

BETTER SAFE THAN SORRY: RISK-AWARE NONLINEAR BAYESIAN ESTIMATION

Dionysios S. Kalogerias, Luiz F. O. Chamon, George J. Pappas, and Alejandro Ribeiro

Department of Electrical & Systems Engineering,
University of Pennsylvania, Philadelphia, PA 19104, USA

ABSTRACT

Despite the simplicity and intuitive interpretation of minimum mean squared error (MMSE) estimators, their effectiveness in certain scenarios is questionable. Indeed, minimizing squared errors on average does not provide any form of stability, as the volatility of the estimation error is left unconstrained. When this volatility is statistically significant, the difference between the average and realized performance of the MMSE estimator can be drastically different. To address this issue, we introduce a new risk-aware MMSE formulation which trades between mean performance and risk by explicitly constraining the expected predictive variance of the involved squared error. We show that, under mild moment boundedness conditions, the corresponding risk-aware optimal solution can be evaluated explicitly, and has the form of an appropriately biased nonlinear MMSE estimator. We further illustrate the effectiveness of our approach via several numerical examples, which also showcase the advantages of risk-aware against risk-neutral MMSE estimation, especially in models involving skewed, heavy-tailed distributions.

Index Terms—MMSE Estimation, Constrained Bayesian Estimation, Risk-Aware Optimization, Risk Measures.

1. INTRODUCTION

Critical applications require that stochastic decisions be made not only on the basis of minimizing average losses, but also safeguarding against less probable, though possibly catastrophic, events. Examples appear naturally in many areas, including wireless industrial control [2], energy [3, 4], finance [5, 6, 7], robotics [8, 9], LIDAR [10], and networking [11]. In such cases, the ultimate goal is to obtain *risk-aware decision policies* that hedge against statistically significant extreme losses, even at the cost of slightly sacrificing performance under nominal operating conditions.

To illustrate this effect, consider the problem of estimating a random state X from observations corrupted by state-dependent noise, namely $Y|X \sim \mathcal{N}(X, 9X^2)$ (see Section 5). In such a setting, either small or large values of Y provide highly ambiguous evidence, since, in both cases, they can come from either small or large values of X . This is corroborated by Fig. 1, which displays the posterior $\mathcal{P}_{X|Y}$ for two values of Y . While the MMSE estimator may incur severe losses, the risk-aware estimator we develop in this work –red lines in Fig. 1– hedges against observation ambiguity, thus avoiding extreme prediction errors. This is achieved by biasing estimates towards the tail of the posterior $\mathcal{P}_{X|Y}$, by the right amount, for each value of Y . Although the risk-aware estimator may incur larger losses on average, it performs statistically more consistently across realizations of Y (also see Fig. 2, Section 5). It is also worth contrasting risk-awareness with statistical robustness, whose goal is to protect against deviations from a nominal model. Robust estimators –green lines in Fig. 1– promote insensitivity to tail events, which they designate as statistically insignificant. On the other hand, estimators resulting

Uncertainty	Frequent	Infrequent
Significant	Model	Risk
Insignificant	Noise	Outliers

Table 1. Classification of statistical uncertainty.




Fig. 1. Comparison between risk-neutral and risk-aware estimates.

from risk-aware formulations treat these events as statistically significant, though relatively infrequent (see Table 1).

Over the last three decades, risk-aware optimization has grown increasingly popular and has been studied in the contexts of both decision making and learning [12, 13, 14, 15, 16, 17, 18, 19]. In risk-aware optimization, expectations are replaced by more general functionals, called *risk measures* [20], whose purpose is to quantify the statistical volatility of random losses, as well as mean performance. Popular examples include mean-variance functionals [5, 20], mean-semideviations [16], and Conditional Value-at-Risk (CVaR) [21].

In Bayesian estimation, risk awareness is typically achieved by replacing the classical quadratic cost with its exponentiation [22, 23, 24, 25, 26]. However, although sometimes effective, this approach is not without limitations. First, the need for finiteness of the moment generating function of the quadratic cost excludes heavy-tailed distributions, which are precisely those that incur in high risk. Second, the exponential approach does not provide an interpretable way to control the trade-off between mean performance and risk, making it hard to use in settings where explicit risk levels must be met. Third, it does not result in a simple, general solution as in classical MMSE estimation, challenging its practical applicability. Finally, it does not effectively quantify *observation-induced risk*, inherent to problems in which measurements provide ambiguous evidence.

