ON UNBIASED ESTIMATION OF SPARSE VECTORS CORRUPTED BY GAUSSIAN NOISE

Alexander Jung[a], Zvika Ben-Haim[b], Franz Hlawatsch[a], and Yonina C. Eldar[b]

[a]Institute of Communications and Radio-Frequency Engineering, Vienna University of Technology
Gusshausstrasse 25/389, A-1040 Vienna, Austria; e-mail: {ajung, fhlawats}@nt.tuwien.ac.at

[b]Technion—Israel Institute of Technology, Haifa 32000, Israel; e-mail: {zvikabh@tx, yonina@ee}.technion.ac.il

ABSTRACT

We consider the estimation of a sparse parameter vector from measurements corrupted by white Gaussian noise. Our focus is on unbiased estimation as a setting under which the difficulty of the problem can be quantified analytically. We show that there are infinitely many unbiased estimators but none of them has uniformly minimum mean-squared error. We then provide lower and upper bounds on the Barankin bound, which describes the performance achievable by unbiased estimators. These bounds are used to predict the threshold region of practical estimators.

Index Terms— Unbiased estimation, sparsity, denoising, Cramér–Rao bound, Barankin bound

1. INTRODUCTION

We consider the estimation of a parameter vector $\mathbf{x}_0 \in \mathbb{R}^N$ which is known to be S-sparse, i.e., at most $S \geq 1$ of its entries are nonzero, where typically $S \ll N$. This will be written as $\mathbf{x}_0 \in \mathcal{X}_S \triangleq \{\mathbf{x} \in \mathbb{R}^N : \|\mathbf{x}\|_0 \leq S\}$, where $\|\mathbf{x}\|_0$ denotes the number of nonzero elements of $\mathbf{x}$. Our focus is the frequentist estimation setting in which the parameter $\mathbf{x}_0$ is modeled as unknown but deterministic [1]. While the sparsity S is assumed known, the positions of the nonzero entries of $\mathbf{x}_0$ are unknown. The S-sparse vector $\mathbf{x}_0$ is corrupted by white Gaussian noise with known variance σ^2, so that the observation is

$$\mathbf{y} = \mathbf{x}_0 + \mathbf{n}, \quad \text{with } \mathbf{x}_0 \in \mathcal{X}_S, \ \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_N). \quad (1)$$

This will be called the *sparse signal-in-noise model* (SSNM).

In this paper, we are interested in unbiased estimators, i.e., estimators $\hat{\mathbf{x}}(\mathbf{y})$ which satisfy $E_{\mathbf{x}}\{\hat{\mathbf{x}}(\mathbf{y})\} = \mathbf{x}$ for all $\mathbf{x} \in \mathcal{X}_S$. (The notation $E_{\mathbf{x}}$ indicates that the expectation uses the distribution of $\mathbf{y}$ induced by the parameter value $\mathbf{x}$.) The assumption of unbiasedness is commonly applied when analyzing estimation performance [1, 2]. Apart from being intuitively appealing, unbiased estimation is known to be optimal when the signal-to-noise ratio (SNR) is high [1]. More practically, to analyze achievable performance, the class of estimators must be restricted in some way in order to exclude approaches such

as $\hat{\mathbf{x}}(\mathbf{y}) \equiv \mathbf{x}_1$, which yield zero error at the specific point $\mathbf{x} = \mathbf{x}_1$ but are evidently undesirable. Consequently, even though many estimators for the SSNM are biased [3], we adopt the commonly used unbiasedness assumption in order to gain insight into the achievable performance in this setting. Note that while $\mathbf{x}_0$ is known to be S-sparse, the estimates $\hat{\mathbf{x}}$ we consider are not constrained to be S-sparse.

Our goal is to characterize the minimum mean-squared error (MSE) achievable by unbiased estimators. This optimal MSE, known as the *Barankin bound* (BB) [4], cannot be calculated analytically. Instead, we prove that the BB is itself bounded above and below by simple, closed-form expressions. The upper bound is obtained by solving a constrained optimization problem, while the lower bound is based on the Hammersley–Chapman–Robbins technique [5]. In a numerical study, we demonstrate the ability of our bounds to roughly predict the threshold SNR of practical estimators.

