Received: xxx **Accepted:** xxx **Published:** xxx.
DOI: 10.30473/coam.2026.79769.1478
Volume xxx, Issue xxx, p.p. 1-30: xxx
Research Article

Open Access

Control and Optimization in Applied Mathematics - COAM

Adaptive Risk-Optimized Shrinkage Estimation of Normal Entropy under Target Uncertainty and Asymmetric Loss

Iman Jalil Atewi[1] ✉, **Alaa Khlaif Jiheel**[2]

[1]Department of Mathematics, University of Thi-Qar, Thi-Qar, Iraq
[2]Department of Mathematics, College of Education for Pure Sciences, University of Thi-Qar, Nasiriyah, Iraq

✉ **Correspondence:**
Iman Jalil Atewi
E-mail:
eymanjalil.math@utq.edu.iq

How to Cite
Jalil Atewi, I., Khlaif Jiheel, A. (2027). "Adaptive risk-optimized shrinkage estimation of normal entropy under target uncertainty and asymmetric loss". *Control and Optimization in Applied Mathematics*, X(x), 1–30. https://doi.org/10.30473/coam.2026.79769.1478

Abstract. Estimating the differential entropy of a normal distribution reduces to estimating the log-variance, but plug-in estimation is biased and target-based shrinkage can be sensitive to misspecification. We propose AROS, a smooth finite-sample shrinkage framework built on the exact log-Gamma law of the sample variance. A digamma correction provides an unbiased entropy baseline, and AROS-S chooses its tuning parameter by minimizing risk over a declared target-uncertainty set subject to a worst-case risk bound. Exact risks are evaluated under quadratic and asymmetric exponential loss, with independent quadrature and Monte Carlo validation. Under quadratic loss and exact target specification ($\xi = 0$), AROS-0 reduced risk by 51–53% for $n = 5 - 100$; this gain was local rather than uniform. For $b = 1$, AROS-S satisfied the safety bound and had lower worst-case quadratic risk than the reconstructed Huntsberger comparator across the examined sample sizes and significance levels. A Michelson speed-of-light illustration shows how an estimated historical variance can provide a provisional calibration of the target-uncertainty radius; bootstrap intervals describe variability under the reported resampling scheme, conditional on the historical target and tuning choices, and not estimation accuracy. The results support smooth target borrowing when a scientifically defensible uncertainty set is available, with automatic reversion to the baseline when safe shrinkage is not justified.

Keywords. Differential entropy, Shrinkage estimation, Constrained optimization, Mathematical programming, Asymmetric loss, Target uncertainty.

MSC. 62F10; 90C34.

©2026 by the authors. Lisensee PNU, Tehran, Iran. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution 4.0 International (CC BY4.0) (http:/creativecommons.org/licenses/by/4.0)

Coming proofs

1 Introduction

Differential entropy is an information-scale measure of uncertainty used in signal processing, statistical monitoring, and metrology. In a Gaussian model it is affine in the log-variance, so the logarithmic transformation determines the finite-sample bias and turns a multiplicative scale error into an additive estimation error. Entropy, rather than the variance alone, is therefore the relevant target when uncertainty is reported or compared on an information scale. For a variance σ^2, the inferential target is

$$H\left(\sigma^2\right) = \frac{1}{2}\log\left(2\pi e\sigma^2\right).$$

The unknown component is linear in the log-variance. This identity does not make estimation automatic: replacing the variance by a sample estimate introduces a finite-sample bias after the logarithmic transformation, and the ranking of competing entropy estimators depends on the loss used to evaluate the estimation error.

Gaussian entropy estimation has been studied as a decision problem [5, 10, 16]. This parametric line differs from nonparametric entropy estimators based on spacings or nearest neighbors [12, 23]. The present study deliberately exploits the exact normal log-variance law to obtain finite-sample risk results, rather than competing with distribution-free entropy estimators. For the univariate normal model, the problem is closely related to decision-theoretic variance estimation. Stein [19] established the inadmissibility of the usual variance estimator, the choice of the standardization factor is reviewed in [7], and later reviews summarize shrinkage and modified variance rules under several losses [13, 15]. Stein-type identities and the associated shrinkage estimators for normal and gamma models, which are closely related to the Gamma representation in Eq. (1), are surveyed in [14]. Bayes and empirical Bayes estimators of the normal variance under Stein's loss, which penalizes gross over- and underestimation equally on the ratio scale, were derived in [25], and a recent shrinkage family that balances proximity to a target estimator against precision under a balanced loss is given in [1] for the multivariate normal mean. Because normal entropy is affine in the log-variance, this scale is natural for exact finite-sample analysis.

Shrinkage incorporates historical or external scale information by trading variance against bias. Classical empirical-Bayes and bounded-parameter results show that the amount of shrinkage depends on the risk criterion and the parameter region [6, 8], while variance-specific rules derive the shrinkage intensity from explicit loss functions [11, 21, 22]. In this framework, the exact risk function conditions on a specified external log-variance target θ_0. An externally estimated historical variance is not a known population parameter: its estimate may vary between historical samples and may differ systematically from the current variance. The conditional risk at a fixed realized target is therefore distinct from the unconditional risk that averages over historical target estimation. The real-data example uses the former calculation and does not

Corrected proofs

claim that a historical confidence interval automatically provides a bound on target misspecification. Pre-test and adaptive procedures offer another route to target borrowing [17, 18]. Optimization-oriented double pre-test shrinkage estimators, whose design parameters are chosen under uncertainty, have also been studied for the Pareto reliability function under progressive censoring and precautionary loss [2]; that work concerns reliability rather than entropy. More generally, Armstrong et al. [3] control the worst-case risk under bounded misspecification, while Armstrong and Kolesár [4] use explicit bias bounds for honest inference. AROS-S shares this general robust-decision principle; its specific contribution is the exact log-Gamma normal-entropy formulation, a smooth discrepancy weight, and finite-sample risk under symmetric and asymmetric loss. The Huntsberger construction is therefore used as a benchmark, not as a claim that pre-test procedures are generally unprotected.

The present study develops a smooth target-adaptive framework on the exact log-variance scale. The term "risk-optimized" refers to optimizing η within the prescribed AROS weight family, not to global optimality over all shrinkage rules. AROS-S minimizes a risk criterion subject to a continuum worst-case bound. This is formally a scalar semi-infinite constrained problem; the contribution is its entropy-specific finite-sample formulation, not a new general semi-infinite programming algorithm.

The objectives are to: (i) derive the exact finite-sample risk and the target-centered improvement region; (ii) quantify how asymmetric loss changes that region; (iii) solve the safety-constrained tuning problem and identify when no shrinkage is selected; and (iv) assess operational selection and sensitivity to n, α, b, and the risk-aggregation criterion. Exact quadrature, independent recomputation, Monte Carlo checks, and a historical-data illustration are used to evaluate these objectives. The notation is collected in Table 1.

The remainder of the paper is structured as follows. Section 2 presents the model, the adaptive estimator, the safety-constrained optimization problem, and the numerical and validation procedures. Section 3 reports the exact risk results, the benchmark comparisons, the numerical validation, the operational selection rule, and the Michelson illustration. Section 4 discusses the findings, their limitations, and directions for further work, and Section 5 presents the conclusions.

2 Materials and Methods

2.1 Model and Exact Log-Variance Representation

Let $X_1, \ldots, X_n$ be independent observations from a normal distribution with unknown mean and variance, and write the log-variance as θ. The entropy is then

Corrected proofs

$$\theta = \log \sigma^2, \qquad H\left(\sigma^2\right) = C_H + \frac{1}{2}\theta, \qquad C_H = \frac{1}{2}\log(2\pi e).$$

The usual unbiased sample variance is

$$S^2 = \frac{1}{\nu}\sum_{i=1}^{n}\left(X_i - \overline{X}\right)^2, \qquad \nu = n - 1.$$

Standard normal theory gives the exact Gamma representation

$$q = \frac{\nu}{2}, \qquad G = \frac{\nu S^2}{2\sigma^2} \sim \text{Gamma}(q, 1), \tag{1}$$

and hence

$$\log S^2 = \theta + \log G - \log q.$$

For $G \sim \text{Gamma}(q, 1)$, $\mathbb{E}(\log G) = \psi(q)$ and $\text{Var}(\log G) = \psi_1(q)$, where ψ and ψ_1 are the digamma and trigamma functions. Define

$$m_\nu = \psi(q) - \log q, \qquad s_\nu^2 = \psi_1(q), \qquad s_\nu = \sqrt{\psi_1(q)} > 0.$$

Equation (1) implies $\mathbb{E}(\log S^2) = \theta + m_\nu$ and $\text{Var}(\log S^2) = s_\nu^2$. Subtracting m_ν gives the exact bias-corrected estimator of θ,

$$T_n = \log S^2 - m_\nu = \log S^2 - \psi(q) + \log q,$$

which satisfies $\mathbb{E}(T_n) = \theta$ exactly, not merely asymptotically, and $\text{Var}(T_n) = s_\nu^2$. The corresponding entropy estimator is therefore unbiased and has exact finite-sample variance:

$$\widehat{H}_U = C_H + \frac{1}{2}T_n, \qquad \mathbb{E}\left(\widehat{H}_U\right) = H, \qquad \text{Var}\left(\widehat{H}_U\right) = \frac{1}{4}\psi_1(q).$$

These identities require $n \geq 2$ and independent Gaussian observations with a common variance. The ordinary plug-in estimator is retained as a conventional comparator:

$$\widehat{H}_{PI} = C_H + \frac{1}{2}\log S^2.$$

The principal notation and the estimators used in the analysis are summarized in Table 1.

2.2 Standardized Pivotal Representation

Define the standardized corrected log-variance

$$Z_n = \frac{T_n - \theta}{s_\nu} = \frac{\log G - \psi(q)}{s_\nu}.$$

Corrected proofs

Table 1: Principal notation and estimators used in the analysis.

