, $\frac{\partial f}{\partial\theta_2}<0$.

Because $\frac{\partial f}{\partial\theta_1}>0$ and $\frac{\partial f}{\partial\theta_2}<0$, $f(\theta_1,\theta_2)$ is maximized when $\theta_1=b$ and $\theta_2=c$, which completes the proof.

ACKNOWLEDGMENTS

The authors thank two anonymous reviewers and the Associate Editor for their thoughtful and constructive comments.

DISCLOSURE

The authors have no conflicts of interest to report.

FUNDING

R. Zhang's research was supported in part by National Science Foundation (NSF) Grant ECCS-2112740. The work of S. Zou was supported in part by the NSF under Grants CCF-1948165, CCF-2106560, and ECCS-2112693.

REFERENCES

Barni, M., and B. Tondi. 2013. "The Source Identification Game: An Information-Theoretic Perspective." *IEEE Transactions on Information Forensics and Security* 8 (3):450–63. doi:10.1109/TIFS.2012.2237397

Bayraktar, E., and L. Lai. 2015. "Byzantine Fault Tolerant Distributed Quickest Change Detection." *SIAM Journal on Control and Optimization* 53 (2):575–91. doi:10.1137/130924445

Cover, T. M., and J. A. Thomas. 2012. *Elements of Information Theory*. New York, NY: John Wiley & Sons.

Fellouris, G., E. Bayraktar, and L. Lai. 2018. "Efficient Byzantine Sequential Change Detection." *IEEE Transactions on Information Theory* 64 (5):3346–60. doi:10.1109/TIT.2017.2755025

Gao, R., L. Xie, Y. Xie, and H. Xu. 2018. "Robust Hypothesis Testing Using Wasserstein Uncertainty Sets." Paper presented at the Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), pp. 7902–7912.

Gül, G., and A. M. Zoubir. 2017. "Minimax Robust Hypothesis Testing." *IEEE Transactions on Information Theory* 63 (9):1–5587. doi:10.1109/TIT.2017.2693198

Huang, Y.-C., S.-C. Lin, and Y.-J. Huang. 2019. "A Tight Converse to the Asymptotic Performance of Byzantine Distributed Sequential Change Detection." Paper presented at the Proceedings of the IEEE International Symposium on Information Theory (ISIT), pp. 2404–2408.

Huang, Y.-J., S.-C. Lin, and Y.-C. Huang. 2019. "On Byzantine Distributed Sequential Change Detection with Multiple Hypotheses." Paper presented at the Proceedings of the IEEE International Symposium on Information Theory (ISIT), pp. 2209–2213.

Huber, P. J. 1965. "A Robust Version of the Probability Ratio Test." *The Annals of Mathematical Statistics* 36 (6):1753–8. doi:10.1214/aoms/1177699803

Huber, P. J. 1981. *Robust Statistics*. New York, NY, USA: Wisley.

Jin, Y., and L. Lai. 2021. "On the Adversarial Robustness of Hypothesis Testing." *IEEE Transactions on Signal Processing* 69:515–30. doi:10.1109/TSP.2020.3045206

Kay, S. M. 1993. *Fundamentals of Statistical Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall.

Levy, B. C. 2009. "Robust Hypothesis Testing with a Relative Entropy Tolerance." *IEEE Transactions on Information Theory* 55 (1):413–21. doi:10.1109/TIT.2008.2008128

Li, Z., Y. Mo, and F. Hao. 2019. "Game Theoretical Approach to Sequential Hypothesis Test with Byzantine Sensors." In *2019 IEEE 58th Conference on Decision and Control (CDC)*, 2654–2659. IEEE.

Marano, S., V. Matta, and L. Tong. 2008. "Distributed Detection in the Presence of Byzantine Attacks." *IEEE Transactions on Signal Processing* 57 (1):16–29. doi:10.1109/TSP.2008.2007335

Mo, Y., J. P. Hespanha, and B. Sinopoli. 2014. "Resilient Detection in the Presence of Integrity Attacks." *IEEE Transactions on Signal Processing* 62 (1):31–43. doi:10.1109/TSP.2013.2284145

Molloy, T. L., and J. J. Ford. 2017. "Misspecified and Asymptotically Minimax Robust Quickest Change Detection." *IEEE Transactions on Signal Processing* 65 (21):5730–42. doi:10.1109/TSP.2017.2740202

Moulin, P., and V. V. Veeravalli. 2018. *Statistical Inference for Engineers and Data Scientists*. Cambridge: Cambridge University Press.

Pandit, C., S. Meyn, and V. V. Veeravalli. 2004. "Asymptotic Robust Neyman-Pearson Hypothesis Testing Based on Moment Classes." Paper presented at the Proceedings of the IEEE International Symposium on Information Theory (ISIT), pp. 220.

Poor, H. V. 2013. *An Introduction to Signal Detection and Estimation*. New York: Springer.

Poor, H. V., and O. Hadjiliadis. 2009. *Quickest Detection*. Cambridge: Cambridge University Press.

Qin, Y., and C. E. Priebe. 2017. "Robust Hypothesis Testing via l_q-Likelihood." *Statistica Sinica* 27 (4):1793–813. doi:10.5705/ss.202015.0441

Reny, P. J. 2008. "Non-Cooperative Games: Equilibrium Existence." In *The New Palgrave Dictionary of Economics*, 2nd ed. Palgrave Macmillan.

Siegmund, D. 1985. *Sequential Analysis: Tests and Confidence Intervals*. New York: Springer.

Tartakovsky, A., I. Nikiforov, and M. Basseville. 2014. *Sequential Analysis: Hypothesis Testing and Changepoint Detection*. New York: Chapman and Hall/CRC.

Vamvoudakis, K. G., J. P. Hespanha, B. Sinopoli, and Y. Mo. 2014. "Detection in Adversarial Environments." *IEEE Transactions on Automatic Control* 59 (12):3209–23. doi:10.1109/TAC.2014.2351671

Veeravalli, V. V., T. Basar, and H. V. Poor. 1994. "Minimax Robust Decentralized Detection." *IEEE Transactions on Information Theory* 40 (1):35–40. doi:10.1109/18.272453

Wald, A. 1945. "Sequential Tests of Statistical Hypotheses." *The Annals of Mathematical Statistics* 16 (2):117–86. doi:10.1214/aoms/1177731118

Wald, A. 2013. *Sequential Analysis*. Mineola: Dover Publications.

Wald, A., and J. Wolfowitz. 1948. "Optimum Character of the Sequential Probability Ratio Test." *The Annals of Mathematical Statistics* 19 (3):326–39. doi:10.1214/aoms/1177730197

Wilcox, R. R. 2011. *Introduction to Robust Estimation and Hypothesis Testing*. Amsterdam: Academic Press.

Yasodharan, S., and P. Loiseau. 2019. "Nonzero-Sum Adversarial Hypothesis Testing Games." Paper presented at the Proc. Advances in Neural Information Processing Systems, pp. 7310–7320.

Zhang, R., and S. Zou. 2020. "A Game-Theoretic Approach to Sequential Detection in Adversarial Environments." Paper presented at the Proceedings of the IEEE International Symposium on Information Theory (ISIT), pp. 1153–1158.