A well-known lower bound on the BB is the Cramér–Rao bound (CRB), which was recently derived for the sparse setting [2, 6]. To obtain the CRB, it is sufficient to assume that the estimator $\hat{\mathbf{x}}(\mathbf{y})$ is unbiased only at a given parameter value $\mathbf{x}_0 \in \mathcal{X}_S$ and in its local neighborhood. By contrast, in this paper we seek to characterize the MSE obtainable by estimators which are unbiased for *all* values $\mathbf{x} \in \mathcal{X}_S$. One may hope to achieve a higher—hence, tighter—lower bound by adopting this more restrictive constraint. Furthermore, the CRB is discontinuous in $\mathbf{x}$ [6], whereas the MSE of an estimator for the SSNM is always continuous [1]; this again suggests that it may be possible to find improved lower bounds.

We note in passing that our results can be easily generalized to the sparse orthonormal signal model

$$\mathbf{y} = \mathbf{H}\mathbf{x}_0 + \mathbf{n}, \quad \text{with } \mathbf{x}_0 \in \mathcal{X}_S, \ \mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_M), \quad (2)$$

where $\mathbf{y} \in \mathbb{R}^M$ with $M \geq N$ and the known matrix $\mathbf{H} \in \mathbb{R}^{M \times N}$ has orthonormal columns, i.e., $\mathbf{H}^T \mathbf{H} = \mathbf{I}_N$. However, for simplicity of notation, we will continue to refer to the model (1). This setting arises, for example, in image denoising based on the assumption of sparsity in the wavelet domain.

2. UNBIASED ESTIMATION

We will denote the set of unbiased estimators by

$$\mathcal{U} \triangleq \{\hat{\mathbf{x}}(\cdot) : E_{\mathbf{x}}\{\hat{\mathbf{x}}(\mathbf{y})\} = \mathbf{x} \text{ for all } \mathbf{x} \in \mathcal{X}_S\}.$$

Before analyzing bounds on the performance of unbiased estimators, we must ask whether such techniques exist at all. This is answered affirmatively in the following theorem. Due to space limitations, the proof of this and subsequent theorems is omitted and will appear in [7]. The MSE $\mathrm{E}_{\mathbf{x}}\{\|\hat{\mathbf{x}}(\mathbf{y}) - \mathbf{x}\|_2^2\}$ of a given estimator $\hat{\mathbf{x}}(\cdot)$ at the parameter value $\mathbf{x}$ will be denoted $\varepsilon(\mathbf{x}; \hat{\mathbf{x}})$.

Theorem 1 *For $S = N$ in the SSNM* (1)*, there is only one unbiased estimator of $\mathbf{x}_0$ (up to modifications having zero measure) which has a bounded MSE. This is the ordinary least-squares (LS) estimator,*

$$\hat{\mathbf{x}}_{\mathrm{LS}}(\mathbf{y}) \triangleq \arg \min_{\mathbf{x} \in \mathbb{R}^N}\{\|\mathbf{y} - \mathbf{x}\|_2^2\} = \mathbf{y}, \qquad (3)$$

whose MSE is $\varepsilon(\mathbf{x}; \hat{\mathbf{x}}_{\mathrm{LS}}) \equiv N\sigma^2$. By contrast, for $S < N$, there are infinitely many unbiased estimators of $\mathbf{x}_0$.

The assumption of a bounded MSE means that there exists a constant C such that $\varepsilon(\mathbf{x}; \hat{\mathbf{x}}) \leq C$ for all $\mathbf{x}$. This assumption is required to exclude certain pathological functions which do not make sense as estimators. Since for $S = N$ the only unbiased estimator with bounded MSE is the LS estimator, we will hereafter assume that $S < N$.