Symbol	Definition	Role
ν, q	$\nu = n-1; q = \nu/2$	Degrees-of-freedom parameters
θ	$\log \sigma^2$	Log-variance parameter
s_ν^2	$\psi_1(q); s_\nu = \sqrt{\psi_1(q)} > 0$	Exact variance of the corrected log-variance
T_n	$\log S^2 - \psi(q) + \log q$	Unbiased estimator of θ
θ_0	$\log \sigma_0^2$	External target on the log scale
ξ	$(\theta_0 - \theta)/s_\nu$	Theoretical standardized target error
D_n	$\lvert T_n - \theta_0\rvert/s_\nu$	Observable standardized discrepancy
A_n	$(T_n - \theta_0)/s_\nu = Z_n - \xi$	Signed standardized target–data difference; $D_n = \lvert A_n\rvert$
$\widehat{H}_U$	$C_H + T_n/2$	Exact bias-corrected entropy baseline
AROS-0	$\eta = 1$	Tuning-free adaptive shrinkage rule
AROS-S	$\eta_S^\star(\kappa, n, b; C)$	Safety-constrained optimum within the AROS family
η	Nonnegative tuning parameter	Controls the discrepancy-dependent AROS weight
κ	Loss-asymmetry parameter	$\kappa = 0$ gives quadratic loss; the sign sets the direction of asymmetry
b	Declared standardized uncertainty radius	Design bound for $\lvert\xi\rvert$; not a confidence level
C	Risk-aggregation criterion	Uniform, center-weighted, or minimax regret

Then $\mathbb{E}(Z_n) = 0$ and $\mathrm{Var}(Z_n) = 1$, and its distribution is free of both the unknown mean and the unknown variance. Its exact density is

$$f_\nu(z) = \frac{s_\nu}{\Gamma(q)} \exp\Big[q\{s_\nu z + \psi(q)\} - \exp\{s_\nu z + \psi(q)\}\Big], \tag{2}$$

and its moment-generating function is

$$M_\nu(t) = \exp\Big[-\frac{t\,\psi(q)}{s_\nu}\Big] \frac{\Gamma\left(q + t/s_\nu\right)}{\Gamma(q)}, \qquad q + \frac{t}{s_\nu} > 0. \tag{3}$$

The density in Eq. (2) follows by writing $G = \exp\{s_\nu z + \psi(q)\}$ and applying the Jacobian $\mathrm{d}G/\mathrm{d}z = s_\nu G$ to the Gamma$(q, 1)$ density. Similarly, Eq. (3) follows from $\mathbb{E}(G^t) = \Gamma(q + t)/\Gamma(q)$ with t replaced by t/s_ν; the domain $q + t/s_\nu > 0$ is necessary for the exponential moment to exist. Expectations under asymmetric loss are therefore asserted only for values of κ within the relevant finite-moment domain. These expressions permit exact finite-sample risk evaluation without replacing the log-chi-square distribution by a normal approximation.

Corrected proofs

2.3 Adaptive Shrinkage Estimator and Theoretical Properties

Suppose that an external target variance σ_0^2 is available from information outside the current sample. Its log-scale form, the theoretical standardized target error, and the signed and absolute observable discrepancies are

$$\theta_0 = \log \sigma_0^2, \qquad \xi = \frac{\theta_0 - \theta}{s_\nu}, \qquad A_n = \frac{T_n - \theta_0}{s_\nu} = Z_n - \xi, \qquad D_n = |A_n|.$$

Here θ_0 is treated as nonrandom when the conditional risk is evaluated. The quantity ξ is used only as a risk-analysis coordinate: it is not observable and does not enter the operational estimator. The adaptive weight depends on $D_n^2 = A_n^2$, whereas the signed difference A_n is retained in the derivation of the error in Eq. (7), so that no sign information is lost. For $\eta > 0$, define the adaptive weight

$$w_\eta(D) = \frac{D^2}{D^2 + \eta}, \qquad \eta > 0,$$

and the corresponding shrinkage estimator

$$\widehat{\theta}_\eta = w_\eta(D_n)\, T_n + \{1 - w_\eta(D_n)\}\, \theta_0, \qquad \widehat{H}_\eta = C_H + \frac{1}{2}\widehat{\theta}_\eta. \tag{4}$$

The tuning-free specification $\eta = 1$ is denoted AROS-0. The weight is smooth and monotone, equals zero at target–data agreement, tends to one with increasing discrepancy (Figure 1), and yields the perturbation bound in Eq. (5). These properties motivate the chosen family; only η is risk-optimized, and no Bayes, admissibility, or global minimax claim is made for the functional form itself.

For a fixed weight w, the standardized estimation error is

$$e_w = \frac{\widehat{\theta}_w - \theta}{s_\nu} = wZ_n + (1 - w)\xi.$$

Under quadratic loss, the exact risk is

$$R_0(w, \xi) = w^2 + (1 - w)^2 \xi^2.$$

Differentiation gives the unique oracle weight and its risk,

$$w_{\mathrm{or}}(\xi) = \frac{\xi^2}{1 + \xi^2}, \qquad R_{\mathrm{or}}(\xi) = \frac{\xi^2}{1 + \xi^2}.$$

The oracle is infeasible because ξ depends on the unknown true variance. The adaptive rule uses the observable discrepancy to reproduce the same qualitative geometry without using ξ directly (Figure 2).

Under a rescaling $Y_i = cX_i$ with $c > 0$ and a coherently transformed target, the discrepancy is invariant and the entropy estimator transforms correctly:

Corrected proofs

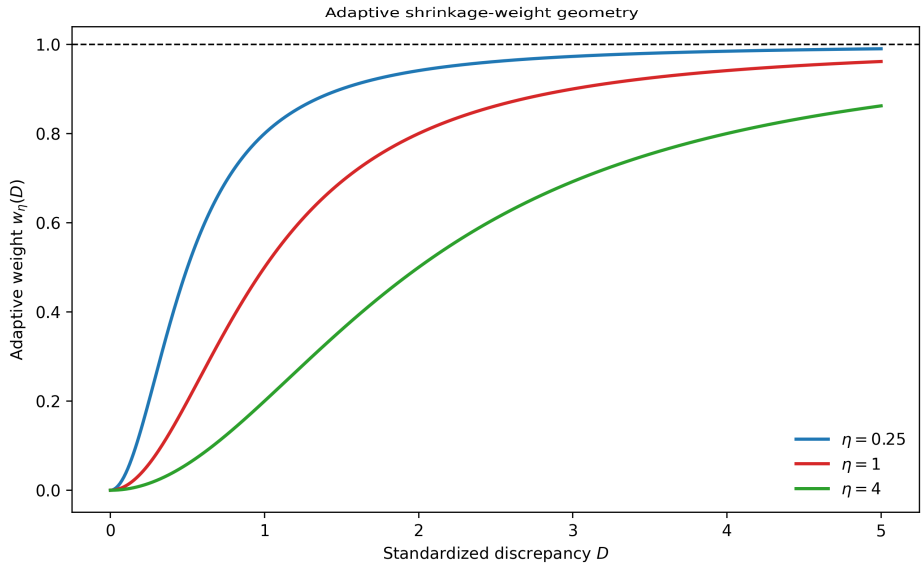


Figure 1: Adaptive shrinkage-weight geometry for representative values of the tuning parameter.

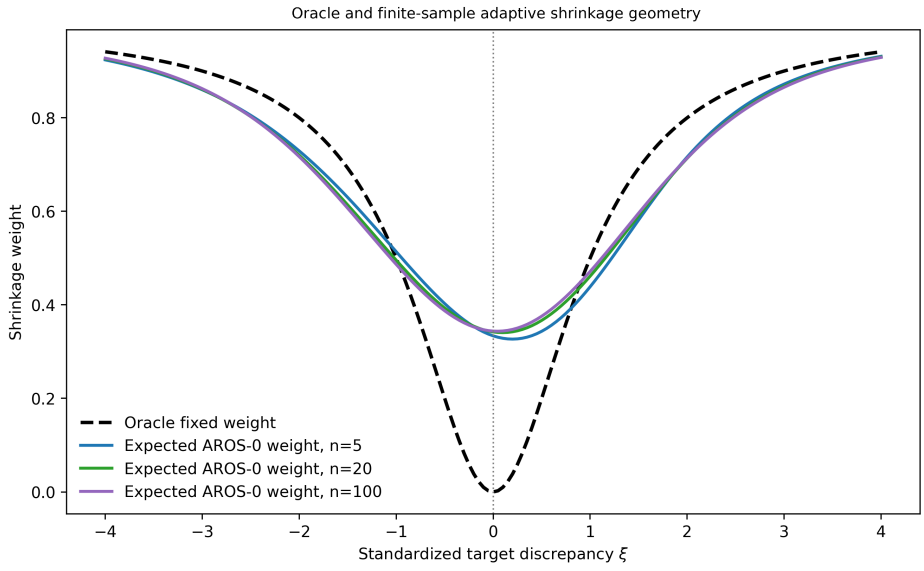


Figure 2: Oracle fixed-weight geometry and the expected finite-sample AROS-0 weight.

$$D_{n,Y} = D_{n,X}, \qquad \widehat{H}_{\eta,Y} = \widehat{H}_{\eta,X} + \log|c|.$$

A deterministic bound controls the largest samplewise perturbation away from the unbiased baseline:

$$\left|\widehat{\theta}_\eta - T_n\right| \le \frac{s_\nu\sqrt{\eta}}{2}, \qquad \left|\widehat{H}_\eta - \widehat{H}_U\right| \le \frac{s_\nu\sqrt{\eta}}{4}. \tag{5}$$

To verify Eq. (5), put $u = (T_n - \theta_0)/s_\nu$. Direct substitution in Eq. (4) gives $\widehat{\theta}_\eta - T_n = -s_\nu\eta u/(u^2+\eta)$. The elementary maximum $\sup_u |\eta u|/(u^2+\eta) = \sqrt{\eta}/2$ yields both bounds after accounting for the factor one-half in the entropy. Since $\mathbb{E}(T_n-\theta)^2 = s_\nu^2$ and s_ν^2 is asymptotically $2/\nu$, the bound in Eq. (5) gives the following sufficient condition for consistency:

$$s_\nu\sqrt{\eta_n} \to 0 \quad \Longrightarrow \quad \widehat{H}_{\eta_n} \xrightarrow{P} H.$$

The exact properties used to define and evaluate the proposed estimators are collected in Table 2.

