

Approximate Minimax Estimation of a Bounded Normal Mean via Stochastic Mirror Ascent^{*}

José Luis Montiel Olea[†] Ekaterina Zubova[‡]

June 4, 2026

Abstract

This paper presents a computational approach to find an approximately minimax estimator for the classical *bounded normal mean* problem. The suggested procedure is the Bayes estimator corresponding to an approximately least-favorable distribution obtained from a stochastic mirror ascent routine for concave maximization. The paper shows that both the approximately least-favorable distribution and the approximately minimax estimator are indeed close (in a sense we make precise) to their desired targets. Simulation evidence suggests that the approximately minimax estimator can yield—with a reasonable amount of compute—risk improvements from 6% to almost 18% relative to the minimax linear estimator (which is known to admit a maximal improvement of 20%). The approximately minimax estimator is then applied to the problem of how to best aggregate the information contained in local projections and vector autoregressions to estimate an impulse response coefficient.

^{*}We would like to thank Tim Armstrong and Paul Delatte for a very helpful discussion that sparked our interest in this problem. We would also like to thank Jeffrey Negrea and Alexander Terenin for useful exchanges regarding different computational aspects of the bounded normal mean problem and several suggestions for future work. We would also like to thank Andres Aradillas-Fernandez, José Blanchet, Chen Qiu, Francesca Molinari, Mikkel Plagborg-Møller, Jörg Stoye, Lezhi Tan, and Christian Wolf. Montiel Olea gratefully acknowledges financial support by the National Science Foundation Grant SES-2315600.

[†]Department of Economics, Cornell University. Email: montiel.olea@gmail.com

[‡]Department of Economics, Cornell University. Email: ez268@cornell.edu

1 Introduction

We propose a numerical algorithm to find an approximately minimax estimator (under quadratic loss) of the mean of a standard normal distribution, when such mean is known to belong to the interval $[-m, m]$. [Casella and Strawderman \[1981\]](#) refer to this classical problem in statistical decision theory as the *bounded normal mean* estimation problem. As noted by [Johnstone and MacGibbon \[1992\]](#), the minimax risk in this problem plays an important role in comparing the behavior of linear to nonlinear estimators in an important class of nonparametric estimation problems; see [Ibragimov and Has'minskii \[1984\]](#) and [Donoho, Liu, and MacGibbon \[1990\]](#). The minimax risk of the bounded normal mean problem has also been used recently by [Armstrong, Kline, and Sun \[2025\]](#) as an input to construct estimators that *adapt to misspecification*.

Our suggested algorithm finds an approximately minimax estimator by first obtaining an approximately *least-favorable distribution*. This strategy is common in the literature; see [Kempthorne \[1987\]](#) and [Chamberlain \[2000\]](#). One point of departure relative to previous work is that we use a *stochastic mirror ascent routine* (with negative entropy as a mirror map) to approximately solve the following concave optimization problem: we maximize Bayes risk over the space of probability distributions supported on a fixed (equally-spaced) grid of I points over the interval $[-m, m]$. In what follows, we let Δ^{I-1} denote the space of such distributions, without making explicit reference to the support points.

Before providing more specific details about our contribution, it is worth noting that our numerical strategy has two sources of approximation error, both of which are explicitly taken into account in our analysis. The first source of error is an obvious *discretization error*, which comes from the fact that we optimize Bayes risk over the simplex Δ^{I-1} and not over *all* probability distributions supported on $[-m, m]$. We show that we can control this discretization error thanks to a special (Lipschitz) continuity property of the risk function (see [Lemma 1](#) and [Lemma 2](#)). The second source of error is what we term *optimization*

error, which comes from the fact that the mirror ascent routine for concave maximization is designed to find an *approximate solution*, and not an *exact* one. An explicit analysis of this type of error is common in theoretical computer science, operations research, and optimization where it is typical “*to relax the requirement of finding an optimal solution, and instead settle for a solution that is good enough*” [Williamson and Shmoys, 2011, p. 14]. We can control this optimization error thanks to the rich literature that has analyzed the theoretical properties of mirror descent.

The methods of mirror descent (or ascent) [Nemirovski and Yudin, 1983, Chapter 3] are a family of iterative procedures recommended in the optimization literature for approximately solving convex problems of high dimensions. These methods are iterative, first-order optimization algorithms, in that they typically require repeated evaluations of the objective function and an estimator of its (sub or super) gradient, but do not exploit any further smoothness information about the objective function.¹ The optimization error associated with these algorithms comes from the fact that the procedure is designed to stop after finitely many iterations.

There are at least four reasons that explain why mirror ascent was the algorithm we chose for finding an approximately least-favorable distribution. *First*, even in the simple bounded normal mean problem, it is difficult to get an explicit characterization of higher-order derivatives of the objective function defining the least-favorable distribution, which means that first-order methods (which only use information concerning the gradient and do not require further smoothness assumptions) are particularly attractive. *Second*, there is also a sense in which our algorithm is essentially the best first-order iterative algorithm for approximately minimizing (maximizing) a convex (concave), Lipschitz function over the simplex (as in the problem defining the approximately least-favorable distribution); see Proposition 4.2 in Ben-Tal, Margalit, and Nemirovski [2001].² *Third*, as our results show, the supergradient of the

¹The algorithm we recommend only requires repeated evaluations of an estimator of the supergradient of the objective function, but not of the objective function itself. See Algorithm 1.

²It is also known that the rate of convergence of mirror descent (or ascent) improves over regular subgradient descent for convex minimization (maximization) problems in the probability simplex [Bubeck, 2015,

objective function at any distribution $\pi \in \Delta^{I-1}$ depends on the risk function $R(d_\pi, \theta)$, where d_π is the Bayes estimator associated with π . This means that even though evaluating the supergradient is potentially costly (as it requires numerical integration), we can rely on the well-known stochastic version of mirror ascent that replaces the supergradient by an unbiased estimator (and in our case, this unbiased estimator can be constructed using only *one* Monte Carlo draw per grid point per iteration, further reducing the computational burden of the algorithm). *Fourth*, since there is a rich theory regarding the convergence analysis of stochastic mirror descent (ascent) [Bubeck, 2015, Chapter 6], we can provide explicit theoretical guarantees about the performance of our suggested algorithm.

The first theoretical result in this paper (Theorem 1) shows that our stochastic mirror ascent algorithm *provably* finds an approximately least-favorable distribution for the bounded normal mean problem after *finitely* many iterations. The theoretical guarantees we provide are expressed in terms of a scalar parameter $\epsilon > 0$. This parameter controls the overall approximation error of our numerical algorithm. Let $\lceil \cdot \rceil$ denote the ceiling function.³ The theorem states that if the number of iterations, $J(\epsilon)$, and the number of points in the grid, $I(\epsilon)$, are chosen to be

$$J(\epsilon) := \left\lceil \frac{2M^2 \log(I(\epsilon))}{(\epsilon/2)^2} \right\rceil, \quad I(\epsilon) := \left\lceil 1 + \frac{3m^3 + 4m^2}{\epsilon/2} \right\rceil, \quad \text{where } M := 4m^2,$$

then, for any $\alpha \in (0, 1)$, with probability at least $1 - \alpha$, the Bayes risk associated with the prior distribution π^ϵ found by the algorithm is close to the *maximin* value of the bounded normal mean problem. And, in particular, it differs at most by an additive factor

$$\epsilon \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}} \right).$$

With the exception of ϵ , our theorem makes explicit recommendations about all of the parameters in Section 4.3].

³The ceiling function returns the smallest integer that is greater than or equal to a given number.

rameters needed to run the algorithm (such as, the initial condition and the learning rate). It is worthwhile to make three remarks about our result. *First*, the probabilistic statement that appears in the theorem comes from the fact that our theoretical analysis acknowledges the fact that the gradient of the objective function needs to be estimated by Monte-Carlo integration. This estimation induces—just as any other stochastic approximation method—an additional source of variation that will make our approximately least-favorable distribution change for different *runs* of the algorithm if the random seed is not fixed. *Second*, the total number of iterations (which scales logarithmically in the number of grid points) should be viewed as *sufficient* to guarantee a good performance of the algorithm, but it is by no means *necessary*. In fact, our simulation results show that the approximation error is likely to be of size ϵ for a considerably smaller number of iterations (which is important, given that $J(\epsilon)$ is already in the order of millions for values of m as small as $m = 3$). *Third*, the convergence guarantee we provide in Theorem 1 is expressed in terms of how well we can approximate the *maximin* value of the bounded normal mean problem, and not in terms of how close the approximate least-favorable distribution is relative to the true least-favorable distribution.⁴

The second theoretical result of this paper (Theorem 2) complements Theorem 1 by showing that π^ϵ will indeed be close (in terms of the *1-Wasserstein distance*) to the least-favorable distribution of the bounded normal mean problem, provided ϵ is small enough. To the best of our knowledge, none of the other existing papers in the literature that have suggested algorithms for the bounded normal mean problem—such as [Feldman and Brown \[1989\]](#), [Donoho et al. \[1990\]](#), and [Gourdin, Jaumard, and MacGibbon \[1994\]](#)—have derived results like ours. This is also true about other existing papers that have suggested algorithms for finding approximate solutions for more general minimax/maximin problems—such as [Kempthorne \[1987\]](#), [Johnstone and MacGibbon \[1992\]](#), [Chamberlain \[2000\]](#), [Aradillas Fernández, Blanchet, Montiel Olea, Qiu, Stoye, and Tan \[2024\]](#); [Aradillas Fernández, Blanchet, Montiel Olea, Qiu, Stoye, and Tan \[2025\]](#), [Elliott, Müller, and Watson \[2015\]](#), and [Guggen-](#)

⁴The least-favorable distribution is known to exist and be unique; see Proposition 4.19 in [Johnstone \[2019\]](#). See also [Delatte \[2026\]](#).

berger and Huang [2025].

The third theoretical result of this paper (Theorem 3) shows that there is a sense in which the Bayes estimator associated with the prior distribution π^ϵ found by the stochastic mirror ascent algorithm is approximately minimax. The notion of approximate minimaxity that we use in this paper extends the classical definition of *minimax sequence* given in the seminal work of Ghosh [1964], p. 1037. Broadly speaking, Ghosh [1964] defines a deterministic sequence of decision rules to be a *minimax sequence* if their worst-case risk converges to the minimax risk of the problem of interest. Since, as we have explained before, our stochastic mirror ascent algorithm introduces noise beyond the statistical uncertainty contained in the estimation problem, we need to allow for the possibility of a stochastic sequence of decision rules that approximate the minimax value of interest. In light of this observation, we define a sequence of decision rules to be a *stochastic minimax sequence* if their worst-case risk converges in probability to the minimax value of interest. This means that the *tail* of a stochastic minimax sequence contains decision rules that are ϵ -minimax—in the sense of Ferguson [1967], Chapter 1, Definition 4, p. 33—with probability at least $1 - \delta$, for any $\epsilon, \delta > 0$.

The last theoretical result of the paper (Theorem 4) shows that the Bayes estimator associated with the prior distribution π^ϵ approximates the *exact* minimax estimator of a bounded normal mean in a stronger sense: for any compact set K of data realizations, there is a small enough ϵ (that could depend on the set) such that the approximate and the exact minimax estimators are *uniformly* close to each other over K , with probability approaching one. This means that our approximate minimax estimator converges—*uniformly over compact sets*—to the minimax estimator of the bounded normal mean problem.

SIMULATIONS: It is known that in the bounded normal mean problem the worst-case risk of the best minimax *linear* estimator is at most 25% above the best, unconstrained, minimax estimator. This is the so-called Ibragimov and Has'minskii [1984] constant, see Theorem 4.7 in Johnstone [2019]. The minimax linear estimator thus provides a computationally feasible

(and trivial to compute benchmark) for the bounded normal mean problem. The value of the [Ibragimov and Has'minskii \[1984\]](#) constant implies that the largest reduction in risk that we could expect to see from our approximately minimax estimator (relative to the minimax linear estimator) is 20%. [Figures 3 and 4](#) show that our stochastic mirror ascent routine generates risk improvements over the minimax linear estimators that range from about 6% to 18%, as m ranges from 1 to 4 (with the largest improvement occurring at $m = 1.6$). In our simulations we choose ϵ to be $\epsilon(m) := (1/5)m^2/(1+m^2)$. We provide a rationale for this choice in [Section 4](#).

In terms of the computational effort required to achieve the 6% to 18% improvement reported above, we note that the largest reduction in risk that we see takes about 45 seconds on a personal computer. The computational time increases with m , and implementing our algorithm for $m = 3$ takes about 2 hours (see [Figure 2](#)). The main bottleneck we face in the implementation of our algorithm is the number of iterations recommended in [Theorem 1](#). We report simulations where we bound the *wall-clock* time of our procedure in 5 and 15 minutes. With the 15 minutes limit, we still see improvements between 6% and 18%, but the improvement over linear estimators when $m = 3$ drops from about 11% without constraints, to almost 6%.

APPLICATION: A common problem in macroeconometrics is the estimation of the dynamic causal effects—or *impulse response functions*—of aggregate shocks to macroeconomic policies or economic fundamentals. The two most common empirical methods currently used by applied macroeconomists for doing so are *structural vector autoregressions* (henceforth interchangeably referred to as VARs or SVARs), going back to the seminal work of [Sims \[1980\]](#), and the so-called *local projections* (LPs), introduced by [Jordà \[2005\]](#). It is now understood that there is a sense in which VARs and LPs share the same estimand [[Plagborg-Møller and Wolf, 2021](#); [Xu, 2023](#)], but that in finite samples these procedures lie on opposite ends of a bias-variance trade-off [[Li, Plagborg-Møller, and Wolf, 2024](#)]: LP estimators have lower bias than VAR estimators, but they also have substantially higher variance at intermediate and

long horizons.

The question of interest is how to best aggregate the information in LP and VAR estimators to estimate a particular impulse response coefficient. We follow closely the work of [Armstrong et al. \[2025\]](#) and consider the problem of constructing a minimax estimator based on quadratic loss, and a bivariate Gaussian model for the VAR and LP estimators.⁵ We impose an a priori bound on the bias of the VAR. As we explain in Section 5, both the bivariate normality and the bound on the bias can be motivated using the *locally misspecified* vector autoregression model used in [Montiel Olea, Plagborg-Møller, Qian, and Wolf \[2026\]](#). In particular, the bound on the bias can be obtained from restricting the fraction of the residual variance that is explained by the lagged shocks in the misspecified part of the VAR model; see [Montiel Olea, Qian, Wolf, and Plagborg-Møller \[2025\]](#), p. 16.