3. OPTIMAL UNBIASED ESTIMATION

Given a fixed parameter value $\mathbf{x}_0 \in \mathcal{X}_S$, one can define an optimal unbiased estimator $\hat{\mathbf{x}}^{(\mathbf{x}_0)}(\cdot)$ for the SSNM (1) as an unbiased estimator which minimizes the MSE at $\mathbf{x}_0$, i.e.,

$$\varepsilon(\mathbf{x}_0; \hat{\mathbf{x}}^{(\mathbf{x}_0)}) = \min_{\hat{\mathbf{x}}(\cdot) \in \mathcal{U}} \mathrm{E}_{\mathbf{x}_0}\{\|\hat{\mathbf{x}}(\mathbf{y}) - \mathbf{x}_0\|_2^2\}. \qquad (4)$$

Solving (4) is equivalent to solving the N individual optimization problems

$$\min_{\hat{x}_k(\cdot) \in \mathcal{U}'} \mathrm{E}_{\mathbf{x}_0}\{|\hat{x}_k(\mathbf{y}) - x_{0,k}|^2\}, \quad k = 1, ..., N, \qquad (5)$$

where $\hat{x}_k(\cdot)$ and $x_{0,k}$ are the kth elements of $\hat{\mathbf{x}}(\cdot)$ and $\mathbf{x}_0$, respectively, and $\mathcal{U}'$ denotes the set of unbiased estimators of the kth element of $\mathbf{x}$.

Since the MSE $\varepsilon(\mathbf{x}; \hat{\mathbf{x}})$ of an estimator $\hat{\mathbf{x}}(\cdot)$ depends on the true parameter value $\mathbf{x} \in \mathcal{X}_S$, an unbiased estimator that has a small MSE for one parameter value $\mathbf{x}_0 \in \mathcal{X}_S$ may have a large MSE for other parameter values in $\mathcal{X}_S$. Sometimes, however, a single estimator solves (4) simultaneously for all $\mathbf{x} \in \mathcal{X}_S$. Such a technique is called a *uniformly minimum variance unbiased* (UMVU) estimator [1]. (Note that for an unbiased estimator the variance equals the MSE, so minimum MSE implies minimum variance.) However, for the SSNM (1), the following negative result can be shown [7].

Theorem 2 *For the SSNM with $S < N$, there does not exist a UMVU estimator.*

Thus, no single unbiased estimator achieves minimum MSE for all values of $\mathbf{x} \in \mathcal{X}_S$. We can still try to find a *locally minimum variance unbiased* (LMVU) estimator by solving

(4) for a *specific* $\mathbf{x}_0 \in \mathcal{X}_S$. This approach yields the unbiased estimator which has minimum MSE for the given parameter value $\mathbf{x}_0$. While the resulting technique may have poor performance for other parameter values, its MSE at $\mathbf{x}_0$ is of interest as it provides a lower bound on the MSE at $\mathbf{x}_0$ of any unbiased estimator. This lower bound is known as the BB [4]. Because it is achieved, by some technique, at each point $\mathbf{x}_0$, the BB is the largest (thus, tightest) lower bound for unbiased estimators. Unfortunately, an analytical expression of the BB is usually not obtainable. Instead, in the following two subsections, we will provide lower and upper bounds on the BB.

3.1. Lower Bound on the MSE

A variety of techniques exist for developing lower bounds on the MSE of unbiased estimators. The simplest of these is the CRB, which was derived for a more general sparse estimation setting in [2, 6]. In our case, the CRB is given by

$$\varepsilon(\mathbf{x}_0; \hat{\mathbf{x}}) \geq \begin{cases} S\sigma^2, & \|\mathbf{x}_0\|_0 = S \\ N\sigma^2, & \|\mathbf{x}_0\|_0 < S. \end{cases} \qquad (6)$$

The CRB assumes unbiasedness only in a neighborhood of $\mathbf{x}_0$. Since we are interested in estimators which are unbiased for all $\mathbf{x} \in \mathcal{X}_S$, we can expect to obtain higher bounds. One such approach is the Hammersley–Chapman–Robbins bound (HCRB) [5], which can be formulated, in our context, as follows. For a given parameter value $\mathbf{x}_0$, consider a set of p "test points" $\{\mathbf{v}_i\}_{i=1}^p$ such that $\mathbf{v}_i + \mathbf{x}_0 \in \mathcal{X}_S$. The HCRB states that the MSE of any unbiased estimator $\hat{\mathbf{x}}(\cdot) \in \mathcal{U}$ satisfies

$$\varepsilon(\mathbf{x}_0; \hat{\mathbf{x}}) \geq \mathrm{Tr}(\mathbf{V}\mathbf{J}^\dagger\mathbf{V}^T),$$

where $\mathbf{V} \triangleq [\mathbf{v}_1 \cdots \mathbf{v}_p] \in \mathbb{R}^{N \times p}$, the elements of the matrix $\mathbf{J} \in \mathbb{R}^{p \times p}$ are given by $(\mathbf{J})_{i,j} \triangleq \exp(\mathbf{v}_i^T \mathbf{v}_j / \sigma^2) - 1$, and $\mathbf{J}^\dagger$ denotes the pseudoinverse of $\mathbf{J}$.

