Controlled Sensing with Corrupted Commands

Meng-Che Chang, Shi-Yuan Wang, Matthieu R. Bloch
School of Electrical and Computer Engineering, Georgia Institute of Technology, Atlanta, GA 30332
Email: {mchang301,shi-yuan.wang,matthieu}@gatech.edu

***Abstract*—We consider a non-adaptive controlled sensing scenario in which the actions of the decision maker are corrupted by an adversary. The objective of the decision maker is to either detect the presence of the corruption or make a correct decision. Accordingly, the performance of a controlled sensing strategy is measured in terms of the error probability when there is no adversary, denoted $P_{E,0}$, and the error probability when an adversary is present, denoted $P_{E,1}$. Our main result is Stein-lemma like characterization of the optimal achievable error exponent of $P_{E,0}$ subject to a constraint on $P_{E,1}$. We also illustrate the result with numerical examples.**

I. Introduction

Given a set of parameters Θ and a set of distributions $\{p_\theta\}_{\theta\in\Theta}$, the goal of hypothesis testing is to design a strategy allowing a decision maker to figure out the true parameter θ by observing a fixed number of samples generated from p_θ. Classical results for this setting including the Chernoff-Stein's lemma for $|\Theta| = 2$, which characterizes the optimal type-I error exponent when the type-II error probability is upper bounded by a non-trivial small value, and the Chernoff information, which characterizes the optimal Bayesian error exponent [1]. The hypothesis testing problem is called *sequential* when the decision maker can observe samples until a stopping time criterion is met, and *active* when the decision maker can take an action a from a set $\mathcal{A}$ to determine the distribution p_θ^a from which the observation is generated [2]–[4]. Active hypothesis testing is also called *controlled sensing*, which is the terminology adopted here. Recent progress in the area include (sometimes partial) characterizations of the error exponents and expected stopping times of controlled sensing in all combinations of adaptive/non-adaptive sequential/non-sequential settings [3], [4], thereby shedding light on the benefits of adaptivity and sequentiality. Other works have also explored variations of the problem, including the characterization of mis-classification probabilities [5]. Unlike these works, our objective is to analyze the performance of controlled sensing in the presence of an adversary.

Depending on the attacker model, different security performance metrics may be defined. In the case of passive eavesdropping attacks, the evasive hypothesis testing problem in [6] attempts to minimize the error exponent of the adversary while ensuring an acceptable performance for the decision maker. The covert hypothesis testing problem in [7], [8] studies how to perform hypothesis testing problem with low probability of detection, by hiding the existence of the test from an eavesdropping adversary. In the case of active attacks, the adversary may actively degrade the performance of the hypothesis testing by varying the observed samples or the controlling action. When the samples are corrupted, [9] investigates hypothesis testing rules that are robust to adversarial attacks. We investigate here the problem of controlled sensing with corruptions on the actions instead of samples, which to the best of our knowledge, has not been studied yet.

When the controlling action is corrupted, the decision maker might become utterly unable to detect the true hypothesis, so that a proper problem formulation is needed. A similar challenge can be found in the problem of authentication [10]–[13]. Therein, a codeword is transmitted through a state-dependent channel while an adversary attempts to degrade the communication quality by varying the state sequence [13]. The objective for the receiver is then to either decode the message correctly or identify the existence of an adversarial attack. Inspired by these works, we study the performance of controlled sensing with corrupted commands by defining two kinds of error probabilities: the error probability when the adversary is not present and the error probability when the adversary is present. As in the Chernoff-Stein lemma, our work characterizes the optimal exponent of the first error probability while ensuring a non-trivial upper bound on the second error probability.

The rest of the paper is organized as follows. In Section II, we review the notation used throughout the paper. In Section III, we introduce the exact model of interest and present our main results, which characterizes an exact error probability exponent. In Section IV, we develop our achievability proof and in Section V we derive a matching converse.

II. Notation

For any discrete set $\mathcal{X}$, $\mathcal{P}_\mathcal{X}$ is the set of all probability distributions on $\mathcal{X}$. For $n \in \mathbb{N}^*$, a sequence of length n is denoted as $x^n \triangleq (x_1, \cdots, x_n) \in \mathcal{X}^n$, where x_i is the ith element of x^n for any $i \in [1;n]$. For any $x^n \in \mathcal{X}^n$, $\hat{p}_{x^n}$ denotes the type of x^n, i.e., $\hat{p}_{x^n}(a) = \frac{1}{n}\sum_{i=1}^n \mathbb{1}(x_i = a)$ for all $a \in \mathcal{X}$. For any type P, $\mathcal{T}^n(P)$ is the corresponding type class with sequence length n, i.e., the set of all sequence $x^n \in \mathcal{X}^n$ such $\hat{p}_{x^n} = P$. We drop the superscript n when there is no ambiguity on the sequence length. Let $\mathcal{Y}$ be another discrete set. We denote $\mathcal{P}_{\mathcal{Y}|\mathcal{X}}$ the set of all conditional distributions. Given a sequence $x^n \in \mathcal{X}^n$ and $y^n \in \mathcal{Y}^n$, we define $\hat{p}_{y^n|a^n}$ as the empirical conditional type, i.e., $\hat{p}_{y^n|a^n}(b|a) = \sum_{i=1}^n \mathbb{1}(x_i = a, y_i = b)/\sum_{i=1}^n \mathbb{1}(x_i = a)$ for all $a \in \mathcal{X}$ and $b \in \mathcal{Y}$. $N_{y^n,x^n}(b,a) \triangleq \sum_{i=1}^n \mathbb{1}(y_i = b, x_i = a)$ for all $a \in \mathcal{X}$ and $b \in \mathcal{Y}$. For any conditional type $P_{Y|X} \in \mathcal{P}_{\mathcal{Y}|\mathcal{X}}$, we also

This work was supported in part by NSF grants 1955401 and NSF 1910859.

define $\mathcal{T}_{x^n}(P_{Y|X})$ as the conditional type class of $P_{Y|X}$, i.e., the set of sequence $y^n \in \mathcal{Y}^n$ such that $\hat{p}_{y^n|x^n} = P_{Y|X}$. Given two conditional distributions $W_{Y|X}$ and $P_{Y|X}$, we set

$$\left|W_{Y|X} - P_{Y|X}\right|_\infty \triangleq \max_{a\in\mathcal{X}, b\in\mathcal{Y}} |W_{Y|X}(b|a) - P_{Y|X}(b|a)|.$$

