

Security versus Privacy

Farhad Farokhi and Peyman Mohajerin Esfahani

Abstract—Linear queries can be submitted to a server containing private data. The server provides a response to the queries systematically corrupted using an additive noise to preserve the privacy of those whose data is stored on the server. The measure of privacy is inversely proportional to the trace of the Fisher information matrix. It is assumed that an adversary can inject a false bias to the responses. The measure of the security, capturing the ease of detecting the presence of the false data injection, is the sensitivity of the Kullback-Leiber divergence to the additive bias. An optimization problem for balancing privacy and security is proposed and subsequently solved. It is shown that the level of guaranteed privacy times the level of security equals a constant. Therefore, by increasing the level of privacy, the security guarantees can only be weakened and *vice versa*. Similar results are developed under the differential privacy framework.

I. INTRODUCTION

Various frameworks, such as differential privacy [1], have been introduced to protect the privacy of individuals whose data is stored in online databases. These methods most often rely on the addition of noises with Laplace or Gaussian distributions to the outcome of queries on the databases containing the private information. More recently, differential privacy has found its way to control systems and signal processing [2]–[5]. In addition to differential privacy, information theoretic methods (using mutual information or Fisher information as a measure of privacy) have been also developed within the control and estimation community for preserving the privacy of individuals [6]–[8]. These methods also rely on the addition of noises which can be tailored for the specific problem at hand in order to protect the private data.

Although privacy preserving, the additive noise might also make it harder for an outsider to be able to use the reported data for identifying malicious behavior. For instance, in the smart meter privacy examples in [6], [7], a battery (which can be modeled as an additive noise with bounded support) is being used to mask the consumption patterns of the household. This ensures the privacy of the household. However, the battery operation will also makes it hard for the power authority to learn about the presence

of malicious agents based on the provided smart meter data. This is because deviations of the smart meter readings from the power authority's expectations (formed on the basis of historical data or models of household consumption) can be attributed equally to the implemented privacy-preserving mechanism or a malicious entity. A systematic analysis of the trade-off between privacy and security is the topic of this paper.

Specifically, a problem setup is considered in which everyone can submit linear queries to an online server containing a vector of private data. The server, in return, provides a systematically corrupted response to the submitted queries. The corruption involves using an additive noise to preserve the privacy of the entries of the database, i.e., the aforementioned vector of private data. The server determines the statistics of the noise so that estimation error of the vector of private data is maximized under a constraint on the quality of the supplied response, captured by the variance of the additive noise. Noting that the estimation error of the private vector is a function of policy used for generating the estimate, the Cramér-Rao bound [9, p. 169] is used to develop a universal measure of privacy which is inversely proportional to the trace of the Fisher information matrix. This measure of privacy is independent of the actions of the eavesdropper and is thus universal. It is assumed that an adversary can inject a bias to the server's response. The ability of users to be able to detect the presence of a bias (and thus raising a security alarm) is related to the Kullback-Leiber divergence of the output distribution with and without the additive bias. This provides a measure of security. The choice is motivated by the Chernoff-Stein Lemma (see, e.g., [10]) relating the probability of false negative (in the sense that a false bias injection attack escaping undetected) when using likelihood ratio hypothesis testing is a decreasing function of the Kullback-Leiber divergence of the output distribution with and without the additive bias. An optimization problem for balancing between privacy and security is proposed and solved. The solution in fact shows that the level of guaranteed privacy times the level of security is upper bounded by a constant. Therefore, by increasing the level of privacy, the security guarantees weaken and *vice versa*. This observation can be generalized to any distribution in fact and is thus a fundamental property of the framework. Subsequently the differential privacy framework is studied for which the same limitation is also observed.

Note that the use of Fisher information as a measure of privacy is not novel [7], [11], [12]; however, a systematic method for balancing privacy and security is completely missing from the literature. This is the topic of the current

F. Farokhi is with the CSIRO's Data61 and the Department of Electrical and Electronic Engineering at the University of Melbourne, Australia. e-mail: ffarokhi@unimelb.edu.au, farhad.farokhi@data61.csiro.au

P. Mohajerin Esfahani is with the Delft Center for Systems and Control at the Delft University of Technology, the Netherlands. e-mail: P.MohajerinEsfahani@tudelft.nl

The work of F. Farokhi was supported by the McKenzie Fellowship from the University of Melbourne, the VESKI Victoria Fellowship from the Victorian State Government, and a grant (MyIP: ID6874) from Defence Science and Technology Group (DSTG).

