

Nonzero-sum Adversarial Hypothesis Testing Games

Sarath Yasodharan, Patrick Loiseau

► **To cite this version:**

Sarath Yasodharan, Patrick Loiseau. Nonzero-sum Adversarial Hypothesis Testing Games. NeurIPS 2019 - Thirty-third Conference on Neural Information Processing Systems, Dec 2019, Vancouver, Canada. pp.1-23. hal-02299451

HAL Id: hal-02299451

<https://hal.inria.fr/hal-02299451>

Submitted on 27 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Nonzero-sum Adversarial Hypothesis Testing Games

Sarath Yasodharan

Department of Electrical Communication Engineering
Indian Institute of Science
Bangalore 560 012, India
sarath@iisc.ac.in

Patrick Loiseau

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG & MPI-SWS
700 avenue Centrale
Domaine Universitaire
38400 St Martin d'Hères, France
patrick.loiseau@inria.fr

Abstract

We study nonzero-sum hypothesis testing games that arise in the context of adversarial classification, in both the Bayesian as well as the Neyman-Pearson frameworks. We first show that these games admit mixed strategy Nash equilibria, and then we examine some interesting concentration phenomena of these equilibria. Our main results are on the exponential rates of convergence of classification errors at equilibrium, which are analogous to the well-known Chernoff-Stein lemma and Chernoff information that describe the error exponents in the classical binary hypothesis testing problem, but with parameters derived from the adversarial model. The results are validated through numerical experiments.

1 Introduction

Classification is a simple but important task that has numerous applications in a variety of domains such as computer vision or security. A traditional assumption that is used in the design of classification algorithms is that the input data is generated without knowledge of the classifier being used, hence the data distribution is independent of the classification algorithm. This assumption is no longer valid in the presence of an adversary, as an adversarial agent can learn the classifier and deliberately alter the data such that the classifier makes an error. This is the case in particular in security applications where the classifier's goal is to detect the presence of an adversary from the data it observes.

Adversarial classification has been studied in two main settings. The first focuses on adversarial versions of a standard classification task in machine learning, where the adversary attacks the classifier (defender/decision maker) by directly choosing vectors from a given set of data vectors; whereas the second focuses on adversarial hypothesis testing, where the adversary (attacker) gets to choose a distribution from a set of distributions and independent data samples are generated from this distribution. The main differences of the latter framework from the former are that: (i) the adversary only gets to choose a distribution (rather than the actual attack vector) and data is generated independently from this distribution, and (ii) the defender makes a decision only once after it observes a whole data sequence instead of making a decision for each individual data sample it receives. Both of these frameworks have applications in a variety of domains, but prior literature has mainly focused on the first setting; see Section 1.1 for a description of the related literature.

In this paper, we focus on the setting of adversarial hypothesis testing. To model the interaction between the attacker and defender, we formulate a nonzero-sum two-player game between the adversary and the classifier where the adversary picks a distribution from a given set of distributions, and data is generated independently from that distribution (a non-attacker always generates data from a fixed distribution). The defender on his side makes a decision based on observation of n data points. Our model can also be viewed as a game-theoretic extension of the classical binary hypothesis testing problem where the distribution under the alternate hypothesis is chosen by an adversary. Based on our game model, we are then able to extend to the adversarial setting the main results of the classical hypothesis testing problem (see Section 2) on the form of the best decision rule and on the rates of decrease of classification errors. More specifically, our contributions can be summarized as follows:

1. We propose nonzero-sum games to model adversarial hypothesis testing problems in a flexible manner.
2. We show existence of mixed strategy Nash equilibria in which the defender employs certain likelihood ratio tests similar to that used in the classical binary hypothesis testing problem.
3. We show that the classification errors under all Nash equilibria for our hypothesis testing games decay exponentially fast in the number of data samples. We analytically obtain these error exponents, and it turns out that they are same as those arising in certain classical hypothesis testing problem, with parameters derived from the adversarial model.
4. We illustrate the results, in particular the importance of some assumptions, using simulations.

Throughout our analysis, an important difficulty lies in that the strategy spaces of both the players are uncountable; we believe, however, that it is an important feature of the model to be realistic.

1.1 Related Work

Adversarial classification and the security of machine learning have been studied extensively in the past decade, see e.g., [9, 21, 4, 15, 19, 24]; here we focus only on game-theoretic approaches to tackle the problem. Note that, besides the adversarial learning problem, game theory has been successfully used to tackle several other security problems such as allocation of monitoring resources to protect targets, see e.g., [8, 18]. We review here only papers relating to classification.

A number of game-theoretic models have appeared in the past decade to study the adversarial classification problem in the classical setting of classification tasks in machine learning. [9] studies the best response in an adversarial classification game, where the adversary is allowed to alter training data. A number of zero-sum game models were also proposed where the attacker is restricted on the amount of modifications he can do to the training set, see [17, 31, 30]. [6] studies the problem of choosing the best linear classifier in the presence of an adversary (a similar model is also studied in [7]) using a nonzero-sum game, and shows the existence of a unique pure strategy Nash equilibrium. Similar to our formulation, the strategy sets in this case are uncountable, and therefore showing the existence and uniqueness of Nash equilibrium needs some work. However, in our formulation, there may not always exist a Nash equilibrium in pure strategies, which makes the subsequent analysis of error exponents more difficult. [20] studies an adversarial classification game where the utilities of the players are defined by using ROC curves. The authors study Nash equilibria for their model and provide numerical discretization techniques to compute the equilibria. [12] studies a nonzero-sum adversarial classification game where the defender has no restriction on the classifier, but the attacker is limited to a finite set of vectors. The authors show that the defender can, at equilibrium, use only a small subset of “threshold classifiers” and characterize the equilibrium through linear programming techniques. In our model, the utility functions share similarities with that of [12], but we work in the hypothesis testing framework and with uncountable action sets, which completely modifies the analysis. Several studies appeared recently on “strategic classification”, where the objective of the attacker(s) is to improve the classification outcome in his own direction, see [14, 11].

On the other hand, adversarial hypothesis testing has been studied by far fewer authors. [2] studies a source identification game in the presence of an adversary, where the classifier needs to distinguish between two source distributions P_0 and P_1 in which the adversary can corrupt samples from P_0 before it reaches the classifier. They show that the game has an asymptotic Nash equilibrium when the number of samples becomes large, and compute the error exponent associated with the false negative probability. [3] and [29] study further extensions of this framework.

A (non game-theoretic) hypothesis testing problem in an adversarial setting has been studied by [5], which is the closest to our work. Here, there are two sets of probability distributions and nature

outputs a fixed number of independent samples generated by using distributions from either one of these two sets. The goal of the classifier is to detect the true state of nature. The authors derive error exponents associated with the classification error, in both Bayesian and Neyman-Pearson frameworks using a worst-case maxmin analysis. Although we restrict to i.i.d. samples and let the non-attacker play a single distribution, we believe that our nonzero-sum game model with flexible utilities can better capture the interaction between adversary and classifier. There also exists extensive prior work within the statistics literature [16] on minimax hypothesis testing, which relates to our paper, but we defer a discussion of how our work differs from it to after we have exposed the details of our model.

Game-theoretic models were also used to study adversarial classification in a sequential setting, see [28, 1, 22], but with very different techniques and results.

2 Basic Setup and Hypothesis Testing Background

In this section, we present the basic setup results in classical binary hypothesis testing.

Throughout the paper, we consider an alphabet set \mathcal{X} that we assume finite. In a classical hypothesis testing problem, we are given two distributions p and q , and a realization of a sequence of independent and identically distributed random variables X_1, \dots, X_n , which are distributed as either p (under hypothesis H_0) or q (under hypothesis H_1). Our goal is to distinguish between the two alternatives:

$$H_0 : X_1, X_2, \dots, X_n \text{ i.i.d. } \sim p \quad \text{versus} \quad H_1 : X_1, X_2, \dots, X_n \text{ i.i.d. } \sim q.$$

In this setting, we could make two possible types of errors: (i) we declare H_1 , whereas the true state of nature is H_0 (Type I error, or false alarm), and (ii) we declare H_0 whereas the true state of nature is H_1 (Type II error, or missed detection). Note that one can make one of these errors arbitrarily small at the expense of the other by always declaring H_0 or H_1 .

The trade-off between the two types of errors can be captured using two frameworks. If we have knowledge on the prior probabilities of the two hypotheses, then we can seek a decision rule that minimizes the average probability of error (this is the Bayesian framework). On the other hand, if we do not have any information on the prior probabilities, then we can fix $\varepsilon > 0$ and seek a decision rule that minimizes the Type II error among all decision rules whose Type I error is at most ε (this is the Neyman-Pearson framework). In both of these frameworks, it can be shown that the optimal test is a likelihood ratio test, i.e., given $\mathbf{x}^n = (x_1, \dots, x_n)$ we compute the likelihood ratio $\frac{q(\mathbf{x}^n)}{p(\mathbf{x}^n)}$ and compare it to a threshold to make a decision (with possible randomization at the boundary in the Neyman-Pearson framework). Here, $p(\mathbf{x}^n)$ (resp. $q(\mathbf{x}^n)$) denotes the probability of observing the n -length word \mathbf{x}^n under the distribution p (resp. q). See Section II.B and II.D in [25] for an introduction to hypothesis testing.

For large enough n , by the law of large numbers, the fraction of i in an observation \mathbf{x}^n is very close to $p(i)$ (resp. $q(i)$) under H_0 (resp. under H_1), for each $i \in \mathcal{X}$. Therefore, one anticipates that the probability of correct decision is very close to 1 for large enough n . Hence, one can study the rate at which the errors go to 0 as n becomes large. It is shown that, under both frameworks, the error decays exponentially in n . In the Bayesian framework, the error exponent associated with the average probability of error is $-\Lambda_0^*(0)$, where $\Lambda_0^*(\cdot)$ is the Fenchel-Legendre transform of the log-moment generating function of the random variable $\frac{q(X)}{p(X)}$ under H_0 , i.e., when $X \sim p$. In the Neyman-Pearson case, the error exponent associated with the Type II error is $-D(p||q)$ where D is the relative entropy functional. The above error exponents are known as Chernoff information and Chernoff-Stein lemma, respectively (see Section 3.4 in [10] for the analysis on error exponents).

In this work, we propose extensions of the classical hypothesis testing framework to an adversarial scenario modeled as a game, both in the Bayesian and in the Neyman-Pearson frameworks; and we investigate how the corresponding results are modified. Due to space constraints, we present only the model and results for the Bayesian framework in the main body of the paper. The corresponding analysis for the Neyman-Pearson framework follows similar ideas and is relegated to Appendix A. The proofs of all results presented in the paper (and in Appendix A) can be found in Appendix B.

3 Hypothesis Testing Game in the Bayesian Framework

In this section, we formulate a one-shot adversarial hypothesis testing game in the Bayesian framework, motivated by security problems where there might be an attacker who modifies the data distribution and a defender who tries to detect the presence of the attacker. Game theoretic modelling of such problems has found great success in understanding the behavior of the agents via equilibrium analysis in many applications, see Section 1.1. We first present the model and then elaborate on its motivations and on how it relates to related works in statistics.

3.1 Problem Formulation

Let $\mathcal{X} = \{0, 1, \dots, d-1\}$ denote the alphabet set with cardinality d , and let $M_1(\mathcal{X})$ denote the space of probability distributions on \mathcal{X} . Fix $n \geq 1$.

The game is played as follows. There are two players: the external agent and the defender. The external agent can either be a non-attacker or an attacker. In the Bayesian framework, we assume that the external agent is an attacker with probability θ , and a non-attacker (normal user) with probability $1 - \theta$. The non-attacker is not strategic and she does not have any adversarial objective. If the external agent is a non-attacker, she generates n samples independently from the distribution p . If the external agent is an attacker, she picks a distribution q from a set of distributions $Q \subseteq M_1(\mathcal{X})$ and generates n samples independently from q . The defender, upon observing the n -length word generated by the external agent, wants to detect the presence of the attacker.

Throughout the paper, a decision rule implemented by the defender is denoted by $\varphi : \mathcal{X}^n \rightarrow [0, 1]$, with the interpretation that $\varphi(\mathbf{x}^n)$ is the probability with which hypothesis H_1 is accepted (i.e., the presence of an adversary is declared) when the defender observes the n -length word $\mathbf{x}^n = (x_1, \dots, x_n)$. We say that a decision rule φ is deterministic if $\varphi(\mathbf{x}^n) \in \{0, 1\}$ for all $\mathbf{x}^n \in \mathcal{X}^n$.

To define the game, let the attacker's strategy set be $Q \subseteq M_1(\mathcal{X})$, and that of the defender be

$$\Phi_n = \{\varphi : \mathcal{X}^n \rightarrow [0, 1]\},$$

which is the set of all randomized decision rules on n -length words.