In this work, we pose risk-aware *functional* Bayesian estimation as a constrained MMSE problem, where squared errors are minimized on average, *subject to* a bound on their expected conditional variance. We show that, under mild conditions, this formulation results in a convex variational problem that admits a closed-form solution. The resulting *optimal risk-aware nonlinear MMSE estimator* is applicable to a wide variety of generative models, including highly skewed and/or heavy-tailed distributions. The effectiveness of our approach is confirmed via numerical examples, also demonstrating its advantages against risk-neutral MMSE estimation.

This work is supported by the National Science Foundation (NSF) under grant CPS 1837253 and the Intel Science & Technology Center for Wireless Autonomous Systems.

Note: Due to space constraints, all proofs are deferred to [1].

2. PROBLEM FORMULATION

Let $(\Omega, \mathcal{F}, \mathcal{P})$ be a probability space, and consider an arbitrary pair of random elements $\mathbf{X} : \Omega \rightarrow \mathbb{R}^M$ and $\mathbf{Y} : \Omega \rightarrow \mathbb{R}^N$ on (Ω, \mathcal{F}) . We are interested in the problem of estimating \mathbf{X} from a *single* realization of \mathbf{Y} in a Bayesian setting, namely by assuming knowledge of the joint probability distribution $\mathcal{P}_{(\mathbf{X}, \mathbf{Y})}$. We may conveniently think of \mathbf{Y} as available *observations*, on the basis of which we would like to make predictions about the *hidden state* \mathbf{X} . Undoubtedly, this general problem is fundamental in many areas, including statistics, signal processing, machine learning, and control, and with numerous interesting applications.

Of course, an established approach to the prediction problem considered is to choose an estimator $\widehat{\mathbf{X}} : \Omega \rightarrow \mathbb{R}^M$ as a solution to the *stochastic variational* MMSE program

$$\begin{aligned} & \underset{\widehat{\mathbf{X}} : \Omega \rightarrow \mathbb{R}^M}{\text{minimize}} && \mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\} \\ & \text{subject to} && \widehat{\mathbf{X}} \text{ is } \mathcal{Y}\text{-measurable} \end{aligned}, \quad (1)$$

where $\mathcal{Y} \equiv \sigma\{\mathbf{Y}\}$ denotes the sub- σ -algebra of \mathcal{F} generated by \mathbf{Y} . Problem (1) is well-understood under rather general conditions. In fact, if we *merely* assume that $\mathbf{X} \in \mathcal{L}_1(\Omega, \mathcal{Y}, \mathcal{P}; \mathbb{R}^M) \equiv \mathcal{L}_{1|\mathcal{Y}}^M$, an optimal solution to (1) is given by any conditional expectation of \mathbf{X} relative to \mathbf{Y} , i.e., $\widehat{\mathbf{X}}^*(\mathbf{Y}) \equiv \mathbb{E}\{\mathbf{X}|\mathbf{Y}\}$.

However, despite the simplicity of MMSE estimation, as well as its intuitive geometric interpretation in Hilbert space whenever $\mathbf{X} \in \mathcal{L}_{2|\mathcal{Y}}^M$, its effectiveness is often questionable. Indeed, minimizing the squared error $\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2$ in expectation does *not* provide stability or robustness, in the sense that statistically significant variability of the resulting *optimal prediction error* is uncontrolled. In other words, the MMSE problem (1) is *risk-neutral*. This has important consequences from a practical perspective, since the error realization $\|\mathbf{X} - \widehat{\mathbf{X}}^*(\mathbf{Y})\|_2^2$ experienced in practice may be far from the expected value $\mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}^*(\mathbf{Y})\|_2^2\}$, or even the predictive statistic $\mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}^*(\mathbf{Y})\|_2^2|\mathbf{Y}\}$. It is then clear that achieving small error variability is at least as desirable as achieving minimal errors on average.

