

Covert Vehicle Misguidance and Its Detection: A Hypothesis Testing Game Over Continuous-Time Dynamics

Takashi Tanaka¹, Senior Member, IEEE, Kenji Sawada¹, Member, IEEE,
Yohei Watanabe¹, Member, IEEE, and Mitsugu Iwamoto¹, Member, IEEE

Abstract—We formulate a stochastic zero-sum game over continuous-time dynamics to analyze the competition between the attacker, who tries to covertly misguide the vehicle to an unsafe region, versus the detector, who tries to detect the attack signal based on the observed trajectory of the vehicle. Based on Girsanov's theorem and the generalized Neyman-Pearson lemma, we show that a constant bias injection attack as the attacker's strategy and a likelihood ratio test as the detector's strategy constitute the unique saddle point of the game. We also derive the first-order and the second-order exponents of the type II error as a function of the data length.

Index Terms—Stochastic optimal control, fault detection, information theory and control.

I. INTRODUCTION

FALSE data injection (FDI) attacks are widely recognized as major threats to control systems. In [1], the authors performed a field experiment to misguide a 65-meter yacht to its unintended destination via GPS spoofing and demonstrated the vulnerability of modern maritime vessels to deceptive sensor data injection. In their experiment, the authors showed that a GPS deception attack, if carefully designed, can be disguised as the effects of natural disturbances such as slowly changing ocean currents and winds, and is difficult to detect unless the controller has an alternative source of reliable sensor data (e.g., radar and visual bearing).

In many circumstances, system faults (including malicious attacks) must be detected and isolated by continuous

Received 7 September 2024; revised 13 November 2024; accepted 29 November 2024. Date of publication 4 December 2024; date of current version 24 December 2024. This work was supported in part by AFOSR under Contract FA9550-20-1-0101; in part by NSF under Award 1944318; and in part by JST CREST under Grant JPMJCR23M2 and Grant JPMJCR23M4. Recommended by Senior Editor S. Dey. (Corresponding author: Takashi Tanaka.)

Takashi Tanaka is with the School of Aeronautics and Astronautics and the Elmore Family School of Electrical and Computer Engineering, Purdue University, Austin, TX 78712 USA (e-mail: tanaka16@purdue.edu).

Kenji Sawada is with the Info-Powered Energy System Research Center, The University of Electro-Communications, Tokyo 182-8585, Japan (e-mail: knj.sawada@uec.ac.jp).

Yohei Watanabe and Mitsugu Iwamoto are with the Department of Informatics, The University of Electro-Communications, Tokyo 182-8585, Japan (e-mail: watanabe@uec.ac.jp; mitsugu@uec.ac.jp).

Digital Object Identifier 10.1109/LCSYS.2024.3511398

monitoring of the sensor readouts. In the vessel misguidance example [1], the spoofed GPS signal may be distinguished from the natural background noise by an appropriate statistical test. However, knowing that the system is continuously monitored, a rational attacker will conduct a covert attack, maximizing the attack's impact while avoiding detection. Hence, a zero-sum game arises between the attacker and the detector, where the detector's purpose is to design the "most effective" statistical test for attack detection, whereas the attacker tries to inject the "most stealthy" attack signal.

Similar games between the attacker and the detector have been studied by many authors in the systems and control community. For example, the works [2], [3], [4] adopted the hypothesis testing theory to characterize covert FDI attacks against control systems. Invoking Stein's lemma, [2] introduced the notion of ϵ -stealthiness as measured by relative entropy. The worst-case degradation of linear control systems attainable by ϵ -stealthy attacks was studied in [3], [4]. Sequential and composite hypothesis testing frameworks have also been proposed (e.g., [5], [6] to name a few) for anomaly detection.

Despite recent progress, existing applications of hypothesis testing frameworks to control systems are limited to discrete-time settings. The goal of this letter is to broaden the scope of the literature by formulating the aforementioned zero-sum game in continuous time based on the generalized Neyman-Pearson theory [7]. We make the following methodological contributions:

- 1) We propose a novel zero-sum game formulation to model the competition between the attacker and the detector over continuous-time dynamics. Instead of taking relative entropy as a stealthiness measure for granted (the operational meaning of relative entropy in continuous-time hypothesis testing scenarios is not well-established in the literature), we use more fundamental quantities, such as probabilities of type I and type II errors and the probability of successful attacks, to formulate the game. While the game considered in this letter is simple, the results and methodologies we present are canonical and allow for various generalizations in future studies.
- 2) We show that a constant bias injection attack as the attacker's strategy and a likelihood ratio test as the detector's strategy constitute the unique saddle point of

the game. The proof is based on Girsanov's theorem and the generalized Neyman-Pearson lemma.

3) We analyze the exponent of the type II error as a function of the horizon length of the game and show that the first-order asymptote coincides with the relative entropy. This result is reminiscent of classical Stein's lemma. We also quantify the second-order asymptote, providing a tighter estimate of the error probability in the finite horizon length regime.

Notation: The normal distribution with mean m and covariance σ^2 is denoted by $\mathcal{N}(m, \sigma^2)$, and the cumulative distribution function of the standard normal distribution is denoted by $\Phi(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x \exp(-\frac{t^2}{2}) dt$. We write $[\cdot]^+ := \max\{0, \cdot\}$. $1_{\{\cdot\}}$ represents the indicator function. $C[0, T]$ is the space of continuous functions $x : [0, T] \rightarrow \mathbb{R}$. The Radon-Nikodym derivative of a probability measure μ with respect to a probability measure ν is denoted by $\frac{d\mu}{d\nu}$.

