

Gaussian Estimation under Attack Uncertainty

Tara Javidi

Yonatan Kaspi

Himanshu Tyagi

Abstract—We consider the estimation of a standard Gaussian random variable under an observation attack where an adversary may add a zero mean Gaussian noise with variance in a bounded, closed interval to an otherwise noiseless observation. A straightforward approach would entail either ignoring the attack and simply using an optimal estimator under normal operation or taking the worst-case attack into account and using a minimax estimator that minimizes the cost under the worst-case attack. In contrast, we seek to characterize the optimal tradeoff between the MSE under normal operation and the MSE under the worst-case attack. Equivalently, we seek a minimax estimator for any fixed prior probability of attack. Our main result shows that a unique minimax estimator exists for every fixed probability of attack and is given by the Bayesian estimator for a least-favorable prior on the set of possible variances. Furthermore, the least-favorable prior is unique and has a finite support. While the minimax estimator is linear when the probability of attack is 0 or 1, our numerical results show that the minimax linear estimator is far from optimal for all other probabilities of attack and a simple nonlinear estimator does much better.

I. INTRODUCTION

Distributed control in the presence of an attacker who can manipulate various components of a control system has received a lot of attention in recent years. As a starting point to address this broad class of problems, we consider the estimation of a signal in the presence of *observation attacks* where an adversary can modify the distribution of the observed signal. A simple approach entails using a minimax estimator that minimizes the worst-case cost under all possible attacks. Such estimators can be constructed using, for instance, the robust estimation techniques in [16], [7]. Such robust estimator, while unavoidable when the system is under attack, might be too pessimistic when the attack is not certain. To remedy this shortcoming, we consider the problem of minimizing the operational cost J_o over all estimators that keep the cost under attack J_a bounded below an acceptable value J_{\max} . Under this formulation, the case when there is a certainty about the presence or the absence of an attack can be handled, respectively, by choosing J_{\max} to be sufficiently small or sufficiently large.

An equivalent formulation is to minimize the Lagrangian $J_o + \lambda J_a$ for a given value of λ , which is further equivalent to a Bayesian formulation where one seeks to minimize $(1 - p_a)J_o + p_a J_a$ for a given prior probability of attack $p_a = \lambda/(1 + \lambda)$. In this paper, for a specific Gaussian estimation problem, we present a structural result for the minimax estimator for any given p_a . Our numerical results

illustrate that there is a marked change in the structure of the minimax estimator as we move from $p_a \in \{0, 1\}$ to the case of attack uncertainty with $p_a \in (0, 1)$.

Specifically, we consider the model where a standard Gaussian signal X is observed under normal operation of the system. However, an attacker may add to X an independent, zero-mean Gaussian noise N of variance in $[\sigma_{\min}^2, \sigma_{\max}^2]$. For a given prior probability of attack p_a , we seek a minimax estimator for the *mean squared error* (MSE) cost where the normal operation cost $J_o(e)$ is given by $\mathbb{E}[(X - e(X))^2]$ and the cost under attack $J_a(e)$ is given by the supremum over $\theta \in [\sigma_{\min}^2, \sigma_{\max}^2]$ of $\mathbb{E}[(X - e(X + N_\theta))^2]$ for $N_\theta \sim \mathcal{N}(0, \theta)$. We show that a unique minimax estimator exists for every fixed p_a and is given by the Bayesian estimator for a least-favorable prior. Furthermore, the least-favorable prior is unique and has a finite support¹. For the case when $p_a = 0$ or 1, this unique minimax estimator is linear. However, our numerical results show that the minimax estimator for the restricted problem where the adversary can attack only with $\theta = \sigma_{\min}^2$ or $\theta = \sigma_{\max}^2$ outperforms the minimax linear estimator for all $p_a \in (0, 1)$.

The problem studied here is closely related to that of parameter estimation when the parameter space is restricted to a compact set (*cf.* [3], [2], [15]). Indeed, one approach for our problem can be to first detect if an attack is present, and if it is present, estimate the noise parameter $\hat{\theta}$ and use the optimal estimator of X for $\hat{\theta}$. An analysis of this estimator that separates detection and estimation is work in progress.