Motivated by the previous discussion, we consider a nontrivial variation of the risk-neutral MMSE problem (1), striking a *balance between mean performance and risk*. Specifically, we introduce and study the *constrained* stochastic variational problem

$$\begin{aligned} & \underset{\widehat{\mathbf{X}} : \Omega \rightarrow \mathbb{R}^M}{\text{minimize}} && \mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\} \\ & \text{subject to} && \mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\}\} \leq \varepsilon, \\ & && \widehat{\mathbf{X}} \text{ is } \mathcal{Y}\text{-measurable} \end{aligned}, \quad (2)$$

where

$$\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\} \triangleq \mathbb{E}\{(\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2 - \mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2|\mathbf{Y}\})^2|\mathbf{Y}\} \quad (3)$$

is the predictive variance of $\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2$ relative to \mathbf{Y} , and $\varepsilon > 0$ is a fixed risk tolerance. In words, problem (2) constrains the *expected predictive variance of the quadratic cost* $\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2$, known in the statistics literature as the *unexplained component* of its variance; the latter is due to the law of total variance. In other words, the constraint quantifies the uncertainty of MMSE-optimally predicting the quadratic cost *achieved by choosing an estimator* $\widehat{\mathbf{X}}(\mathbf{Y})$, on the basis of the observations \mathbf{Y} . Of course, $\mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\}\}$ is a

measure of risk. Therefore, we suggestively refer to the task fulfilled by problem (2) as *risk-aware MMSE estimation*.

Problem (2) confines the search for an optimal MMSE estimator within the family of estimators exhibiting risk (in the sense described above) within tolerance ε ; thus, problem (2) is well-motivated. Naturally, an optimal solution to the risk-aware problem (2) in general achieves larger MSE as compared to the risk-neutral problem (1). However, the statistical variability of the squared errors achieved by the former will be explicitly controlled, according to the tunable tolerance ε , resulting in more stable statistical prediction.

3. CONVEX VARIATIONAL QCQP REFORMULATION

As it turns out, in such a general form, the risk-aware MMSE problem (2) is rather challenging to study, let alone solve. Therefore, in the following, we will consider a slightly more constrained version of (2), by enforcing square integrability on the decision $\widehat{\mathbf{X}}$, namely,

$$\begin{aligned} & \underset{\widehat{\mathbf{X}} : \Omega \rightarrow \mathbb{R}^M}{\text{minimize}} && \mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\} \\ & \text{subject to} && \mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X} - \widehat{\mathbf{X}}\|_2^2\}\} \leq \varepsilon, \\ & && \widehat{\mathbf{X}} \in \mathcal{L}_{2|\mathcal{Y}}^M \end{aligned}, \quad (4)$$

Of course, the additional \mathcal{L}_2 constraint in problem (4) may not be in favor of generality, per se, but it is harmless for almost every practical consideration. Further, in the following we make use of the following regularity condition on the statistical behavior of (\mathbf{X}, \mathbf{Y}) .

Assumption 1. *It is true that $\mathbb{E}\{\|\mathbf{X}\|_2^3|\mathbf{Y}\} \in \mathcal{L}_{2|\mathcal{Y}}^1$.*

In words, Assumption 1 simply says that the third-order moment filter $\mathbb{E}\{\|\mathbf{X}\|_2^3|\mathbf{Y}\}$ is of finite energy. Using Assumption 1, problem (4) may be conveniently reformulated, as the next result suggests.

Lemma 1. (QCQP Reformulation of Problem (4)) *Suppose that Assumption 1 is in effect, and define the posterior covariance*

$$\Sigma_{\mathbf{X}|\mathbf{Y}} \triangleq \mathbb{E}\{(\mathbf{X} - \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})(\mathbf{X} - \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})^T|\mathbf{Y}\} \succeq \mathbf{0}. \quad (5)$$

Then, problem (4) is well-defined and equivalent to the convex variational Quadratically Constrained Quadratic Program (QCQP)

$$\begin{aligned} & \underset{\widehat{\mathbf{X}} : \Omega \rightarrow \mathbb{R}^M}{\text{minimize}} && \frac{1}{2} \mathbb{E}\{\|\widehat{\mathbf{X}}\|_2^2 - 2(\mathbb{E}\{\mathbf{X}|\mathbf{Y}\})^T \widehat{\mathbf{X}} + \mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\}\} \\ & \text{subject to} && \mathbb{E}\{\widehat{\mathbf{X}}^T \Sigma_{\mathbf{X}|\mathbf{Y}} \widehat{\mathbf{X}} - (\mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \widehat{\mathbf{X}} \\ & && \quad - \mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})^T \widehat{\mathbf{X}}\} \\ & && \leq \frac{\varepsilon - \mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X}\|_2^2\}\}}{4} \\ & && \widehat{\mathbf{X}} \in \mathcal{L}_{2|\mathcal{Y}}^M \end{aligned}, \quad (6)$$

where all expectations and involved operations are well-defined.