II. PROBLEM FORMULATION

Inspired by the vessel misguidance [1], we formulate a stochastic zero-sum game modeling the competition between an attacker, who tries to covertly misguide the vehicle to an unsafe region, versus a detector, who tries to detect the existence of the attack based on the observed trajectory of the vehicle. For simplicity, we model the trajectory of the vehicle (deviation from the nominal trajectory) as a continuous-time, scalar-valued Ito process x_t over the time interval $0 \leq t \leq T$ defined by the following stochastic differential equation:

$$dx(t) = \theta(t)dt + dw(t), x(0) = 0. \quad (1)$$

Here, $w(t)$ is the standard Brownian motion in the underlying probability space $(\Omega, \mathcal{F}, \mu)$. For each $0 \leq t \leq T$, we denote by $\mathcal{F}(t) \subseteq \mathcal{F}$ the filtration of the process $w(t)$. We call the drift term $\theta : [0, T] \rightarrow \mathbb{R}$ the attack signal, which is chosen by the attacker. The attack signal θ is assumed to be a Borel measurable function such that $\int_0^T |\theta(t)|dt < \infty$ to guarantee the existence of the strong solution to (1).

As shown in Fig. 1, we consider the terminal condition such that $x(T) > Td$ unsafe, where $d > 0$ is a given constant. When there is no attack (i.e., $\theta(t) = 0, \forall t \in [0, T]$), we have $x(T) \sim \mathcal{N}(0, T)$. Therefore, the probability of the terminal state being unsafe is $\Phi(-\sqrt{T}d)$. This probability can be altered by injecting a non-zero attack signal.

In this letter (except in Section V-A where we consider feedback policies), we restrict the attacker's strategy to the class of open-loop policies. That is, the attacker must fix an attack signal $\theta : [0, T] \rightarrow \mathbb{R}$ *a priori* without observing the state $x(t)$. We also restrict the attacker to the space of pure strategies, i.e., the action is not randomized. With this setup, we have $x(T) \sim \mathcal{N}(m, T)$ with $m = \int_0^T \theta(t)dt$. Consequently, the probability of the terminal state being unsafe is $\Phi(\frac{m}{\sqrt{T}} - \sqrt{T}d)$. In the sequel, we call $\gamma(\theta) := \Phi(\frac{m}{\sqrt{T}} - \sqrt{T}d)$ the attack success rate. We assume that, whenever the attack is applied, the attacker must ensure that the attack success rate is beyond a given threshold c , i.e., $\gamma(\theta) \geq c$. Notice that requiring $\gamma(\theta) \geq c$ is equivalent to imposing a constraint

$$\int_0^T \theta(t)dt \geq \sqrt{T}\Phi^{-1}(c) + Td \quad (2)$$

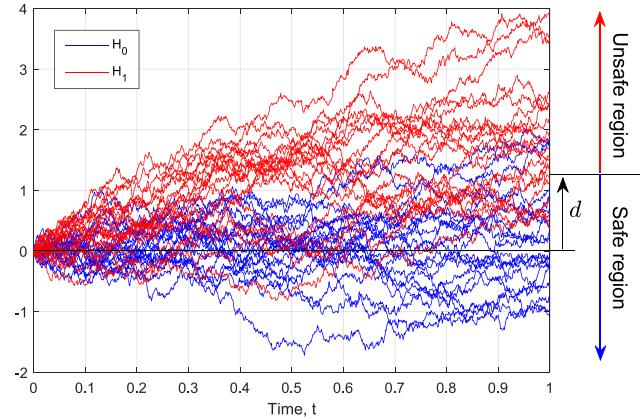


Fig. 1. Sample paths of (1) with $\theta(t) = 0$ (blue) and sample paths (1) with $\theta(t) = 2$ (red). We assume $T = 1$ in this plot.

on the attack signal θ . In the sequel, we assume $c > \frac{1}{2}$ and $d > 0$, which ensures that the quantity (2) is positive.

From the detector's viewpoint, it is not known in advance if the vehicle's operation is nominal or under attack. Therefore, the detector's task is to determine if the observed trajectory x is generated by (1) with $\theta = 0$, or if a non-zero attack signal θ is injected. The former scenario is the null hypothesis H_0 , whereas the latter is the alternative H_1 :

$$H_0 : \theta(t) = 0 \quad \forall t \in [0, T] \quad (3)$$

$$H_1 : \int_0^T \theta(t)dt \geq \sqrt{T}\Phi^{-1}(c) + Td. \quad (4)$$

Notice that this is a composite hypothesis testing problem since, while the null hypothesis H_0 is simple, the alternative H_1 contains a family of functions θ .

The role of the detector is to design a hypothesis testing algorithm $\phi : C[0, T] \rightarrow \{0, 1\}$ such that

$$\phi(x) = \begin{cases} 0 & \text{"No alarm"} \\ 1 & \text{"Alarm".} \end{cases} \quad (5)$$

The detector's decision (5) is made *a posteriori* after observing the entire state trajectory $x(t), 0 \leq t \leq T$.

The quality of a testing algorithm ϕ is measured in terms of the probability $\alpha(\phi)$ of a false alarm (also known as Type I error) and the probability $\beta(\theta, \phi)$ of a detection failure (also known as Type II error):

$$\alpha(\phi) := \Pr\{\phi(x) = 1 \mid H_0 \text{ is true}\} \quad (6)$$

$$\beta(\theta, \phi) := \Pr\{\phi(x) = 0 \mid H_1 \text{ is true}\}. \quad (7)$$

We say that a testing algorithm ϕ is admissible if $\alpha(\phi) \leq \epsilon$ for some given constant $\epsilon \in (0, \frac{1}{2})$.

