

Adaptive bias-variance trade-off: a robust Bayesian approach

Paul Delatte[†]

October 23, 2026

The latest version is available [here](#).

Abstract

We tackle the problem of combining an unbiased but noisy estimator with a biased but informative estimator when the bias is unknown. As a main application, we aim to improve a causal effect estimator from a small-scale randomized controlled trial (RCT) using an estimator from a large-scale observational study. We show that the problem can be solved comprehensively within a robust Bayesian framework in which Nature adversarially chooses a prior on the parameter and the bias. We prove that a unique robust Bayes estimator exists for any fixed set of priors. To achieve adaptation when the set is unknown, we develop both a minimax regret approach that minimizes worst-case regret relative to an oracle knowing the set and an empirical Bayes approach that calibrates the set directly from the data. The solutions provide a transparent foundation for the problem, nesting prominent adaptive estimators as special cases and yielding computationally advantageous extensions from simple Bayesian primitives. We illustrate the practical utility of the results by combining estimators of the effect of class-size reduction on student test scores obtained from small-scale RCTs in Tennessee (Project STAR) and large-scale observational studies in California and Connecticut.

KEYWORDS: bias-variance trade-off; adaptive estimation; robust Bayes; causal inference

[†]University of Southern California, delatte@usc.edu. We thank Hyungsik Roger Moon, Timothy Armstrong, Geert Ridder, Nicolas Lambert, Gilles Stoltz, José Luis Montiel Olea, Yuehao Bai, Mengsi Gao, Julian Duggan, Vittorio Bassi, Xun Tang, and Yingyao Hu for discussions that helped improve this paper.

1 Introduction

1.1 Motivation and results

Causal inference is characterized by a fundamental trade-off between the likelihood of capturing true causal effects and the cost of doing so. This trade-off has led to the separation of randomized experiments and observational studies on two sides of the economist’s toolkit. Randomized controlled trials (RCTs) are generally viewed as the *gold standard* for the identification of causal effects. However, RCTs are known to be costly and time-intensive with hard-to-answer ethical and legal questions as soon as treatments are given or withheld. These budget and ethical constraints often force economists to run small-scale RCTs with limited sample sizes, leading to wide confidence intervals and underpowered tests. When large-scale experiments are not feasible, economists regularly leverage a vast arsenal of quasi-experimental methods to infer causal effects from observational data at a fraction of the cost. While these observational datasets are often large and granular enough to yield sufficient statistical power, quasi-experimental methods lack the credibility of RCTs in identifying causal effects. Our work aims to answer a simple question: by combining an estimator from a small-scale RCT and an estimator from a large-scale observational study, can we get the best of both worlds – that is, statistically precise estimates that remain causally credible?

In econometric terms, the problem reduces to combining two estimators that stand on opposite ends of a simple bias–variance trade-off: an unbiased, large-variance estimator from the RCT and a biased, small-variance estimator from the observational study. The bias in this causal inference context stems both from the failure of the quasi-experimental methods to capture causal effects (i.e., internal validity) and from the difference in causal effects across the experimental and observational populations (i.e., external validity). While the sampling variances can be consistently estimated from each dataset (and treated as known in practice), the bias is fundamentally unobserved and imperfectly captured by the difference between the two estimators. The core challenge is then one of adaptation: how to optimally navigate the trade-off without imposing unwarranted assumptions on the bias? Our main contribution in this paper is to frame and solve this adaptive trade-off problem within a robust Bayesian framework. The encompassing solutions that we obtain, nesting several well-known adaptive estimators as special cases, demonstrate that this approach can be used as a unifying and transparent foundation for the problem.

The main idea behind our robust Bayesian approach is to view Nature as adversarially choosing a prior on the causal effect and the bias from a set of possible priors. The objective is then to find, if it exists, an estimator of the causal effect with minimal integrated risk over the worst-case prior chosen by Nature. Our first contribution is to show that this problem admits a unique and simple solution. By leveraging asymptotic and equivariance arguments, we show that the problem reduces to finding a robust Bayes estimator in a normal mean model where Nature only picks a prior on the bias. From this reduction, the existence of a unique solution follows from standard minimax arguments. That is, for any set of priors on the bias, there exists a unique robust Bayes estimator of the causal effect. We also obtain from the reduction a computationally tractable expression of the solution in terms of the robust Bayes estimator of a normal mean. Under mild regularity conditions, these estimators can be recovered as the Bayes estimators for the least favorable priors in the corresponding problems.

Our second contribution is to solve the adaptation problem directly within the robust Bayesian framework. We show that adapting to the unknown bias can be done in two distinct ways. We first develop a minimax regret approach where we let Nature adversarially pick the set of priors. For this purpose, we define a notion of regret against the robust Bayes estimator when the set of priors is known (i.e., the oracle estimator) and aim to find the estimator that minimizes the worst-case regret. We show that under regularity conditions on the class of sets that Nature can pick from, the problem reduces to a simple frequentist minimax game that can be solved by numerical methods. We prove that the conditions are satisfied for the sets of priors with bounded supports and the sets of priors with bounded second-order moments. The first case recovers the adaptive estimator of [Armstrong, Kline, and Sun \(2025\)](#). The second case gives a new minimax regret estimator based on a much simpler computational routine. We then develop a second approach to adaptation that favors expediency over worst-case regret guarantees. The starting point differs: we now view the original robust Bayesian decision problem as being solved by an oracle for an unknown but fixed set of priors. We then solve the adaptation problem by simply calibrating the set of priors from the data. Under mild regularity conditions, we show that this is equivalent to an empirical Bayes approach with respect to a family of least favorable priors. When the family of least favorable priors is parametric, as for sets of priors with bounded second-order moments, we obtain explicit closed-form solutions for the adaptive estimators. In the multivariate

extension of the problem, we show that this recovers exactly the optimal estimator of [Green and Strawderman \(1991\)](#). To extend the approach to nonparametric families of least favorable priors, we build an algorithmic solution that leverages the fact that least favorable priors can be well approximated by discrete mixtures under general conditions. This solution significantly extends the set of Bayesian primitives from which adaptive estimators can be derived.

Our robust Bayesian approach thus provides a transparent foundation for the combination problem, anchoring and connecting adaptive solutions to explicit, interpretable sets of priors. The approach not only recovers prominent estimators as special cases ([Armstrong, Kline, and Sun \(2025\)](#), [Green and Strawderman \(1991\)](#)), but also yields new solutions with direct computational advantages. To compare these adaptive estimators, we perform extensive Monte Carlo simulations based on designs capturing standard failures of quasi-experimental methods. The simulation results confirm that all the estimators from the robust Bayes approach adapt successfully, outperforming conventional benchmarks such as pretesting estimators. The second outcome of the simulations is to uncover a clear grouping of the solutions based on the adaptation approach: the minimax regret estimators favor variance reduction while the robust empirical Bayes estimators favor bias protection. The Bayesian primitives only have secondary effects on performance. This sharp clustering informs our main recommendation for practitioners: rather than relying on a single adaptive estimator, researchers should report results based on solutions derived from both adaptation approaches (preferably derived from the same Bayesian primitives).

We demonstrate the practical relevance of these results by reanalyzing the effect of class-size reduction on student test scores using experiments conducted in Tennessee from 1985 to 1989 (Project STAR). Because Project STAR is a large-scale RCT, we can perform a simple simulation exercise where we artificially create small-scale RCTs by subsampling the original data. This gives us direct control over the noise of the RCT estimators and allows us to operate as if we knew the true causal effect of the treatment by fixing it as the estimated effect in the full-sample dataset. We then combine the estimators in the small-scale RCTs with estimators from two large-scale observational studies conducted in California and Connecticut, respectively. Because we know the true effect, the two observational studies correspond to two opposite scenarios: a small-bias observational study (California) and a large-bias observational study (Connecticut). The results of the combination exercise

corroborate the overall performance of the adaptive estimators and confirm the practicality of our recommendation. With the California study, all the solutions yield substantial precision gains with a natural advantage for the minimax regret ones. With the Connecticut study, the robust empirical Bayes estimators guard against severe bias distortion, underlining clearly the risk of using the minimax regret estimators in this case. Evaluating both types of adaptive estimators side by side provides a simple and intuitive way to transparently and efficiently navigate the bias–variance trade-off.

All the results that we derive in this paper generalize beyond the causal inference application. As soon as two estimators are governed by a bias–variance trade-off similar to the one we study, they can be combined using the robust Bayesian approach introduced in this paper. In the Appendix, we show that the results generalize to the case where the two estimators are dependent by simple reweighting according to the covariance structure. This allows us to address with the exact same methods the problem of combining two estimators obtained from the same data but from two nested models. The solutions then automatically adapt to the misspecification bias in the more parametric model. In an Online Supplementary Appendix, we consider some of these extensions, including a macroeconomic application for the estimation of dynamic causal effects based on local projection estimators and structural vector autoregression estimators. Similarly, our results generalize beyond the univariate case. We show in the Appendix how these extensions are obtained and consider as a direct application the recovery of the estimator in [Green and Strawderman \(1991\)](#). Additional possible extensions, including variance estimation and the problem of inference, are introduced and discussed in different parts of the paper but their resolution is left for future research.

The rest of the paper is organized as follows. Section [2](#) introduces our robust Bayesian approach and contains the existence result for robust Bayes estimators (Proposition [2.1](#)). The minimax regret approach to adaptation is introduced and solved in Section [3](#) with Theorem [3.1](#) serving as the main proof mechanism. The empirical Bayes approach to adaptation is introduced and solved in Section [4](#). The Monte Carlo simulations are collected in Section [5](#) and Appendix [C](#). The main empirical application is tackled in Section [6](#). All the proofs are collected in Appendix [A](#). This Appendix also contains the multivariate extension (Appendix [A.3](#)), including Proposition [A.4](#) for the equivalence with [Green and Strawderman \(1991\)](#), and the misspecification extension (Appendix [A.4](#)).

1.2 Literature review

The problem of adaptively combining estimators bound by some bias–variance trade-off has a long history in econometrics and statistics. Abstracting from the causal inference context, the problem has been tackled in three general ways: a pretesting approach, a plug-in approach, and a decision-theoretic approach. The pretesting approach is based on the simple idea of running a hypothesis test before choosing the estimator to use – see [Bancroft and Han \(1977\)](#) and [Giles and Giles \(1993\)](#). This approach leads to hard-thresholding estimators that are intuitive but underperform – the mean squared error is inflated for values of the bias where the hard threshold hits (see also [Sclove, Morris, and Radhakrishnan \(1972\)](#)). The plug-in approach is also heuristic but leads to estimators with better overall performance. The idea is to start from a simple linear combination of the initial estimators and to find the weights that perform best. An intuitive solution is to plug in weights that minimize the mean squared error (MSE) of the combination estimator. This strategy leads to the adaptive estimator of [Green and Strawderman \(1991\)](#) with direct antecedents in [Stein \(1966\)](#) and extensions in [Judge and Mittelhammer \(2004\)](#). In the univariate case, [de Chaisemartin and D’Haultfœuille \(2020\)](#) considers weights obtained directly by empirical MSE minimization. The expression of the solution in [Green and Strawderman \(1991\)](#) leads, moreover, to ad hoc alternatives based on different shrinkage functions. An example of this is the solution considered in [Chen, Li, Zhang, and Ye \(2025\)](#) based on soft-thresholding – see also [Armstrong, Kline, and Sun \(2025\)](#). The decision-theoretic approach improves on the plug-in one by providing a proper foundation from which adaptive estimators with inbuilt optimality properties can be derived. A main strategy framed in terms of minimax regret can be traced back to the elegant results of [Bickel \(1984\)](#), building on earlier contributions of [Hodges and Lehmann \(1952\)](#). These results are extended significantly in [Armstrong, Kline, and Sun \(2025\)](#) to cover the entire bias–variance trade-off. Our results further extend this minimax regret approach by bringing it to the Bayesian realm. This extension is of interest in itself by enlarging the set of decision-theoretic primitives that can be imposed on the bias, but it is also of interest conceptually by providing a simple bridge with the plug-in approach through an empirical Bayes lens with respect to a family of least favorable priors.

In the particular context of causal inference, the problem of merging RCTs with observational studies extends beyond the combination problem for estimation. Many proposed solutions directly attempt to merge the datasets before estimation by modeling the individual

populations. These solutions naturally shift the focus from combination to identification – see, for instance, [Kallus, Puli, and Shalit \(2018\)](#), [Athey, Chetty, and Imbens \(2020\)](#), and [Ghassami, Liu, Yang, Richardson, Shpitser, and Tchetgen \(2022\)](#). These options belong to a larger literature on data fusion – see [Bareinboim and Pearl \(2016\)](#) in a causal inference context and [Ridder and Moffitt \(2007\)](#) more generally. Naturally, the problem of merging RCTs and observational studies can also be tackled as a combination problem for estimation. Most works that have tackled this problem in causal inference have leaned heavily on the plug-in approach, either applying or extending the existing generic methods – see, for instance, [Rosenman, Basse, Owen, and Baiocchi \(2023\)](#), [Chen, Li, Zhang, and Ye \(2025\)](#), [Wu, Salazar, Green, and Blei \(2026\)](#), and [Lin, Bickel, and Ding \(2026\)](#). As far as we know, very few works have attempted to apply or extend the decision-theoretic approach in this context. Two exceptions are [Epanomeritakis and Viviano \(2026\)](#), with the caveat that they address the experimental design problem rather than estimation, and our own work. For a broader review of the combination problem in causal inference, we refer to [Rosenman \(2025\)](#).

From a technical viewpoint, our work relates to a large literature on robust Bayes estimators and the normal mean problem – see [Berger \(1985\)](#), [Johnstone \(2019\)](#), and [Delatte \(2026a\)](#). The connection between empirical Bayes and robust Bayes estimators also has antecedents – see, in particular, [Berger and Berliner \(1984\)](#) and [Feldman \(1991\)](#). The robust Bayesian approach has also been applied to other econometric problems – see, for instance, [Giacomini, Kitagawa, and Read \(2026\)](#). For most of our solutions, we need to numerically approximate least favorable priors in the normal mean problem. An extensive literature on the problem exists – see [Nelson \(1966\)](#), [Kempthorne \(1987\)](#), [Chamberlain \(2000\)](#), [Noubiap and Seidel \(2001\)](#), [Montiel Olea and Zubova \(2026\)](#), to name only a few. Additional references on the problem are introduced in [Appendix B](#).

Notation. Let $p \in \mathbb{N}$. We denote the one-point compactification of the Euclidean space $\mathbb{R}^p$ by $\mathbb{R}^p \cup \{\infty\}$. We denote by $\mathcal{N}(0, 1)$ the standard normal distribution on $\mathbb{R}$ and by $\phi(\cdot)$ its density. We extend the notation to the multivariate normal distribution on $\mathbb{R}^p$. For any Euclidean subspace $\Theta \subseteq \mathbb{R}^p$, we denote by $\mathcal{M}_1(\Theta)$ the space of Borel probability measures on Θ and by δ_θ the Dirac measure at $\theta \in \Theta$. For any random variable X in a measurable space $\mathcal{X}$ and any measurable function $g: \mathcal{X} \rightarrow \mathbb{R}^p \cup \{\infty\}$, we call g and $g(X)$ an estimator and use both notations interchangeably. We denote by $\mathcal{D}_p(\mathcal{X}) := \{g: \mathcal{X} \rightarrow \mathbb{R}^p \cup \{\infty\} : g \text{ measurable}\}$ the set of all estimators based on X for the estimation of a parameter in any Euclidean subset

$\Theta \subseteq \mathbb{R}^p$ and write $\mathcal{D}_1(\mathcal{X}) = \mathcal{D}(\mathcal{X})$. We denote the risk of g for the estimation of $\theta \in \Theta$ based on the observation X by

$$R(g, \theta) = R(g, \theta; L^2) = \mathbb{E} [\|g(X) - \theta\|^2],$$

where $\|\cdot\|$ is the Euclidean norm on $\mathbb{R}^p$ with $\|g(x) - \theta\|^2 = \infty$ whenever $g(x) = \infty$. We denote the integrated risk of g with respect to a distribution $P \in \mathcal{M}_1(\Theta)$ on θ by

$$B(g, P) = B(g, P; L^2) = \int_{\Theta} \mathbb{E} [\|g(X) - \theta\|^2] dP(\theta) = \int_{\Theta} R(g, \theta) dP(\theta),$$

and call P a prior on the parameter. An estimator $g^* \in \mathcal{D}_p(\mathcal{X})$ is said to be frequentist minimax over Θ if it satisfies $\sup_{\theta \in \Theta} R(g^*, \theta) = \inf_{g \in \mathcal{D}_p(\mathcal{X})} \sup_{\theta \in \Theta} R(g, \theta)$. For any set $\Gamma \subseteq \mathcal{M}_1(\Theta)$, an estimator $g^* \in \mathcal{D}_p(\mathcal{X})$ is said to be robust Bayes over Γ (or Γ -minimax) if it satisfies $\sup_{\pi \in \Gamma} B(g^*, \pi) = \inf_{g \in \mathcal{D}_p(\mathcal{X})} \sup_{\pi \in \Gamma} B(g, \pi)$. For any $\Gamma \subseteq \mathcal{M}_1(\Theta)$, a prior $\pi^* \in \Gamma$ is said to be a least favorable prior over Γ (or Γ -maximin prior) if it satisfies $\inf_{g \in \mathcal{D}_p(\mathcal{X})} B(g, \pi^*) = \sup_{\pi \in \Gamma} \inf_{g \in \mathcal{D}_p(\mathcal{X})} B(g, \pi)$. We denote by $R^*(\Theta) = \inf_{g \in \mathcal{D}_p(\mathcal{X})} \sup_{\theta \in \Theta} R(g, \theta)$ the frequentist minimax risk over Θ and by $B^*(\Gamma) = \inf_{g \in \mathcal{D}_p(\mathcal{X})} \sup_{\pi \in \Gamma} B(g, \pi)$ the robust Bayes risk over Γ (or Γ -minimax risk).

2 Existence and uniqueness of robust Bayes estimators

2.1 Gaussian limit experiment and equivariance reduction

We consider the estimation of some unknown causal effect $\tau \in \mathbb{R}$ based on two independent statistics: an unbiased but noisy estimator $\hat{\tau}_{\text{RCT}}$ from a small-scale RCT and a biased but informative estimator $\hat{\tau}_{\text{OS}}$ from a large-scale observational study. We denote by $b \in \mathbb{R}$ the unknown bias of $\hat{\tau}_{\text{OS}}$ and remain agnostic about its origin (e.g., unobserved confounding, population differences). The trade-off between the two estimators can be parsimoniously captured by a bivariate Gaussian limit experiment given by

$$\begin{bmatrix} \hat{\tau}_{\text{RCT}} \\ \hat{\tau}_{\text{OS}} \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \tau \\ \tau + b \end{bmatrix}, \begin{bmatrix} \sigma_{\text{RCT}}^2 & 0 \\ 0 & \sigma_{\text{OS}}^2 \end{bmatrix} \right),$$

where the asymptotic sampling variances σ_{RCT}^2 and σ_{OS}^2 satisfy $0 < \sigma_{\text{OS}}^2 < \sigma_{\text{RCT}}^2$ and are treated as known¹. For convenience, we can isolate all the information about the nuisance parameter $b \in \mathbb{R}$ by defining the difference

$$D := \hat{\tau}_{\text{OS}} - \hat{\tau}_{\text{RCT}}.$$

We can then work more simply with the pair of statistics $(\hat{\tau}_{\text{RCT}}, D)$ whose distribution is given by

$$\begin{bmatrix} \hat{\tau}_{\text{RCT}} \\ D \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \tau \\ b \end{bmatrix}, \begin{bmatrix} \sigma_{\text{RCT}}^2 & -\sigma_{\text{RCT}}^2 \\ -\sigma_{\text{RCT}}^2 & \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2 \end{bmatrix} \right).$$

When the location of the causal effect $\tau \in \mathbb{R}$ is arbitrary, it is natural to restrict attention to estimators $g: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R} \cup \{\infty\}$ that are equivariant under translations in the sense that $g(\hat{\tau}_{\text{RCT}} + a, D) = g(\hat{\tau}_{\text{RCT}}, D) + a$ for any $a \in \mathbb{R}$. Under this property, an estimator g can be uniquely represented as

$$g(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + \mathcal{G}(D),$$

for some measurable function $\mathcal{G}: \mathbb{R} \rightarrow \mathbb{R} \cup \{\infty\}$. In Remark 2.4, we show that, under general conditions, equivariance can be assumed without loss of generality for our problem.

Let $G \subseteq \mathcal{D}(\mathbb{R}^2)$ then denote the set of all estimators $g: \mathbb{R}^2 \rightarrow \mathbb{R} \cup \{\infty\}$ that are equivariant under translation. By leveraging the joint normality of $\hat{\tau}_{\text{RCT}}$ and D , we can show that the risk of any estimator $g \in G$ satisfies a strong invariance property. Indeed, for any $g \in G$ and any $(\tau, b) \in \mathbb{R}^2$, we have that

$$\begin{aligned} R(g, \tau) &= \mathbb{E} \left[\mathbb{E} \left[\left(\hat{\tau}_{\text{RCT}} - \tau + \mathcal{G}(D) \right)^2 \mid D \right] \right] \\ &= \mathbb{E} \left[w\sigma_{\text{OS}}^2 + \left(-w(D - b) + \mathcal{G}(D) \right)^2 \right], \end{aligned}$$

where $w := \sigma_{\text{RCT}}^2 / (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. By introducing the measurable function $h: \mathbb{R} \rightarrow \mathbb{R} \cup \{\infty\}$ defined by $h(D) = D - w^{-1}\mathcal{G}(D)$, we then obtain that

$$R(g, \tau) = w\sigma_{\text{OS}}^2 + w^2 R(h, b),$$

where $R(h, b)$ is the risk of h for the estimation of b based on $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. In

¹In practice, we plug in consistent estimators of the asymptotic variances derived from each sample.

other words, the risk of any estimator $g \in G$ is invariant with respect to the causal effect τ , depends only on the unknown bias b , and can be simply expressed in terms of the risk of an estimator in a simple univariate normal mean problem.