The work of P. Mohajerin Esfahani was supported by the Swiss National Science Foundation under the grant P2EZP2_165264.

paper.

Recently, in [13], it was shown that differential privacy noise can prevent detection of integrity attack in dynamical systems. This is because the additive noise of differential privacy provides new avenues for an attacker to inject false information without raising suspicion. The results of this paper, although having similar interpretations, are different from [13]. Most importantly, using the Fisher information as a measure of privacy and the Kullback-Leiber divergence as a measure of security, we can develop a more fundamental understanding of the trade-off between security and privacy without restricting the framework to differential privacy.

The rest of the paper is organized as follows. First, the problem formulation introducing the measures of privacy and security is presented in Section II. The results capturing the trade-off between privacy and security are then developed in Section III. Finally, the paper is concluded in Section IV.

II. PROBLEM FORMULATION

Consider the communication block diagram in Figure 2. A trustworthy server has access to a vector $x \in \mathcal{X} \subseteq \mathbb{R}^n$ whose entries must be kept private. Any agent, including those with an interest on infringing on the privacy of the individuals whose data is stored on the server, can submit a linear query of the form Cx to the server with observation matrix $C \in \mathbb{R}^{m \times n}$.

Assumption 1: C has full row rank.

The server returns a response to the query of the form $z = Cx + w$, where $w \in \mathbb{R}^m$ is an additive privacy-preserving noise with probability density function $\gamma : \mathbb{R}^m \rightarrow \mathbb{R}_{\geq 0}$.

Assumption 2: γ is twice continuously differentiable and $\text{supp}(\gamma) := \{w \in \mathbb{R}^m \mid \gamma(w) > 0\}$ may only differ from \mathbb{R}^m over a Lebesgue measure zero set.

These are technical assumptions that allow us to efficiently capture the optimal trade-off between security and privacy. The first part of Assumption 2 simplifies the search for the optimal privacy-preserving policy by allowing the use of tools available from the calculus of variations [14]. The second part of Assumption 2 ensures that the Fisher information matrix is well-defined and its trace is a convex function of γ [12]. The set of all such probability density functions is denoted by Γ .

A. Measure of Privacy

In this paper, the Fisher information is utilized as a measure of privacy. In fact, the server aims at increasing

$$\mathcal{P}(\gamma) := 1/\text{Tr}(W\mathcal{I}), \quad (1)$$

where the weighting matrix W is a positive definite matrix and \mathcal{I} is the Fisher information matrix defined as

$$\mathcal{I} := \int \frac{\partial \log(\gamma(w))}{\partial w} \frac{\partial \log(\gamma(w))}{\partial w}^\top \gamma(w) dw.$$

Note that the Fisher information matrix is a function of the probability density function γ . This measure has been recently utilized within privacy literature; see, e.g., [7], [12]. The motivation behind this selection is given in what follows.

The server wishes to keep the entries of the vector x private. Therefore, it aims to select a probability density function $\gamma \in \Gamma$ to maximize $\mathbb{E}\{\|\Pi_x(x - \hat{x}(y))\|_2^2\}$, where Π_x is a weighting matrix and $\hat{x}(y)$ is an estimator that an eavesdropper may use to estimate the value of the vector x based on the received message y .

Noting that the term $\mathbb{E}\{\|\Pi_x(x - \hat{x}(y))\|_2^2\}$ is a function of $\hat{x}(y)$, which makes the privacy measure depending on the eavesdropper (whose actions may not be known in advance), a lower bound of this term based on the Fisher information matrix is optimized. Using the Cramér-Rao bound [15], under mild assumptions, it can be shown that

$$\begin{aligned} \mathbb{E}\{\|\Pi_x(x - \hat{x}(y))\|_2^2\} &= \text{Tr}(\Pi_x^\top \Pi_x \mathbb{E}\{(x - \hat{x}(y))(x - \hat{x}(y))^\top\}) \\ &\geq \text{Tr}(\Pi_x^\top \Pi_x ((g(x) - x)(g(x) - x)^\top \\ &\quad + G(x)\mathcal{I}_x^\dagger G(x)^\top)) \\ &\geq \text{Tr}(\Pi_x^\top \Pi_x (g(x) - x)(g(x) - x)^\top) \\ &\quad + \text{Tr}(\Pi_x^\top \Pi_x \mathcal{I}_x^\dagger) \lambda_{\min}(G(x)^\top G(x)) \end{aligned}$$