In this letter, we model the interaction of the detector and the attacker as a zero-sum game. Since the false alarm rate $\alpha(\phi)$ does not depend on the attacker's policy θ , it is convenient to use $\alpha(\phi)$ as a constraint on the detector's policy. Similarly, since the attack success rate $\gamma(\theta)$ does not depend on the detector's policy ϕ , it is convenient to use $\gamma(\theta) \geq c$ as a constraint on the attacker's policy. In contrast, the detection failure rate $\beta(\theta, \phi)$ depends on both parties' policies which is minimized by the detector and maximized by the attacker. Therefore, we formulate a mini-max game:

$$p^* = \min_{\phi : \alpha(\phi) \leq \epsilon} \max_{\theta : \gamma(\theta) \geq c} \beta(\theta, \phi) \quad (8)$$

and its dual:

$$d^* = \max_{\theta : \gamma(\theta) \geq c} \min_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta, \phi). \quad (9)$$

By the weak duality, $p^* \geq d^*$ holds trivially. In this letter, we will provide a unique pair of policies (θ^*, ϕ^*) that constitutes a saddle point of the game, satisfying

$$\beta(\theta, \phi^*) \leq \beta(\theta^*, \phi^*) \leq \beta(\theta^*, \phi) \quad (10)$$

for all ϕ with $\alpha(\phi) \leq \epsilon$ and for all θ with $\gamma(\theta) \geq c$. Consequently, the strong duality $p^* = d^*$ will be established, and the value of the game $\beta(\theta^*, \phi^*)$ will be computed.

III. PRELIMINARIES

This section summarizes the mathematical ingredients needed to derive the main result.

A. Girsanov's Theorem

Given two random processes x and w related by (1), we have already defined the probability measure μ as the one in which w is the standard Brownian motion. In μ , x is not the Brownian motion. However, Girsanov's theorem [8, Th. 8.6.3] [9, Th. 6.3] states that there exists an alternative measure μ_θ in which x is the standard Brownian motion. Specifically, for each sample path x , the likelihood ratio $\frac{d\mu}{d\mu_\theta}(x)$ is given by

$$\frac{d\mu}{d\mu_\theta}(x) = \exp \left\{ \int_0^T \theta(t) dw(t) + \frac{1}{2} \int_0^T \theta^2(t) dt \right\} \quad (11a)$$

$$= \exp \left\{ \int_0^T \theta(t) dx(t) - \frac{1}{2} \int_0^T \theta^2(t) dt \right\}. \quad (11b)$$

That is, observing a particular sample path x as an outcome of (1) (this occurs with probability $\alpha(\mu(x))$) is $\frac{d\mu}{d\mu_\theta}(x)$ times more likely than observing the same sample path x as a realization of the standard Brownian motion (this occurs with probability $\alpha(\mu_\theta(x))$).

The false alarm rate $\alpha(\phi)$ is defined as the probability of $\phi(x) = 1$ when there is no attack, i.e., when x is the standard Brownian motion. Since x is the standard Brownian motion in μ_θ , this quantity can be expressed as $\alpha(\phi) = \mathbb{E}^{\mu_\theta}[\phi(x)]$. Despite the appearance of θ on the right-hand side, the false alarm rate does not depend on θ . In contrast, the detection failure rate $\beta(\theta, \phi) = 1 - \mathbb{E}^\mu[\phi(x)]$ depends on θ , as the distribution of x depends on θ under the measure μ in which w is the standard Brownian motion.

B. Neyman-Pearson Lemma

Consider the binary hypothesis testing problem in which the simple null hypothesis (3) is to be discriminated from the simple alternative (4) with a fixed θ satisfying $\gamma(\theta) \geq c$.

Lemma 1: The testing algorithm $\phi : C[0, T] \rightarrow \{0, 1\}$ that minimizes $\beta(\theta, \phi)$ subject to the constraint $\alpha(\phi) \leq \epsilon$ is given by

$$\phi(x) = \begin{cases} 0 & \text{if } z_\theta(x, T) \leq \lambda^* \\ 1 & \text{if } z_\theta(x, T) > \lambda^* \end{cases} \quad \mu\text{-almost surely} \quad (12)$$

where

$$z_\theta(x, t) = \exp \left\{ \int_0^t \theta(s) dx(s) - \frac{1}{2} \int_0^t \theta^2(s) ds \right\} \quad (13)$$

and $\lambda^* > 0$ is a constant satisfying $\alpha(\phi) = \epsilon$.

Proof: We accept the Neyman-Pearson lemma [7], which states that the optimal hypothesis test to discriminate a null hypothesis $x \sim \mu_\theta(x)$ from an alternative $x \sim \mu(x)$ is in general given by a randomized policy of the form

$$\phi(x) = 1_{\{z_\theta(x, T) > \lambda^*\}} + b \cdot 1_{\{z_\theta(x, T) = \lambda^*\}} \quad (14)$$

μ -almost surely, where $\lambda^* = \inf\{\lambda \geq 0 : \alpha(\phi) \leq \epsilon\}$ and b is an appropriate binary random variable. To complete the proof based on this result, it is sufficient to show that λ^* in our setup attains $\alpha(\phi) = \epsilon$, and $\Pr\{z_\theta(x, T) = \lambda^*\} = 0$.

Notice that

$$\alpha(\phi) = \mathbb{E}^{\mu_\theta}[\phi(x)] = \mathbb{E}^{\mu_\theta}[1_{\{z_\theta(x, T) > \lambda\}}]. \quad (15)$$

Since x is the standard Brownian motion in μ_θ , $\alpha(\phi)$ can be written in term of the standard Brownian motion x as

$$\begin{aligned} \alpha(\phi) &= \Pr \left\{ \exp \left\{ \int_0^T \theta(t) dx(t) - \frac{1}{2} \int_0^T \theta^2(t) dt \right\} > \lambda \right\} \\ &= \Pr \left\{ \int_0^T \theta(t) dx(t) > \frac{1}{2} \int_0^T \theta^2(t) dt + \log \lambda \right\}. \end{aligned} \quad (16)$$