To define the utilities, consider the attacker first. We assume that there is a cost associated with choosing a distribution from Q which we model using a cost function $c : Q \rightarrow \mathbb{R}_+$. The goal of the attacker is to fool the defender as much as possible, i.e., he wants to maximize the probability that the defender classifies an n -length word as coming from the non-attacker whereas it is actually being generated by the attacker. To capture this, the utility of the attacker when she plays the pure strategy $q \in Q$ and the defender plays the pure strategy $\varphi \in \Phi_n$ is defined as

$$u_n^A(q, \varphi) = \sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n))q(\mathbf{x}^n) - c(q), \quad (3.1)$$

where $q(\mathbf{x}^n)$ denotes the probability of observing the n -length word \mathbf{x}^n when the symbols are generated independently from the distribution q .

For the defender, the goal is to minimize the classification error. Similar to the classical hypothesis testing problem, there could be two types of errors: (i) the external agent is actually a non-attacker whereas the defender declares that there is an attack (Type I error, or false alarm), and (ii) the external agent is an attacker whereas the defender declares that there is no attack (Type II error, or missed detection). The goal of the defender is to minimize a weighted sum of the above two types of errors. After suitable normalization, we define the utility of the defender as

$$u_n^D(q, \varphi) = - \left(\sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n))q(\mathbf{x}^n) + \gamma \sum_{\mathbf{x}^n} \varphi(\mathbf{x}^n)p(\mathbf{x}^n) \right), \quad (3.2)$$

where $\gamma > 0$ is a constant that captures the exogenous probability of attack (i.e., θ), as well as the relative weights given to the error terms.

We denote our Bayesian hypothesis testing game with utility functions (3.1) and (3.2) by $\mathcal{G}^B(d, n)$. With a slight abuse of notation, we denote by $u_n^A(\sigma_n^A, \sigma_n^D)$ and $u_n^D(\sigma_n^A, \sigma_n^D)$, the utility of the players under a mixed strategy (σ_n^A, σ_n^D) , where $\sigma_n^A \in M_1(Q)$, and $\sigma_n^D \in M_1(\Phi_n)$.

For our analysis of game $\mathcal{G}^B(d, n)$, we will make use of the following assumptions:

- (A1) Q is a closed subset of $M_1(\mathcal{X})$, and $p \notin Q$.
- (A2) $p(i) > 0$ for all $i \in \mathcal{X}$. Furthermore, for each $q \in Q$, $q(i) > 0$ for all $i \in \mathcal{X}$.
- (A3) c is continuous on Q , and there exists a unique $q^* \in Q$ such that

$$q^* = \arg \min_{q \in Q} c(q).$$

- (A4) The point p is distant from the set Q relative to the point q^* , i.e.,

$$\{\mu \in M_1(\mathcal{X}) : D(\mu||p) \leq D(\mu||q^*)\} \cap Q = \emptyset,$$

where $D(\mu||\nu) = \sum_{i \in \mathcal{X}} \mu(i) \log \frac{\mu(i)}{\nu(i)}$, $\mu, \nu \in M_1(\mathcal{X})$, denotes the relative entropy between the distributions μ and ν .

Note that (A1) and (A2) are very natural. In (A2), if $p(i) = 0$ for some $i \in \mathcal{X}$ and $q(i) > 0$ for some $q \in Q$, then the adversary will never pick q , as the defender can easily detect the presence of the attacker by looking for element i . On the other hand, if $p(i) = 0$ and $q(i) = 0$ for all $q \in Q$, we may consider a new alphabet set without i . In (A3), continuity of the cost function c is natural and we do not assume any extra condition other than the requirement that there is a unique minimizer. Assumption (A4) is used to show certain property of the equilibrium of the defender, which is later used in the study of error exponents associated with classification error. Specifically, Assumption (A4) is used in the proofs of Lemma 4.4, Lemma 4.5 and Theorem 4.1; all other results are valid without this assumption. We will further discuss the role of (A3) and (A4) in Section 4.3 after Theorem 4.1.

3.2 Model discussion

Our setting is that of adversarial hypothesis testing, where the attacker chooses a distribution and points are then generated i.i.d. according to it. This is a reasonable model in applications such as multimedia forensics (where one tries to determine if an attacker has tampered with an image from signals that can be modeled as random variables following an image-dependent distribution) or biometrics (where again one tries to detect from random signals whether the perceived signals do come from the characteristics of a given individual or they come from tampered characteristics)—see more details about these applications in [2, 3, 29]. In such applications, it is reasonable that different ways of tampering have different costs for the attacker and that one can estimate those costs for a given application at least to some extent. Modeling the attacker’s utility via a cost function is classical in other settings, for instance in adversarial classification [12, 28, 6] and experiments with real-world applications where a reasonable cost function can be estimated has been done, for instance, in [6].

Our setting is very similar to that of a composite hypothesis testing framework where nature picks a distribution from a given set and generates independent samples from it. However, in such problems, one does not model a utility function for the nature/statistician and one is often interested in existence and properties of uniformly most powerful test or locally most powerful test (depending on the Bayesian or frequentist approach; see Section II.E in [25]). In contrast, here, we specifically model the utility functions for the agents and investigate the behavior at Nash equilibrium using very different analysis, which is more natural in adversarial settings where two rational agents interact.

Our setting also coincides with the well-studied setting of minimax testing [16] when $c(q) = 0$ for all $q \in Q$ (and hence every q is a minimizer of c). Note, however, that this case is not included in our model due to Assumption (A3)—rather we study the opposite extreme where c has a unique minimizer. Our results are not an easy extension of the classical results because our game is now a nonzero-sum game (whereas the minimax setting corresponds to a zero-sum game). We can therefore not inherit any of the nice properties of zero-sum games; in particular we cannot compute the NE and we instead have to prove properties of the NE (e.g., concentration) without being able to explicitly compute it. In fact, our results too are quite different since we show that the error rate is the same as a simple test where H_1 would contain only q^* , which is different from the classical minimax case.

Finally, in our model we fix the sample size n , i.e., the defender makes a decision only after observing all n samples. We restrict to this simpler setting since it has applications in various domains (see Section 1.1), and understanding the equilibrium of such games leads to interesting and non-trivial results. We leave the study of a sequential model where the defender has the flexibility to choose the number of samples for decision making as future work.

4 Main Results

4.1 Mixed Strategy Nash Equilibrium for $\mathcal{G}^B(d, n)$

We first examine the Nash equilibrium for $\mathcal{G}^B(d, n)$. Note that the strategy sets of both the attacker and the defender are uncountable, hence it is a priori not clear whether our game has a Nash equilibrium.

Towards this, we equip the set Φ_n of all randomized decision rules with the sup-norm metric, i.e.,

$$d_n(\varphi_1, \varphi_2) = \max_{\mathbf{x}^n \in \mathcal{X}^n} |\varphi_1(\mathbf{x}^n) - \varphi_2(\mathbf{x}^n)|,$$

for $\varphi_1, \varphi_2 \in \Phi_n$. It is easy to see that the set Φ_n endowed with the above metric is a compact metric space. We also equip $M_1(\mathcal{X})$ with the usual Euclidean topology on \mathbb{R}^d , and equip Q with the subspace topology. Also, for studying the mixed extension of the game, we equip the spaces $M_1(Q)$ and $M_1(\Phi_n)$ with their corresponding weak topologies. Product spaces are always equipped with the corresponding product topology.

We begin with a simple continuity property of the utility functions.

Lemma 4.1. *Assume (A1)-(A3). Then, the utility functions u_n^A and u_n^D are continuous on $Q \times \Phi_n$.*

We now show the main result of this subsection, namely existence and partial characterization of a NE for our hypothesis testing game.

Proposition 4.1. *Assume (A1)-(A3). Then, there exists a mixed strategy Nash equilibrium for $\mathcal{G}^B(d, n)$. If $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a NE, then so is $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ where $\hat{\varphi}_n$ is the likelihood ratio test given by*

$$\hat{\varphi}_n(\mathbf{x}^n) = \begin{cases} 1, & \text{if } q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) > 0, \\ \varphi_{\hat{\sigma}_n^D}, & \text{if } q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) = 0, \\ 0, & \text{if } q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) < 0, \end{cases} \quad (4.1)$$

where $q_{\hat{\sigma}_n^A}(\mathbf{x}^n) = \int q(\mathbf{x}^n) \hat{\sigma}_n^A(dq)$, and $\varphi_{\hat{\sigma}_n^D} = \int \varphi(\mathbf{x}^n) \hat{\sigma}_n^D(d\varphi)$.

The existence of a NE follows from Glicksberg's fixed point theorem (see e.g., Corollary 2.4 in [26]); for the form of the defender's equilibrium strategy, we have to examine the utility function u_n^D .

Remark 4.1. Note that we have considered randomization over Φ_n to show existence of a NE. Once this is established, we can then show the form of the strategy of the defender $\hat{\varphi}_n$ at equilibrium; the existence of a NE is not clear if we do not consider randomization over Φ_n .

Remark 4.2. Note that the distribution $q_{\hat{\sigma}_n^A}$ on \mathcal{X}^n cannot necessarily be written as an n -fold product distribution of some element from $M_1(\mathcal{X})$. Therefore, the test $\hat{\varphi}_n$ is slightly different from the usual likelihood ratio test that appears in the classical hypothesis testing problem where samples are generated independently.

Remark 4.3. Apart from the conditions of the above proposition, a sufficient condition for existence of pure strategy Nash equilibrium is that the utilities are individually quasiconcave, i.e., $u_n^A(\cdot, \varphi)$ is quasiconcave for all $\varphi \in \Phi_n$, and $u_n^D(q, \cdot)$ is quasiconcave for all $q \in Q$. However, it is easy to check that the Type II error term is not quasiconcave in the attacker's strategy, and hence the utility of the attacker is not quasiconcave. Hence, a pure strategy Nash equilibrium is not guaranteed to exist—see numerical experiments in Section C.1.

Remark 4.4. Proposition 4.1 does not provide any information about the structure of the attacker's strategy at a NE. We believe that obtaining the complete structure of a NE and computing it is a difficult problem in general because the strategy spaces of both players are uncountable (and there is no pure-strategy NE in general), and we cannot use the standard techniques for finite games. However, we emphasize that we are able to obtain error exponents at an equilibrium (see Theorem 4.1) without explicitly computing the structure of a NE. Also, one could study Stackelberg equilibrium for our game $\mathcal{G}^B(d, n)$ to help solve computational issues, although we note that most of the security games literature using Stackelberg games assumes finite action spaces (see, for example, [18]); however we do not address the study of Stackelberg equilibrium in this paper.

4.2 Concentration Properties of Equilibrium

We now study some concentration properties of the mixed strategy Nash equilibrium for the game $\mathcal{G}^B(d, n)$ for large n . The results in this section will be used later to show the exponential convergence of the classification error at equilibrium.

Let e_n denote the classification error, i.e., $e_n(q, \varphi) = -u_n^D(q, \varphi)$, $q \in Q$, $\varphi \in \Phi_n$. We begin with the following lemma, which asserts that the error at equilibrium is small for large enough n .

Lemma 4.2. *Assume (A1)-(A3). Let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be a sequence such that, for each $n \geq 1$, $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a mixed strategy Nash equilibrium for $\mathcal{G}^B(d, n)$. Then, $e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \rightarrow 0$ as $n \rightarrow \infty$.*

The main idea in the proof is to let the defender play a decision rule whose acceptance set is a small neighborhood around the point p , and then bound $e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ using the error of the above strategy.

We now show that the mixed strategy profile of the attacker $\hat{\sigma}_n^A$ converges weakly to the point mass at q^* (denoted by δ_{q^*}) as $n \rightarrow \infty$. This is a consequence of the fact that q^* is the minimizer of c , and hence for large enough n , the attacker does not gain much by deviating from the point q^* .

Lemma 4.3. *Assume (A1)-(A3), and let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be as in Lemma 4.2. Then, $\hat{\sigma}_n^A \rightarrow \delta_{q^*}$ weakly as $n \rightarrow \infty$.*

Note that it is not clear from the above lemma that the equilibrium strategy of the attacker $\hat{\sigma}_n^A$ is supported on a small neighborhood around q^* for large enough n . By playing a strategy q that is far from q^* we could still have $u_n^A(q, \hat{\sigma}_n^D) = u_n^A(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$, since the error term in u_n^A could compensate for the possible loss of utility from the cost term. We now proceed to show that this cannot happen under Assumption (A4). We first argue that the equilibrium error is small even when the attacker deviates from her equilibrium strategy.

Lemma 4.4. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be as in Lemma 4.2. Then,*

$$\sup_{q \in Q} e_n(q, \hat{\sigma}_n^D) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

We are now ready to show the concentration of the attacker's equilibrium:

Lemma 4.5. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be as in Lemma 4.2. Let $(q_n)_{n \geq 1}$ be a sequence such that $q_n \in \text{supp}(\hat{\sigma}_n^A)$ for each $n \geq 1$. Then, $q_n \rightarrow q^*$ as $n \rightarrow \infty$.*

The above concentration phenomenon is a consequence of the uniqueness of q^* and Assumption (A4). The main idea in the proofs of Lemma 4.4 and Lemma 4.5 is to essentially show that, for large enough n , the acceptance region of H_0 under (any) mixed strategy Nash equilibrium does not intersect the set Q . If we do not assume (A4), then the decision region at equilibrium could intersect Q , and we may not have the concentration property in the above lemma (we will still have the convergence property in Lemma 4.3 though, which does not use (A4)).