Under a convenient invariance assumption, the minimax invariant estimator is a *debiased* version of the VAR estimator. The optimal debiasing function depends on the data only through the standardized difference between the VAR and LP estimators, which can be interpreted as a [Hausman \[1978\]](#) statistic for misspecification and which we denote as Δ . As shown in [Armstrong et al. \[2025\]](#), the optimal debiasing function turns out to correspond to the minimax estimator of a bounded normal mean problem, where we use Δ to try to estimate the magnitude of the bias of the VAR estimator.

We report our estimator for four influential macroeconomic papers, where we allow the fraction of the residual variance that is explained by the lagged shocks in the misspecified part of the VAR model to be at most 1%. The four papers we consider are [Gertler and Karadi \[2015\]](#), focusing on the impulse response of log industrial production to a monetary policy shock; [Romer and Romer \[2010\]](#), focusing on the impulse response of log of real GDP per capita to a tax shock; [Ramey \[2011\]](#), focusing on the impulse response of log of real GDP per capita to a government spending shock (based on military news); and [Francis, Owyang, Roush, and DiCecio \[2014\]](#), focusing on the impulse response of log of

⁵In the bivariate Gaussian model, the LP estimator is unbiased but has potentially high variance and the VAR estimator is biased but with smaller variance than LP; see Equation (18).

labor productivity to a technology shock. Even though the amount of misspecification that we allow is relatively small (which should favor the VAR estimator), in all applications the approximately minimax estimator is very close to the LP estimator. This is also consistent with the minimax linear estimator reported by [Montiel Olea et al. \[2026\]](#) (Corollary 4.2), which under our assumptions, places about 70% of the weight over the LP estimator. In our applications, for medium and long horizons the *shrinkage* towards the LP estimator is a bit more aggressive than the one implied by the linear estimator (and the gains in excess risk relative to an oracle that knows the bias are about 15%).

The remainder of the paper is organized as follows. Section 2 introduces the bounded normal mean problem and its minimax formulation. Section 3 presents the algorithm and the main theoretical results. Section 4 reports simulation results. Section 5 applies the method to estimation of impulse responses.

2 The Bounded Normal Mean Problem

Fix $m > 0$ and consider the statistical model

$$Y \sim \mathcal{N}(\theta, 1), \quad \theta \in [-m, m].$$

The problem of interest is to estimate θ under the loss function:

$$\mathcal{L}(a, \theta) = (a - \theta)^2.$$

[Casella and Strawderman \[1981\]](#) refer to this problem as the *bounded normal mean estimation problem*. An excellent and detailed treatment of the problem is given in Chapter 4.6 of [Johnstone \[2019\]](#).

A large literature in statistics has studied the theoretical properties of different estimators for the bounded normal mean problem; see, for example, [Bickel \[1981\]](#), [Levit \[1981\]](#), [Donoho](#)

et al. [1990], and the references therein. As we will explain below, an important difference of our paper relative to existing literature is that we make it our goal to suggest a concrete numerical strategy—based on the methods of *stochastic mirror descent (ascent)*—to obtain an *approximate* minimax estimator (in a sense we make precise). The main theoretical results in our paper will be focused on showing that our numerical strategy indeed achieves the desired objective.

We note that we have deliberately set the variance of the outcome Y to equal one. As we discuss later, this is without loss of generality as the bounded normal mean problem treats the variance as known. See also the scale-invariance relation in Equation 4.35, Chapter 4.6 of Johnstone [2019]. It is also without loss of generality to restrict the action space to be $[-m, m]$. The reason is that, under squared-error loss, truncating any real-valued estimator to $[-m, m]$ weakly reduces risk pointwise in θ .⁶

Let \mathcal{D} be the class of all measurable decision rules $d : \mathbb{R} \rightarrow [-m, m]$. As usual, the *risk* of a rule d is

$$R(d, \theta) := \mathbb{E}_\theta[(d(Y) - \theta)^2],$$

where the expectation is taken assuming $Y \sim \mathcal{N}(\theta, 1)$. The first result of this paper establishes an important property of the risk function that will be exploited by our numerical strategy; namely, that $R(d, \cdot)$ is *Lipschitz continuous* in θ for any $d \in \mathcal{D}$.

Lemma 1 (The Risk Function is Lipschitz in θ) For any $d \in \mathcal{D}$ and any $\theta, \theta' \in [-m, m]$

$$|R(d, \theta) - R(d, \theta')| \leq (3m^2 + 4m) |\theta - \theta'|. \quad (1)$$

Proof. See Appendix A.1. ■

This first lemma shows that the Lipschitz constant is independent of d , but dependent on the bound m . To the best of our knowledge this result has not appeared previously in

⁶One benchmark is the projection of Y onto $[-m, m]$, $d_{\text{trunc}}(y) = \min\{\max\{y, -m\}, m\}$, which is measurable, valued in $[-m, m]$, and, by truncation, weakly dominates the unconstrained estimator $d(y) = y$ under squared loss as $(d_{\text{trunc}}(y) - \theta)^2 \leq (y - \theta)^2$ for all $\theta \in \Theta$ pointwise.

the literature. The key argument we use to establish this result is an upper bound on the total variation distance between two normal distributions with unit variance; see Theorem 1.3 in [Devroye, Mehrabian, and Reddad \[2018\]](#).⁷

2.1 Approximate Minimax Estimation in the Bounded Normal Mean Problem

Given $m > 0$, we say that an estimator d^* is minimax—or *m-minimax* adopting the terminology of [Armstrong et al. \[2025\]](#)—if it satisfies

$$\sup_{\theta \in [-m, m]} R(d^*, \theta) = \bar{v}^*(m) := \inf_{d \in \mathcal{D}} \sup_{\theta \in [-m, m]} R(d, \theta). \quad (2)$$

We refer to $\bar{v}^*(m)$ as the *minimax value* of the bounded normal mean estimation problem associated with the bound $m > 0$.

It is known that when the bound m is sufficiently small, there is a closed-form solution for the problem in (2); see Theorems 3.1 and 4.2 in [Casella and Strawderman \[1981\]](#). More generally, the problem of finding a minimax estimator can be viewed as an infinite dimensional convex programming problem; see p. 1419 and 1420 in [Donoho et al. \[1990\]](#). This means that finding a minimax estimator must necessarily involve some type of numerical approximation.

To be explicit about the fact that any given numerical strategy will only approximately solve the minimax problem in (2), we build on the seminal work of [Ghosh \[1964\]](#) and define a *stochastic minimax sequence* as follows. Let $\{P_n\}_{n=1}^\infty$ be a sequence of probability distributions over the space \mathcal{D} and let $\{d_n^*\}_{n=1}^\infty$ be a stochastic sequence of decision rules in which each decision rule d_n^* is drawn from P_n .

⁷For normal distributions with common covariance, Theorem 1 in [Barsov and Ulyanov \[1987\]](#) implies the sharper bound:

$$\frac{1}{2} \int_{-\infty}^{\infty} |\phi(y; \theta) - \phi(y; \theta')| dy \leq \frac{|\theta - \theta'|}{\sqrt{2\pi}}.$$

Using this bound in [Appendix A.1](#) would give the sharper Lipschitz constant equal to $3\sqrt{\frac{2}{\pi}}m^2 + 4m$.

Definition 1 (Stochastic Minimax Sequence) *A stochastic sequence of decision rules $\{d_n^*\}_{n=1}^\infty$ is said to be a stochastic minimax sequence for the bounded normal mean problem if for any $\varepsilon > 0$*

$$P_n \left(\sup_{\theta \in [-m, m]} R(d_n^*, \theta) \leq \bar{v}^*(m) + \varepsilon \right) \rightarrow 1, \quad (3)$$

as $n \rightarrow \infty$.

Definition 1 extends the classical definition of *minimax sequence* given in Ghosh [1964], p. 1037. Ghosh [1964] defines a deterministic sequence of decision rules $\{d_n\}_{n=1}^\infty \subseteq \mathcal{D}$ to be a minimax sequence if

$$\sup_{\theta \in [-m, m]} R(d_n, \theta) \rightarrow \bar{v}^*(m). \quad (4)$$

Since the numerical procedure that we will recommend (which is akin to stochastic gradient descent for convex minimization problems) introduces noise beyond the statistical uncertainty contained in the estimation problem, we need to allow for the possibility that d_n is stochastic. To this end, Definition 1 replaces the “ \rightarrow ” relation in (4) by convergence in probability, where convergence in probability is assessed relative to the distributions P_n .

Definition 1 is also related to the notion of ε -minimax decision rules, which is a common way of formalizing the notion of an approximate minimax solution in statistical decision problems; see Ferguson [1967], Chapter 1, Definition 4, p. 33.⁸ Note that if $\{d_n^*\}_{n=1}^\infty$ is a stochastic minimax sequence, then for every pair $\varepsilon, \delta > 0$ there exists $n(\varepsilon, \delta)$ such that whenever $n \geq n(\varepsilon, \delta)$

$$P_n \left(\bar{v}^*(m) \leq \sup_{\theta \in [-m, m]} R(d_n^*, \theta) \leq \bar{v}^*(m) + \varepsilon \right) \geq 1 - \delta.$$

This means that the *tail* of a stochastic minimax sequence contains decision rules that are ε -minimax with probability at least $1 - \delta$.

⁸See also Aradillas Fernandez et al. [2024] for recent applications and an algorithmic approach for finding ε -minimax decision rules in a class of statistical decision problems.

2.2 Approximate Least-Favorable Distribution for a Discretized Version of the Bounded Normal Mean Problem

Our strategy for constructing a stochastic minimax sequence for the problem in (2) consists of finding a sequence of priors that approximate the problem's *least-favorable distribution*. We do this by discretizing the parameter space. The strategy of discretizing the parameter space to find an approximate least-favorable distribution has been recommended before in the econometrics literature by Chamberlain [2000] and it has been used recently by Armstrong et al. [2025]. To the best of our knowledge, neither paper provides theoretical results that guarantee that such a strategy indeed succeeds in finding an approximate minimax rule (or even an approximate least-favorable distribution). The closest result that we are aware of is Lemma 1 of Guggenberger and Huang [2025], which provides a set of sufficient conditions that guarantee that the value of a discretized version of a minimax problem converges to the minimax value of the original problem.

For the sake of exposition, this section presents a quick overview of the typical relation between a minimax decision rule and a least-favorable distribution. We also present the discretized bounded normal mean problem and the definition of an approximate least-favorable distribution for this problem.

Let τ denote the standard (subspace) topology over $[-m, m]$. It is known that the topological space $([-m, m], \tau)$ is a Polish space.⁹ In order to define probability distributions over $[-m, m]$, we endow this space with the σ -algebra generated by τ ; i.e., its Borel σ -algebra. Let $\Delta([-m, m])$ denote the set of all (Borel) probability measures over $[-m, m]$.

Let π be an element of $\Delta([-m, m])$. The *Bayes risk* of d with respect to π is defined as

$$r(d, \pi) := \mathbb{E}_\pi[\mathbb{E}_\theta(d(Y) - \theta)^2] = \int_{[-m, m]} R(d, \theta) d\pi.$$

⁹This follows from the fact that τ is metrizable by the absolute value metric, and the space $([-m, m], |\cdot|)$ is a complete and separable metric space.

The rule d_π is said to be a Bayes rule with respect to π if it satisfies

$$r(d_\pi, \pi) = \inf_{d \in \mathcal{D}} r(d, \pi). \quad (5)$$

Define a *least-favorable distribution* as a distribution π^* that satisfies

$$\inf_{d \in \mathcal{D}} r(d, \pi^*) = \underline{v}^*(m) := \sup_{\pi \in \Delta([-m, m])} \inf_{d \in \mathcal{D}} r(d, \pi). \quad (6)$$

In Appendix C we show that the maximin problem that defines the least-favorable distribution can be viewed as a concave maximization problem over the infinite-dimensional space $\Delta([-m, m])$. The objective function in this problem—which corresponds to the function $f(\pi) = \inf_{d \in \mathcal{D}} r(d, \pi)$ —is Lipschitz continuous in π , with respect to the 1-Wasserstein metric over $\Delta([-m, m])$.

It is known that

$$\bar{v}^*(m) = \underline{v}^*(m). \quad (7)$$

See Proposition 4.19 in [Johnstone \[2019\]](#) and also their Theorem A.5.¹⁰ It is also known that if (7) holds, any minimax rule d^* must be a Bayes rule with respect to the least-favorable distribution π^* .

Discretized Maximin Problem: As mentioned before, our numerical strategy to construct a stochastic minimax sequence starts by discretizing the parameter space $\Theta := [-m, m]$ using an equally spaced grid of $I > 1$ points $\{\theta_1, \dots, \theta_I\}$ in the interval $[-m, m]$. This means that the i -th point in the grid equals:

$$\theta_i := -m + \frac{2m(i-1)}{I-1}.$$

¹⁰There are two additional interesting properties that we will reference later. First, the least-favorable distribution is unique. Second, the least-favorable distribution is supported on finitely many points (and all support points have the same risk). See, again, Proposition 4.19 in [Johnstone \[2019\]](#).

The gap between any consecutive points θ_{i+1}, θ_i is

$$\Delta(m, I) := \theta_{i+1} - \theta_i = \frac{2m}{I-1}. \quad (8)$$

Define the $(I-1)$ -dimensional simplex $\Delta^{I-1} := \{\pi \in [0, 1]^I : \sum_i \pi_i = 1\}$ and consider then the *discretized* version of the maximin problem in (6)

$$\underline{v}(m; I) := \max_{\pi \in \Delta^{I-1}} \inf_{d \in \mathcal{D}} r(d, \pi), \quad \text{where} \quad r(d, \pi) := \sum_{i=1}^I \pi_i R(d, \theta_i). \quad (9)$$

There is an evident gap between the value of the original maximin problem in (6)—which we denoted as $\underline{v}^*(m)$ —and the value of the discretized maximin problem in (9)—which we are denoting as $\underline{v}(m; I)$. Obviously, $\underline{v}(m; I) \leq \underline{v}^*(m)$, since the outer maximum in the discretized maximin problem is a subset of $\Delta([-m, m])$. The second result of this paper establishes a concrete quantitative relation between the maximin values of the bounded normal mean problem and its discretized version.