We also let

$$C\left(P_{Y|X}\middle\|W_{Y|X}\middle|Q\right) \triangleq \max_{\ell\in[0,1]} -\sum_x Q(x) \times \log\left(\sum_y P_{Y|X}(y|x)^\ell W_{Y|X}(y|x)^{1-\ell}\right)$$

denote the Chernoff information given an input type Q. We also define the relative entropy given an input type as

$$D\left(P_{Y|X}\middle\|W_{Y|X}\middle|Q\right) \triangleq \mathbb{E}_Q\left[D\left(P_{Y|X}(\cdot|X)\middle\|W_{Y|X}(\cdot|X)\right)\right].$$

III. PROBLEM FORMULATION AND MAIN RESULTS

We consider a controlled sensing problem in which the actions chosen by the decision maker are corrupted. Let Θ be the finite set of hypotheses, $\mathcal{A}$ be the finite set of actions and n be the time horizon. Without loss of generality, we assume $\mathcal{A}$ is also a group with the addition operator $+$. In a conventional controlled sensing problem, at each time $t \in [1;n]$, the observation y_t is generated from a known distribution $P^\theta_{Y|X}(\cdot|a_t)$ that depends on the true hypothesis $\theta \in \Theta$ and the chosen action a_t. We consider here a setting in which the action chosen by the decision maker is potentially corrupted and shifted by an amount w drawn according an unknown distribution $P_W \in \mathcal{P}_\mathcal{A}$. To avoid confusing terminology, we call *commands* the sequence of random variables $\{A_t\}_{t=1}^n$ determined by the decision maker, and *actions* the sequence of random variables $\{X_t\}_{t=1}^n$ that effectively influences the distribution of observations. We denote by C the variable that indicates the corruption status. When $C = -1$, the action corresponding to a command A_t is $X_t \sim \sum_{w\in\mathcal{A}} P_W(w)(A_t + w)$, where the operator $+$ is the addition over the group $\mathcal{A}$; when $C = 0$, the action is equal to the command and $X_t = A_t$. Note that when $C = -1$, the effective kernel from which observations are drawn after corruption is $(P^\theta_{Y|X} \circ P_W)(y_t|a_t) \triangleq \sum_{w\in\mathcal{A}} P_W(w) P^\theta_{Y|X}(y_t|a_t + w)$. In our model, the set of actions $\mathcal{X}$ is equal to $\mathcal{A}$.

We restrict our focus to deterministic non-adaptive (also called open-loop) strategies, for which the command sequence is $A^n = a^n$ for some $a^n \in \mathcal{A}^n$ determined *without* the knowledge of Y^n. The probability of receiving $y^n \in \mathcal{Y}^n$ given the hypothesis $\theta \in \Theta$ and the command sequence a^n without corruption is given by

$$P(y^n|\theta, a^n, 0) \triangleq \prod_{t=1}^n P^\theta_{Y|X}(y_t|a_t).$$

Similarly, when $C = -1$, the probability of receiving $y^n \in \mathcal{Y}^n$ under the corruption probability P_W is given by

$$P(y^n|\theta, a^n, -1) \triangleq \prod_{t=1}^n \left(\sum_{w\in\mathcal{A}} P_W(w) P^\theta_{Y|X}(y_t|a_t + w)\right).$$

After obtaining observations, the decision function $\phi : \mathcal{Y}^n \mapsto \Theta \cup \{-1\}$ either maps the observations to the estimated hypothesis or -1 to declare an existence of corruptions. A deterministic open loop policy $\lambda \triangleq (a^n, \phi)$ is a tuple characterized by a^n and ϕ. Given the policy λ, we define two kinds of error probabilities $P^n_{E,0}$ and $P^n_{E,-1}$ as follows.

$$P^n_{E,0}(\lambda) \triangleq \max_{\theta\in\Theta} P(y^n \notin \phi^{-1}(\{\theta\})|\theta, a^n, 0),$$
$$P^n_{E,-1}(\lambda) \triangleq \max_{\theta\in\Theta} P(y^n \notin \phi^{-1}(\{\theta, -1\})|\theta, a^n, -1).$$

The definition of $P^n_{E,-1}$ implies that, in the presence of corruptions, the decision maker makes a correct decision either by identifying the correct hypothesis or identifying the presence of corruptions. The reason behind this definition is elaborated in Remark 1. We have the following definition of achievability.

Definition 1 (Achievability)**.** *Policy λ achieves the exponent γ if for any $\epsilon > 0$ it holds that*

$$\lim_{n\to\infty} -\frac{1}{n} \log P^n_{E,0}(\lambda) \geqslant \gamma \tag{1}$$
$$\lim_{n\to\infty} P^n_{E,-1}(\lambda) < \epsilon \tag{2}$$

for any P_W. We denote $\gamma(\lambda)$ the supremum of all achievable exponent for a given policy λ.

This formulation of achievable exponents is inspired by Stein's lemma, which analyzes the maximum type-I error exponent under a fixed constraint on the type-II error probability. The asymmetry introduced in Definition 1 implies that we do not demand an exponentially decaying error probability when corruptions happen. The objective of the present work is to analyze the supremum of all achievable exponents

$$\gamma^* \triangleq \sup_\lambda \gamma(\lambda). \tag{3}$$

Remark 1. *When there are corruptions on the action sequence, the decision maker might not be able to identify the true hypothesis correctly. We illustrate this by the following example of binary channels. Let $\theta \in \{0, 1\}$, $\mathcal{A} = \mathcal{X} = \{0, 1\}$ and $\mathcal{Y} = \{0, 1\}$. The channels are specified in Table I. In this example, one can observe that sending $X_t = 1$ for all $t \in [1;n]$ is the best for distinguish the two channels in the conventional setting without adversaries because $C\left(P^0_{Y|X}(\cdot|1)\middle\|P^1_{Y|X}(\cdot|1)\right) = \infty$. Estimators that have the form $\phi(y^n) = \mathbb{1}(\hat{p}_{y^n}(0) > T)$ for any $0 < T < 1$ can perfectly detect the hypothesis without any error. However, there exists some P_W such that either $P^0_{Y|X} \circ P_W(0|1) > T$ or $P^1_{Y|X} \circ P_W(0|1) < T$ for all $0 < T < 1$. This means that*

TABLE I
TABLE FOR $P^\theta_{Y|X}(0|x)$ FOR ALL $x \in \{0, 1\}$ AND $\theta \in \{0, 1\}$.

x θ	0	1
0	0.9	0
1	0.1	1

no detector can correctly identify the true hypothesis with a non-zero error exponent for all P_W if the decision maker must choose a hypothesis. However, one can nevertheless identify the existence of an adversary by examining the empirical type of y^n. In this example, any P_W with non-zero $P_W(1)$ will be detected because y^n can only be an all 1 or an all 0 sequence in the conventional setting. Therefore, we allow the decision maker to output the -1 symbol to declare the presence of an adversary.