Since the random variable $X := \int_0^T \theta(t) dx(t)$ has a continuous cumulative distribution function, $\alpha(\phi)$ is a continuous, non-increasing function of λ such that $\alpha(\phi) \rightarrow 1$ as $\lambda \rightarrow 0$ and $\alpha(\phi) \rightarrow 0$ as $\lambda \rightarrow +\infty$. Therefore, for each $\epsilon \in (0, \frac{1}{2})$, there exists $\lambda^* > 0$, which is the smallest constant satisfying $\alpha(\phi) \leq \epsilon$. Using the fact that X has a continuous cumulative distribution function, we also have $\Pr\{z_\theta(x, T) = \lambda^*\} = \Pr\{X = \frac{1}{2} \int_0^T \theta^2(t) dt + \log \lambda^*\} = 0$. ■

For a fixed θ such that $\gamma(\theta) \geq c$, let ϕ be the optimal test given by Lemma 1. Then, the detection failure rate can be written as

$$\begin{aligned} \beta(\theta, \phi) &= 1 - \mathbb{E}^\mu[\phi(x)] \\ &= 1 - \mathbb{E}^{\mu_\theta}[\phi(x) z_\theta(x, T)] \\ &= 1 - \lambda^* \mathbb{E}^{\mu_\theta}[\phi(x)] - \mathbb{E}^{\mu_\theta}[\phi(x)(z_\theta(x, T) - \lambda^*)] \\ &= 1 - \lambda^* \epsilon - \mathbb{E}^{\mu_\theta}[z_\theta(x, T) - \lambda^*]^+ \end{aligned} \quad (17)$$

In the last step, we used the fact that $\mathbb{E}^{\mu_\theta}[\phi(x)] = \alpha(\phi) = \epsilon$ (Lemma 1) and (12).

IV. MAIN RESULT

The main result of this letter is the following:

Theorem 1: The following pair of policies form a saddle point of the zero-sum game (8) and (9):

$$\theta^*(t) = \bar{\theta} := \frac{1}{\sqrt{T}} \Phi^{-1}(c) + d \quad \forall t \in [0, T] \quad (18)$$

$$\phi^*(x) = \begin{cases} 0 & \text{if } x(T) \leq \sqrt{T} \Phi^{-1}(1 - \epsilon) \\ 1 & \text{if } x(T) > \sqrt{T} \Phi^{-1}(1 - \epsilon). \end{cases} \quad (19)$$

Moreover, the saddle point (θ^*, ϕ^*) is unique in the sense that if (θ', ϕ') is another saddle point, then $\theta^*(t) = \theta'(t)$ holds almost everywhere in $[0, T]$ and $\phi^*(x) = \phi'(x)$ holds μ -almost surely. Furthermore, the value of the game is

$$\beta(\theta^*, \phi^*) = \Phi(\Phi^{-1}(1 - \epsilon) - \Phi^{-1}(c) - \sqrt{T}d). \quad (20)$$

Specifically, if $1 - \epsilon = c$, then $\beta(\theta^*, \phi^*) = \Phi(-\sqrt{T}d)$.

Remark 1: Theorem 1 states that the max-min policy (the most covert attack) is a constant bias injection $\theta(t) = \bar{\theta}$, where the constant $\bar{\theta}$ is chosen to be the smallest value satisfying

$\gamma(\theta) \geq c$. Conversely, the minimax policy $\phi^*(x)$ (i.e., the most powerful hypothesis test) only examines the final value $x(T)$ of the observed sample path x . As we will see below, $\phi^*(x)$ can be viewed as the Neyman-Pearson type binary hypothesis testing algorithm that discriminates $H_0 : dx_t = dw_t$ from $H_1 : dx_t = \bar{\theta}dt + dw_t$.

Remark 2: Suppose both (θ^*, ϕ^*) and (θ', ϕ') are saddle points of the game. Then, Theorem 1 states that they can only differ on a set with measure zero. By the interchangeability of saddle points of two-person zero-sum games [10], it also follows that (θ^*, ϕ^*) , (θ^*, ϕ') , (θ', ϕ^*) , and (θ', ϕ') are all saddle points, and they attain the same value.

To prove Theorem 1, notice that for any fixed θ^* such that $\gamma(\theta^*) \geq c$ and ϕ^* such that $\alpha(\phi^*) \leq \epsilon$, we have

$$\inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta^*, \phi) \leq \sup_{\theta : \gamma(\theta) \geq c} \inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta, \phi) \quad (21a)$$

$$\leq \inf_{\phi : \alpha(\phi) \leq \epsilon} \sup_{\theta : \gamma(\theta) \geq c} \beta(\theta, \phi) \quad (21b)$$

$$\leq \sup_{\theta : \gamma(\theta) \geq c} \beta(\theta, \phi^*) \quad (21c)$$

where (21b) follows from the max-min inequality. Hence, if the pair (θ^*, ϕ^*) satisfies the saddle point condition

$$\sup_{\theta : \gamma(\theta) \geq c} \beta(\theta, \phi^*) = \beta(\theta^*, \phi^*) = \inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta^*, \phi) \quad (22)$$

then the chain of inequalities (21) holds with equality and the strong duality $p^* = d^* = \beta(\theta^*, \phi^*)$ is implied.

Therefore, in Section IV-A below, we show that the pair (θ^*, ϕ^*) given by (18) and (19) indeed satisfies the saddle point condition (22). However, such an argument is insufficient to prove that (θ^*, ϕ^*) is the unique saddle point. To establish the uniqueness result, notice that the first inequality (21a) implies that if (θ', ϕ') is a saddle point, then θ' must be the max-min solution that attains

$$\sup_{\theta : \gamma(\theta) \geq c} \inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta, \phi) = \inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta', \phi). \quad (23)$$

and that ϕ' is the best response to θ' . In Section IV-B below, we show that (23) is attained uniquely by $\theta' = \theta^*$, and that ϕ^* is the unique best response to θ^* . This will establish the uniqueness of the saddle point (θ^*, ϕ^*) .