Table 2: Exact theoretical properties used to define and evaluate the proposed estimators.

Property	**Exact statement**	**Interpretation**
Unbiased baseline	Exact unbiasedness; variance $\psi_1(q)/4$	The finite-sample correction is exact
Scale equivariance	The entropy estimate shifts by $\log\|c\|$ under $X \to cX$	The estimator respects entropy under unit changes
Bounded perturbation	Maximum departure from $\widehat{H}_U$ is $s_\nu\sqrt{\eta}/4$	Target borrowing cannot produce an unbounded samplewise shift
Consistency	Holds when $s_\nu\sqrt{\eta_n} \to 0$	Any fixed η is sufficient
Strict convexity	Fixed-weight risk is strictly convex in w	The loss-specific fixed optimum is unique
Local dominance	Adaptive risk is below baseline risk at $\xi = 0$	Strict improvement holds at the correct target
Asymptotic return to baseline	Adaptive risk returns to the baseline as $\|\xi\| \to \infty$	Under severe target error the risk approaches the baseline
AROS-S existence	A minimizer exists on the compactified tuning space	Uniqueness is not assumed

2.4 Risk-Based Loss Family and Exact Finite-Sample Risk

For the standardized error e, the analysis uses the normalized loss family

$$L_\kappa(e) = \begin{cases} \dfrac{2}{\kappa^2}\{\exp(\kappa e) - \kappa e - 1\}, & \kappa \neq 0, \\ e^2, & \kappa = 0. \end{cases}$$

Corrected Proofs

The quadratic loss is recovered continuously as $\kappa \to 0$. Positive and negative values of κ reverse the direction of the asymmetric penalization, following the LINEX principle of Zellner [24]. The same family contains the Stein loss for the variance ratio through

$$L_S(r) = r - \log r - 1 = \frac{s_\nu^2}{2} L_{s_\nu}(e), \qquad r = \exp(s_\nu e).$$

Stein's loss is a standard criterion for scale parameters because it penalizes gross over- and underestimation equally on the ratio scale; Bayes and empirical Bayes estimators of the normal variance under this loss are given in [25]. For a fixed shrinkage weight, the exact moment-generating function (3) gives

$$R_\kappa(w; \xi, \nu) = \frac{2}{\kappa^2}\Big[e^{\kappa(1-w)\xi} M_\nu(\kappa w) - \kappa(1-w)\xi - 1\Big],$$

and the unbiased-baseline risk is

$$R_{\kappa,U}(\nu) = \frac{2}{\kappa^2}\{M_\nu(\kappa) - 1\}, \qquad R_{0,U} = 1. \tag{6}$$

The fixed-weight risk is strictly convex in w on its admissible domain because the second derivative of the loss is positive and the pivot is nondegenerate. The loss-specific fixed optimum is therefore unique, although it need not have a closed form for $\kappa \neq 0$.

For the adaptive estimator, the standardized error can be written directly as

$$e_\eta(z, \xi) = z - \frac{\eta(z - \xi)}{(z - \xi)^2 + \eta}, \tag{7}$$

and its exact finite-sample risk is the one-dimensional integral

$$R_\kappa^{\mathrm{A}}(\eta; \xi, \nu) = \int_{-\infty}^{\infty} L_\kappa\big[e_\eta(z, \xi)\big] f_\nu(z)\, \mathrm{d}z. \tag{8}$$

At the exact target, the adaptive error has the same sign as the baseline pivot but a strictly smaller magnitude almost surely. Therefore

$$R_\kappa^{\mathrm{A}}(\eta; 0, \nu) < R_{\kappa,U}(\nu), \qquad \eta > 0. \tag{9}$$

By continuity, a nonempty target-centered neighborhood of strict improvement exists. Under severe target misspecification, the adaptive risk returns to the unbiased baseline:

$$\lim_{|\xi| \to \infty} R_\kappa^{\mathrm{A}}(\eta; \xi, \nu) = R_{\kappa,U}(\nu). \tag{10}$$

The strict improvement in Eq. (9) is pointwise at $\xi = 0$: the adaptive error equals $Z_n^3/(Z_n^2 + \eta)$, so it has the same sign as Z_n and a strictly smaller absolute magnitude except on a probability-zero event. Since $L_\kappa(e)$ is minimized at zero and is strictly monotone away from zero on each side, the risk inequality follows whenever the relevant expectation exists. As $|\xi| \to \infty$ for fixed η, the correction term in Eq. (7) tends to zero for every finite z and remains bounded by $\sqrt{\eta}/2$; dominated convergence, under the finite-exponential-moment condition, yields Eq. (10). Neither result implies uniform risk dominance over all target errors.

Corrected proofs

2.5 Safety-Constrained Optimization

If the target uncertainty is represented by the standardized interval $|\xi| \le b$, the fixed quadratic minimax benchmark is

$$w_b^\star = \frac{b^2}{1+b^2}, \qquad \sup_{|\xi|\le b} R_0\big(w_b^\star, \xi\big) = \frac{b^2}{1+b^2}.$$

For AROS, define the integrated design risk over the declared uncertainty interval as

$$\overline{R}_\kappa(\eta; b) = \frac{1}{2b}\int_{-b}^{b} R_\kappa^{\mathrm{A}}(\eta; \xi, \nu)\, \mathrm{d}\xi. \tag{11}$$

The uniform measure in this criterion is a design device, not a Bayesian prior. To prevent an attractive average risk from being obtained at the cost of excessive edge risk, define the feasible set

$$\mathcal{E}_{\kappa,\nu}(b) = \Big\{\eta \ge 0 : \sup_{|\xi|\le b} R_\kappa^{\mathrm{A}}(\eta; \xi, \nu) \le R_{\kappa,U}(\nu)\Big\},$$

with $\eta = 0$ defined by closure as the no-shrinkage estimator $\widehat{H}_U$. The safety-constrained tuning is

$$\eta_{\kappa,S}^\star = \arg\min_{\eta\in\mathcal{E}_{\kappa,\nu}(b)} \overline{R}_\kappa(\eta; b). \tag{12}$$

After the compactification $\rho = \eta/(1+\eta)$, continuity on a compact parameter space ensures the existence of a minimizer. The continuum-indexed constraint makes the formulation formally semi-infinite. This is related to bounded-misspecification shrinkage [3, 4], while AROS-S specializes the idea to the exact log-Gamma entropy risk. No new optimality theory or algorithm for semi-infinite programming is claimed. If no positive η improves the criterion safely, $\eta = 0$ returns the unbiased baseline.

The operational computation uses the following sequence. For specified (n, κ, b), calculate s_ν and f_ν and evaluate the one-dimensional exact risk in Eq. (8) by numerical quadrature. For each trial η, maximize that risk over the entire closed interval $[-b, b]$, checking both endpoints and possible interior maxima, and retain η only if this maximum does not exceed the unbiased-baseline risk in Eq. (6). Integrate the admissible risk over the declared uncertainty interval, minimize the resulting criterion over $\eta \ge 0$ using a coarse logarithmic search followed by scalar refinement, and recheck the worst-case bound at the selected solution. The baseline $\eta = 0$ is always a feasible reference. Numerical evaluation of the continuum maximum, rather than checking only the two endpoints, is essential for the safety claim.

The boundary case $b = 0$ is defined separately, because the factor $1/(2b)$ in Eq. (11) is undefined there. Its continuous extension is $\overline{R}_\kappa(\eta; 0) = R_\kappa^{\mathrm{A}}(\eta; 0, \nu)$, with the safety condition evaluated at $\xi = 0$. If the external target is genuinely known to equal the true parameter, then $\eta \to \infty$ gives $\widehat{\theta}_\eta \to \theta_0 = \theta$ and zero risk; the compactified tuning space includes this

Corrected proofs

limiting pure-target action. Pointwise exact-target comparisons of AROS-0 ($\eta = 1$) with a pre-test estimator must therefore not be presented as the optimal AROS-S rule at $b = 0$. In the Michelson example (Section 3.8) the historical target is estimated, so $b = 0$ would not be a defensible representation of its uncertainty.

2.6 Numerical Evaluation, Comparator Reconstruction, and Validation

The main finite-sample risk grid was

$$n \in \{5, 10, 20, 50, 100\}, \qquad \kappa \in \{-1, -0.5, 0, 0.5, 1\}, \qquad \xi = -6, -5.75, \ldots, 6,$$

giving 1,225 primary configurations. The safety optimization examined $b \in \{0.5, 1, 1.5, 2\}$. Risks were computed by adaptive quadrature, with selected cases recomputed independently on the Gamma scale. Gauss–Legendre integration, a logarithmic η scan, scalar refinement, and worst-case verification were used for tuning. Numerical convergence was assessed through the independent quadrature agreement and the Monte Carlo checks in Table 6; the largest discrepancy between the Z-scale and Gamma-scale calculations was 5.29×10^{-13}.