Our main results is the following characterization of γ^*.

Theorem 2. *The optimum exponent γ^* is*

$$\gamma^* = \sup_Q \min_{\theta\in\Theta} \min_{\theta'\neq\theta} \min \left(\min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right), \right.$$
$$\left. C\left(P^{\theta}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \right). \quad (4)$$

Recall that a channel W' is degraded with respect to (w.r.t.) another channel W if there exists a distribution P such that $W' = W \circ P$. Then, we have the following corollary.

Corollary 3. *If there exists a pair (θ, θ'), where $\theta \in \Theta$ and $\theta' \neq \theta$, such that $P^{\theta'}_{Y|X}$ is a degraded channel w.r.t $P^{\theta}_{Y|X}$, then $\gamma^* = 0$.*

Remark 2. *We illustrate Corollary 3 by the following example. Let $\theta \in \{0,1\}$, $\mathcal{A} = \{0,1\}$ and $\mathcal{Y} = \{0,1\}$. When $\theta = 0$, the transition probability $P^0_{Y|X}$ is a symmetric channel with cross-over probability p. When $\theta = 1$, $P^1_{Y|X}$ is a symmetric channel with cross-over probability q. Assume without loss of generality that $1/2 > q > p > 0$. When $P_W(1) = \frac{q-p}{1-2p}$, $P^0_{Y|X} \circ P_W = P^1_{Y|X}$, and the decision maker is unable to correctly declare the existence of corruptions.*

Before proving Theorem 2, we provide a high level intuition about the proof and the meaning of this theorem, also illustrated in Fig. 1. Without loss of generality, we fix Q as the type of the command sequence in the discussion below. When an adversary exists, given two hypotheses θ and θ', the corresponding corrupted distributions $P^{\theta}_{Y|X} \circ P_W$ and $P^{\theta'}_{Y|X} \circ P_W$ vary according to the choice of P_W. We illustrate this by plotting the red dash line in Fig. 1. Let $\mathcal{L}_\theta \triangleq \{P^{\theta}_{Y|X} \circ P_W : P_W \in \mathcal{P}_{\mathcal{A}}\}$ and $\mathcal{L}_{\theta'} \triangleq \{P^{\theta'}_{Y|X} \circ P_W : P_W \in \mathcal{P}_{\mathcal{A}}\}$ be the trajectories of $P^{\theta}_{Y|X} \circ P_W$ and $P^{\theta'}_{Y|X} \circ P_W$. When $\mathcal{L}_\theta$ and $\mathcal{L}_{\theta'}$ overlap, i.e., there exists some P'_W and P''_W such that $P^{\theta}_{Y|X} \circ P'_W = P^{\theta'}_{Y|X} \circ P''_W$, the decision maker cannot distinguish between

$$\mathcal{H} \triangleq \text{ true parameter } \theta \text{ with corruption distribution } P'_W;$$
$$\mathcal{H}' \triangleq \text{true parameter } \theta' \text{ with corruption distribution } P''_W.$$

This means

$$\mathbb{P}\left(\phi(Y^n) \neq \theta | \mathcal{H}\right) + \mathbb{P}\left(\phi(Y^n) \neq \theta' | \mathcal{H}'\right) \geqslant 1, \quad (5)$$

and hence there is no detector that can achieve a non-zero error exponent for all $\theta \in \Theta$ and for all P_W when an adversary exists and $\mathcal{L}_\theta \cap \mathcal{L}_{\theta'} \neq \emptyset$ for some $\theta' \neq \theta$. The intersection of $\mathcal{L}_\theta$ and $\mathcal{L}_{\theta'}$ is marked as the blue dash line region in Fig. 1. Therefore, we allow the decision maker to output the -1 symbol when $\hat{p}_{y^n|a^n} \in \mathcal{L}_\theta \cap \mathcal{L}_{\theta'}$ for some $\theta' \neq \theta$, and the decision maker only outputs an estimation of the hypothesis when $D\left(\hat{p}_{y^n|a^n} \middle\| P^{\theta}_{Y|X} \middle| Q\right) < \Lambda$ for some $\theta \in \Theta$. By choosing Λ smaller than

$$\min_{\theta\in\Theta} \min_{\theta'\neq\theta} \min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right),$$

we can guarantee that $\hat{p}_{y^n|a^n} \notin \mathcal{L}_\theta \cap \mathcal{L}_{\theta'}$ for all $\theta' \neq \theta$ when $D\left(\hat{p}_{y^n|a^n} \middle| P^{\theta}_{Y|X} \middle| Q\right) < \Lambda$ for some $\theta \in \Theta$. Finally, the term

$$\sup_Q \min_{\theta\in\Theta} \min_{\theta'\neq\theta} C\left(P^{\theta}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \quad (6)$$

in Theorem 2 corresponds to the best achievable exponent when there is no adversary.

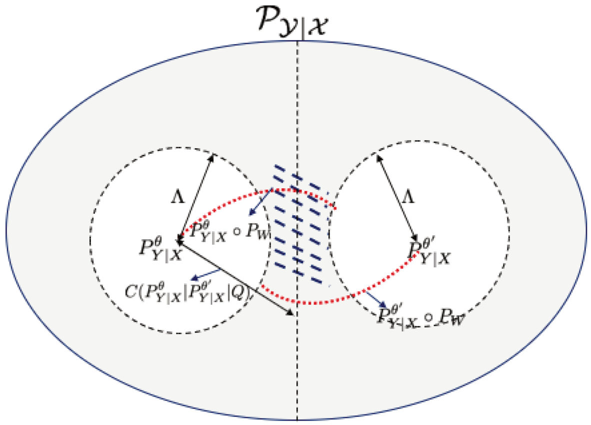
$\mathcal{P}_{\mathcal{Y}|\mathcal{X}}$
Λ
$P^{\theta}_{Y|X} \circ P_W$
$P^{\theta}_{Y|X}$
$C(P^{\theta}_{Y|X}|P^{\theta'}_{Y|X}|Q)$
Λ
$P^{\theta'}_{Y|X}$
$P^{\theta'}_{Y|X} \circ P_W$

Fig. 1. The space of all conditional types. Each location in the figure is a specific conditional type. The sets $\mathcal{L}_\theta$ and $\mathcal{L}_{\theta'}$ are the red dash lines. The blue dash line region contains the conditional types in $\mathcal{L}_\theta \cap \mathcal{L}_{\theta'}$. The decision maker output the -1 symbol when the conditional empirical type is located in the area with grey background.