A. Saddle Point Condition

We prove (22) by showing the first equality (optimality of θ^*) and the second equality (optimality of ϕ^*) separately.

1) Optimality of θ^* : We first prove that $\beta(\theta, \phi^*) \leq \beta(\theta^*, \phi^*)$ holds for all θ such that $\gamma(\theta) \geq c$. Notice that the function $\phi^*(x)$ in (19) only depends on the terminal state $x(T)$. Moreover, under any admissible attack strategy, we have $x(T) \sim \mathcal{N}(m, T)$, where $m = \int_0^T \theta(t)dt \geq \sqrt{T}\Phi^{-1}(c) + Td$. Therefore, to maximize the detection failure rate, it is optimal for the attacker to choose a strategy that attains the smallest admissible value of m . Hence, any function $\theta : [0, T] \rightarrow \mathbb{R}$ such that $\int_0^T \theta(t)dt = \sqrt{T}\Phi^{-1}(c) + Td$ is a best response to ϕ^* . Since such a class of functions contains θ^* , we have $\beta(\theta, \phi^*) \leq \beta(\theta^*, \phi^*)$.

2) Optimality of ϕ^* : We next prove that $\beta(\theta^*, \phi^*) \leq \beta(\theta^*, \phi)$ holds for all ϕ such that $\alpha(\phi) \leq \epsilon$. To this end, let

the attacker's policy be fixed to θ^* in (18). Then, by Lemma 1, the optimal test ϕ is given by

$$\phi(x) = \begin{cases} 0 & \text{if } z_{\theta^*}(x, T) \leq \lambda^* \\ 1 & \text{if } z_{\theta^*}(x, T) > \lambda^* \end{cases} \quad (24)$$

where

$$z_{\theta^*}(x, T) = \exp \left\{ \int_0^T \bar{\theta} dx(t) - \frac{1}{2} \int_0^T \bar{\theta}^2 dt \right\} \quad (25a)$$

$$= \exp \left\{ \bar{\theta} x(T) - \frac{T}{2} \bar{\theta}^2 \right\}. \quad (25b)$$

Hence, it is sufficient to show that (24) is equivalent to (19).

Since $x(T) \sim \mathcal{N}(0, T)$ under μ_θ ,

$$\alpha(\phi) = \mathbb{E}^{\mu_\theta} [\phi(x)] \quad (26a)$$

$$= \mu_\theta \left(\left\{ \exp \left(\bar{\theta} x(T) - \frac{T}{2} \bar{\theta}^2 \right) > \lambda^* \right\} \right) \quad (26b)$$

$$= 1 - \Phi \left(\frac{\sqrt{T} \bar{\theta}}{2} + \frac{\log \lambda^*}{\sqrt{T} \bar{\theta}} \right). \quad (26c)$$

Solving $\alpha(\phi) = \epsilon$, we obtain

$$\lambda^* = \exp \left(\sqrt{T} \bar{\theta} \Phi^{-1}(1 - \epsilon) - \frac{T}{2} \bar{\theta}^2 \right). \quad (27)$$

Substituting (27) into (24), we obtain (19).

B. Uniqueness of the Saddle Point

We now solve the max-min problem on the left-hand side of (23). Let the function $\theta : [0, T] \rightarrow \mathbb{R}$ be fixed. Then, according to the Neyman-Pearson lemma, the best response is a threshold-based policy of the form:

$$\phi(x) = \begin{cases} 0 & \text{if } z_\theta(x, T) \leq \lambda^* \\ 1 & \text{if } z_\theta(x, T) > \lambda^* \end{cases} \quad (28)$$

where

$$z_\theta(x, t) = \exp \left\{ \int_0^t \theta(s) dx(s) - \frac{1}{2} \int_0^t \theta^2(s) ds \right\} \quad (29)$$

and $\lambda^* > 0$ is a constant that satisfies $\alpha(\theta) = \epsilon$. Assuming the best response (28) to the attack signal θ , we obtain from (17) that

$$\begin{aligned} & \sup_{\theta : \gamma(\theta) \geq c} \inf_{\phi : \alpha(\phi) \leq \epsilon} \beta(\theta, \phi) \\ &= \sup_{\theta : \gamma(\theta) \geq c} 1 - \lambda^* \epsilon - \mathbb{E}^{\mu_\theta} [z_\theta(x, T) - \lambda^*]^+. \end{aligned} \quad (30)$$

We will show that this supremum is attained by a constant function $\theta^*(t) = \bar{\theta}$ given by (18), and that any θ that attains the supremum must coincide with (18) almost everywhere in $t \in [0, T]$. Recall that the last term in (30) means

$$\begin{aligned} & \mathbb{E}^{\mu_\theta} [z_\theta(x, T) - \lambda^*]^+ \\ &= \mathbb{E}^{\mu_\theta} \left[\exp \left\{ \int_0^T \theta(t) dx(t) - \frac{1}{2} \int_0^T \theta^2(t) dt \right\} - \lambda^* \right]^+ \end{aligned} \quad (31)$$

Since $x(t)$ is the standard Brownian motion in μ_θ , and since $w(t)$ is the standard Brownian motion in μ , the quantity (31) can also be written as

$$\mathbb{E}^\mu \left[\exp \left\{ \int_0^T \theta(t) dw(t) - \frac{1}{2} \int_0^T \theta^2(t) dt \right\} - \lambda^* \right]^+. \quad (32)$$

Introducing

$$\zeta_\theta(t) := \exp \left\{ \int_0^t \theta(s) dw(s) - \frac{1}{2} \int_0^t \theta^2(s) ds \right\}, \quad (33)$$

Eq. (32) can also be written as $\mathbb{E}^\mu[\zeta_\theta(T) - \lambda^*]^+$. Therefore, it is left to prove the following lemma:

Lemma 2: Let θ^* be given by (18). Then,

$$\inf_{\theta : \gamma(\theta) \geq c} \mathbb{E}^\mu[\zeta_\theta(T) - \lambda^*]^+ = \mathbb{E}^\mu[\zeta_{\theta^*}(T) - \lambda^*]^+. \quad (34)$$

Moreover, if θ' attains the infimum on the left-hand side, then $\theta'(t) = \theta^*(t)$ almost everywhere in $t \in [0, T]$.