4.3 Error Exponents

With the results on concentration properties of the equilibrium from the previous section, we are now ready to examine the error exponent associated with the classification error at equilibrium. Let Λ_0 denote the log-moment generating function of the random variable $\frac{q^*(X)}{p(X)}$ under H_0 , i.e., when $X \sim p$: $\Lambda_0(\lambda) = \log \sum_{i \in \mathcal{X}} \exp\left(\lambda \frac{q^*(i)}{p(i)}\right) p(i)$, $\lambda \in \mathbb{R}$. Define its Fenchel-Legendre transform

$$\Lambda_0^*(x) = \sup_{\lambda \in \mathbb{R}} \{\lambda x - \Lambda_0(\lambda)\}, \quad x \in \mathbb{R}.$$

Our main result in the paper (for the Bayesian case) is the following theorem.

Theorem 4.1. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be as in Lemma 4.2. Then,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) = -\Lambda_0^*(0).$$

Our approach to show this result is via obtaining asymptotic lower and upper bounds for the classification error at equilibrium $e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$. Since we do not have much information about the structure of the equilibrium, we first let one of the players deviate from their equilibrium strategy, so that we can estimate the error corresponding to the new pair of strategies, and then use these estimates to compute the error rate at equilibrium. The lower bound easily follows by letting the attacker play the strategy q^* , and using the error exponent in the classical hypothesis testing problem between p and q^* . For the upper bound, we let the defender play a specific deterministic decision rule, and make use of the concentration properties of the equilibrium of the attacker in Section 4.2.

Thus, we see that the error exponent is the same as that for the classical hypothesis testing problem of X_1, \dots, X_n i.i.d. $\sim p$ versus X_1, \dots, X_n i.i.d. $\sim q^*$ (see Corollary 3.4.6 in [10]). That is, for large values of n , the adversarial hypothesis testing game is not much different from the above classical setting (whose parameters are derived from the adversarial setting) in terms of classification error. We emphasize that we have not used any property of the specific structure of the mixed strategy Nash equilibrium in obtaining the error exponent associated with the classification error, and hence Theorem 4.1 is valid for any NE. We believe that obtaining the actual structure of a NE is a difficult problem, as the strategy spaces are infinite, and the utility functions do not possess any monotonicity properties in general. For numerical computation of error exponents in a simple case, see Section 5.

We conclude this section by discussing the role of Assumptions (A3) and (A4). We used (A4) to obtain the concentration of equilibrium in Lemma 4.5. Without this assumption, Theorem 4.1 is not valid; see Section 5 for numerical counter-examples. Also, in our setting, unlike the classical minimax testing, it is not clear whether it is always true that the error goes to 0 as the number of samples becomes large, and whether the attacker should always play a point close to q^* at equilibrium. It could be that playing a point far from q^* is better if she can compensate the loss from c from the error term. In fact, that is what happens when (A4) is not satisfied, since there is partial overlap of the decision region of the defender with the set Q . Regarding (A3), when c has multiple minimizers, our analysis can only tell us that the equilibrium of the attacker is supported around the set of minimizers for large enough n ; to study error exponents in such cases, one has to do a finer analysis of characterizing the attacker’s equilibrium. All in all, using (A3) and (A4) allows us to establish interesting concentration properties of the equilibrium (which is not a priori clear) and error exponents associated with classification error without characterizing a NE, hence we believe that these assumptions serve as a good starting point.

5 Numerical Experiments

In this section, we present two numerical examples in the Bayesian formulation to illustrate the result in Theorem 4.1 and the importance of Assumption (A4). Due to space limitations, additional experiments in the Bayesian formulation are relegated to Appendix C, which illustrate (a) best response strategies of the players, (b) existence of pure strategy Nash equilibrium for large values on n as suggested by Lemma 4.5, and (c) importance of Assumption (A4).¹

We illustrate the result in Theorem 4.1 numerically in the following setting. We fix $\mathcal{X} = \{0, 1\}$ (i.e. $d = 2$) and each probability distribution on \mathcal{X} is represented by the probability that it assigns to the symbol 1, and hence $M_1(\mathcal{X})$ is viewed as the unit interval. We fix $p = 0.5$. For numerical computations, we discretize the set Q into 100 equally spaced points, and we only consider deterministic threshold-based decision rules for the defender. To compute a NE, we solve the linear programs associated with attacker as well as the defender for the zero-sum equivalent game of $\mathcal{G}^B(2, n)$.

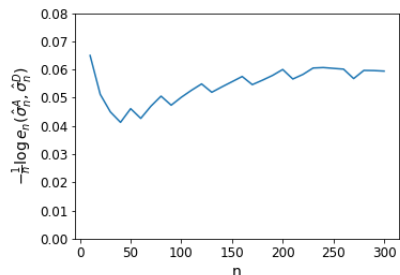
For the function $c(q) = |q - q^*|$ with $q^* = 0.8$, Figure 1(a) shows the error exponent at the NE computed by the above procedure as a function of the number of samples, from $n = 10$ to $n = 300$ in steps of 10. As suggested by Theorem 4.1, we see that the error exponents approach the value $\Lambda_0^*(0) = 0.054$ (the boundary of the decision region is around the point $q = 0.66$, and $D(q||p) \approx D(q||q^*) \approx 0.054$).

We now consider an example which demonstrates that, the result on error exponent in Theorem 4.1 may not be valid if Assumption (A4) is not satisfied. In this experiment, we consider the case where $Q = [0.6, 0.9]$ and $q^* = 0.9$. Note that the present setting does not satisfy Assumption (A4). Figure 1(b) shows the error exponent at the equilibrium as a function of n , from $n = 100$ to $n = 400$ in steps of 100, for the cost function $c(q) = 3|q - q^*|$. From this plot, we see that, the error exponents converge to somewhere around 0.032, whereas $\Lambda_0^*(0) \approx 0.111$.

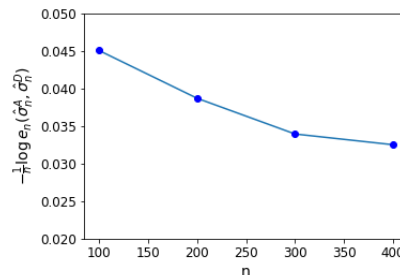
6 Concluding Remarks

In this paper, we studied hypothesis testing games that arise in the context of adversarial classification. We showed that, at equilibrium, the strategy of the classifier is to use a likelihood ratio test. We also

¹Appendix C also contains numerical experiments in the Neyman-Pearson formulation presented in Appendix A. The code used for our simulations is available at https://github.com/sarath1789/ahgtg_neurips2019.



(a) $Q = [0.7, 0.9], c(q) = |q - 0.8|$



(b) $Q = [0.6, 0.9], c(q) = 3|q - 0.9|$

Figure 1: Error exponents as a function of n

examined the exponential rate of decay of classification error at equilibrium and showed that it is same as that of a classical testing problem with parameters derived from the adversarial model.

Throughout the paper, we assumed that the alphabet \mathcal{X} is finite. This is a reasonable assumption in applications that deal with digital signals such as image forensics (an important application for adversarial hypothesis testing); and it is also a good starting point because even in this case, our analysis of the error exponents is nontrivial. Making \mathcal{X} countable/uncountable will make the space $M_1(\mathcal{X})$ infinite dimensional, and the analysis of error exponents will become more difficult (e.g., the continuity of relative entropy is no longer true in this case, which we crucially use in our analysis), but the case of general state space \mathcal{X} is an interesting future direction.

Finding the exact structure of the equilibrium for our hypothesis testing games is a challenging future direction. This will also shed some light on the error exponent analysis for the case when Assumption (A4) is not satisfied. Another interesting future direction is to examine the hypothesis testing game in the sequential detection context where the defender can also decide the number of data samples for classification. In such a setting, an important question is to understand whether the optimal strategy of the classifier is to use a standard sequential probability ratio test.

Acknowledgments

The first author is partially supported by the Cisco-IISc Research Fellowship grant. The work of the second author was supported in part by the French National Research Agency (ANR) through the “Investissements d’avenir” program (ANR-15-IDEX-02) and through grant ANR-16- TERC0012; and by the Alexander von Humboldt Foundation.

References

- [1] Bao, N., Kreidl, P., and Musacchio, J. (2011). Binary hypothesis testing game with training data. In *Proceedings of the 2nd International Conference on Game Theory for Networks (GameNets)*, pages 265–280.
- [2] Barni, M. and Tondi, B. (2013). The source identification game: An information-theoretic perspective. *IEEE Transactions on Information Forensics and Security*, 8(3):450–463.
- [3] Barni, M. and Tondi, B. (2014). Binary hypothesis testing game with training data. *IEEE Transactions on Information Theory*, 60(8):4848–4866.
- [4] Barreno, M., Nelson, B., Joseph, A. D., and Tygar, J. D. (2010). The security of machine learning. *Machine Learning*, 81(2):121–148.
- [5] Brandão, F. G. S. L., Harrowy, A. W., Leez, J. R., and Peres, Y. (2014). Adversarial hypothesis testing and a quantum stein’s lemma for restricted measurements. In *Proceedings of the 5th Innovations in Theoretical Computer Science conference (ITCS)*, pages 183–194.
- [6] Brückner, M., Kanzow, C., and Scheffer, T. (2012). Static prediction games for adversarial learning problems. *The Journal of Machine Learning Research*, 13(1):2617–2654.

- [7] Brückner, M. and Scheffer, T. (2011). Stackelberg games for adversarial prediction problems. In *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 547–555.
- [8] Chen, L. and Leneutre, J. (2009). A game theoretical framework on intrusion detection in heterogeneous networks. *IEEE Transactions on Information Forensics and Security*, 4(2):165–178.
- [9] Dalvi, N., Domingos, P., Mausam, Sanghai, S., and Verma, D. (2004). Adversarial classification. In *Proceedings of the 10th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 99–108.
- [10] Dembo, A. and Zeitouni, O. (2010). *Large Deviations: Techniques and Applications*. Springer-Verlag, Berlin Heidelberg, 2nd edition.
- [11] Dong, J., Roth, A., Schutzman, Z., Waggoner, B., and Wu, Z. S. (2018). Strategic classification from revealed preferences. In *Proceedings of the 2018 ACM Conference on Economics and Computation (EC)*, pages 55–70.
- [12] Dritsoula, L., Loiseau, P., and Musacchio, J. (2017). A game-theoretic analysis of adversarial classification. *IEEE Transactions on Information Forensics and Security*, 12(12):3094–3109.
- [13] Ferguson, T. S. (1967). *Mathematical Statistics: A Decision Theoretic Approach*. Academic Press, New York.
- [14] Hardt, M., N, M., Papadimitriou, C., and Wootters, M. (2016). Strategic classification. In *Proceedings of the 7th Innovations in Theoretical Computer Science conference (ITCS)*, pages 111–122.
- [15] Huang, L., Joseph, A. D., Nelson, B., Rubinstein, B. I. P., and Tygar, J. D. (2011). Adversarial machine learning. In *Proceedings of the 2011 ACM Workshop on Artificial Intelligence and Security (AISec)*, pages 43–58.
- [16] Ingster, Y. I. and Suslina, I. A. (2003). *Nonparametric Goodness-of-Fit Testing Under Gaussian Models*, volume 169 of *Lecture Notes in Statistics*. Springer-Verlag, New York.
- [17] Kantarcioglu, M., Xi, B., and Clifton, C. (2011). Classifier evaluation and attribute selection against 344 active adversaries. *Data Mining and Knowledge Discovery*, 22(1):291–335.
- [18] Korzhyk, D., Yin, D., Kiekintveld, C., Conitzer, V., and Tambe, M. (2011). Stackelberg vs. nash in security games: An extended investigation of interchangeability, equivalence, and uniqueness. *Journal of Artificial Intelligence Research*, 41(2):297–327.
- [19] Li, B. and Vorobeychik, Y. (2015). Scalable optimization of randomized operational decisions in adversarial classification settings. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 599–607.
- [20] Lisý, V., Kessler, R., and Pevý, T. (2014). Randomized operating point selection in adversarial classification. In *Proceedings of the European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML PKDD)*, pages 240–255.
- [21] Lowd, D. and Meeck, C. (2005). Adversarial learning. In *Proceedings of the 11th ACM SIGKDD International Conference on Knowledge Discovery in Data Mining (KDD)*, pages 641–647.
- [22] Lye, K. and Wing, J. M. (2005). Game strategies in network security. *International Journal of Information Security*, 4(1-2):71–86.
- [23] Osborne, M. (2003). *An Introduction to Game Theory*. Oxford University Press, 1st edition.
- [24] Papernot, N., McDaniel, P., Sinha, A., and Wellman, M. (2018). Towards the science of security and privacy in machine learning. In *Proceedings of the 3rd IEEE European Symposium on Security and Privacy*.
- [25] Poor, H. (1994). *An Introduction to Signal Detection and Estimation*. Springer-Verlag, New York, 2nd edition.
- [26] Reny, P. J. (2005). Non-cooperative games: Equilibrium existence. In *The New Palgrave Dictionary of Economics*. Palgrave Macmillan.
- [27] Shiryaev, A. N. (2016). *Probability-1*. Springer-Verlag, New York, 3rd edition.
- [28] Soper, B. and Musacchio, J. (2015). A non-zero-sum, sequential detection game. In *Proceedings of the Allerton Conference on Communication, Control*, pages 361–371.