Both the test points $\mathbf{v}_i$ and their number p are arbitrary and can depend on $\mathbf{x}_0$. Clearly, the challenge is to choose test points which result in a tight but analytically tractable bound. To this end, we note the following facts. First, it can be shown [7] that the CRB is obtained as a limit of HCRBs by choosing the test points $\mathbf{v}_i$ as[1]

$$\{t\mathbf{e}_i\}_{i \in \mathrm{supp}(\mathbf{x}_0)}, \quad \text{if } \|\mathbf{x}_0\|_0 = S \qquad (7a)$$

$$\{t\mathbf{e}_i\}_{i \in \{1,...,N\}}, \quad \text{if } \|\mathbf{x}_0\|_0 < S \qquad (7b)$$

and letting $t \to 0$. Here, $\mathbf{e}_i$ represents the ith column of the $N \times N$ identity matrix and $\mathrm{supp}(\mathbf{x}_0)$ denotes the set of indices of all nonzero elements of $\mathbf{x}_0$. Second, the CRB (6) is tight whenever $\|\mathbf{x}_0\|_0 < S$, since in this case the bound is achieved by the LS estimator (3). The CRB is also tight when $\|\mathbf{x}_0\|_0 = S$ and all components of $\mathbf{x}_0$ are much larger (in magnitude) than σ [6]. Any laxity in the CRB will therefore arise when $\mathbf{x}_0$ contains exactly S nonzero components, one or more of which are small. For such $\mathbf{x}_0$, it follows from (7a) that only S test points are employed; yet when the small components

---

[1]Hereafter, with a slight abuse of notation, the index i of $\mathbf{v}_i$ is allowed to take on p different values from the set $\{1, \ldots, N\}$.

of $\mathbf{x}_0$ are replaced with zero, then according to (7b), a much larger number N of test points is used. In light of this, one could attempt to improve the CRB by increasing the number of test points when $\|\mathbf{x}_0\|_0 = S$. This should be done such that when the smallest component tends to zero, the test points converge to (7b), for which the CRB is tight. In doing so, one still needs to ensure that $\mathbf{v}_i + \mathbf{x}_0 \in \mathcal{X}_S$.

A reasonable choice which satisfies these requirements is

$$\mathbf{v}_i = \begin{cases} t\mathbf{e}_i, & i \in \mathrm{supp}(\mathbf{x}_0) \\ -x_0^{(S)}\mathbf{e}_k + t\mathbf{e}_i, & i \notin \mathrm{supp}(\mathbf{x}_0) \end{cases} \quad (8)$$

for $i = 1, \dots, N$. Here, $x_0^{(S)}$ is the S-largest (in magnitude) entry of $\mathbf{x}_0$ (note that $x_0^{(S)} = 0$ if $\|\mathbf{x}_0\|_0 < S$). The test points (7), which yield the CRB, are a subset of the choice of test points in (8). Consequently, it can be shown that the resulting bound will always be at least as tight as the CRB.

In analogy to the CRB, a simple lower bound can be obtained from (8) by letting $t \to 0$. Some rather tedious calculations [7] yield the following result.

Theorem 3 *The MSE of any unbiased estimator $\hat{\mathbf{x}}(\cdot) \in \mathcal{U}$ is bounded below by*

$$\varepsilon(\mathbf{x}_0; \hat{\mathbf{x}}) \geq \begin{cases} S\sigma^2 + (N{-}S{-}1)\sigma^2 e^{-\xi_0^2/\sigma^2}, & \|\mathbf{x}_0\|_0 = S \\ N\sigma^2, & \|\mathbf{x}_0\|_0 < S, \end{cases}$$
(9)
where $\xi_0 \triangleq \min_{k \in \mathrm{supp}(\mathbf{x}_0)} |x_{0,k}|$.