By elementary properties of the integral, the invariance of $R(g, \tau)$ directly extends to the invariance of the integrated risk $B(g, P)$ with respect to the marginal distribution of τ for any joint prior P on $(\tau, b) \in \mathbb{R}^2$. We collect this invariance property of the integrated risk in the lemma below. In what follows, we denote by $\Lambda \subseteq \mathcal{M}_1(\mathbb{R}^2)$ a set of joint probability distributions for (τ, b) and by $\Gamma = \{P(\mathbb{R} \times \cdot) : P \in \Lambda\} \subseteq \mathcal{M}_1(\mathbb{R})$ the set of marginal distributions on the bias obtained from Λ .

Lemma 2.1. *For any joint prior $P \in \Lambda$ and any estimator $g \in G$ such that $g(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + \mathcal{G}(D)$ for some measurable function $\mathcal{G}: \mathbb{R} \rightarrow \mathbb{R} \cup \{\infty\}$, it holds that*

$$B(g, P) = w\sigma_{\text{OS}}^2 + w^2 B(h, \pi),$$

where $\pi = P(\mathbb{R} \times \cdot)$ is the marginal of P with respect to b and $B(h, \pi)$ is the integrated risk over π of the estimator $h(D) = D - w^{-1}\mathcal{G}(D)$ based on the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$.

Remark 2.1. The model can be extended in several ways without complications. Multivariate extensions are introduced in Appendix A.3. Generalization to dependent estimators is covered in Appendix A.4. The extension to other losses than the L^2 loss is possible but less straightforward as many of the computational simplifications break down – see Remark A.1. Similarly, relaxing normality and knowledge of the variances is possible, but it leads to a number of new computational challenges that go beyond the scope of this paper.

2.2 Robust Bayes estimators: existence, uniqueness, computation

The limit experiment and the equivariance reduction create a direct bridge between the combination problem and the univariate normal mean problem. This bridge opens the way for general solutions based on a robust Bayesian approach. This decision-theoretic framework is general, interpretable, and modular – it nests both the Bayesian and frequentist minimax criteria as special cases (see Remark 2.3) and allows us to entertain rich sets of beliefs on the causal effect and the bias without committing to any single subjective prior.

For any fixed set $\Lambda \subseteq \mathcal{M}_1(\mathbb{R}^2)$ and any estimator $g \in G$, we thus view Nature as adversarially choosing a prior $P \in \Lambda$ on the causal effect and the bias (τ, b) that maximizes

the integrated risk $B(g, P)$. From Lemma 2.1, the effect of P on $B(g, P)$ only depends on the marginal $\pi = P(\mathbb{R} \times \cdot)$. Therefore, the problem for Nature reduces to the choice of an adversarial prior π on the bias b that maximizes $B(h, \pi)$ and hence $B(g, P)$. For this reason, we can take the set $\Gamma = \{P(\mathbb{R} \times \cdot) : P \in \Lambda\} \subseteq \mathcal{M}_1(\mathbb{R})$ of priors on the bias as a primitive and write $B(g, P) = B(g, \pi)$ when $g \in G$. The choice of an estimator in G can then be based on minimizing the worst-case integrated risk $\sup_{\pi \in \Gamma} B(g, \pi)$. That is,

$$g_\Gamma^* \in \arg \min_{g \in G} \sup_{\pi \in \Gamma} B(g, \pi).$$

We call an estimator g_Γ^* that satisfies this property a robust Bayes estimator of τ over Γ (and, more simply, a robust Bayes estimator when there is no risk of confusion). As discussed in Remark 2.4, the restriction from $\mathcal{D}(\mathbb{R}^2)$ to G is without loss of generality as soon as no location constraint is imposed on the marginal priors over the causal effect.

From the reduction of Lemma 2.1, we can show that there exists a unique robust Bayes estimator g_Γ^* for any fixed set of priors $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ on the bias. This leads to Proposition 2.1 below, which is proved in Appendix A.1 using standard minimax arguments and the direct connection with the univariate normal mean problem.

Proposition 2.1. *For any set $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$, the robust Bayes estimator of τ over Γ exists, is unique, and takes the form*

$$g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w\left[\hat{\tau}_{\text{OS}} - h_\Gamma^*(D)\right],$$

where $h_\Gamma^* = \arg \min_{h \in \mathcal{D}(\mathbb{R})} \sup_{\pi \in \Gamma} B(h, \pi)$ is the robust Bayes estimator of b over Γ based on the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$.

The robust Bayes estimator of the causal effect thus admits a simple expression as the inverse-variance weighted combination of $\hat{\tau}_{\text{RCT}}$ and a bias-corrected version of $\hat{\tau}_{\text{OS}}$ given by $\hat{\tau}_{\text{OS}} - h_\Gamma^*(D)$. The expression of g_Γ^* in terms of the robust Bayes estimator h_Γ^* in the univariate normal mean problem then provides a direct way to compute g_Γ^* for different choices of Γ . When Γ is convex, a minimax theorem holds for the univariate normal mean problem – see Johnstone (2019) and Delatte (2026a) – and guarantees that $\sup_{\pi \in \Gamma} B(h_\Gamma^*, \pi) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{\pi \in \Gamma} B(h, \pi) = \sup_{\pi \in \Gamma} \inf_{h \in \mathcal{D}(\mathbb{R})} B(h, \pi)$. When, moreover, there exists a unique least favorable prior π_Γ^* over Γ , then we can recover h_Γ^* as the Bayes

estimator with respect to π_Γ^* . In this case, the estimator h_Γ^* is simply given as the posterior mean, that is,

$$h_\Gamma^*(D) = \int_{\mathbb{R}} b d\pi_\Gamma^*(b \mid D),$$

where $\pi_\Gamma^*(\cdot \mid D)$ denotes the posterior distribution of b given D for the prior π_Γ^* . From the normality of D , we then obtain explicit expressions of h_Γ^* in terms of the least favorable prior π_Γ^* . These explicit formulations can be either exactly computed or approximated depending on the structure of the least favorable prior π_Γ^* . Existence results for particular choices of Γ can then be converted into explicit expressions for g_Γ^* – see Section 2.4 for some examples.

Remark 2.2. Under the same conditions for the existence of a least favorable prior, we also have access to a general expression for g_Γ^* that leverages Tweedie’s formula and highlights a different side of the robust Bayes estimator as a simple correction to $\hat{\tau}_{\text{RCT}}$ based on the marginal score function for the difference D . This formulation is collected in Corollary A.1 in Appendix A.1.

Remark 2.3. The robust Bayesian optimality criterion encompasses both the Bayesian and the frequentist minimax ones for appropriate choices of Γ . The reduction to the Bayesian one under singletons is immediate. The reduction to the frequentist minimax one follows whenever Γ contains all Dirac measures on $\mathbb{R}$. If we restrict the parameter space for the bias by taking $b \in \mathcal{B}$ for some set $\mathcal{B} \subset \mathbb{R}$, then the frequentist minimax reduction holds as soon as Γ contains all Dirac measures on $\mathcal{B}$. We use these facts to state a simple minimax equivalence in Corollary A.2 in Appendix A.1. These results extend to $\mathcal{T} \times \mathcal{B}$ without change if one entertains a priori constraints on the effect $\tau \in \mathcal{T}$ for some set $\mathcal{T} \subset \mathbb{R}$.

Remark 2.4. To obtain Lemma 2.1 and Proposition 2.1, we focused on equivariant estimators, leading to the restriction from $\mathcal{D}(\mathbb{R}^2)$ to G . For a large collection of sets Λ , this restriction is without loss of generality in the sense that $\min_{g \in G} \sup_{P \in \Lambda} B(g, P) = \min_{g \in \mathcal{D}(\mathbb{R}^2)} \sup_{P \in \Lambda} B(g, P)$. In other words, under mild invariance conditions on the set of priors, there always exists a robust Bayes estimator that is equivariant. This holds, for instance, when $\Lambda = \mathcal{M}_1(\mathbb{R}) \otimes \Gamma$ where $\mathcal{M}_1(\mathbb{R}) \otimes \Gamma := \{\lambda \otimes \pi : \lambda \in \mathcal{M}_1(\mathbb{R}), \pi \in \Gamma\}$ for any given set $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$. In this case, Corollary A.2 and an application of the Hunt–Stein theorem directly yield the equality of the robust Bayes risks over G and over $\mathcal{D}(\mathbb{R}^2)$. The same argument extends to more general sets, including dependent priors, as soon as these sets

do not impose restrictions on the location of the causal effect. It should be noted that these results do not imply that all robust Bayes estimators in $\mathcal{D}(\mathbb{R}^2)$ are equivariant. In particular, the uniqueness of the robust Bayes estimator proved in Proposition 2.1 only holds over G . While uniqueness extends to $\mathcal{D}(\mathbb{R}^2)$ under mild regularity conditions in the univariate case, it is generally false in dimensions strictly greater than 2.

2.3 Sets of priors and bias–variance trade-off

The previous existence result (Proposition 2.1) maps a set of priors on the bias onto a single robust Bayes estimator of the causal effect. From a statistical viewpoint, the choice of a set of priors is the exact mechanism by which the bias–variance trade-off is navigated. To make this clear, we first show that the extreme sets $\Gamma_0 = \{\delta_0\}$ and $\Gamma_\infty = \mathcal{M}_1(\mathbb{R})$ lead to robust Bayes estimators that stand at the two endpoints of the bias–variance trade-off. We then show that any set Γ such that $\Gamma_0 \subseteq \Gamma \subseteq \Gamma_\infty$ leads to a robust Bayes estimator g_Γ^* that is a data-dependent affine combination of $g_{\Gamma_0}^*$ and $g_{\Gamma_\infty}^*$.

An application of Proposition 2.1 yields simple expressions for the robust Bayes estimators of τ over Γ_0 and Γ_∞ . Indeed, by directly plugging in the robust Bayes estimators of b over Γ_0 and Γ_∞ , we obtain that

$$g_{\Gamma_0}^*(\widehat{\tau}_{\text{RCT}}, D) = (1 - w)\widehat{\tau}_{\text{RCT}} + w\widehat{\tau}_{\text{OS}},$$

and

$$g_{\Gamma_\infty}^*(\widehat{\tau}_{\text{RCT}}, D) = \widehat{\tau}_{\text{RCT}}.$$

The robust Bayes estimator $g_{\Gamma_0}^*$ is thus the inverse-variance weighted estimator of τ when $b = 0$ and corresponds to the maximum likelihood estimator in this case. The robust Bayes estimator $g_{\Gamma_\infty}^*$ is simply the RCT estimator, which also corresponds to the frequentist minimax estimator over $\mathbb{R}^2$ by Corollary A.2. We show in the following proposition that these two estimators stand at opposite ends of the bias–variance trade-off: over all equivariant estimators, $g_{\Gamma_0}^*$ is the estimator with minimal variance and $g_{\Gamma_\infty}^*$ is the estimator with minimal worst-case bias.

Proposition 2.2. 1. For any estimator $g \in G$ and any $(\tau, b) \in \mathbb{R}^2$,

$$\text{Var}(g(\hat{\tau}_{\text{RCT}}, D)) \geq \text{Var}(g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D)) = \frac{1}{\sigma_{\text{RCT}}^{-2} + \sigma_{\text{OS}}^{-2}}.$$

Moreover, $g_{\Gamma_0}^*$ is the unique estimator in G that is unbiased at $b = 0$ and achieves this bound.

2. For any estimator $g \in G$ and any $\tau \in \mathbb{R}$,

$$\sup_{b \in \mathbb{R}} \left| \text{Bias}(g(\hat{\tau}_{\text{RCT}}, D)) \right| \geq \sup_{b \in \mathbb{R}} \left| \text{Bias}(g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D)) \right| = 0.$$

Moreover, $g_{\Gamma_\infty}^*$ is the unique estimator in G that achieves this bound.

For any set Γ such that $\Gamma_0 \subseteq \Gamma \subseteq \Gamma_\infty$, the robust Bayes estimator g_Γ^* over Γ falls in between these two extreme ends of the bias–variance trade-off. To make this clear, we first break down the mean squared risk of g_Γ^* – this leads to Proposition 2.3. We then show that g_Γ^* can be expressed as an affine combination of the two extremes $g_{\Gamma_0}^*$ and $g_{\Gamma_\infty}^*$ – this is collected in Proposition 2.4 below. Both propositions are proved in Appendix A.1.

Proposition 2.3. For any $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ and any $(\tau, b) \in \mathbb{R}^2$, the bias and variance of the robust Bayes estimator g_Γ^* are respectively given by

$$\text{Bias}(g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D)) = -w \text{Bias}(h_\Gamma^*(D)),$$

and

$$\text{Var}(g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D)) = \text{Var}(g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D)) + w^2 \text{Var}(h_\Gamma^*(D)).$$

Proposition 2.4. For any $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$, the robust Bayes estimator g_Γ^* rewrites as the affine combination

$$g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = W(D)g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D) + (1 - W(D))g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D),$$

with data-dependent weight $W(D) = 1 - h_\Gamma^*(D)/D$. Moreover, the combination is convex in the sense that $0 \leq W(D) \leq 1$ if and only if $h_\Gamma^*(D)$ shrinks D toward the origin.

Remark 2.5. When $\Gamma = \{\delta_c\}$ for $c \in \mathbb{R}$, we recover a naive form of bias correction. The robust Bayes estimator for the normal mean problem is $h_\Gamma^*(D) = c$ so that the robust Bayes

estimator g_{Γ}^* is the inverse-variance weighted combination of $\widehat{\tau}_{\text{RCT}}$ and the bias-corrected estimator $\widehat{\tau}_{\text{OS}} - c$.

2.4 Examples: moment bound, support bound, contamination

For illustration and as a preliminary step towards adaptation, we consider in this section three general classes of sets of priors and derive explicit expressions for the robust Bayes estimators over these sets. These classes take a parametric form $\{\Gamma(\eta) : \eta \in S\}$, where $\eta \in S$ is some univariate state with $S \subseteq \mathbb{R}$, and cover the entire bias–variance trade-off in the sense that $\bigcap_{\eta \in S} \Gamma(\eta) = \{\delta_0\}$ and $\text{cl}(\bigcup_{\eta \in S} \Gamma(\eta)) = \mathcal{M}_1(\mathbb{R})$. In each case, the scalar parameter $\eta \in S$ becomes the mechanism through which the bias–variance trade-off is arbitrated. We define these classes by considering, respectively:

1. all priors satisfying the moment constraint $\mathbb{E}_{\pi}[b^2] \leq M$ for $M \geq 0$, that is,

$$\Gamma_1(M) = \left\{ \pi \in \mathcal{M}_1(\mathbb{R}) : \mathbb{E}_{\pi}[b^2] \leq M \right\};$$

2. all priors supported on $[-m, m]$ for $m \geq 0$, that is,

$$\Gamma_2(m) = \left\{ \pi \in \mathcal{M}_1(\mathbb{R}) : \text{supp}(\pi) \subseteq [-m, m] \right\};$$

3. the Huber ε -contamination class with a Dirac measure at 0 for $\varepsilon \in (0, 1]$, that is,

$$\Gamma_3(\varepsilon) = \left\{ \varepsilon \delta_0 + (1 - \varepsilon) \pi : \pi \in \mathcal{M}_1(\mathbb{R}) \right\}.$$

For each set in these classes, it is known that a unique least favorable prior for the corresponding univariate normal mean problem exists – see [Delatte \(2026a\)](#). This allows us to recover explicit expressions for the robust Bayes estimators of b over these sets. From Proposition 2.1, we can then directly obtain explicit expressions for the robust Bayes estimators of τ over $\Gamma_1(M)$, $\Gamma_2(m)$, and $\Gamma_3(\varepsilon)$ as shown below.

For $\Gamma_1(M)$ with $M > 0$, the least favorable prior is given by $\pi_{\Gamma_1(M)}^* = \mathcal{N}(0, M)$ – see [Feldman \(1991\)](#). The robust Bayes estimator of b is then a simple linear shrinkage function of D given by $h_{\Gamma_1(M)}^*(D) = MD/(M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. The robust Bayes estimator of τ is then

an inverse-variance weighted combination of $\hat{\tau}_{\text{RCT}}$ and a bias-corrected version of $\hat{\tau}_{\text{OS}}$ based on the linear shrinkage $h_{\Gamma_1(M)}^*(D)$. That is,

$$g_{\Gamma_1(M)}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w \left[\hat{\tau}_{\text{OS}} - \frac{M}{M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2} D \right].$$

This expression simplifies to an inverse-variance weighted combination of $\hat{\tau}_{\text{RCT}}$ and $\hat{\tau}_{\text{OS}}$ without bias correction but where the variance of $\hat{\tau}_{\text{OS}}$ is inflated by M . That is,

$$g_{\Gamma_1(M)}^*(\hat{\tau}_{\text{RCT}}, D) = \frac{(M + \sigma_{\text{OS}}^2)}{\sigma_{\text{RCT}}^2 + (M + \sigma_{\text{OS}}^2)} \hat{\tau}_{\text{RCT}} + \frac{\sigma_{\text{RCT}}^2}{\sigma_{\text{RCT}}^2 + (M + \sigma_{\text{OS}}^2)} \hat{\tau}_{\text{OS}}.$$

For $\Gamma_2(m)$ and $\Gamma_3(\varepsilon)$ with $m > m_1$ and $\varepsilon \in (0, 1)$, the least favorable priors exist but do not have a closed form – see [Johnstone \(2019\)](#) and [Delatte \(2026a\)](#) for a review of the existence results and see Remark 2.7 for the definition of the threshold m_1 . This prevents us from deriving a closed-form expression for the robust Bayes estimators $g_{\Gamma_2(m)}^*$ and $g_{\Gamma_3(\varepsilon)}^*$. Since, in both cases, the priors are symmetric and discretely supported (with bounded support for $\Gamma_2(m)$ and unbounded support for $\Gamma_3(\varepsilon)$), the optimal solutions admit a common expression that can be made explicit in terms of the support points and weights of the least favorable priors. Let Γ admit a symmetric and discretely supported least favorable prior π_Γ^* . Then we can always write $\pi_\Gamma^* = \sum_{j=-J}^J v_j \delta_{b_j}$ where $b_0 = 0$ and $b_{-j} = -b_j$, and $v_{-j} = v_j$ (and $J = +\infty$ if the support is unbounded). If we denote by $p_j(D)$ the posterior probability that the bias is equal to b_j given D , then the robust Bayes estimator of b over Γ is $h_\Gamma^*(D) = \sum_{j=-J}^J p_j(D) b_j$ and the robust Bayes estimator of τ over Γ takes the simple form

$$g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w \left[\hat{\tau}_{\text{OS}} - \sum_{j=-J}^J p_j(D) b_j \right].$$

This generic expression applies to both $\Gamma_2(m)$ and $\Gamma_3(\varepsilon)$ for the corresponding locations and weights of the least favorable priors (with $J(\varepsilon) = +\infty$ when $\Gamma = \Gamma_3(\varepsilon)$). To compute these estimators in practice, the least favorable priors $\pi_{\Gamma_2(m)}^*$ and $\pi_{\Gamma_3(\varepsilon)}^*$ need to be numerically approximated. Algorithms for this purpose are introduced in Appendix B. Because the algorithms return approximating distributions that are discrete and symmetric, the expression above can be directly used to approximate $g_{\Gamma_2(m)}^*$ and $g_{\Gamma_3(\varepsilon)}^*$ by simply plugging in the numerically approximated locations and weights.

Remark 2.6. From Corollary A.2, the robust Bayes estimator $h_{\Gamma_2(m)}^*$ is also the frequentist minimax estimator of b over the bounded parameter space $[-m, m]$ based on the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. For this reason, the robust Bayes estimator $g_{\Gamma_2(m)}^*$ is also the frequentist minimax estimator of τ over $\mathbb{R} \times [-m, m]$. This minimax estimator has been independently studied before – see, for instance, [Armstrong, Kline, and Sun \(2025\)](#) and [Montiel Olea and Zubova \(2026\)](#).

Remark 2.7. In general, the least favorable prior for the normal mean problem under $\Gamma_2(m)$ has no closed form. When $m < m_1 \approx 1.05(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)^{1/2}$, however, it does and it is equal then to the discrete mixture $\pi_{\Gamma_2(m)}^* = \frac{1}{2}\delta_{-m} + \frac{1}{2}\delta_m$ – see [Casella and Strawderman \(1981\)](#). This directly leads to the bounded nonlinear shrinkage function $h_{\Gamma_2(m)}^*(D) = m \tanh(mD/(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2))$ and the corresponding robust Bayes estimator $g_{\Gamma_2(m)}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w[\hat{\tau}_{\text{OS}} - m \tanh(mD/(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2))]$.




Figure 1: Mean Squared Error of Robust Bayes Estimators

Remark 2.8. The robust Bayes estimators $g_{\Gamma_1(M)}^*$, $g_{\Gamma_2(m)}^*$, and $g_{\Gamma_3(\varepsilon)}^*$ exhibit distinct behaviors that are entirely captured by the bias corrections $h_{\Gamma_1(M)}^*$, $h_{\Gamma_2(m)}^*$, and $h_{\Gamma_3(\varepsilon)}^*$. Two sets of properties are of particular interest. First, $h_{\Gamma_1(M)}^*$ is linear, $h_{\Gamma_2(m)}^*$ is nonlinear but always bounded (indeed, $|h_{\Gamma_2(m)}^*(D)| \leq m$), and $h_{\Gamma_3(\varepsilon)}^*$ is nonlinear and unbounded. Secondly, $h_{\Gamma_1(M)}^*$ is a proper shrinkage in the sense that $h_{\Gamma_1(M)}^*(D)/D = M/(M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2) \in (0, 1)$, but $h_{\Gamma_2(m)}^*$ and $h_{\Gamma_3(\varepsilon)}^*$ are not. The consequences of these properties on the performance of $g_{\Gamma_1(M)}^*$, $g_{\Gamma_2(m)}^*$, and $g_{\Gamma_3(\varepsilon)}^*$ are illustrated in Figure 1 for the arbitrarily chosen values $M = 1$, $m = 1$,

and $\varepsilon = 0.65$. It can be noted, moreover, that in all three cases, the parameters M , m , and ε entirely control the bias–variance trade-off. In each case, the robust Bayes estimators converge to $g_{\Gamma_0}^*$ and $g_{\Gamma_\infty}^*$ in the limits. Indeed, we have: $g_{\Gamma_1(M)}^* \rightarrow g_{\Gamma_0}^*$ when $M \rightarrow 0$ and $g_{\Gamma_1(M)}^* \rightarrow g_{\Gamma_\infty}^*$ when $M \rightarrow \infty$; $g_{\Gamma_2(m)}^* \rightarrow g_{\Gamma_0}^*$ when $m \rightarrow 0$ and $g_{\Gamma_2(m)}^* \rightarrow g_{\Gamma_\infty}^*$ when $m \rightarrow \infty$; and, $g_{\Gamma_3(\varepsilon)}^* \rightarrow g_{\Gamma_0}^*$ when $\varepsilon \rightarrow 1$ and $g_{\Gamma_3(\varepsilon)}^* \rightarrow g_{\Gamma_\infty}^*$ when $\varepsilon \rightarrow 0$.