- [29] Tondi, B., Merhav, N., and Barni, M. (2019). Detection games under fully active adversaries. *Entropy*, 21(1):23.
- [30] Zhou, Y. and Kantarcioglu, M. (2014). Adversarial learning with bayesian hierarchical mixtures of experts. In *Proceedings of the 2014 SIAM International Conference on Data Mining (SDM)*, pages 929–937.
- [31] Zhou, Y., Kantarcioglu, M., Thuraisingham, B., and Xi, B. (2012). Adversarial support vector machine learning. In *Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 1059–1067.

A Hypothesis Testing Game: Neyman-Pearson Formulation

In this section, we study the Neyman-Pearson version of the hypothesis testing problem. The presentation of results in this section is similar to the Bayesian case, and we will also use the same notation as in the Bayesian formulation.

A.1 Problem Formulation

In the Neyman-Pearson point of view, we do not assume any knowledge on the probability that the external agent is an attacker. Fix $\varepsilon > 0$. As before, the strategy set of the attacker is the set Q . For the defender, motivated by the Neyman-Pearson approach for the classical hypothesis testing problem, we define the strategy set as the set of all randomized decision rules on n -length words whose false alarm probability is at most ε , i.e.,

$$\Phi_n = \{\varphi : \mathcal{X}^n \rightarrow [0, 1] : P_n^{FA}(\varphi) \leq \varepsilon\},$$

where $P_n^{FA}(\varphi) = \sum_{\mathbf{x}^n} \varphi(\mathbf{x}^n)p(\mathbf{x}^n)$ denotes the false alarm probability under the decision rule φ .

We now define the utilities. Similar to the Bayesian framework, the utility of the attacker is defined as

$$u_n^A(q, \varphi) = \sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n))q(\mathbf{x}^n) - c(q). \quad (\text{A.1})$$

Since we have already constrained the strategy set of the defender by imposing an upper bound on the Type I error, the utility of the defender is defined as

$$u_n^D(q, \varphi) = - \left(\sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n))q(\mathbf{x}^n) \right), \quad (\text{A.2})$$

which captures the Type II error.

We denote our two-player Neyman-Pearson hypothesis testing game with utility functions (A.1) and (A.2) by $\mathcal{G}^{NP}(\varepsilon, d, n)$. Similar to the Bayesian case, we will use Assumptions (A1)-(A4) throughout the analysis of the Neyman-Pearson case.

A.2 Mixed Strategy Nash Equilibrium for $\mathcal{G}^{NP}(\varepsilon, d, n)$

We now examine the mixed strategy Nash equilibrium for the game $\mathcal{G}^{NP}(\varepsilon, d, n)$. Similar to the Bayesian framework, we endow the set Q with the standard Euclidean topology on \mathbb{R}^d and the set Φ_n with the sup-norm metric. The following lemma asserts compactness of the strategy space of the defender.

Lemma A.1. *The set Φ_n equipped with the metric d_n is a compact metric space.*

We also have joint continuity of the utility functions of both the players, and it can be proved similar to Lemma 4.1.

Lemma A.2. *Assume (A1)-(A3). Then, the utility functions u_n^A and u_n^D are continuous on $Q \times \Phi_n$.*

We are now ready to show the existence of a mixed strategy Nash equilibrium.

Proposition A.1. *Assume (A1)-(A3). Then, there exists a mixed strategy Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, d, n)$. If $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a NE, then $\hat{\sigma}_n^D$ is the point mass at the most powerful ε -level Neyman-Pearson test for X_1, \dots, X_n i.i.d. $\sim p$ versus $(X_1, \dots, X_n) \sim q_{\hat{\sigma}_n^A}$.*

The existence of a mixed strategy equilibrium easily follows from the compactness of strategy spaces and continuity of utility functions. To show the specific form of defender’s equilibrium strategy, we appeal to the Neyman-Pearson lemma.

Let $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ denote a NE given by Proposition A.1.

A.3 Characterization of Equilibrium in the Binary case

Consider the game $\mathcal{G}^{NP}(\varepsilon, d, n)$ in the binary case, i.e., $d = 2$. Here, there are some interesting monotonicity properties of the utility functions that allow us to get a pure strategy Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, 2, n)$, in which the defender plays a threshold-based test, i.e., declares the presence of an adversary whenever the number of 1's in the observation exceeds a threshold:

Lemma A.3. *Assume (A1)-(A3). Then, the defender admits a strictly dominant strategy, and there exists a pure strategy Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, 2, n)$.*

Remark A.1. The monotonicity alluded to above is a consequence of the fact that u_n^D captures just the Type II error. In the Bayesian framework, we do not have this monotonicity in u_n^D , due to the presence of both the Type I and Type II errors in u_n^D , and hence, existence of a pure strategy Nash equilibrium in the binary case cannot be guaranteed in the Bayesian framework.

A.4 Concentration Properties of Equilibrium

In this section, we study some concentration properties of the equilibrium. We have the following two lemmas, which can be proved similar to the corresponding Lemmas for the Bayesian formulation in Section 4.2.

Lemma A.4. *Assume (A1)-(A3). Let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be a sequence such that, for each $n \geq 1$, $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ is a mixed strategy Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, d, n)$. Then, $e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) \rightarrow 0$ as $n \rightarrow \infty$.*

Lemma A.5. *Assume (A1)-(A3), and let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be as in Lemma A.4. Then $\hat{\sigma}_n^A \rightarrow \delta_{q^*}$ weakly as $n \rightarrow \infty$.*

We also have that the error at equilibrium goes to 0 even when the attacker deviates from her equilibrium strategy.

Lemma A.6. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be as in Lemma A.4. Then,*

$$\sup_{q \in Q} e_n(q, \hat{\varphi}_n) \rightarrow 0 \text{ as } n \rightarrow \infty.$$

The main idea in the proof of the above lemma is to show that the acceptance region of H_0 under any equilibrium does not intersect the set Q . With this lemma at hand, we now have the following concentration property for the support of the equilibrium strategy of the attacker, which can be proved similar to Lemma 4.5 in the Bayesian formulation.

Lemma A.7. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be as in Lemma A.4. Let $(q_n)_{n \geq 1}$ be a sequence such that $q_n \in \text{supp}(\hat{\sigma}_n^A)$ for each $n \geq 1$. Then, $q_n \rightarrow q^*$ as $n \rightarrow \infty$.*

Remark A.2. Note that, in one dimension ($d = 1$), the acceptance region of an optimal Neyman-Pearson test for a fixed alternative q will be a “vanishingly small neighborhood of the null distribution p ” and that while it can still intersect Q for finite n , it may not for large-enough n ; so that Lemma A.6 may always hold. However, it is unclear how this might generalize to higher dimension. Intuitively, in higher dimension, the acceptance region may become close to p only from certain directions. We also note that our proof of Lemma A.6 actually uses Assumption (A4) and not a weaker version of it—see the expression of Γ_n in the proof of Lemma A.6. Therefore, we believe that (A4) is needed in higher dimensions even for the Neyman-Pearson case; although it is possible that a weaker assumption will suffice in one dimension.

A.5 Error Exponents

Our main result in the Neyman-Pearson formulation is the following theorem.

Theorem A.1. *Assume (A1)-(A4), and let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be as in Lemma A.4. Then,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) = -D(p||q^*).$$

Again, we note that the error exponent is the same as that of the classical Neyman-Pearson hypothesis testing problem between p and q^* .

B Proofs

B.1 Proof of Lemma 4.1

Since we are on a metric space, it suffices to show sequential continuity. Let $\{(q_k, \varphi_k), k \geq 1\}$ be a sequence such that $(q_k, \varphi_k) \rightarrow (q, \varphi)$ as $k \rightarrow \infty$. First, consider u_n^D . Notice that, for each \mathbf{x}^n , we have $q_k(\mathbf{x}^n) \rightarrow q(\mathbf{x}^n)$, and $\varphi_k(\mathbf{x}^n) \rightarrow \varphi(\mathbf{x}^n)$ as $k \rightarrow \infty$. Therefore, we have that $q_k(\mathbf{x}^n)\varphi_k(\mathbf{x}^n) \rightarrow q(\mathbf{x}^n)\varphi(\mathbf{x}^n)$, which yields that,

$$\lim_{k \rightarrow \infty} \sum_{\mathbf{x}^n} q_k(\mathbf{x}^n)\varphi_k(\mathbf{x}^n) = \sum_{\mathbf{x}^n} q(\mathbf{x}^n)\varphi(\mathbf{x}^n).$$

Similarly, we also have

$$\lim_{k \rightarrow \infty} \sum_{\mathbf{x}^n} p(\mathbf{x}^n) \varphi_k(\mathbf{x}^n) = \sum_{\mathbf{x}^n} p(\mathbf{x}^n) \varphi(\mathbf{x}^n).$$

Therefore, we have that $u_n^D(q_k, \varphi_k) \rightarrow u_n^D(q, \varphi)$ as $k \rightarrow \infty$ which proves continuity of the utility of the defender. Using similar arguments, and by using the continuity of the cost function c on Q , we see that $u_n^A(q_k, \varphi_k) \rightarrow u_n^A(q, \varphi)$ as $k \rightarrow \infty$, which shows the continuity of the utility of the attacker. \square

B.2 Proof of Proposition 4.1

$\mathcal{G}^B(d, n)$ is a two-player game with compact strategy spaces. Also, by Lemma 4.1, the utilities (in pure strategies) of both the attacker and the defender are jointly continuous on $Q \times \Phi_n$. Therefore, an application of the Glicksberg fixed point theorem (see, for example, Corollary 2.4 in [26]) tell us that there exists a mixed strategy Nash equilibrium (NE) for the adversarial hypothesis testing game $\mathcal{G}^B(d, n)$.

We now show the structure of the equilibrium strategy of the defender. Note that, for any $\varphi \in \Phi_n$, $u_n^D(\hat{\sigma}_n^A, \varphi) = -1 + \sum_{\mathbf{x}^n} \varphi(\mathbf{x}^n) (q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n))$, where $q_{\hat{\sigma}_n^A}(\mathbf{x}^n) = \int q(\mathbf{x}^n) \hat{\sigma}_n^A(dq)$. Therefore, using the characterization of a NE (see Proposition 140.1 in [23]), it follows that for any $\varphi \in \text{supp}(\hat{\sigma}_n^D)$, we have

$$\varphi(\mathbf{x}^n) = \begin{cases} 1, & \text{if } q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) > 0, \\ 0, & \text{if } q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) < 0. \end{cases}$$

Now, define $\hat{\varphi}_n$ such that $\hat{\varphi}_n(\mathbf{x}^n) = \int \varphi(\mathbf{x}^n) \hat{\sigma}_n^D(d\varphi)$ whenever \mathbf{x}^n is such that $q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) = 0$, and that satisfies the above condition when \mathbf{x}^n is such that $q_{\hat{\sigma}_n^A}(\mathbf{x}^n) - \gamma p(\mathbf{x}^n) \neq 0$. Consider the strategy profile $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ where the defender plays the pure strategy $\hat{\varphi}_n$. By the choice of $\hat{\varphi}_n$, we see that $u_n^A(q, \hat{\varphi}_n) = u_n^A(q, \hat{\sigma}_n^D)$ for all $q \in Q$, and $u_n^D(\hat{\sigma}_n^A, \varphi) \leq u_n^D(\hat{\sigma}_n^A, \hat{\varphi}_n)$ for any $\varphi \in \Phi_n$. Therefore, using the characterization of a NE, we see that $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ is a NE. This completes the proof of the Proposition. \square

B.3 Proof of Lemma 4.2

By Assumption (A1), there exist a $\delta > 0$ such that $B(p, \delta) \cap Q = \emptyset$, where $B(p, \delta)$ denotes an open ball of radius δ centered at p . Let φ^δ denote the deterministic decision rule whose rejection region is the set $B(p, \delta)^c$, i.e., $\varphi^\delta(\mathbf{x}^n) = 1$ whenever $\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)^c$ and $\varphi^\delta(\mathbf{x}^n) = 0$ otherwise, where $\mathcal{P}_{\mathbf{x}^n} \in M_1(\mathcal{X})$ denotes the type of \mathbf{x}^n , i.e., $\mathcal{P}_{\mathbf{x}^n}(i) = \frac{1}{n} \sum_{k=1}^n 1_{\{x_k=i\}}$. Since $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a Nash equilibrium, and $e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) = -u_n^D(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$, we see that

$$e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \leq e_n(\hat{\sigma}_n^A, \varphi^\delta), \quad (\text{B.1})$$

where $(\hat{\sigma}_n^A, \varphi^\delta)$ denotes the strategy profile where the attacker plays the mixed strategy $\hat{\sigma}_n^A$ and the defender plays the pure strategy φ^δ .