where $g(x) = \mathbb{E}\{\hat{x}(y)\}$, $G(x)$ is the Jacobian of $g(x)$, $\mathcal{I}_x = C^\top \mathcal{I} C$, and X^\dagger denotes the Moore-Penrose pseudo-inverse of any matrix X . Note that, if $G(x)$ is a full rank matrix (e.g., for all unbiased estimators), $\text{Tr}(\Pi_x^\top \Pi_x \mathcal{I}_x^\dagger)$ can be utilized as a measure of privacy that is independent of the behavior of the adversary. This is because by increasing $\text{Tr}(\Pi_x^\top \Pi_x \mathcal{I}_x^\dagger)$, the estimation error also increases. Noting that $\text{Tr}(\Pi_x^\top \Pi_x \mathcal{I}_x^\dagger)$ is not a concave function of γ , the measure of privacy can be replaced with $1/\text{Tr}((\Pi_x^\top \Pi_x)^\dagger \mathcal{I}_x)$ because¹ $\text{Tr}(\Pi_x^\top \Pi_x \mathcal{I}_x^\dagger) \geq 1/\text{Tr}((\Pi_x^\top \Pi_x)^\dagger \mathcal{I}_x)$. Interestingly, $1/\text{Tr}((\Pi_x^\top \Pi_x)^\dagger \mathcal{I}_x)$ is a concave function of the probability density function γ because $\text{Tr}((\Pi_x^\top \Pi_x)^\dagger \mathcal{I}_x)$ is a convex function [7]. Thus maximizing $1/\text{Tr}((\Pi_x^\top \Pi_x)^\dagger \mathcal{I}_x)$ is a more computationally-friendly task. Note that, for this motivational example, the weighting function in \mathcal{P} is given by $W = C(\Pi_x^\top \Pi_x)^\dagger C^\top$.

Remark 1 (Worst-case analysis): In the preceding discussion, it is assumed that $G(x)$ is full rank, which although sensible (as estimators, such as least mean square, meet this condition), might not be desirable. In [16], it was shown that $1/\text{Tr}(CC^\top \mathcal{I})$ can be proved to be a measure of privacy by studying worst-case privacy violations. In worst-case privacy attack, an eavesdropper has access to all the entries of the vector x except one of them (the subject of the privacy infringement or eavesdropping attack) and it would like to infer the value of that entry based on the response to the submitted query. In that case, the weighting function in \mathcal{P} is given by $W = CC^\top$.

B. Measure of Performance

Noting that the error $\mathbb{E}\{\|\Pi_x(x - \hat{x}(y))\|_2^2\}$ can be made potentially unbounded (since there is no prior on x and the server can add a Gaussian noise with increasing covariance), the server also aims at maintaining a sensible level

¹Note that, for any non-zero semi-definite matrix A , it can be deduced that $\text{Tr}(A^\dagger)\text{Tr}(A) \geq \text{Tr}(A^\dagger A) \geq 1$ while implies that $\text{Tr}(A) \geq 1/\text{Tr}(A^\dagger)$.




Fig. 1. Communication diagram.

of performance by enforcing that

$$\mathcal{Q}(\gamma) := \mathbb{E}\{\|y - Cx\|_2^2\} \quad (2)$$

remains below a certain level η , i.e., it is desired that the variance of the probability density function γ is less than the provided upper bound by ensuring that

$$\int w^\top w \gamma(w) \leq \eta.$$

C. Measure of Security

The communication channel can be infiltrated by an adversary, which may inject the bias $d \in \mathbb{R}^p$. Thus, the final output is given by $y = Cx + w + Fd$. Therefore, the server may also wish to make it possible for potential users to identify bias injection attacks performed by an adversary (to raise an alarm). This can be achieved by enforcing a constraint on an appropriately selected measure of security.

In this paper, the sensitivity of Kullback-Leibler divergence between the probability density functions $\gamma(y - Cx)$ and $\gamma(y - Cx - Fd)$ is used as a measure of security. This shows that how easy it is for to distinguish between the probability density functions $\gamma(y - Cx)$ and $\gamma(y - Cx - Fd)$ for small offset term d (which an adversary may use to avoid being detected). Therefore, the measure of security is given by

$$\mathcal{S}(\gamma) := \min_{\xi \in \mathbb{R}^p} \lim_{d = \xi, \epsilon \rightarrow 0} \frac{\mathcal{KL}}{\|d\|_2^2},$$

where

$$\mathcal{KL} := \int \gamma(y - Cx) \log \left(\frac{\gamma(y - Cx)}{\gamma(y - Cx - Fd)} \right) dy.$$

In this framework, it is desired to ensure that $\mathcal{S}(\gamma) \geq \alpha$, where $\alpha > 0$ is an appropriately selected constant.