The rest of the paper is organized as follows. In the next section, we present a general saddle-point theorem that will be used to show the optimality of Bayesian estimators. Our main results on the structure of the minimax estimator for Gaussian signal under observation attack are given in Section III. The final section contains numerical results comparing the minimax linear estimator and the minimax estimator for the restricted problem where the adversary can only choose noise variance to be σ_{\min}^2 or σ_{\max}^2 .

II. SADDLE POINT THEOREMS

In this section, we prove a general saddle point theorem and apply it to the specific case of estimation under the MSE criterion to show the existence of a robust *minimum mean squared error* (MMSE) estimator.

While a general result claiming the optimality of Bayesian policies for the minimax risk problem was shown in [8], it is

Authors are with the Department of Electrical and Computer Engineering, University of California, San Diego, La Jolla, CA 92093, USA. Email: {tjavidi, ykaspi, htyagi}@ucsd.edu

¹Analogous results for the case when the adversary controls the mean of the noise are omitted due to lack of space.

not directly applicable to the specific case of robust MMSE estimation for two reasons: First, the quadratic loss function is unbounded, and second, it is unclear how to cast the robust MMSE problem as a minimax risk problem that satisfies the assumptions in [8]. Another plausible approach entails using the saddle point theorem in [17, Theorem 2.1]. However, the ‘‘regular point condition’’ required in that result does not hold for our problem. Various other saddle point theorems for estimation problems with restricted parameter spaces (*cf.* [10]) are not applicable either. In contrast, our saddle point theorem below is tailored to the cases where additional structure is available for the set of Bayesian policies².

A. General result

Let Ω be a compact metric space. Consider the minimax problem

$$\mathcal{L}(\mathcal{F}, \Omega) = \inf_{f \in \mathcal{F}} \sup_{\omega \in \Omega} \ell(f, \omega),$$

where $\ell : \mathcal{F} \times \Omega \rightarrow [0, \infty)$ is the cost function to be minimized. Further, let $\mathcal{P}(\Omega)$ denote the set of all probability measures on Ω , and let $L(f, \pi)$ denote the average cost

$$L(f, \pi) = \mathbb{E}_\pi[\ell(f, W)],$$

where the expectation is taken over the random variable W with distribution $\pi \in \mathcal{P}(\Omega)$. Then, the minimax problem $\mathcal{L}(\mathcal{F}, \Omega)$ can be reexpressed as its convexified version

$$\mathcal{L}(\mathcal{F}, \mathcal{P}) = \inf_{f \in \mathcal{F}} \sup_{\pi \in \mathcal{P}(\Omega)} L(f, \pi) = \mathcal{L}(\mathcal{F}, \Omega).$$

Recall the following definition.

Definition 1. A policy $f^* \in \mathcal{F}$ and a prior $\pi^* \in \mathcal{P}(\Omega)$ constitute a *saddle point* of $\mathcal{L}(\mathcal{F}, \mathcal{P})$ if for all f, π

$$L(f^*, \pi) \leq L(f^*, \pi^*) \leq L(f, \pi^*),$$

namely f^* is an optimal policy for the prior π^* and π^* is a worst prior when the policy is fixed to be f^* .

For a saddle point (f^*, π^*) , we have

$$\begin{aligned} \mathcal{L}(\mathcal{F}, \mathcal{P}) &= \min_{f \in \mathcal{F}} \max_{\pi \in \mathcal{P}(\Omega)} L(f, \pi) \\ &= \max_{\pi \in \mathcal{P}(\Omega)} \min_{f \in \mathcal{F}} L(f, \pi) \\ &= L(f^*, \pi^*). \end{aligned}$$

Therefore, π^* is a *least favorable prior* for $\mathcal{L}(\mathcal{F}, \Omega)$, and f^* is a *minimax policy* for it.