We now proceed to bound the error term $e_n(\hat{\sigma}_n^A, \varphi^\delta)$. We have

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \varphi^\delta) &= \int \left[\sum_{\mathbf{x}^n} (1 - \varphi^\delta(\mathbf{x}^n)) q(\mathbf{x}^n) + \gamma \varphi^\delta(\mathbf{x}^n) p(\mathbf{x}^n) \right] \hat{\sigma}_n^A(dq) \\ &= \int q(\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)) \hat{\sigma}_n^A(dq) + \gamma p(\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)^c). \end{aligned}$$

We bound the first term above using a simple upper bound for the probability of observing a given type under a given distribution (see, for example, Lemma 2.1.9 in [10]). Let \mathcal{P}_n denote the set of all possible types of an n -length word. For any $q \in Q$, we have that

$$\begin{aligned} q(\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)) &\leq \sum_{\nu \in B(p, \delta) \cap \mathcal{P}_n} e^{-nD(\nu||q)} \\ &\leq |B(p, \delta) \cap \mathcal{P}_n| e^{-n \inf_{\nu \in B(p, \delta)} D(\nu||q)} \\ &\leq (n+1)^d e^{-n \inf_{\nu \in B(p, \delta), q \in Q} D(\nu||q)}, \end{aligned}$$

where the last inequality follows since $|\mathcal{P}_n| \leq (n+1)^d$. Therefore,

$$e_n(\hat{\sigma}_n^A, \varphi^\delta) \leq (n+1)^d e^{-n \inf_{\nu \in B(p, \delta), q \in Q} D(\nu||q)} + \gamma p(\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)^c).$$

The first term above goes to 0 as $n \rightarrow \infty$, since $\inf_{\nu \in B(p, \delta), q \in Q} D(\nu||q) > 0$. Also, by the weak law of large numbers, we see that $\mathcal{P}_{\mathbf{x}^n}$ converges to p in probability under the null hypothesis H_0 . Therefore,

$$p(\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)^c) \rightarrow 0.$$

Hence, we conclude that $e_n(\hat{\sigma}_n^A, \varphi^\delta) \rightarrow 0$ as $n \rightarrow \infty$. Combining this with (B.1) completes the proof of the Lemma. \square

B.4 Proof of Lemma 4.3

From Lemma 4.2, we have $e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \rightarrow 0$ as $n \rightarrow \infty$. Since $u_n^A(q^*, \hat{\sigma}_n^D) \geq -c(q^*)$ and since $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a NE for all $n \geq 1$, it follows that

$$\int c(q)\hat{\sigma}_n^A(dq) \rightarrow c(q^*)$$

as $n \rightarrow \infty$. Since $(\hat{\sigma}_n^A)_{n \geq 1}$ is a sequence of probability measures on the compact space Q , by Prohorov's theorem (see Theorem 1, Section 2 in Chapter 3 of [27]), there exists a weakly convergent subsequence (say $(n_k)_{k \geq 1}$). Let μ denote the weak limit of $(\hat{\sigma}_{n_k}^A)_{k \geq 1}$. Then, we have,

$$\begin{aligned} c(q^*) &= \lim_{k \rightarrow \infty} \int c(q)\hat{\sigma}_{n_k}^A(dq) \\ &= \int c(q)\mu(dq), \end{aligned} \tag{B.2}$$

where the last equality follows from weak convergence.

We now claim that $\mu = \delta_{q^*}$. Suppose not. Then, there exists $\varepsilon > 0$ such that $\mu(B(q^*, \varepsilon)^c) > 0$. By Assumption (A3), for the above ε , there exists a $\delta > 0$ such that $c(q) > c(q^*) + \delta$ whenever $q \in B(q^*, \varepsilon)^c$. Therefore,

$$\begin{aligned} \int_Q c(q)\mu(dq) &= \int_{B(q^*, \varepsilon)} c(q)\mu(dq) + \int_{B(q^*, \varepsilon)^c} c(q)\mu(dq) \\ &\geq c(q^*)\mu(B(q^*, \varepsilon)) + (c(q^*) + \delta)\mu(B(q^*, \varepsilon)^c) \\ &= c(q^*) + \delta\mu(B(q^*, \varepsilon)^c) \\ &> c(q^*), \end{aligned}$$

which contradicts (B.2). Therefore, it follows that $\mu(B(q^*, \varepsilon)^c) = 0$ for every $\varepsilon > 0$ and hence $\mu = \delta_{q^*}$. Since μ is independent of the subsequence $(n_k)_{k \geq 1}$, it follows that the whole sequence $(\hat{\sigma}_n^A)_{n \geq 1}$ converges to δ_{q^*} (see Lemma 1, Section 3 in Chapter 3 of [27]). This completes the proof of the lemma. \square

To prove Lemma 4.4, we need the following lemma, which asserts uniform convergence of integrals of the relative entropy functional w.r.t. the equilibrium strategy of the attacker.

Lemma B.1. *Let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)_{n \geq 1}$ be as in Lemma 4.2. Then,*

$$\sup_{\mu \in M_1(\mathcal{X})} \left| \int D(\mu||q)\hat{\sigma}_n^A(dq) - D(\mu||q^*) \right| \rightarrow 0 \text{ as } n \rightarrow \infty.$$

Proof. Fix $\varepsilon > 0$ and $\mu \in M_1(\mathcal{X})$. Then, using the uniform continuity of the relative entropy function on $M_1(\mathcal{X}) \times Q$, there exists $\delta > 0$ such that

$$\sup_{q \in Q} |D(\mu||q) - D(\mu'||q)| < \varepsilon \text{ for all } \mu' \in B(\mu, \delta).$$

Therefore, for all $\mu' \in B(\mu, \delta)$, we have

$$\begin{aligned} \left| \int D(\mu'||q)\hat{\sigma}_n^A(dq) - \int D(\mu||q)\hat{\sigma}_n^A(dq) \right| &\leq \int |D(\mu'||q) - D(\mu||q)|\hat{\sigma}_n^A(dq) \\ &\leq \varepsilon, \text{ for all } n \geq 1. \end{aligned}$$

Also, using weak convergence of $(\hat{\sigma}_n^A)_{n \geq 1}$ to the point mass at q^* , there exists $N_\mu \geq 1$ such that

$$\left| \int D(\mu||q)\hat{\sigma}_n^A(dq) - D(\mu||q^*) \right| \leq \varepsilon \text{ for all } n \geq N_\mu.$$

Note that, the sets $B(\mu, \delta)_{\mu \in Q}$ is an open cover for $M_1(\mathcal{X})$. By compactness of the space $M_1(\mathcal{X})$, extract a finite subcover $B(\mu_i, \delta)$, $1 \leq i \leq k$. Put $N = \max\{N_{\mu_1}, \dots, N_{\mu_k}\}$. Then, for all $n \geq N$, we have

$$\begin{aligned} \left| \int D(\mu||q)\hat{\sigma}_n^A(dq) - D(\mu||q^*) \right| &\leq \left| \int D(\mu||q)\hat{\sigma}_n^A(dq) - \int D(\mu_i||q)\hat{\sigma}_n^A(dq) \right| + \\ &\quad \left| \int D(\mu_i||q)\hat{\sigma}_n^A(dq) - D(\mu_i||q^*) \right| + \\ &\quad |D(\mu_i||q^*) - D(\mu||q^*)| \\ &\leq 3\varepsilon, \end{aligned}$$

where μ_i is such that $\mu \in B(\mu_i, \delta)$. The result now follows since ε and μ are arbitrary. \square

B.5 Proof of Lemma 4.4

Recall the decision rule $\hat{\varphi}_n$ from Proposition 4.1. Note that if H_0 is accepted under $\hat{\varphi}_n$ when the defender observes \mathbf{x}^n , then we have

$$\frac{\int q(\mathbf{x}^n) \hat{\sigma}_n^A(dq)}{p(\mathbf{x}^n)} \leq \gamma$$

(note that there could be randomization when equality holds above). By Proposition 4.1, notice that $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ is a Nash equilibrium, and $e_n(q, \hat{\sigma}_n^D) = e_n(q, \hat{\varphi}_n)$ for all $q \in Q$. Therefore it suffices to show that $\sup_{q \in Q} e_n(q, \hat{\varphi}_n) \rightarrow 0$ as $n \rightarrow \infty$.

Note that, the acceptance region of H_0 under the decision rule $\hat{\varphi}_n$ is type-based, i.e., for every n -length word \mathbf{x}^n , $\hat{\varphi}_n(\mathbf{x}^n)$ depends only on $\mathcal{P}_{\mathbf{x}^n}$. Therefore, if H_0 is accepted when the defender observes \mathbf{x}^n , the type $\mathcal{P}_{\mathbf{x}^n}$ must belong to the following subset of $M_1(\mathcal{X})$:

$$\left\{ \mathcal{P}_{\mathbf{x}^n} : \log \int \frac{q(\mathbf{x}^n)}{p(\mathbf{x}^n)} \hat{\sigma}_n^A(dq) \leq \log \gamma \right\}.$$

Define

$$\Gamma_n = \left\{ \mathcal{P}_{\mathbf{x}^n} : \int \log \left(\frac{q(\mathbf{x}^n)}{p(\mathbf{x}^n)} \right) \hat{\sigma}_n^A(dq) \leq \log \gamma \right\}.$$

Notice that, by Jensen's inequality, the acceptance region of H_0 under the decision rule $\hat{\varphi}_n$ is a subset of the above set Γ_n . Also, it is easy to check that,

$$\Gamma_n = \{ \mu \in M_1(\mathcal{X}) : D(\mu||p) - \int D(\mu||q) \hat{\sigma}_n^A(dq) \leq \frac{\log \gamma}{n} \} \cap \mathcal{P}_n.$$

We now show that the set Γ_n does not intersect the set Q for large enough n . First, notice that, the set $\{ \mu \in M_1(\mathcal{X}) : D(\mu||p) \leq D(\mu||q^*) \}$ is closed in $M_1(\mathcal{X})$. Therefore, by Assumption (A4), there exists $\eta > 0$ such that $Q^\eta \cap \{ \mu \in M_1(\mathcal{X}) : D(\mu||p) \leq D(\mu||q^*) \} = \emptyset$, where $Q^\eta = \{ \mu \in M_1(\mathcal{X}) : \inf_{q \in Q} \|\mu - q\| \leq \eta \}$ is the η -expansion of the set Q .

We show that there exists $N \geq 1$ such that $Q^\eta \cap \Gamma_n = \emptyset$ for all $n \geq N$. Suppose not, then we can find a sequence $(\mu_n)_{n \geq 1}$ such that $\mu_n \in Q^\eta$ and $\mu_n \in \Gamma_n$ for all $n \geq 1$. Since Q^η is compact, we can find a subsequence $(n_k)_{k \geq 1}$ along which μ_n converges, and let $\mu = \lim_{k \rightarrow \infty} \mu_{n_k} \in Q^\eta$. Since $\mu_n \in \Gamma_n$ for all $n \geq 1$, using Lemma B.1, we see that μ satisfies $D(\mu||p) \leq D(\mu||q^*)$. This contradicts the fact that $Q^\eta \cap \Gamma_n = \emptyset$, and hence, there exists $N \geq 1$ such that $Q^\eta \cap \Gamma_n = \emptyset$ for all $n \geq N$.

By the law of large numbers, we have

$$\sup_{q \in Q} q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) \rightarrow 0,$$

and

$$p(\mathcal{P}_{\mathbf{x}^n} \notin \Gamma_n) \rightarrow 0$$

as $n \rightarrow \infty$. But, notice that

$$\begin{aligned} e_n(q, \hat{\varphi}_n) &\leq q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma_n) + \gamma p(\mathcal{P}_{\mathbf{x}^n} \notin \Gamma_n) \\ &\leq q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) + \gamma p(\mathcal{P}_{\mathbf{x}^n} \notin \Gamma_n) \end{aligned}$$

for all $q \in Q$ and $n \geq N$. Therefore,

$$\begin{aligned} \sup_{q \in Q} e_n(q, \hat{\varphi}_n) &\leq \sup_{q \in Q} q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) + \gamma p(\mathcal{P}_{\mathbf{x}^n} \notin \Gamma_n) \\ &\rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. □

B.6 Proof of Lemma 4.5

Fix $\varepsilon > 0$. By Lemma 4.4, there exists N_ε such that

$$e_n(q_n, \hat{\sigma}_n^D) \leq \varepsilon$$

for all $n \geq N_\varepsilon$. Therefore,

$$u_n^A(q_n, \hat{\sigma}_n^D) \leq \varepsilon - c(q_n)$$

for all $n \geq N_\varepsilon$. However, by playing the pure strategy q^* , the attacker utility is

$$u_n^A(q^*, \hat{\sigma}_n^D) \geq -c(q^*)$$

for all $n \geq 1$. Since $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a Nash equilibrium, and since $q_n \in \text{supp}(\hat{\sigma}_n^A)$, we must have $u_n^A(q_n, \hat{\sigma}_n^D) \geq u_n^A(q^*, \hat{\sigma}_n^D)$ for all $n \geq N_\varepsilon$. That is,

$$c(q_n) \leq c(q) + \varepsilon$$

for all $n \geq N_\varepsilon$. Thus, it follows that, $c(q_n) \rightarrow c(q^*)$ as $n \rightarrow \infty$. Using the definition of q^* , we see that $q_n \rightarrow q^*$ as $n \rightarrow \infty$. □

B.7 Proof of Theorem 4.1

First, we obtain the asymptotic lower bound. Towards this, we shall consider an equivalent zero-sum game for $\mathcal{G}^B(d, n)$. For $q \in Q$ and $\varphi \in \Phi_n$, define

$$u_n^{eq}(q, \varphi) = \sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n))q(\mathbf{x}^n) + \gamma \sum_{\mathbf{x}^n} \varphi(\mathbf{x}^n)p(\mathbf{x}^n) - c(q).$$

Observe that, as far as the attacker is concerned, for any $\varphi \in \Phi_n$, maximizing $u_n^A(\cdot, \varphi)$ is the same as maximizing $u_n^{eq}(\cdot, \varphi)$, as the extra term present in u_n^{eq} does not depend on the attacker strategy. Similarly, for any $q \in Q$, maximizing the defender's utility function $u_n^D(q, \cdot)$ is the same as minimizing $u_n^{eq}(q, \cdot)$, as the cost function c does not depend on the defender's strategy. Therefore, we see that $\mathcal{G}^B(d, n)$ is best-response equivalent to a two-player zero sum game (with attacker being first player and defender being second player) with the above utility for the first player. Hence, any equilibrium for the original game is also going to be an equilibrium for the zero-sum equivalent game with utility function u_n^{eq} (see Definition 4 in [12] and the remarks before Theorem 2).