3 A minimax regret approach to adaptation

The estimators derived in Section 2 are robust Bayes for the respective sets $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ of priors on the bias. When Γ is known and fixed, the robust Bayes estimator g_Γ^* comes with sharp worst-case guarantees for the integrated risk. However, researchers may be unwilling to commit to any particular set Γ of priors on the bias. Moreover, the results of Section 2.3 make clear that leaving this choice to a malicious researcher opens the way for statistical manipulation where any bias can be injected with maximal variance reduction. The problem is thus one of adaptation: can we find estimators with worst-case guarantees without knowing Γ ?

We develop in this section a solution to the adaptation problem based on a minimax regret approach. The adaptive estimators that we obtain come with worst-case regret guarantees over a parametrized class of sets relative to an oracle knowing the set of priors. This procedure extends the idea of Bickel (1984) and Armstrong, Kline, and Sun (2025) to the robust Bayesian realm – see Appendix A.2 for a review and a comparison. The success of the procedure depends on a reduction result – Theorem 3.1 – which is proved in Appendix A.2. In Section 3.2, we show that the conditions for the reduction result are verified by $\{\Gamma_1(M) : M \in [0, \infty)\}$ and $\{\Gamma_2(m) : m \in [0, \infty)\}$. The application to the latter exactly recovers the adaptive estimator introduced in Armstrong, Kline, and Sun (2025). The application to the former provides a new minimax regret estimator that benefits from a much simpler computational routine.

3.1 Bayesian adaptation regret and frequentist reduction

The first step in developing a minimax regret approach to the adaptation problem is to define a notion of regret with respect to an oracle knowing the set of priors. For this purpose, we restrict attention to sets of priors $\Gamma(\eta)$ that can be parametrized by a scalar state $\eta \in S$ with $S \subseteq \mathbb{R}$. We assume, moreover, that the resulting class of sets $\{\Gamma(\eta) : \eta \in S\}$ covers the entire bias–variance trade-off in the sense that $\bigcap_{\eta \in S} \Gamma(\eta) = \{\delta_0\}$ and $\text{cl}(\bigcup_{\eta \in S} \Gamma(\eta)) = \mathcal{M}_1(\mathbb{R})$.

The adaptation problem then transfers from the choice of $\Gamma(\eta)$ to the choice of the state $\eta \in S$. For any $\eta \in S$, we then define the regret of using the estimator $g_0 \in G$ when the set of priors is $\Gamma(\eta)$ as the ratio of the worst-case integrated risk of g_0 over $\Gamma(\eta)$ against the worst-case integrated risk of the robust Bayes estimator $g_{\Gamma(\eta)}^*$. We denote this quantity by $\mathcal{A}(g_0, \Gamma(\eta))$ and refer to it as the Bayesian adaptation regret. That is,

$$\mathcal{A}(g_0, \Gamma(\eta)) = \frac{\sup_{\pi \in \Gamma(\eta)} B(g_0, \pi)}{\sup_{\pi \in \Gamma(\eta)} B(g_{\Gamma(\eta)}^*, \pi)},$$

where $\sup_{\pi \in \Gamma(\eta)} B(g_{\Gamma(\eta)}^*, \pi) = \inf_{g \in G} \sup_{\pi \in \Gamma(\eta)} B(g, \pi)$ is the robust Bayes risk over $\Gamma(\eta)$ for estimators in G . By Lemma 2.1, the Bayesian adaptation regret can be simplified. Indeed, we have that $\sup_{\pi \in \Gamma(\eta)} B(g_0, \pi) = w\sigma_{\text{OS}}^2 + w^2 \sup_{\pi \in \Gamma(\eta)} B(h_0, \pi)$ with $h_0(D) = D - w^{-1}\mathcal{G}_0(D)$ and $\sup_{\pi \in \Gamma(\eta)} B(g_{\Gamma(\eta)}^*, \pi) = w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(\eta))$ where $B^*(\Gamma(\eta))$ denotes the robust Bayes risk for the normal mean problem based on the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. It follows then that

$$\mathcal{A}(g_0, \Gamma(\eta)) = \frac{\sup_{\pi \in \Gamma(\eta)} B(g_0, \pi)}{w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(\eta))} = \frac{\sigma_{\text{OS}}^2 + w \sup_{\pi \in \Gamma(\eta)} B(h_0, \pi)}{\sigma_{\text{OS}}^2 + w B^*(\Gamma(\eta))}.$$

The quantity $(\mathcal{A}(g_0, \Gamma(\eta)) - 1) \times 100$ then gives the percentage increase in worst-case integrated risk over $\Gamma(\eta)$ if one uses the estimator g_0 instead of the robust Bayes estimator $g_{\Gamma(\eta)}^*$. In terms of adaptation, this has a natural interpretation as a regret relative to an oracle who knows the state $\eta \in S$.

Based on the notion of Bayesian adaptation regret $\mathcal{A}(g_0, \Gamma(\eta))$, we can now consider a principled approach to the adaptation problem within the robust Bayesian framework. For this purpose, we define a new decision problem where:

1. Nature adversarially chooses the state $\eta \in S$ for the prior class;
2. the statistician selects the estimator in G that minimizes the worst-case Bayesian adaptation regret.

This leads to the new objective function given by

$$\inf_{g \in G} \sup_{\eta \in S} \mathcal{A}(g, \Gamma(\eta)).$$

We call an estimator that reaches the minimum a minimax regret estimator over $\{\Gamma(\eta) : \eta \in S\}$.

By Lemma 2.1 again, we have the correspondence

$$\inf_{g \in G} \sup_{\eta \in S} \mathcal{A}(g, \Gamma(\eta)) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{\eta \in S} \frac{\sigma_{\text{OS}}^2 + w \sup_{\pi \in \Gamma(\eta)} B(h, \pi)}{\sigma_{\text{OS}}^2 + w B^*(\Gamma(\eta))}.$$

This problem is not a robust Bayesian decision problem of the form considered in Section 2. In this case, the problem is considerably harder as we have to solve two nested maximization problems: one over the state $\eta \in S$ and one over the priors $\pi \in \Gamma(\eta)$. In particular, we cannot find a solution in the same way as we did in Section 2.2. The fact that the problem may admit a solution is actually far from obvious. The next two results show that if $\{\Gamma(\eta) : \eta \in S\}$ is regular enough, then solutions exist and can be computed. Theorem 3.1 is the main proof mechanism: it provides a fundamental reduction of the original adaptation problem to a frequentist minimax game that admits a computable solution.

Theorem 3.1. *Let $\{\Gamma(\eta) : \eta \in [0, \infty)\}$ be an increasing family of sets such that $B^*(\Gamma(\eta)) > 0$ for each $\eta > 0$. For any $b \in \mathbb{R}$, we define the map s by $s(b) \leq \eta$ if and only if $\delta_b \in \Gamma(\eta)$. If for every $\eta > 0$ and every $\pi \in \Gamma(\eta)$, it holds that $\int_{\mathbb{R}} B^*(\Gamma(s(b))) d\pi(b) \leq B^*(\Gamma(\eta))$, then*

$$\inf_{g \in G} \sup_{\eta \geq 0} \mathcal{A}(g, \Gamma(\eta)) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{b \in \mathbb{R}} \frac{\sigma_{\text{OS}}^2 + w R(h, b)}{\sigma_{\text{OS}}^2 + w B^*(\Gamma(s(b)))},$$

where $B^*(\Gamma(s(b)))$ is the robust Bayes risk over $\Gamma(s(b))$ for the estimation of b_0 based on an observation distributed according to $\mathcal{N}(b_0, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$.

The result is proved in Appendix A.2. The right-hand side of the equation in Theorem 3.1 directly rewrites as the function $W_1(b) + W_2(b)R(h, b)$ with $W_1(b) = \sigma_{\text{OS}}^2 / (\sigma_{\text{OS}}^2 + w B^*(\Gamma(s(b))))$ and $W_2(b) = w / (\sigma_{\text{OS}}^2 + w B^*(\Gamma(s(b))))$. For any fixed $b \in \mathbb{R}$, this function is an affine transformation of $R(h, b)$. For classes that satisfy the regularity conditions of Theorem 3.1, the reduction result is enough to ensure that a minimax regret estimator exists. This follows again from standard minimax arguments in the univariate normal mean problem. This leads to the existence result below whose proof is collected in Appendix A.2. The computational aspect of the problem is tackled in the next section.

Corollary 3.2. *Suppose that $\{\Gamma(\eta) : \eta \in [0, \infty)\}$ satisfies the conditions of Theorem 3.1. Then there exists a minimax regret estimator over $\{\Gamma(\eta) : \eta \in [0, \infty)\}$.*

Remark 3.1. As in Section 2, the restriction to equivariant estimators $g \in G$ is without loss of generality for a large collection of sets $\Lambda(\eta) \subseteq \mathcal{M}_1(\mathbb{R}^2)$ of joint priors. An argument based on the Hunt–Stein theorem similar to the one invoked in Remark 2.4 allows us to recover the equality of $\inf_{g \in G} \sup_{\eta \in \mathcal{S}} \mathcal{A}(g, \Lambda(\eta))$ and $\inf_{g \in \mathcal{D}(\mathbb{R}^2)} \sup_{\eta \in \mathcal{S}} \mathcal{A}(g, \Lambda(\eta))$ as soon as $\Lambda(\eta)$ does not impose restrictions on the location of the causal effect, where $\mathcal{A}(g, \Lambda(\eta))$ is naturally defined over $\Lambda(\eta)$ by extension.

Remark 3.2. These results can again be viewed as a generalization to the robust Bayesian world of the results in Armstrong, Kline, and Sun (2025) – in particular, Lemma 1 and Theorem 1 (pp. 1989–1990). These extensions increase the set of decision-theoretic conditions on the bias over which adaptation is possible, leading to simpler adaptive estimators that still benefit from worst-case regret guarantees. The notion of regret defined by $\mathcal{A}(g_0, \Gamma(\eta))$ extends the ideas of Bickel (1984) to the robust Bayesian world. This choice trades off scale-invariance for sensitivity to small values of $\sigma_{\text{OS}}^2 + wB^*(\Gamma(\eta))$.

3.2 Minimax regret estimators: existence and computation

We start by showing that the reduction given by Theorem 3.1 applies to the two classes of priors given by $\{\Gamma_1(M) : M \in [0, \infty)\}$ and $\{\Gamma_2(m) : m \in [0, \infty)\}$. In both cases, this reduction guarantees the existence of a minimax regret estimator over the corresponding class and gives a way for its computation. For the class defined by $\Gamma_2(m)$, we exactly recover the estimator introduced in Armstrong, Kline, and Sun (2025). For the class defined by $\Gamma_1(M)$, we obtain a new minimax regret estimator with a clear computational advantage. The simplification is a direct consequence of conjugacy for the least favorable priors and is discussed at more length in Remark 3.3.

Corollary 3.3. *For $\{\Gamma_2(m) : m \in [0, \infty)\}$, it holds that*

$$\inf_{g \in G} \sup_{m \geq 0} \mathcal{A}(g, \Gamma_2(m)) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{b \in \mathbb{R}} \frac{\sigma_{\text{OS}}^2 + wR(h, b)}{\sigma_{\text{OS}}^2 + wR^*([-|b|, |b|])},$$

where $R^*([-|b|, |b|])$ is the frequentist minimax risk over $[-|b|, |b|]$ for the estimation of b_0 based on an observation distributed according to $\mathcal{N}(b_0, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. Moreover, there exists a minimax regret estimator over $\{\Gamma_2(m) : m \in [0, \infty)\}$.

This result can be derived by verifying the conditions of Theorem 3.1 as done in Ap-

pendix A.2. It can also be derived from the results in [Armstrong, Kline, and Sun \(2025\)](#) combined with Corollary A.2. The following result, however, has no counterpart in the frequentist world and is proved in Appendix A.2.

Corollary 3.4. *For $\{\Gamma_1(M) : M \in [0, \infty)\}$, it holds that*

$$\inf_{g \in G} \sup_{M \geq 0} \mathcal{A}(g, \Gamma_1(M)) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{b \in \mathbb{R}} \frac{\sigma_{\text{OS}}^2 + wR(h, b)}{\sigma_{\text{OS}}^2 + wB^*(\Gamma_1(b^2))},$$

where $B^*(\Gamma_1(b^2)) = b^2(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)/(b^2 + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$ is the robust Bayes risk over $\Gamma_1(b^2)$ for the estimation of b_0 based on an observation distributed according to $\mathcal{N}(b_0, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. Moreover, there exists a minimax regret estimator over $\{\Gamma_1(M) : M \in [0, \infty)\}$.

In both cases, the minimax regret estimators can be recovered by first solving the normal mean problems based on the respective affine transformations of $R(h, b)$ on the right-hand-side of each equation in Corollary 3.3 and Corollary 3.4. Accounting for the simple transformations of the risk, these problems are standard frequentist minimax games for the estimation of b over the unrestricted space $\mathbb{R}$ based on an observation distributed according to $\mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. They can be solved by approximating the least favorable prior and by taking the Bayes estimator with respect to it. Implementation details for both classes are introduced in Remark 3.4. Given this computational recipe, the advantage of the minimax regret estimator for $\{\Gamma_1(M) : M \in [0, \infty)\}$ is clear: the affine transformation has a simple closed form. This is not the case for the minimax regret estimator for $\{\Gamma_2(m) : m \in [0, \infty)\}$: the quantity $R^*([-|b|, |b|])$ needs to be approximated for any value of $b \in \mathbb{R}$ beyond a small threshold. Even with the improvements provided in Appendix B, the task remains difficult and computationally intensive.

Remark 3.3. The difference between the minimax regret estimators based on $\Gamma_2(m)$ and $\Gamma_1(M)$ depends only on the difference between the weights $B^*(\Gamma_2(|b|)) = R^*([-|b|, |b|])$ and $B^*(\Gamma_1(b^2))$. Using Algorithm [BNM-LFP](#) in Appendix B, we can approximate $R^*([-|b|, |b|])$ for $\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2 = 1$. The difference between the weights is minimal for small and large values of b and always bounded by 0.17 for all values of $b \in \mathbb{R}$. This difference is plotted in Figure 2 below.

Remark 3.4. In the empirical applications of Section 5 and Section 6, we compute the minimax regret estimator for $\{\Gamma_1(M) : M \in [0, \infty)\}$ by first solving the frequentist minimax




Figure 2: Difference between $R^*([-|b|, |b|])$ and $B^*(\Gamma_1(b^2))$ for $\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2 = 1$

problem of Corollary 3.4. For this purpose, we approximate the least favorable prior in this problem using Algorithm GNM-LFP and then recover the minimax regret estimator from the Bayes estimator with respect to the approximation. The same approach can be used to compute the minimax regret estimator for $\{\Gamma_2(m) : m \in [0, \infty)\}$ based on first approximating $R^*([-|b|, |b|])$ for a large grid of values for b . In the applications, however, we do not use this approach but directly approximate the minimax estimator for $\{\Gamma_2(m) : m \in [0, \infty)\}$ using the adaptive soft-thresholding estimator considered in Armstrong, Kline, and Sun (2025) and based on the bias correction $h(D) = \text{sgn}(D) \max(|D| - \lambda(w)(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)^{1/2}, 0)$ with $\lambda(w) = 0.45 - 0.24 \ln(1 - w)$. The reason for this is two-fold: the soft-thresholding adaptive estimator is a close approximation of the minimax regret estimator; moreover, the soft-thresholding estimator is of independent interest (a version of this estimator is considered, for instance, by Chen, Li, Zhang, and Ye (2025) based on heuristic arguments).

Remark 3.5. The conditions of Theorem 3.1 can be verified for other classes than those defined by $\Gamma_1(M)$ and $\Gamma_2(m)$. Delineating the exact range of sets of priors covered by these conditions is left for future research. Moreover, Theorem 3.1 only gives sufficient conditions for the frequentist reduction. More general conditions may exist; characterizing these conditions is another fruitful avenue for research.

4 An empirical Bayes approach to adaptation

The approach to adaptation developed in Section 3 is not the only one that can be considered within the robust Bayesian framework. We develop in this section another strategy, leading to adaptive solutions that have a natural interpretation as empirical Bayes estimators with respect to families of least favorable priors. Given a class $\{\Gamma(\eta) : \eta \in S\}$ of sets of priors, the idea is to take the robust Bayesian decision problem of Section 2 as being solved by an oracle for some set $\Gamma(\eta_0)$ known only to them and Nature. We can then consider approximations of the oracle estimator $g_{\Gamma(\eta_0)}^*$ obtained by directly plugging in estimators of the unknown state η_0 where the bias is viewed as being generated from the optimal strategy of Nature, that is, from the least favorable prior $\pi_{\Gamma(\eta_0)}^*$ when it exists.

We start by making explicit the structure of this empirical Bayes approach to adaptation. We then isolate the class $\{\Gamma_1(M) : M \in [0, \infty)\}$ by virtue of the parametric family of least favorable priors it generates and show that it leads to simple estimators based on adaptive shrinkage. In Appendix A.3, we prove that a multivariate extension of this solution exactly recovers the adaptive estimator of Green and Strawderman (1991). We finally provide a general algorithm to implement the strategy for any class where the least favorable priors exist but may not have a parametric expression or even a closed form. We apply it to obtain adaptive estimators from the classes $\{\Gamma_2(m) : m \in [0, \infty)\}$ and $\{\Gamma_3(\varepsilon) : \varepsilon \in (0, 1]\}$.

Remark 4.1. This approach to adaptation is related to a long tradition of linear combination estimators where the weights are obtained by plugging in some well-chosen estimators – see, for instance, Green and Strawderman (1991), de Chaisemartin and D’Haultfoeulle (2020), or Chen, Li, Zhang, and Ye (2025). In all these approaches, both the form of the estimators as linear combinations and the choice of the weights are based on heuristic arguments. Our approach differs in an important way: both the form of the estimators and the choice of the plug-in weights derive from the robust Bayesian decision problem. The former was proved in Section 2; the latter is shown in this section.

4.1 Robust empirical Bayes estimators

We consider again a class $\{\Gamma(\eta) : \eta \in S\}$ of sets of priors that is parametrized by $S \subseteq \mathbb{R}$ and that covers the entire bias–variance trade-off in the sense that $\bigcap_{\eta \in S} \Gamma(\eta) = \{\delta_0\}$ and $\text{cl}(\bigcup_{\eta \in S} \Gamma(\eta)) = \mathcal{M}_1(\mathbb{R})$. We assume, moreover, that for each $\eta \in S$, there exists a unique

least favorable prior $\pi_{\Gamma(\eta)}^*$ for the univariate normal mean problem corresponding to $\Gamma(\eta)$. To develop our adaptation strategy, we now assume that an oracle solves the robust Bayesian decision problem of Section 2 for some set $\Gamma(\eta_0)$ where the state $\eta_0 \in S$ is unknown to everybody except Nature and the oracle. Our goal is then to approximate the oracle estimator $g_{\Gamma(\eta_0)}^*$ using an estimator of the unknown state η_0 from the available data. The key point is that the bias can now be viewed as being generated from the optimal strategy of Nature, that is, the least favorable prior $\pi_{\Gamma(\eta_0)}^*$. In this case, the marginal density of the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$ is given by

$$\psi(D) = \int_{\mathbb{R}} \phi_0(D - b) d\pi_{\Gamma(\eta_0)}^*(b),$$

where ϕ_0 is the density of $\mathcal{N}(0, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. The approximation problem is then equivalent to that of recovering $\pi_{\Gamma(\eta_0)}^*$ in the family of least favorable priors $\{\pi_{\Gamma(\eta)}^* : \eta \in S\}$. For any estimator $\hat{\eta}$ of η_0 , a simple solution is to take the plug-in least favorable prior $\pi_{\Gamma(\hat{\eta})}^*$ as the truth and compute the bias correction $h_{\Gamma(\hat{\eta})}^*$ as the Bayes estimator with respect to $\pi_{\Gamma(\hat{\eta})}^*$. That is,

$$h_{\Gamma(\hat{\eta})}^*(D) = \int_{\mathbb{R}} b d\pi_{\Gamma(\hat{\eta})}^*(b \mid D).$$

The approximation to the oracle estimator $g_{\Gamma(\eta_0)}^*$ is then obtained by plugging in the empirical Bayes estimator $h_{\Gamma(\hat{\eta})}^*$ into the expression obtained in Proposition 2.1. That is,

$$g_{\Gamma(\hat{\eta})}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w[\hat{\tau}_{\text{OS}} - h_{\Gamma(\hat{\eta})}^*(D)].$$

We call $g_{\Gamma(\hat{\eta})}^*$ a robust empirical Bayes estimator of τ over $\{\Gamma(\eta) : \eta \in S\}$ based on the estimator $\hat{\eta}$. The terminology is warranted: by construction, the estimator $g_{\Gamma(\hat{\eta})}^*$ is empirical Bayes relative to the family of least favorable priors obtained from the robust Bayesian decision problems over $\{\Gamma(\eta) : \eta \in S\}$.

Remark 4.2. If the family $\{\pi_{\Gamma(\eta)}^* : \eta \in S\}$ of least favorable priors is parametric, then the choice of $\hat{\eta}$ can be motivated from the normality of the observation and the parametric constraints (through maximum likelihood or unbiasedness, for instance). This is what we do for the class $\{\Gamma_1(M) : M \in [0, \infty)\}$ as shown in Section 4.2. This is also the path we follow for the discrete approximations considered in Section 4.3. It should be noted, however, that, even in these cases, the choice of $\hat{\eta}$ is not entirely constrained by the robust Bayesian

framework and several legitimate candidates usually exist. In [Delatte \(2026b\)](#), we use regret in the robust Bayesian problem to help discipline the choice.