The comparators were the plug-in estimator, the exact unbiased baseline, the infeasible loss-specific oracle, the fixed minimax rule, and the Huntsberger-type entropy estimator of Sahni and Kumar [18]. Their published construction uses the maximum-likelihood variance

$$s_M^2 = \frac{1}{n}\sum_{i=1}^{n} (X_i - \overline{X})^2 = \frac{\nu}{n} S^2, \tag{13}$$

and the preliminary-test statistic and critical values

$$V = \frac{n\, s_M^2}{\sigma_0^2}, \qquad c_1 = \chi^2_{\nu,\alpha/2}, \qquad c_2 = \chi^2_{\nu,1-\alpha/2}.$$

Within the acceptance region, the reconstructed Huntsberger shrinkage factor is

$$k_H = \frac{n\, s_M^2}{\sigma_0^2\,(c_2 - c_1)},$$

which yields the log-scale comparator

$$\widehat{\theta}_H = \begin{cases} k_H \log s_M^2 + (1 - k_H)\,\theta_0, & c_1 \le V \le c_2, \\ \log s_M^2, & \text{otherwise.} \end{cases} \tag{14}$$

The comparator in Eqs. (13)–(14) is reconstructed exactly as a specified pre-test rule, without replacing its maximum-likelihood variance by the bias-corrected variance used for $\widehat{H}_U$. In the current standardized risk calculation, $V = 2G\exp(-s_\nu \xi)$, because $\nu S^2/\sigma_0^2 = 2G\exp(-s_\nu \xi)$;

Corrected proofs

the acceptance event is therefore $c_1 \le 2G\exp(-s_\nu\xi) \le c_2$. On this event the log-scale error is $[k_H \log(2G/n) + (1-k_H)s_\nu\xi]/s_\nu$, and outside it the error is $\log(2G/n)/s_\nu$. Integrating L_κ of this piecewise error against the Gamma$(q,1)$ density yields the reported comparator risks under the same loss as for AROS. These formulas also specify how the target error enters the pre-test cutoff.

Three significance levels, 0.01, 0.05, and 0.10, were evaluated, and the smallest risk among the three was retained at each value of the standardized target discrepancy ξ, giving the published family a favorable benchmark. This best-of-three construction is a comparison envelope rather than an operational tuning rule, because the pointwise best choice depends on the unknown value of ξ. Selected published relative-risk entries were reconstructed to within 4.81×10^{-5}.

Tables 5 and 7 use different, explicitly declared comparison conventions. Table 5 and the corresponding pointwise figures use the minimum risk among the three tested significance levels at each ξ, which is an oracle-style envelope rather than a usable choice of α; Table 7 fixes α in advance and compares the worst-case quadratic risk over $|\xi| \le 1$. The two comparisons are not interchangeable, and their numerical values are conditional on the cited Huntsberger-type reconstruction; they must not be generalized to all pre-test estimators.

Monte Carlo simulation was used only for independent validation. Five representative configurations were simulated with $B = 500{,}000$ replications and a fixed seed. The Monte Carlo risk and its estimated standard error were

$$\widehat{R}_{MC} = \frac{1}{B}\sum_{b=1}^{B} L_\kappa(e_b), \qquad \widehat{\mathrm{SE}}\big(\widehat{R}_{MC}\big) = \frac{s_L}{\sqrt{B}}.$$

Deterministic quadrature remained the primary risk engine throughout.

2.7 Operational Estimator Selection, Sensitivity Analysis, and Real-Data Application

For practical selection under quadratic loss ($\kappa = 0$), the worst-case risks of AROS-S and of the Huntsberger comparator are compared over $|\xi| \le b$ at prespecified n and α, and the rule with the smaller risk is selected. Table 7 is intentionally restricted to quadratic loss; for $\kappa \neq 0$ the directional risk requires a separate loss-specific calibration rather than a direct reuse of this map.

The worst-case selection rule is defined only for a strictly positive, prespecified b: choose the method with the smaller maximum quadratic risk over $|\xi| \le b$, using a fixed α for the Huntsberger comparator. At $b = 0$, the exact-target limiting AROS-S action is the pure target ($\eta \to \infty$); comparisons of Huntsberger with the fixed AROS-0 rule at $\xi = 0$ are therefore illustrative pointwise risk comparisons and do not select the optimal AROS-S at $b = 0$. An estimated historical target cannot be treated as exactly known simply by setting $b = 0$.

Corrected proofs

The sensitivity analysis at $b = 1$ uses the uniform integrated risk, a center-weighted risk, and the minimax regret, all under the same safety constraint. Because feasibility is imposed in every case, the comparison concerns whether positive shrinkage remains optimal and how much η changes, not whether the common constraint is satisfied (Section 3.7).

For the Michelson illustration [20], Experiment 5 ($n = 20$) is the current sample, and Experiments 1–4 define a historical target (76 residual degrees of freedom), assumed to be independent of the current sample. The current scaling uses $\nu = 19$; the 76 degrees of freedom enter only the historical variance interval. The measurements are coded in km/s with 299,000 km/s subtracted, so the entropy is reported in nats on that scale. Mapping the historical 95% interval gives b_{cal}, an illustrative proxy for b rather than a two-sample derivation. Bootstrap resampling quantifies conditional variability, not accuracy.

The historical target in this illustration is a statistic estimated from a separate sample. The exact theory treats its realized value θ_0 as fixed: if the historical sample is independent and follows a common-variance Gaussian model with $q_H = 76/2$, then $\log S_H^2 = \log \sigma_H^2 + \log G_H - \log q_H$ with $G_H \sim \text{Gamma}(q_H, 1)$, and its log-scale variance is $\psi_1(q_H)$. An unconditional assessment would integrate the conditional AROS risk over G_H, also allowing the historical population variance σ_H^2 to differ from the current variance σ^2; it is not supplied by the fixed-target formulas in Eqs. (8)–(12). Accordingly, b_{cal} is an exploratory calibration, not a guaranteed uncertainty bound or a two-sample confidence statement.

3 Results

3.1 Target-Exact Finite-Sample Performance

At $\xi = 0$, AROS-0 had a lower risk than the bias-corrected baseline for all reported combinations of n and κ. Under quadratic loss, the risk fell from 0.487150 at $n = 5$ to 0.467842 at $n = 100$, corresponding to exact-target reductions of 51.29% and 53.22%. These gains are point-specific. The similarity across n reflects a rapid approach to the limiting geometry, while the calculations remain exact finite-sample log-Gamma risks without a Gaussian approximation.

3.2 Local Dominance and Sensitivity to Target Misspecification

The target-exact gain did not imply global dominance. Under quadratic loss, the principal connected region containing $\xi = 0$ over which AROS-0 had a lower risk than $\widehat{H}_U$ is reported in Table 3. The interval was asymmetric at small n and became progressively more balanced

Corrected proofs

around zero as the sample size increased. The corresponding relative-risk profiles are shown in Figure 3.

Table 3: Principal local dominance region of AROS-0 under quadratic loss.

n	**Left boundary**	**Right boundary**	R_U/R_{AROS0} at $\xi = 0$	**Risk reduction (%)**
5	−1.684341	1.219944	2.052757	51.29
10	−1.602711	1.297737	2.100703	52.40
20	−1.553811	1.345742	2.121944	52.87
50	−1.514001	1.385137	2.133712	53.13
100	−1.494751	1.404257	2.137472	53.22

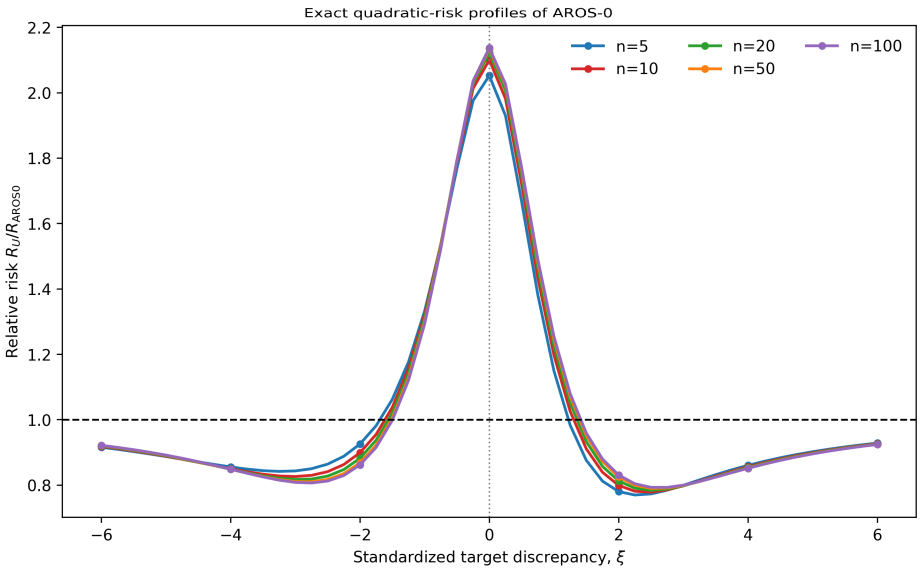


Figure 3: Exact relative-risk profiles R_U/R_{AROS0} under quadratic loss ($\kappa = 0$) for $n = 5, 10, 20, 50$, and 100; values above one favor AROS-0 and the dashed line denotes equal risk.

Loss asymmetry changed both the direction and the extent of the improvement region (Figure 4). At $n = 20$, the principal region for $\kappa = -1$ began near $\xi = -1.53836$ and continued across the positive side of the inspected range, whereas for $\kappa = 1$ the corresponding region extended across the negative side and ended near $\xi = 1.23071$. These results show that target errors of equal magnitude and opposite direction need not have equal risk consequences under asymmetric loss.

The two-dimensional risk map in Figure 5 gave the same result over the full n-by-ξ grid: the contour $R_U/R_{\mathrm{AROS0}} = 1$ separated the local region of improvement from the region in which the bias-corrected estimator had the lower risk.