Proof: The proof strategy is inspired by [7, Sec. 5], which is further attributed to [11].

Let $f : \mathbb{R} \rightarrow [0, \infty)$ be a convex function satisfying the linear growth condition. We will show that

$$\mathbb{E}^\mu f(\zeta_{\theta^*}(T)) \leq \mathbb{E}^\mu f(\zeta_\theta(T)) \quad (35)$$

for all θ such that $\gamma(\theta) \geq c$. The claim (34) follows from (35) by choosing $f(z) = [z - \lambda^*]^+$.

For each θ such that $\gamma(\theta) \geq c$, observe that $\Lambda(t) := \int_0^t \theta^2(\tau) / \bar{\theta}^2 d\tau$ is a non-decreasing function. Moreover, consider the inner product $\langle v_1, v_2 \rangle := \int_0^T v_1(t) v_2(t) dt$ of $v_1(t) = \theta(t)$ and $v_2(t) = 1$. Since $\|v_1\|^2 = \int_0^T \theta^2(t) dt$, $\|v_2\|^2 = T$, and $\gamma(\theta) \geq c$ implies $\langle v_1, v_2 \rangle = \int_0^T \theta(t) dt \geq T\bar{\theta}$, it follows from the Cauchy-Schwarz inequality that $\int_0^T \theta^2(t) dt \geq T\bar{\theta}^2$. The equality holds if and only if $\theta(t) = \bar{\theta}$ almost everywhere. Hence, we have $\Lambda(T) \geq T$. Therefore, if we define a right inverse Λ^{-1} of Λ by $\Lambda^{-1}(s) := \inf\{t : \Lambda(t) > s\}$, we have $\Lambda^{-1}(T) \leq T$. Now, notice that a time-changed process

$$\hat{w}(s) := \int_0^{\Lambda^{-1}(s)} \frac{\theta(s)}{\bar{\theta}} dw(s), \quad 0 \leq s \leq T \quad (36)$$

is a martingale of the filtration $\hat{\mathcal{F}}(s) := \mathcal{F}(\Lambda^{-1}(s))$, $0 \leq s \leq T$ such that

$$\mathbb{E}\hat{w}^2(s) = \int_0^{\Lambda^{-1}(s)} \frac{\theta^2(s)}{\bar{\theta}^2} ds = \Lambda(\Lambda^{-1}(s)) = s. \quad (37)$$

Therefore, $\hat{w}(s)$ is a Brownian motion with respect to $\hat{\mathcal{F}}(s)$. Moreover, considering the time change $\tau = \Lambda^{-1}(s)$,

$$1 + \int_0^s \zeta_\theta(\Lambda^{-1}(\tau)) \bar{\theta} d\hat{w}(\tau) \quad (38a)$$

$$= 1 + \int_0^{\Lambda^{-1}(s)} \zeta_\theta(\tau) \bar{\theta} d\hat{w}(\Lambda(\tau)) \quad (38b)$$

$$= 1 + \int_0^{\Lambda^{-1}(s)} \zeta_\theta(\tau) \bar{\theta} \frac{\theta(\tau)}{\bar{\theta}} dw(\tau) \quad (38c)$$

$$= 1 + \int_0^{\Lambda^{-1}(s)} \zeta_\theta(\tau) \theta(\tau) dw(\tau) \quad (38d)$$

$$= \zeta_\theta(\Lambda^{-1}(s)). \quad (38e)$$

In the last step, we used the fact that $\zeta_\theta(t)$ as defined in (33) satisfies the stochastic differential equation $d\zeta_\theta(t) = \zeta_\theta(t) \theta(t) dw(t)$ (see, e.g., [8, Exercise 4.4]). Since

$$\zeta_{\theta^*}(t) = 1 + \int_0^t \zeta_{\theta^*}(\tau) \bar{\theta} dw(\tau), \quad (39)$$

comparing (38) and (39), we conclude that processes $\zeta_{\theta}(\cdot)$ and $\zeta_\theta(\Lambda^{-1}(\cdot))$ have the same distribution. From this observation, and from the optional sampling theorem, we have

$$\mathbb{E}^\mu f(\zeta_{\theta^*}(t)) = \mathbb{E}^\mu f(\zeta_\theta(\Lambda^{-1}(t))) \quad (40)$$

for $0 \leq t \leq T$. Also, since $\Lambda^{-1}(T) \leq T$, and since $f(\zeta_\theta(\cdot))$ is a submartingale (a consequence of Jensen's inequality), we obtain

$$\mathbb{E}^\mu f(\zeta_\theta(\Lambda^{-1}(T))) \leq \mathbb{E}^\mu f(\zeta_\theta(T)). \quad (41)$$

From (40) and (41), we obtain (35). \blacksquare

C. Value of the Game

The result (20) follows directly as follows:

$$\beta(\theta^*, \phi^*) = \mathbb{E}^\mu[1 - \phi(x)] \quad (42a)$$

$$= \mathbb{E}^\mu \left[1_{\{x(T) \leq \sqrt{T}\Phi^{-1}(1-\epsilon)\}} \right] \quad (42b)$$

$$= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\Phi^{-1}(1-\epsilon)} \exp \left\{ -\frac{1}{2} (x - \sqrt{T}\bar{\theta})^2 \right\} dx \quad (42c)$$

$$= \Phi(\Phi^{-1}(1-\epsilon) - \sqrt{T}\bar{\theta}) \quad (42d)$$

$$= \Phi(\Phi^{-1}(1-\epsilon) - \Phi^{-1}(c) - \sqrt{T}\bar{\theta}). \quad (42e)$$

We used the fact that $\frac{1}{\sqrt{T}}x(T) \sim \mathcal{N}(\sqrt{T}\bar{\theta}, 1)$ in step (42c).