Remark 4.3. The notation $g_{\Gamma(\hat{\eta})}^*$ is a plug-in expression, not a robust Bayes estimator for a data-dependent set of priors $\Gamma(\hat{\eta})$ – this notion and the resulting decision problems would be ill-posed.

4.2 Parametric case: moment constraint and adaptive shrinkage

We apply the empirical Bayes approach introduced in [Section 4.1](#) to the class $\{\Gamma_1(M) : M \in [0, \infty)\}$ of sets of priors with second-order moment bounded by $M \geq 0$. This class allows us to illustrate how the adaptive approach can be implemented with a parametric family of least favorable priors. This example leads, moreover, to simple robust empirical Bayes estimators based on adaptive shrinkage that are of independent interest. In [Appendix A.3](#), we show that the robust empirical Bayes estimator for a multivariate extension of the class exactly recovers the optimal estimator of [Green and Strawderman \(1991\)](#).

For any $M \in (0, \infty)$, we already know that the least favorable prior for the normal mean problem based on $\Gamma_1(M)$ is $\pi_{\Gamma_1(M)}^* = \mathcal{N}(0, M)$ – see again [Feldman \(1991\)](#). The class $\{\Gamma_1(M) : M \in [0, \infty)\}$ thus leads to the parametric family of least favorable priors $\{\mathcal{N}(0, M) : M \in [0, \infty)\}$ (where we slightly abuse notation by writing $\mathcal{N}(0, 0) = \delta_0$). We now follow the approach of [Section 4.1](#): we assume that an oracle solves the robust Bayesian decision problem based on $\Gamma_1(M_0)$ for some unknown state $M_0 \geq 0$ and view the bias as being generated from the least favorable prior $\pi_{\Gamma_1(M_0)}^*$. Under this prior, the marginal distribution of D is now given by $\mathcal{N}(0, M_0 + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. In this parametric case, a natural estimator $\hat{M}$ of M_0 is the maximum likelihood estimator based on the marginal distribution of D , that is,

$$\hat{M} = \max(0, D^2 - \sigma_{\text{RCT}}^2 - \sigma_{\text{OS}}^2) = (D^2 - \sigma_{\text{RCT}}^2 - \sigma_{\text{OS}}^2)_+.$$

We can then use $\hat{M}$ to select the class $\Gamma_1(\hat{M})$ and, equivalently, the least favorable prior $\pi_{\Gamma_1(\hat{M})}^*$. The Bayes estimator in this case takes the simple form of an adaptive shrinkage estimator, namely

$$h_{\Gamma_1(\hat{M})}^*(D) = \left(1 - \frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{D^2}\right)_+ D.$$

We can then plug this estimator into [Proposition 2.1](#) to obtain a closed-form expression for

the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ over $\{\Gamma_1(M) : M \in [0, \infty)\}$ based on $\hat{M}$. We collect this expression in the proposition below.

Proposition 4.1. *The robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ over $\{\Gamma_1(M) : M \in [0, \infty)\}$ based on the maximum marginal likelihood estimator $\hat{M}$ is given by*

$$g_{\Gamma_1(\hat{M})}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w \left[\hat{\tau}_{\text{OS}} - \left(1 - \frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{D^2} \right)_+ D \right].$$

By rearranging the terms, we can show that the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ admits another simple expression that more directly highlights its structure as an adaptive thresholding estimator. The following result, which follows from Corollary A.3 proved in Appendix A.3, can be appropriately compared with the non-adaptive result of Proposition 2.4.

Corollary 4.2. *1. If $D^2 \leq \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, then*

$$g_{\Gamma_1(\hat{M})}^*(\hat{\tau}_{\text{RCT}}, D) = g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D).$$

2. If $D^2 > \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, then

$$g_{\Gamma_1(\hat{M})}^*(\hat{\tau}_{\text{RCT}}, D) = \left[\frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{D^2} \right] g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D) + \left[1 - \frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{D^2} \right] g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D).$$

This result makes explicit the thresholding behavior of $g_{\Gamma_1(\hat{M})}^*$. When $D^2 \leq \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, the estimator attributes D to sampling noise and recovers the optimal inverse-variance weighted estimator $g_{\Gamma_0}^*$ under $b = 0$. When $D^2 > \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, the estimator resorts to the convex combination of $g_{\Gamma_0}^*$ and $g_{\Gamma_\infty}^*$ with convex weight $W(D) = (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)/D^2 \in (0, 1)$. For large enough values of D , the estimator recovers the estimator with minimal worst-case bias, that is, $g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}}$.

Remark 4.4. The choice of the maximum marginal likelihood estimator $\hat{M}$ to derive the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ is natural given the parametric constraints, but it is not entirely forced by the robust Bayesian framework. The choice of an optimal estimator of M_0 in terms of an appropriate criterion is an open question. We show in a separate manuscript – see Delatte (2026b) – that the estimator $g_{\Gamma_1(\hat{M})}^*$ has bounded absolute and relative regret with respect to the oracle estimator $g_{\Gamma_1(M_0)}^*$ (in terms of integrated risk for the least favorable prior as well as for an adversarial prior for the estimator). In the same manuscript, we show

that the choice of an estimator of M_0 can be partly disciplined by these regret bounds. This difficulty (which is inherent to the empirical Bayes approach) can be appropriately contrasted with the clear optimality guarantees derived from the minimax regret approach of Section 3.

Remark 4.5. In Appendix A.3, we show that the multivariate extension of the approach opens the way for another natural estimator $\widetilde{M}$ of M_0 based on unbiasedness for the whole shrinkage factor. This approach, which is unfeasible in the univariate case, yields in dimensions strictly greater than 2 an equivalence of independent interest: the robust empirical Bayes estimator over $\{\Gamma_1(M) : M \in [0, \infty)\}$ based on $\widetilde{M}$ is the optimal estimator of Green and Strawderman (1991). This is proved in Proposition A.4 and Corollary A.5 in Appendix A.3.

4.3 Nonparametric case: a general approximation algorithm

The empirical Bayes approach to adaptation developed in Section 4.1 does not require the family of least favorable priors to be parametric. We show in this section that the approach can be implemented generally with nonparametric families and, in particular, with families of least favorable priors that do not admit a closed form. Our solution leverages the fact that least favorable priors can be well approximated by finitely supported distributions under general conditions – these conditions and algorithms for this purpose are introduced in Appendix B. The outcome is a general algorithm – Algorithm ReBA – that returns approximations of the robust empirical Bayes estimators based on these discrete approximations.

We consider again a class $\{\Gamma(\eta) : \eta \in S\}$ of sets of priors where $S \subseteq \mathbb{R}$. We assume that for each $\eta \in S$, there exists a unique least favorable prior in the normal mean problem based on $\Gamma(\eta)$. We view again the bias as being generated from the robust Bayesian decision problem solved by the oracle for $\Gamma = \Gamma(\eta_0)$ for some unknown state $\eta_0 \in S$. For each $\eta \in S$, we then assume that we have access to a discretely supported approximation of the least favorable prior. We can generally write this approximation as $\pi_{\Gamma(\eta)}^* = \sum_{k=1}^{K(\eta)} v_k(\eta) \delta_{b_k(\eta)}$ where $v_k(\eta) > 0$, $\sum_{k=1}^{K(\eta)} v_k(\eta) = 1$, and $b_k(\eta) \in \mathbb{R}$. From the discreteness of the approximation, the marginal density of D for $\eta \in S$ is then a convex combination of Gaussian densities, that is,

$$\psi_\eta(D) = \sum_{k=1}^{K(\eta)} v_k(\eta) \phi_0\left(D - b_k(\eta)\right),$$

where ϕ_0 is again the density of $\mathcal{N}(0, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. By virtue of this parametric constraint, a

natural approach for the estimation of η_0 is to approximate the maximum marginal likelihood estimator. For this purpose, we discretize S into a grid $S_0 = \{\eta_1, \dots, \eta_L\}$, evaluate $\psi_\eta(D)$ for each $\eta \in S_0$, and return as an estimator of η_0 the value in S_0 that maximizes $\psi_\eta(D)$ (with any tie-breaking choice in case of equality). That is,

$$\hat{\eta} \in \arg \max_{\eta \in S_0} \psi_\eta(D).$$

We can then leverage the results of Section 2.4 (without symmetry) to approximate a robust empirical Bayes estimator over $\{\Gamma(\eta) : \eta \in S\}$. This estimator is the direct outcome of applying the empirical Bayes approach of Section 4.1 based on the estimator $\hat{\eta}$ for the finite family $\{\pi_{\Gamma(\eta)}^* : \eta \in S_0\}$ of approximations of the least favorable priors. We denote this approximation by $g_{\Gamma(\hat{\eta})}^*$ and have that

$$g_{\Gamma(\hat{\eta})}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w \left[\hat{\tau}_{\text{OS}} - \sum_{k=1}^{K(\hat{\eta})} p_k(D; \hat{\eta}) b_k(\hat{\eta}) \right],$$

where $p_k(D; \hat{\eta}) = (\psi_{\hat{\eta}}(D))^{-1} v_k(\hat{\eta}) \phi_0(D - b_k(\hat{\eta}))$ is the posterior probability that the bias is equal to $b_k(\hat{\eta})$ given D . We collect this procedure in Algorithm ReBA below. Given a class $\{\Gamma(\eta) : \eta \in S\}$, the algorithm takes a finite grid S_0 for S and the approximations of the least favorable priors over S_0 as inputs and returns the approximation $g_{\Gamma(\hat{\eta})}^*(\hat{\tau}_{\text{RCT}}, D)$.

ReBA Algorithm Robust empirical Bayes Approximation Algorithm

Input Grid: $S_0 = \{\eta_1, \dots, \eta_L\}$

Approximate LFPs: $\pi_{\Gamma(\eta)}^* = \sum_{k=1}^{K(\eta)} v_k(\eta) \delta_{b_k(\eta)}$ for every $\eta \in S_0$

Output ReBA estimator: $g_{\Gamma(\hat{\eta})}^*(\hat{\tau}_{\text{RCT}}, D)$

- 1: Evaluate the marginal densities $\psi_\eta(D) = \sum_{k=1}^{K(\eta)} v_k(\eta) \phi_0(D - b_k(\eta))$ for each $\eta \in S_0$;
 - 2: Find $\hat{\eta} \in \arg \max_{\eta \in S_0} \psi_\eta(D)$ that maximizes the marginal densities over S_0 ;
 - 3: Evaluate $p_k(D; \hat{\eta}) = (\psi_{\hat{\eta}}(D))^{-1} v_k(\hat{\eta}) \phi_0(D - b_k(\hat{\eta}))$ for each $k \in \{1, \dots, K(\hat{\eta})\}$;
 - 4: Return the estimator $g_{\Gamma(\hat{\eta})}^*(\hat{\tau}_{\text{RCT}}, D) = (1 - w)\hat{\tau}_{\text{RCT}} + w(\hat{\tau}_{\text{OS}} - \sum_{k=1}^{K(\hat{\eta})} p_k(D; \hat{\eta}) b_k(\hat{\eta}))$.
-

Remark 4.6. Conditions under which least favorable priors can be approximated are reviewed in Appendix B. It holds, in particular, that weakly converging approximations are available

as soon as the sets $\Gamma(\eta)$ are defined through generalized moment constraints. We show in Appendix B that this holds, for instance, for the classes defined by $\Gamma_1(M)$, $\Gamma_2(m)$, and $\Gamma_3(\varepsilon)$.

As a sanity check, we first use Algorithm ReBA over $\{\Gamma_1(M) : M \in [0, \infty)\}$ and compare the approximation to the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ derived in Section 4.2. The results of the comparison are collected in Example B.1 in Appendix B and show that the approximation almost perfectly recovers the closed-form solution. We then apply Algorithm ReBA over the classes defined by $\Gamma_2(m)$ and $\Gamma_3(\varepsilon)$, which lead to families of least favorable priors without closed forms. The approximations $g_{\Gamma_2(\hat{m})}^*$ and $g_{\Gamma_3(\hat{\varepsilon})}^*$ we obtain are used in the Monte Carlo simulations of Section 5 and the empirical applications of Section 6. In Example 4.1 and Example 4.2 below, we make clear the implementation choices we use in Algorithm ReBA to obtain them.

Example 4.1. The least favorable prior in the normal mean problem for $\Gamma_2(m)$ is symmetric and discretely supported on $[-m, m]$ with a number of support points that grows linearly with m . We consider in Appendix B a modified version of the algorithm introduced in Kempthorne (1987) – Algorithm BNM-LFP – that is seen to outperform existing solutions. This iterative procedure allows us to accurately recover the support points and weights of the least favorable priors, hence yielding stable approximations with few support points. For the empirical applications, we apply Algorithm ReBA with $S_0 = \{m(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)^{1/2} : m = 0.1, 0.2, \dots, 19.9, 20\}$ and least favorable prior approximations obtained from Algorithm BNM-LFP.

Example 4.2. The least favorable prior in the normal mean problem for $\Gamma_3(\varepsilon)$ is discretely supported but on a countably infinite number of points. This prevents an exact recovery (up to numerical errors) of the least favorable prior (unlike $\Gamma_2(m)$). We introduce in Appendix B a solution that leads to stable approximations based on truncating the support and approximating the solutions for the problem based on $\Gamma_3(\varepsilon) \cap \Gamma_2(m_{\text{max}})$ for some $m_{\text{max}} > 0$. This leads to the new Algorithm HNM-LFP. For the empirical applications, we apply Algorithm ReBA with $S_0 = \{0.01, 0.02, \dots, 0.99\}$ and least favorable prior approximations obtained from Algorithm HNM-LFP with $m_{\text{max}} = 20(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)^{1/2}$.

5 Monte Carlo simulations

In this section, we perform Monte Carlo simulations to evaluate the performance of the adaptive estimators derived in Section 3 and Section 4. The results are collected in Table 1 and Table 9 and discussed in Section 5.3.

5.1 Simulation designs and objectives

We consider two designs where each design captures a type of bias commonly encountered in causal inference with observational data. The first design models an observational study with two-sided non-compliance and heterogeneous treatment effects across the compliance groups – Section 5.2. The second design models an observational study with selection on unobservables – Appendix C.

In each design, we consider an RCT with a perfectly randomized treatment and perfect compliance among the experimental units. We take the average treatment effect (ATE) as the causal effect of interest $\tau \in \mathbb{R}$, and normalize the designs so that the ATE is always null, that is, $\tau = 0$. We estimate the ATE with a simple difference-in-means estimator and denote it by $\hat{\tau}_{\text{RCT}}$. The bias is then controlled by a single scalar $b \in \mathbb{R}$ that stems in both designs from the impossibility of recovering the ATE in the observational study. In each design, for the observational study, we consider an estimator $\hat{\tau}_{\text{OS}}$ based on a quasi-experimental strategy targeting a well-identified quantity that can be viewed as an approximation to the ATE. In each case, the variances σ_{RCT}^2 and σ_{OS}^2 are set equal to standard variance estimators of $\hat{\tau}_{\text{RCT}}$ and $\hat{\tau}_{\text{OS}}$. All the estimators are chosen so as to mimic standard empirical practices among economists. The distributional assumptions and the estimators used in each design are introduced in Section 5.2 and Appendix C, respectively.

Our objective is then to combine $\hat{\tau}_{\text{RCT}}$ and $\hat{\tau}_{\text{OS}}$ to estimate $\tau = 0$. For this purpose, we consider eight estimators. From the minimax regret approach, we include the minimax regret estimator over $\{\Gamma_1(M) : M \in [0, \infty)\}$ and the soft-thresholding approximation of the minimax regret estimator over $\{\Gamma_2(m) : m \in [0, \infty)\}$ (“*Moment Regret*” and “*Support Regret*”, respectively). From the empirical Bayes approach, we include the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ (“*Moment Plug-in*”), and the approximations $g_{\Gamma_2(\hat{m})}^*$ and $g_{\Gamma_3(\hat{\varepsilon})}^*$ from Algorithm ReBA (“*Support Plug-in*” and “*Huber Plug-in*”, respectively). See Remark 3.4, Example 4.1, and Example 4.2 for computational details. For the sake of comparison, we

include the bias–variance trade-off endpoints $g_{\Gamma_\infty}^*$ (“*RCT*”) and $g_{\Gamma_0}^*$ (“*MLE*”) as well as a simple pretesting estimator (“*Pretesting*”) based on a Z -test ($c = 1.96$).

The sample size of the RCT is set to $N = 200$ and the sample size of the observational study to $n = 20,000$. We run 1,000 simulations and estimate the bias, variance, and mean squared error of each estimator for $b \in \{0, 0.1, \dots, 0.9, 1\}$. The results of the simulations are discussed in Section 5.3. Additional visualizations of the results are provided in Appendix C.3.

5.2 Design 1: Two-sided noncompliance and heterogeneous effects

The observational study in this design exhibits a standard form of two-sided noncompliance with heterogeneous treatment effects. This type of issue is prevalent in observational causal inference – see, for instance, [Imbens and Angrist \(1994\)](#). In this case, the average treatment effect (ATE) for the full population is not identified – only the average treatment effect for the compliers or local average treatment effect (LATE) is. The bias in this case is commensurate with the heterogeneity in the treatment effects and is captured by the difference between the LATE and the ATE. We collect in this section the distributional assumptions and estimators used in this design.

The sample sizes of the RCT and the observational study are denoted by N and n , respectively. For each $i \in \{1, \dots, N, N+1, \dots, N+n\}$, the outcome Y_i is determined by the vector $(Z_i, S_i, U_i, \varepsilon_i)$ of random variables, where Z_i is observable and

$$Z_i \sim \text{Bernoulli}(0.5),$$

while $S_i \in \{C, AT, NT\}$, U_i , and ε_i are unobservable and

$$S_i \sim \text{Multinoulli}(0.5, 0.25, 0.25),$$

$$U_i \mid S_i \sim \mathcal{N}(\mu(S_i), 1),$$

$$\varepsilon_i \sim \mathcal{N}(0, 1),$$

with $\mu(S_i) = 0 \cdot \mathbb{1}_{(S_i=C)} + 2\mathbb{1}_{(S_i=AT)} - 2\mathbb{1}_{(S_i=NT)}$. For each $i \in \{1, \dots, N, N+1, \dots, N+n\}$,

the observable outcome Y_i is then given by

$$Y_i = \begin{cases} bT_i + U_i + \varepsilon_i & \text{if } S_i = C, \\ -bT_i + U_i + \varepsilon_i & \text{if } S_i = AT, NT. \end{cases}$$

This leads to the potential outcomes $Y_i(0) = U_i + \varepsilon_i$ and $Y_i(1) = b + U_i + \varepsilon_i$ if $S_i = C$ and $Y_i(1) = -b + U_i + \varepsilon_i$ otherwise. In the RCT, we assume that the treatment is randomized independently of the subgroups, that is, for $i \in \{1, \dots, N\}$,

$$T_i = Z_i.$$

In this case, the ATE is identified and equal to zero. Indeed, for any $i \in \{1, \dots, N\}$, we have

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0)] = \mathbb{E}[Y_i|T_i = 1] - \mathbb{E}[Y_i|T_i = 0] = 0.$$

In the observational study, we assume that the treatment is randomized for the subgroup of compliers $S_i = C$ while the two other subgroups either always receive the treatment or never receive it. That is, for $i \in \{m+1, \dots, m+n\}$,

$$T_i = Z_i \mathbb{1}_{(S_i=C)} + 1 \cdot \mathbb{1}_{(S_i=AT)} + 0 \cdot \mathbb{1}_{(S_i=NT)}.$$

In this case, the ATE is not identified. However, the LATE for the subgroup of compliers $S_i = C$ is identified and equal to b . Indeed, for any $i \in \{N+1, \dots, N+n\}$, we have

$$\mathbb{E}[Y_i(1) - Y_i(0)|S_i = C] = \frac{\mathbb{E}[Y_i|Z_i = 1] - \mathbb{E}[Y_i|Z_i = 0]}{\mathbb{E}[T_i|Z_i = 1] - \mathbb{E}[T_i|Z_i = 0]} = b.$$

We estimate the ATE in the RCT using the standard difference-in-means estimator, that is,

$$\hat{\tau}_{\text{RCT}} = \frac{1}{N_1} \sum_{i=1}^N T_i Y_i - \frac{1}{N_0} \sum_{i=1}^N (1 - T_i) Y_i,$$

where $N_1 = |\{i \in \{1, \dots, N\} : T_i = 1\}|$ and $N_0 = |\{i \in \{1, \dots, N\} : T_i = 0\}|$. We estimate

the LATE in the observational study using the Wald estimator, that is,

$$\hat{\tau}_{\text{OS}} = \left(\frac{1}{n_1} \sum_{i=N+1}^{N+n} Z_i T_i - \frac{1}{n_0} \sum_{i=N+1}^{N+n} (1 - Z_i) T_i \right)^{-1} \left(\frac{1}{n_1} \sum_{i=N+1}^{N+n} Z_i Y_i - \frac{1}{n_0} \sum_{i=N+1}^{N+n} (1 - Z_i) Y_i \right),$$

where $n_1 = |\{i \in \{N+1, \dots, N+n\} : Z_i = 1\}|$ and $n_0 = |\{i \in \{N+1, \dots, N+n\} : Z_i = 0\}|$.