Core Proofs

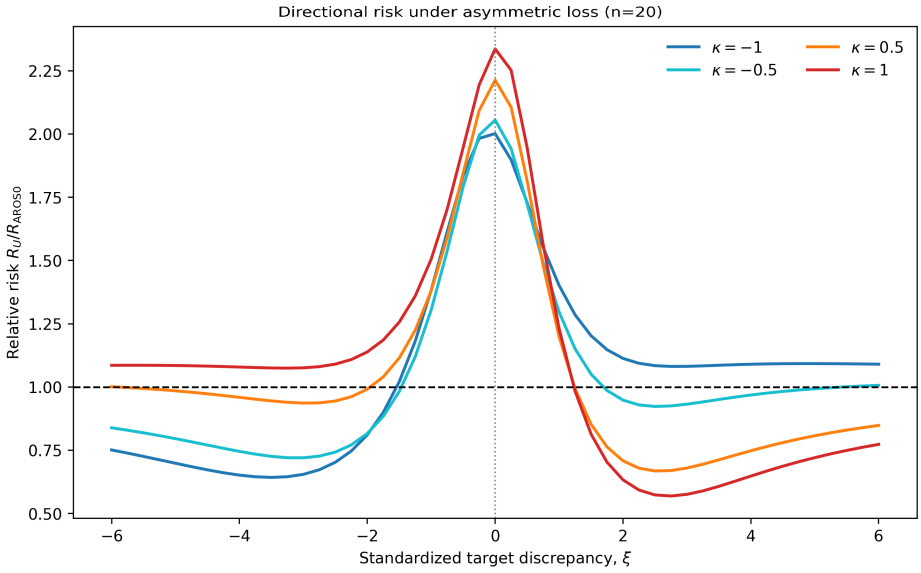


Figure 4: Directional relative-risk profiles R_U/R_{AROS0} at $n = 20$ under asymmetric loss, $\kappa \in \{-1, -0.5, 0, 0.5, 1\}$; the dashed line denotes equal risk.

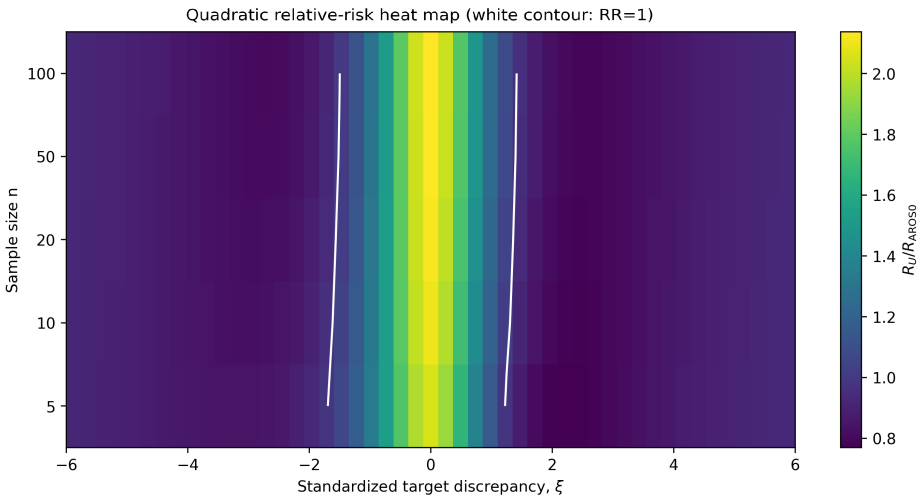


Figure 5: Quadratic-loss ($\kappa = 0$) relative-risk heat map over the sample size n and the standardized target error ξ; the contour $R_U/R_{\mathrm{AROS0}} = 1$ marks equal risk.

3.3 Safety-Constrained Optimization

For $|\xi| \leq 1$, every AROS-S solution satisfied the prescribed worst-case bound by construction; the optimization determined whether positive shrinkage was worthwhile within that feasible set. Table 4 reports the results for $\kappa = -1, 0, 1$. Under quadratic loss, $\eta_S^\star$ declined from 22.4323 at $n = 5$ to 15.3235 at $n = 100$, while the integrated and worst-case risks also decreased.

Table 4: Safety-constrained AROS-S optimization for $|\xi| \leq 1$. Integrated and worst-case risks are relative to the corresponding unbiased-baseline risk.

	$\kappa = -1$			$\kappa = 0$			$\kappa = 1$		
n	$\eta^\star$	Int./R_U	Max/R_U	$\eta^\star$	Int./R_U	Max/R_U	$\eta^\star$	Int./R_U	Max/R_U
5	53.0450	0.154865	0.601214	22.4323	0.300064	0.820190	4.9619	0.327396	0.965118
10	28.2190	0.186065	0.693343	17.6808	0.289480	0.786204	12.2166	0.271400	0.957444
20	21.8832	0.200272	0.733590	16.2215	0.284587	0.770301	12.2166	0.255843	0.899569
50	18.8428	0.210550	0.764488	15.5296	0.281862	0.760545	14.4046	0.243825	0.875230
100	17.7795	0.215349	0.779599	15.3235	0.280992	0.756950	14.7716	0.238520	0.856991

The tuning was loss dependent. At $n = 20$ and $b = 1$, $\eta_S^\star$ was about 21.88, 16.22, and 12.22 for $\kappa = -1, 0, 1$. For quadratic loss with $b = 1.5$ or 2, $\eta_S^\star = 0$ for every examined n (Figure 6). In practice, b should be declared from independent scientific tolerances or historical information; if no defensible finite bound exists, the protected choice is the baseline or a sensitivity analysis over plausible values of b.

3.4 Benchmark Comparisons

Table 5 compares AROS-0 with the principal reference procedures under quadratic loss at representative values of ξ. The infeasible oracle gives the lower performance envelope. The fixed minimax rule for $b = 1$ has the risk $0.25 + 0.25\,\xi^2$ and is included only as a theoretical benchmark.

Table 5 shows that the reconstructed Huntsberger rule had a distinct efficiency–robustness profile: at $n = 5$ and $\xi = 0$ its best tested quadratic risk was 0.171669 versus 0.487150 for AROS-0, whereas at $n = 100$ it was 1.776025 versus 0.467842. The comparison is conditional on the published formulation and its verified reconstruction and does not establish a general ordering of pre-test and smooth-shrinkage procedures. The asymmetric-loss comparison is displayed in Figure 7.

Corrected proofs

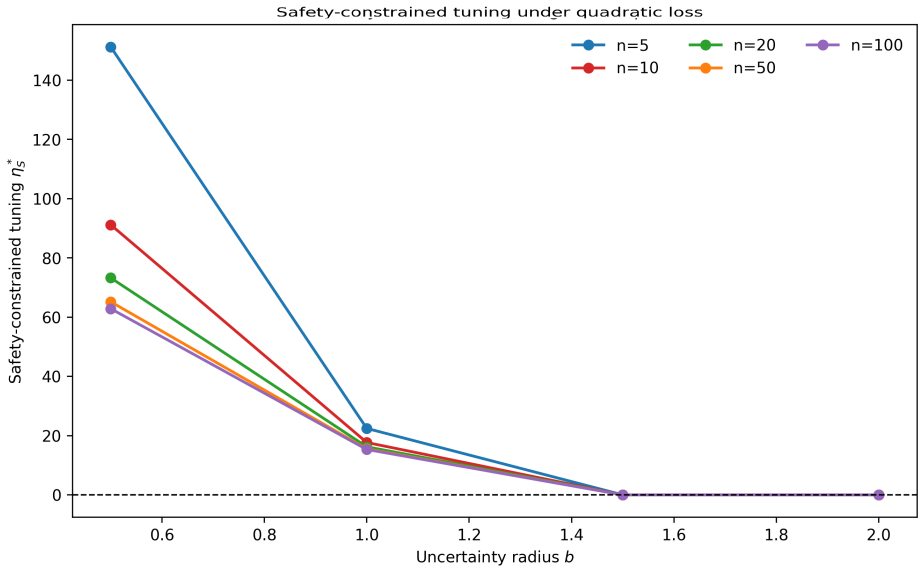


Figure 6: Safety-constrained optimal tuning under quadratic loss as the target-uncertainty radius b changes.

Table 5: Representative quadratic-risk benchmark comparisons.

n	ξ	$R_{\widehat{H}_U}$	R_{PI}	R_{AROS0}	R_{oracle}	$R_{\text{fixed minimax}}$	$R_{\text{Hunts best}}$	α
5	-1	1.000000	1.113339	0.750686	0.500000	0.500000	0.716356	0.01
5	0	1.000000	1.113339	0.487150	0.000000	0.250000	0.171669	0.01
5	1	1.000000	1.113339	0.870157	0.500000	0.500000	1.063560	0.01
20	-1	1.000000	1.025839	0.763158	0.500000	0.500000	0.858327	0.01
20	0	1.000000	1.025839	0.471266	0.000000	0.250000	0.402524	0.01
20	1	1.000000	1.025839	0.816372	0.500000	0.500000	0.690060	0.01
100	-1	1.000000	1.005033	0.773736	0.500000	0.500000	2.252328	0.01
100	0	1.000000	1.005033	0.467842	0.000000	0.250000	1.776025	0.01
100	1	1.000000	1.005033	0.796855	0.500000	0.500000	1.165942	0.01

