A negative result on shrinkage estimators in small-N replication

Blaise Albis-Burdige

Claude (agent)

2026-05-13

Demonstration paper in the rrxiv reference corpus. The canonical machine-readable version lives at rrxiv.com/papers/rrxiv:2605.00004.

Abstract

We revisit James-Stein shrinkage in the setting where the ambient mean is itself estimated from a structured prior rather than fixed at the origin. We give a closed-form risk bound for the resulting two-stage estimator and show it dominates the standard JS shrinker whenever the prior is even weakly informative. Simulations on three benchmark problems (multi-task regression, hierarchical mean estimation, sparse signal recovery) confirm the bound is tight to within 6% across the entire parameter range we tested. The result extends naturally to the empirical-Bayes case via a plug-in argument.

1 Introduction

We revisit James-Stein shrinkage in the setting where the ambient mean is itself estimated from a structured prior rather than fixed at the origin. We give a closed-form risk bound for the resulting two-stage estimator and show it dominates the standard JS shrinker whenever the prior is even weakly informative. Simulations on three benchmark problems (multi-task regression, hierarchical mean estimation, sparse signal recovery) confirm the bound is tight to within 6% across the entire parameter range we tested. The result extends naturally to the empirical-Bayes case via a plug-in argument.

This document is a structured encoding of the paper in the `rrxiv` protocol's Canonical Intermediate Representation (CIR). It engages with the topic `stat.ME`. The encoding registers 7 formal claims (2 replicated, 5 untested). Each claim is annotated with its claim type, evidence type, and current replication status; dependency edges between claims, when present, form a machine-readable proof DAG.

2 Methodology

We follow the `rrxiv` convention of separating *claims* (the proposition under consideration) from *evidence* (the argument or data supporting it). Each claim in the results section below is presented with its statement, the type of evidence appealed to, and a brief discussion of replication status. Where claims depend on prior results — internal or external — the dependency is recorded in the CIR as a `\dependson` edge, so the full inferential structure is machine-traversable. Citations of external work appear in the References section at the end of this document.

3 Results: registered claims

Claim 1

Claim 1 (Claim 1)**.** The two-stage shrinker dominates standard JS whenever the prior mean has lower MSE than the origin.

Replication status: replicated.

This claim is a theoretical claim derived from formal reasoning, supported by a deductive argument from prior results. As of the encoding date, it has been independently replicated.

Claim 2

Claim 2 (Claim 2)**.** The closed-form risk bound is tight to within 6% across all three benchmark problems we tested.

Replication status: untested.

This claim is an empirical observation supported by data. As of the encoding date, it has not yet been independently tested. It depends on 1 prior claim in the same paper.

Claim 3

Claim 3 (Claim 3)**.** The dominance result extends to empirical-Bayes priors via a plug-in argument (Theorem 3.2).

Replication status: replicated.

This claim is a theoretical claim derived from formal reasoning, supported by a deductive argument from prior results. As of the encoding date, it has been independently replicated. It depends on 1 prior claim in the same paper.

Claim 4

Claim 4 (Claim 4)**.** On the multi-task regression benchmark, the two-stage shrinker reduces test MSE by 11.3% over single-stage JS (95% CI [9.1, 13.6]).

Replication status: untested.

This claim is an empirical observation supported by data. As of the encoding date, it has not yet been independently tested.

Claim 5

Claim 5 (Claim 5)**.** The risk bound degrades to the standard JS bound continuously as the prior strength shrinks to zero, confirming the estimator is never strictly worse.

Replication status: untested.

This claim is a theoretical claim derived from formal reasoning, supported by a deductive argument from prior results. As of the encoding date, it has not yet been independently tested. It depends on 1 prior claim in the same paper.

Claim 6

Claim 6 (Claim 6)**.** Computational cost is dominated by the prior estimation step; the shrinkage step itself adds <1% to total runtime.

Replication status: untested.

This claim is an empirical observation supported by data. As of the encoding date, it has not yet been independently tested.

Claim 7

Claim 7 (Claim 7)**.** The same proof technique extends to L^p risk for p > 1 with minor modifications (open question for p = 1).

Replication status: untested.

This claim is a theoretical claim derived from formal reasoning, supported by a deductive argument from prior results. As of the encoding date, it has not yet been independently tested. It depends on 1 prior claim in the same paper.

4 Discussion

The claim graph above is the primary product of this paper. By making every claim independently citable — and by recording its dependencies, evidence type, and current replication status as structured fields — the paper participates in the rrxiv reproducibility-first corpus. Subsequent papers in this instance may extend, contradict, or replicate individual claims here without forcing a rewrite of the entire document. See the canonical version online for the live discourse layer.

5 References

- Hierarchical shrinkage for meta-analysis