We estimate the variance of $\hat{\tau}_{\text{RCT}}$ with the Neyman estimator,

$$\hat{\sigma}_{\text{RCT}}^2 = \frac{1}{N_1(N_1 - 1)} \sum_{i=1}^N T_i (Y_i - \bar{Y}_1)^2 + \frac{1}{N_0(N_0 - 1)} \sum_{i=1}^N (1 - T_i) (Y_i - \bar{Y}_0)^2,$$

where $\bar{Y}_1 = N_1^{-1} \sum_{i=1}^N T_i Y_i$ and $\bar{Y}_0 = N_0^{-1} \sum_{i=1}^N (1 - T_i) Y_i$. We estimate the variance of $\hat{\tau}_{\text{OS}}$ using a heteroskedastic-robust variance estimator given by

$$\begin{aligned} \hat{\sigma}_{\text{OS}}^2 = \frac{1}{(\bar{T}_1 - \bar{T}_0)^2} & \left(\frac{1}{n_1(n_1 - 1)} \sum_{i=N+1}^{N+n} Z_i (Y_i - \tilde{Y}_1 - \hat{\tau}_{\text{OS}}(T_i - \bar{T}_1))^2 \right. \\ & \left. + \frac{1}{n_0(n_0 - 1)} \sum_{i=N+1}^{N+n} (1 - Z_i) (Y_i - \tilde{Y}_0 - \hat{\tau}_{\text{OS}}(T_i - \bar{T}_0))^2 \right), \end{aligned}$$

where $\bar{T}_1 = n_1^{-1} \sum_{i=N+1}^{N+n} Z_i T_i$, $\bar{T}_0 = n_0^{-1} \sum_{i=N+1}^{N+n} (1 - Z_i) T_i$, $\tilde{Y}_1 = n_1^{-1} \sum_{i=N+1}^{N+n} Z_i Y_i$, and $\tilde{Y}_0 = n_0^{-1} \sum_{i=N+1}^{N+n} (1 - Z_i) Y_i$. All these estimators are standard choices in economics practice.

5.3 Results of the Monte Carlo simulations

The results of the simulations for Design 1 are collected in Table 1 and plotted in Figure 3. The results of the simulations for Design 2 are collected in Table 9 and plotted in Figure 6 in Appendix C.2. Appendix C.3 contains additional visualizations of the same results, including close-ups on the cases $b \in \{0, 0.2, 0.5, 1\}$ where we decompose the MSE in squared bias and variance and plot the empirical distributions of the estimators over the simulations – see Figure 9 and Figure 10.

The results of the simulations are consistent across the two designs and highlight three main facts. First, the pretesting estimator exhibits a suboptimal hump for moderate values of the bias due to the hard threshold. This hump is expected and leads to it being dominated by all the minimax regret and robust empirical Bayes estimators for a large range of bias values. Secondly, the estimators are seen to navigate the bias–variance trade-off as two


Table 1: Monte Carlo – Bias, Variance, and MSE of estimators in Design 1

	LATE (b)	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Panel A: Benchmark Estimators												
RCT	Bias ²	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
	Variance	0.0785	0.0785	0.0788	0.0792	0.0799	0.0808	0.0818	0.0831	0.0846	0.0863	0.0883
	MSE	0.0785	0.0785	0.0788	0.0793	0.0799	0.0808	0.0819	0.0832	0.0847	0.0864	0.0883
MLE	Bias ²	0.0000	0.0096	0.0379	0.0851	0.1514	0.2369	0.3418	0.4662	0.6101	0.7736	0.9567
	Variance	0.0029	0.0028	0.0027	0.0026	0.0025	0.0024	0.0023	0.0022	0.0022	0.0021	0.0021
	MSE	0.0029	0.0124	0.0406	0.0877	0.1539	0.2393	0.3441	0.4684	0.6123	0.7757	0.9588
Pretesting	Bias ²	0.0001	0.0051	0.0159	0.0302	0.0396	0.0344	0.0274	0.0193	0.0085	0.0043	0.0024
	Variance	0.0244	0.0305	0.0490	0.0762	0.1109	0.1470	0.1736	0.1860	0.1715	0.1581	0.1472
	MSE	0.0245	0.0356	0.0650	0.1064	0.1505	0.1813	0.2010	0.2052	0.1801	0.1623	0.1495
Panel B: Minimax Regret Estimators												
Moment Regret	Bias ²	0.0000	0.0057	0.0207	0.0411	0.0624	0.0803	0.0924	0.0992	0.1028	0.1055	0.1094
	Variance	0.0100	0.0116	0.0172	0.0269	0.0397	0.0539	0.0673	0.0773	0.0823	0.0829	0.0818
	MSE	0.0100	0.0173	0.0378	0.0680	0.1021	0.1341	0.1597	0.1766	0.1851	0.1884	0.1912
Support Regret	Bias ²	0.0000	0.0061	0.0225	0.0456	0.0709	0.0954	0.1159	0.1317	0.1441	0.1542	0.1626
	Variance	0.0096	0.0110	0.0152	0.0220	0.0314	0.0424	0.0531	0.0625	0.0703	0.0766	0.0812
	MSE	0.0096	0.0171	0.0376	0.0676	0.1024	0.1378	0.1690	0.1942	0.2144	0.2308	0.2438
Panel C: Robust Empirical Bayes Estimators												
Moment Plug-in	Bias ²	0.0000	0.0026	0.0092	0.0158	0.0207	0.0224	0.0210	0.0185	0.0158	0.0133	0.0113
	Variance	0.0281	0.0319	0.0420	0.0563	0.0732	0.0880	0.0981	0.1045	0.1075	0.1080	0.1076
	MSE	0.0281	0.0345	0.0512	0.0721	0.0939	0.1105	0.1191	0.1230	0.1234	0.1213	0.1189
Support Plug-in	Bias ²	0.0000	0.0020	0.0071	0.0124	0.0175	0.0186	0.0187	0.0181	0.0161	0.0140	0.0124
	Variance	0.0330	0.0354	0.0440	0.0559	0.0692	0.0800	0.0892	0.0967	0.1012	0.1036	0.1065
	MSE	0.0330	0.0373	0.0511	0.0683	0.0867	0.0986	0.1079	0.1148	0.1174	0.1176	0.1189
Huber Plug-in	Bias ²	0.0000	0.0042	0.0140	0.0250	0.0310	0.0312	0.0295	0.0221	0.0171	0.0135	0.0109
	Variance	0.0246	0.0299	0.0427	0.0627	0.0859	0.1081	0.1254	0.1259	0.1255	0.1237	0.1211
	MSE	0.0246	0.0340	0.0567	0.0877	0.1169	0.1393	0.1549	0.1480	0.1425	0.1372	0.1320

distinct groups defined by the adaptation approach: the minimax regret estimators favor variance reduction whereas the robust empirical Bayes estimators favor bias protection. This separation is clear through the MSE profiles in Figure 3 and Figure 6. The same results also follow immediately from the MSE decompositions in Appendix C.3. In other words, even if both types of estimators are robust, adaptive, and derived from the same primitives, the adaptation approach fundamentally determines how the bias–variance trade-off is navigated in practice. The minimax regret estimators trade off improved performance for small values of the bias against worse performance for larger values of it, and vice versa for the robust empirical Bayes estimators. The consequences of these results for econometric practices are discussed in Section 6 through a real-data application in education economics. Finally, the Bayesian primitives and the approximations to compute the estimators are seen to have

only a secondary effect on the performance of the estimators. This holds for both the robust empirical Bayes estimators computed using Algorithm [ReBA](#) and the minimax regret estimators. The close proximity of the MSE profiles for the two minimax regret estimators corroborates, in particular, the results and subsequent comments of Remark [3.3](#). Given this equivalence, the natural recommendation is to favor Bayesian primitives that favor simpler adaptive estimators (i.e., $\Gamma_1(M)$).

Figure 3: Monte Carlo – MSE of the estimators in Design 1



6 Example in education economics (Project STAR)

In this section, we use real data to illustrate the use of the adaptive estimators derived in Section [3](#) and Section [4](#). We focus on a central problem in education economics that consists of ascertaining the effects of class sizes on students' test scores. This problem has generated an extensive body of work and has led to one of the most influential, large-scale randomized controlled experiments within economics and public policy: Project STAR (Student-Teacher Achievement Ratio). The experiment was conducted in 79 schools across Tennessee over four years (1985–1989) and randomly assigned one cohort of 11,571 students and their teachers to classrooms of different sizes from Kindergarten to Grade 3. This large-scale RCT unearthed statistically significant positive effects of class size reduction on test scores. This experiment was followed by many large-scale observational studies performed in different school systems across the world – see, for instance, [Angrist and Lavy \(1999\)](#) and [Angrist, Lavy, Leder-Luis, and Shany \(2019\)](#) in Israel, [Hoxby \(2000\)](#) in Connecticut, or [Leuven, Oosterbeek, and Rønning \(2008\)](#) in Norway – pointing (with a few exceptions) to effects not significantly different from zero. This difference between Project STAR and these large-scale observational studies has

since generated a large and opinionated discussion on the real effect of class size on student achievement. For a detailed review of the experiment and the results, we refer to [Krueger \(1999\)](#) and [Schanzenbach \(2006\)](#).

Our objective is not to arbitrate the debate with new evidence but to show how the effect from a small-scale RCT can be appropriately combined with the effect from a large-scale observational study using the adaptive estimators of [Section 3](#) and [Section 4](#). The large-scale RCT from Project STAR provides an ideal laboratory for this purpose. We use it to perform a simple exercise where we randomly drop a significant fraction of the original experimental dataset (between 80% and 98%), compute the effect of class size reduction on the resulting small-scale RCT, and use the adaptive estimators to combine the RCT estimates with estimated effects from observational studies available in the literature. We average the results by repeating the exercise across random draws and use the estimated effect from the original full-sample RCT as the true causal effect $\tau \in \mathbb{R}$ that we try to recover from the small-scale RCT.

Given the purpose of our exercise, we consider standard specifications and estimators that are directly drawn from the existing literature. In the RCTs, we focus on the effect of class size reduction for Grade 3 and use a simple difference-in-means estimator. We compute it by OLS and use Huber-White robust standard errors (HC1). The effects are standardized by dividing the estimates by the standard deviations of the test scores. To obtain comparable results, we consider two observational studies in the United States within ten years of Project STAR’s conclusion. The first one is taken from [Stock and Watson \(2015\)](#) and is based on all K–6 and K–8 districts in California with data available for 1999. The treatment is measured by the district-wide student-to-teacher ratio and the test scores are the district-wide averages of reading and math scores in Grade 5. The second is taken from [Hoxby \(2000\)](#) for all elementary schools in Connecticut with data available between 1992 and 1998. The treatment is measured by predicting the class size using residual enrollment and the test scores are the school-wide averages in reading scores in Grade 4 divided by the standard deviation in test scores at the student level across Connecticut. The effects in the California study are estimated by OLS with a few controls (English learners, free lunch eligibility, and average district income) and standardized by the standard deviation of the test scores across students in California. The effects in the Connecticut study are estimated with an instrumental variable (enrollment residual) by 2SLS with fixed effects and observable characteristics at

each stage as given by Equations 7 and 8 in [Hoxby \(2000\)](#). In both cases, we compute the effect of a class size reduction comparable to the average class size reduction experienced in Tennessee during Project STAR. The estimated effects and corresponding standard errors are collected in [Table 2](#) below.

Table 2: Estimates and Standard Errors from Full-Sample RCT and Large-Scale OS

	Full-sample STAR RCT	California 1999 OS	Connecticut 1992–1998 OS
Estimated Effect	0.213	0.145	0.033
Standard Error	(0.033)	(0.051)	(0.034)

To compare the estimators of [Section 3](#) and [Section 4](#), we repeatedly and randomly draw a limited subsample from the experimental data and estimate the effect and the standard deviation of the estimator on each subsample using the same empirical strategy as for the full-sample large-scale RCT. We then use the same eight estimators used in the Monte Carlo simulations of [Section 5](#) to combine the estimators $\hat{\tau}_{\text{RCT}}$ of the subsampled small-scale RCTs with each of the two estimators $\hat{\tau}_{\text{OS}}$ from the large-scale observational studies. Over the subsamples, we compute for each estimator (with respect to the full-sample RCT estimated effect), the average bias, the empirical variance, the average MSE, the average coverage and width of 95% confidence intervals obtained by the parametric bootstrap, and the empirical probability that the intervals do not contain 0. We use 500 repetitions for the subsamples and 500 bootstrap draws for each repetition. The results are collected in [Table 3](#) and [Table 4](#) below for subsamples where 80%, 95%, and 98% of the original experimental dataset is dropped. This corresponds to a respective reduction of the sample size from 5,967 students to 1,193, to 298, and to 119 students on average with a respective increase in the standard error for the RCT estimator from 0.033 to 0.073, to 0.147, and to 0.233.

Within the scope of this exercise, the two observational studies play opposite roles. Because we take the true causal effect to be the estimated effect in the full-sample RCT, the California study corresponds to a small-bias scenario and the Connecticut study to a large-bias scenario. The results in [Table 3](#) and [Table 4](#) are then consistent with the Monte Carlo simulations – see [Section 5.3](#) – and can be used to highlight two additional facts.

The first fact is that the difficulty of the combination task depends crucially on the signal-to-noise ratio for the bias. Across both scenarios, the combination results based on the minimax regret estimators and the robust empirical Bayes estimators only separate when a

Table 3: Combination results with the *small-bias* California study

	Estimator	Point Estimate	Error Decomposition			Confidence Interval (CI)		
		Estimate	MSE	Bias ²	Variance	$P(0 \notin \text{CI})$	Coverage	CI Width
80% dropped SE = 0.073 N = 1,193	RCT	0.216	0.005	0.000	0.005	86.2%	95.4%	0.284
	MLE	0.168	0.003	0.002	0.001	100.0%	95.2%	0.162
	Pretesting	0.178	0.004	0.001	0.003	92.0%	95.2%	0.268
	Moment Regret	0.196	0.003	0.000	0.003	92.6%	95.6%	0.244
	Support Regret	0.189	0.003	0.001	0.003	96.2%	95.6%	0.229
	Moment Plug-in	0.188	0.004	0.001	0.003	94.2%	96.0%	0.249
	Support Plug-in	0.192	0.004	0.000	0.003	93.4%	96.8%	0.249
	Huber Plug-in	0.181	0.004	0.001	0.003	93.2%	96.0%	0.254
95% dropped SE = 0.147 N = 298	RCT	0.220	0.019	0.000	0.019	34.6%	95.8%	0.570
	MLE	0.153	0.004	0.004	0.000	100.0%	98.8%	0.187
	Pretesting	0.170	0.009	0.002	0.007	69.4%	98.0%	0.484
	Moment Regret	0.178	0.006	0.001	0.005	65.4%	95.4%	0.383
	Support Regret	0.174	0.006	0.002	0.004	66.6%	96.2%	0.358
	Moment Plug-in	0.183	0.009	0.001	0.008	57.4%	95.8%	0.457
	Support Plug-in	0.189	0.009	0.001	0.009	50.6%	96.4%	0.463
	Huber Plug-in	0.174	0.009	0.002	0.007	61.8%	97.8%	0.460
98% dropped SE = 0.233 N = 119	RCT	0.227	0.051	0.000	0.051	16.4%	94.6%	0.904
	MLE	0.149	0.004	0.004	0.000	100.0%	99.8%	0.193
	Pretesting	0.170	0.018	0.002	0.016	65.2%	97.8%	0.724
	Moment Regret	0.169	0.009	0.002	0.007	60.8%	94.0%	0.495
	Support Regret	0.166	0.009	0.002	0.007	58.6%	95.4%	0.470
	Moment Plug-in	0.183	0.020	0.001	0.019	39.6%	95.4%	0.706
	Support Plug-in	0.189	0.022	0.001	0.022	32.4%	95.4%	0.717
	Huber Plug-in	0.172	0.018	0.002	0.017	51.4%	96.6%	0.702

sufficiently large share of the experimental data is dropped (that is, when the noise of the RCT estimator prevents us from isolating the bias). The second fact is that the respective profiles of the adaptive estimators matter for reaching valid conclusions. From the Monte Carlo simulations, we know that the minimax regret estimators favor variance reduction whereas the robust empirical Bayes estimators favor bias protection. The effects of this split are clearly exhibited in Table 3 and Table 4. In the small-bias California case, the minimax regret estimators outperform the robust empirical Bayes estimators at no cost on the inferential side. In the large-bias Connecticut case, however, the bias becomes too significant for the minimax regret estimators to perform satisfactorily. When the noise of the RCT estimator is too large, the minimax regret estimators lean too heavily on the observational study estimator: the bias of the estimators increases and affects the coverage of the corresponding confidence intervals. This defect is not shared by the robust empirical Bayes estimators, which lean more heavily on the RCT estimator. This fact has direct and important consequences in

Table 4: Combination results with the *large-bias* Connecticut study

	Estimator	Point Estimate	Error Decomposition			Confidence Interval (CI)		
		Estimate	MSE	Bias ²	Variance	$P(0 \notin \text{CI})$	Coverage	CI Width
80% dropped SE = 0.073 N = 1,193	RCT	0.216	0.005	0.000	0.005	86.2%	95.4%	0.284
	MLE	0.066	0.022	0.022	0.000	66.8%	0.0%	0.120
	Pretesting	0.182	0.012	0.001	0.011	81.2%	96.0%	0.329
	Moment Regret	0.173	0.006	0.002	0.005	90.2%	94.2%	0.259
	Support Regret	0.160	0.007	0.003	0.004	89.6%	90.6%	0.255
	Moment Plug-in	0.185	0.007	0.001	0.006	88.8%	94.8%	0.292
	Support Plug-in	0.185	0.006	0.001	0.005	88.2%	95.0%	0.288
	Huber Plug-in	0.182	0.008	0.001	0.007	83.0%	95.8%	0.307
95% dropped SE = 0.147 N = 298	RCT	0.220	0.019	0.000	0.019	34.6%	95.8%	0.570
	MLE	0.043	0.029	0.029	0.000	0.4%	0.0%	0.128
	Pretesting	0.117	0.033	0.009	0.024	13.2%	92.6%	0.530
	Moment Regret	0.106	0.020	0.011	0.009	37.8%	82.8%	0.375
	Support Regret	0.101	0.020	0.012	0.008	30.0%	81.2%	0.360
	Moment Plug-in	0.148	0.021	0.004	0.017	34.0%	94.6%	0.491
	Support Plug-in	0.156	0.019	0.003	0.016	34.6%	94.8%	0.495
	Huber Plug-in	0.132	0.026	0.007	0.019	21.6%	93.6%	0.497
98% dropped SE = 0.233 N = 119	RCT	0.227	0.051	0.000	0.051	16.4%	94.6%	0.904
	MLE	0.037	0.031	0.031	0.000	0.0%	0.0%	0.130
	Pretesting	0.105	0.048	0.012	0.036	5.2%	81.4%	0.754
	Moment Regret	0.077	0.028	0.018	0.010	14.4%	71.2%	0.460
	Support Regret	0.076	0.028	0.019	0.009	10.8%	76.2%	0.440
	Moment Plug-in	0.137	0.036	0.006	0.030	17.4%	91.0%	0.726
	Support Plug-in	0.146	0.036	0.004	0.032	17.4%	92.8%	0.736
	Huber Plug-in	0.113	0.041	0.010	0.031	8.8%	86.2%	0.718

terms of empirical practice. When adaptively combining estimators, particular attention should be paid to the adaptation approach used, the known guarantees it provides, and the way it empirically navigates the bias–variance trade-off. To improve transparency, we recommend that researchers perform a simple robustness check by reporting the results of adaptive estimators obtained from both the minimax regret approach and the empirical Bayes approach (preferably derived from the same Bayesian primitives). When reporting the results side-by-side, large discrepancies between the point estimates and the (estimated) sampling variances can be used to detect cases where adaptation may come at the cost of the causal credibility of the combined results.

References

- ANGRIST, J. D., AND V. LAVY (1999): “Using Maimonides’ rule to estimate the effect of class size on scholastic achievement,” *The Quarterly Journal of Economics*, 114(2), 533–575.
- ANGRIST, J. D., V. LAVY, J. LEDER-LUIS, AND A. SHANY (2019): “Maimonides rule redux,” *American Economic Review: Insights*, 1(3), 309–324.
- ARMSTRONG, T. B., P. KLINE, AND L. SUN (2025): “Adapting to misspecification,” *Econometrica*, 93(6), 1981–2005.
- ATHEY, S., R. CHETTY, AND G. IMBENS (2020): “Using experiments to correct for selection in observational studies,” *arXiv preprint arXiv:2006.09676*.
- BANCROFT, T., AND C.-P. HAN (1977): “Inference based on conditional specification: a note and a bibliography,” *International Statistical Review/Revue Internationale de Statistique*, pp. 117–127.
- BARANCHIK, A. J. (1964): “Multiple regression and estimation of the mean of a multivariate normal distribution,” *Stanford University Technical Report*, 51.
- BAREINBOIM, E., AND J. PEARL (2016): “Causal inference and the data-fusion problem,” *Proceedings of the National Academy of Sciences*, 113(27), 7345–7352.
- BERGER, J., AND L. M. BERLINER (1984): “Bayesian input in Stein estimation and a new minimax empirical Bayes estimator,” *Journal of Econometrics*, 25(1-2), 87–108.
- BERGER, J. O. (1985): “Statistical decision theory and Bayesian analysis,” *Springer Series in Statistics*.
- BICKEL, P. J. (1984): “Parametric robustness: small biases can be worthwhile,” *The Annals of Statistics*, 12(3), 864–879.
- BROWN, L. D. (1971): “Admissible estimators, recurrent diffusions, and insoluble boundary value problems,” *The Annals of Mathematical Statistics*, 42(3), 855–903.
- (1974): “Notes on decision theory,” *Lecture notes*.

- CASELLA, G., AND W. E. STRAWDERMAN (1981): “Estimating a bounded normal mean,” *The Annals of Statistics*, 9(4), 870–878.
- CHAMBERLAIN, G. (2000): “Econometric applications of maxmin expected utility,” *Journal of Applied Econometrics*, 15(6), 625–644.
- CHEN, S., S. LI, B. ZHANG, AND T. YE (2025): “Minimax rates and adaptivity in combining experimental and observational data,” *Journal of Causal Inference*, 13(1), 20240024.
- DE CHAISEMARTIN, C., AND X. D’HAULTFŒUILLE (2020): “Empirical MSE minimization to estimate a scalar parameter,” *arXiv preprint arXiv:2006.14667*.
- DELATTE, P. (2026a): “Existence of maximin priors and statistical minimax theorems,” *Manuscript*.
- (2026b): “Regret bounds for empirical Bayes combination estimators,” *Manuscript*.
- EICHENAUER, J., AND J. LEHN (1989): “Computation of gamma-minimax estimators for a bounded normal mean under squared error loss,” *Statistics & Risk Modeling*, 7(1-2), 37–62.
- EPANOMERITAKIS, A., AND D. VIVIANO (2026): “Learning what to learn: experimental design when combining experimental with observational evidence,” *arXiv preprint arXiv:2510.23434*.
- FELDMAN, I. (1991): “Constrained minimax estimation of the mean of the normal distribution with known variance,” *The Annals of Statistics*, 19(4), 2259–2265.
- GHASSAMI, A., C. LIU, A. YANG, D. RICHARDSON, I. SHPITSER, AND E. T. TCHETGEN (2022): “Combining experimental and observational data for identification and estimation of long-term causal effects,” *arXiv preprint arXiv:2201.10743*.
- GIACOMINI, R., T. KITAGAWA, AND M. READ (2026): “Robust Bayesian Analysis for Econometrics,” in *Advances in Economics and Econometrics: Twelfth World Congress*, Econometric Society Monographs, p. 117–157. Cambridge University Press.
- GILES, J. A., AND D. E. GILES (1993): “Pre-test estimation and testing in econometrics: recent developments,” *Journal of Economic Surveys*, 7(2), 145–197.