IV. Achievability Proof of Theorem 2

We first specify the policy as follows. Let the sequence of command $A^n = a^n$ for some $a^n \in \mathcal{A}^n$ and have the type $\hat{p}_{a^n} = Q$. The decision function ϕ is defined as

$$\phi(y^n) = \begin{cases} -1 & \text{if } D\left(\hat{p}_{y^n|a^n} \middle\| P^{\theta}_{Y|X} \middle| Q\right) \geqslant \Lambda \quad \forall \theta \in \Theta \\ \arg\max_{\theta\in\Theta} P(y^n|\theta, a^n, 0) & \text{otherwise.} \end{cases}$$

The policy λ is then well defined. We also define $\tilde{\phi}$ as the ML estimator, i.e.,

$$\tilde{\phi}(y^n) = \arg\max_{\theta\in\Theta} P(y^n|\theta, a^n, 0)$$

for all $y^n \in \mathcal{Y}^n$. The probability $P^n_{E,-1}(\lambda)$ can be upper bounded by

$$P^n_{E,-1}(\lambda) \leqslant \max_{\theta\in\Theta} \min \Big(P(y^n \notin \tilde{\phi}^{-1}(\{\theta\}) | \theta, a^n, -1),$$

$$P(y^n \notin \phi^{-1}(\{-1\})|\theta, a^n, -1)\Big), \quad (7)$$

where we have use the fact that

$$\{y^n : \phi(y^n) \in \Theta \setminus \{\theta\}\} \subset \left\{y^n : \tilde{\phi}(y^n) \in \Theta \setminus \{\theta\}\right\}.$$

Note that for any $\theta \in \Theta$,

$$\begin{aligned} &P(y^n \notin \tilde{\phi}^{-1}(\{\theta\})|\theta, a^n, -1) \\ &\leqslant \sum_{\theta' \neq \theta} \sum_{y^n \in \mathcal{Y}^n} \mathbb{1}\left(\frac{P(y^n|\theta', a^n, 0)}{P(y^n|\theta, a^n, 0)} \geqslant 1\right) P(y^n|\theta, a^n, -1). \end{aligned}$$

We show in Appendix A that the ML estimator

$$P(y^n|\theta', a^n, 0) \underset{\theta}{\overset{\theta'}{\gtrless}} P(y^n|\theta, a^n, 0)$$

is equivalent to

$$D\left(\hat{p}_{y^n|a^n} \middle\| P^{\theta}_{Y|X} \middle| Q\right) \underset{\theta}{\overset{\theta'}{\gtrless}} D\left(\hat{p}_{y^n|a^n} \middle\| P^{\theta'}_{Y|X} \middle| Q\right). \quad (8)$$

We define $\mathcal{E}_{\theta \to \theta'}$ as the set of all conditional types in $\mathcal{P}(Y|X)$ that satisfies (8), i.e.,

$$\mathcal{E}_{\theta \to \theta'} \triangleq \left\{\bar{P}_{Y|X} : \left(D\left(\bar{P}_{Y|X} \middle\| p^{\theta}_{Y|X} \middle| Q\right) \geqslant D\left(\bar{P}_{Y|X} \middle\| p^{\theta'}_{Y|X} \middle| Q\right)\right)\right\}.$$

Moreover, for any type $\bar{P}_{Y|X} \in \mathcal{P}_{\mathcal{Y}|\mathcal{X}}$, the probability of receiving any y^n in the type class $\mathcal{T}_{a^n}(\bar{P}_{Y|X})$ under the command sequence a^n and the corruption probability P_W is given by

$$P(y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X})|\theta, a^n, -1) \leqslant e^{-nD\left(\bar{P}_{Y|X} \| P^{\theta}_{Y|X} \circ P_W | Q\right)}.$$

Then, the probability $P(y^n \notin \tilde{\phi}^{-1}(\{\theta\})|\theta, a^n, -1)$ can be upper bounded by

$$\begin{aligned} &P(y^n \notin \tilde{\phi}^{-1}(\{\theta\})|\theta, a^n, -1) \\ &\leqslant \sum_{\theta' \neq \theta} \sum_{\bar{P}_{Y|X} \in \mathcal{E}_{\theta \to \theta'}} e^{-nD\left(\bar{P}_{Y|X} \| P^{\theta}_{Y|X} \circ P_W | Q\right)} \\ &\leqslant \text{poly}(n) e^{-n \min_{\theta' \neq \theta} \min_{\bar{P}_{Y|X} \in \mathcal{E}_{\theta \to \theta'}} D\left(\bar{P}_{Y|X} \| P^{\theta}_{Y|X} \circ P_W | Q\right)} \end{aligned}$$

If the corruption distribution P_W is such that

$$\min_{\theta' \neq \theta} \min_{\bar{P}_{Y|X} \in \mathcal{E}_{\theta \to \theta'}} D\left(\bar{P}_{Y|X} \middle\| P^{\theta}_{Y|X} \circ P_W \middle| Q\right) > 0, \quad (9)$$

then $\lim_{n\to\infty} P(y^n \notin \tilde{\phi}^{-1}(\{\theta\})|\theta, a^n, -1) = 0$. On the other hand, if (9) does not hold, then it means there exists some $\tilde{\theta} \in \Theta \setminus \{\theta\}$ and some $\bar{P}_{Y|X} \in \mathcal{E}_{\theta \to \tilde{\theta}}$ satisfying $\bar{P}_{Y|X} = P^{\theta}_{Y|X} \circ P_W$. This implies $P^{\theta}_{Y|X} \circ P_W \in \mathcal{E}_{\theta \to \tilde{\theta}}$, and

$$D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta}_{Y|X} \middle| Q\right) \geqslant C\left(P^{\theta}_{Y|X} \middle\| P^{\tilde{\theta}}_{Y|X} \middle| Q\right) \quad (10)$$

as shown in Appendix B. Moreover, for all $\theta' \neq \theta$,

$$\begin{aligned} &D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \\ &\qquad \geqslant \min_{P_W} \min_{\theta'' \neq \theta} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta''}_{Y|X} \middle| Q\right). \quad (11) \end{aligned}$$