Lemma 1 is very useful, because it shows the equivalence of problem (4) to the convex QCQP (6), which is well-defined and favorably structured. In particular, this reformulation will allow us to effectively study problem (4) by looking at its variational Lagrangian dual. Actually, as we discuss next, working in the dual domain will allow us to *solve* problem (4) in *closed-form*. Of course, such a closed form is important, not only because it provides an analytical, textbook-level solution to a functional risk-aware problem, which happens rather infrequently, but also because, as we will see, the solution itself provides intuition, highlights connections and enables comparison of problem (4) with its risk-neutral counterpart (1).

4. RISK-AWARE MMSE ESTIMATORS

In our development, we exploit a variational version of Slater's condition, which is one the most widely used constraint qualifications in both deterministic and stochastic optimization.

Assumption 2. Given $\varepsilon > 0$, problem (4) satisfies Slater's condition, i.e., there exists $\widehat{\mathbf{X}}_{\dagger} \in \mathcal{L}_{2|\mathcal{Y}}^M$, such that $\mathbb{E}\{\|\mathbf{X} - \widehat{\mathbf{X}}_{\dagger}\|_2^2\} < \infty$ and $\mathcal{L}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X} - \widehat{\mathbf{X}}_{\dagger}\|_2^2\}\} < \varepsilon$.

Under both Assumptions 1 and 2, it follows that the QCQP (6) satisfies Slater's condition, as well. Then, it must be the case that $\mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X}\|_2^2\}\} < \infty$; if not, Assumption 2 is impossible to hold. Further, problem (6) must be feasible, with convex effective domain

$$\mathcal{F}_{2|\mathcal{Y}}^M \triangleq \{\widehat{\mathbf{X}} \in \mathcal{L}_{2|\mathcal{Y}}^M \mid \mathbb{E}\{\widehat{\mathbf{X}}^T \Sigma_{\mathbf{X}|\mathbf{Y}} \widehat{\mathbf{X}}\} < \infty\}. \quad (7)$$

Next, if Assumption 1 holds, define the *variational Lagrangian* of the primal problem (6) $\mathcal{L}: \mathcal{L}_{2|\mathcal{Y}}^M \times \mathbb{R}_+ \rightarrow (-\infty, \infty]$ as

$$\begin{aligned} \mathcal{L}(\widehat{\mathbf{X}}, \mu) &\triangleq \frac{1}{2} \mathbb{E}\{\|\widehat{\mathbf{X}}\|_2^2\} - 2(\mathbb{E}\{\mathbf{X}|\mathbf{Y}\})^T \widehat{\mathbf{X}} + \mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \\ &\quad + \mu \mathbb{E}\{\widehat{\mathbf{X}}^T \Sigma_{\mathbf{X}|\mathbf{Y}} \widehat{\mathbf{X}} - (\mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \\ &\quad - \mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})^T \widehat{\mathbf{X}}\} \\ &\quad - \mu \frac{\varepsilon - \mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X}\|_2^2\}\}}{4}, \end{aligned} \quad (8)$$

where $\mu \in \mathbb{R}_+$ is a multiplier associated with the constraint of (6). The dual function $\mathcal{D}: \mathbb{R}_+ \rightarrow [-\infty, \infty]$ is accordingly defined as

$$\mathcal{D}(\mu) \triangleq \inf_{\widehat{\mathbf{X}} \in \mathcal{F}_{2|\mathcal{Y}}^M} \mathcal{L}(\widehat{\mathbf{X}}, \mu). \quad (9)$$

If $\mathcal{P}^* \in [0, \infty]$ denotes the optimal value of problem (6), it is true that $\mathcal{D} \leq \mathcal{P}^*$ on \mathbb{R}_+ . Then, the optimal value of the dual problem

$$\begin{aligned} &\text{maximize} && \mathcal{D}(\mu) \\ &\text{subject to} && \mu \geq 0 \end{aligned}, \quad (10)$$

$\mathcal{D}^* \triangleq \sup_{\mu \geq 0} \mathcal{D}(\mu) \in [-\infty, \infty]$, is the tightest under-estimate of \mathcal{P}^* , when knowing only \mathcal{D} .

Exploiting Assumptions 1 and 2, we may now formulate the following fundamental result. Although not presented here, the proof invokes ([27], Section 8.3, Theorem 1) (for instance).