- GOURDIN, E., B. JAUMARD, AND B. MACGIBBON (1994): “Global optimization decomposition methods for bounded parameter minimax risk evaluation,” *SIAM Journal on Scientific Computing*, 15(1), 16–35.
- GREEN, E. J., AND W. E. STRAWDERMAN (1991): “A James-Stein type estimator for combining unbiased and possibly biased estimators,” *Journal of the American Statistical Association*, 86(416), 1001–1006.
- HODGES, J. L., AND E. L. LEHMANN (1952): “The use of previous experience in reaching statistical decisions,” *The Annals of Mathematical Statistics*, pp. 396–407.
- HOXBY, C. M. (2000): “The effects of class size on student achievement: New evidence from population variation,” *The Quarterly Journal of Economics*, 115(4), 1239–1285.
- IMBENS, G. W., AND J. D. ANGRIST (1994): “Identification and estimation of local average treatment effects,” *Econometrica*, 62(2), 467–475.
- JOHNSTONE, I. M. (2019): “Gaussian estimation: Sequence and wavelet models,” *Lecture notes (Version: September 16, 2019)*.
- JUDGE, G. G., AND R. C. MITTELHAMMER (2004): “A semiparametric basis for combining estimation problems under quadratic loss,” *Journal of the American Statistical Association*, 99(466), 479–487.
- KALLUS, N., A. M. PULI, AND U. SHALIT (2018): “Removing hidden confounding by experimental grounding,” *Advances in neural information processing systems*, 31.
- KEMPTHORNE, P. J. (1987): “Numerical specification of discrete least favorable prior distributions,” *SIAM Journal on Scientific and Statistical Computing*, 8(2), 171–184.
- KRUEGER, A. B. (1999): “Experimental estimates of education production functions,” *The Quarterly Journal of Economics*, 114(2), 497–532.
- LE CAM, L. (1955): “An extension of Wald’s theory of statistical decision functions,” *The Annals of Mathematical Statistics*, 26(1), 69–81.
- LEUVEN, E., H. OOSTERBEEK, AND M. RØNNING (2008): “Quasi-experimental estimates of the effect of class size on achievement in Norway,” *The Scandinavian Journal of Economics*, 110(4), 663–693.

- LIN, Z., P. J. BICKEL, AND P. DING (2026): “Introducing the b-value: combining unbiased and biased estimators from a sensitivity analysis perspective,” *arXiv preprint arXiv:2602.16310*.
- MONTIEL OLEA, J. L., AND E. ZUBOVA (2026): “Approximate minimax estimation of a bounded normal mean via Stochastic Mirror Ascent,” *arXiv preprint arXiv:2607.05350*.
- NELSON, W. (1966): “Minimax solution of statistical decision problems by iteration,” *The Annals of Mathematical Statistics*, 37(6), 1643 – 1657.
- NOUBIAP, R. F., AND W. SEIDEL (2001): “An algorithm for calculating Γ -minimax decision rules under generalized moment conditions,” *Annals of statistics*, pp. 1094–1116.
- RIDDER, G., AND R. MOFFITT (2007): “The econometrics of data combination,” *Handbook of econometrics*, 6, 5469–5547.
- ROSENBAUM, P. R., AND D. B. RUBIN (1983): “Assessing sensitivity to an unobserved binary covariate in an observational study with binary outcome,” *Journal of the Royal Statistical Society: Series B (Methodological)*, 45(2), 212–218.
- ROSENMAN, E. T. (2025): “Methods for combining observational and experimental causal estimates: A review,” *Wiley Interdisciplinary Reviews: Computational Statistics*, 17(2), e70027.
- ROSENMAN, E. T., G. BASSE, A. B. OWEN, AND M. BAIOCCHI (2023): “Combining observational and experimental datasets using shrinkage estimators,” *Biometrics*, 79(4), 2961–2973.
- SCHANZENBACH, D. W. (2006): “What have researchers learned from Project STAR?,” *Brookings Papers on Education Policy*, (9), 205–228.
- SCLOVE, S. L., C. MORRIS, AND R. RADHAKRISHNAN (1972): “Non-optimality of preliminary-test estimators for the mean of a multivariate normal distribution,” *The Annals of Mathematical Statistics*, pp. 1481–1490.
- STEIN, C. (1966): “An approach to the recovery of interblock information in balanced incomplete block designs,” *Research paper in statistics: Festschrift for J. Neyman*, pp. 351–366.

STOCK, J., AND M. WATSON (2015): *Introduction to Econometrics*. Pearson.

WU, B., S. SALAZAR, D. P. GREEN, AND D. M. BLEI (2026): “The illusion of learning from observational data: an empirical Bayes perspective,” *arXiv preprint arXiv:2604.08853*.

Appendix A Proofs and additional results

A.1 Proofs and results for Section 2

We start by proving the decomposition of the integrated risk under the L^2 loss for equivariant estimators and use it to prove Proposition 2.1, which is one of our main existence results – it guarantees that unique robust Bayes estimators exist for any given set of priors and provides a simple formula for their computation. We derive in Corollary A.1 another expression of the robust Bayes estimator based on Tweedie’s formula. In Corollary A.2, we derive a strict equivalence result between frequentist minimax and robust Bayes estimators when the sets of priors are rich enough. The rest of the section provides the proofs of the three propositions in Section 2.3 clarifying the relation between the sets of priors and the performance of the robust Bayes estimators in terms of (squared) bias and variance. We conclude the section with Remark A.1 where we make explicit the extent to which the results of Section 2 generalize to other losses than the L^2 loss.

Proof of Lemma 2.1. From the risk expansion $R(g, \tau) = w\sigma_{\text{OS}}^2 + w^2 R(h, b)$ for $g \in G$, we directly have, for any $P \in \Lambda$, that $\int_{\mathbb{R}^2} R(g, \tau) dP(\tau, b) = w\sigma_{\text{OS}}^2 + w^2 \int_{\mathbb{R}^2} R(h, b) dP(\tau, b) = w\sigma_{\text{OS}}^2 + w^2 \int_{\mathbb{R}} R(h, b) d\pi(b)$ where the last equality holds without integrability conditions by non-negativity of the integrand. $\square$

Proof of Proposition 2.1. The continuity and strict convexity of the L^2 loss on $\mathbb{R}$ guarantee the existence of a unique (a.e.) robust Bayes estimator $h_{\Gamma}^* \in \mathcal{D}(\mathbb{R})$ for the normal mean problem – see Delatte (2026a) for a rigorous construction for decision rules with values in $\mathbb{R} \cup \{\infty\}$ based on the regular topology on decision rules introduced by Le Cam (1955) and revisited by Brown (1974). From Lemma 2.1, we have for any $g \in G$,

$$\sup_{\pi \in \Gamma} B(g, \pi) = w\sigma_{\text{OS}}^2 + w^2 \sup_{\pi \in \Gamma} B(h, \pi),$$

where $h \in \mathcal{D}(\mathbb{R})$ is of the usual form. Since $w > 0$, the function $\mathcal{L}: G \rightarrow \mathcal{D}(\mathbb{R})$ that maps $g(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + \mathcal{G}(D)$ to $h(D) = D - w^{-1}\mathcal{G}(D)$ is a well-defined bijection. It also follows that the function $x \mapsto w\sigma_{\text{OS}}^2 + w^2 x$ is strictly increasing and so

$$\arg \min_{g \in G} \sup_{\pi \in \Gamma} B(g, \pi) = \mathcal{L}^{-1} \left(\arg \min_{h \in \mathcal{D}(\mathbb{R})} \sup_{\pi \in \Gamma} B(h, \pi) \right).$$

From the existence and uniqueness of h_Γ^* , we thus directly obtain that the robust Bayes estimator g_Γ^* exists and is unique with $\mathcal{G}_\Gamma^* = wD - wh_\Gamma^*(D)$. $\square$

Corollary A.1. *Let $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ be a convex set such that a least favorable prior π_Γ^* for the estimation of b from the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$ exists and is unique. Then the robust Bayes estimator g_Γ^* rewrites as*

$$g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} - \sigma_{\text{RCT}}^2 \frac{d}{dD} \log \psi_{\pi_\Gamma^*}(D),$$

where $\psi_{\pi_\Gamma^*}(D) = \int_{\mathbb{R}} \phi(D; b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2) d\pi_\Gamma^*(b)$ is the marginal density of D under the prior π_Γ^* .

Proof. Tweedie's formula directly yields $h_\Gamma^*(D) = D + (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2) \psi'_{\pi_\Gamma^*}(D) / \psi_{\pi_\Gamma^*}(D)$ – see, for instance, (1.2.2) in [Brown \(1971\)](#). By plugging it in Proposition 2.1, we directly obtain that $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} - \sigma_{\text{RCT}}^2 \psi'_{\pi_\Gamma^*}(D) / \psi_{\pi_\Gamma^*}(D)$. $\square$

Corollary A.2. *Let $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ and $\mathcal{B} \subseteq \mathbb{R}$. Suppose that Γ contains all Dirac measures on $\mathcal{B}$ and that for any $\pi \in \Gamma$, $\text{supp}(\pi) \subseteq \mathcal{B}$. Then the robust Bayes estimator h_Γ^* is the frequentist minimax estimator of b over the restricted parameter space $\mathcal{B}$ based on the observation $D \sim \mathcal{N}(b, \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. In this case, the robust Bayes estimator g_Γ^* is the frequentist minimax estimator of τ over $\mathbb{R} \times \mathcal{B}$.*

Proof. Since $\text{supp}(\pi) \subseteq \mathcal{B}$, we have $B(h, \pi) \leq \sup_{b \in \mathcal{B}} R(h, b)$ and so $\sup_{\pi \in \Gamma} B(h, \pi) \leq \sup_{b \in \mathcal{B}} R(h, b)$. Since Γ contains all Dirac measures on $\mathcal{B}$, we have $\sup_{\pi \in \Gamma} B(h, \pi) \geq \sup_{b \in \mathcal{B}} B(h, \delta_b) = \sup_{b \in \mathcal{B}} R(h, b)$. It follows that $\sup_{\pi \in \Gamma} B(h, \pi) = \sup_{b \in \mathcal{B}} R(h, b)$. We conclude directly from Lemma 2.1 and Proposition 2.1 since the risk invariance holds for any $b \in \mathbb{R}$ and the existence result holds for any $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$. $\square$

Proof of Proposition 2.2. 1. We have $\text{Cov}(g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D), D) = 0$. It follows that $g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D)$ and D are independent from the joint normality of D and $\hat{\tau}_{\text{RCT}}$. Since any estimator $g \in G$ can be rewritten immediately as $g(\hat{\tau}_{\text{RCT}}, D) = g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D) + \bar{\mathcal{G}}(D)$ where $\bar{\mathcal{G}}(D) = \mathcal{G}(D) - wD$, we have $\text{Var}(g(\hat{\tau}_{\text{RCT}}, D)) = \text{Var}(g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D)) + \text{Var}(\bar{\mathcal{G}}(D))$ by independence, yielding the inequality. For equality, we must have $\bar{\mathcal{G}}(D)$ constant (a.e.). Since $g_{\Gamma_0}^*$ is unbiased at $b = 0$, it follows that $\mathbb{E}[\bar{\mathcal{G}}(D)] = 0$ at $b = 0$, and so $\mathcal{G}(D) = wD$ (a.e.), which is enough to ensure the uniqueness of $g_{\Gamma_0}^*$ among estimators achieving the bound.

2. The estimator $g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}}$ satisfies $\text{Bias}(g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D)) = 0$ for all $b \in \mathbb{R}$. For any estimator $g \in G$, we have $\text{Bias}(g(\hat{\tau}_{\text{RCT}}, D)) = \mathbb{E}[\mathcal{G}(D)]$, and so g achieves the bound if and only if $\mathbb{E}[\mathcal{G}(D)] = 0$ for all $b \in \mathbb{R}$. By completeness of the normal family, we have that $\mathcal{G}(D) = 0$ (a.e.), which is enough to ensure the uniqueness of $g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}}$ among estimators achieving the bound. $\square$

Proof of Proposition 2.3. From Proposition 2.1, we have $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + wD - wh_\Gamma^*(D)$. Taking the expectation yields $\mathbb{E}[g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D)] - \tau = wb - w\mathbb{E}[h_\Gamma^*(D)]$ from which the bias result follows. For the variance, we have that $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D) - wh_\Gamma^*(D)$. We already know that $g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D)$ and D are independent, and so the result follows directly. $\square$

Proof of Proposition 2.4. From Proposition 2.1, we have $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + wD - wh_\Gamma^*(D)$, and so $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + wD(1 - h_\Gamma^*(D)/D)$ (a.e.). From the expressions of $g_{\Gamma_0}^*$ and $g_{\Gamma_\infty}^*$, we directly have $g_\Gamma^*(\hat{\tau}_{\text{RCT}}, D) = W(D)g_{\Gamma_0}^*(\hat{\tau}_{\text{RCT}}, D) + (1 - W(D))g_{\Gamma_\infty}^*(\hat{\tau}_{\text{RCT}}, D)$ (a.e.) with $W(D) = 1 - h_\Gamma^*(D)/D$ (a.e.). The fact that $0 \leq W(D) \leq 1$ (a.e.) is directly seen to be equivalent to $0 \leq h_\Gamma^*(D)/D \leq 1$ (a.e.). $\square$

Remark A.1. A natural question is whether the results in the main text generalize to loss functions other than the L^2 loss. As one may expect, the answer is partly negative. The decomposition of the risk and the integrated risk obtained in Lemma 2.1 does not hold for any other loss. This is a direct and immediate consequence of the Pythagorean theorem: the risk decomposes exactly into squared bias and variance only under the L^2 loss. Without the decomposition of Lemma 2.1, we cannot isolate the conditional variance of $\hat{\tau}_{\text{RCT}}$ given D from the risk of estimating b . As a result, we have to abandon the equivalence with the univariate normal mean problem and all the closed-form solutions we obtained in the main text. Nonetheless, provided the loss remains translation-invariant, we can preserve the equivariance of the estimators, $g(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + \mathcal{G}(D)$, and thus the reduction from a bivariate robust Bayesian decision problem to a univariate one. In contrast to the L^2 case, however, the optimal choice of $\mathcal{G}(D)$ depends in this more general setting on the conditional variance of $\hat{\tau}_{\text{RCT}}$ given D . Moreover, to solve these problems, existence results for the normal mean problem cannot be used anymore; specific existence and uniqueness arguments need to be invoked – see, for instance, Delatte (2026a). The solutions have then to be computed numerically using extensions of Algorithm GNM-LFP.

A.2 Proofs for Section 3

We start by reviewing the frequentist notion of adaptation regret considered in [Armstrong, Kline, and Sun \(2025\)](#) that motivates our Bayesian extension. We then prove our main reduction result that allows us to solve the adaptation problem – Theorem 3.1 – and use it to prove the existence result for minimax regret estimators – Corollary 3.2. We then prove Corollary 3.3 and Corollary 3.4 where we verify that the conditions of the theorem apply for the classes based on $\Gamma_2(m)$ and $\Gamma_1(M)$, respectively.

Remark A.2. Following earlier contributions of [Hodges and Lehmann \(1952\)](#) and [Bickel \(1984\)](#), [Armstrong, Kline, and Sun \(2025\)](#) considers a solution to the adaptation problem that parallels the approach in Section 3.1 but for a minimax problem based on the frequentist risk and the hypothesis of a bounded parameter space $[-m, m]$ for the bias (as compared to the integrated risk and the hypothesis of a set of priors on the bias). This leads to a frequentist notion of adaptation regret, which, for our problem and with our notations, reads as

$$A(g_0, m) = \frac{\sup_{(\tau, b) \in \mathbb{R} \times [-m, m]} R(g_0, \tau)}{R^*(\mathbb{R} \times [-m, m])},$$

for any $g_0 \in G$ and any $m \geq 0$. In the adaptation problem then, Nature does not choose a set of priors on the parameter, but a bound on the parameter space. This leads to the minimax problem

$$\inf_{g \in G} \sup_{m \geq 0} A(g, m).$$

Based on similar arguments as in Section 2, [Armstrong, Kline, and Sun \(2025\)](#) obtains the following reduction

$$\inf_{g \in G} \sup_{m \geq 0} A(g, m) = \inf_{h \in \mathcal{D}(\mathbb{R})} \sup_{b \in \mathbb{R}} \frac{\sigma_{\text{OS}}^2 + wR(h, b)}{\sigma_{\text{OS}}^2 + wR^*([-|b|, |b|])},$$

where $R^*([-|b|, |b|])$ is the frequentist minimax risk in the univariate bounded normal mean problem with bounds $[-|b|, |b|]$. This corresponds to an application of Theorem 1 in [Armstrong, Kline, and Sun \(2025\)](#) to our problem with our notations. As in our setting, the restriction of the problem to G can be upgraded to the whole space $\mathcal{D}(\mathbb{R}^2)$ by minimaxity. Our Bayesian extension introduced in Section 3 serves two main purposes: it allows us to consider richer constraints on the bias than only bounds on the parameter space; it provides solutions with

clear computational advantages by removing the need to approximate $R^*([-m, m])$ for each value of m – see Corollary 3.4 for an example. The caveat is that it comes with a slightly more difficult reduction argument as shown in the proof of Theorem 3.1.

Proof of Theorem 3.1. For any $g \in G$, let $C(g) = \sup_{b \in \mathbb{R}} R(g, \tau) / (w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(s(b))))$. By definition of s , we have $\delta_b \in \Gamma(s(b))$ for all $b \in \mathbb{R}$. It follows that for any $g \in G$, $C(g) \leq \sup_{\eta \in S} \mathcal{A}(g, \Gamma(\eta))$. Suppose now without loss of generality that $C(g) < \infty$. Then $R(g, \tau) \leq C(g)(w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(s(b))))$ for all $b \in \mathbb{R}$. By standard properties of the integral, we have for any $\eta \in S$ and any $\pi \in \Gamma(\eta)$, $B(g, \pi) \leq C(g)(w\sigma_{\text{OS}}^2 + w^2 \int_{\mathbb{R}} B^*(\Gamma(s(b))) d\pi(b))$. From the integral inequality, we obtain that $B(g, \pi) \leq C(g)(w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(\eta)))$, and so $B(g, \pi) / (w\sigma_{\text{OS}}^2 + w^2 B^*(\Gamma(\eta))) \leq C(g)$. Taking the supremum over $\pi \in \Gamma(\eta)$ and the supremum over $\eta \in S$ yields $\sup_{\eta \in S} \mathcal{A}(g, \Gamma(\eta)) \leq C(g)$. An application of Lemma 2.1 as in the proof of Proposition 2.1 allows us to conclude. $\square$

Proof of Corollary 3.2. From the construction and equivalence given by Theorem 3.1, the proof of Corollary 3.2 is almost identical to the proof of Proposition 2.1. We only need to guarantee the existence of a frequentist minimax estimator for the problem defined by $W_1(b) + W_2(b)R(h, b)$. Since for any $b \in \mathbb{R}$, $W_1(b) + W_2(b)R(h, b)$ is a positive affine transformation of $R(h, b)$, the lower semi-continuity of $R(h, b)$ (with respect to the regular topology on the compact space of randomized decision rules) is preserved. From the arguments in Johnstone (2019) or Delatte (2026a), this is enough to guarantee that a frequentist minimax estimator for the rescaled risk function $W_1(b) + W_2(b)R(h, b)$ exists. The existence of a minimax regret estimator then follows as in the proof of Proposition 2.1. $\square$

Proof of Corollary 3.3. For $\Gamma_2(m)$, we have $\delta_b \in \Gamma_2(m)$ if and only if $|b| \leq m$, and so $s(b) = |b|$. Since $\{\Gamma_2(m) : m > 0\}$ is an increasing family of sets, we have that $m \mapsto B^*(\Gamma_2(m))$ is non-decreasing. For any $\pi \in \Gamma_2(m)$, we have $|b| \leq m$ (π -a.s.), and so $B^*(\Gamma_2(|b|)) \leq B^*(\Gamma_2(m))$ (π -a.s.). It follows that $\int_{\mathbb{R}} B^*(\Gamma_2(|b|)) d\pi(b) \leq B^*(\Gamma_2(m))$ and so Theorem 3.1 applies. The fact that $B^*(\Gamma_2(|b|)) = R^*([-|b|, |b|])$ from Corollary A.2 allows us to conclude. $\square$

Proof of Corollary 3.4. For $\Gamma_1(M)$, we have $\delta_b \in \Gamma_1(M)$ if and only if $b^2 \leq M$, and so $s(b) = b^2$. The function $M \mapsto B^*(\Gamma_1(M)) = M(\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2) / (M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$ is strictly concave and non-decreasing on $(0, \infty)$. For any $\pi \in \Gamma_1(M)$, we have $\int_{\mathbb{R}} b^2 d\pi(b) \leq M$. By Jensen's inequality and monotonicity, we then have $\int_{\mathbb{R}} B^*(\Gamma_1(b^2)) d\pi(b) \leq B^*(\Gamma_1(\int_{\mathbb{R}} b^2 d\pi(b))) \leq B^*(\Gamma_1(M))$ and so Theorem 3.1 applies. $\square$

A.3 Proofs and results for Section 4

We start with a simple equivalence and its proof for the robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ over $\{\Gamma_1(M) : M \in [0, \infty)\}$. We then extend the results of Section 2 to the multivariate case and use these extensions to consider the application of the empirical Bayes approach based on an extension of the class defined by $\Gamma_1(M)$. The main result is the recovery of the estimator considered in [Green and Strawderman \(1991\)](#) as a particular application of our robust Bayesian approach – see Proposition A.4 and Corollary A.5.