This choice is motivated by that, at least for discrete random variables, it can be proved that the probability of false negative in the sense that a bias injection attack remains undetected when using likelihood ratio hypothesis testing is a decreasing function of \mathcal{KL} [10, Chernoff-Stein Lemma]. Although such a result may not hold in general, this observation can be used a motivation for the use of the Kullback-Leibler divergence as a measure of security. Thus, to keep the probability of false negatives small, a constraint of the form $\mathcal{KL}/\|d\|_2^2 \geq \alpha$. Note that \mathcal{KL} grows as a function of d so a scaled version of the Kullback-Leibler divergence $\mathcal{KL}/\|d\|_2^2$ is considered. Considering that the adversary intends to not be detected (and the fact that the identification becomes easier as $\|d\|$ grows), it would be of interest to study small bias vectors d .

D. Balancing Privacy and Security

With the definitions of the measures of privacy and security in hand, it is now time to pose the problem mathematically.

Problem 1: Find privacy preserving policy

$$\gamma^* \in \arg \max_{\gamma \in \Gamma} \mathcal{P}(\gamma), \quad (3a)$$

$$\text{s.t. } \mathcal{S}(\gamma) \geq \alpha, \quad (3b)$$

$$\mathcal{Q}(\gamma) \leq \eta. \quad (3c)$$

A popular framework for studying privacy is differential privacy; see, e.g., [17]. The server's response is ϵ -differentially private if

$$\mathbb{P}\{y \in \mathcal{Y} | x'\} \leq \exp(\epsilon) \mathbb{P}\{y \in \mathcal{Y} | x\} \quad (4)$$

for all $x, x' \in \mathcal{X}$ that are only different in maximum on entry and \mathcal{Y} is a Lebesgue-measurable subset of \mathbb{R}^m .

Problem 2: Find ϵ -differentially private $\gamma^* \in \Gamma$ such that $\mathcal{S}(\gamma) \geq \alpha$.

Studying Problem 2 allows us to see if we can observe the same results as in [13] within this setup. Furthermore, it can be investigated that if such results are in agreement with the optimal additive noise extracted from solving Problem 1.

III. MAIN RESULTS

The first result of this paper, formalized in the following theorem, states that the Gaussian additive noise with an appropriately selected co-variance matrix provides the best balance between privacy and security requirements according to Problem 1.

Theorem 1: The solution to Problem 1 is given by

$$\gamma^*(w) = \frac{1}{\sqrt{\det(2\pi V_{ww})}} \exp \left(-\frac{1}{2} w^\top V_{ww}^{-1} w \right),$$

where

$$V_{ww} = \frac{\eta}{\text{Tr}(W^{1/2})} W^{1/2},$$

if $\text{Tr}(W^{1/2}) \lambda_{\min}(F^\top W^{-1/2} F) \geq 2\eta\alpha$.

Proof: By eliminating the security constraint, Problem 1 can be relaxed into

$$\gamma^* \in \arg \min_{\gamma \in \Gamma} \text{Tr}(W\mathcal{I}), \quad (5a)$$

$$\text{s.t. } \mathcal{Q}(\gamma) \leq \eta. \quad (5b)$$

Note that the duality gap in (5) is zero [18]. Therefore, the constraint on the variance can be added to the cost function using a Lagrange multiplier, which transforms the problem into

$$\max_{\lambda \geq 0} \min_{\gamma \in \Gamma} \text{Tr}(W\mathcal{I}) + \lambda(\mathcal{Q}(\gamma) - \eta). \quad (6)$$

Following the same line of reasoning as in [12], the solution of the inner problem in (6) is given by $\gamma^*(w) = u(w)^2$,

where

$$\begin{cases} \text{Tr}(WD^2u(w)) \\ \quad + (\mu - (\lambda/4)w^\top w)u(w) = 0, & w \in \mathcal{W}, \\ u(w) = 0, & w \in \partial\mathcal{W}, \\ u(w) \neq 0, & w \in \text{int}\mathcal{W}, \\ \int_{w \in \mathcal{W}} u(w)^2 dw = 1. \end{cases} \quad (7)$$