Lemma 2 (Distance between the original and discretized maximin problems) *For any $m > 0$ and $I > 1$*

$$\underline{v}(m; I) \leq \underline{v}^*(m) \leq \underline{v}(m; I) + (3m^3 + 4m^2)/(I-1).$$

Proof. See Appendix A.2. ■

The key argument to establish Lemma 2 is the Lipschitz continuity of the risk function (which was established in Lemma 1) and the minimax theorem applied to both the discretized and the original problem. As expected, the gap between the discretized maximin problem and the original problem decreases as a function of I , and that larger values of m will require a finer grid to attain a good approximation. We note that Lemma 2 readily yields a formula to select I in order to guarantee that the gap between the two maximin problems is small. For example, to guarantee a gap of at most $\epsilon/2$ it suffices to set the number of points in the

uniform grid equal to

$$I(\epsilon) := \left\lceil 1 + \frac{3m^3 + 4m^2}{\epsilon/2} \right\rceil. \quad (10)$$

The function $\lceil \cdot \rceil$ is the ceiling function: the function that returns the smallest integer that is greater than or equal to a given number. In light of Lemma 2 our numerical strategy will focus on obtaining an approximate solution to the discretized maximin problem. The following definition is entirely analogous to the notion of ϵ -minimax decision rules—see Ferguson [1967], Chapter 1, Definition 4, p. 33—but applied to the maximin problem.

Definition 2 (ϵ -least-favorable distribution for the discretized maximin problem)

We say that $\pi^{(\epsilon, I)} \in \Delta^{I-1}$ is an ϵ -least-favorable distribution for the discretized maximin problem in (9) if

$$\underline{v}(m; I) - \epsilon \leq \inf_{d \in \mathcal{D}} r(d, \pi^{(\epsilon, I)}) \leq \underline{v}(m; I).$$

We say that a distribution $\pi^I \in \Delta^{I-1}$ is an approximately least-favorable distribution for the discretized maximin problem in (9), if there exists $0 < \epsilon < 4m^2$ for which π^I is an ϵ -least-favorable distribution.

Let $I(\epsilon)$ be defined as in (10). Note that an important implication of Lemma 2 is that if $\pi^{(\epsilon/2, I(\epsilon))}$ is an $(\epsilon/2)$ -least-favorable distribution for the discretized problem with $I(\epsilon)$ equally spaced points, then $\pi^{(\epsilon/2, I(\epsilon))}$ is an ϵ -least-favorable distribution of the *original* bounded normal mean problem. Our strategy to construct a stochastic minimax sequence for the bounded normal problem relies on this observation. More concretely, we will consider a sequence $\{(\epsilon_n, I_n)\}_{n=1}^{\infty}$, where $\epsilon_n \rightarrow 0$ and $I_n \rightarrow \infty$. We will show that we can build an associated stochastic sequence of priors $\pi^n \equiv \pi^{(\epsilon_n, I_n)}$ that are guaranteed to be—with high probability—approximately the least-favorable distribution for the original problem. In the next section we will show how to obtain such a stochastic sequence of priors. Our main result will be that the sequence of Bayes decision rules corresponding to the sequence of priors $\{\pi^n\}_{n=1}^{\infty}$ is a stochastic minimax sequence.

Note also that without an upper bound on ϵ , any distribution $\pi \in \Delta^{I-1}$ is approximately least-favorable. The risk function in the bounded normal mean problem is bounded by $4m^2$, so we are only interested in considering approximation errors below this upper bound.

3 Main Results

This section presents our main results. First, we show that a *stochastic mirror ascent* algorithm for concave maximization over the $(I(\epsilon) - 1)$ -dimensional simplex can be used to provably obtain—with high probability—an approximate least-favorable distribution for the maximin problem in (6). Second, we show that the Bayes rules associated with the sequence of approximately least-favorable priors obtained via stochastic mirror ascent constitute a stochastic minimax sequence.

3.1 Stochastic Mirror Ascent Algorithm for the Discretized Maximin Problem

For any $\pi \in \Delta([-m, m])$, define the function $f : \Delta([-m, m]) \rightarrow \mathbb{R}_+$ given by

$$f(\pi) := \inf_{d \in \mathcal{D}} r(d, \pi) = \inf_{d \in \mathcal{D}} \int_{[-m, m]} R(d, \theta) d\pi.$$

The first key observation here is that the function $f(\cdot)$ is concave in π over $\Delta([-m, m])$, and also over Δ^{I-1} for any $I > 1$. This result has been established more generally before; see, for example, Chamberlain [2000], Equation 5, p. 630, and the discussion therein. For the sake of completeness, we formalize this observation in Lemma 7 in Appendix C. Lemma 7 also shows that for any prior $\pi \in \Delta^{I-1}$, the vector $g(\pi) := (R(d_\pi, \theta_1), \dots, R(d_\pi, \theta_I))$ is a supergradient of $f(\cdot)$ at π (where d_π is the Bayes rule corresponding to π). Lastly, since $|R(d, \theta)| \leq 4m^2$ for all $(d, \theta) \in \mathcal{D} \times \Theta$, the function $f : \Delta^{I-1} \rightarrow \mathbb{R}$ is Lipschitz (with respect to L_1 -norm) with Lipschitz constant $M := 4m^2$.

A popular algorithmic approach for *approximately* solving concave maximization problems—in particular, those of high dimensions—is to use the methods of (stochastic) mirror descent of Nemirovski and Yudin [1983]. The methods of mirror descent are a family of iterative procedures designed for finding an *approximate solution* of convex minimization (or concave maximization) problems in high dimensions. These methods require repeated (but finitely many) evaluations of the subgradient of the objective function or a noisy but unbiased estimator of the subgradient. These methods exploit the geometry of the optimization domain, but they do not exploit any higher-order smoothness information about the optimization problem. Mirror descent is known to outperform gradient descent in some optimization domains, such as the probability simplex; see Section 4.3 of Bubeck [2015]. See also Section 5.1.1 in Nemirovski, Juditsky, Lan, and Shapiro [2009].

We consider the iterative algorithm described below for maximizing $f(\cdot)$ over the simplex Δ^{I-1} . The stochastic component of this algorithm comes from the evaluation of the supergradient. The pseudo-code that describes Algorithm 1 below takes the grid $\{\theta_i\}_{i=1}^I$ as given. Another input is the *step size* (denoted as η) which controls the adjustment of the prior at each epoch. We also take the number of epochs (J) as given.

The next theorem shows that—with high probability—the stochastic mirror ascent routine outputs an approximate least-favorable distribution for the bounded normal mean problem.

Theorem 1 (Approximately least-favorable distribution) *Fix $\epsilon > 0$ and $m > 0$. Let π^ϵ be the output of Algorithm 1 with inputs*

$$\eta(\epsilon) := \frac{\epsilon/2}{M^2}, \quad J(\epsilon) := \left\lceil \frac{2M^2 \log(I(\epsilon))}{(\epsilon/2)^2} \right\rceil, \quad I(\epsilon) := \left\lceil 1 + \frac{3m^3 + 4m^2}{\epsilon/2} \right\rceil, \quad (11)$$

where $M := 4m^2$. Then, for any $\alpha \in (0, 1)$, with probability at least $1 - \alpha$ we have

$$\inf_{d \in \mathcal{D}} r(d, \pi^\epsilon) \geq \underline{v}^*(m) - \epsilon \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}} \right). \quad (12)$$

Algorithm 1 Stochastic Mirror Ascent to Find Approximate Least-Favorable Prior

- 1: **Input:** Step-size $\eta > 0$; number of epochs $J \in \mathbb{N}$; equally-spaced grid $\{\theta_i\}_{i=1}^I \subset \Theta$.
- 2: Initialize $w_0 \in \mathbb{R}^I$ by setting $w_{i,0} = 1$ for all $i \in \{1, \dots, I\}$.
- 3: **for** $j = 1, 2, \dots, J$ **do**
- 4: Compute $\phi_j := \sum_{i=1}^I w_{i,j-1}$.
- 5: For each $i \in \{1, \dots, I\}$, compute

$$\pi_{i,j} := \frac{w_{i,j-1}}{\phi_j}.$$

- 6: Compute the Bayes decision rule d_{π_j} under prior π_j ¹¹:

$$d_{\pi_j}(y) = E_{\pi_j}[\theta \mid y] = \frac{\sum_{i=1}^I \theta_i \pi_{i,j} \phi(y; \theta_i)}{\sum_{i=1}^I \pi_{i,j} \phi(y; \theta_i)},$$

where $\phi(y; \theta_i)$ is the p.d.f. of $\mathcal{N}(\theta_i, 1)$.

- 7: For each $i \in \{1, \dots, I\}$, draw an independent $y_{i,j} \sim \mathcal{N}(\theta_i, 1)$ and compute the stochastic supergradient estimate

$$\tilde{g}_{i,j} = (d_{\pi_j}(y_{i,j}) - \theta_i)^2.$$

- 8: Multiplicative weights update: for each $i \in \{1, \dots, I\}$,

$$w_{i,j} := w_{i,j-1} \cdot \exp(\eta \tilde{g}_{i,j}).$$

9: **end for**

- 10: **Output:** $\frac{1}{J} \sum_{j=1}^J \pi_j = \left(\frac{1}{J} \sum_{j=1}^J \pi_{1j}, \dots, \frac{1}{J} \sum_{j=1}^J \pi_{Ij} \right)$.
-

Proof. See Appendix A.3. ■

Theorem 1 shows that our suggested stochastic mirror ascent algorithm—with the choice of tuning parameters in (11)—succeeds in finding an approximate least-favorable distribution for the maximin problem (6). Since Algorithm 1 relies on an estimator of the supergradient (see Step 7), it is important to quantify this uncertainty in our theoretical guarantees. The theorem provides an explicit description of the effects of the stochastic gradient on the optimization procedure. In particular, the theorem says that with probability at least $1 - \alpha$, the Bayes risk of π^ϵ is at most within the additive factor

$$\epsilon \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}} \right)$$

of the maximin value of interest. Note that this term is decreasing in ϵ . However, a smaller value of ϵ increases the computational complexity of the procedure because it increases the dimension of the concave optimization problem via its effect on $I(\epsilon)$.

To complement Theorem 1 we also show that there is a sense in which π^ϵ is close to the (unique) least-favorable distribution of the original maximin problem, which we have denoted as π^* . To show this, take an arbitrary sequence $\{\epsilon_n\}_{n=1}^\infty \subseteq \mathbb{R}_{++}$ such that $\epsilon_n \rightarrow 0$. For each ϵ_n in the sequence, we obtain the corresponding approximate least-favorable prior $\pi_n := \pi^{\epsilon_n}$ as described in Theorem 1. We show that π_n converges (in 1-Wasserstein distance, as defined below) to the least-favorable distribution π^* .

The 1-Wasserstein Distance. For arbitrary Borel probability distributions $\pi, \pi' \in \Delta([-m, m])$ define $\Gamma(\pi, \pi')$ the set of all couplings of (π, π') . Define the 1-Wasserstein distance (with cost function $|\cdot|$) as

$$W(\pi, \pi') := \inf_{\gamma \in \Gamma(\pi, \pi')} \int_{[-m, m]^2} |\theta - \theta'| d\gamma.$$

It is known that $W(\pi, \pi')$ defines a metric over all Borel probability measures over $\Delta([-m, m])$; see Theorem 7.3 in Villani [2021]. Moreover, it is common to refer to $W(\cdot, \cdot)$ as the *Kantorovich-Rubinstein* distance. It is also known that the Kantorovich-Rubinstein duality formula implies that

$$W(\pi, \pi') = \sup_{\varphi} \left(\int_{[-m, m]} \varphi d\pi - \int_{[-m, m]} \varphi d\pi' \right) \quad \text{s.t.} \quad \|\varphi\|_{\text{Lip}} \leq 1, \quad (13)$$

where

$$\|\varphi\|_{\text{Lip}} := \sup_{\theta \neq \theta'} \frac{|\varphi(\theta) - \varphi(\theta')|}{|\theta - \theta'|}.$$

See Remark 7.5, p. 207 in Villani [2021].

Theorem 2 (π_n converges to π^* in 1-Wasserstein Distance) *Fix $m > 0$. Take an arbitrary sequence $\{\epsilon_n\}_{n=1}^\infty \subseteq \mathbb{R}_{++}$ such that $\epsilon_n \rightarrow 0$. Let $\pi_n := \pi^{\epsilon_n}$ be the output of Algorithm 1 with the parameters specified in Theorem 1. Let π^* be the unique least-favorable prior of*

the bounded normal mean problem. Then, for any $\xi > 0$

$$P_n(W(\pi_n, \pi^*) > \xi) \rightarrow 0, \text{ as } n \rightarrow \infty.$$

Proof. See Appendix A.4. ■

3.2 Stochastic Minimax Sequence for the Bounded Normal Mean Problem

Theorem 1 and 2 have shown that the stochastic mirror ascent algorithm can be used to generate an approximate least-favorable prior. In this section, we use these results to construct a stochastic minimax sequence.

Theorem 3 Fix $m > 0$. Take an arbitrary sequence $\{\epsilon_n\}_{n=1}^\infty \subseteq \mathbb{R}_{++}$ such that $\epsilon_n \rightarrow 0$. Let $\pi_n := \pi^{\epsilon_n}$ be the output of Algorithm 1 with the parameters specified in Theorem 1. Let d_n denote the Bayes rule associated with π_n . The sequence $\{d_n\}_{n=1}^\infty$ is a stochastic minimax sequence.

Proof. See Appendix A.5. ■

Theorem 4 Fix $m > 0$ and let $S \subset \mathbb{R}$ be any compact subset. Under the same conditions as Theorem 3, for any $\xi > 0$

$$P_n\left(\sup_{y \in S} |d_n(y) - d_{\pi^*}(y)| > \xi\right) \rightarrow 0, \text{ as } n \rightarrow \infty.$$

Proof. See Appendix A.6. ■

Theorems 3 and 4 translate prior convergence into estimator convergence. As the approximate least-favorable priors converge, the induced Bayes rules converge to the Bayes rule under the least-favorable distribution. Since that rule is minimax, the rules generated by the algorithm form a stochastic minimax sequence. This is the least-favorable-prior logic behind

numerical minimax procedures such as in [Chamberlain \[2000\]](#), with stochastic convergence reflecting additional randomness from Monte Carlo estimation of supergradients.

4 Simulations

This section explores the numerical performance of the stochastic mirror ascent algorithm in the bounded normal mean problem. We compare the risk of the approximately minimax estimator to the linear minimax benchmark. We also assess the computational cost of the suggested procedure.

4.1 Selection of the parameter ϵ

As discussed in the introduction—with the exception of ϵ —our [Theorem 1](#) makes explicit recommendations about all of the parameters that need to be specified in order to run the stochastic mirror ascent algorithm. We now argue that in the bounded normal mean problem, there is a simple way to choose ϵ (as a function of m).