Strictly speaking, this bound is obtained as a limit of HCRBs. However, for simplicity we will continue to refer to (9) as an HCRB. The bound is higher (tighter) than the CRB (6), except when the CRB itself is already tight. Furthermore, while both bounds are discontinuous in the transition between $\|\mathbf{x}_0\|_0 = S$ and $\|\mathbf{x}_0\|_0 < S$, the jump in (9) is much smaller than that in (6). This discontinuity can be eliminated altogether by adding more test points, but the resulting bound no longer has a simple closed form and can only be evaluated numerically.

3.2. Upper Bound on LMVU Performance

We next derive an upper bound on the BB by approximately solving (5) for $k = 1, \dots, N$. When $k \in \mathrm{supp}(\mathbf{x}_0)$, it can be shown [7] that the solution of (5) is given by the kth component of the LS estimator in (3), i.e.,

$$\hat{x}_k(\mathbf{y}) = y_k. \quad (10)$$

It remains to consider $k \notin \mathrm{supp}(\mathbf{x}_0)$. For such k, let us denote $\hat{x}_k(\mathbf{y}) = y_k + \hat{x}'_k(\mathbf{y})$. To obtain an upper bound on the solution of the optimization problem (5), we will restrict the set of allowed estimators $\hat{x}_k(\cdot)$ by requiring the correction component $\hat{x}'_k(\cdot)$ to satisfy the following two properties.

- *Odd symmetry* with respect to index k and all indices in $\mathrm{supp}(\mathbf{x}_0)$: For all $l \in \{k\} \cup \mathrm{supp}(\mathbf{x}_0)$,

$$\hat{x}'_k(\dots, -y_l, \dots) = -\hat{x}'_k(\dots, y_l, \dots). \quad (11)$$

- *Independence* with respect to all other indices: For all $l \notin \{k\} \cup \mathrm{supp}(\mathbf{x}_0)$,

$$\hat{x}'_k(\dots, y_l, \dots) = \hat{x}'_k(\dots, 0, \dots). \quad (12)$$

We introduce these constraints (which are rather ad-hoc) because under them a closed-form solution of (5) can be obtained. This constrained solution is given by [7]

$$\hat{x}_{\mathrm{c},k}(\mathbf{y}; \mathbf{x}_0) = y_k \left[1 - \prod_{l \in \mathrm{supp}(\mathbf{x}_0)} \tanh\left(\frac{y_l x_{0,l}}{\sigma^2}\right) \right] \quad (13)$$

for $\mathbf{y} \in \mathbb{R}_+^N$. For other values of $\mathbf{y}$, the solution can be obtained via the properties (11) and (12).

Combining (10) and (13), we obtain a constrained solution of (4) which will be denoted $\hat{\mathbf{x}}_{\mathrm{c}}(\mathbf{y}; \mathbf{x}_0)$. The MSE of $\hat{\mathbf{x}}_{\mathrm{c}}(\mathbf{y}; \mathbf{x}_0)$ at the parameter value $\mathbf{x}_0$ is the tightest lower bound on the MSE at $\mathbf{x}_0$ of any unbiased estimator that satisfies the constraints (11) and (12). It will thus be called the *constrained BB* and denoted by $\mathrm{BB}_{\mathrm{c}}(\mathbf{x}_0)$. Note that $\mathrm{BB}_{\mathrm{c}}(\mathbf{x}_0)$ is an upper bound on the BB. A closed-form expression of $\mathrm{BB}_{\mathrm{c}}(\mathbf{x}_0)$ is given in the next theorem.