V. DISCUSSION

A. Feedback Information Structure

The results so far are restricted to games (8) and (9) in which the attacker must choose the attack signal $\theta(t)$ in an open-loop manner. We now consider a modified setup in which the attacker is allowed to choose a state-dependent attack signal. Specifically, in (1), suppose that $\theta(t)$ is an $\mathcal{F}(t)$ -adapted function satisfying $\Pr\{\int_0^T |\theta(t)| dt \leq \infty\} = 1$. Such a class of functions includes feedback policies and gives an advantage to the attacker. We keep the strategy space for the detector the same. The detector's policy is a hypothesis testing algorithm $\phi : C[0, T] \rightarrow \{0, 1\}$.

We now demonstrate that the pair of policies (θ^*, ϕ^*) provided in Theorem 1 is no longer a saddle point in this modified information structure. To see this, it is sufficient to construct a feedback policy θ' such that $\beta(\theta', \phi^*) > \beta(\theta^*, \phi^*)$. To be concrete, assume $T = 1$, $d = 1.5$, and $c = 1 - \epsilon = 0.95$. In this case, the region $x(1) > 1.5$ is considered unsafe and ϕ^* triggers the alarm if $x(1) \geq \Phi^{-1}(1-\epsilon) \approx 1.645$ as shown in Fig. 2. Consider a feedback policy $\theta'(t) = \frac{b-x(t)}{1-t}$, where b is a constant satisfying $d < b < \Phi^{-1}(1-\epsilon)$. In this case, (1) becomes

$$dx(t) = \frac{b-x(t)}{1-t} dt + dw(t), \quad x(0) = 1. \quad (43)$$

The solution to (43) is known as the Brownian bridge, and satisfies $\lim_{t \rightarrow 1} x(t) = b$ μ -almost surely (Fig. 2). This means that the feedback policy θ' attains $\gamma(\theta') = 1$ and $\beta(\theta', \phi^*) = 1$. That is, the attacker wins most dramatically.

We are currently unaware of the saddle point strategies under the feedback information structure.

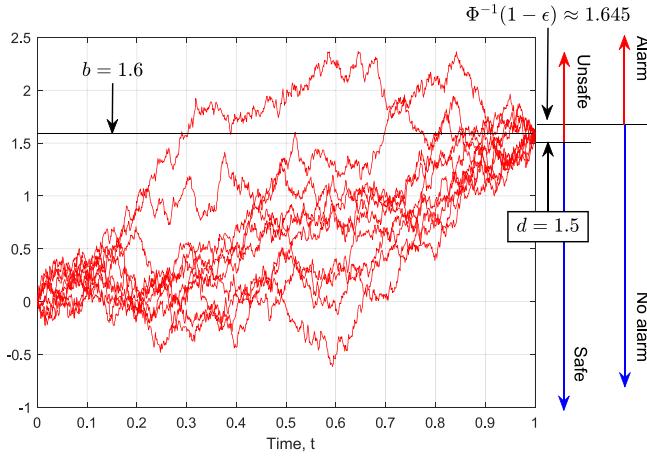


Fig. 2. Sample paths of the Brownian bridge.

B. The First and the Second Order Error Exponents

It is apparent from (20) that the detection failure rate $\beta(\theta^*, \phi^*)$ diminishes to zero as $T \rightarrow \infty$. While (20) already provides a compact formula, it is insightful to characterize it in terms of the first and the second-order error exponents (i.e., the coefficients of the T and \sqrt{T} terms in $\log \beta(\theta^*, \phi^*)$). To this end, let the attack signal $\theta^*(t) = \bar{\theta}$ be fixed. Introduce the relative entropy rate $\bar{D}(\mu \parallel \mu_{\theta^*}) := \limsup_{T \rightarrow \infty} \frac{1}{T} D(\mu \parallel \mu_{\theta^*})$ and the variance rate $\bar{V}(\mu \parallel \mu_{\theta^*}) := \limsup_{T \rightarrow \infty} \frac{1}{T} V(\mu \parallel \mu_{\theta^*})$, where

$$D(\mu \parallel \mu_{\theta^*}) := \mathbb{E}^\mu [\log z_{\theta^*}] \quad (44)$$

$$V(\mu \parallel \mu_{\theta^*}) := \mathbb{E}^\mu [(D(\mu \parallel \mu_{\theta^*}) - \log z_{\theta^*})^2] \quad (45)$$

with $\log z_{\theta^*} = \int_0^T \theta^*(t) dw(t) + \frac{1}{2} \int_0^T \theta^*(t)^2 dt$. Using $\theta^*(t) = \bar{\theta}$ and the fact that $w(t)$ is the standard Brownian motion in μ , we obtain

$$\bar{D}(\mu \parallel \mu_{\theta^*}) = \frac{1}{2} \bar{\theta}^2, \quad \bar{V}(\mu \parallel \mu_{\theta^*}) = \bar{\theta}^2. \quad (46)$$

Now, using (24) and (27), $\beta(\theta^*, \phi^*)$ can be written as

$$\begin{aligned} \beta(\theta^*, \phi^*) &= \mu(\{\log z_{\theta^*} \leq \log \lambda^*\}) \\ &= \mu\left(\left\{\bar{\theta}w(T) + \frac{T}{2}\bar{\theta}^2 \leq \sqrt{T}\bar{\theta}\Phi^{-1}(1 - \epsilon) - \frac{T}{2}\bar{\theta}^2\right\}\right) \\ &= \mu\left(\left\{\frac{w(T)}{\sqrt{T}} \leq \Phi^{-1}(1 - \epsilon) - \sqrt{T}\bar{\theta}\right\}\right) \\ &= \mu\left(\left\{\frac{w(T)}{\sqrt{T}} \geq \sqrt{T}\bar{\theta} + \Phi^{-1}(\epsilon)\right\}\right). \end{aligned} \quad (47)$$