Note first that the minimax risk over linear estimators in the bounded normal mean problem is

$$\bar{v}_{\text{linear}}(m) := \frac{m^2}{1 + m^2}, \tag{14}$$

see Equation 4.28 in [Johnstone \[2019\]](#). It is also known that

$$\bar{v}^*(m) \leq \bar{v}_{\text{linear}}(m) \leq 1.25\bar{v}^*(m), \tag{15}$$

where the upper bound is due to [Donoho et al. \[1990\]](#); see also Theorem 4.17 in [Johnstone \[2019\]](#).¹²

If ϵ were chosen to satisfy $\epsilon > 0.25\bar{v}^*(m)$, then the minimax linear estimator would be ϵ -minimax: $\bar{v}_{\text{linear}}(m) \leq 1.25\bar{v}^*(m) < \bar{v}^*(m) + \epsilon$. We are interested in values of ϵ associated

¹²The factor 1.25 is typically referred to as the [Ibragimov and Has'minskii \[1984\]](#) constant.

with smaller approximation errors (to avoid having provable ϵ -minimax guarantees for the minimax linear estimator). Thus, using (15) we recommend choosing ϵ as

$$\epsilon(m) := \frac{1}{5} \frac{m^2}{1+m^2} = \frac{0.25}{1.25} \bar{v}_{\text{linear}}(m) \leq 0.25 \bar{v}^*(m). \quad (16)$$

Figure 1 displays both the required number of iterations of the stochastic mirror ascent algorithm and the implied grid size when we choose $\epsilon(m)$ as in (16) and $m \in [0.1, 9]$. The figure illustrates that both the number of iterations and the grid size grow (polynomially) with m . The dominant cost is the number of iterations, since $J(m) := J(\epsilon(m))$ grows as $m^4 \log(m)$ while the number of grid points $I(m) := I(\epsilon(m))$ as m^3 . For example, at $m = 1$, $I(m) = 141$ and $J(m) \approx 63,345$ (yielding a total time of about 3 seconds for the algorithm to run); while at $m = 3$, $I(m) = 1,301$ and $J(m) \approx 2.3$ million (which makes the runtime of the algorithm to be about 2 hours).




Figure 1: Number of iterations and grid size as functions of m for $m \in [0.1, 9]$.

Figure 2 illustrates the total time required to run the stochastic mirror ascent routine and evaluate the risk with 1,000 MC draws and the values of $\epsilon(m)$ for $m \in \{1, \dots, 3\}$, implied by the formula in (16).




Figure 2: Computation time and $\epsilon(m)$.

4.2 Improvement of the approximately minimax estimator relative to the minimax linear estimator

Figure 3 compares the risk of the minimax estimator obtained via stochastic mirror ascent and the linear minimax estimator for $m \in \{1, \dots, 4\}$. The upper bound (gray dashed line) is $\bar{v}_{\text{linear}}(m) = m^2/(1 + m^2)$ (gray line), the worst-case risk of the minimax linear estimator. The lower bound (orange dashed line) is $0.8\bar{v}_{\text{linear}}(m)$, obtained from the [Ibragimov and Has'minskii \[1984\]](#) constant. The difference between this upper and lower bound equals $(1/5)m^2/(1 + m^2)$, which is our choice of $\epsilon(m)$.

We obtain the worst-case risk of the estimator generated by the stochastic mirror ascent algorithm (red line, with circle markers) in two steps. In the first step, we evaluate the risk function of the approximately minimax estimator over a finite grid of 200,000 equally-spaced points in the interval $[-m, m]$. The risk function is evaluated by Monte Carlo simulation using 1,000 draws.¹³ In the second step, we acknowledge the fact that the worst-case risk

¹³We simulate $N = 1,000$ independent draws $Z_j \sim N(0, 1)$ and compute, for each grid point θ_i ,

$$\hat{R}_N(d, \theta_i) = \frac{1}{N} \sum_{j=1}^N [d(\theta_i + Z_j) - \theta_i]^2. \quad (17)$$

evaluated on the finite grid understates the estimator’s worst-case risk over the parameter space (since the worst-case risk may be attained between grid points). In Appendix E.2 we show how the Lipschitz continuity of the risk function in Lemma 1 can be used to translate the evaluation of the risk over a finite grid into a *certifiable* upper bound for the worst-case risk of the approximately minimax estimator—and, consequently, for the problem’s minimax value. The red line (with circle markers) in Figure 3 is such an upper bound.

The blue line (with circle markers) is the average risk—computed according to the approximate least-favorable distribution—of the estimator generated by the stochastic mirror ascent algorithm. Since, by construction, the approximately minimax estimator is a Bayes decision rule with respect to the approximately least-favorable prior, the blue line (with circle markers) is a legitimate lower bound for both the problem’s minimax value and for the estimator’s worst-case risk.

Since the worst-case risk of d_{π^ϵ} —the estimator generated by the stochastic mirror ascent algorithm—lies in between the red and blue lines (both with circle markers), we conclude that d_{π^ϵ} is ϵ -minimax optimal for a value of ϵ considerably smaller (the difference between the lines with circle markers) than the target $\epsilon(m)$ (the difference between the dashed lines). Figure 3 is thus graphical evidence of the success of our stochastic mirror ascent algorithm.




Figure 3: Bounds on the Worst-case risk of d_{π^ϵ} .

To further illustrate the success of our stochastic mirror ascent algorithm, Figure 4 reports a lower bound on the percentage improvement in the worst-case risk of d_{π^ϵ} relative to the worst-case risk of the minimax linear estimator. That is, we would like to report the risk improvement computed as

$$100 \times \left(1 - \frac{\sup_{\theta \in [-m, m]} R(d_{\pi^\epsilon}, \theta)}{\bar{v}_{\text{linear}}(m)} \right),$$

but we replace $\sup_{\theta \in [-m, m]} R(d_{\pi^\epsilon}, \theta)$ by the certified upper bound reported in Figure 3.




Figure 4: Lower Bound on Percentage improvement in worst-case risk of d_{π^ϵ} over the worst-case risk of the minimax linear estimator.

For comparison, in Appendix E.2 we also report the worst-case risk of two alternative estimators: truncated maximum likelihood estimator $d_{\text{trunc}}(y) = \max\{\min\{y, m\}, -m\}$ and *clipped* linear estimator (purple line) $d_{\text{linear, clipped}} = \max\{\min\{c \cdot y, m\}, -m\}$, where $c = m^2/(1 + m^2)$.

4.3 Comparison of d_{π^ϵ} and the minimax linear estimator

Figure 5 plots the approximately minimax estimator and the minimax linear estimator for $m \in \{1, 2, 4\}$. Note that for $m = 2$ and $m = 4$, the most visible differences between d_{π^ϵ} and the minimax linear estimator concentrate near the boundaries.




Figure 5: Comparison of different estimators for the bounded normal mean problem.

4.4 Simulations with a time limit

Suppose that we are only willing to run the algorithm for at most x minutes. That is, we will run the algorithm either until the required number of iterations is reached or the time budget is exhausted, whichever happens first. These types of simulations are typically referred to as *fixed wall-clock time* simulations.¹⁴

Figure 6 displays the number of iterations completed within a five-minute limit for each $m \in \{1, \dots, 4\}$, versus the theoretical target implied by $\epsilon(m) = \frac{1}{5} \frac{m^2}{1+m^2}$. As illustrated by this figure, the five-minute time limit starts binding after $m = 2$.

Figure 7 displays different risk values: analytical bounds (\bar{v} and $\frac{1}{1.25}\bar{v}$), and the grid-evaluated worst-case risk of the stochastic mirror ascent algorithm.

Figure 8 illustrates the percentage improvements in risk, analogous to Figure 4, but under time restrictions: a five-minute limit (solid red line) and a fifteen-minute limit (dashed red line). Even with the five-minute limit binding for $m > 2.2$, the proposed algorithm still delivers a meaningful improvement in the worst-case risk up to $m \approx 3$. Beyond this point, the five-minute constraint is no longer sufficient to certify any improvement, and the SMA-based decision rule is dominated by the linear minimax rule. Increasing the limit to fifteen

¹⁴See Section 7 of Blanchet, Kang, Olea, Nguyen, and Zhang [2023] for an example of numerical simulations that also fix the wall-clock time.




Figure 6: Required iterations vs completed within five-minute limit for $m \in \{1, \dots, 4\}$.




Figure 7: Risk of d_{π^ϵ} vs linear minimax (five-minute limit).

minutes pushes the break-even point out to $m \approx 3.3$ and recovers additional risk improvement throughout the upper range of m . However, at large m the problem becomes hard enough that this time is not sufficient to beat the linear rule.




Figure 8: Risk improvement of proposed estimator vs linear (five- and fifteen-minute limit)

5 Application

5.1 Motivation: Estimation of Impulse Response Coefficients

A common problem in macroeconometrics is the estimation of the dynamic causal effects—or *impulse response functions*—of aggregate shocks to macroeconomic policies or economic fundamentals. The two most common empirical methods currently used by applied macroeconomists for doing so are *structural vector autoregressions* (henceforth interchangeably referred to as VARs or SVARs), going back to the seminal work of Sims [1980], and the so-called *local projections* (LPs), introduced by Jordà [2005].

Understanding the properties of these two estimation procedures has been an active area of recent research; see the references and discussion in Montiel Olea and Plagborg-Møller [2021], Montiel Olea et al. [2025], and Jordà and Taylor [2025]. It is now understood that there is a sense in which VARs and LPs share the same estimand [Plagborg-Møller and Wolf, 2021; Xu, 2023], but that in finite samples these procedures lie on opposite ends of a bias-variance trade-off [Li et al., 2024]: LP estimators have lower bias than VAR estimators, but they also have substantially higher variance at intermediate and long horizons.

Let θ_h denote the impulse response coefficient of interest at horizon h . For the sake of

notation, we omit any explicit reference to the horizon of interest in the estimators and the covariance matrix (although both of these clearly depend on h). We consider a set-up where an applied macroeconomist has access to both LP and VAR estimators of θ_h . We assume that these estimators follow the simple parametric model

$$\begin{pmatrix} \hat{\theta}^{\text{LP}} \\ \hat{\theta}^{\text{VAR}} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \theta_h \\ \theta_h + b \end{pmatrix}, \Sigma \right), \quad \Sigma \equiv \begin{pmatrix} \sigma_{\text{LP}}^2 & \rho \cdot \sigma_{\text{LP}} \sigma_{\text{VAR}} \\ \rho \cdot \sigma_{\text{LP}} \sigma_{\text{VAR}} & \sigma_{\text{VAR}}^2 \end{pmatrix}. \quad (18)$$

We treat the matrix Σ as known, and assume that the variance of the LP estimator (σ_{LP}^2) is strictly larger than the variance of the VAR estimator (σ_{VAR}^2). The unknown parameter b is used to represent the potential bias of the VAR estimator.

While the exact bivariate normality of the estimators is unlikely to hold in finite samples, the asymptotic normality of the estimators can be established by using the *locally misspecified* vector autoregression model used in Montiel Olea et al. [2026].¹⁵ In particular, their Proposition 3.1 provides an asymptotic expansion that shows that the large-sample distribution of the LP estimator is unaffected by local misspecification (and so, there is a sense in which LPs are asymptotically unbiased). In contrast, their Proposition 3.2 shows that local misspecification typically leads to asymptotic bias in the VAR estimator. This bias is captured by the scalar parameter b in (18), which will depend on the magnitude of the local misspecification. When the local misspecification is of order $O_P(1/\sqrt{T})$, a martingale central limit theorem can then be applied to show that the right-hand side of (18) provides a valid asymptotic approximation to the finite-sample distribution of the VAR and LP estimators under the locally misspecified VAR model.

We make two additional remarks about the structure of the bivariate normal distribution in (18). First, when the bias parameter b equals zero, the estimator $\alpha \hat{\theta}^{\text{LP}} + (1 - \alpha) \hat{\theta}^{\text{VAR}}$ is also unbiased for the parameter of interest θ_h (this holds for any positive or negative

¹⁵Montiel Olea et al. [2026] consider stationary data generating processes that are well approximated by a finite-order SVAR model, but subject to local misspecification in the form of an asymptotically vanishing moving average process of potentially infinite order.

$\alpha \in \mathbb{R}$). Since we want to treat the VAR estimator as the efficient estimator in the absence of potential bias, this means that the linear estimator with the smallest possible variance must coincide with the VAR estimator. This is enforced by setting $\rho = \sigma_{\text{VAR}}/\sigma_{\text{LP}}$ in (18), which implies that the covariance between the two estimators is σ_{VAR}^2 . We note this result can also be established more generally under the locally misspecified vector autoregression model used in Montiel Olea et al. [2026]; see their Corollary A.2.

Second, throughout the rest of this section, we will assume that there is a known bound on the bias parameter b that enters the model (18). The motivation for this bound is the following. Let T denote the sample size used to produce the VAR estimator. Montiel Olea et al. [2026] showed that if the researcher is willing to place an upper bound on the fraction of the variance of the moving average residual that is explained by the moving average component of their locally misspecified VAR (denote this parameter by \mathcal{M}), then the worst-case bias (in absolute value) for the VAR estimator takes the following form:

$$\underbrace{\sqrt{T \times \mathcal{M}}}_{:=m} \sigma_{\Delta}, \quad \text{where} \quad \sigma_{\Delta} := \sqrt{\sigma_{\text{LP}}^2 - \sigma_{\text{VAR}}^2}, \quad (19)$$

See Proposition 4.1 of Montiel Olea et al. [2026] and also Equation 3.3 in Montiel Olea et al. [2025].

5.2 Minimax combinations of VAR and LP estimators

The question of interest is how to best aggregate the information in LP and VAR estimators to estimate the impulse response coefficient θ_h . This question has been considered recently by Nemtyrev and Boldea [2026] and Chen, Pesavento, and Vonnák [2026]. More generally, the question of how to best aggregate the information in two estimators—one of which is potentially biased but more efficient—has been analyzed by Armstrong et al. [2025] and Lin, Bickel, and Ding [2026].

We follow closely the work of Armstrong et al. [2025] and consider the problem of con-

structuring a minimax estimator (or decision rule) based on quadratic loss, and the bivariate Gaussian model in (18). We want to impose the bound (19) on the bias parameter $|b| \leq m \cdot \sigma_\Delta$.