Note that the cost function and the constraint set are convex, the stationarity condition in (7) is sufficient for optimality. Further, if multiple density functions satisfy the conditions, they all exhibit the same cost. It can be shown that the following satisfies the stationarity condition:

$$u(w) = \frac{1}{\sqrt[4]{\det(2\pi V)}} \exp\left(-\frac{1}{4}w^\top V^{-1}w\right),$$

where $V = W^{1/2}/\sqrt{\lambda}$. This shows that

$$\min_{\gamma \in \Gamma} \text{Tr}(W\mathcal{I}) + \lambda(\mathcal{Q}(\gamma) - \eta) = \text{Tr}(WV^{-1}) + \lambda(\text{Tr}(V) - \eta).$$

Therefore, the outer optimization problem in (6) can be rewritten as

$$\max_{\lambda \geq 0} 2\text{Tr}(W^{1/2})\sqrt{\lambda} - \lambda\eta,$$

and as a result $\lambda^* = \text{Tr}(W^{1/2})^2/\eta^2$. Using [19], it can be shown that

$$\lim_{d=\varrho\xi, \varrho \rightarrow 0} \frac{\mathcal{K}\mathcal{L}}{\|d\|_2^2} = \frac{1}{2} \frac{\xi^\top \mathcal{I}_d \xi}{\xi^\top \xi}.$$

where $\mathcal{I}_d := F^\top \mathcal{I} F$. Thus,

$$\mathcal{S}(\gamma) = \min_{\xi \in \mathbb{R}^p} \lim_{d=\varrho\xi, \varrho \rightarrow 0} \frac{\mathcal{K}\mathcal{L}}{\|d\|_2^2} = \frac{1}{2} \lambda_{\min}(\mathcal{I}_d)$$

For γ^* , it can be seen that

$$\mathcal{I}_d = \sqrt{\lambda^*} F^\top W^{-1/2} F = \text{Tr}(W^{1/2}) F^\top W^{-1/2} F / \eta.$$

If $\mathcal{S}(\gamma) = (1/2)\lambda_{\min}(\mathcal{I}_d) \geq \alpha$, the solution of (5) is also a solution of (6). This concludes the proof. \blacksquare

Theorem 1 presents the solution of Problem 1 in the case where the constraint $\mathcal{Q}(\gamma) \leq \eta$ is active and $\mathcal{S}(\gamma) \geq \alpha$ is inactive. The following theorem extends this results to the case where the constraint $\mathcal{S}(\gamma) \geq \alpha$ is active and $\mathcal{Q}(\gamma) \leq \eta$ is inactive.

Theorem 2: Let

$$\mathcal{V} := \arg \min_{X \succeq 0} \text{Tr}(WX), \quad (8a)$$

$$\text{s.t. } F^\top XF \succeq 2\alpha I. \quad (8b)$$

The solution to Problem 1 is given by

$$\gamma(w) = \frac{1}{\sqrt{\det(2\pi V_{ww})}} \exp\left(-\frac{1}{2}w^\top V_{ww}^{-1}w\right) \quad (9)$$

if there exists $V_{ww}^{-1} \in \mathcal{V}$ such that $\text{Tr}(V_{ww}) \leq \eta$.

Proof: Note that $\mathcal{I}_d = F^\top \mathcal{I} F$. Assume that each $\mathcal{I} \succeq 0$ is realizable, i.e., there exists $\gamma(w)$ that results in it. Thus, Problem 1 can be transformed into the semi-definite program in (8). It remains to find a density function that has

a Fisher information equal to the solution of (8). This is in fact possible using a multivariate normal distribution with covariance matrix \mathcal{I}^{-1} . This concludes the proof. \blacksquare

For scalar queries, such as averaging, the solution to Problem 1 can be greatly simplified. This is shown in the following corollary.

Corollary 1: For scalar queries (i.e., $m = 1$), the solution to Problem 1 is given by

$$\gamma(w) = \frac{1}{\sqrt{2\pi V_{ww}}} \exp\left(-\frac{w^2}{2V_{ww}}\right), \quad (10)$$

where

$$V_{ww} = \begin{cases} \eta, & \eta \leq \lambda_{\min}(F^\top F)/\alpha, \\ 1/\alpha, & \text{otherwise.} \end{cases}$$

Proof: If $\eta \leq \lambda_{\min}(F^\top F)/\alpha$, the results of Theorem 1 can be used. Otherwise, the results of Theorem 2 should be utilized in which case it can be seen that $\text{Tr}(WX) = WX$ (since both X and W are scalars) and $F^\top XF = (F^\top F)X$ (again because X is a scalar). Hence, the optimization problem in (8) can be transformed into $V_{ww}^{-1} \in \arg \min_{X \geq \alpha} X$. Thus, $V_{ww} = 1/\alpha$. This concludes the proof. \blacksquare

Corollary 2: For the optimal probability density function in Corollary 1, $\mathcal{S}(\gamma)\mathcal{P}(\gamma) = \lambda_{\min}(F^\top F)/(2W)$.