Theorem 4 *The BB (i.e., the minimum MSE achievable by any unbiased estimator) at $\mathbf{x}_0$ is bounded above by*

$$\mathrm{BB}_{\mathrm{c}}(\mathbf{x}_0) = S\sigma^2 + (N{-}S)\,\sigma^2 \left[1 - \prod_{l \in \mathrm{supp}(\mathbf{x}_0)} g(x_{0,l}; \sigma^2) \right],$$
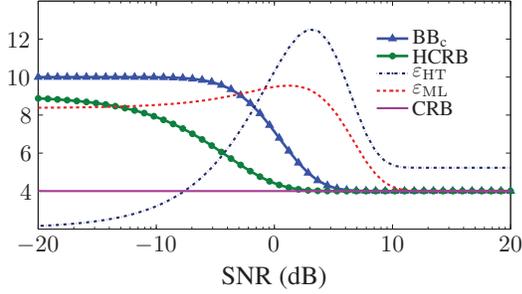
with

$$g(x; \sigma^2) \triangleq$$
$$\sqrt{\frac{2}{\pi\sigma^2}} \int_0^\infty e^{-(y^2+x^2)/(2\sigma^2)} \sinh\left(\frac{yx}{\sigma^2}\right) \tanh\left(\frac{yx}{\sigma^2}\right) dy.$$

We note that for all parameter values $\mathbf{x}_0$ with $\|\mathbf{x}_0\|_0 < S$, the lower and upper bounds of Theorem 3 and 4 coincide and equal $N\sigma^2$, which is therefore also the BB. This value is identical to the unconstrained CRB. We refer to [2, 6] for a discussion of this result.

4. NUMERICAL RESULTS

In Section 3, we bounded the achievable performance of unbiased estimators as a means for quantifying the difficulty of estimation in the SSNM. One use of this analysis is in the identification of the threshold region, a range of SNR values which constitutes a transition between low-noise and high-noise behavior. Specifically, the performance of estimators can often be calculated analytically when the SNR $\|\mathbf{x}_0\|_2^2/(N\sigma^2)$ is either very high or very low. It is then important to identify the threshold which separates these two regimes.

The lower and upper bounds on the BB which were derived above also exhibit a transition between a high-SNR region and a low-SNR region. In the high-SNR region, the lower bound (assuming $\|\mathbf{x}_0\|_0 = S$) and the upper bound both converge to $S\sigma^2$, while in the low-SNR region, both bounds are on the order of $N\sigma^2$. The true BB therefore also displays such a threshold. Since the BB is itself a lower bound on the MSE

Fig. 1. MSE of the HT and ML estimators compared with the performance bounds BB_c, HCRB and CRB (all normalized by σ^2), as a function of the SNR.

Fig. 2. Ratio $BB_c(\mathbf{x}_0)/\mathrm{HCRB}(\mathbf{x}_0)$ for different sets of parameter vectors $\mathbf{x}_0$.

of unbiased estimators, one would expect that the transition region of the BB occurs at slightly lower SNR values than that of actual estimators.

To test this hypothesis, we compared the bounds of Section 3 with the MSE of two well-known *biased* estimation schemes. First, we considered the maximum likelihood (ML) estimator, which can be shown to equal

$$\hat{x}_{\mathrm{ML},k}(\mathbf{y}) = \begin{cases} y_k, & \text{if } k \in \mathcal{L} \\ 0, & \text{else} \end{cases} \tag{14}$$

where $\mathcal{L}$ denotes the indices of the S largest (in magnitude) entries of $\mathbf{y}$. We also considered the hard-thresholding (HT) estimator

$$\hat{x}_{\mathrm{HT},k}(\mathbf{y}) = \begin{cases} y_k, & \text{if } |y_k| \geq T \\ 0, & \text{else} \end{cases} \tag{15}$$

with the commonly employed threshold $T = \sigma\sqrt{2\log N}$ [3].

We used $N = 10$ and $S = 4$, and chose 50 parameter vectors $\mathbf{x}_0$ from the set $\mathcal{R} \triangleq \{c\,(1,1,1,1,0,0,0,0,0,0)^T\}_{c\in\mathbb{R}_+}$, with different values of c to obtain a wide range of SNR values. The results are plotted in Fig. 1. Although there is some gap between the lower bound (HCRB) and the upper bound (BB_c), a rough indication of the behavior of the BB is conveyed. As expected, the SNR threshold predicted by these bounds is somewhat lower than that of practical estimators. Specifically, the transition region of the BB can be seen to occur at SNR values between -5 and 5 dB, while the transition of the ML and HT estimators is at SNR values between 0 and 10 dB. Another effect which is visible in Fig. 1 is the convergence of the ML estimator to the BB at high SNR; this is a consequence of the well-known fact that the ML approach is asymptotically unbiased and asymptotically minimizes the MSE at high SNR. Furthermore, our performance bounds in Fig. 1 suggest that at intermediate and high SNR (above 0 dB), there may exist unbiased estimators that outperform the ML and HT estimators. However, at low SNR, both estimators (ML and HT) are better than the best unbiased estimator. This agrees with the general rule that unbiased estimators perform poorly at low SNR.