Theorem 1. (QCQP (6): Zero Duality Gap) Suppose that Assumptions 1 and 2 are in effect. Then, strong duality holds for problem (6), that is, $0 \leq \mathcal{D}^* \equiv \mathcal{P}^* < \infty$. Additionally, the set of dual optimal solutions, $\arg \max_{\mu \geq 0} \mathcal{D}(\mu)$, is nonempty. Further, if $\widehat{\mathbf{X}}_*$ is primal optimal for (6), it follows that $\widehat{\mathbf{X}}_* \equiv \widehat{\mathbf{X}}_*(\mu_*) \in \arg \min_{\widehat{\mathbf{X}} \in \mathcal{F}_{2|\mathcal{Y}}^M} \mathcal{L}(\widehat{\mathbf{X}}, \mu_*)$, where $0 \leq \mu_* \in \arg \max_{\mu \geq 0} \mathcal{D}(\mu)$.

Leveraging Theorem 1, it is possible to show that, under Assumptions 1 and 2, the QCQP (6) and, therefore, the original \mathcal{L}_2 risk-aware MMSE problem (4), admit a common closed form solution. In this respect, we have the next theorem, which constitutes the main result of this paper.

Theorem 2. (QCQP (6): Closed-Form Solution) Suppose that Assumptions 1 and 2 are in effect. Then, an optimal solution to problem (6) may be expressed as

$$\widehat{\mathbf{X}}_*(\mu_*) \equiv (\mathbf{I} + 2\mu_* \Sigma_{\mathbf{X}|\mathbf{Y}})^{-1} (\mathbb{E}\{\mathbf{X}|\mathbf{Y}\} + \mu_* (\mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} - \mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})), \quad (11)$$

with $\widehat{\mathbf{X}}_*(\mu_*) \in \arg \min_{\widehat{\mathbf{X}} \in \mathcal{F}_{2|\mathcal{Y}}^M} \mathcal{L}(\widehat{\mathbf{X}}, \mu_*)$, and where $\mu_* \in \mathbb{R}_+$ is an optimal solution to the concave dual problem

$$\begin{aligned} \sup_{\mu \geq 0} \mathcal{D}(\mu) &\equiv \sup_{\mu \geq 0} \inf_{\widehat{\mathbf{X}} \in \mathcal{F}_{2|\mathcal{Y}}^M} \mathcal{L}(\widehat{\mathbf{X}}, \mu) \\ &\equiv \frac{1}{2} \mathbb{E}\{\|\mathbf{X}\|_2^2\} + \frac{1}{4} \sup_{\mu \geq 0} \left\{ \mu \mathbb{E}\{\mathbb{V}_{\mathbf{Y}}\{\|\mathbf{X}\|_2^2\}\} \right. \\ &\quad \left. - 2\mathbb{E}\{\widehat{\mathbf{X}}_*^T(\mu)(\mathbf{I} + 2\mu \Sigma_{\mathbf{X}|\mathbf{Y}}) \widehat{\mathbf{X}}_*(\mu)\} - \mu \varepsilon \right\}. \end{aligned} \quad (12)$$

Additionally, the optimal risk-aware filter $\widehat{\mathbf{X}}_*(\mu_*)$ is unique, almost everywhere relative to \mathcal{P} .

Theorem 2 completely solves problem (4) by providing a closed-form expression for the risk-aware MMSE estimator $\widehat{\mathbf{X}}_*$, defined in terms of the dual optimal solution μ_* (a number). The latter *always* exists, thanks to Theorem 1, and may be computed by leveraging our knowledge of the distribution $\mathcal{P}_{(\mathbf{X}, \mathbf{Y})}$ and via either some gradient-based method, or even empirically. Note that the dual function \mathcal{D} is merely a concave function on the positive line.