It is without loss of generality to focus on estimators $d_h(\hat{\theta}^{LP}, \Delta)$ that depend on the VAR and LP estimator only through the statistics

$$\begin{pmatrix} \hat{\theta}^{LP} \\ \Delta \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \theta_h \\ b/\sigma_\Delta \end{pmatrix}, \begin{pmatrix} \sigma_{LP}^2 & -\sigma_\Delta \\ -\sigma_\Delta & 1 \end{pmatrix} \right), \quad \Delta \equiv (\hat{\theta}_{VAR} - \hat{\theta}_{LP}) / \sigma_\Delta,$$

where the absolute value of Δ can be viewed as a [Hausman \[1978\]](#) test of correct specification of the VAR model that compares the VAR and LP impulse response estimates. A test of this kind was proposed by [Stock and Watson \[2018\]](#) for testing the invertibility of a vector autoregressive model. It will be convenient to introduce an auxiliary random variable $V \sim \mathcal{N}(0, \sigma_{LP}^2 - \sigma_\Delta^2)$ assumed to be independent of Δ . We can use this random variable to construct a simple representation of the estimation error in the LP estimator. In particular, if we define $U := V - \sigma_\Delta(\Delta - b/\sigma_\Delta) \sim \mathcal{N}(0, \sigma_{LP}^2)$, then (U, Δ) has the same joint distribution as $(\hat{\theta}_{LP} - \theta_h, \Delta)$.

We follow [Armstrong et al. \[2025\]](#) and consider the following *invariance* restriction on the estimators $d : \mathbb{R}^2 \rightarrow \mathbb{R}$:

$$d_h(s + s', t) = d_h(s, t) + s' \quad \forall \quad s, s', t \in \mathbb{R}. \quad (20)$$

This restriction means that if we shift the unbiased LP estimator by s' units (but we keep the Hausman statistic constant), then the estimator d should also be adjusted by s' units. Note that, by definition of Δ , any estimator of θ_h that tries to additively “bias-adjust” the VAR estimator by means of a function $\hat{b}(\cdot)$ that only depends on Δ satisfies the invariance

restriction in (20). To see this, note that

$$\begin{aligned}
\hat{\theta}_{\text{VAR}} - \hat{b}(\Delta) &= \hat{\theta}_{\text{LP}} + \sigma_{\Delta}\Delta - \hat{b}(\Delta) \\
&\quad (\text{by definition of } \Delta) \\
&= \hat{\theta}_{\text{LP}} - \left(\hat{b}(\Delta) - \sigma_{\Delta}\Delta \right).
\end{aligned}$$

Note also that any invariant estimator $d_h(\hat{\theta}_{\text{LP}}, \Delta)$ can be written as a “bias-adjusted” VAR estimator:

$$d_h(\hat{\theta}_{\text{LP}}, \Delta) = d_h(0, \Delta) + \hat{\theta}_{\text{LP}} = \hat{\theta}_{\text{VAR}} + \underbrace{(-\sigma_{\Delta}\Delta + d(0, \Delta))}_{:= -\hat{b}(\Delta)}$$

The risk function of estimators satisfying the invariance restriction in (20) admits a convenient simplification:

$$\begin{aligned}
\mathcal{R}(d_h, \theta_h, b) &:= \mathbb{E}_{\theta, b}[(d_h(\hat{\theta}_{\text{LP}}, \Delta) - \theta_h)^2] \\
&= \mathbb{E}_{\theta, b}[(d_h(0, \Delta) + \hat{\theta}_{\text{LP}} - \theta_h)^2] \\
&= \mathbb{E}_{\theta, b}[(d_h(0, \Delta) + V - \sigma_{\Delta}(\Delta - b/\sigma_{\Delta}))^2], \quad V \perp \Delta, \quad V \sim \mathcal{N}(0, \sigma_{\text{VAR}}^2) \\
&= \sigma_{\text{VAR}}^2 + \sigma_{\Delta}^2 \mathbb{E}_{b/\sigma_{\Delta}} \left[\left(\left[\Delta - \frac{d_h(0, \Delta)}{\sigma_{\Delta}} \right] - \frac{b}{\sigma_{\Delta}} \right)^2 \right]. \tag{21}
\end{aligned}$$

Consequently, the worst-case risk of any invariant estimator does not depend on θ_h , and can be expressed in terms of the minimax risk of the bounded normal mean problem we discussed in Section 2. In particular, let \mathcal{D}_I be the set of all invariant estimators for θ_h . Then,

$$\bar{v}_h^*(m; \Sigma) := \inf_{d_h \in \mathcal{D}_I} \sup_{\theta_h, |b| \leq m \cdot \sigma_{\Delta}} \mathcal{R}(d_h, \theta_h, b) = \sigma_{\text{VAR}}^2 + \sigma_{\Delta}^2 \bar{v}^*(m). \tag{22}$$

Moreover, if $d_{\text{BNM-m}}^*$ denotes the minimax estimator for the bounded normal mean problem

with bound equal to m , the minimax invariant estimator of the parameter θ_h becomes

$$\begin{aligned} d_{\text{h-minimax-inv}}^*(\hat{\theta}_{\text{LP}}, \Delta) &= \hat{\theta}_{\text{LP}} + \sigma_{\Delta} \Delta - \sigma_{\Delta} d_{\text{BNM-m}}^*(\Delta) \\ &= \hat{\theta}_{\text{VAR}} - \sigma_{\Delta} d_{\text{BNM-m}}^*(\Delta). \end{aligned} \tag{23}$$

That is, the minimax invariant estimator for the impulse response coefficient θ_h , is simply a debiased VAR estimator, where the debiasing function is based on a minimax estimator of b/σ_{Δ} (imposing the bound m). The estimator of b/σ_{Δ} relies only on the marginal distribution of Δ .

By the same logic as in Section 3, the stochastic mirror ascent algorithm applied to the bounded normal mean problem with bound m induces, via (23), a stochastic minimax sequence of invariant estimators for θ_h , with worst-case risk converging to $\bar{v}_h^*(m; \Sigma)$. A formal statement is given in Appendix B.3.

5.3 Data

To illustrate our results, we consider four examples of impulse response estimation from four influential macroeconomic papers. In each application, impulse response estimates are obtained using both local projection and vector autoregression methods, including the same set of observable variables and controls. These estimates are based on the replication files provided by Montiel Olea et al. [2026]¹⁶.

1. Gertler and Karadi [2015]: This paper studies how monetary policy shocks transmit to economic and financial variables when shocks are allowed to include forward guidance. The authors identify monetary policy shocks using high-frequency interest-rate surprises around FOMC announcements as external instruments. The shock is a one-standard-deviation monetary policy tightening, which corresponds to a roughly

¹⁶Source: [GitHub repository with replication files for Montiel Olea et al. \[2026\]](#). The original data is from Ramey [2016].

25-basis-points increase in the one-year government bond rate on impact. We focus on the impulse response of log industrial production to this shock, using monthly data from January 1990 to June 2012 (270 data points) and including two lags.

2. [Romer and Romer \[2010\]](#): The authors study the macroeconomic effects of tax changes by using historical records to distinguish tax actions driven by policy motives from those responding to the state of the economy. Tax shocks are identified using the narrative classification of legislated tax changes. We focus on the impulse response of log real GDP per capita to a positive exogenous tax shock—measured as an increase in tax liabilities relative to GDP—using quarterly data from 1950:Q1 to 2007:Q2 (230 data points), including four lags and a quadratic trend.
3. [Ramey \[2011\]](#): This study shows that standard VAR identification may miss the timing of government-spending news, since agents may respond when news about future military spending arrives, before spending is realized. The authors identify government spending shocks using narrative measures of military spending news. We focus on the impulse response of log real GDP per capita to a government spending shock, corresponding to an increase in the expected present discounted value of future government spending associated with military events relative to previous-quarter GDP, using quarterly data from 1947:Q1 to 2013:Q2 (266 data points), including four lags and a quadratic trend.
4. [Francis et al. \[2014\]](#): They revisit the identification of technology shocks in structural VARs by replacing infinite-horizon long-run restrictions with a finite-horizon criterion. The authors propose the “Max Share” technique, which identifies the technology shock as the structural shock that maximizes the explained share of the forecast-error variance of labor productivity at a long but finite horizon. We focus on the impulse response of log labor productivity to a positive unit-variance technology shock selected by this “Max Share” criterion, using quarterly data from 1949:Q1 to 2009:Q2 (242 data points),

including four lags and a quadratic trend.

Figure 9 shows the estimated impulse responses along with bootstrap standard errors for both the LP and VAR estimators.

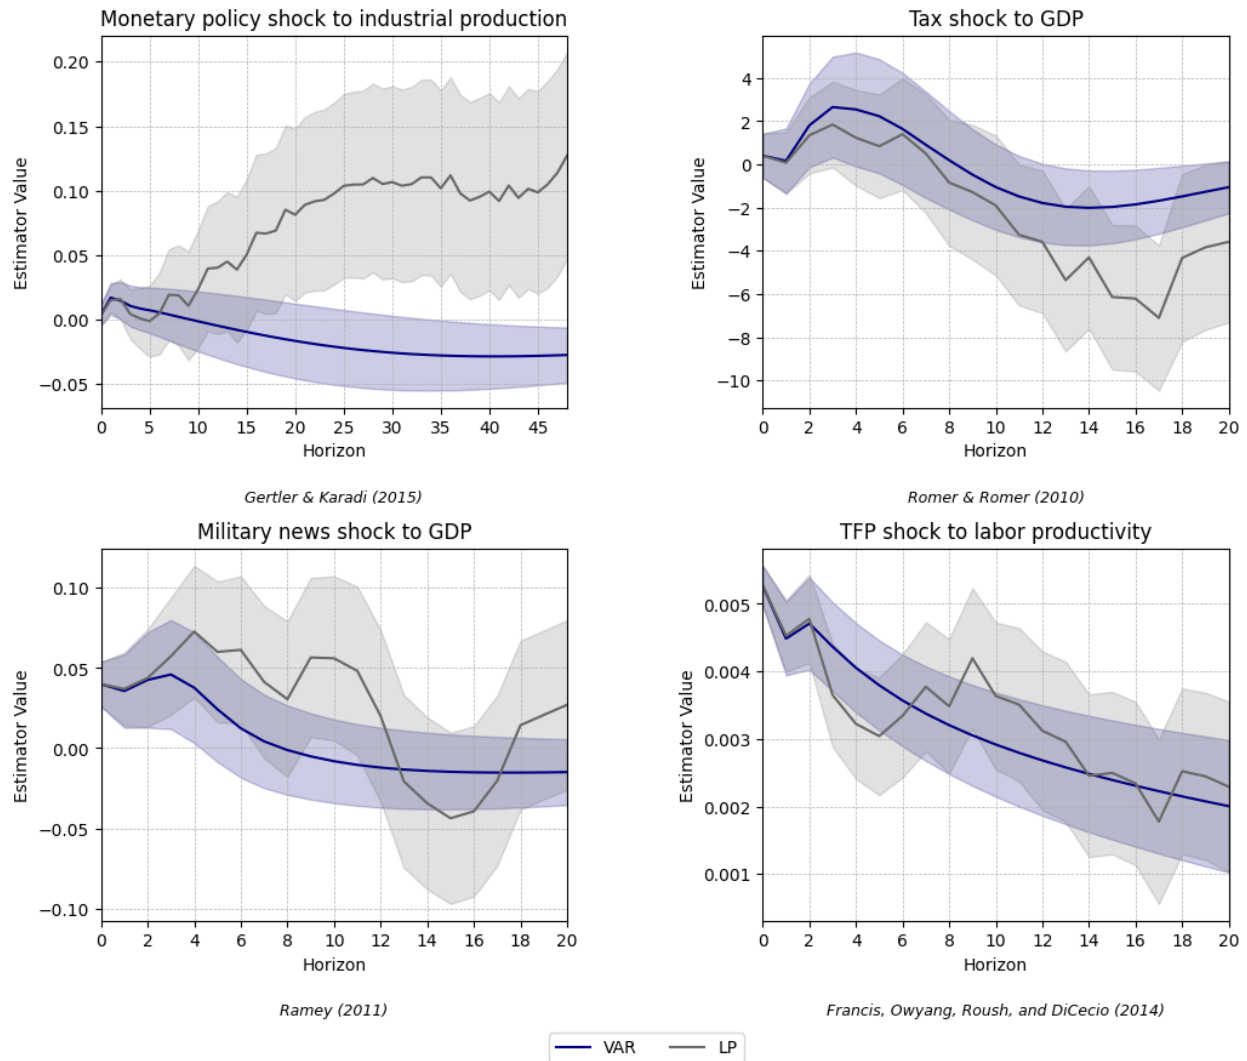


Figure 9: LP vs. VAR impulse responses in four applications.

Notes: Each panel reports point estimates from vector autoregressions (VAR, blue) and local projections (LP, gray). Shaded areas are bootstrapped standard errors. Top-left (monthly horizons): monetary policy shock to industrial production from [Gertler and Karadi \[2015\]](#). Other panels (quarterly horizons): tax shock to real GDP per capita from [Romer and Romer \[2010\]](#), government spending shocks to real GDP per capita from [Ramey \[2011\]](#), and unanticipated TFP shocks to labor productivity from [Francis et al. \[2014\]](#).

5.4 Implementation

We now apply the estimator in (23) to the four applications described above. For each application and horizon h , we compute: (1) the Hausman [1978]’s specification test statistic $|\Delta_h|$, (2) the bias correction $\sigma_{\Delta_h} d_{\text{BNM-m}}^*(\Delta_h)$, (3) the resulting minimax estimator as described by Equation (23), and (4) the percentage risk improvement relative to the minimax-optimal linear combination of LP and VAR. For reference, Table 1 reports the implied values of m and $\epsilon(m)$ across applications. We use those values to run the stochastic mirror ascent algorithm as described in Section 3 and construct an approximate minimax decision rule corresponding to Δ_h .

Table 1: Values of m across applications ($\mathcal{M} = 1\%$).

Application	T	$m = \sqrt{T} \times \mathcal{M}$	$\epsilon(m)$	Time (sec)
Gertler and Karadi [2015]	270	1.64	0.146	54
Romer and Romer [2010]	230	1.52	0.139	32
Ramey [2011]	266	1.63	0.145	54
Francis et al. [2014]	242	1.56	0.142	40

Figure 10 displays the value of the Hausman specification test statistic $|\Delta_h|$ across horizon in each application. We set the Hausman test statistic to be equal to zero at horizon $h = 0$ (although it theoretically equals 0 divided by 0, since the LP and VAR estimator coincide upon impact).

For each application, we compute the approximate rule $d_{\text{BNM-m}}^*$ in the corresponding minimax bounded normal mean problem using stochastic mirror ascent. The bias correction of the VAR estimator is nonlinear in Δ_h and is given by $\sigma_{\Delta_h} d_{\text{BNM-m}}^*(\Delta_h)$.

Figure 11 plots the implied nonlinear bias correction across horizons for the four different applications.

We then construct the bias-corrected VAR estimator for impulse responses in each of these applications. Figure 12 illustrates the comparison of this estimator to the underlying VAR and LP estimators.