We show that a saddle point exists for $\mathcal{L}(\mathcal{F}, \mathcal{P})$ under the following assumption:

There exists a reflexive Banach space \mathcal{R} and its bounded, closed, and convex subset $\mathcal{F}_0 \subset \mathcal{F} \cap \mathcal{R}$ such that

- (A1) $\ell(f, \omega)$ is continuous ω for each fixed policy $f \in \mathcal{F}_0$;
- (A2) $\ell(f, \omega)$ is continuous and convex in $f \in \mathcal{F}_0$ with respect to \mathcal{R} , uniformly in $\omega \in \Omega$;

²For instance, under mild conditions, Bayesian policies for the MMSE problem lie in a Hilbert space.

- (A3) for each $\pi \in \mathcal{P}(\Omega)$ there is a policy $f_\pi \in \mathcal{F}_0$ that is optimal for π , *i.e.*, for each π there is a f_π satisfying

$$f_\pi \in \left\{ \arg \min_{f \in \mathcal{F}} L(f, \pi) \right\} \cap \mathcal{F}_0.$$

Theorem 1. *If the assumptions (A1)-(A3) hold, then there exists a saddle point (f^*, π^*) for $\mathcal{L}(\mathcal{F}, \mathcal{P})$.*

The proof is omitted due to lack of space.

Note that we do not require the class \mathcal{F} to satisfy any structure. We only make assumptions about the structure of \mathcal{F}_0 , a set containing optimal policies for all priors.

B. Robust Minimum Mean Squared Error Estimation

Let X, Y , and N_θ , $\theta \in \Theta$, be \mathbb{R}^n -valued random variables, and let Θ be a compact metric space. Consider the problem of estimating a random variable X by observing its noisy version

$$Y = X + N_\theta,$$

where the noise N_θ is independent of X and has a distribution P_θ , $\theta \in \Theta$. For brevity, we denote the expectation with respect to the measure P_θ on noise N_θ by $\mathbb{E}_\theta[\cdot]$, and with a slight abuse of notation, the expectation with respect to a random θ drawn according to a prior π by $\mathbb{E}_\pi[\cdot]$. For an estimator $e : \mathbb{R}^n \rightarrow \mathbb{R}^n$ and a parameter $\theta \in \Theta$, the cost function is given by

$$\ell(e, \theta) = \mathbb{E}_\theta[\|X - e(Y)\|_2^2].$$

We are interested in finding an optimal estimator $e^*(Y)$ that minimizes the worst-case cost $\sup_{\theta \in \Theta} \ell(e, \theta)$ over the class \mathcal{E} of estimators e satisfying

$$\sup_{\theta \in \Theta} \mathbb{E}_\theta[\|e(Y)\|_2^2] < \infty.$$

Recall that, for each prior π on Θ , the *unique* estimator $e_\pi^* \in \mathcal{E}$ which minimizes the cost

$$L(e, \pi) = \mathbb{E}_\pi[\ell(e, T)]$$

is given by

$$e_\pi^*(y) = \mathbb{E}_\pi[X | Y = y].$$

We make the following assumptions:

- (B1) The total variation distance $d_{TV}(P_\theta, P_{\theta'})$ goes to zero in the limit $\theta \rightarrow \theta'$;
- (B2) there exists a $\theta_0 \in \Theta$ and a constant C_0 such that for all Borel measurable sets \mathcal{A}

$$P_{\theta_0}(\mathcal{A}) \leq C_0 P_{\theta_0}(\mathcal{A});$$

- (B3) the optimal estimators $e_\pi^*(y) = \mathbb{E}_\pi[X | Y = y]$ satisfy

$$\sup_{\pi \in \mathcal{P}(\Theta)} \sup_{\theta \in \Theta} \mathbb{E}_\theta[\|e_\pi^*(Y)\|_2^2] < \infty.$$

We have the following corollary of Theorem 1. The proof is omitted due to lack of space.

Theorem 2. *If (B1)-(B3) hold, then there exists a saddle point (e^*, π^*) for the robust MMSE problem $\mathcal{L}(\mathcal{E}, \mathcal{P}(\Theta))$. Furthermore, the saddle point estimator e^* is unique.*

While the minimax estimator is shown to be unique in the proof above, the least favorable prior may not be unique. Specifically, it is possible that two least favorable priors π_1^* and π_2^* both lead to the same minimax estimator e^* . In the next section, we will rule out this possibility for the specific case of Gaussian noise.