The fact that a closed-form optimal solution to problem (6) exists is remarkable, and it provides insight into the intrinsic structure of *constrained* Bayesian risk-aware estimation, also enabling an explicit comparison of the optimal risk-aware filter $\widehat{\mathbf{X}}_*$ with its risk-neutral counterpart. Indeed, by looking at the explicit form of the optimal risk-aware filter $\widehat{\mathbf{X}}_*$, we readily see that it is a function involving the MMSE estimator $\mathbb{E}\{\mathbf{X}|\mathbf{Y}\}$, its predictive covariance matrix $\Sigma_{\mathbf{X}|\mathbf{Y}}$, as well as the second and third order filters $\mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\}$ and $\mathbb{E}\{\|\mathbf{X}\|_2^2|\mathbf{Y}\} \mathbb{E}\{\mathbf{X}|\mathbf{Y}\}$. All these quantities are elementary and, in principle, they can be evaluated by utilizing a single observation of \mathbf{Y} , and by exploiting our knowledge of the conditional measure $\mathcal{P}_{\mathbf{X}|\mathbf{Y}}$, just as in risk-neutral MMSE estimation.

Also, we see that $\widehat{\mathbf{X}}_*$ may be regarded as a *biased MMSE estimator*, drawing parallels to *James-Stein estimators*, in another statistical context. Through the effect of bias, while James-Stein estimators achieve lower mean squared error, $\widehat{\mathbf{X}}_*$ achieves lower risk. Therefore, *optimal risk aversion, in the sense of problem (2), may be interpreted as the result of bias injection in MMSE estimators.*

Additionally, we observe that the solution is *regularized*, in the sense that the term $2\mu_* \Sigma_{\mathbf{X}|\mathbf{Y}}$ is diagonally loaded with an identity matrix; as a result, $\widehat{\mathbf{X}}_*$ is always well-defined and numerically stable. In fact, whenever $\mu_* \equiv 0$ (depending on the magnitude of the tolerance ε), it follows that $\widehat{\mathbf{X}}_* \equiv \mathbb{E}\{\mathbf{X}|\mathbf{Y}\}$. But this is not the only case where the two estimators turn out to be the same. The next result confirms that there exists a certain family of models for which risk-neutral and risk-aware MMSE estimation actually coincide; in such cases, posing (4) is redundant.

Theorem 3. (When Risk-Neutral/Aware Filters Coincide?) Suppose that the conditional measure $\mathcal{P}_{\mathbf{X}|\mathbf{Y}}$ is such that

$$\mathbb{E}\{(X_i - \mathbb{E}\{X_i|\mathbf{Y}\})^2 (\mathbf{X} - \mathbb{E}\{\mathbf{X}|\mathbf{Y}\})|\mathbf{Y}\} \equiv \mathbf{0}, \quad \forall i \in \mathbb{N}_M^+. \quad (13)$$

Then, in the notation of Theorem 2, it is true that, for every choice of $\mu \geq 0$, $\widehat{\mathbf{X}}_*(\mu) \equiv \mathbb{E}\{\mathbf{X}|\mathbf{Y}\}$. In other words, under Assumptions 1 and 2, risk-neutral MMSE estimation is also risk-aware, for every qualifying value of $\varepsilon > 0$. In particular, this is the case whenever $\mathcal{P}_{\mathbf{X}|\mathbf{Y}}$ is joint Gaussian.

In the next section, we put the risk-aware MMSE estimator to work, as well as numerically evaluate its performance in comparison with that of the usual, risk-neutral MMSE estimator.

5. NUMERICAL SIMULATIONS

We evaluate the behavior of the estimator in (11) in two different scenarios. The first consists of the problem of estimating an exponentially distributed hidden state X , $\mathbb{E}\{X\} = 2$, from the observation $Y = X + v$, where v is a zero-mean Gaussian random variable independent of X whose variance is given by $\mathbb{E}\{v^2\} = 9X^2$; in this case, v constitutes a state-dependent noise. In the second scenario, the goal is to jointly estimate the random vector $X = [z \ h]^T$ from the observation $Y = hz + w$, where z is a zero-mean Gaussian random variable with variance $\mathbb{E}\{z^2\} = 2$, h has a Rayleigh distribution with rate 2, and w is a zero-mean Gaussian noise with variance $\mathbb{E}\{w^2\} = 10^{-1}$. This scenario is prototypical for estimation problems in communications, where z is the signal of interest and h represents the channel fading. Throughout the simulations, we also show results for the risk-neutral MMSE estimator and the Minimum Mean Absolute Error (MMAE) estimator, or, equivalently, the conditional median relative to the respective observations, the latter being used as an example of a robust location parameter estimator.