By defining

$$\begin{aligned} \tilde{\Lambda} \triangleq \min_{\theta \in \Theta} \min_{\theta' \neq \theta} \min \Bigg(&\min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right), \\ &C\left(P^{\theta}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \Bigg), \end{aligned}$$

one can observe that

$$D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \geqslant \tilde{\Lambda} \quad (12)$$

for all $\theta' \in \Theta$. We define the threshold Λ as

$$\Lambda \triangleq \tilde{\Lambda} - \eta \quad (13)$$

for some $\eta > 0$. Then, $\forall \epsilon > 0$ and for all P_W that does not satisfies (9), the probability $P(y^n \notin \phi^{-1}(\{-1\}|\theta, a^n, -1)$ can be upper bounded by

$$\begin{aligned} &P(\phi(y^n) \neq -1|\theta, a^n, -1) \\ &\leqslant P\left(\left|\hat{p}_{y^n|a^n} - P^{\theta}_{Y|X} \circ P_W\right|_{\infty} > \epsilon \middle| \theta, a^n, -1\right) \\ &+ 1\left(\min_{\theta'} \min_{\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P_W, \epsilon)} D\left(\bar{P}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) < \Lambda\right). \end{aligned} \quad (14)$$

The first term on the right hand side (14) goes to zero when $n \to \infty$ for any $\epsilon > 0$ by concentration inequalities, and the second term is greater than $\Lambda = \tilde{\Lambda} - \eta$ by the continuity property of divergence when we choose ϵ sufficiently small. Therefore, for any $\eta > 0$, there exists some $\epsilon > 0$ such that $\lim_{n\to\infty} P(\phi(y^n) \neq -1|\theta, a^n, -1) = 0$ whenever P_W does not satisfy (9). Then, we have shown that $\lim_{n\to\infty} P^n_{E,-1}(\lambda) = 0$ for all P_W. We next proceed to derive the exponent for $P^n_{E,0}$ as follows.

$$\begin{aligned} P^n_{E,0}(\lambda) \leqslant \max_{\theta \in \Theta} \Bigg(&P(\phi(y^n) = -1|\theta, a^n, 0) \\ &+ \sum_{\theta' \neq \theta} P(\tilde{\phi}(y^n) = \theta'|\theta, a^n, 0) \Bigg). \end{aligned} \quad (15)$$

By denoting

$$\mathcal{E}_{\theta \to -1} \triangleq \left\{\bar{P}_{Y|X} : D\left(\bar{P}_{Y|X} \middle\| P^{\theta}_{Y|X} \middle| Q\right) \geqslant \Lambda\right\},$$

the first term inside the maximization of (15) can be upper bounded by

$$\begin{aligned} &P(\phi(y^n) = -1|\theta, a^n, 0) \\ &\leqslant P\left(D\left(\hat{p}_{y^n|a^n} \middle\| P^{\theta}_{Y|X} | Q\right) \geqslant \Lambda \middle| \theta, a^n, 0\right) && (16) \\ &\leqslant \sum_{\bar{P}_{Y|X} \in \mathcal{E}_{\theta \to -1}} e^{-nD\left(\bar{P}_{Y|X} \| P^{\theta}_{Y|X} | Q\right)} && (17) \\ &\leqslant \text{poly}(n) e^{-n\Lambda} && (18) \end{aligned}$$

for all $\theta \in \Theta$. Moreover, it is known that

$$\max_{\theta' \neq \theta} P(\tilde{\phi}(y^n) = \theta'|\theta, a^n, 0) \leqslant e^{-n \min_{\theta' \neq \theta} C\left(P^{\theta}_{Y|X} \| P^{\theta'}_{Y|X} | Q\right)},$$

and hence

$$\lim_{n\to\infty} -\frac{1}{n} \log P^n_{E,0} \geqslant \Lambda. \quad (19)$$

by the fact that the error probability is dominated by the term with the minimum exponent. Finally, the achievability is proved by choosing the type Q that maximize $\tilde{\Lambda}$ and make η arbitrarily small.

V. Converse Proof of Theorem 2

We prove the converse of Theorem 2 by contradiction. Let's fix the pair (θ, θ'), where $\theta \in \Theta$ and $\theta' \neq \theta$ and also fix the type of the command sequence Q. We assume that there exists some $\lambda^{\#} = (\phi^{\#}, a^n)$, where $\hat{p}_{a^n} = Q$, such that

$$\gamma(\lambda^{\#}) > \min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right). \tag{20}$$

For all $W_{Y|X} \in \mathcal{P}_{\mathcal{Y}|\mathcal{X}}$ and $\epsilon > 0$, we define

$$\mathcal{B}(W_{Y|X}, \epsilon) \triangleq \left\{\bar{P}_{Y|X} \in \mathcal{P}_{\mathcal{Y}|\mathcal{X}} : \left|\bar{P}_{Y|X} - W_{Y|X}\right|_{\infty} \leqslant \epsilon\right\}.$$

Then, by the continuity property of the divergence, there exists some $P^{\theta,\theta'}_W$ and ϵ small enough such that

$$\begin{aligned} &\min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \\ &\qquad < D\left(\bar{P}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) < \gamma(\lambda^{\#}) \end{aligned} \tag{21}$$

for all $\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)$. Note that for any $y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X})$ and for all $\bar{P}_{Y|X}$, the probability of receiving y^n when the true hypothesis is θ' can be express as

$$e^{-n\left(H\left(\bar{P}_{Y|X} \middle| Q\right) + D\left(\bar{P}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right)\right)}$$

This means for all $\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)$, the number of sequence in $\mathcal{T}_{a^n}(\bar{P}_{Y|X})$ that are decoded into the set $\{-1\} \cup \Theta \setminus \{\theta'\}$ is less than $e^{nH(\bar{P}_{Y|X}) - \xi}$ for some $\xi > 0$, otherwise,

$$\begin{aligned} &P(\phi(y^n) \neq \theta' | \theta', a^n, 0) \\ &\geqslant \sum_{\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)} \ \sum_{y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X})} \mathbb{1}(\phi^{\#}(y^n) \in \{-1\} \cup \Theta \setminus \{\theta'\}) \\ &\times e^{-n\left(H\left(\bar{P}_{Y|X} \middle\| Q\right) + D\left(\bar{P}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right)\right)} \\ &\geqslant e^{-nD\left(\bar{P}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right)} \\ &> e^{-n\gamma(\lambda^{\#})}, \end{aligned}$$

which contradicts with the assumption that the policy achieves the exponent $\gamma^{\#}$. Therefore, we conclude that for all $\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)$,

$$\frac{\left|\left\{y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X}) : \phi(y^n) \in \{-1\} \cup \Theta \setminus \{\theta'\}\right\}\right|}{|\mathcal{T}_{a^n}(\bar{P}_{Y|X})|} \leqslant e^{-n\xi}$$

for some $\xi > 0$. Finally, when the true hypothesis is θ, the value of $P^n_{E,-1}$ under the corruption distribution $P^{\theta,\theta'}_W$ is lower bounded by

$$\begin{aligned} P^n_{E,-1} \geqslant\ & 1 - P(\hat{p}_{y^n|a^n} \notin \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon) | \theta, a^n, -1) \\ & - \sum_{\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)} P(\mathcal{T}_{a^n}(\bar{P}_{Y|X}) | \theta, a^n, -1) \\ & \times \frac{|\{y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X}) : \phi(y^n) \neq \theta'\}|}{|\mathcal{T}_{a^n}(\bar{P}_{Y|X})|}. \end{aligned} \tag{22}$$