Proof: For the optimal policy in Corollary 1, it can be seen that $\mathcal{K}\mathcal{L} = \frac{1}{2}(Fd)^2V_{ww}^{-1}$, and, as a result, $\mathcal{S}(\gamma) = \frac{1}{2}\lambda_{\min}(F^\top F)/V_{ww}$. On the other hand, $\mathcal{P}(\gamma) = V_{ww}/W$. \blacksquare

Proposition 2 shows that by increasing $\mathcal{P}(\gamma)$ to achieve a higher privacy guarantee, $\mathcal{S}(\gamma)$ decreases, which makes the system more vulnerable to bias injection attacks. In fact, in lay terms, it can be expressed that

$$\boxed{\text{“privacy} \times \text{security} = \text{constant”}}. \quad (\star)$$

In what follows, it is shown that Corollary 2 and its interpretation in (\star) hold for any probability density function $\gamma(w)$ if $m = 1$ (and not necessarily the solution of Problem 1)

Proposition 1 (Trade-off between Privacy and Security):

For $m = 1$, $\mathcal{S}(\gamma)\mathcal{P}(\gamma) = \lambda_{\min}(F^\top F)/(2W)$ for any $\gamma \in \Gamma$.

Proof: For any density function, it can be seen that

$$\mathcal{S}(\gamma) = \lim_{d=\varrho\xi, \varrho \rightarrow 0} \frac{\mathcal{K}\mathcal{L}}{\|d\|_2^2} = \frac{1}{2} \frac{\xi^\top F^\top \mathcal{I} F \xi}{\xi^\top \xi} = \frac{1}{2} \lambda_{\min}(F^\top \mathcal{I} F).$$

Thus, $\mathcal{P}(\gamma)\mathcal{S}(\gamma) = \lambda_{\min}(F^\top \mathcal{I} F)/(2\text{Tr}(W\mathcal{I}))$. For $m = 1$, it can be shown that $\mathcal{P}(\gamma)\mathcal{S}(\gamma) = \lambda_{\min}(F^\top F)/(2W)$ because \mathcal{I} is scalar. \blacksquare

Figure 2 illustrates the trade-off between measure of privacy $\mathcal{P}(\gamma)$ for the optimal policy in Corollary 1 versus the lower bound on the measure of security α for various quality of response guarantees η . In this numerical example, $m = 1$, $F = 1$, and $W = 1$. The plateau on the achievable privacy guarantee for small values of α is caused by the constraint on the quality of measurement $\mathcal{Q}(\gamma)$. The gray area denotes the cases for which $\mathcal{P}(\gamma)\alpha \leq 1$. All these cases are achievable for various values of η . Note that this not in contrast with the results of Propositions 2 and 1 as they

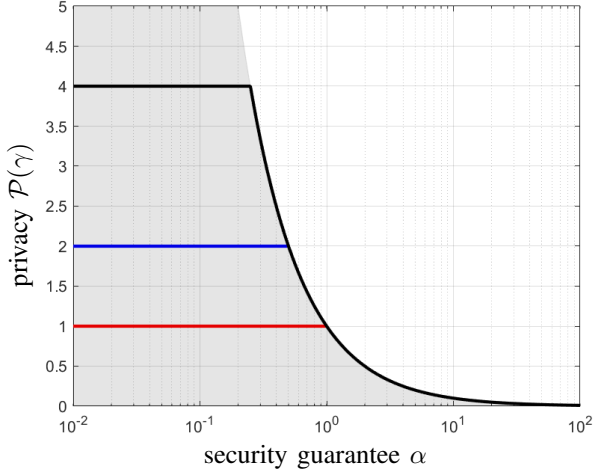


Fig. 2. The trade-off between measure of privacy $\mathcal{P}(\gamma)$ for the optimal policy in Corollary 1 versus the lower bound on the measure of security α for various response quality guarantees $\eta = 1$ (solid red —), $\eta = 2$ (solid red —), and $\eta = 4$ (solid green —). The plateau on the achievable privacy guarantee for small α is caused by the constraint on the quality of measurement $\mathcal{Q}(\gamma)$. The gray area denotes the cases for which $\mathcal{P}(\gamma)\alpha \leq 1$.

explore $\mathcal{P}(\gamma)\mathcal{S}(\gamma)$ (which is not necessarily equal to $\mathcal{P}(\gamma)\alpha$ as the constraint $\mathcal{S}(\gamma) \geq \alpha$ is not always active).