III. GAUSSIAN ESTIMATION UNDER ATTACK UNCERTAINTY

We now move to the main subject of this paper, namely the estimation of scalar Gaussian signals in Gaussian noise under the MSE criterion when the noise parameters may be controlled by an adversary. We begin with a formal description of the setup, followed by a saddle point theorem which shows that there is a unique least favorable prior which leads to the minimax estimator. Finally, we show that the unique least favorable prior has a discrete support.

We restrict to an attack only on the variance of the noise. Analogous results for attack on the mean of the noise hold. Furthermore, the uniqueness of least favorable prior shown below holds in general for \mathbb{R}^n -valued observations and also under joint mean and variance uncertainty. We omit these general results due to the lack of space and illustrate all our observations in the context of an adversary who controls the variance of the noise.

A. Formulation and the minimax estimator

Consider the robust MMSE formulation of Section II-B with $n = 1$ for the case when X is a standard Gaussian random variable and the parametric family P_θ is given by

$$N_\theta = \begin{cases} 0, & \text{with prob. } 1 - p_a \\ \mathcal{N}(0, \theta), & \text{with prob. } p_a, \end{cases}$$

where $\mathcal{N}(0, \theta)$ denotes the zero mean Gaussian distribution with variance θ taking values in the compact set $\Theta = [\sigma_{\max}^2, \sigma_{\min}^2]$. Therefore, the goal is to obtain an estimator e^* that minimizes the worst-case cost

$$\sup_{\Theta} \mathbb{E}_\theta [\|X - e(Y)\|_2^2] = (1 - p_a) \mathbb{E} [\|X - e(X)\|_2^2] + p_a \sup_{\theta \in \Theta} \mathbb{E} [\|X - e(Y_\theta)\|_2^2], \quad (1)$$

where p_a is the given probability of attack and Y_θ denotes the random variable $X + N_\theta$. As discussed in Section I, the formulation here corresponds to minimizing the Lagrangian of the constrained optimization problem where one seeks to minimize the MSE cost under normal operation while ensuring a reasonable performance under attack.

When $p_a = 0$, the (trivial) linear estimator $e(y) = y$ is optimal. Interestingly, even for $p_a = 1$, the minimax estimator is linear³ and corresponds to the optimal estimator for noise variance σ_{\max}^2 . However, a linear estimator is far from optimal for any other choice of p_a ; see numerical results in the next section.

³The simple proof follows upon noting that σ_{\max}^2 is the worst variance for the optimal estimator for $\theta = \sigma_{\max}^2$.

As in the previous section, we consider the Bayesian relaxation of the problem and argue by Theorem 2 that there is a unique minimax estimator for this problem, leading to the following corollary.

Corollary 3. *There exists a saddle point (e^*, π^*) for the cost (1) for estimating Gaussian signals under attack uncertainty. Furthermore, the minimax estimator e^* is unique.*

Proof: It suffices to verify conditions (B1)-(B3). For (B1), note that the Kullback-Leibler divergence between P_{θ_1} and P_{θ_2} is given by

$$D(P_{\theta_1} \| P_{\theta_2}) = (1 - p_a) \log \frac{(1 - p_a) + c(\theta_1)p_a}{(1 - p_a) + c(\theta_2)p_a} + p_a \log \frac{c(\theta_1)}{c(\theta_2)} + p_a \theta_1^2 \left(\frac{\theta_1 - \theta_2}{2\theta_1\theta_2} \right),$$

where $c(\theta) = (2\pi\theta)^{-1/2}$. Then, (B1) follows from Pinsker's inequality (cf. [4]). Also, (B2) holds with $\theta_0 = \sigma_{\max}^2$ and $C_0 = \sigma_{\max}/\sigma_{\min}$. For (B3), note that for any prior π the optimal estimator is given by

$$e_\pi^*(y) = \mathbb{E}_\pi [X|Y = y] = \mathbb{E}_\pi [\mathbb{E}[X|\theta = S, Y = y] | Y = y] = y \mathbb{E}_\pi \left[\frac{1}{1 + S} | Y = y \right],$$

where the random parameter S takes the value 0 with probability $(1 - p_a)$, takes values in Θ according to the distribution $p_a\pi$, and $Y = X + N_S$. Consequently,

$$\mathbb{E}_\theta [e_\pi^*(Y_\theta)^2] \leq \theta \leq \sigma_{\max}^2,$$

which gives (B3). ■

B. Uniqueness of the least favorable prior

While we established the uniqueness of the minimax estimator in Corollary 3, there might be multiple priors that lead to it. The next result says that, in fact, there is a unique least favorable prior.