Corollary A.3. *The robust empirical Bayes estimator $g_{\Gamma_1(\hat{M})}^*$ rewrites as*

$$g_{\Gamma_1(\hat{M})}^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + \min \left(w, \frac{\sigma_{\text{RCT}}^2}{D^2} \right) D.$$

Proof. From Proposition 2.1, we have $g_{\Gamma_1(\hat{M})}^*(\hat{\tau}_{\text{RCT}}, D) = \hat{\tau}_{\text{RCT}} + wD - wh_{\Gamma_1(\hat{M})}^*(D)$. If $D^2 \leq \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, then $wD - wh_{\Gamma_1(\hat{M})}^*(D) = wD$. If $D^2 > \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, then $wD - wh_{\Gamma_1(\hat{M})}^*(D) = \sigma_{\text{RCT}}^2/D$. Since $wD^2 \leq \sigma_{\text{RCT}}^2$ if and only if $D^2 \leq \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$, we can directly conclude. $\square$

We now extend the results of Section 2 derived in the one-dimensional case $p = 1$ to the multivariate case $p > 1$. We show that most of the results generalize straightforwardly for appropriate choices of covariance matrices in the multivariate normal reduction. We now assume that $\hat{\boldsymbol{\tau}}_{\text{RCT}} = (\hat{\tau}_{\text{RCT},1}, \dots, \hat{\tau}_{\text{RCT},p})$ and $\hat{\boldsymbol{\tau}}_{\text{OS}} = (\hat{\tau}_{\text{OS},1}, \dots, \hat{\tau}_{\text{OS},p})$ are independent where $\hat{\boldsymbol{\tau}}_{\text{RCT}} \sim \mathcal{N}(\boldsymbol{\tau}, \Sigma_{\text{RCT}})$ and $\hat{\boldsymbol{\tau}}_{\text{OS}} \sim \mathcal{N}(\boldsymbol{\tau} + \mathbf{b}, \Sigma_{\text{OS}})$ with $\mathbf{b} = (b_1, b_2, \dots, b_p) \in \mathbb{R}^p$, $\boldsymbol{\tau} = (\tau_1, \tau_2, \dots, \tau_p) \in \mathbb{R}^p$, and $\Sigma_{\text{RCT}}, \Sigma_{\text{OS}} \in \mathbb{R}^{p \times p}$ known positive-definite matrices with $\Sigma_{\text{RCT}} - \Sigma_{\text{OS}}$ itself being positive (semi-)definite. We can define the vector $\mathbf{D} = \hat{\boldsymbol{\tau}}_{\text{OS}} - \hat{\boldsymbol{\tau}}_{\text{RCT}}$ where $\mathbf{D} \sim \mathcal{N}(\mathbf{b}, \Sigma_{\text{RCT}} + \Sigma_{\text{OS}})$. For an equivariant estimator $g(\hat{\boldsymbol{\tau}}_{\text{RCT}}, \mathbf{D}) = \hat{\boldsymbol{\tau}}_{\text{RCT}} + \mathcal{G}(\mathbf{D})$, the risk decomposition leading to Lemma 2.1 then directly extends to

$$R(g, \boldsymbol{\tau}) = \text{tr}(\Sigma_{\text{RCT}} - \mathbf{w}\Sigma_{\text{RCT}}) + \mathbb{E} \left[(h(\mathbf{D}) - \mathbf{b})^T \mathbf{w}^T \mathbf{w} (h(\mathbf{D}) - \mathbf{b}) \right],$$

with $\mathbf{w} = \Sigma_{\text{RCT}}(\Sigma_{\text{RCT}} + \Sigma_{\text{OS}})^{-1}$ and $h(\mathbf{D}) = \mathbf{D} - \mathbf{w}^{-1}\mathcal{G}(\mathbf{D})$. For $\Gamma \subseteq \mathcal{M}_1(\mathbb{R}^p)$, the robust Bayes estimator then takes the form

$$g_{\Gamma}^*(\hat{\boldsymbol{\tau}}_{\text{RCT}}, \mathbf{D}) = \left(I_p - \mathbf{w} \right) \hat{\boldsymbol{\tau}}_{\text{RCT}} + \mathbf{w} \left(\hat{\boldsymbol{\tau}}_{\text{OS}} - h_{\Gamma}^*(\mathbf{D}) \right),$$

where h_{Γ}^* is the robust Bayes estimator for the normal mean problem under the L^2 loss

weighted by $\mathbf{w}^T \mathbf{w}$. When the errors are spherical, the loss reduces to the standard L^2 one.

Remark A.3. Regardless of the structure of the errors, the problem can be solved numerically using multivariate extensions of Algorithm [GNM-LFP](#). This solution works in low dimensions. For large dimensions, other algorithms need to be considered to avoid the dimensionality bottleneck. The problem is of independent interest and is left for future research.

We now use these multivariate extensions to implement the empirical Bayes approach to adaptation for a multivariate extension of the class based on $\Gamma_1(M)$. In this case, we assume that $p \geq 3$ and directly consider spherical errors, that is, $\hat{\boldsymbol{\tau}}_{\text{RCT}} \sim \mathcal{N}(\boldsymbol{\tau}, \sigma_{\text{RCT}}^2 I_p)$ and $\hat{\boldsymbol{\tau}}_{\text{OS}} \sim \mathcal{N}(\boldsymbol{\tau} + \mathbf{b}, \sigma_{\text{OS}}^2 I_p)$. We can then form the vector $\mathbf{D} = \hat{\boldsymbol{\tau}}_{\text{OS}} - \hat{\boldsymbol{\tau}}_{\text{RCT}}$ where $\mathbf{D} \sim \mathcal{N}(\mathbf{b}, (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2) I_p)$. For the bias vector, we extend $\Gamma_1(M)$ to the set of (multivariate) priors given by

$$\Gamma_1^p(M) := \left\{ \pi \in \mathcal{M}_1(\mathbb{R}^p) : \mathbb{E}_\pi[\|\mathbf{b}\|^2] \leq pM \right\}$$

For $\Gamma = \Gamma_1^p(M)$, the robust Bayes estimator in the normal mean problem is then given by

$$h_{\Gamma_1^p(M)}^*(\mathbf{D}) = \frac{M}{M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2} \mathbf{D}.$$

The empirical Bayes approach to adaptation introduced in Section [4.1](#) then extends straightforwardly by estimating the unknown state $M_0 > 0$ generating the data. The maximum marginal likelihood estimator of M_0 can still be considered in this case. However, the extension to higher dimensions opens the way for an alternative estimator based on unbiasedness. The idea is to directly target the constant factor for $\mathbf{D}$. Indeed, we can rewrite $h_{\Gamma_1^p(M)}^*(\mathbf{D}) = (1 - (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)/(M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2))\mathbf{D}$ and denote the shrinkage factor as $\mu_M = (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)/(M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. Since it holds marginally that $\|\mathbf{D}\|^2 \sim (M + \sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)\chi_p^2$, taking the expectation of the inverse yields an unbiased estimator for the shrinkage factor

$$\tilde{\mu}_M = (p - 2) \frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{\|\mathbf{D}\|^2}$$

with implicit estimator $\widetilde{M} = \|\mathbf{D}\|^2/(p - 2) - (\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2)$. By directly plugging $\widetilde{M}$ into $h_{\Gamma_1^p(M)}^*$, we obtain a simple expression for the Bayes estimator based on the empirical least

favorable prior $\pi_{\Gamma_1^p(\widetilde{M})}^*$, that is,

$$h_{\Gamma_1^p(\widetilde{M})}^*(\mathbf{D}) = \left(1 - (p-2) \frac{\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2}{\|\mathbf{D}\|^2}\right) \mathbf{D}.$$

This estimator is exactly the James–Stein estimator for the estimation of normal means. The possibility to obtain a robust Bayesian foundation for the James–Stein estimator was already hinted at in [Feldman \(1991\)](#). From the multivariate extension of Proposition 2.1 stated above, we obtain a simple expression for the robust empirical Bayes estimator over $\{\Gamma_1^p(M) : M \in [0, \infty)\}$ based on $\widetilde{M}$.

Proposition A.4. *The robust empirical Bayes estimator over $\{\Gamma_1^p(M) : M \in [0, \infty)\}$ based on $\widetilde{M}$ is given by*

$$g_{\Gamma_1^p(\widetilde{M})}^*(\widehat{\boldsymbol{\tau}}_{\text{RCT}}, \mathbf{D}) = \widehat{\boldsymbol{\tau}}_{\text{RCT}} + \left(\frac{(p-2)\sigma_{\text{RCT}}^2}{\|\mathbf{D}\|^2}\right) \mathbf{D}.$$

This recovers exactly the combination estimator introduced in [Green and Strawderman \(1991\)](#), which was obtained as a data-driven approximation of the best linear combination of $\widehat{\boldsymbol{\tau}}_{\text{RCT}}$ and $\widehat{\boldsymbol{\tau}}_{\text{OS}}$. We collect this equivalence for reference in the corollary below.

Corollary A.5. *The robust empirical Bayes estimator over $\{\Gamma_1^p(M) : M \in [0, \infty)\}$ based on $\widetilde{M}$ is equal to the estimator introduced in [Green and Strawderman \(1991\)](#).*

This equivalence is of interest on both sides: it gives a robust Bayesian origin to the estimator of [Green and Strawderman \(1991\)](#); it allows us to recover a number of strong properties for $g_{\Gamma_1^p(\widetilde{M})}^*$. We notably inherit a domination result over the RCT estimator $\widehat{\boldsymbol{\tau}}_{\text{RCT}}$. This domination result, which is not possible when $p = 1$, only becomes available because of the multivariate extension to $p \geq 3$ (as typical with James–Stein-type results). If we abandon the purely frequentist perspective and focus on average regret performance, it is then natural to wonder if there exist better estimators of M_0 than $\widetilde{M}$. Remark A.4 below briefly explores the question and hints at improvements over $\widetilde{M}$.

Remark A.4. As a simple alternative to $\widetilde{M}$, it is natural to consider the restriction of $\widetilde{M}$ to its positive part $\widetilde{M}_+ = \max(0, \widetilde{M})$. This was already suggested by [Green and Strawderman \(1991\)](#) based on the result of [Baranchik \(1964\)](#). From our perspective, the restriction is immediately motivated by the restriction $M_0 \geq 0$. It is also immediate from a multivariate extension of

Proposition 2.4 and by noting that the untruncated estimator $\widetilde{M}$ yields non-convex weights. From Baranchik (1964), this restriction is also seen to yield improvements in terms of average regret performance. As another alternative, we can consider the estimator $\hat{M}$ obtained by maximum marginal likelihood. While the adaptive estimator based on $\hat{M}$ does not exhibit the same optimal frequentist properties as the ones based on $\widetilde{M}$ or $\widetilde{M}_+$, it can be shown that it dominates both in terms of average regret for specific values of M_0 and $\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$. These results are of independent interest and are explored at more length in another manuscript – see Delatte (2026b).

A.4 Extension to the misspecification problem

We focused in the main text on a combination problem for independent estimators in a causal inference context. All the results naturally extend beyond the causal inference setting. More importantly, the results extend even if the two estimators we aim to combine are not independent. This generalization to the dependent case allows us to optimally combine estimators based on the same dataset but derived from two nested models. The goal is then to adaptively trade misspecification bias against model efficiency gains – see, for instance, Bickel (1984) and Armstrong, Kline, and Sun (2025). We show below that all our results extend to this setting. We consider the univariate case $p = 1$ for simplicity, but multivariate extensions as in the previous section pose no difficulty. To avoid confusion, we consider new notations and introduce the estimators $\hat{\tau}_0$ and $\hat{\tau}_1$, which are assumed to satisfy

$$\begin{bmatrix} \hat{\tau}_0 \\ \hat{\tau}_1 \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \tau \\ \tau + b \end{bmatrix}, \begin{bmatrix} \sigma_0^2 & \rho\sigma_0\sigma_1 \\ \rho\sigma_0\sigma_1 & \sigma_1^2 \end{bmatrix} \right),$$

where σ_0^2 , σ_1^2 , and ρ are known and satisfy $0 < \sigma_1^2 < \sigma_0^2$ and $\rho \in (-1, 1)$. We again introduce the difference $D := \hat{\tau}_1 - \hat{\tau}_0$, but it now holds that

$$\begin{bmatrix} \hat{\tau}_0 \\ D \end{bmatrix} \sim \mathcal{N} \left(\begin{bmatrix} \tau \\ b \end{bmatrix}, \begin{bmatrix} \sigma_0^2 & \rho\sigma_0\sigma_1 - \sigma_0^2 \\ \rho\sigma_0\sigma_1 - \sigma_0^2 & \sigma_0^2 + \sigma_1^2 - 2\rho\sigma_0\sigma_1 \end{bmatrix} \right).$$

To extend the results in the main text, it suffices to appropriately reweight all the solutions so as to take into account the dependence structure. This can be done by simply propagating

the following weights in all the results. The variance $\sigma_{\text{RCT}}^2 + \sigma_{\text{OS}}^2$ of the difference becomes

$$\text{Var}(D) = \sigma_0^2 + \sigma_1^2 - 2\rho\sigma_0\sigma_1,$$

the regression coefficient w becomes

$$\bar{w} = \frac{\sigma_0^2 - \rho\sigma_0\sigma_1}{\sigma_0^2 + \sigma_1^2 - 2\rho\sigma_0\sigma_1},$$

and the baseline variance $w\sigma_{\text{OS}}^2$ becomes

$$\text{Var}(\hat{\tau}_0 \mid D) = \frac{\sigma_0^2\sigma_1^2(1 - \rho^2)}{\sigma_0^2 + \sigma_1^2 - 2\rho\sigma_0\sigma_1} = \bar{w}\sigma_1^2 + (1 - \bar{w})\rho\sigma_0\sigma_1.$$

From these weights, we directly obtain that, for any estimator $g \in G$ and any $\pi \in \mathcal{M}_1(\mathbb{R})$,

$$B(g, \pi) = \bar{w}\sigma_1^2 + (1 - \bar{w})\rho\sigma_0\sigma_1 + \bar{w}^2 B(h, \pi),$$

and so, for any $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$, the unique robust Bayes estimator of τ over Γ is given by

$$\begin{aligned} g_\Gamma^*(\hat{\tau}_0, D) &= (1 - \bar{w})\hat{\tau}_0 + \bar{w}(\hat{\tau}_1 - h_\Gamma^*(D)) \\ &= W(D)g_{\Gamma_0}^*(\hat{\tau}_0, D) + (1 - W(D))g_{\Gamma_\infty}^*(\hat{\tau}_0, D), \end{aligned}$$

where h_Γ^* is the robust Bayes estimator of b over Γ based on the observation $D \sim \mathcal{N}(b, \sigma_0^2 + \sigma_1^2 - 2\rho\sigma_0\sigma_1)$, $g_{\Gamma_0}^*(\hat{\tau}_0, D) = (1 - \bar{w})\hat{\tau}_0 + \bar{w}\hat{\tau}_1$, $g_{\Gamma_\infty}^*(\hat{\tau}_0, D) = \hat{\tau}_0$, and $W(D) = 1 - h_\Gamma^*(D)/D$. These results are generalizations of Lemma 2.1, Proposition 2.1, and Proposition 2.4. All the subsequent results in the main text then generalize from these new expressions without complications.

Remark A.5. In generalizing the results, some attention should be paid to the regret case where the cancellation of the weight w in the frequentist reduction in Theorem 3.1 for the independent case does not generalize to the dependent case (as seen from the new expression for $B(g, \pi)$). However, a simple expansion from the primitives leads to a similar reduction without difficulty.

Appendix B Approximation of least favorable priors

The solutions in the main text call for the numerical recovery of least favorable priors in the normal mean problem. These approximations are needed to: compute some of the robust Bayes estimators of Section 2.4; compute the minimax risk $R^*([-|b|, |b|])$ and solve the frequentist minimax problem for the minimax regret approach to adaptation in Section 3; compute the approximate robust empirical Bayes estimators in Section 4.3. In this section, we review and provide new solutions to this numerical problem. We start with a generic solution that applies to any set $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$ satisfying some generalized moment conditions. We then provide two specific, iterative solutions for $\Gamma_2(m)$ and $\Gamma_3(\varepsilon)$, respectively. Without loss of generality, we consider the standardized observation $D \sim \mathcal{N}(b, 1)$ where $b \in \mathbb{R}$.

Remark B.1. Under regularity conditions on Γ , it is possible to show that there exists a unique least favorable prior π_Γ^* in the univariate normal mean problem that satisfies a minimax theorem with the (a.e.) unique robust Bayes estimator (which we denoted by h_Γ^* in the main text) – see, for instance, Johnstone (2019) or Delatte (2026a). These results allow us to recover h_Γ^* and its robust Bayes risk $B^*(\Gamma)$ by computing the Bayes estimator with respect to the least favorable prior π_Γ^* . The same result can also be used to compute the frequentist minimax risk $R^*([-m, m])$ in the bounded normal mean problem by virtue of the equivalence $R^*([-m, m]) = B^*(\Gamma_2(m))$ – see Corollary A.2 for the argument. In general, least favorable priors (and hence the corresponding Bayes estimators) do not have a closed form, hence the need for the approximations considered in this section.

B.1 Generic solution under generalized moment conditions

We start with a generic solution to the problem that can be applied to any convex set Γ defined by generalized moment conditions – see Remark B.2 for a definition. This includes in particular the sets $\Gamma_2(m)$, $\Gamma_1(M)$, and $\Gamma_3(\varepsilon)$ – see again Remark B.2. The idea is general, simple, and computationally tractable. Its origin can be traced back to the early developments of the field – see, for instance, Nelson (1966). It has been popularized in econometrics by Chamberlain (2000) for the unconstrained case.

The solution is to approximate Γ using finite mixtures of Dirac measures $\sum_{k=1}^K v_k \delta_{b_k}$ and to find the weights v in the mixture that maximize the expected risk for the respective Bayes estimator. The moment constraints in the definition of Γ can be enforced on the

mixtures, leading to a set of convex constraints on the weights and thus to a convex subset L_K of the simplex $\Sigma_K = \{v \in [0, 1]^K : \sum_{k=1}^K v_k = 1\}$. The maximization program over the convex set L_K is concave in the weights and can be solved by any concave optimization routine under convex constraints. To form an approximation to the least favorable prior, the maximizing weights can then be plugged into the mixture of Dirac measures. This leads to Algorithm [GNM-LFP](#), which is collected below. In Example [B.1](#), we use Algorithm [GNM-LFP](#) to recover the known least favorable prior $\mathcal{N}(0, M)$ in the problem based on $\Gamma_1(M)$.

Remark B.2. A set Γ is said to satisfy generalized moment conditions if Γ can be characterized by a finite number of inequalities $\int_{\mathbb{R}} f(b) d\pi(b) \leq c$ where f is a measurable function and c is a constant. This applies to $\Gamma_2(m)$, $\Gamma_1(M)$, and $\Gamma_3(\varepsilon)$ with $f(b) = \mathbb{1}_{\{|b|>m\}}$, $f(b) = b^2$, and $f(b) = \mathbb{1}_{\{b=0\}}$. These constraints lead to the respective convex polytopes $L_K(m) = \Sigma_K$, $L_K(M) = \{v \in \Sigma_K : \sum v_k b_k^2 \leq M\}$, and $L_K(\varepsilon) = \{v \in \Sigma_K : v_{k_0} \geq \varepsilon\}$ for $b_{k_0} = 0$. For the case $\Gamma_3(\varepsilon)$, the grid G_0 thus needs to include zero. In what follows, we assume that the grid is always chosen so that L_K can be properly defined.

GNM-LFP Algorithm Generic Normal Mean Least Favorable Prior Algorithm

Input Convex set defined by generalized moment conditions: $\Gamma \subseteq \mathcal{M}_1(\mathbb{R})$

Grid: $G_0 = \{b_1, b_2, \dots, b_K\} \subseteq \text{cl}(\bigcup_{\pi \in \Gamma} \text{supp}(\pi))$

Output Approximation of the LFP: $\pi_\Gamma^* = \sum_{k=1}^K v_k^* \delta_{b_k}$

- 1: Define the convex constraints on the simplex $\Sigma_K = \{v \in [0, 1]^K : \sum_{k=1}^K v_k = 1\}$ by forming the convex polytope

$$L_K = \left\{ v \in \Sigma_K : \sum_{k=1}^K v_k \delta_{b_k} \in \Gamma \right\};$$

- 2: For any weight $v = (v_1, v_2, \dots, v_K) \in L_K$, compute the Bayes estimator with respect to the discrete prior $\sum_{k=1}^K v_k \delta_{b_k}$:

$$h(D; v) = \frac{\sum_{k=1}^K v_k b_k \phi(D - b_k)}{\sum_{k=1}^K v_k \phi(D - b_k)};$$

- 3: Maximize the Bayes risk over L_K , that is, solve the concave program:

$$v^* \in \arg \max_{v \in L_K} \sum_{k=1}^K v_k \int_{\mathbb{R}} (h(x; v) - b_k)^2 \phi(x - b_k) dx;$$

- 4: Return the approximation of the least favorable prior $\pi_\Gamma^* = \sum_{k=1}^K v_k^* \delta_{b_k}$.
-

Remark B.3. As far as we know, no general convergence guarantees under grids getting dense in $\text{cl}(\bigcup_{\pi \in \Gamma} \text{supp}(\pi))$ exist for the solution – some guarantees can still be derived from [Noubiap](#)

and Seidel (2001) under compactness or finiteness conditions.

Remark B.4. Under the condition that a unique least favorable prior exists for the frequentist reduction of the regret problems of Section 3.2, a version of Algorithm GNM-LFP can be directly applied to compute a solution. The only modification is a reweighting of the Bayes estimator and the Bayes risk. The structure of the algorithm remains identical. The set L_K in this case is the unrestricted simplex Σ_K . In practice, when computing the approximation, we take a sufficiently large and fine grid (61 points equally spaced across $[-9, 9]$). A similar remark as in Armstrong, Kline, and Sun (2025) holds: the output of the algorithm remains stable for finer and larger grids. However, both the question of the existence of a least favorable prior and the question of the convergence of the algorithm remain open.