Consider the strategy profile $(q^*, \hat{\sigma}_n^D)$, i.e., the attacker plays the pure strategy q^* and the defender plays the mixed strategy $\hat{\sigma}_n^D$ that comes from the equilibrium for the original game, and the equivalence of $\mathcal{G}^B(d, n)$ with the above zero-sum game, we have

$$u_n^{eq}(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \geq u_n^{eq}(q^*, \hat{\sigma}_n^D). \quad (\text{B.3})$$

$(u_n^{eq}(\hat{\sigma}_n^A, \hat{\sigma}_n^D))$ denotes the utility in mixed extension of the equivalent zero-sum game.

Define the deterministic decision rule φ_n^* by

$$\varphi_n^*(\mathbf{x}^n) = \begin{cases} 1, & \text{if } \frac{q^*(\mathbf{x}^n)}{p(\mathbf{x}^n)} \geq \gamma, \\ 0, & \text{otherwise.} \end{cases}$$

It is easy to see that φ_n^* minimizes $e_n(q^*, \cdot)$. Writing the probabilities $p(\mathbf{x}^n)$ and $q(\mathbf{x}^n)$ in terms of $\mathcal{P}_{\mathbf{x}^n}$, it is easy to check that, the acceptance region of φ_n^* is given by

$$\Gamma_n^* = \{ \nu \in M_1(\mathcal{X}) \cap \mathcal{P}_n : D(\nu || q^*) - D(\nu || p) > \frac{\log \gamma}{n} \},$$

i.e., $\varphi_n^*(\mathbf{x}^n) = 0$ whenever $\mathcal{P}_{\mathbf{x}^n} \in \Gamma_n^*$, and $\varphi_n^*(\mathbf{x}^n) = 1$ otherwise. Noting that $u_n^{eq}(q, \varphi) = e_n(q, \varphi) - c(q)$, (B.3) becomes

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) &\geq \int \sum_{\mathbf{x}^n} ((1 - \varphi(\mathbf{x}^n))q^*(\mathbf{x}^n) + \gamma\varphi(\mathbf{x}^n)p(\mathbf{x}^n)) \hat{\sigma}_n^D(d\varphi) - c(q^*) + \int c(q)\hat{\sigma}_n^A(dq) \\ &\geq \int \sum_{\mathbf{x}^n} ((1 - \varphi(\mathbf{x}^n))q^*(\mathbf{x}^n) + \gamma\varphi(\mathbf{x}^n)p(\mathbf{x}^n)) \hat{\sigma}_n^D(d\varphi) \\ &\geq \sum_{\mathbf{x}^n} ((1 - \varphi_n^*(\mathbf{x}^n))q^*(\mathbf{x}^n) + \gamma\varphi_n^*(\mathbf{x}^n)p(\mathbf{x}^n)), \end{aligned} \quad (\text{B.4})$$

where the second inequality follows from the definition of q^* , and the last inequality follows from the optimality of φ_n^* . The quantity in the RHS of the last inequality is the minimum Bayesian error for the following standard binary hypothesis testing problem: under the null hypothesis, each symbol in \mathbf{x}^n is generated independently from p , and under the alternate hypothesis, each symbol is generated independently from q^* . It is well known that (see, for example, Corollary 3.4.6 in [10]),

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log e_n(q^*, \varphi_n^*) \geq -\Lambda_0^*(0),$$

and hence, from (B.3) and (B.4), it follows that,

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \geq -\Lambda_0^*(0). \quad (\text{B.5})$$

We now proceed to show the upper bound. Define the decision rule φ_n' by

$$\varphi_n'(\mathbf{x}^n) = \begin{cases} 1, & \text{if } \frac{q^*(\mathbf{x}^n)}{p(\mathbf{x}^n)} \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

Similar to the decision rule φ_n^* , the acceptance region of φ_n' can be written as

$$\Gamma' = \{ \nu \in M_1(\mathcal{X}) : D(\nu || q^*) - D(\nu || p) > 0 \},$$

i.e., $\varphi_n'(\mathbf{x}^n) = 0$ if $\mathcal{P}_{\mathbf{x}^n} \in \Gamma'$, and $\varphi_n'(\mathbf{x}^n) = 1$ otherwise. By the definition of a Nash equilibrium, and noting that $u_n^D(\hat{\sigma}_n^A, \hat{\sigma}_n^D) = -e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$, we have

$$e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \leq e_n(\hat{\sigma}_n^A, \varphi_n'), \quad (\text{B.6})$$

where $(\hat{\sigma}_n^A, \varphi'_n)$ denotes the strategy profile where the attacker plays the mixed strategy $\hat{\sigma}_n^A$ that comes from the equilibrium, and the defender plays the pure strategy φ'_n . We have,

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \varphi'_n) &= \int \sum_{\mathbf{x}^n} ((1 - \varphi'_n(\mathbf{x}^n))q(\mathbf{x}^n) + \gamma\varphi'_n(\mathbf{x}^n)p(\mathbf{x}^n)) \sigma_n^A(dq) \\ &= \int q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma') \hat{\sigma}_n^A(dq) + p(\mathcal{P}_{\mathbf{x}^n} \in (\Gamma')^c). \end{aligned}$$

Consider the first term. Using the upper bound on the probability of observing a type under a given distribution (Lemma 2.1.9 in [10]), we have

$$q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma') \leq (n+1)^d e^{-n \inf_{\nu \in \Gamma'} D(\nu||q)}.$$

Fix $\varepsilon > 0$. Since the relative entropy is jointly uniformly continuous on $\Gamma' \times Q$, there exists a $\delta > 0$ such that

$$D(\nu||q) \geq D(\nu||q^*) - \varepsilon$$

for all $\nu \in \Gamma'$ whenever $\|q - q^*\|_2 < \delta$. For the above δ , by Lemma 4.5, there exists N_δ such that $\|q - q^*\|_2 < \delta$ whenever $q \in \text{supp}(\hat{\sigma}_n^A)$ for all $n \geq N_\delta$. Therefore, we see that, for all $n \geq N_\delta$ and $\nu \in \Gamma'$,

$$D(\nu||q) \geq D(\nu||q^*) - \varepsilon \text{ for all } q \in \text{supp}(\hat{\sigma}_n^A).$$

Therefore, for all $n \geq N_\delta$, we have

$$q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma') \leq (n+1)^d e^{-n(\inf_{\nu \in \Gamma'} D(\nu||q^*) - \varepsilon)}$$

for all $q \in \text{supp}(\hat{\sigma}_n^A)$. For the second term, using Lemma 2.1.9 in [10], we have

$$p(\mathcal{P}_{\mathbf{x}^n} \in (\Gamma')^c) \leq e^{-n \inf_{\nu \in (\Gamma')^c} D(\nu||p)}.$$

It can be easily shown that (for example, see Exercise 3.4.14(b) in [10]), $\inf_{\nu \in \Gamma'} D(\nu||q^*) = \inf_{\nu \in (\Gamma')^c} D(\nu||p) = \Lambda_0^*(0)$. Hence, the above implies that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \varphi'_n) \leq -\Lambda_0^*(0) + \varepsilon.$$

Letting $\varepsilon \rightarrow 0$, we get

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \varphi'_n) \leq -\Lambda_0^*(0).$$

Therefore, from (B.6) and the above inequality, we have

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) \leq -\Lambda_0^*(0). \quad (\text{B.7})$$

The theorem now follows from (B.5) and (B.7). \square

B.8 Proof of Lemma A.1

We show sequential compactness of Φ_n . Let $(\varphi_k)_{k \geq 1}$ be a sequence in Φ_n . Let $\mathbf{x}_1^n, \dots, \mathbf{x}_l^n$ denote the elements of \mathcal{X}^n . Since $\varphi_n(\mathbf{x}^n) \in [0, 1]$ for all $\mathbf{x}^n \in \mathcal{X}^n$, there exists a subsequence $(k_l^{(1)})_{l \geq 1}$ along which $\varphi(\mathbf{x}_1^n)$ converges. We can then extract a further subsequence $(k_l^{(2)})_{l \geq 1}$ of $(k_l^{(1)})_{l \geq 1}$ along which $\varphi(\mathbf{x}_2^n)$ converges. Repeating the above procedure 2^n times, we see that, there exists a subsequence $(k_l)_{l \geq 1}$ along which $\varphi(\mathbf{x}^n)$ converges for all $\mathbf{x}^n \in \mathcal{X}^n$. Put

$$\varphi(\mathbf{x}^n) = \lim_{l \rightarrow \infty} \varphi_{k_l}(\mathbf{x}^n), \quad \mathbf{x}^n \in \mathcal{X}^n.$$

It is clear that $d_n(\varphi_{k_l}, \varphi) \rightarrow 0$ as $l \rightarrow \infty$, and we have

$$\begin{aligned} P_n^{FA}(\varphi) &= \sum_{\mathbf{x}^n} \varphi(\mathbf{x}^n) p(\mathbf{x}^n) \\ &= \sum_{\mathbf{x}^n} \lim_{l \rightarrow \infty} (\varphi_{k_l}(\mathbf{x}^n)) p(\mathbf{x}^n) \\ &= \lim_{l \rightarrow \infty} \sum_{\mathbf{x}^n} \varphi_{k_l}(\mathbf{x}^n) p(\mathbf{x}^n) \\ &= \lim_{l \rightarrow \infty} P^{FA}(\varphi_{k_l}) \\ &\leq \varepsilon, \end{aligned}$$

since $P^{FA}(\varphi_{k_l}) \leq \varepsilon$ for all $l \geq 1$. This shows that the space Φ_n equipped with the metric d_n is sequentially compact, and hence compact. \square

B.9 Proof of Proposition A.1

By Lemma A.1, the strategy space of the defender is compact. Also, the strategy space of the attacker is compact under the standard Euclidean topology on \mathbb{R}^d . By Lemma A.2, we see that the utility functions of both players are jointly continuous. Therefore, by the Glicksberg fixed point theorem (see, for example, Corollary 2.4. in [26]), there exists a mixed strategy Nash equilibrium for the game $\mathcal{G}(\varepsilon, d, n)$.

We now show the structure of the equilibrium strategy of the defender. Let $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ denote a mixed strategy Nash equilibrium of $\mathcal{G}(\varepsilon, d, n)$. By the property of Nash equilibrium, we have that $e_n(\hat{\sigma}_n^A, \varphi) = e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ for all $\varphi \in \text{supp}(\hat{\sigma}_n^D)$. We claim that $P^{FA}(\varphi) = \varepsilon$ for all $\varphi \in \text{supp}(\hat{\sigma}_n^D)$. If there exists $\varphi \in \text{supp}(\hat{\sigma}_n^D)$ with $P^{FA}(\varphi) < \varepsilon$, then we can find $\mathbf{x}_0^n \in \mathcal{X}^n$ such that $\varphi(\mathbf{x}_0^n) = 0$ and $\delta > 0$ such that the decision rule defined by $\varphi'(\mathbf{x}^n) = \varphi(\mathbf{x}^n)$ for all $\mathbf{x}^n \neq \mathbf{x}_0^n$, and $\varphi'(\mathbf{x}_0^n) = \delta$ has the property that $P_n^{FA}(\varphi') \leq \varepsilon$ and $e_n(\hat{\sigma}_n^A, \varphi') < e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$. This contradicts the fact that, $(\hat{\sigma}_n^A, \hat{\sigma}_n^D)$ is a Nash equilibrium, which proves our claim.