Lemma 4. *There is a unique least favorable prior for (1) and consequently, also a unique saddle point.*

Proof: In view of Corollary 3, it suffices to show that if $e_{\pi_1}^* = e_{\pi_2}^*$, then the priors π_1 and π_2 must be the same. We use a method that was introduced in [5] for a different purpose. Specifically, on reparametrizing by $\alpha = (1 + \theta)^{-1}$ which takes the value 1 with probability $1 - p_a$ and values on the interval $[(1 + \sigma_{\max}^2)^{-1}, (1 + \sigma_{\min}^2)^{-1}]$ according to the distribution $p_a\pi$, we get

$$e_{\pi_1}^* = \frac{\int \alpha^{\frac{3}{2}} e^{-\alpha y^2/2} \pi(d\alpha)}{\int \alpha^{\frac{1}{2}} e^{-\alpha y^2/2} \pi(d\alpha)}.$$

On changing the distribution of α to π' given by $(d\pi'/d\pi) = \sqrt{\alpha}$ and denoting $z = -y^2$, we get

$$e_{\pi_1}^*(y) = \frac{\int \alpha e^{z\alpha} \pi'(d\alpha)}{\int e^{z\alpha} \pi'(d\alpha)} = \frac{d}{dz} \log \left(\int e^{z\alpha} \pi'(d\alpha) \right) = \frac{d}{dz} \log \phi_{\pi'}(z),$$

where $\phi_{\pi'}(z)$ is the moment generating function of π' at z . Since $\phi_{\pi'}(0) = 1$, $e_{\pi_1}^* = e_{\pi_2}^*$ implies that $\phi_{\pi_1'}(z) = \phi_{\pi_2'}(z)$ for all $z \leq 0$, and since α is supported on a compact set, the analytic extension of $\phi_{\pi_1'}$ equals that of $\phi_{\pi_2'}$. In particular, π_1' and π_2' have the same characteristic functions. Therefore, $\pi_1' = \pi_2'$ (cf. [6]) and so, $\pi_1 = \pi_2$. ■

C. The least favorable prior has a finite support

We close this section with a structural result about the least favorable prior, namely that it is discrete and is supported on finitely many points. Similar structure has been observed in several parameter estimation problems when the parameter set is restricted to a compact set; see [15, Section 4.2] for a literature review. Our proof is based on the analytic extension of the cost function which was introduced by Smith in [14] to show that the capacity achieving distribution in an amplitude constrained Gaussian channel has a finite support. Specifically, we rely on the well-known result in parametric estimation (cf. [1, Chapter 5, Theorem 19]) which states that if the parameter set is compact and the Bayesian risk for a least favorable prior is analytic in the parameter, then either the least-favorable prior has a finite support or the corresponding Bayesian estimator has constant risk for all values of the parameter.

Definition 2. Let π be a (Borel) probability measure on a compact metric space Θ . The *support* of π is the set of all points $x \in \Theta$ such that $\pi(\mathcal{O}) > 0$ for all open sets \mathcal{O} containing x .

In preparation for our main result, we make the following observation about the support of a least favorable prior π^* .

Lemma 5. Given $\sigma_{\min} > 0$, for every $M > 0$ there exists $\sigma_{\max}^2 \geq \sigma_{\min}^2$ such that the support of a least favorable prior for the parameter set $[\sigma_{\min}^2, \sigma_{\max}^2]$ contains a point $\theta > M$.