Now, we are ready to explore the solution of Problem 2 regarding the balance between privacy and security in the differential privacy framework.

Theorem 3: For scalar problems, i.e., $m = 1$, and $\epsilon \geq \Delta\sqrt{2\alpha}$, the solution to Problem 2 is given by

$$\gamma(w) = \frac{1}{2\Delta/\epsilon} \exp\left(-\frac{|w|}{\Delta/\epsilon}\right), \quad (11)$$

where $\Delta := \sup_{x, x' \in \mathcal{X}: \|x-x'\|_0 \leq 1} |C(x-x')|$.

Proof: Note that

$$\begin{aligned} \frac{p(y|x')}{p(y|x)} &= \exp\left(\frac{|y-Cx| - |y-Cx'|}{\Delta/\epsilon}\right) \\ &\leq \exp\left(\frac{|C(x'-x)|}{\Delta/\epsilon}\right) \\ &\leq \exp(\epsilon), \end{aligned} \quad (12)$$

where the first inequality follows from that $|y-Cx| = |y-Cx + Cx' - Cx'| \leq |y-Cx'| + |C(x'-x)|$. Integrating both sides of (12) concludes the proof. Furthermore, γ meets $\mathcal{I}_d = \epsilon^2/\Delta^2$. Thus, $\mathcal{I}_d \geq 2\alpha$ if and only if $\epsilon \geq \Delta\sqrt{2\alpha}$. ■

For the ϵ -differentially private distribution in Theorem 3, the following can be proved:

$$\begin{aligned} \mathcal{KL} &= \int \frac{1}{2\Delta/\epsilon} \exp\left(-\frac{|y-Cx|}{\Delta/\epsilon}\right) \\ &\quad \times \left(\frac{|y-Cx-Fd|}{\Delta/\epsilon} - \frac{|y-Cx|}{\Delta/\epsilon}\right) dy \\ &= \int \frac{1}{2\Delta/\epsilon} \exp\left(-\frac{|\bar{y}|}{\Delta/\epsilon}\right) \left(\frac{|\bar{y}-Fd|}{\Delta/\epsilon} - \frac{|\bar{y}|}{\Delta/\epsilon}\right) d\bar{y} \\ &= \exp(-|Fd|\epsilon/\Delta) - 1 + |Fd|\epsilon/\Delta. \end{aligned}$$




Fig. 3. The trade-off between measure of privacy $1/\epsilon$ and the measure of security $\mathcal{S}(\gamma)$ for the differentially-private policy in Theorem 3.

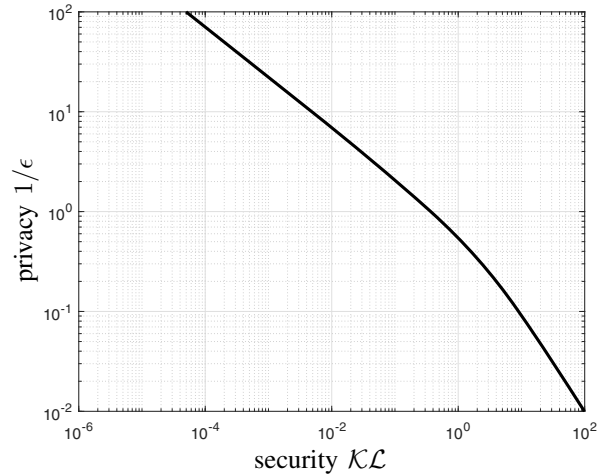


Fig. 4. The trade-off between measure of privacy $1/\epsilon$ and \mathcal{KL} for the differentially-private policy in Theorem 3.

Therefore

$$\mathcal{S}(\alpha) = \min_{\xi} \lim_{d=\epsilon\xi, \epsilon \rightarrow 0} \frac{\mathcal{KL}}{\|d\|_2^2} = \frac{\lambda_{\min}(F^T F)\epsilon^2}{2\Delta^2}.$$

This implies that by increasing the privacy guarantee (which is inversely proportional to ϵ), the security level decreases and *vice versa*. This is a similar observation to that of (\star) .