By the law of large number

$$\lim_{n\to\infty} P(\hat{p}_{y^n|a^n} \notin \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon) | \theta, a^n, -1) = 0$$

for any $\epsilon > 0$, and

$$\begin{aligned} \lim_{n\to\infty} & \sum_{\bar{P}_{Y|X} \in \mathcal{B}(P^{\theta}_{Y|X} \circ P^{\theta,\theta'}_W, \epsilon)} P(\mathcal{T}_{a^n}(\bar{P}_{Y|X}) | \theta, a^n, -1) \\ & \times \frac{|\{y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X}) : \phi(y^n) \neq \theta'\}|}{|\mathcal{T}_{a^n}(\bar{P}_{Y|X})|} = 0 \end{aligned}$$

by the fact that there are at most a polynomial number of types and $\frac{|\{y^n \in \mathcal{T}_{a^n}(\bar{P}_{Y|X}) : \phi(y^n) \neq \theta'\}|}{|\mathcal{T}_{a^n}(\bar{P}_{Y|X})|} < e^{-n\xi}$. Hence, we have

$$\lim_{n\to\infty} P^n_{E,-1} = 1, \tag{23}$$

which contradicts with the achievability definition of the exponent. This argument is true for any (θ, θ') pairs. So, for all λ that has input type Q, it holds that

$$\gamma(\lambda) \leqslant \min_{\theta \in \Theta} \min_{\theta' \neq \theta} \min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| Q\right). \tag{24}$$

We also know that

$$\gamma(\lambda) \leqslant \min_{\theta \in \Theta} \min_{\theta' \neq \theta} C\left(P^{\theta}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| Q\right) \tag{25}$$

from the result of multi-hypotheses testing. Finally, the converse proof to the Theorem is done by choosing a type Q that maximizes the minimum on the right hand side of (24) and (25).

VI. Numerical Examples

In this section, we present two numerical examples to illustrate our main result. We consider a set of binary channels as in Table II and III, where $\mathcal{A} = \mathcal{X} = \{0, 1\}$, $\mathcal{Y} = \{0, 1\}$ and $\Theta = \{0, 1, 2\}$. The commands chosen by the decision maker are the input of the channel. Any distribution $Q \in \mathcal{P}_{\mathcal{X}}$ is a Bernoulli distribution and is characterized by a parameter $q \in [0, 1]$ and we denote it by $\mathcal{B}(q)$. Then, we define

$$C_{\theta\theta'}(q) \triangleq C\left(P^{\theta}_{Y|X} \middle\| P^{\theta'}_{Y|X} \middle| \mathcal{B}(q)\right)$$

$$E_{\theta\theta'}(q) \triangleq \min\left(\min_{P_W} D\left(P^{\theta}_{Y|X} \circ P_W \middle\| P^{\theta'}_{Y|X} \middle| \mathcal{B}(q)\right), C_{\theta\theta'}(q)\right),$$

and $\bar{\gamma}(q) \triangleq \min_{\theta \in \Theta} \min_{\theta' \neq \theta} \min E_{\theta\theta'}(q)$ is the best exponent we can obtain by using the policy with input type $\mathcal{B}(q)$.

In our first example, i.e., channels specified in Table II, one can observe that the decision maker can not estimate the true hypothesis by keeping choosing the same input. The input 0 is the best for distinguishing $\theta = 2$ from $\theta = 0$ and $\theta = 1$, but this input is not helpful for distinguish $\theta = 0$ from $\theta = 1$. This phenomenon can be seen from Fig.2, where $C_{0,2}$ and $C_{1,2}$ decrease with the parameter q while $C_{0,1}$ increases with q. In the conventional multi-hypotheses scenario, the optimum value of q needs to strike a balance in the aforementioned trade-off and has the value approximately equal to 0.83 in our first example.

When the adversary exists, the optimal value of q changes. To see this point, we first observe that $q = 0$ is supposed to be the value that is good for distinguish between $\theta = 0$ and $\theta = 2$. However, under such value of q, when $P_W(1) = 1/3$, it holds that

$$\begin{aligned}
&P^0_{Y|X}(0|0) \times P_W(0) + P^0_{Y|X}(0|1) \times (1 - P_W(0)) \\
&= 0.9 \times \frac{2}{3} + 0.3 \times \frac{1}{3} = 0.7 \\
&= P^2_{Y|X}(0|0),
\end{aligned}$$

which leads to a zero value of $E_{0,2}$. A similar phenomenon happens to distinguishing $\theta = 1$ and $\theta = 2$ by using $q = 1$ as well. There exists some P_W such that

$$\begin{aligned}
&P^1_{Y|X}(0|1) \times P_W(0) + P^1_{Y|X}(0|0) \times (1 - P_W(0)) \\
&= P^2_{Y|X}(0|1)
\end{aligned}$$

in our example. From the E_{12} curve in Fig. 2, it can be observed that a high value of q has a bad performance in distinguishing $\theta = 1$ and $\theta = 2$ when an adversary exists, and this leads to a decreased value of the optimal q. The optimal γ^* happens when $q \approx 0.73$, which is lower than its counterpart in the conventional setting.

In Fig. 2, there exists a unique positive γ^*. In contrast, the best achievable exponent γ^* in our second example, i.e., channels specified in Table III, is zero as shown in Fig.3. This result comes from Corollary 3 and the fact that $P^2_{Y|X}$ is a degraded channel w.r.t. $P^1_{Y|X}$.