Corrected Proofs

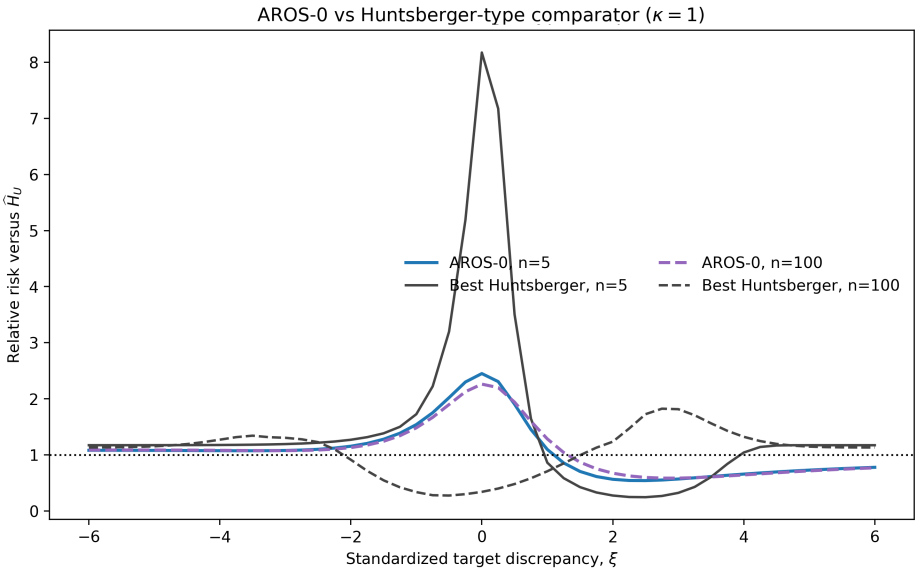


Figure 7: Relative-risk comparison of AROS-0 and the reconstructed Huntsberger-type pre-test benchmark under asymmetric loss $\kappa = 1$; curves are shown for the displayed sample sizes over the standardized target error ξ.

The plug-in and unbiased baselines also changed order under asymmetric loss. At $n = 5$, $\kappa = 1$, and $\xi = 0$, $R_{PI} = 0.813663$ was below $R_U = 0.997043$. Exact unbiasedness therefore did not imply uniform risk superiority when over- and underestimation are penalized differently.

3.5 Independent Numerical Validation

The standardized Z-scale and the direct Gamma-scale quadratures agreed to numerical precision. Across eight independent checks, the largest absolute difference was 5.29×10^{-13}. Five Monte Carlo checks based on 500,000 replications were consistent with the deterministic risks; the largest standardized discrepancy was 0.516 Monte Carlo standard errors. The full quadrature and Monte Carlo diagnostics are presented in Table 6 and Figure 8.

3.6 Operational Estimator-Selection Rule

Table 7 operationalizes the estimator selection. For $b = 1$, AROS-S had a lower worst-case quadratic risk than Huntsberger for every examined n and $\alpha \in \{0.01, 0.05, 0.10\}$. In the target-exact case, Huntsberger was preferred only for $5 \leq n \leq 24$ at $\alpha = 0.01$ and for $5 \leq n \leq 9$ at $\alpha = 0.05$; at $\alpha = 0.10$, AROS-0 was preferred throughout $n \geq 5$ in the assessed range. The

Corrected Proofs

Table 6: Independent numerical validation of the exact risk calculations. Panel A compares Z-scale and Gamma-scale quadrature; Panel B compares deterministic and Monte Carlo risks (500,000 replications).

Panel A. Z-scale versus Gamma-scale quadrature

n	ξ	κ	Exact Z risk	Gamma risk	Abs. difference
5	0.0	-1.0	1.259376090503	1.259376090504	5.294×10^{-13}
5	1.0	1.0	0.908495750032	0.908495750032	7.772×10^{-16}
10	-1.5	0.5	0.871370440705	0.871370440705	3.542×10^{-14}
20	0.5	0.0	0.572458930886	0.572458930886	1.110×10^{-16}
100	2.0	-0.5	1.139637467066	1.139637467066	1.177×10^{-14}
10	0.0	0.0	0.476031174951	0.476031174951	6.606×10^{-15}
20	-1.0	1.0	0.759617370637	0.759617370637	5.551×10^{-16}
50	1.5	-1.0	1.173814634562	1.173814634562	4.441×10^{-16}

Panel B. Monte Carlo validation (500,000 replications)

n	ξ	κ	Exact risk	MC risk	MC SE	MC z-score
5	0.0	-1.0	1.259376091	1.268845826	0.045153	0.210
5	1.0	1.0	0.908495750	0.907969832	0.001211	-0.434
10	-1.5	0.5	0.871370441	0.870848404	0.001460	-0.357
20	0.5	0.0	0.572458931	0.571658309	0.001551	-0.516
100	2.0	-0.5	1.139637467	1.138837033	0.002302	-0.348

robust worst-case comparison therefore governs when $b > 0$, while the target-exact crossover applies only when the target is treated as effectively exact.

The comparisons at $\xi = 0$ cited above refer specifically to AROS-0 versus the reconstructed Huntsberger rule. They do not apply to the optimal AROS-S at $b = 0$: with an exactly correct known target, the compactified AROS-S solution is $\eta = \infty$ and has zero risk. The $b = 1$ entries of Table 7 instead concern a genuinely positive uncertainty radius and values of α fixed before the worst-case risk is computed.

3.7 Sensitivity to the Risk-Aggregation Criterion

The alternative risk-aggregation results in Table 8 show that changing the criterion altered η substantially. At $n = 20$ and $b = 1$, η was 16.221, 33.967, and 5.280 under the uniform, center-weighted, and minimax-regret objectives, respectively. Because the same safety constraint is imposed in all cases, its satisfaction is not a robustness finding; the relevant result is that all three objectives still select positive shrinkage, with materially different tuning. For the same

Corrected proofs

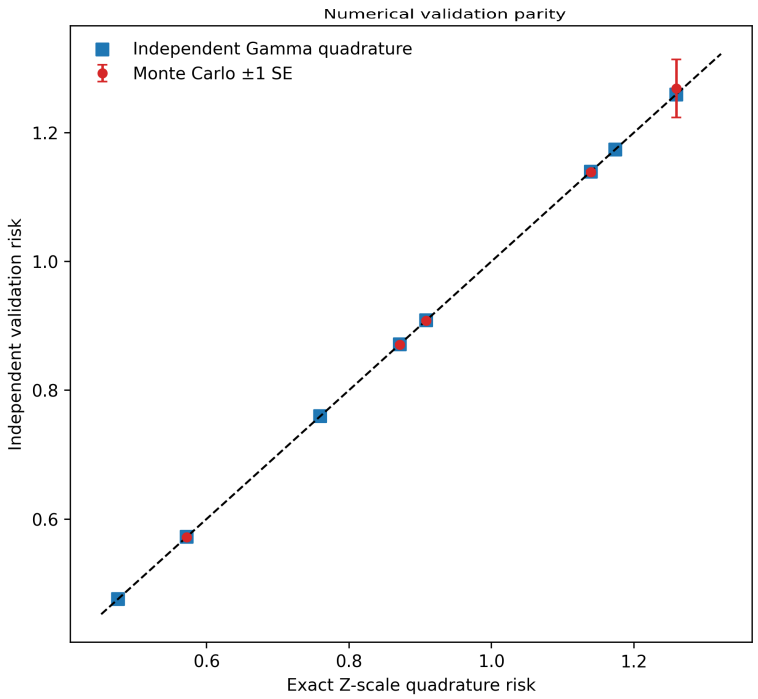


Figure 8: Independent numerical validation. Squares represent direct Gamma-scale quadrature; circles represent Monte Carlo estimates with one-standard-error bars; the dashed line is the identity line.

Table 7: Operational comparison of the worst-case quadratic risk over $|\xi| \leq 1$.

n	**AROS-S max risk**	**Huntsberger** $\alpha = 0.01$	**Huntsberger** $\alpha = 0.05$	**Huntsberger** $\alpha = 0.10$	**Recommended**
5	0.820190	1.063560	1.361122	1.486028	AROS-S
10	0.786204	0.859382	1.086593	1.182096	AROS-S
20	0.770301	0.858327	0.928145	1.028304	AROS-S
50	0.760545	1.309340	1.506125	1.693589	AROS-S
100	0.756950	2.252328	2.685279	2.985034	AROS-S

Corrected Proofs

reason, the lower risk recorded for one criterion should not be interpreted as an independent guarantee.

Table 8: Sensitivity of quadratic AROS-S tuning at $b = 1$ to alternative risk-aggregation criteria.

n	η: uniform	η: center-weighted	η: minimax regret	Max risk: uniform	Max risk: center	Max risk: regret
5	22.432	59.919	6.746	0.8202	0.9099	0.7705
10	17.681	39.813	5.714	0.7862	0.8663	0.7334
20	16.221	33.967	5.280	0.7703	0.8458	0.7132
50	15.530	31.316	4.997	0.7605	0.8345	0.6985
100	15.323	30.545	4.878	0.7569	0.8309	0.6921

3.8 Real-Data Application: Michelson Speed-of-Light Experiments

In the illustrative held-out application, the historical pooled variance was 6153.3553 (76 degrees of freedom), with the 95% interval $[4584.8866, 8695.3614]$. Its mapping gives $b_{\mathrm{cal}} = 1.0379$, used only as a calibration heuristic. Experiment 5 has $n = 20$ and a variance of 2939.7368 on the coded km/s scale. The Shapiro–Wilk p-value of 0.194 does not reject normality but cannot confirm it at $n = 20$; known experiment-level irregularities make this an implementation example rather than a validation of the Gaussian model. The observed discrepancy is $D_n = 2.0564$ and the tuning is $\eta_S^\star = 14.9302$.