One may argue that considering only parameter values in the set $\mathcal{R}$ is not representative, since $\mathcal{R}$ covers only a small part of the parameter space $\mathcal{X}_S$. However, it can be
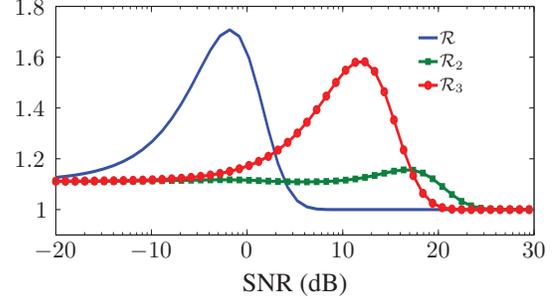
shown that the choice for $\mathcal{R}$ is conservative in that the maximum deviation between the HCRB and the BB_c is largest when the non-zero entries of $\mathbf{x}_0$ have approximately the same magnitude, which is the case for each element of $\mathcal{R}$. To illustrate this fact, we considered the two additional sets $\mathcal{R}_2 \triangleq \{c\,(0.1,1,1,1,0,0,0,0,0,0)^T\}_{c\in\mathbb{R}_+}$ and $\mathcal{R}_3 \triangleq \{c\,(10,1,1,1,0,0,0,0,0,0)^T\}_{c\in\mathbb{R}_+}$, in which the four non-zeros are not all equal. Fig. 2 depicts the ratio BB_c/HCRB versus the SNR for the three sets $\mathcal{R}$, $\mathcal{R}_2$, and $\mathcal{R}_3$. It is seen that the maximum value of BB_c/HCRB is indeed highest when $\mathbf{x}_0$ is in $\mathcal{R}$.

5. CONCLUSION

We considered unbiased estimation of a sparse parameter vector in white Gaussian noise. The Barankin bound (i.e., the minimum MSE achievable by any unbiased estimator) was characterized via upper and lower bounds. Our numerical results suggest that these bounds may also give information about the success of more general, biased techniques, for example by providing a lower bound on the threshold region. A subject for further study is the extension of these results to a model of the form $\mathbf{y} = \mathbf{A}\mathbf{x}_0 + \mathbf{n}$ in which $\mathbf{x}_0$ is sparse and $\mathbf{A}$ is an arbitrary matrix with full column rank.

6. REFERENCES

[1] E. L. Lehmann and G. Casella, *Theory of Point Estimation*, Springer, 2003.

[2] Z. Ben-Haim and Y. C. Eldar, "Performance bounds for sparse estimation with random noise," in *Proc. IEEE-SP Workshop on Statistical Signal Processing*, pp. 225–228, Cardiff, Wales, UK, Aug. 2009.

[3] S. G. Mallat, *A Wavelet Tour of Signal Processing*, Academic Press, 3rd edition, 2009.

[4] E. W. Barankin, "Locally best unbiased estimates," *Ann. Math. Statist.*, vol. 20, pp. 477–501, 1949.

[5] J. D. Gorman and A. O. Hero, "Lower bounds for parametric estimation with constraints," *IEEE Trans. Inf. Theory*, vol. 36, no. 6, pp. 1285–1301, Nov. 1990.

[6] Z. Ben-Haim and Y. C. Eldar, "The Cramér–Rao bound for sparse estimation," *IEEE Trans. Signal Processing*, May 2009, submitted. Available: http://arxiv.org/abs/0905.4378

[7] A. Jung, Z. Ben-Haim, F. Hlawatsch, and Y. C. Eldar, "Unbiased estimation of a sparse vector in white Gaussian noise," in preparation.