Figure 3 illustrates the trade-off between measure of privacy $1/\epsilon$ versus the measure of security $\mathcal{S}(\gamma)$ for the differentially-private policy in Theorem 3. Here, $m = 1$, $\Delta = 1$, and $F = 1$. Recalling that $\mathcal{S}(\gamma)$ is motivated by small biases d , we also explore \mathcal{KL} for differentially-private policies. This relationship is shown in Figure 4. Clearly, the same trend regarding the inverse relationship of the privacy and security can still be observed.

IV. CONCLUSIONS AND FUTURE WORK

A framework was developed in which linear query can be submitted to a server containing private data. The server

provides a response to the query corrupted using an additive noise to preserve the privacy of those whose data is on the server. It is shown that the level of guaranteed privacy times the level of security is always upper bounded by a constant and, as a result, higher privacy guarantees dictates weakened security guarantees. Future work can focus on dynamic problems.

REFERENCES

- [1] C. Dwork, "Differential privacy," in *Encyclopedia of Cryptography and Security* (H. C. A. van Tilborg and S. Jajodia, eds.), Boston, MA: Springer US, 2011.
- [2] Z. Li and T. J. Oechtering, "Privacy-constrained parallel distributed neyman-pearson test," *IEEE Transactions on Signal and Information Processing over Networks*, vol. 3, no. 1, pp. 77–90, 2017.
- [3] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Transactions on Automatic Control*, vol. 59, no. 2, pp. 341–354, 2014.
- [4] H. Sandberg, G. Dán, and R. Thobaben, "Differentially private state estimation in distribution networks with smart meters," in *Proceedings of the 54th IEEE Conference on Decision and Control*, pp. 4492–4498, 2015.
- [5] Z. Huang, Y. Wang, S. Mitra, and G. E. Dullerud, "On the cost of differential privacy in distributed control systems," in *Proceedings of the 3rd International Conference on High Confidence Networked Systems*, pp. 105–114, ACM, 2014.
- [6] S. R. Rajagopalan, L. Sankar, S. Mohajer, and H. V. Poor, "Smart meter privacy: A utility-privacy framework," in *Proceedings of the IEEE International Conference on Smart Grid Communications (SmartGrid-Comm)*, pp. 190–195, IEEE, 2011.
- [7] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, 2017. In Press.
- [8] S.-Y. Kung, "Compressive privacy: From information/estimation theory to machine learning [lecture notes]," *IEEE Signal Processing Magazine*, vol. 34, no. 1, pp. 94–112, 2017.
- [9] J. Shao, *Mathematical Statistics*. Springer Texts in Statistics, Springer-Verlag New York, 2003.
- [10] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. Wiley, 2012.
- [11] H. Anderson, "Efficiency versus protection in a general randomized response model," *Scandinavian Journal of Statistics*, pp. 11–19, 1977.
- [12] F. Farokhi and H. Sandberg, "Optimal constrained additive noise distribution minimizing Fisher information for ensuring privacy," in *Proceedings of the 56th IEEE Conference on Decision and Control*, pp. 2692–2697, 2017.
- [13] J. Giraldo, A. A. Cardenas, and M. Kantarcioglu, "Security vs. privacy: How integrity attacks can be masked by the noise of differential privacy," in *Proceedings of the American Control Conference*, pp. 1679–1684, 2017.
- [14] D. E. Kirk, *Optimal Control Theory: An Introduction*. Dover Books on Electrical Engineering Series, Dover Publications, 2004.
- [15] A. O. Hero, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Transactions on Signal Processing*, vol. 44, no. 8, pp. 2026–2041, 1996.
- [16] F. Farokhi and H. Sandberg, "Ensuring privacy with constrained additive noise by minimizing fisher information." Submitted, 2017.
- [17] C. Dwork, "Differential privacy: A survey of results," in *Theory and Applications of Models of Computation: 5th International Conference, TAMC 2008, Xi'an, China, April 25-29, 2008. Proceedings* (M. Agrawal, D. Du, Z. Duan, and A. Li, eds.), pp. 1–19, Berlin, Heidelberg: Springer Berlin Heidelberg, 2008.
- [18] V. Jeyakumar and H. Wolkowicz, "Zero duality gaps in infinite-dimensional programming," *Journal of Optimization Theory and Applications*, vol. 67, no. 1, pp. 87–108, 1990.
- [19] F. Critchley, P. Marriott, and M. Salmon, "Preferred point geometry and the local differential geometry of the Kullback-Leibler divergence," *The Annals of Statistics*, pp. 1587–1602, 1994.