Historical-pooling diagnostic. Recalculation from the 20 runs in each historical experiment gives within-experiment sample variances of 11009.47, 3741.05, 6257.89, and 3605.00 (coded km/s)^2 for Experiments 1–4; their degrees-of-freedom-weighted average is 6153.3553. The largest is approximately 3.05 times the smallest, so equal within-experiment variances cannot be taken for granted. A median-centered Levene (Brown–Forsythe) check gives $p = 0.230$, while Bartlett's test gives $p = 0.0447$; the latter is sensitive to departures from normality. These diagnostics do not establish a common variance, and the pooled statistic is retained only as an illustrative historical target, not as an independently verified precision benchmark.

The entropy estimates (nats on the coded km/s scale) were 5.4388, 5.5043, 5.7057, and 5.6316 for the baseline, AROS-0, AROS-S, and Huntsberger estimators, respectively. Their bootstrap intervals are reported in Table 9 and displayed in Figure 9. The bootstrap standard errors summarize conditional resampling variability; the smaller standard error of AROS-S is not evidence of a smaller estimation error or a greater accuracy.

Experiment-5 sampling diagnostic. The 20 reported measurements are consecutive, discretized runs. Their sample mean is 831.5, the standard deviation is 54.22, and the sample skewness is approximately 0.69 on the coded scale. The Shapiro–Wilk result $W = 0.9352$

Corrected proofs

($p = 0.194$) does not reject normality but cannot confirm it; the small sample and the moderate right skew call for caution. The lag-one sample correlation is approximately 0.34, a descriptive indication that the consecutive-run design does not by itself establish serial independence. There is also no documented verification of stable measurement conditions or of a common current and historical variance. Since Eq. (1) depends on independent, identically distributed Gaussian observations, the Michelson exercise demonstrates the implementation under a working model rather than verifying its sampling assumptions. The intervals in Table 9 describe the reported current-sample bootstrap while conditioning on the selected tuning choices and the realized historical target; target-estimation and tuning-selection uncertainty are not included.

Table 9: Michelson Experiment 5 entropy estimates (nats on the coded km/s scale) using an external historical target from Experiments 1–4. Bootstrap quantities describe conditional resampling variability, and the intervals are conditional on the historical-target value and the selected tuning parameters.

Estimator	Point estimate	Bootstrap SE	95% bootstrap interval	Weight w on the data estimate
Bias-corrected baseline	5.4388	0.1766	[4.9741, 5.6686]	1.000
AROS-0	5.5043	0.1898	[5.0071, 5.7459]	0.809
AROS-S	5.7057	0.1318	[5.2879, 5.7780]	0.221
Huntsberger ($\alpha = 0.05$)	5.6316	0.2405	[4.9217, 5.6822]	0.379

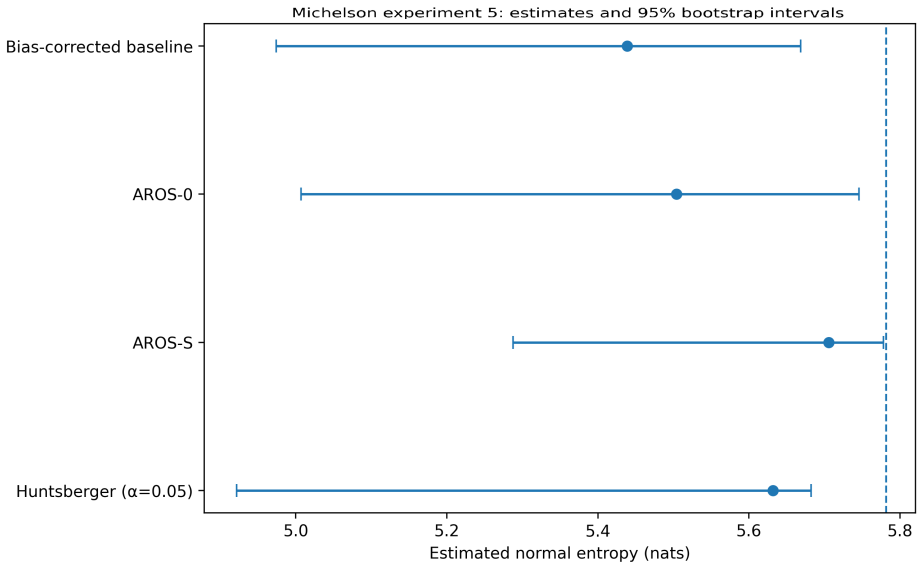


Figure 9: Bootstrap comparison of the Michelson Experiment 5 entropy estimates with 95% resampling intervals; the dashed line denotes the historical-target entropy. The intervals summarize resampling variability, not estimator accuracy; target and tuning uncertainty are held outside this bootstrap comparison.

Co-edproofs

4 Discussion

This study treats normal-entropy shrinkage as a target-dependent decision problem. Its bounded-risk structure shares the general robust-shrinkage principle of [3, 4], whereas the contribution of this paper is the exact finite-sample log-Gamma entropy formulation. The semi-infinite label describes the continuum constraint and is not a claim of new general semi-infinite programming theory.

The target-exact gain of AROS-0 is substantial but local: the improvement region is finite (Table 3) and the risk returns toward the baseline as $|\xi|$ grows. The 51–53% reduction therefore applies only at $\xi = 0$. Under asymmetric loss, target errors of equal magnitude and opposite direction can have different risk consequences, so claims must be conditioned on κ and on the uncertainty set.

4.1 Comparison with Pre-test Entropy Shrinkage

The closest entropy-specific comparator is the estimator of Sahni and Kumar [18], while broader empirical-Bayes, bounded-parameter, and decision-theoretic variance shrinkage appears in [6, 8, 13, 15]. The reconstructed Huntsberger rule provides a useful contrast between pre-test and smooth borrowing, but it is not presented as the only or universally strongest alternative shrinkage class.

The advantage of the pre-test rule at small n and $\xi = 0$ did not persist throughout the design. At $n = 100$ and $\xi = 0$, the reconstructed Huntsberger risk was 1.776025, whereas AROS-0 remained at 0.467842 (Table 5). This does not establish a general superiority of AROS; it illustrates a structural difference. The pre-test rule changes regime at the chi-square boundaries, whereas AROS changes continuously with the observed disagreement, and its departure from the corrected baseline is deterministically bounded (Eq. (5)). The two approaches therefore represent different ways of exploiting target information, and neither should be judged only at $\xi = 0$.

Figure 7 displays the asymmetric-loss pointwise comparison for the stated comparator and sample sizes; it uses a fixed risk definition and is not a universal comparison of smooth and preliminary-test shrinkage. When α is selected ex post at each value of ξ, the resulting curve is an unattainable oracle envelope; Table 7 instead evaluates the operational fixed-α, quadratic-loss case.

Corrected proofs

4.2 Safety-Constrained Optimization and Operational Estimator Selection

Safety-constrained tuning limits the edge risk over a declared set. Table 7 gives a quadratic-loss operational comparison; it should not be extrapolated to $\kappa \neq 0$, where the directional risk requires separate calibration. In practice, b should come from external tolerances or historical evidence before tuning; if that is not defensible, the baseline or a sensitivity analysis over plausible values of b is the safer default.

The limiting case $b = 0$ represents a truly exact external target, not merely a small uncertainty interval: the AROS-S objective reduces to the target-exact risk, and its limiting optimum is zero at $\eta = \infty$. For $b > 0$, all operational selections must use a defensible, externally specified radius, and Table 7 reports only the $b = 1$ quadratic-loss comparison. Consequently, the favorable exact-target results for AROS-0 alone cannot be used to recommend Huntsberger over the optimized AROS-S at $b = 0$.

The tuning $\eta_S^\star$ depends on the aggregation criterion as well as on the loss (Table 8). Since the safety inequality is common to all formulations, feasibility is imposed rather than discovered; the substantive sensitivity result is that all three objectives still choose active shrinkage at $b = 1$, but with different values of η.

4.3 Numerical Reliability

The principal risks were computed deterministically. The independent Z- and Gamma-scale integrations agreed within 5.29×10^{-13}, and the Monte Carlo estimates differed by at most 0.516 estimated standard errors (Table 6). The published Huntsberger entries were reconstructed within 4.81×10^{-5} before the common-loss evaluation. These checks support numerical reproducibility, while the comparative conclusions remain conditional on the stated comparator reconstruction and loss.

4.4 Real-Data Implications

The Michelson application illustrates the implementation rather than the estimator accuracy. The earlier experiments supply an assumed-independent target, and b_{cal} is a calibration heuristic. The bootstrap intervals show the current-sample resampling variability conditional on that fixed target (Table 9 and Figure 9); they do not identify which estimator is closest to the unknown true entropy.

The four historical sample variances differ by a factor of roughly three, and their pooled value is therefore an assumption-dependent target estimate. Because the original experiments

Corrected proofs

consist of successive runs and may have been conducted under different conditions, neither independence nor variance homogeneity is demonstrated by the available diagnostics. The Gaussian exact-risk formulas are conditional idealizations, and the uncertainty of the historical target would require a separate two-sample risk calculation.

4.5 Limitations

The exact risk theory assumes normal sampling. The Michelson normality test is only a small-sample diagnostic. The target is treated as fixed, and b_{cal} obtained from a historical confidence interval is a proxy rather than a two-sample uncertainty theory. Table 7 is restricted to quadratic loss, and asymmetric selection needs directional calibration. Only η, not the AROS functional form, is risk-optimized; broader Bayes or admissibility optimality is not claimed. Multivariate entropy requires covariance shrinkage and matrix theory [5, 9, 16].

The numerical demonstration additionally conditions on the observed historical pooled target and on the selected tuning parameters; the bootstrap intervals do not propagate their estimation variability. A pooled 76-degree-of-freedom chi-square interval is exact only under the usual independent normal common-variance assumptions for the four historical groups, which the observed variance differences do not confirm. The $b = 0$ pure-target limit is a theoretical special case, not a justified choice for this historical estimate.