Proof: Assume the contrary. Then, there exists $M > 0$ such that the support of least favorable priors for all intervals $[\sigma_{\min}^2, \sigma_{\max}^2]$ is contained in $[0, M]$. Note that for any prior π with support contained in $[0, M]$

$$\frac{y^2}{(1+M)^2} \leq e_{\pi}^*(y)^2 \leq y^2.$$

It follows that

$$\begin{aligned} L(e_{\pi^*}^*, \theta) &= \mathbb{E}_{\theta} [X - e_{\pi^*}^*(Y)]^2 \\ &= 1 + \mathbb{E}_{\theta} [e_{\pi^*}^*(Y)^2] - 2\mathbb{E}_{\theta} [Xe_{\pi^*}^*(Y)] \\ &\geq 1 + \mathbb{E}_{\theta} [e_{\pi^*}^*(Y)^2] - 2\sqrt{\mathbb{E}_{\theta} [e_{\pi^*}^*(Y)^2]} \\ &\geq 1 + \frac{1+\theta}{(1+M)^2} - 2\sqrt{1+\theta}, \end{aligned}$$

where the first inequality is by Cauchy-Schwartz inequality. Therefore, in the limit as σ_{\max} goes to infinity, the worst-case cost $\sup_{\theta \in \Theta} L(e_{\pi^*}^*, \theta)$ also goes to infinity. However, in view of Corollary 3 and Lemma 4, for the least favorable prior π^* , the estimator $e_{\pi^*}^*$ is the unique minimax estimator. But the estimator $e(y) = 0$ attains the cost $k+1$ and outperforms $e_{\pi^*}^*$ for sufficiently large σ_{\max} , which is a contradiction. ■

Our main result on the structure of the least favorable prior is the following.

Theorem 6. For every compact parameter set Θ , the least favorable prior has a finite support.

Proof: Note that the function $L(e_{\pi^*}^*, \theta)$ is analytic in θ in the right-half complex plane away from zero (cf. [9, Theorem 2.7.1]). Therefore, by [1, Chapter 5, Theorem 19] (see also [14]) either the support \mathcal{S} of the least favorable prior contains finitely many points or $L(e_{\pi^*}^*, \theta)$ must equal $\mathcal{L}(\mathcal{E}, \Theta)$ in the entire right-half complex plane away from zero. But if the latter holds then $e_{\pi^*}^*$ will be a minimax estimator for every parameter set Θ' that contains Θ and π^* will be the least favorable prior for every Θ' containing π^* . However, this contradicts Lemma 5. Thus, the set \mathcal{S} must contain only finitely many points. ■

IV. NUMERICAL RESULTS

In this final section, we continue with the setup of the previous section and present numerical results to illustrate the utility of estimators designed with the possibility of an observation attack in mind. Note that for $p_a = 0$ the optimal estimator $e(y) = y$ is linear. Also, since a minimax linear estimator exists for the robust MMSE problem with uncertain noise [16], even for $p_a = 1$ the minimax estimator is linear. However, when there is some uncertainty about the attack, i.e., $p_a \notin \{0, 1\}$, a linear estimator is no longer minimax. To illustrate this, we compare the performance of a minimax linear estimator with the minimax estimator for the restricted problem when the adversary can attack only with one of the two extreme variances σ_{\min}^2 or σ_{\max}^2 . We begin by describing the minimax estimators for the two cases.

A. Minimax linear estimator e_l

Let \mathcal{E}_l be the class of all linear estimators $e(y) = ay$. Consider the minimax problem $\mathcal{L}(\mathcal{E}_l, [\sigma_{\min}^2, \sigma_{\max}^2])$ for the cost function (1). For a given $\theta \in [\sigma_{\min}^2, \sigma_{\max}^2]$, the optimal linear estimator $e_{\theta}(y)$ is given by $e_{\theta}(y) = (1 + p_a\theta)^{-1}y$. Furthermore, the MSE cost for a parameter is linear in θ . Therefore, the minimax linear estimator e_l is given by $e_{\theta^*}(y)$ for $\theta^* = \sigma_{\max}^2$.

B. Minimax estimator e_r for the restricted problem

Consider a restricted problem where the choice of noise variances for the adversary are restricted to $\Theta_r = \{\sigma_{\min}^2, \sigma_{\max}^2\}$. For the minimax problem $\mathcal{L}(\mathcal{E}, \mathcal{P}(\Theta_r))$, a minimax estimator e_r exists and can be determined explicitly. Specifically, the following holds (cf. [11]).