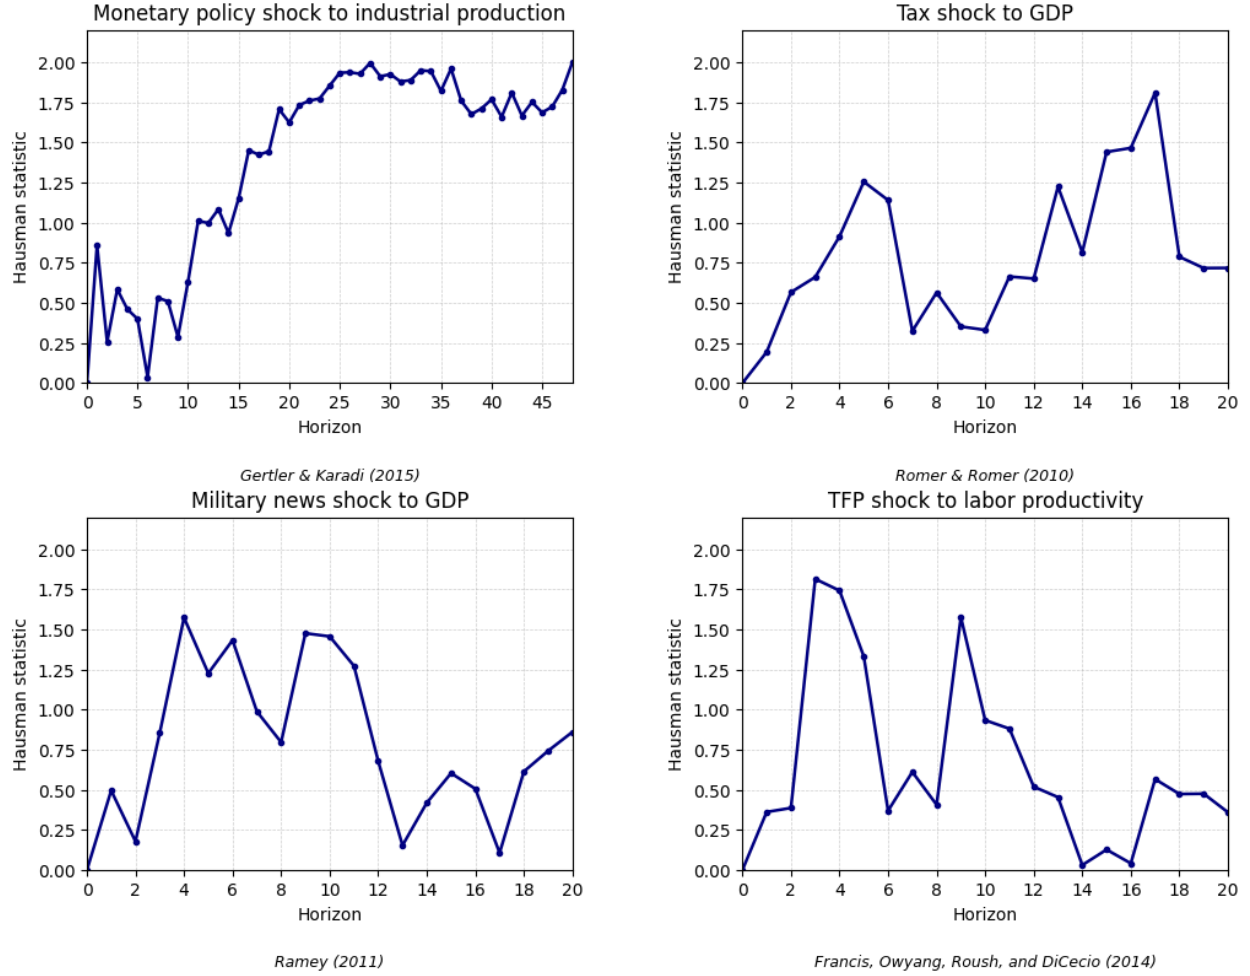


Figure 10: Hausman statistic across horizons.

To benchmark the nonlinear bias-corrected VAR, we compare it with the minimax-optimal linear combination of LP and VAR. That is, a linear combination of VAR and LP estimators, with weights chosen to minimize the worst-case risk. As shown in [Montiel Olea et al. \[2026, p. 40\]](#), the worst-case risk for such a linear combination is given by

$$(1 - \omega)^2 m^2 (\sigma_{LP}^2 - \sigma_{VAR}^2) + \sigma_{VAR}^2 + \omega^2 (\sigma_{LP}^2 - \sigma_{VAR}^2).$$

As shown in [Montiel Olea et al. \[2026\]](#), the optimal weight that minimizes this expression is

$$\omega^* = \frac{m^2}{1 + m^2},$$

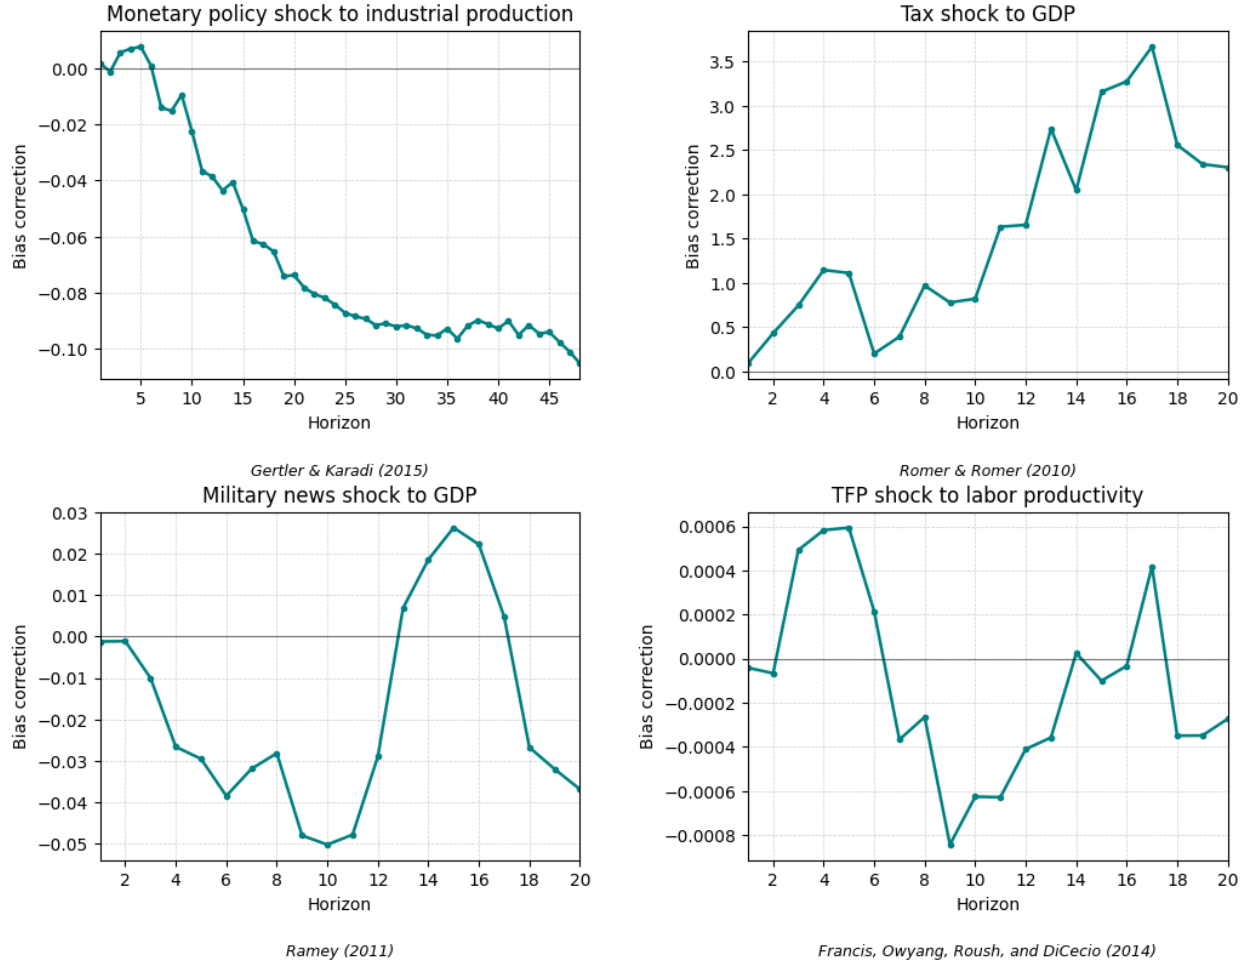


Figure 11: Bias correction.

leading to the estimator

$$d_{\text{linear}}^* = \hat{\theta}^{\text{VAR}} - \sigma_{\Delta} \omega^* \Delta.$$

In the following exercise, we compare these values with $\mathcal{R}(d_h, \theta_h, b)$, which uses the (grid-evaluated) worst-case risk of the optimal decision rule in the bounded normal mean problem (1,000 Monte Carlo draws).

The figure above compares the estimated worst-case risk of the nonlinear minimax correction with the worst-case risk of this optimal linear combination. Positive values indicate that the nonlinear rule improves on the best linear rule.

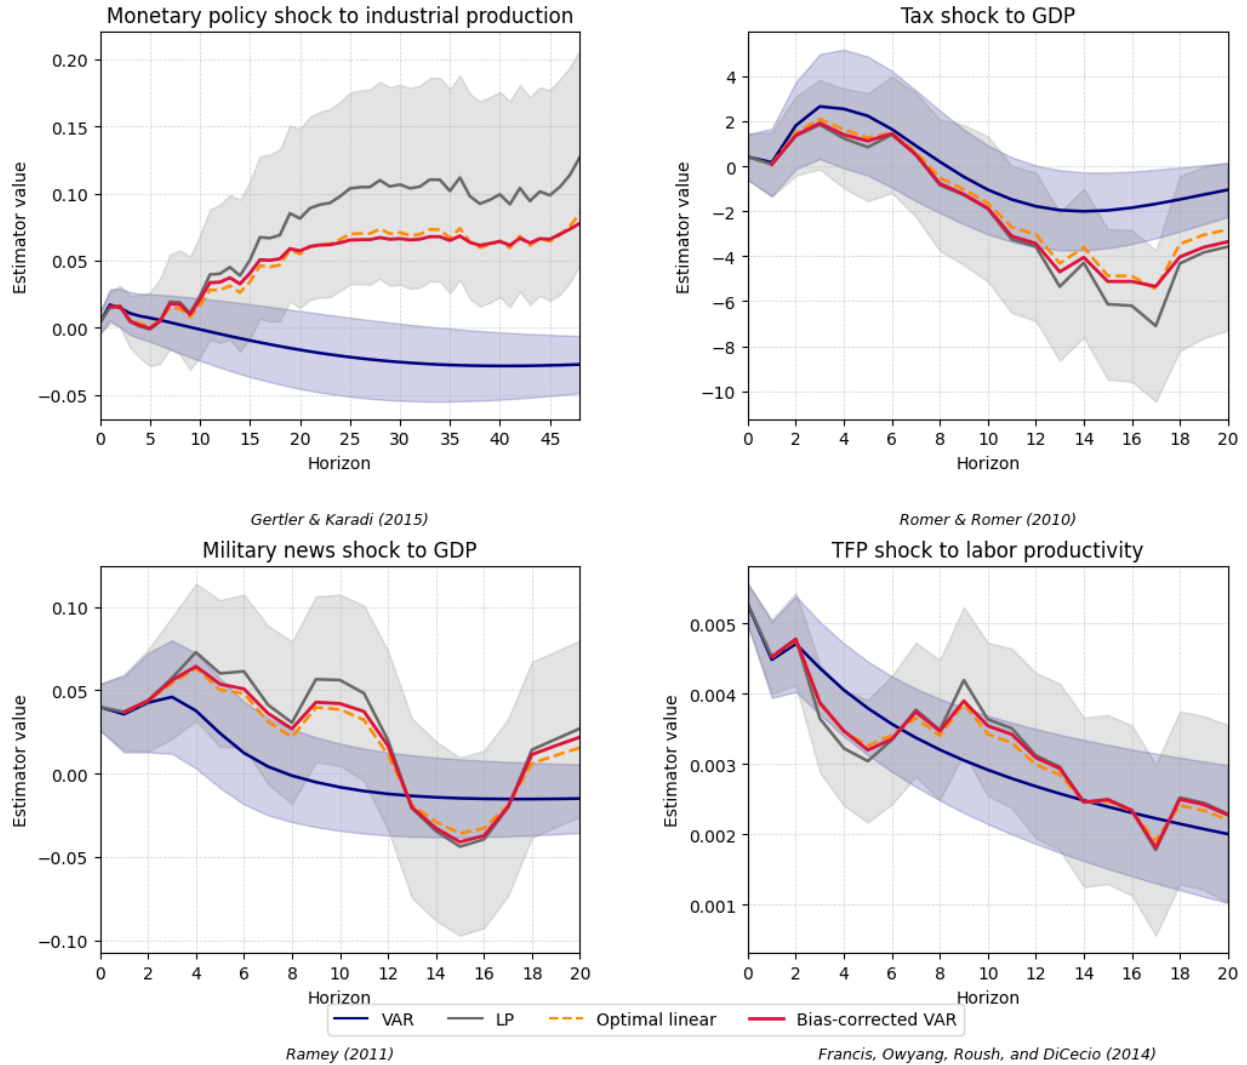


Figure 12: Comparison of VAR, LP and proposed bias-corrected VAR estimators.