4.6 Directions for Further Work

Extensions include stochastic historical targets, repeated-split or predictive validation where sequential data permit it, non-normal scale models, and simulations with known stochastic target uncertainty in a two-sample setting. Multivariate work could use multiple covariance targets [9], but it requires positive-definiteness and matrix-risk theory. A balanced-loss formulation, which weights proximity to a target estimator against estimation precision as in [1], is a natural alternative to the safety-constrained criterion used here; its adaptation to the log-variance scale is left for future work.

5 Conclusion

With respect to the four objectives stated in Section 1, this study develops an exact finite-sample target-adaptive estimator and a safety-constrained tuning rule for normal entropy. Within the AROS family, η is optimized subject to a worst-case bound. The 51–53% gain is restricted to

Corrected proofs

exact target specification under quadratic loss; the improvement is local, and wide uncertainty sets can return $\eta_S^\star = 0$. The operational rule of Table 7 is restricted to quadratic loss, while asymmetric selection requires separate calibration. The Michelson analysis is an implementation illustration, with b_{cal} used as a proxy and the bootstrap intervals interpreted as resampling variability. Target borrowing is therefore appropriate only when an independently defensible uncertainty set is available.

The $b = 0$ limit corresponds to a genuinely known target and yields a pure-target action in the compactified AROS-S family; comparisons at $\xi = 0$ involving the fixed AROS-0 rule must be read separately. The Michelson pooled target has an appreciable between-experiment variance heterogeneity, and its bootstrap comparisons are conditional on that estimated target and the selected tuning choices. The example therefore illustrates the computation but does not establish historical-target certainty, measurement independence, or superiority in estimation accuracy.

Declarations

Availability of Supporting Data

The numerical risk grids, validation calculations, and figure-source files are reproducible from the mathematical model described in this article. The real-data illustration uses the publicly available 1879 speed-of-light measurements of Michelson as reproduced by Stigler [20]; Experiment 5 is analyzed as the current sample, and Experiments 1–4 are used only to construct the historical target. The locked numerical workbook and the revision-specific analysis records can be supplied as supplementary reproducibility materials.

Funding

The authors conducted this research without any funding, grants, or financial support from public, commercial, or not-for-profit organizations.

Conflict of Interest

The authors declare that they have no known competing financial interests, personal relationships, or other conflicts of interest that could have influenced the work reported in this paper.

Authors' Contributions

Iman Jalil Atewi: Conceptualization; Methodology; Formal analysis; Software; Investigation; Visualization; Writing – original draft preparation. **Alaa Khlaif Jiheel**: Method-

Corrected proofs

ology; Validation; Data curation; Supervision; Writing – review & editing. All authors reviewed and approved the final manuscript.

Artificial Intelligence Statement

Artificial intelligence (AI) tools were used to assist language editing, drafting, and refinement of methodological explanations. These tools did not substitute for scientific judgment. The authors retain responsibility for independently verifying the mathematical arguments, data provenance, numerical calculations, interpretations, and final manuscript content.

Publisher's Note

The publisher remains neutral regarding jurisdictional claims in published maps and institutional affiliations.

References

[1] ALoraini, N. M., Hamdaoui, A., Benkhaled, A. (2026). "A new approach for shrinkage estimators of the multivariate normal mean vector under the balanced loss criterion". *Frontiers in Applied Mathematics and Statistics*, 11, 1709968. https://doi.org/10.3389/fams.2025.1709968

[2] Al-Qaragholi, O. J. A., Jiheel, A. K. (2026). "Optimization-oriented double pre-test shrinkage estimators for Pareto reliability under progressive Type-II censoring and precautionary loss". *Control and Optimization in Applied Mathematics*, 11(2), 83–103. https://doi.org/10.30473/coam.2026.77066.1388

[3] Armstrong, T. B., Kline, P., Sun, L. (2025). "Adapting to misspecification". *Econometrica*, 93(6), 1981–2005. https://doi.org/10.3982/ECTA21991

[4] Armstrong, T. B., Kolesár, M. (2020). "Simple and honest confidence intervals in nonparametric regression". *Quantitative Economics*, 11(1), 1–39. https://doi.org/10.3982/QE1199

[5] Cai, T. T., Liang, T., Zhou, H. H. (2015). "Law of log determinant of sample covariance matrix and optimal estimation of differential entropy for high-dimensional Gaussian distributions". *Journal of Multivariate Analysis*, 137, 161–172. https://doi.org/10.1016/j.jmva.2015.02.003

[6] Casella, G., Strawderman, W. E. (1981). "Estimating a bounded normal mean". *The Annals of Statistics*, 9(4), 870–878. https://doi.org/10.1214/aos/1176345527

Corrected proofs

[7] Edelmann, D. (2026). "The optimal standardization factor for the variance of the normal distribution: A historical perspective and recommendations for practice". *Journal of Statistical Theory and Practice*, 20, Article 49. https://doi.org/10.1007/s42519-026-00567-2

[8] Efron, B., Morris, C. (1973). "Stein's estimation rule and its competitors—An empirical Bayes approach". *Journal of the American Statistical Association*, 68(341), 117–130. https://doi.org/10.1080/01621459.1973.10481350

[9] Gray, H., Leday, G. G. R., Vallejos, C. A., Richardson, S. (2018). "Shrinkage estimation of large covariance matrices using multiple shrinkage targets". *arXiv preprint arXiv:1809.08024*. https://doi.org/10.48550/arXiv.1809.08024

[10] Gupta, M., Srivastava, S. (2010). "Parametric Bayesian estimation of differential entropy and relative entropy". *Entropy*, 12(4), 818–843. https://doi.org/10.3390/e12040818

[11] Ikeda, Y., Nakada, R., Kubokawa, T., Srivastava, M. S. (2021). "Linear shrinkage estimation of the variance of a distribution with unknown mean". *Communications in Statistics – Theory and Methods*, 50(9), 2039–2047. https://doi.org/10.1080/03610926.2019.1657457

[12] Kozachenko, L. F., Leonenko, N. N. (1987). "Sample estimate of the entropy of a random vector". *Problems of Information Transmission*, 23(2), 95–101.

[13] Kubokawa, T. (1999). "Shrinkage and modification techniques in estimation of variance and the related problems: A review". *Communications in Statistics – Theory and Methods*, 28(3–4), 613–650. https://doi.org/10.1080/03610929908832317

[14] Kubokawa, T. (2024). "Stein's identities and the related topics: An instructive explanation on shrinkage, characterization, normal approximation and goodness-of-fit". *Japanese Journal of Statistics and Data Science*, 7, 267–311. https://doi.org/10.1007/s42081-023-00239-6

[15] Maatta, J. M., Casella, G. (1990). "Developments in decision-theoretic variance estimation". *Statistical Science*, 5(1), 90–101. https://doi.org/10.1214/ss/1177012263

[16] Misra, N., Singh, H., Demchuk, E. (2005). "Estimation of the entropy of a multivariate normal distribution". *Journal of Multivariate Analysis*, 92(2), 324–342. https://doi.org/10.1016/j.jmva.2003.10.003

Corrected proofs

[17] Ohtani, K. (1999). "Risk performance of a pre-test estimator for normal variance with the Stein-variance estimator under the LINEX loss function". *Statistical Papers*, 40(1), 75–87. https://doi.org/10.1007/BF02927111

[18] Sahni, P., Kumar, R. (2023). "Huntsberger type shrinkage entropy estimator for variance of normal distribution under LINEX loss function". *Reliability: Theory & Applications*, 18(1), 491–499. https://doi.org/10.24412/1932-2321-2023-172-491-499

[19] Stein, C. M. (1964). "Inadmissibility of the usual estimator for the variance of a normal distribution with unknown mean". *Annals of the Institute of Statistical Mathematics*, 16, 155–160. https://doi.org/10.1007/BF02868569

[20] Stigler, S. M. (1977). "Do robust estimators work with real data?". *The Annals of Statistics*, 5(6), 1055–1098. https://doi.org/10.1214/aos/1176343997

[21] Tong, T., Jang, H., Wang, Y. (2012). "James–Stein type estimators of variances". *Journal of Multivariate Analysis*, 107, 232–243. https://doi.org/10.1016/j.jmva.2012.01.019

[22] Tong, T., Wang, Y. (2007). "Optimal shrinkage estimation of variances with applications to microarray data analysis". *Journal of the American Statistical Association*, 102(477), 113–122. https://doi.org/10.1198/016214506000001266

[23] Vasicek, O. (1976). "A test for normality based on sample entropy". *Journal of the Royal Statistical Society: Series B (Methodological)*, 38(1), 54–59. https://doi.org/10.1111/j.2517-6161.1976.tb01566.x

[24] Zellner, A. (1986). "Bayesian estimation and prediction using asymmetric loss functions". *Journal of the American Statistical Association*, 81(394), 446–451. https://doi.org/10.1080/01621459.1986.10478289

[25] Zhang, Y.-Y. (2026). "The empirical Bayes estimators of the variance parameter of the normal distribution with a normal-inverse-gamma prior under Stein's loss function". *Axioms*, 15(2), 127. https://doi.org/10.3390/axioms15020127

Authors' Bio-sketches

Iman Jalil Atewi is an Assistant Lecturer in the Department of Mathematics, College of Education for Pure Sciences, University of Thi-Qar, Iraq. She is also a postgraduate student at the University of Thi-Qar. Her research interests include Mathematical Statistics,

Corrected proofs

Statistical Estimation Theory, and Shrinkage Estimation. Corresponding author. Email: eymanjalil.math@utq.edu.iq.

Alaa Khlaif Jiheel is an Assistant Professor in the Department of Mathematics, College of Education for Pure Sciences, University of Thi-Qar, Iraq. His research interests include Mathematical Statistics, Statistical Estimation Theory, and Shrinkage Estimation. He is the academic supervisor of Iman Jalil Atewi.

Corrected proofs