Lemma 7. Given a probability of attack p_a , let $e_p(y)$ be the optimal estimator for the Bayesian case when the adversary chooses σ_{\min}^2 with probability p and σ_{\max}^2 with probability $1-p$. Then, one of the following must hold:

- 1) $e_{\sigma_{\min}^2}$ is minimax for the restricted problem;
- 2) $e_{\sigma_{\max}^2}$ is minimax for the restricted problem;
- 3) there exists a p^* such that $L(e_{p^*}, \sigma_{\min}^2) = L(e_{p^*}, \sigma_{\max}^2)$, and therefore, e_{p^*} is minimax for the restricted problem.

Note that $\mathcal{L}(\mathcal{E}, \Theta_r)$ constitutes a lower bound for the original minimax problem $\mathcal{L}(\mathcal{E}, [\sigma_{\min}^2, \sigma_{\max}^2])$. Furthermore, the performance of the minimax estimator e_r over $[\sigma_{\min}^2, \sigma_{\max}^2]$ yields an upper bound for cost in the original minimax problem.

C. Comparison of performances of e_l and e_r

In Figure 1, we compare the performance of the minimax restricted estimator e_r and the minimax linear estimator e_l for the original problem $\mathcal{L}(\mathcal{E}, [\sigma_{\min}^2, \sigma_{\max}^2])$. We plot the cost pair (J_o, J_a) for e_r and e_l for different values of p_a . As a lower bound, we plot the optimal cost region for the restricted problem. It can be seen that the linear estimator is far from optimal when there is an attack uncertainty. As an attempt to explain this disparity, we compare the estimators e_r and e_l in Figure 2. Heuristically, the nonlinear estimator e_r outperforms the minimax linear estimator due to its ability to “detect” an attack. In particular, note that for small p_a the estimator e_r detects normal operation for small y and mimics $e(y) = y$. On the other hand, for large p_a , it detects an attack if y is large and mimics $e_l(y)$.

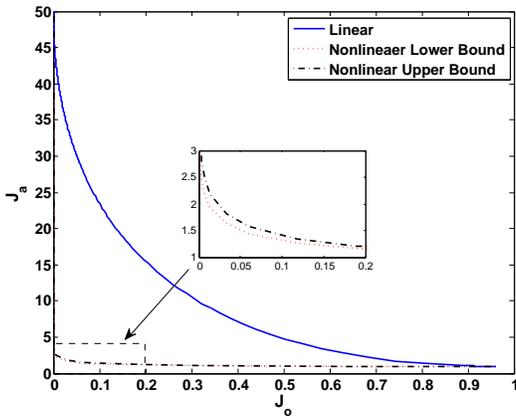


Fig. 1: Cost tradeoff curve for $\sigma_{\min}^2 = 5$ and $\sigma_{\max}^2 = 50$ comparing the minimax linear estimator and the minimax estimator for the restricted problem with only two possible attacks. The lower bound corresponds to the optimal performance for the restricted problem.

Another interesting observation from Figure 1 is that the estimator e_r is close to optimal. In fact, we observed no difference between the lower bound and the actual performance of e_r for the case when σ_{\min}^2 and σ_{\max}^2 are close to each other. It can also be observed (numerically) that the function $\ell(e_r, \theta)$ is monotone when σ_{\min} and σ_{\max} are close to each other, suggesting the minimaxity of e_r . This observation is along the lines of [3] where a similar result was proved for the problem of estimating Gaussian mean. Proving this numerical observation rigorously and a comparison of the minimax estimator with a simple “detect and estimate” approach is work in progress.

ACKNOWLEDGMENTS

The authors would like to thank Young-Han Kim for suggesting the lower bound in Figure 1.