Example B.1. We conclude this section by using Algorithm GNM-LFP to approximate the least favorable prior, $\mathcal{N}(0, M)$, in the robust Bayesian decision problem for $\Gamma_1(M)$. We collect the approximated robust Bayes risk in Table 5 below. We also plot in Figure 4 two approximations of the least favorable prior for $M = 2$ using Algorithm GNM-LFP based on a grid with 20 points and a grid with 101 points. We then use the approximation in combination with Algorithm ReBA to approximate the robust empirical Bayes estimator over $\{\Gamma_1(M) : M \in [0, \infty)\}$ we obtained in Proposition 4.1. We compare the estimators by re-using the simulations of Section 5.2. The results are recorded in Table 6 below.


Table 5: GNM-LFP approximation and exact robust Bayes risk for $\Gamma_1(M)$

Moment bound (M)	0	0.5	1	1.5	2	2.5	3	3.5	4	10	20
GNM-LFP $B^*(\Gamma_1(M))$	0.0000	0.3333	0.5000	0.6000	0.6667	0.7143	0.7500	0.7778	0.8000	0.9091	0.9524
Exact $B^*(\Gamma_1(M))$	0.0000	0.3333	0.5000	0.6000	0.6667	0.7143	0.7500	0.7778	0.8000	0.9091	0.9524

Table 6: Monte Carlo – Bias, Variance, and MSE in Design 1

Estimator	LATE (b)	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
GNM-LFP + ReBA	Bias ²	0.0000	0.0027	0.0092	0.0158	0.0207	0.0224	0.0210	0.0186	0.0159	0.0135	0.0116
	Variance	0.0282	0.0318	0.0421	0.0563	0.0732	0.0882	0.0980	0.1043	0.1073	0.1073	0.1062
	MSE	0.0282	0.0345	0.0513	0.0721	0.0939	0.1106	0.1191	0.1229	0.1233	0.1208	0.1178
Moment Plug-in	Bias ²	0.0000	0.0026	0.0092	0.0158	0.0207	0.0224	0.0210	0.0185	0.0158	0.0133	0.0113
	Variance	0.0281	0.0319	0.0420	0.0564	0.0732	0.0880	0.0981	0.1045	0.1076	0.1080	0.1076
	MSE	0.0281	0.0345	0.0512	0.0721	0.0939	0.1105	0.1191	0.1230	0.1234	0.1213	0.1189

Figure 4: Approximation of the LFP for $\Gamma_1(M)$ with $M = 2$ using Algorithm [GNM-LFP](#)



B.2 Improved iterative solution for $\Gamma_2(m)$

In the normal mean problem for $\Gamma_2(m)$, the least favorable prior has no closed form but is known to satisfy some strong regularity constraints that can be leveraged to improve on Algorithm [GNM-LFP](#). The improvements take the form of iterative algorithms leveraging the shape of the least favorable prior as well as the frequentist reduction of the minimax problem – see again Corollary [A.2](#). Solutions of this form have been popularized by [Kempthorne \(1987\)](#). Improvements have been suggested – see, for instance, [Eichenauer and Lehn \(1989\)](#) or [Gourdin, Jaumard, and MacGibbon \(1994\)](#). We provide one encompassing implementation of this iterative solution in Algorithm [BNM-LFP](#) below. We use this algorithm to approximate $\pi_{\Gamma_2(m)}^*$ and $B^*(\Gamma_2(m))$ for different values of m and collect the results in Table [7](#). The values of $B^*(\Gamma_2(m))$ are plotted against the values of $B^*(\Gamma_1(M))$ in Figure [2](#) when $m^2 = M$.

Remark B.5. The algorithm leverages the fact that the least favorable prior under $\Gamma_2(m)$ is discretely supported with a number of support points that grows linearly with m (and so even when $B^*(\Gamma_2(m))$ is close to $B^*(\mathcal{M}_1(\mathbb{R}))$). The strategy is then to build a solution iteratively by adding support points sequentially and to maximize the Bayes risk at each step over both the location points and the weights. The algorithm comes with an explicit stopping criterion derived from minimax inequalities – see Remark [B.6](#). The only defect of this algorithm is that the main program is not globally concave anymore. However, the program is still semi-concave and thus amenable to smooth optimization approaches. Moreover, the iterative nature of the algorithm comes with a massive reduction in the dimensionality of the search space as compared to Algorithm [GNM-LFP](#). The practical implication is that, even for comparably

large values of m , the program is solved almost instantly compared to Algorithm [GNM-LFP](#).

BNM-LFP Algorithm Bounded Normal Mean Least Favorable Prior Algorithm

Input Support bound: $m > 0$

Tolerance for minimax gap: $e > 0$

Output Approximation of the LFP: $\pi_\Gamma^\star = \sum_{k=1}^{K^\star} v_k^\star \delta_{b_k^\star}$

1: Initialize the support size $K = 1$ and set $\text{Gap} = \infty$;

2: **while** $\text{Gap} > e$ **do**

3: For any $b = (b_1, b_2, \dots, b_K) \in [-m, m]^K$ and any $v = (v_1, v_2, \dots, v_K) \in \Sigma_K = \{v \in [0, 1]^K : \sum_{k=1}^K v_k = 1\}$, compute the Bayes estimator with respect to $\pi(v, b) = \sum_{k=1}^K v_k \delta_{b_k}$:

$$h(D; v, b) = \frac{\sum_{k=1}^K v_k b_k \phi(D - b_k)}{\sum_{k=1}^K v_k \phi(D - b_k)};$$

4: Maximize the Bayes risk over Σ_K and $[-m, m]^K$, that is, solve the non-convex program:

$$(v^\star, b^\star) \in \arg \max_{(v, b) \in \Sigma_K \times [-m, m]^K} \sum_{k=1}^K v_k \int_{\mathbb{R}} (h(x; v, b) - b_k)^2 \phi(x - b_k) dx;$$

5: Compute the stopping criterion:

$$\text{Gap} = \max_{b \in [-m, m]} R(h(\cdot; v^\star, b^\star), b) - B(h(\cdot; v^\star, b^\star), \pi(v^\star, b^\star));$$

6: **if** $\text{Gap} \leq e$ **then**

7: **break**

8: **end if**

9: Increment the support size $K = K + 1$;

10: **end while**

11: Return the approximate least favorable prior $\pi_\Gamma^\star = \pi(v^\star, b^\star)$ with $K = K^\star$.

Remark B.6. The stopping criterion comes from the minimax inequalities

$$B(h_\pi, \pi) \leq B^\star(\Gamma_2(m)) = R^\star([-m, m]) \leq \sup_{b \in [-m, m]} R(h_\pi, b)$$

where h_π denotes the Bayes estimator for any prior $\pi \in \Gamma$. The optimization of the frequentist risk to evaluate the stopping criterion in Step 5 is done in practice by discretizing $[-m, m]$.

Remark B.7. Because the least favorable prior for $\Gamma_2(m)$ is symmetric around zero with support always containing $\{-m, m\}$, we can directly enforce these constraints in the choice of the location points and weights in Step 3 of the algorithm. This leads to another significant reduction in the dimensionality of the search space and makes the solutions much more stable.

Remark B.8. [Montiel Olea and Zubova \(2026\)](#) recently revisited the bounded normal mean problem and considered a version of Algorithm [GNM-LFP](#) to approximate the minimax

solutions (maximizing the Bayes risk directly over the simplex in the spirit of [Chamberlain \(2000\)](#)). For a specific optimization routine, the authors obtained convergence guarantees in the weak convergence sense for the approximate least favorable prior. While these convergence results are valuable, Algorithm [BNM-LFP](#) provides faster and more stable solutions in our experiments. The reason is that for reasonable support bounds $m > 0$ the least favorable prior for $\Gamma_2(m)$ has few support points whose locations can be well-recovered by Algorithm [BNM-LFP](#). This suggests that convergence guarantees in the much stronger total variation sense may be obtained for Algorithm [BNM-LFP](#). We leave the question unanswered – one difficulty is that the number of support points for any value of m does not admit a simple closed-form solution (as it requires solving a difficult system of transcendental equations whose solution does not seem to be of much interest beyond this particular problem).

Table 7: Least Favorable Priors for $\mathcal{N}(b, 1)$ with $b \in [-m, m]$

m	$R^*([-m, m])$	supp	Non-negative support points ($b_k \geq 0$)	Weights (v_k)
1	0.4496	2	{ 1 }	{ 0.5 }
2	0.6448	3	{ 0, 2 }	{ 0.42, 0.29 }
3	0.7506	5	{ 0, 1.12, 3 }	{ 0.11, 0.27, 0.18 }
4	0.8151	6	{ 0.60, 2.13, 4 }	{ 0.18, 0.2, 0.12 }
5	0.8573	7	{ 0, 1.46, 3.12, 5 }	{ 0.18, 0.18, 0.15, 0.08 }
6	0.8865	8	{ 0.78, 2.39, 4.10, 6 }	{ 0.17, 0.15, 0.12, 0.06 }
10	0.9445	14	{ 0.58, 1.85, 3.26, 4.77, 6.37, 8.10, 10 }	{ 0.09, 0.1, 0.1, 0.08, 0.07, 0.04, 0.02 }

B.3 Improved iterative solution for $\Gamma_3(\varepsilon)$

The problem of approximating the least favorable prior in the normal mean problem for the Huber contamination class $\Gamma_3(\varepsilon)$ has received much less attention. As far as we know, no specific solution has been devised for this class. In this case, we know that the least favorable prior $\pi_{\Gamma_3(\varepsilon)}^*$ is symmetric and supported on a countably infinite number of points all across the real line. The generic solution provided by Algorithm [GNM-LFP](#) can be applied with an increasing grid. However, for speed and stability, we favor an alternative solution based on computing the least favorable prior in $\Gamma_3(\varepsilon) \cap \Gamma_2(m)$ for increasing values of $m > 0$. Because, in this case, the least favorable prior is also discretely supported with support points growing linearly with m , we can amend Algorithm [BNM-LFP](#) and obtain an algorithm that executes rapidly and yields stable solutions. The only difficulty comes from finding the right stopping criterion for the class $\Gamma_3(\varepsilon) \cap \Gamma_2(m)$. This is handled in Remark [B.9](#) and leads to

Algorithm [HNM-LFP](#), which we collect below. Because the approximation comes both from the truncation $[-m, m]$ and the optimization error, no guarantees for convergence can be provided. The same dimensionality reduction of the search space as stated in Remark [B.7](#) for Algorithm [BNM-LFP](#) can be enforced onto Algorithm [HNM-LFP](#). We apply the algorithm to approximate $B^*(\Gamma_3(\varepsilon))$ by computing $B^*(\Gamma_3(\varepsilon) \cap \Gamma_2(m))$ for different values of $\varepsilon \in (0, 1)$ and $m = 20$. The results are recorded in Table [8](#). In Figure [5](#), we plot these values as a function of ε and record the shape of the approximations for the least favorable priors.

HNM-LFP Algorithm Huber's Normal Mean Least Favorable Prior Algorithm

Input Support bound: $m > 0$

Contamination level: $\varepsilon \in (0, 1)$

Tolerance for minimax gap: $e > 0$

Output Approximation of the LFP: $\pi_\Gamma^* = \varepsilon\delta_0 + (1 - \varepsilon) \sum_{k=1}^{K^*} v_k^* \delta_{b_k^*}$

1: Initialize the support size $K = 1$ and set $\text{Gap} = \infty$;

2: **while** $\text{Gap} > e$ **do**

3: For any $b = (b_1, \dots, b_K) \in [-m, m]^K$ and any $v = (v_1, \dots, v_K) \in \Sigma_K$, define the prior $\pi(v, b) = \varepsilon\delta_0 + (1 - \varepsilon) \sum_{k=1}^K v_k \delta_{b_k}$, and compute the Bayes estimator with respect to $\pi(v, b)$:

$$h(D; v, b) = \frac{(1 - \varepsilon) \sum_{k=1}^K v_k b_k \phi(D - b_k)}{\varepsilon \phi(D) + (1 - \varepsilon) \sum_{k=1}^K v_k \phi(D - b_k)};$$

4: Maximize the Bayes risk over Σ_K and $[-m, m]^K$, that is, solve the non-convex program:

$$(v^*, b^*) \in \arg \max_{(v, b) \in \Sigma_K \times [-m, m]^K} \left[\varepsilon \int_{\mathbb{R}} h(x; v, b)^2 \phi(x) dx + (1 - \varepsilon) \sum_{k=1}^K v_k \int_{\mathbb{R}} (h(x; v, b) - b_k)^2 \phi(x - b_k) dx \right];$$

5: Compute the stopping criterion:

$$\text{Gap} = \varepsilon R(h(\cdot; v^*, b^*), 0) + (1 - \varepsilon) \max_{b \in [-m, m]} R(h(\cdot; v^*, b^*), b) - B(h(\cdot; v^*, b^*), \pi(v^*, b^*));$$

6: **if** $\text{Gap} \leq e$ **then**

7: **break**

8: **end if**

9: Increment the support size $K = K + 1$;

10: **end while**

11: Return the approximate least favorable prior $\pi_\Gamma^* = \pi(v^*, b^*)$ with $K = K^*$.

Remark B.9. To compute a least favorable prior in the class $\Gamma_3(\varepsilon) \cap \Gamma_2(m)$ iteratively, we need to find the right stopping criterion. By standard minimax arguments, it holds that

$$B(h_\pi, \pi) \leq B^*(\Gamma_3(\varepsilon) \cap \Gamma_2(m)) \leq \varepsilon R(h_\pi, 0) + (1 - \varepsilon) \max_{b \in [-m, m]} R(h_\pi, b),$$


where h_π is the Bayes estimator for any prior $\pi \in \Gamma_3(\varepsilon) \cap \Gamma_2(m)$. We can then use the difference $\text{Gap} = \varepsilon R(h_\pi, 0) + (1 - \varepsilon) \max_{b \in [-m, m]} R(h_\pi, b) - B(h_\pi, \pi)$ as our new stopping

criterion. From the minimax inequalities above, the magnitude Gap provides direct guarantees for the approximate least favorable prior and its Bayes estimator in terms of approximating the robust Bayes risk.

Table 8: Approximation of the robust Bayes risk $B^*(\Gamma_3(\varepsilon))$ using Algorithm [HNM-LFP](#)

ε	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
$B^*(\Gamma_3(\varepsilon))$	0.9812	0.9636	0.9273	0.8752	0.8071	0.7215	0.6148	0.4802	0.3025

Figure 5: Robust Bayes risk and LFP for $\Gamma_3(\varepsilon)$ using Algorithm [HNM-LFP](#)



Appendix C Supplementary Monte Carlo simulations

C.1 Design 2: Selection into treatment on unobservables

We consider for the second design an observational study with omitted variable bias due to the fact that the units select into treatment based on some unobservable that cannot be controlled for using the observables. Selection on unobservables is another canonical issue that plagues observational causal inference – see [Rosenbaum and Rubin \(1983\)](#). The average treatment effect (ATE) cannot be identified in the observational study; only the integrated difference in expectations conditional on the observable controls (IDE) can. We take the ATE as the causal effect of interest $\tau \in \mathbb{R}$ and consider distributional assumptions so that $\tau = 0$. The bias in this case is caused by the unobservable confounder (i.e., omitted variable bias) and is entirely captured by the difference between the IDE and the ATE.

We recall that N and n denote the sample sizes of the RCT and the observational study, respectively. The simulation design is based on the following distributional assumptions. For any $i \in \{1, \dots, N, N+1, \dots, N+n\}$, the outcome Y_i is determined by the vector $(X_i, U_i, \varepsilon_i)$ where

$$\begin{aligned} X_i &\sim \mathcal{N}(0, 1), \\ U_i &\sim \text{Bernoulli}(0.5), \\ \varepsilon_i &\sim \mathcal{N}(0, 1), \end{aligned}$$

and

$$Y_i = Y_i(0) = Y_i(1) = X_i + 2bU_i + \varepsilon_i.$$

We assume that X_i is observable while U_i and ε_i are unobservable. In this design, the potential outcomes are unaffected by the treatment so that the individual and average treatment effects are all equal to 0. In the RCT, we assume that the treatment is independent and for any $i \in \{1, \dots, N\}$,

$$T_i \sim \text{Bernoulli}(0.5).$$

In the observational study, we assume that the treatment is independent of X_i but depends

on the unobservable confounder U_i . That is, for any $i \in \{N + 1, \dots, N + n\}$,

$$T_i \mid U_i \sim \text{Bernoulli}(0.25 + 0.5U_i).$$

The randomization of the treatment in the RCT ensures that the ATE is identified. That is, for any $i \in \{1, \dots, N\}$,

$$\tau = \mathbb{E}[Y_i(1) - Y_i(0)] = \mathbb{E}[Y_i|T_i = 1] - \mathbb{E}[Y_i|T_i = 0] = 0.$$

In the observational study, we have that $Y_i(0), Y_i(1) \not\perp\!\!\!\perp T_i \mid X_i$ and so we can only recover the integrated difference in expectations conditional on the observed characteristics, that is, for any $i \in \{N + 1, \dots, N + n\}$,

$$\mathbb{E} \left[\mathbb{E}[Y_i|T_i = 1, X_i] - \mathbb{E}[Y_i|T_i = 0, X_i] \right] = b.$$


We estimate the ATE in the RCT using the difference-in-means estimator $\hat{\tau}_{\text{RCT}}$ and its variance using the Neyman variance estimator $\hat{\sigma}_{\text{RCT}}^2$ as in Section 5.2. For the estimation of the IDE in the observational study, we can simply run a linear regression of Y on T and X and use the OLS coefficient for T as $\hat{\tau}_{\text{OS}}$. For the variance estimator $\hat{\sigma}_{\text{OS}}^2$ of $\hat{\tau}_{\text{OS}}$, we use the HC2 variant of the Eicker–Huber–White heteroskedasticity-robust variance estimator. Since the design is completely homoskedastic by construction, we could consider more efficient variance estimators, but the simplicity of the design makes the choice of a variance estimator practically immaterial. The results of the simulations for this design are collected in Table 9 and Figure 6. Additional visualizations are provided in Figure 8 and Figure 10.

C.2 Monte Carlo results for Design 2 (selection on unobservables)

Table 9: Monte Carlo – Bias, Variance, and MSE of the estimators in Design 2

	Bias (b)	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	1
Panel A: Benchmark Estimators												
RCT	Bias ²	0.0002	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
	Variance	0.0388	0.0390	0.0395	0.0404	0.0418	0.0435	0.0456	0.0481	0.0510	0.0543	0.0580
	MSE	0.0390	0.0391	0.0396	0.0406	0.0419	0.0436	0.0457	0.0482	0.0511	0.0544	0.0581
MLE	Bias ²	0.0000	0.0099	0.0397	0.0893	0.1587	0.2478	0.3568	0.4854	0.6339	0.8020	0.9899
	Variance	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0002	0.0003	0.0003	0.0003	0.0003
	MSE	0.0002	0.0101	0.0399	0.0895	0.1589	0.2481	0.3570	0.4857	0.6341	0.8023	0.9902
Pretesting	Bias ²	0.0001	0.0037	0.0114	0.0157	0.0128	0.0078	0.0033	0.0009	0.0002	0.0000	0.0000
	Variance	0.0097	0.0197	0.0395	0.0670	0.0904	0.1022	0.0993	0.0872	0.0789	0.0692	0.0633
	MSE	0.0098	0.0234	0.0509	0.0827	0.1032	0.1100	0.1026	0.0881	0.0792	0.0692	0.0634
Panel B: Minimax Regret Estimators												
Moment Regret	Bias ²	0.0000	0.0080	0.0293	0.0569	0.0822	0.1001	0.1099	0.1145	0.1182	0.1241	0.1341
	Variance	0.0007	0.0020	0.0061	0.0144	0.0272	0.0417	0.0529	0.0590	0.0601	0.0584	0.0564
	MSE	0.0007	0.0099	0.0355	0.0713	0.1093	0.1417	0.1629	0.1735	0.1783	0.1824	0.1905
Support Regret	Bias ²	0.0000	0.0077	0.0283	0.0556	0.0821	0.1042	0.1207	0.1330	0.1435	0.1535	0.1638
	Variance	0.0010	0.0022	0.0055	0.0116	0.0199	0.0294	0.0379	0.0448	0.0502	0.0547	0.0588
	MSE	0.0010	0.0098	0.0338	0.0672	0.1020	0.1336	0.1587	0.1778	0.1936	0.2081	0.2226
Panel C: Robust Empirical Bayes Estimators												
Moment Plug-in	Bias ²	0.0001	0.0016	0.0056	0.0086	0.0091	0.0077	0.0061	0.0048	0.0040	0.0034	0.0031
	Variance	0.0123	0.0175	0.0286	0.0419	0.0527	0.0582	0.0602	0.0612	0.0620	0.0638	0.0665
	MSE	0.0124	0.0191	0.0343	0.0505	0.0618	0.0660	0.0663	0.0660	0.0660	0.0672	0.0696
Support Plug-in	Bias ²	0.0001	0.0011	0.0041	0.0072	0.0077	0.0073	0.0062	0.0053	0.0045	0.0039	0.0035
	Variance	0.0145	0.0190	0.0280	0.0390	0.0478	0.0536	0.0569	0.0599	0.0625	0.0649	0.0679
	MSE	0.0146	0.0201	0.0321	0.0462	0.0555	0.0609	0.0631	0.0651	0.0670	0.0688	0.0713
Huber Plug-in	Bias ²	0.0001	0.0028	0.0093	0.0136	0.0127	0.0102	0.0066	0.0046	0.0033	0.0025	0.0020
	Variance	0.0103	0.0171	0.0323	0.0505	0.0651	0.0729	0.0712	0.0694	0.0676	0.0672	0.0699
	MSE	0.0103	0.0199	0.0416	0.0642	0.0778	0.0831	0.0778	0.0741	0.0709	0.0697	0.0719

Figure 6: Monte Carlo – MSE of the estimators in Design 2



C.3 Supplementary visualizations for the Monte Carlo simulations

Figure 7: Monte Carlo – MSE of the estimators in Design 1




Figure 8: Monte Carlo – MSE of the estimators in Design 2




Figure 9: Close-up – Empirical distribution and MSE decomposition in Design 1




Figure 10: Close-up – Empirical distribution and MSE decomposition in Design 2