But, note that

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \hat{\sigma}_n^D) &= \int \sum_{\mathbf{x}^n} (1 - \varphi(\mathbf{x}^n)) q(\mathbf{x}^n) \hat{\sigma}_n^A(dq) \hat{\sigma}_n^D(d\varphi) \\ &= \sum_{\mathbf{x}^n} \left[\int (1 - \varphi(\mathbf{x}^n)) q(\mathbf{x}^n) \hat{\sigma}_n^A(dq) \hat{\sigma}_n^D(d\varphi) \right] \\ &= \sum_{\mathbf{x}^n} \left[\int (1 - \varphi(\mathbf{x}^n)) q_{\hat{\sigma}_n^A}(\mathbf{x}^n) \hat{\sigma}_n^D(d\varphi) \right] \end{aligned}$$

where $q_{\hat{\sigma}_n^A} \in M_1(\mathcal{X}^n)$ is given by

$$q_{\hat{\sigma}_n^A}(\mathbf{x}^n) = \int q(\mathbf{x}^n) \hat{\sigma}_n^A(dq).$$

That is, when the attacker plays the Nash equilibrium $\hat{\sigma}_n^A$, the defender faces the problem of distinguishing between the two alternatives: (i) (X_1, \dots, X_n) is generated by i.i.d. p , versus (ii) (X_1, \dots, X_n) is generated by $q_{\hat{\sigma}_n^A}$. By the Neyman-Pearson lemma, we know that there exists a Neyman-Pearson decision rule $\hat{\varphi}_n \in \Phi_n$ with the property that $P^{FA}(\hat{\varphi}_n) = \varepsilon$ and $e_n(\hat{\sigma}_n^A, \cdot)$ is minimized by $\hat{\varphi}_n$ on Φ_n . Since every $\varphi \in \text{supp}(\hat{\sigma}_n^D)$ minimizes $e_n(\hat{\sigma}_n^A, \cdot)$, and $P^{FA}(\varphi) = \varepsilon$, and since each $\mathbf{x}^n \in \mathcal{X}^n$ has positive probability of observing under both H_0 and H_1 , an application of the uniqueness part in Neyman-Pearson lemma (see, for example, Section 5.1 in [13]) yields that that $\hat{\sigma}_n^D = \delta_{\hat{\varphi}_n}$. This completes the proof. \square

B.10 Proof of Lemma A.3

Recall the definition of a Neyman-Pearson decision rule. In the binary case, since the comparison of the ratio $\frac{q(\mathbf{x}^n)}{p(\mathbf{x}^n)}$ to a threshold is the same as comparison of the number of 1's in the n -length word \mathbf{x}^n to some other threshold, we see that any Neyman-Pearson decision rule φ must necessarily be of the following form:

$$\varphi(\mathbf{x}^n) = \begin{cases} 0, & \text{if } \mathcal{P}_{\mathbf{x}^n}(1) \in \{0, \frac{1}{n}, \dots, \frac{k}{n}\}, \\ \pi, & \text{if } \mathcal{P}_{\mathbf{x}^n}(1) = \frac{k+1}{n}, \\ 1, & \text{if } \mathcal{P}_{\mathbf{x}^n}(1) \in \{\frac{k+2}{n}, \dots, 1\}, \end{cases} \quad (\text{B.8})$$

for some $\pi \in [0, 1]$ and $0 \leq k \leq n$. Here, $\mathcal{P}_{\mathbf{x}^n}(1)$ denotes the fraction of 1's in \mathbf{x}^n . The false alarm probability of the above decision rule is

$$P_n^{FA}(\varphi) = p \left(\mathcal{P}_{\mathbf{x}^n}(1) \in \{0, \frac{1}{n}, \dots, \frac{k}{n}\} \right) + \pi p \left(\mathcal{P}_{\mathbf{x}^n}(1) = \frac{k+1}{n} \right).$$

Since every n -length word \mathbf{x}^n has positive probability under the distribution p , we see that, there exists a unique k and π such that $P_n^{FA}(\varphi) = \varepsilon$. Let $\hat{\varphi}_n$ denote the above Neyman-Pearson decision rule. Then, by the Neyman-Pearson lemma (see, for example, Proposition II.D.1 in [25]), we see that,

$$\hat{\varphi}_n = \arg \max_{\varphi \in \Phi_n} u_n^D(q, \varphi) \text{ for all } q \in Q.$$

Thus, the defender has a unique strictly dominant strategy. Using the continuity of c , and the continuity of the Type II error term in the attacker's strategy, we see that $u_n^A(\cdot, \hat{\varphi}_n)$ is continuous on Q , and hence there exist a maximum. Therefore, letting the attacker play a pure strategy \hat{q}_n that maximizes $u_n^A(\cdot, \hat{\varphi}_n)$ yields a pure strategy Nash equilibrium $(\hat{q}_n, \hat{\varphi}_n)$. \square

To prove Lemma A.6, we need the following lemma, which can be proved similar to Lemma B.1.

Lemma B.2. Let $(\hat{\sigma}_n^A, \hat{\varphi}_n)_{n \geq 1}$ be as in Lemma A.4. Then,

$$\sup_{\mu \in M_1(\mathcal{X})} \left| \int D(\mu \| q) \hat{\sigma}_n^A(dq) - D(\mu \| q^*) \right| \rightarrow 0 \text{ as } n \rightarrow \infty.$$

B.11 Proof of Lemma A.6

Let γ_n denote the threshold and π denote the randomization used in the decision rule $\hat{\varphi}_n$, i.e., $\hat{\varphi}_n$ is of the form

$$\hat{\varphi}_n(\mathbf{x}^n) = \begin{cases} 1, & \text{if } \frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} > \gamma_n, \\ \pi, & \text{if } \frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} = \gamma_n, \\ 0, & \text{if } \frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} < \gamma_n. \end{cases}$$

We first claim that $\limsup_{n \rightarrow \infty} \gamma_n \leq 1$. Since $P^{FA}(\hat{\varphi}_n) = \varepsilon$, we have that

$$p \left(\frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} \geq \gamma_n \right) \geq \varepsilon \quad (\text{B.9})$$

But, using the probability of observing an n -length word under a distribution (see, for example, Lemma 2.1.6 in [10]), we have

$$\begin{aligned} q_{\hat{\sigma}_n^A}(\mathbf{x}^n) &= \int q(\mathbf{x}^n) \hat{\sigma}_n^A(dq) \\ &= \int e^{-n(H(\mathcal{P}_{\mathbf{x}^n}) + D(\mathcal{P}_{\mathbf{x}^n} \| q))} \hat{\sigma}_n^A(dq), \end{aligned}$$

and

$$p(\mathbf{x}^n) = e^{-n(H(\mathcal{P}_{\mathbf{x}^n}) + D(\mathcal{P}_{\mathbf{x}^n} \| p))}.$$

Therefore,

$$\begin{aligned} p \left(\frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} \geq \gamma_n \right) &= p \left(\int e^{n(D(\mathcal{P}_{\mathbf{x}^n} \| p) - D(\mathcal{P}_{\mathbf{x}^n} \| q))} \geq \gamma_n \right) \\ &\leq p \left(e^{n(D(\mathcal{P}_{\mathbf{x}^n} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^n} \| q))} \geq \gamma_n \right) \end{aligned}$$

By Assumption (A1), we can choose $\delta > 0$ such that $D(\mu \| p) < \inf_{q \in Q} D(\mu \| q)$ for all $\mu \in B(p, \delta)$. Thus,

$$\begin{aligned} &p \left(e^{n(D(\mathcal{P}_{\mathbf{x}^n} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^n} \| q))} \geq \gamma_n \right) \\ &= p \left(e^{n(D(\mathcal{P}_{\mathbf{x}^n} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^n} \| q))} \geq \gamma_n, \mathcal{P}_{\mathbf{x}^n} \in B(p, \delta) \right) \\ &+ p \left(e^{n(D(\mathcal{P}_{\mathbf{x}^n} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^n} \| q))} \geq \gamma_n, \mathcal{P}_{\mathbf{x}^n} \notin B(p, \delta) \right) \end{aligned}$$

By law of large numbers, $p(\mathcal{P}_{\mathbf{x}^n} \notin B(p, \delta)) \rightarrow 0$, and hence the second term above goes to 0 as $n \rightarrow \infty$. Suppose that $\limsup_{n \rightarrow \infty} \gamma_n > 1$, then there exists a subsequence $(n_k)_{k \geq 1}$ such that $\gamma_{n_k} > 1$ for all $k \geq 1$. Therefore, along this subsequence, the first term above becomes

$$\begin{aligned} &p(e^{n_k(D(\mathcal{P}_{\mathbf{x}^{n_k}} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^{n_k}} \| q))} \geq \gamma_{n_k}, \mathcal{P}_{\mathbf{x}^{n_k}} \in B(p, \delta)) \\ &\leq p(e^{n(D(\mathcal{P}_{\mathbf{x}^{n_k}} \| p) - \inf_{q \in Q} D(\mathcal{P}_{\mathbf{x}^{n_k}} \| q))} > 1, \mathcal{P}_{\mathbf{x}^{n_k}} \in B(p, \delta)) \end{aligned}$$

which goes to 0 as $n \rightarrow \infty$ by the choice of δ . This implies that, $p \left(\frac{q_{\hat{\sigma}_n^A}(\mathbf{x}^n)}{p(\mathbf{x}^n)} \geq \gamma_n \right) \rightarrow 0$ as $n \rightarrow \infty$, which contradicts (B.9). Therefore, we must have $\limsup_{n \rightarrow \infty} \gamma_n \leq 1$.

We now argue that, for some $\eta > 0$, the acceptance set of H_0 under $\hat{\varphi}_n$ does not intersect the set Q^η . Towards this, consider the set

$$\Gamma_n = \left\{ \mathcal{P}_{\mathbf{x}^n} : \int \log \left(\frac{q(\mathbf{x}^n)}{p(\mathbf{x}^n)} \right) \hat{\sigma}_n^A(dq) \leq \log \gamma_n \right\}.$$

Notice that, by Jensen's inequality, the acceptance region of H_0 under the decision rule $\hat{\varphi}_n$ is a subset of the above set Γ_n . Also, it is easy to check that,

$$\Gamma_n = \left\{ \mu \in M_1(\mathcal{X}) : D(\mu \| p) - \int D(\mu \| q) \hat{\sigma}_n^A(dq) \leq \frac{\log \gamma_n}{n} \right\} \cap \mathcal{P}_n.$$

We now show that the set Γ_n does not intersect the set Q for large enough n . First, notice that, the set $\{\mu \in M_1(\mathcal{X}) : D(\mu||p) \leq D(\mu||q^*)\}$ is closed in $M_1(\mathcal{X})$. Therefore, by Assumption (A4) there exists $\eta > 0$ such that $\{\mu \in M_1(\mathcal{X}) : D(\mu||p) \leq D(\mu||q^*)\} \cap Q^\eta = \emptyset$.

We show that there exists $N \geq 1$ such that $Q^\eta \cap \Gamma_n = \emptyset$ for all $n \geq N$. Suppose not, then we can find a sequence $(\mu_n)_{n \geq 1}$ such that $\mu_n \in Q^\eta$ and $\mu_n \in \Gamma_n$ for all $n \geq 1$. Since Q^η is compact, we can find a subsequence $(n_k)_{k \geq 1}$ along which μ_n converges, and let $\mu = \lim_{k \rightarrow \infty} \mu_{n_k} \in Q^\eta$. Since $\mu_n \in \Gamma_n$ for all $n \geq 1$, using Lemma B.2 and the fact that $\limsup_{n \rightarrow \infty} \gamma_n = 0$, we see that μ satisfies $D(\mu||p) \leq D(\mu||q^*)$. This contradicts the fact that $Q^\eta \cap \Gamma_n = \emptyset$, and hence, there exists $N \geq 1$ such that $Q^\eta \cap \Gamma_n = \emptyset$ for all $n \geq N$.