TABLE II
TABLE FOR $P^{\theta}_{Y|X}(0|x)$ FOR ALL $x \in \{0, 1\}$ AND $\theta \in \{0, 1, 2\}$.

x θ	0	1
0	0.9	0.3
1	0.9	0.2
2	0.7	0.3

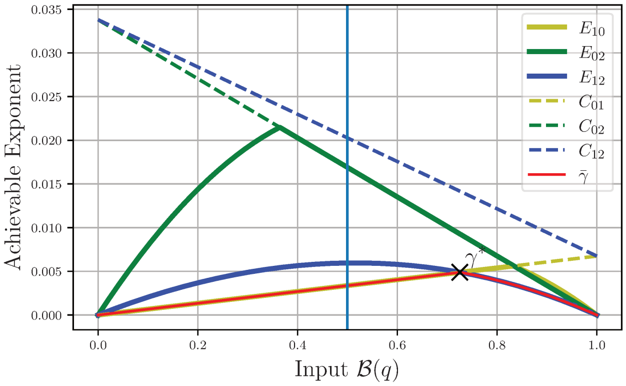

Fig. 2. Trajectories of $\bar{\gamma}(q)$, $E_{\theta,\theta'}(q)$ and $C_{\theta,\theta'}(q)$ for some (θ, θ') pairs when channels are specified in Table II.

TABLE III
TABLE FOR $P^{\theta}_{Y|X}(0|x)$ FOR ALL $x \in \{0, 1\}$ AND $\theta \in \{0, 1, 2\}$.

x θ	0	1
0	0.9	0.3
1	0.8	0.2
2	0.7	0.3

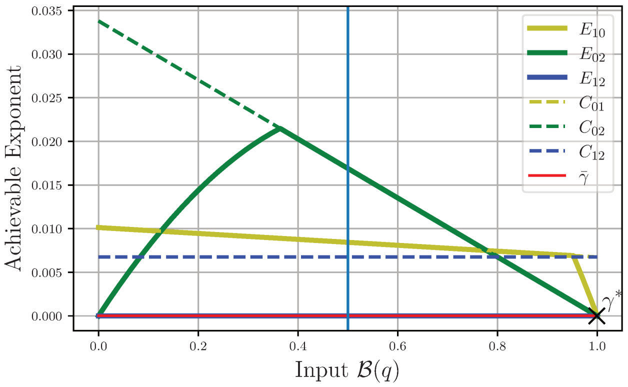

Fig. 3. Trajectories of $\bar{\gamma}(q)$, $E_{\theta,\theta'}(q)$ and $C_{\theta,\theta'}(q)$ for some (θ, θ') pairs when channels are specified in Table III.

VII. Conclusion

In this work, we study the problem of controlled sensing with corruption commands and analyze the error exponent in a setting inspired by the works in authentication and Stein's lemma. A characterization of the optimal exponent is given in our main theorem, and the rigorous proof to the theorem is also given. We also provide some numerical examples to illustrate the theorem. We have assume throughout the work that the adversary influences the hypothesis testing by generating the noise sequence identically and independently from a certain corruption distribution. However, a stronger adversary who has the information of decision maker's observations can be considered. Our future work is to analyze such scenario.

Appendix A
Proof of (8)

The event $P(y^n|\theta', a^n, 0) > P(y^n|\theta, a^n, 0)$ implies

$$\log \frac{P(y^n|\theta', a^n, 0)}{P(y^n|\theta, a^n, 0)} > 0,$$

where the log likelihood ratio can be written as

$$\begin{aligned}
&\log \frac{P(y^n|\theta', a^n, 0)}{P(y^n|\theta, a^n, 0)} \\
&= \log \frac{\prod_{t=1}^n W^{\theta'}_{Y|X}(y_t|a_t)}{\prod_{t=1}^n W^{\theta}_{Y|X}(y_t|a_t)} \\
&= \log \frac{\prod_{a\in\mathcal{A}} \prod_{y\in\mathcal{Y}} \left(W^{\theta'}_{Y|X}(y|a)\right)^{N_{y^n,a^n}(y,a)}}{\prod_{a\in\mathcal{A}} \prod_{y\in\mathcal{Y}} \left(W^{\theta}_{Y|X}(y|a)\right)^{N_{y^n,a^n}(y,a)}}
\end{aligned}$$

$$
\begin{aligned}
&= \sum_{a\in\mathcal{A}}\sum_{y\in\mathcal{Y}} N_{y^n,a^n}(y,a)\log\frac{W^{\theta'}_{Y|X}(y|a)}{W^{\theta}_{Y|X}(y|a)}\\
&= n\sum_{a\in\mathcal{A}} Q(a)\sum_{y\in\mathcal{Y}}\hat{p}_{y^n|a^n}(y|a)\log\frac{W^{\theta'}_{Y|X}(y|a)}{W^{\theta}_{Y|X}(y|a)}\\
&= n\sum_{a\in\mathcal{A}} Q(a)\sum_{y\in\mathcal{Y}}\hat{p}_{y^n|a^n}(y|a)\log\frac{W^{\theta'}_{Y|X}(y|a)}{W^{\theta}_{Y|X}(y|a)}\log\frac{\hat{p}_{y^n|a^n}(y|a)}{\hat{p}_{y^n|a^n}(y|a)}\\
&= n\left(D\left(\hat{p}_{y^n|a^n}\middle\|W^{\theta}_{Y|X}\middle|Q\right)-D\left(\hat{p}_{y^n|a^n}\middle\|W^{\theta'}_{Y|X}\middle|Q\right)\right).
\end{aligned}
$$

Therefore, we have shown that the event $P(y^n|\theta',a^n,0) > P(y^n|\theta,a^n,0)$ is equivalent to

$$
D\left(\hat{p}_{y^n|a^n}\middle\|W^{\theta}_{Y|X}\middle\|Q\right)-D\left(\hat{p}_{y^n|a^n}\middle\|W^{\theta'}_{Y|X}\middle|Q\right)>0.
$$

APPENDIX B
PROOF OF (10)

We derive the minimum of $D\left(\bar{P}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)$ when $\bar{P}_{Y|X}\in\mathcal{E}_{\theta\to\theta'}$ by using the method Lagrange multiplier as follows. Let $J(\bar{P}_{Y|X},\zeta,\{c_x\}_{x\in\mathcal{X}})$ be defined as

$$
\begin{aligned}
&J(\bar{P}_{Y|X},\zeta,\{c_x\}_{x\in\mathcal{X}})\\
&= D\left(\bar{P}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)\\
&+\zeta\left(D\left(\bar{P}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)-D\left(\bar{P}_{Y|X}\middle\|P^{\theta'}_{Y|X}\middle|Q\right)\right)\\
&+\sum_{x\in\mathcal{X}} c_x\left(\sum_{y\in\mathcal{Y}}\bar{P}_{Y|X}(y|x)-1\right).
\end{aligned}
$$