6 Conclusion

We presented a computational approach to find an approximately minimax estimator (under quadratic loss) for the bounded normal mean problem of [Casella and Strawderman \[1981\]](#). Our suggested algorithm considers a fixed (equally-spaced) grid of I points that discretize the parameter space $[-m, m]$. We then approximately maximize Bayes risk over the space of probability distributions supported on the I points of the grid using a stochastic mirror ascent routine for concave maximization. We use this approximately least-favorable distribution to

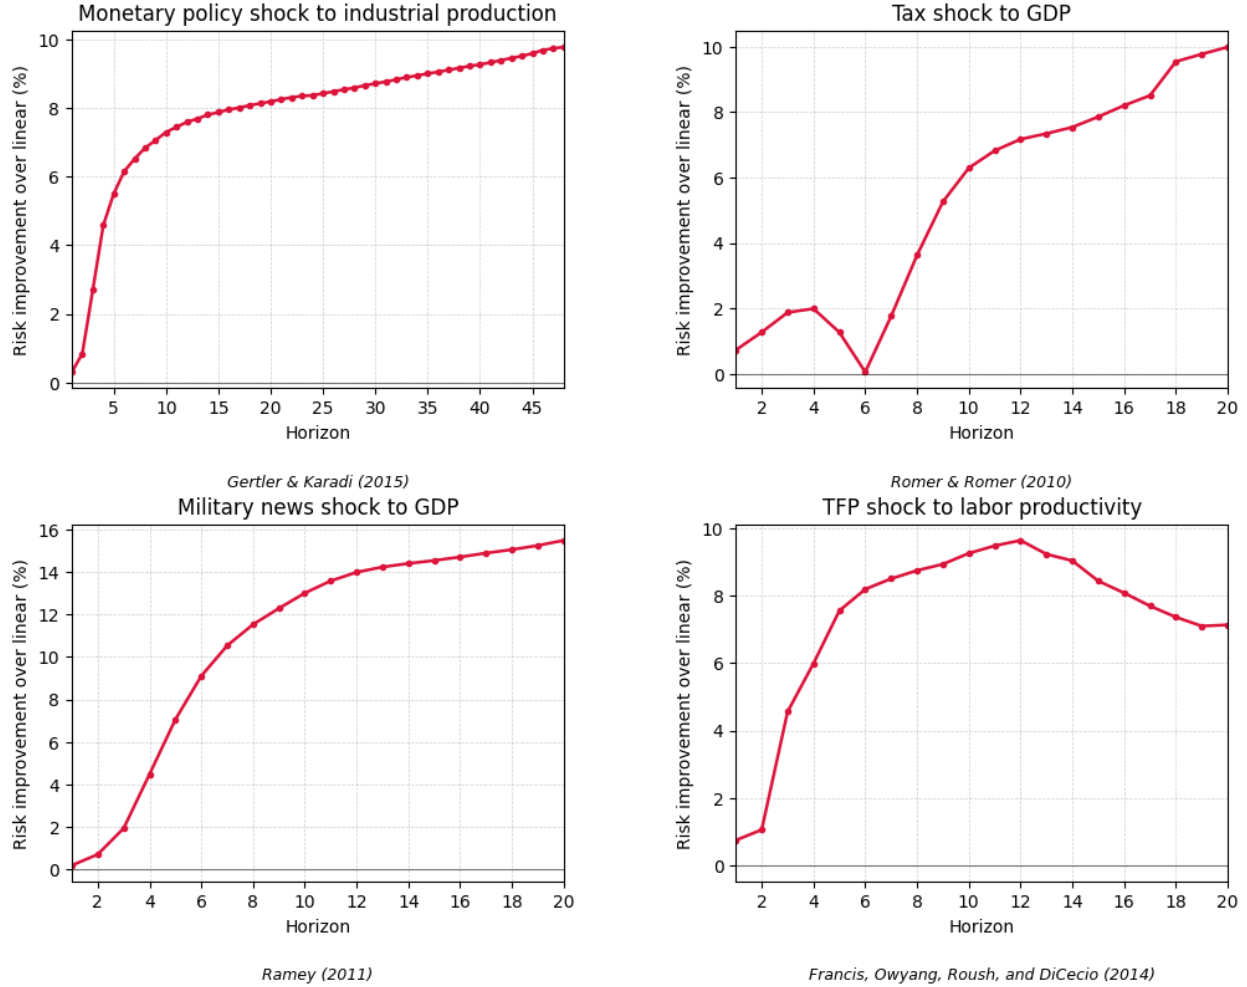


Figure 13: Risk improvement over linear combination.

find the corresponding Bayes estimator. Our main results (Theorems 1-4) showed that the approximately least-favorable distribution and the corresponding Bayes estimator are indeed close (in a sense we made precise) to their desired targets.

As we explained in the paper, our algorithm has two sources of approximation error: discretization and optimization error. The Lipschitz continuity property of the risk function that we established in Lemma 1 controls the first component. The well-known theoretical properties of stochastic mirror descent allow us to control the second component. Taking into account these two sources of approximation error, we showed that the approximately least-favorable distribution obtained from our algorithm converges in 1–Wasserstein distance to

the exact least-favorable distribution of the bounded normal mean problem. In addition, the approximately minimax estimator converges (uniformly over compacts) to the exact minimax estimator.

We think that several questions remain open. First, as we have explained in the paper, our algorithm relies on an ex ante discretization of the parameter space $[-m, m]$. This makes the problem tractable and allows the use of mirror ascent on a finite-dimensional probability simplex, but it also introduces discretization error that needs to be accounted for. A natural direction for future work would be to explore the possibility of approximating the least-favorable distribution by looking directly at the concave problem of maximizing Bayes risk over all probability distributions on the interval $[-m, m]$. An approach for this problem could be based on the algorithm presented in the work of [Kent, Li, Blanchet, and Glynn \[2021\]](#) for optimizing infinite dimensional functionals of probability measures.¹⁷ Another possibility would be to use recent advances in online learning with continuous action spaces; see [Krichene, Balandat, Tomlin, and Bayen \[2015\]](#) and [Negrea, Bilodeau, Campolongo, Orabona, and Roy \[2021\]](#). Second, some of the theoretical results we derived in order to control optimization error are based on constants/bounds that are conservative. For example, the Lipschitz constant for the risk function we derived in Lemma 1 could perhaps be tighter (which could decrease the number of iterations in Theorem 1). Similarly, the number of iterations for the mirror-ascent algorithm is designed to work for any (Lipschitz) concave optimization problem over the simplex. Since we are only interested in solving the bounded normal mean problem, it is perhaps possible to sharpen the analysis and reduce the number of iterations. Third, the least-favorable distribution in the bounded normal mean problem is known to have a very particular structure: the support of the distribution is finite. The stochastic mirror ascent algorithm that we use is initialized at the uniform distribution in Δ^{I-1} , and by design, the output of the algorithm places positive mass on every point (although as our simulations show, the mass associated with some of these points is

¹⁷In Appendix C we derived different properties of the general maximin problem that could be helpful for this purpose.

numerically very close to zero). It could be useful to consider other algorithms for finding the least-favorable distribution that increase the number of support points in every iteration, as in the recent work of Guggenberger and Huang [2025].

We view our paper as part of a larger research program that emphasizes the importance of providing a careful analysis of the computational aspects surrounding problems in econometrics and statistics; see, for example, the recent work of Cattaneo, Klusowski, and Shigida [2024]; Cattaneo and Shigida [2025a,b]. As mentioned by Chandrasekaran and Jordan [2013] *“Although classical statistics gave little consideration to computational complexity, computational issues have come increasingly to the fore in modern high-dimensional statistics.”* There are many examples now within statistics and machine learning that illustrate the benefits of bringing *computational awareness* to the forefront of these fields.¹⁸ Advances in econometrics in this regard have not followed the same pace, but as mentioned recently by Patrick Kline in one of his interventions at the Online Chamberlain Seminar reflecting on the future of econometrics *“There is probably going to be a more important role for computation moving forward. I think we have not rewarded computational advances enough in the profession, and we ought to be more serious about that ... And it is important for procedures that are proposed to actually be implementable. And in some cases ... Having some guarantees on computability of different procedures is important and I think we ought to learn more computer science and we will.”*¹⁹

Finally, we hope that the methods outlined in this paper could make the application of statistical decision theory in econometrics more frequent. As noted by Manski [2021] the *“primary challenge to use of statistical decision theory is computational.”*

¹⁸See, for illustration, the lectures of the *“Computational Complexity of Statistical Inference Boot Camp* of the Simons Institute for the Theory of Computing: <https://www.youtube.com/playlist?list=PLgKuh-1Kre13obZNGYdY-25ihfDOPIyJ->

¹⁹See minute 19:58 at https://www.youtube.com/watch?v=ifp_BUK13ig

References

- ARADILLAS FERNANDEZ, A., J. BLANCHET, J. L. MONTIEL OLEA, C. QIU, J. STOYE, AND L. TAN (2024): “ ε -Minimax Solutions of Statistical Decision Problems,” Working Paper.
- ARADILLAS FERNÁNDEZ, A., J. BLANCHET, J. L. MONTIEL OLEA, C. QIU, J. STOYE, AND L. TAN (2025): “Approximate least-favorable distributions and nearly optimal tests via stochastic mirror descent,” *arXiv preprint arXiv:2511.16925*.
- ARMSTRONG, T. B., P. KLINE, AND L. SUN (2025): “Adapting to Misspecification,” *Econometrica*, 93, 1981–2005.
- BARSOV, S. S. AND V. V. ULYANOV (1987): “Estimates of the Proximity of Gaussian Measures,” *Soviet Mathematics Doklady*, 34, 462–466.
- BEN-TAL, A., T. MARGALIT, AND A. NEMIROVSKI (2001): “The ordered subsets mirror descent optimization method with applications to tomography,” *SIAM Journal on Optimization*, 12, 79–108.
- BICKEL, P. J. (1981): “Minimax estimation of the mean of a normal distribution when the parameter space is restricted,” *The Annals of Statistics*, 9, 1301–1309.
- BLACKWELL, D. A. AND M. A. GIRSHICK (1954): *Theory of games and statistical decisions*, John Wiley, New York.
- BLANCHET, J., Y. KANG, J. L. M. OLEA, V. A. NGUYEN, AND X. ZHANG (2023): “Dropout Training is Distributionally Robust Optimal,” *Journal of Machine Learning Research*, 24, 1–60.
- BOUCHERON, S., G. LUGOSI, AND P. MASSART (2013): *Concentration Inequalities: A Nonasymptotic Theory of Independence*, Oxford University Press.

- BUBECK, S. (2015): “Convex Optimization: Algorithms and Complexity,” *Foundations and Trends in Machine Learning*, 8, 231–357.
- CASELLA, G. AND W. E. STRAWDERMAN (1981): “Estimating a Bounded Normal Mean,” *The Annals of Statistics*, 9, 870–878.
- CATTANEO, M. D., J. M. KLUSOWSKI, AND B. SHIGIDA (2024): “On the Implicit Bias of Adam,” in *Forty-first International Conference on Machine Learning*.
- CATTANEO, M. D. AND B. SHIGIDA (2025a): “How Memory in Optimization Algorithms Implicitly Modifies the Loss,” in *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.
- (2025b): “Modified Loss of Momentum Gradient Descent: Fine-Grained Analysis,” *arXiv preprint arxiv:2509.08483*.
- CHAMBERLAIN, G. (2000): “Econometric applications of maxmin expected utility,” *Journal of Applied Econometrics*, 15, 625–644.
- CHANDRASEKARAN, V. AND M. I. JORDAN (2013): “Computational and statistical trade-offs via convex relaxation,” *Proceedings of the National Academy of Sciences*, 110, E1181–E1190.
- CHEN, C., E. PESAVENTO, AND B. VONNÁK (2026): “Estimator Averaging of Local Projection and VAR Impulse Responses,” *arXiv preprint arXiv:2605.05456*.
- DELATTE, P. (2026): “Existence of maximin priors and statistical minimax theorems,” *Working Paper*.
- DEVROYE, L., A. MEHRABIAN, AND T. REDDAD (2018): “The total variation distance between high-dimensional Gaussians with the same mean,” *arXiv preprint arXiv:1810.08693*.
- DONOHO, D. L., R. C. LIU, AND B. MACGIBBON (1990): “Minimax Risk Over Hyperrectangles, and Implications,” *The Annals of Statistics*, 18, 1416–1437.

- ELLIOTT, G., U. K. MÜLLER, AND M. W. WATSON (2015): “Nearly optimal tests when a nuisance parameter is present under the null hypothesis,” *Econometrica*, 83, 771–811.
- FELDMAN, I. AND L. D. BROWN (1989): “The minimax risk for estimating a bounded normal mean,” Tech. rep., Technical report, Cornell Univ., Cornell Statistical Center Series.
- FERGUSON, T. S. (1967): *Mathematical Statistics: A Decision Theoretic Approach*, New York: Academic Press.
- FRANCIS, N., M. T. OWYANG, J. E. ROUSH, AND R. DICECIO (2014): “A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks,” *Review of Economics and Statistics*, 96, 638–647.
- GERTLER, M. AND P. KARADI (2015): “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics*, 7, 44–76.
- GHOSH, M. (1964): “Uniform approximation of minimax point estimates,” *The Annals of Mathematical Statistics*, 1031–1047.
- GOURDIN, E., B. JAUMARD, AND B. MACGIBBON (1994): “Global optimization decomposition methods for bounded parameter minimax risk evaluation,” *SIAM Journal on Scientific Computing*, 15, 16–35.
- GUGGENBERGER, P. AND J. HUANG (2025): “On the numerical approximation of minimax regret rules via fictitious play,” *arXiv preprint arXiv:2503.10932*.
- HAUSMAN, J. A. (1978): “Specification tests in econometrics,” *Econometrica: Journal of the econometric society*, 1251–1271.
- IBRAGIMOV, I. A. AND R. Z. HAS’MINSKII (1984): “On nonparametric estimation of the value of a linear functional in Gaussian white noise,” *Theory of Probability & Its Applications*, 29, 18–32.

- JOHNSTONE, I. M. (2019): “Gaussian Estimation: Sequence and Wavelet Models,” Draft version, September 16, 2019.
- JOHNSTONE, I. M. AND K. B. MACGIBBON (1992): “Minimax estimation of a constrained Poisson vector,” *The Annals of Statistics*, 20, 807–831.
- JORDÀ, Ò. AND A. M. TAYLOR (2025): “Local projections,” *Journal of Economic Literature*, 63, 59–110.
- JORDÀ, (2005): “Estimation and Inference of Impulse Responses by Local Projections,” *American Economic Review*, 95, 161–182.
- KEMPTHORNE, P. J. (1987): “Numerical specification of discrete least favorable prior distributions,” *SIAM Journal on Scientific and Statistical Computing*, 8, 171–184.
- KENT, C., J. LI, J. BLANCHET, AND P. W. GLYNN (2021): “Modified Frank Wolfe in probability space,” *Advances in Neural Information Processing Systems*, 34, 14448–14462.
- KRICHENE, W., M. BALANDAT, C. TOMLIN, AND A. BAYEN (2015): “The hedge algorithm on a continuum,” in *International Conference on Machine Learning*, PMLR, 824–832.
- LEVIT, B. Y. (1981): “On asymptotic minimax estimates of the second order,” *Theory of Probability & Its Applications*, 25, 552–568.
- LI, D., M. PLAGBORG-MØLLER, AND C. K. WOLF (2024): “Local projections vs. VARs: Lessons from thousands of DGPs,” *Journal of Econometrics*, 244, 105722.
- 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*.
- MANSKI, C. F. (2021): “Econometrics for decision making: Building foundations sketched by Haavelmo and Wald,” *Econometrica*, 89, 2827–2853.