By the law of large numbers, we have

$$\sup_{q \in Q} q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) \rightarrow 0$$

as $n \rightarrow \infty$. But, notice that

$$\begin{aligned} e_n(q, \hat{\varphi}_n) &\leq q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma_n) \\ &\leq q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) \end{aligned}$$

for all $q \in Q$ and $n \geq N$. Therefore,

$$\sup_{q \in Q} e_n(q, \hat{\varphi}_n) \leq \sup_{q \in Q} q(\mathcal{P}_{\mathbf{x}^n} \notin B(q, \eta)) \rightarrow 0$$

as $n \rightarrow \infty$. □

B.12 Proof of Theorem A.1

We proceed through similar steps as in the proof of Theorem 4.1. To show the lower bound, we let the attacker play the pure strategy q^* instead of the her equilibrium strategy $\hat{\sigma}_n^A$ for all $n \geq 1$. Since $u_n^A(q, \varphi) = e_n(q, \varphi) - c(q)$, and since $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ is a Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, d, n)$, we see that

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) &\geq e_n(q^*, \hat{\varphi}_n) - c(q^*) + \int c(q) \hat{\sigma}_n^A(dq) \\ &\geq e_n(q^*, \hat{\varphi}_n) \\ &\geq e_n(q^*, \varphi_n^*) \end{aligned}$$

where φ_n^* denotes the best ε -level Neyman-Pearson test for distinguishing p from q^* from n independent samples. Here, the second inequality follows from the definition of q^* , and the last inequality follows from the optimality of Neyman-Pearson test φ_n^* . Hence, using Stein's lemma (see, for example, Lemma 3.4.7 in [10]), we see that

$$\liminf_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) \geq -D(p||q^*). \quad (\text{B.10})$$

We now show the upper bound. Fix $0 < \delta < 1$ such that $B(p, \delta) \cap Q = \emptyset$, and consider the deterministic decision rule φ^δ with acceptance region $B(p, \delta)$, i.e., $\varphi^\delta(\mathbf{x}^n) = 0$ whenever $\mathcal{P}_{\mathbf{x}^n} \in B(p, \delta)$ and $\varphi^\delta(\mathbf{x}^n) = 1$ otherwise. To obtain the upper bound, we let the defender play the strategy φ_n^δ for all $n \geq 1$. Since $(\hat{\sigma}_n^A, \hat{\varphi}_n)$ is a Nash equilibrium, and $u_n^D(q, \varphi) = -e_n(q, \varphi)$, we have

$$\begin{aligned} e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) &\leq e_n(\hat{\sigma}_n^A, \varphi^\delta) \\ &= \int q(\mathcal{P}_{\mathbf{x}^n} \in \Gamma^\delta) \hat{\sigma}_n^A(dq) \\ &\leq \int (n+1)^d e^{-n \inf_{\nu \in \Gamma^\delta} D(\nu||q)} \hat{\sigma}_n^A(dq), \end{aligned} \quad (\text{B.11})$$

where the last inequality follows from the upper bound in Lemma 2.1.9 in [10]. By Lemma A.7 and by the uniform continuity of $D(\cdot||\cdot)$ on $\Gamma^\delta \times Q$, there exists $N_\delta \geq 1$ such that

$$D(\nu||q) \geq D(\nu||q^*) - \delta \text{ for all } \nu \in \Gamma^\delta, q \in \text{supp}(\hat{\sigma}_n^A) \text{ and } n \geq N_\delta.$$

Therefore, (B.11) implies that

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) \leq - \inf_{\nu \in \Gamma^\delta} D(\nu||q^*) + \delta.$$

Letting $\delta \rightarrow 0$ and using the continuity of $D(\cdot||q^*)$ on $M_1(\mathcal{X})$, we get

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \log e_n(\hat{\sigma}_n^A, \hat{\varphi}_n) \leq -D(p||q^*). \quad (\text{B.12})$$

The result now follows from (B.10) and (B.12). □

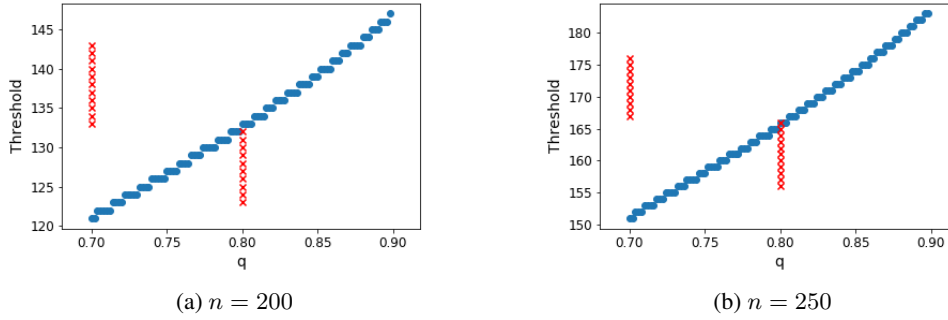
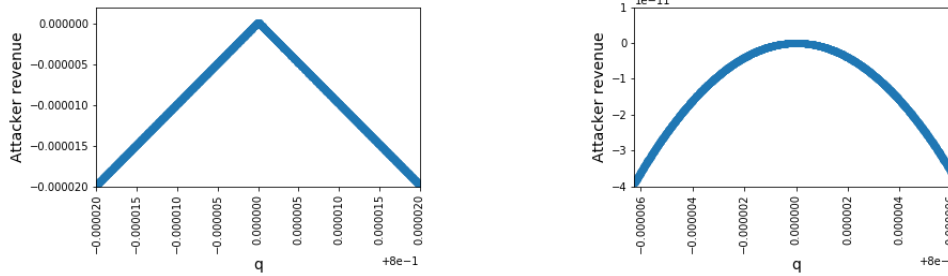


Figure 2: Best response plots for $c(q) = |q - 0.8|$



(a) $c(q) = |q - 0.8|$, $n = 250$, defender plays threshold 166 (b) $c(q) = (q - 0.8)^2$, $n = 800$, defender plays threshold 529

Figure 3: Finer plots of attacker revenue around q^* for specific defender thresholds

C Additional Numerical Experiments

C.1 Bayesian Formulation

As explained in Section 5, we fix $\mathcal{X} = \{0, 1\}$ and $p = 0.5$. For numerical computations, we discretize the set Q into 100 equally spaced points, and we only consider deterministic threshold-based decision rules for the defender.

We first examine the best response of the players. We fix $Q = [0.7, 0.9]$ and the cost function to be $c(q) = |q - q^*|$ where $q^* = 0.8$. Figure 2(a) shows the best response of the players for $n = 200$. The x -axis shows the strategy space of the attacker and y -axis shows the defender's threshold. The blue curve plots the best response of the defender, and the red curve plots the best response of the attacker for 20 thresholds around the best response threshold corresponding to q^* . As we see from the figure, the two curves do not intersect (the best threshold for 0.8 is 133, whereas the best value of q for threshold 133 is 0.7) and hence this suggests that there is no pure strategy equilibrium in this case. Figure 2(b) plots the best response curves for $n = 250$. We see that the two curves intersect (the point of intersection is when the attacker plays 0.8 and the defender plays the threshold 166). However, this does not mean that there is a pure equilibrium, as our discretization may not capture the exact value of the attacker strategy. To see whether this is the case, we plot the attacker revenue when the defender plays the threshold 166 over a finer grid around q^* (1000 equally sized points on the interval of length $1/(100 * n)$ around q^*), which is shown in Figure 3(a). From this, we observe that the maximum of the attacker utility is indeed attained at the point $q = 0.8$. This suggests that there is a pure strategy Nash equilibrium when the attacker plays 0.8 and defender plays the threshold 166, though we could not prove this analytically.

Similar to the best response plots in Figure 2, we plot the best response plots for the quadratic cost function $c(q) = (q - q^*)^2$ where $q^* = 0.8$. Figure 4(a) shows the best response plots for $n = 700$. Here, they don't intersect (the best threshold for 0.8 is 463 whereas the best value of q for threshold 463 is 0.7), which shows that there is no pure equilibrium for the game. Figure 4(b) shows the best response plots for $n = 800$. Here, we see that the curves intersect (when the attacker plays 0.8 and defender plays the threshold 529). As before, Figure 3(b) shows a finer plot of the utility of the attacker around the point q^* . We see that the utility of the attacker is indeed maximized at 0.8, which suggests that there is a pure strategy equilibrium when the attacker plays the strategy 0.8 and defender plays the threshold 529.

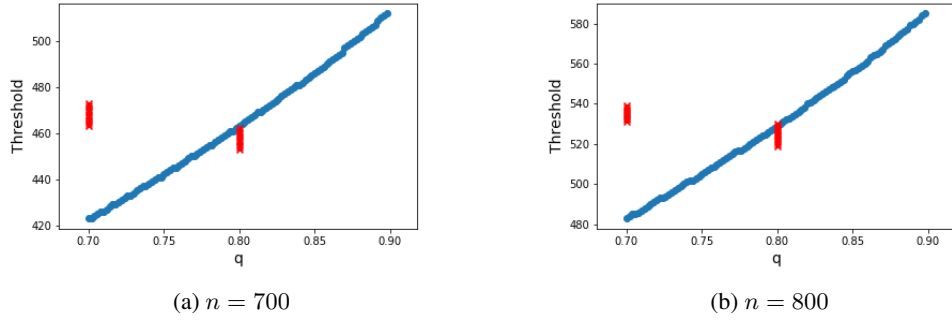


Figure 4: Best response plots for the cost function $c(q) = (q - 0.8)^2$

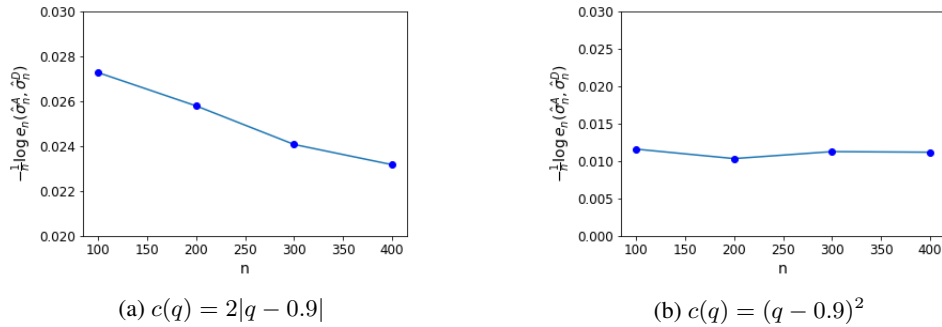


Figure 5: Error exponents as a function of n

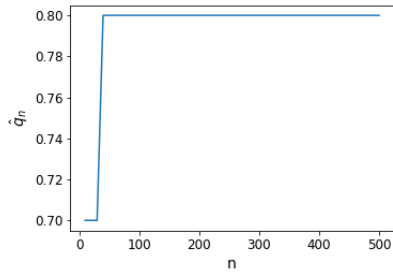
From these experiments for the linear as well as quadratic cost functions, as we expect, there is no incentive for the attacker to deviate much from the point q^* , since for large values of n , the error term in the utility of the attacker does not contribute to the overall revenue. However, in the second case, since the cost function has zero derivative at q^* , it is not clear whether a slight deviation from the point q^* can increase the error term compared to the decrease in the cost function, so that the overall utility of the attacker increases. Therefore, the existence of a pure strategy Nash equilibrium with the attacker strategy being equal to q^* is surprising in this case. However, in the first case, since the left and right derivatives of the cost function at q^* are non-zero, the decrease in the cost function is much larger compared to the possible increase in the error term as we slightly deviate from q^* , and hence it is reasonable to expect the existence of a pure strategy equilibrium at q^* for large n .

Comparing the two cost functions, a much higher value of n is needed in the second case for us to have a pure equilibrium at q^* , since the increase in the cost function is much slower in the second case as we move away from the point q^* .


We now give two more examples that does not satisfy Assumption (A4) whose error exponents are different from Theorem 4.1. As before, $Q = [0.6, 0.9]$ and $q^* = 0.9$, and recall that Assumption (A4) is not satisfied in this case. We consider the linear cost function $c(q) = 2|q - q^*|$ and the quadratic cost function $c(q) = (q - q^*)^2$. Figures 5(a) and 5(b) show the error exponent at the equilibrium as a function of n , from $n = 100$ to $n = 400$ in steps of 100. From these plots, we see that, the error exponents converge to somewhere around 0.023 and 0.011 respectively for the above two cost functions, whereas the value of $\Lambda_0^*(0)$ is around 0.111.

C.2 Neyman-Pearson Formulation

We fix $\varepsilon = 0.1$ and consider the piecewise linear cost function $c(q) = |q - q^*|$ on $Q = [0.7, 0.8]$ where $q^* = 0.8$. As suggested by Lemma A.3, there exists a pure strategy Nash equilibrium for $\mathcal{G}^{NP}(\varepsilon, 2, n)$ for each $n \geq 1$. We first compute the dominant decision rule of the defender by finding the appropriate value of threshold and randomization. Once this is done, we compute the equilibrium by finding the best response of the attacker corresponding to this dominant strategy of the defender (as before, we discretize the set Q into 100 equally-spaced points). We repeat the experiment for different values of n , and for the quadratic cost function $c(q) = 0.001 * (q - 0.8)^2$. Figure 6 and Figure 7 shows the results for the above two cost functions.

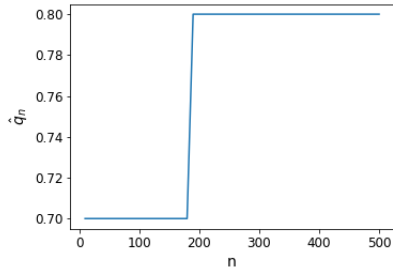


(a) Attacker's strategy as a function of n




(b) Error exponent as a function of n

Figure 6: $c(q) = |q - 0.8|$



(a) Attacker's strategy as a function of n



(b) Error exponent as a function of n

Figure 7: $c(q) = 0.001 * (q - 0.8)^2$

Since the former cost function increases much faster than the latter as we move away from the point q^* , we see that the attacker has much more incentive to play a strategy that is away from q^* in the second case compared to the first. This is reflected in the equilibrium strategy of the attacker; from Figures 6(a) and 7(a), we see that it takes much larger values of n for the equilibrium strategy of the attacker to become equal to q^* in the second case compared to the first. From Figures 6(b) and 7(b), we see that the error exponents approach the limiting value $D(p||q^*) = 0.223$.