In Fig. 1, we saw that the risk-aware estimator yields larger estimates than the MMSE estimator, in order to account for statistical ambiguities of the state-dependent noise model. Though this difference may seem extreme in some instances, e.g., for small values of Y (as in Fig. 1), it is in fact quite effective in reducing the conditional variance. Indeed, for $Y = 0.1$, the risk-aware estimator in Fig. 1 optimally reduces the (conditional) risk by approximately 26% as compared to the risk of the risk-neutral estimator; this is achieved by sacrificing average performance, also by a factor of 26%. Of course, this is only one of the operation points of the risk-aware estimator. In Fig. 2, we show results for different values of μ , where we average over the distribution of Y . Observe that the risk-aware estimator obtained via the constrained problem (2) achieves a sharp trade-off between average performance (i.e., MSE) and risk, which can be tuned according to the needs of the application. Further, note that the decrease in risk is considerably faster than the increase in MSE.

Interestingly, a similar phenomenon is observed in the communication scenario (Fig. 3). Again, the estimator risk displays a much faster initial rate of decrease with respect to μ than the rate at which the MSE increases. This is more pronounced in the estimation of the component z , for which the risk-aware estimator can provide reductions of almost 60% in risk for a 35% increase in average squared error. Note that, as per Theorem 3, the Gaussian noise has indeed no bearing on risk-awareness, as evidenced by the performance in the noiseless case, i.e., for $w = 0$ (dashed lines). To achieve the behavior of Fig. 3, the risk-aware estimator overestimates both z and h as compared to the MMSE estimator, as illustrated in Fig. 4. In fact, for small values of Y , the former hedges against the event of a deep fade ($h \approx 0$) by maintaining its estimates for z away from zero.

6. CONCLUSIONS

We derived a risk-aware MMSE estimator that accounts for statistical model volatility by hedging against extreme losses. We did so by formulating a Bayesian risk-aware MMSE estimation problem that minimizes squared errors on average, subject to an explicit tolerance on their expected predictive variance. We then showed that this problem admits an analytical solution under mild moment boundedness assumptions, and results in a risk-aware, biased nonlinear MMSE estimator. The effectiveness of our approach was confirmed via several numerical simulations. Future work includes further analysis of the properties of the proposed estimator, as well as study of other constrained risk-aware formulations.




Fig. 2. Mean squared error and risk for different values of μ in the state-dependent noise scenario.




Fig. 3. Mean squared error and risk for different values of μ in the communication scenario.




Fig. 4. Risk-aware, MMSE, and robust estimates of z and h in the communication scenario for different values of Y .