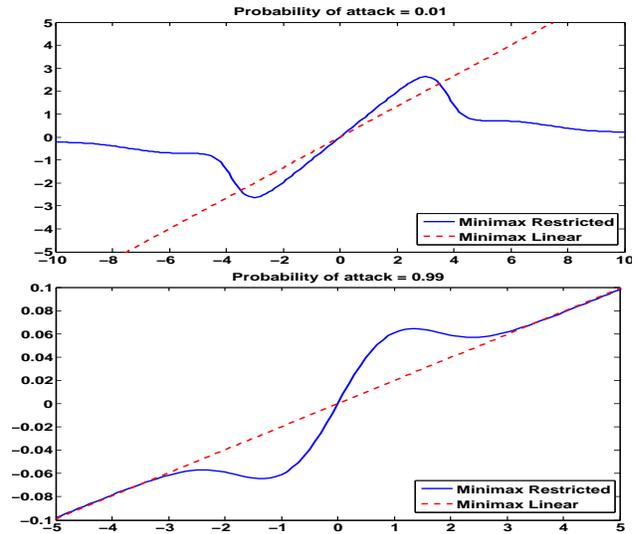


Fig. 2: The minimax linear estimator and the minimax estimator for the restricted problem with only two possible attacks with $\sigma_{\min}^2 = 5$ and $\sigma_{\max}^2 = 50$.

REFERENCES

- [1] J. O. Berger, *Statistical decision theory and Bayesian analysis*, 2nd ed., ser. Springer series in statistics. Springer, 1985.
- [2] P. J. Bickel, “Minimax estimation of the mean of a normal distribution when the parameter space is restricted,” *Annals of Statistics*, vol. 9, no. 6, pp. 1301–1309, 1981.
- [3] G. Casella and W. E. Strawderman, “Estimating a bounded normal mean,” *Annals of Statistics*, vol. 9, no. 4, pp. 870–878, July 1981.
- [4] I. Csiszár and J. Körner, *Information theory: Coding theorems for discrete memoryless channels. 2nd edition*. Cambridge University Press, 2011.
- [5] Y. Eldar and N. Merhav, “A competitive minimax approach to robust estimation of random parameters,” *Signal Processing, IEEE Transactions on*, vol. 52, no. 7, pp. 1931–1946, July 2004.
- [6] W. Feller, *An Introduction to Probability Theory and its Applications, Volume II. 2nd edition*. John Wiley & Sons Inc., UK, 1971.
- [7] S. Kassam and H. Poor, “Robust techniques for signal processing: A survey,” *Proceedings of the IEEE*, vol. 73, no. 3, pp. 433–481, March 1985.
- [8] L. Le Cam, *Asymptotic methods in statistical decision theory*, ser. Springer series in statistics. Springer, 1986.
- [9] E. L. Lehmann and J. P. Romano, *Testing statistical hypotheses*, 3rd ed., ser. Springer Texts in Statistics. Springer, 2005.
- [10] E. Marchand and W. E. Strawderman, “A unified minimax result for restricted parameter spaces,” *Bernoulli*, vol. 18, no. 2, pp. 635–643, 2012.
- [11] H. V. Poor, *An Introduction to Signal Detection and Estimation (2nd Ed.)*. Springer-Verlag New York, Inc., 1994.
- [12] H. Royden and P. M. Fitzpatrick, *Real analysis. 4th ed.* International Edition. New York, NY: Prentice Hall., 2010.
- [13] M. Sion, “On general minimax theorems,” *Pacific Journal of Mathematics*, vol. 8, no. 1, pp. 171–176, 1958.
- [14] J. G. Smith, “The information capacity of amplitude- and variance-constrained scalar gaussian channels,” *Information and Control*, vol. 18, no. 3, pp. 203 – 219, 1971.
- [15] C. van Eeden, *Restricted Parameter Space Estimation Problems: Admissibility and Minimality Properties*, ser. Lecture Notes in Statistics. Springer, 2006.
- [16] S. Verdú and H. Poor, “Minimax linear observers and regulators for stochastic systems with uncertain second order statistics,” *IEEE Trans. Inf. Theory*, vol. AC-29, no. 6, pp. 499–511, Jun 1984.
- [17] —, “On minimax robustness: A general approach and applications,” *IEEE Trans. Inf. Theory*, vol. 30, no. 2, pp. 328–340, Mar 1984.