Since $w(T)/\sqrt{T} \sim \mathcal{N}(0, 1)$ in μ , by Hoeffding's inequality for sub-Gaussian distributions, we have $\beta(\theta^*, \phi^*) \leq \exp\{-\frac{1}{2}(\sqrt{T}\bar{\theta} + \Phi^{-1}(\epsilon))^2\}$, or

$$-\log \beta(\theta^*, \phi^*) \geq \frac{1}{2}T\bar{\theta}^2 + \sqrt{T}\bar{\theta}\Phi^{-1}(\epsilon) + \text{const.} \quad (48)$$

Using (46), this inequality can be expressed as

$$\begin{aligned} -\log \beta(\theta^*, \phi^*) &\geq T\bar{D}(\mu \parallel \mu_{\theta^*}) + \sqrt{T}\sqrt{\bar{V}(\mu \parallel \mu_{\theta^*})}\Phi^{-1}(\epsilon) + \text{const.} \end{aligned} \quad (49)$$

The appearance of $\bar{D}(\mu \parallel \mu_{\theta^*})$ on the right-hand side of (49) is a reminiscent of Stein's lemma [12, Th. 11.8.3]. The inequality above shows the achievability of the relative entropy rate

as the first-order asymptotes. The second term provides a tighter estimate in the regime of finite T . Notably, (49) is consistent with the known characterization of the second-order asymptotes (e.g., [13], [14], [15]), despite the major difference between prior works on discrete-time hypothesis tests and our study on continuous-time counterparts. The appearance of higher-order terms in (49) implies that stealthiness measured by relative entropy alone [2], [3], [4] may not be accurate for moderate values of T .

VI. FUTURE WORK

While the scope of this letter is restricted to a simple system model (1), the approach we introduced in this letter can be generalized to high-dimensional and nonlinear system models. By substituting the function $\gamma(\theta)$ with other cost functions, the proposed framework accommodates a broader class of attack scenarios. Numerical approaches to compute the saddle point solutions (e.g., [16]) in these generalized settings are important research topics in the future. Saddle point solutions under the feedback information structure need further investigation. Finally, non-asymptotic (finite sample) analysis of the saddle point value for a broader class of games in view of the recent progress [13], [14], [15] in information theory literature will also be a fruitful research direction.

REFERENCES

- [1] J. Bhatti and T. E. Humphreys, "Hostile control of ships via false GPS signals: Demonstration and detection," *NAVIGATION J. Inst. Navig.*, vol. 64, no. 1, pp. 51–66, 2017.
- [2] C.-Z. Bai, F. Pasqualetti, and V. Gupta, "Data-injection attacks in stochastic control systems: Detectability and performance trade-offs," *Automatica*, vol. 82, pp. 251–260, Aug. 2017.
- [3] Z. Guo, D. Shi, K. H. Johansson, and L. Shi, "Worst-case stealthy innovation-based linear attack on remote state estimation," *Automatica*, vol. 89, pp. 117–124, Mar. 2018.
- [4] J. Shang, H. Yu, and T. Chen, "Worst-case stealthy innovation-based linear attacks on remote state estimation under Kullback–Leibler divergence," *IEEE Trans. Autom. Control*, vol. 67, no. 11, pp. 6082–6089, Nov. 2022.
- [5] J. Zhang and I. C. Paschalidis, "Statistical anomaly detection via composite hypothesis testing for Markov models," *IEEE Trans. Signal Process.*, vol. 66, no. 3, pp. 589–602, Feb. 2018.
- [6] S. Salimi, S. Dey, and A. Ahlén, "Sequential detection of deception attacks in networked control systems with watermarking," in *Proc. 18th Eur. Control Conf. (ECC)*, 2019, pp. 883–890.
- [7] J. Cvitanic and I. Karatzas, "Generalized Neyman-Pearson lemma via convex duality," *Bernoulli*, vol. 7 no. 1, pp. 79–97, 2001.
- [8] B. Øksendal, *Stochastic Differential Equations*. Berlin, Germany: Springer, 2003.
- [9] R. S. Liptser and A. N. Shiryaev, *Statistics of Random Processes I: General Theory*, vol. 394, New York, NY, USA: Springer, 2013.
- [10] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory*. Philadelphia, PA, USA: SIAM, 1998.
- [11] G.-L. Xu and S. E. Shreve, "A duality method for optimal consumption and investment under short-selling prohibition. II. Constant market coefficients," *Ann. Appl. Probab.*, vol. 2, pp. 314–328, May 1992.
- [12] T. Cover and J. Thomas, *Elements of Information Theory*. Hoboken, NJ, USA: Wiley, 1999.
- [13] K. Li, "Second-order asymptotics for quantum hypothesis testing," *Ann. Statist.*, vol. 42, pp. 171–189, Feb. 2014.
- [14] S. Watanabe, "Second-order optimal test in composite hypothesis testing," in *Proc. Int. Symp. Inf. Theory Appl. (ISITA)*, 2018, pp. 722–726.
- [15] V. Lungu and I. Kontoyiannis, "The optimal finite-sample error probability in asymmetric binary hypothesis testing," in *Proc. IEEE Int. Symp. Inf. Theory (ISIT)*, 2024, pp. 831–836.
- [16] A. Patil, M. O. Karabag, T. Tanaka, and U. Topcu, "Simulator-driven deceptive control via path integral approach," in *Proc. 62nd IEEE Conf. Decis. Control (CDC)*, 2023, pp. 271–277.