Take derivative of J w.r.t. each $\bar{P}_{Y|X}(y|x)$ and ζ, we have

$$
\begin{cases}
Q(x)\left(\log\frac{\bar{P}_{Y|X}(y|x)}{P^{\theta}_{Y|X}}+1+\zeta\log\frac{P^{\theta'}_{Y|X}}{P^{\theta}_{Y|X}}\right)+c_x=0 & \forall x,y\\
D\left(\bar{P}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)-D\left(\bar{P}_{Y|X}\middle\|P^{\theta'}_{Y|X}\middle|Q\right)=0.
\end{cases}
$$

By solving the above equality and using the fact that $\sum_{y\in\mathcal{Y}}\bar{P}_{Y|X}(y|x)=1$ for all $x\in\mathcal{X}$, we obtain that the optimum $\bar{P}^{*}_{Y|X}(y|x)$ has the following form

$$
\bar{P}_{\ell}(y|x)=\frac{P^{\theta}_{Y|X}(y|x)^{\ell}P^{\theta'}_{Y|X}(y|x)^{1-\ell}}{\sum_{y\in\mathcal{Y}}P^{\theta}_{Y|X}(y|x)^{\ell}P^{\theta'}_{Y|X}(y|x)^{1-\ell}}
$$

where ℓ is chosen such that

$$
D\left(\bar{P}_{\ell}\middle\|P^{\theta}_{Y|X}\middle|Q\right)=D\left(\bar{P}_{\ell}\middle\|P^{\theta'}_{Y|X}\middle|Q\right). \tag{26}
$$

Let $\bar{P}^{*}_{Y|X}(y|x)=\bar{P}_{\ell^*}$ for some ℓ^* satisfying (26). Then, the minimum divergence becomes

$$
\begin{aligned}
&D\left(\bar{P}^{*}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)\\
&=\sum_{x\in\mathcal{X}}Q(x)\sum_{y\in\mathcal{Y}}\bar{P}^{*}_{Y|X}(y|x)\log\frac{\bar{P}^{*}_{Y|X}(y|x)}{P^{\theta}_{Y|X}(y|x)}\\
&=(\ell^*-1)\left(D\left(\bar{P}_{\ell^*}\middle\|P^{\theta'}_{Y|X}\middle|Q\right)-D\left(\bar{P}_{\ell^*}\middle\|P^{\theta}_{Y|X}\middle|Q\right)\right)\\
&-\sum_{x\in\mathcal{X}}Q(x)\log\sum_{y\in\mathcal{Y}}P^{\theta}_{Y|X}(y|x)^{\ell^*}P^{\theta'}_{Y|X}(y|x)^{1-\ell^*}\\
&=-\sum_{x\in\mathcal{X}}Q(x)\log\sum_{y\in\mathcal{Y}}P^{\theta}_{Y|X}(y|x)^{\ell^*}P^{\theta'}_{Y|X}(y|x)^{1-\ell^*}.\\
&=-\min_{\ell\in[0,1]}\sum_{x\in\mathcal{X}}Q(x)\log\sum_{y\in\mathcal{Y}}P^{\theta}_{Y|X}(y|x)^{\ell}P^{\theta'}_{Y|X}(y|x)^{1-\ell},
\end{aligned}
$$

where the last equality comes from the fact the value of ℓ that minimizes

$$
\sum_{x\in\mathcal{X}}Q(x)\log\sum_{y\in\mathcal{Y}}P^{\theta}_{Y|X}(y|x)^{\ell}P^{\theta'}_{Y|X}(y|x)^{1-\ell}
$$

lies within $[0,1]$ and is exactly the one satisfying (26). So, we have shown that

$$
\min_{\bar{P}_{Y|X}\in\mathcal{E}_{\theta\to\theta'}}D\left(\bar{P}_{Y|X}\middle\|P^{\theta}_{Y|X}\middle|Q\right)=C\left(P^{\theta}_{Y|X}\middle\|P^{\theta'}_{Y|X}\middle|Q\right).
$$

REFERENCES

[1] J. A. T. Thomas M. Cover, *Elements of Information Theory*. Wiley John + Sons, 2006.

[2] H. Chernoff, "Sequential design of experiments," *The Annals of Mathematical Statistics*, vol. 30, no. 3, pp. 755–770, 09 1959.

[3] M. Naghshvar and T. Javidi, "Active sequential hypothesis testing," *The Annals of Statistics*, vol. 41, no. 6, dec 2013.

[4] S. Nitinawarat, G. K. Atia, and V. V. Veeravalli, "Controlled sensing for multihypothesis testing," *IEEE Transactions on Automatic Control*, vol. 58, no. 10, pp. 2451–2464, oct 2013.

[5] D. Kartik, A. Nayyar, and U. Mitra, "Active hypothesis testing: Beyond chernoff-stein." Paris, France: IEEE, Jul. 2019.

[6] M.-C. Chang and M. R. Bloch, "Evasive active hypothesis testing," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 2, pp. 735–746, 2021.

[7] M. Tahmasbi and M. R. Bloch, "Active covert sensing," in *2020 IEEE International Symposium on Information Theory (ISIT)*, 2020, pp. 840–845.

[8] M.-C. Chang and M. R. Bloch, "Covert sequential hypothesis testing," in *2021 IEEE Information Theory Workshop (ITW)*, 2021, pp. 1–6.

[9] Y. Jin and L. Lai, "On the adversarial robustness of hypothesis testing," *IEEE Transactions on Signal Processing*, vol. 69, pp. 515–530, 2021.

[10] L. Lai, H. E. Gamal, and H. V. Poor, "Authentication over noisy channels," *IEEE Transactions on Information Theory*, vol. 55, no. 2, pp. 906–916, feb 2009.

[11] W. Tu and L. Lai, "Keyless authentication and authenticated capacity," *IEEE Transactions on Information Theory*, vol. 64, no. 5, pp. 3696–3714, may 2018.

[12] O. Kosut and J. Kliewer, "Authentication capacity of adversarial channels," in *2018 IEEE Information Theory Workshop (ITW)*. IEEE, nov 2018.

[13] A. Beemer, O. Kosut, J. Kliewer, E. Graves, and P. Yu, "Authentication against a myopic adversary," in *2019 IEEE Conference on Communications and Network Security (CNS)*. IEEE, jun 2019.