7. REFERENCES

- [1] D. S. Kalogerias, L. F. Chamon, G. J. Pappas, and A. Ribeiro, "Risk-Aware MMSE Estimation," *arXiv preprint, arXiv:1912.02933*, 2019.
- [2] A. Ahlén, J. Akerberg, M. Eriksson, A. L. F. J. Isaksson, T. Iwaki, K. H. Johansson, S. Knorn, T. Lindh, and H. Sandberg, "Toward Wireless Control in Industrial Process Automation: A Case Study at a Paper Mill," *Control Systems, IEEE*, vol. 39, no. 5, pp. 36–57, 2019.
- [3] S. Bruno, S. Ahmed, A. Shapiro, and A. Street, "Risk Neutral and Risk-Averse Approaches to Multistage Renewable Investment Planning under Uncertainty," *European Journal of Operational Research*, vol. 250, no. 3, pp. 979–989, may 2016.
- [4] S. Moazeni, W. B. Powell, and A. H. Hajimiragha, "Mean-Conditional Value-at-Risk Optimal Energy Storage Operation in the Presence of Transaction Costs," *IEEE Transactions on Power Systems*, vol. 30, no. 3, pp. 1222–1232, may 2015.
- [5] H. Markowitz, "Portfolio Selection," *The Journal of Finance*, vol. 7, no. 1, pp. 77–91, mar 1952.
- [6] H. Föllmer and A. Schied, "Convex Measures of Risk and Trading Constraints," *Finance and Stochastics*, vol. 6, no. 4, pp. 429–447, oct 2002.
- [7] D. Shang, V. Kuzmenko, and S. Uryasev, "Cash Flow Matching with Risks Controlled by Buffered Probability of Exceedance and Conditional Value-at-Risk," *Annals of Operations Research*, vol. 260, no. 1-2, pp. 501–514, jan 2018.
- [8] S.-K. Kim, R. Thakker, and A.-a. Agha-mohammadi, "Bi-directional Value Learning for Risk-aware Planning Under Uncertainty," *IEEE Robotics and Automation Letters*, vol. 4, no. 3, pp. 2493–2500, jul 2019.
- [9] A. A. Pereira, J. Binney, G. A. Hollinger, and G. S. Sukhatme, "Risk-Aware Path Planning for Autonomous Underwater Vehicles using Predictive Ocean Models," *Journal of Field Robotics*, vol. 30, no. 5, pp. 741–762, sep 2013.
- [10] A. S. Bedi, A. Koppel, and K. Rajawat, "Nonparametric Compositional Stochastic Optimization," *arXiv preprint, arXiv:1902.06011*, feb 2019.
- [11] W.-J. Ma, C. Oh, Y. Liu, D. Dentcheva, and M. M. Zavlanos, "Risk-Averse Access Point Selection in Wireless Communication Networks," *IEEE Transactions on Control of Network Systems*, vol. 5870, no. c, pp. 1–1, 2018.
- [12] P. L. A. and M. Fu, "Risk-Sensitive Reinforcement Learning: A Constrained Optimization Viewpoint," *arXiv preprint, arXiv:1810.09126*, oct 2018.
- [13] A. R. Cardoso and H. Xu, "Risk-Averse Stochastic Convex Bandit," in *International Conference on Artificial Intelligence and Statistics*, apr 2019, vol. 89, pp. 39–47.
- [14] W. Huang and W. B. Haskell, "Risk-Aware Q-learning for Markov Decision Processes," in *2017 IEEE 56th Annual Conference on Decision and Control, CDC 2017*, dec 2018, vol. 2018-Janua, pp. 4928–4933, IEEE.
- [15] D. R. Jiang and W. B. Powell, "Risk-Averse Approximate Dynamic Programming with Quantile-Based Risk Measures," *Mathematics of Operations Research*, vol. 43, no. 2, pp. 554–579, nov 2018.
- [16] D. S. Kalogerias and W. B. Powell, "Recursive Optimization of Convex Risk Measures: Mean-Semideviation Models," *arXiv preprint, arXiv:1804.00636*, apr 2018.
- [17] A. Tamar, Y. Chow, M. Ghavamzadeh, and S. Mannor, "Sequential Decision Making with Coherent Risk," *IEEE Transactions on Automatic Control*, vol. 62, no. 7, pp. 3323–3338, jul 2017.
- [18] C. A. Vitt, D. Dentcheva, and H. Xiong, "Risk-Averse Classification," *Annals of Operations Research*, aug 2019.
- [19] L. Zhou and P. Tokekar, "An Approximation Algorithm for Risk-averse Submodular Optimization," *arXiv preprint, arXiv:1807.09358*, jul 2018.
- [20] A. Shapiro, D. Dentcheva, and A. Ruszczyński, *Lectures on Stochastic Programming: Modeling and Theory*, Society for Industrial and Applied Mathematics, 2nd edition, 2014.
- [21] R. T. Rockafellar and S. Uryasev, "Optimization of Conditional Value-at-Risk," *Journal of Risk*, vol. 2, pp. 21–41, 1997.
- [22] P. Whittle, "Risk-Sensitive Linear/Quadratic/Gaussian Control," *Adv. Appl. Prob.*, vol. 13, pp. 764–777, 1981.
- [23] J. Speyer, C.-H. Fan, and R. Banavar, "Optimal Stochastic Estimation with Exponential Cost Criteria," in *Proceedings of the 31st IEEE Conference on Decision and Control*, aug 1992, pp. 2293–2298, Institute of Electrical and Electronics Engineers (IEEE).
- [24] J. B. Moore, R. J. Elliott, and S. Dey, "Risk-Sensitive Generalizations of Minimum Variance Estimation and Control," *Journal of Mathematical Systems, Estimation, and Control*, vol. 7, no. 1, pp. 123–126, 1997.
- [25] S. Dey and J. B. Moore, "Risk-Sensitive Filtering and Smoothing via Reference Probability Methods," *IEEE Transactions on Automatic Control*, vol. 42, no. 11, pp. 1587–1591, 1997.
- [26] S. Dey and J. B. Moore, "Finite-Dimensional Risk-Sensitive Filters and Smoothers for Discrete-Time Nonlinear Systems," *IEEE Transactions on Automatic Control*, vol. 44, no. 6, pp. 1234–1239, 1999.
- [27] D. G. Luenberger, *Optimization by Vector Space Methods*, Wiley, 1968.