- MONTIEL OLEA, J. L., M. PLAGBORG-MØLLER, E. QIAN, AND C. K. WOLF (2026): “Double Robustness of Local Projections and Some Unpleasant VARithmetic,” *Econometrica*, forthcoming.
- MONTIEL OLEA, J. L. AND M. PLAGBORG-MØLLER (2021): “Local Projection Inference is Simpler and More Robust Than You Think,” *Econometrica*, 89, 1789–1823.
- MONTIEL OLEA, J. L., E. QIAN, C. K. WOLF, AND M. PLAGBORG-MØLLER (2025): “Local Projections or VARs? A Primer for Macroeconomists,” *NBER Macroeconomics Annual*, forthcoming.
- NEGREA, J., B. BILODEAU, N. CAMPOLONGO, F. ORABONA, AND D. ROY (2021): “Minimax optimal quantile and semi-adversarial regret via root-logarithmic regularizers,” *Advances in Neural Information Processing Systems*, 34, 26237–26249.
- NEMIROVSKI, A., A. JUDITSKY, G. LAN, AND A. SHAPIRO (2009): “Robust stochastic approximation approach to stochastic programming,” *SIAM Journal on optimization*, 19, 1574–1609.
- NEMIROVSKI, A. AND D. YUDIN (1983): *Problem Complexity and Method Efficiency in Optimization*, New York: Wiley.
- NEMTYREV, A. AND O. BOLDEA (2026): “Targeted Local Projections,” *arXiv preprint arXiv:2603.00248*.
- PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica*, 89, 955–980.
- RAMEY, V. A. (2011): “Identifying Government Spending Shocks: It’s all in the Timing,” *Quarterly Journal of Economics*, 126, 1–50.
- (2016): “Macroeconomic Shocks and Their Propagation,” in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Elsevier, vol. 2, chap. 2, 71–162.

- ROMER, C. D. AND D. H. ROMER (2010): “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review*, 100, 763–801.
- SIMS, C. A. (1980): “Macroeconomics and Reality,” *Econometrica*, 48, 1–48.
- STOCK, J. H. AND M. W. WATSON (2018): “Identification and estimation of dynamic causal effects in macroeconomics using external instruments,” *The Economic Journal*, 128, 917–948.
- VILLANI, C. (2021): *Topics in optimal transportation*, vol. 58, American Mathematical Soc.
- WILLIAMSON, D. P. AND D. B. SHMOYS (2011): *The design of approximation algorithms*, Cambridge university press.
- XU, K.-L. (2023): “Local projection based inference under general conditions,” *Available at SSRN 4372388*.

APPENDIX

Appendix A. Proofs of Main Results

A.1 Proof of Lemma 1

Let $\phi(\cdot; \theta)$ denote the p.d.f. of a univariate normal random variable with mean θ and unit variance. By definition

$$\begin{aligned} R(d, \theta) &= \int_{-\infty}^{\infty} (d(y) - \theta)^2 \phi(y; \theta) dy \\ &= \int_{-\infty}^{\infty} d(y)^2 \phi(y; \theta) dy - 2\theta \int_{-\infty}^{\infty} d(y) \phi(y; \theta) dy + \theta^2. \end{aligned}$$

Algebra then shows that for any decision rule $d \in \mathcal{D}$, and any $\theta, \theta' \in [-m, m]$ we have

$$\begin{aligned} R(d, \theta) - R(d, \theta') &= \int_{-\infty}^{\infty} d(y)^2 (\phi(y; \theta) - \phi(y; \theta')) dy \\ &\quad - 2\theta \int_{-\infty}^{\infty} d(y) (\phi(y; \theta) - \phi(y; \theta')) dy \\ &\quad + 2(\theta' - \theta) \int_{-\infty}^{\infty} d(y) \phi(y; \theta') dy \\ &\quad + (\theta - \theta')(\theta + \theta'). \end{aligned}$$

Consequently,

$$|R(d, \theta) - R(d, \theta')| \leq 3m^2 \int_{-\infty}^{\infty} |\phi(y; \theta) - \phi(y; \theta')| dy + 4m|\theta - \theta'|.$$

Note that

$$\frac{1}{2} \int_{-\infty}^{\infty} |\phi(y; \theta) - \phi(y; \theta')| dy$$

is the total variation distance between two univariate normals, with variance 1 and means θ and θ' , respectively. Therefore, Theorem 1.3 in [Devroye et al. \[2018\]](#) implies that

$$\frac{1}{2} \int_{-\infty}^{\infty} |\phi(y; \theta) - \phi(y; \theta')| dy \leq \frac{|\theta - \theta'|}{2}.$$

Consequently,

$$|R(d, \theta) - R(d, \theta')| \leq (3m^2 + 4m) |\theta - \theta'|.$$

A.2 Proof of Lemma 2

Note first that

$$\underline{v}(m; I) := \max_{\pi \in \Delta^{I-1}} \inf_{d \in \mathcal{D}} r(d, \pi) \leq \underline{v}^*(m) := \max_{\pi \in \Delta([-m, m])} \inf_{d \in \mathcal{D}} r(d, \pi),$$

since $\Delta^{I-1} \subseteq \Delta([-m, m])$. Thus, the remaining part of the proof focuses on establishing the upper bound.

For every $\theta \in [-m, m]$, let $N_I(\theta)$ denote a nearest neighbor of θ in the grid $\{\theta_i\}_{i=1}^I$. That is, $N_I(\theta)$ is any element that solves

$$\min_{i=1, \dots, I} |\theta_i - \theta|.$$

Because the grid is equally spaced with step $\Delta(m, I) = 2m/(I - 1)$, the maximum distance from any $\theta \in [-m, m]$ to its nearest neighbor is given by

$$\sup_{\theta \in \Theta} \min_{i=1, \dots, I} |\theta_i - \theta| = \frac{1}{2} \Delta(m, I).$$

Let d_I^* be a minimax estimator for the discretized problem, in the sense that

$$\sup_{i=1, \dots, I} R(d_I^*, \theta_i) = \bar{v}(m, I) := \inf_{d \in \mathcal{D}} \sup_{i=1, \dots, I} R(d, \theta_i).$$

By the Lipschitz continuity of the risk function established in Lemma 1 we have that for any $\theta \in [-m, m]$

$$R(d_I^*, \theta) \leq R(d_I^*, N_I(\theta)) + (3m^2 + 4m)\frac{1}{2}\Delta(m, I),$$

where $\Delta(m, I)$ is the gap between consecutive elements of the grid defined in (8). Then,

$$\begin{aligned} \sup_{\theta \in [-m, m]} R(d_I^*, \theta) &\leq \sup_{\theta \in [-m, m]} R(d_I^*, N_I(\theta)) + (3m^2 + 4m)\frac{1}{2}\Delta(m, I), \\ &= \sup_{i=1, \dots, I} R(d_I^*, \theta_i) + (3m^2 + 4m)\frac{1}{2}\Delta(m, I) \\ &= \bar{v}(m, I) + (3m^2 + 4m)\frac{1}{2}\Delta(m, I). \end{aligned}$$

Let d^* be the minimax estimator in the original Bounded Normal Mean Problem. This means that

$$\sup_{\theta \in [-m, m]} R(d^*, \theta) = \bar{v}^*(m) \leq \sup_{\theta \in [-m, m]} R(d_I^*, \theta) \leq \bar{v}(m, I) + (3m^2 + 4m)\frac{1}{2}\Delta(m, I).$$

The minimax theorem holds for the Bounded Normal Mean problem $\bar{v}^*(m) = \underline{v}^*(m)$; see Proposition 4.19 in Johnstone [2019] and also their Theorem A.5. And note also that the minimax theorem holds for the discretized problem (see Theorem 2.4.2 in Blackwell and Girshick [1954] and also the discussion regarding S-games in Appendix C of Aradillas Fernandez et al. [2024]). Since the minimax theorem for the discretized problem holds, we have

$$\underline{v}(m; I) = \inf_{d \in \mathcal{D}} \sup_{i=1, \dots, I} R(d, \theta_i).$$

The result then follows.

A.3 Proof of Theorem 1

Let $\pi^{*,I}$ be any maximizer of the discretized maximin problem, so that

$$\underline{v}(m; I) = \inf_{d \in \mathcal{D}} r(d, \pi^{*,I}).$$

In a slight abuse of notation, for each epoch $j = 1, \dots, J$, let g_j denote the supergradient vector $g(\pi_j)$. Let \tilde{g}_j be the estimator of the supergradient defined in Algorithm 1. By definition of supergradient, we have that for every $\pi \in \Delta^{I-1}$:

$$\inf_{d \in \mathcal{D}} r(d, \pi_j) \geq \inf_{d \in \mathcal{D}} r(d, \pi) + g_j^\top (\pi_j - \pi).$$

Taking $\pi = \pi^{*,I}$ and rearranging yields

$$\underline{v}(m; I) - \inf_{d \in \mathcal{D}} r(d, \pi_j) \leq g_j^\top (\pi^{*,I} - \pi_j).$$

Averaging over $j = 1, \dots, J$ gives

$$\underline{v}(m; I) - \frac{1}{J} \sum_{j=1}^J \inf_{d \in \mathcal{D}} r(d, \pi_j) \leq \frac{1}{J} \sum_{j=1}^J g_j^\top (\pi^{*,I} - \pi_j). \quad (24)$$

Decompose the right-hand side as

$$\begin{aligned} g_j^\top (\pi^{*,I} - \pi_j) &= \tilde{g}_j^\top (\pi^{*,I} - \pi_j) + (g_j - \tilde{g}_j)^\top (\pi^{*,I} - \pi_j) \\ &= \tilde{g}_j^\top (\pi^{*,I} - \pi_j) + (\tilde{g}_j - g_j)^\top (\pi_j - \pi^{*,I}). \end{aligned}$$

Substituting into (24) yields

$$\underline{v}(m; I) - \frac{1}{J} \sum_{j=1}^J \inf_{d \in \mathcal{D}} r(d, \pi_j) \leq \frac{1}{J} \sum_{j=1}^J \tilde{g}_j^\top (\pi^{*,I} - \pi_j) + \frac{1}{J} \sum_{j=1}^J (\tilde{g}_j - g_j)^\top (\pi_j - \pi^{*,I}). \quad (25)$$

By inequality (4.10) in [Bubeck \[2015\]](#) applied to the optimization domain Δ^{I-1} ,

$$\sum_{j=1}^J (-\tilde{g}_j)^\top (\pi_j - \pi) \leq \frac{R^2}{\eta} + \frac{\eta}{2} \sum_{j=1}^J \|\tilde{g}_j\|_\infty^2 \quad \forall \pi \in \Delta^{I-1},$$

where η is the step-size used in the stochastic mirror descent routine and $R^2 := \log(I)$.

Taking $\pi = \pi^{*,I}$, using $\|\tilde{g}_j\|_\infty = \|\tilde{g}_j\|_\infty \leq M$ almost surely, and dividing by J yields

$$\frac{1}{J} \sum_{j=1}^J \tilde{g}_j^\top (\pi^{*,I} - \pi_j) \leq \frac{\log(I)}{\eta J} + \frac{\eta M^2}{2}. \quad (26)$$

Next, Lemma 3 in Appendix A applied with $\pi = \pi^{*,I}$ implies that for any $\delta > 0$,

$$\mathbb{P} \left(\frac{1}{J} \sum_{j=1}^J (\tilde{g}_j - g_j)^\top (\pi_j - \pi^{*,I}) < \delta \right) \geq 1 - \exp \left(-\frac{\delta^2 J}{128m^4} \right). \quad (27)$$

Thus, combining (25) and (26) gives

$$\frac{1}{J} \sum_{j=1}^J \inf_{d \in \mathcal{D}} r(d, \pi_j) \geq \underline{v}(m; I) - \frac{\log(I)}{\eta J} - \frac{\eta M^2}{2} - \delta,$$

with probability at least $1 - \exp \left(-\frac{\delta^2 J}{128m^4} \right)$. Finally, by concavity of Bayes risk (Lemma 7) and Jensen's inequality,

$$\inf_{d \in \mathcal{D}} r \left(d, \frac{1}{J} \sum_{j=1}^J \pi_j \right) \geq \frac{1}{J} \sum_{j=1}^J \inf_{d \in \mathcal{D}} r(d, \pi_j),$$

Since the step-size η was chosen as $\eta(\epsilon) = (\epsilon/2)/M^2$ and the number of iterations, J , was set to $J(\epsilon) = 2M^2 \log(I)/(\epsilon/2)^2$

$$\frac{\log(I)}{\eta J} + \frac{\eta M^2}{2} = \frac{\epsilon}{4} + \frac{\epsilon}{4} = \frac{\epsilon}{2},$$

which yields

$$\inf_{d \in \mathcal{D}} r(d, \pi^\epsilon) \geq \underline{v}(m; I) - \epsilon/2 - \delta.$$

with probability at least $1 - \exp\left(-\frac{\delta^2 J(\epsilon)}{128m^4}\right)$. By Lemma 2 and the choice of $I(\epsilon)$ in the Theorem,

$$\underline{v}(m; I(\epsilon)) \geq \underline{v}^*(m) - \epsilon/2,$$

so the previous expression becomes

$$\inf_{d \in \mathcal{D}} r(d, \pi^\epsilon) \geq \underline{v}^*(m) - \epsilon - \delta.$$

To finish the proof, pick $\alpha \in (0, 1)$ and set

$$\delta(\alpha, \epsilon) := \sqrt{\frac{128m^4 \ln(1/\alpha)}{J(\epsilon)}}$$

Note that

$$\exp\left(-\frac{\delta(\alpha, \epsilon)^2 J(\epsilon)}{128m^4}\right) = \alpha.$$

Algebra shows that

$$\begin{aligned} \delta(\alpha, \epsilon) &\leq \frac{\epsilon}{2} \sqrt{\frac{128m^4 \ln(1/\alpha)}{2M^2 \ln(I(\epsilon))}} \\ &= \frac{8m^2 \epsilon}{M} \frac{1}{2} \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}} \\ &= \epsilon \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}}, \end{aligned}$$

where the last line uses the definition of $M := 4m^2$. We conclude that with probability at least $1 - \alpha$

$$\inf_{d \in \mathcal{D}} r(d, \pi^\epsilon) \geq \underline{v}^*(m) - \epsilon \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon))}}\right).$$

A.4 Proof of Theorem 2

Lemma 6 shows that the map

$$f : \Delta([-m, m]) \rightarrow \mathbb{R}_+, \quad f(\pi) := \inf_{d \in \mathcal{D}} r(d, \pi),$$

is Lipschitz continuous with respect to the 1-Wasserstein distance $W(\cdot, \cdot)$. Fix $\xi > 0$ and define the set

$$S(\xi) := \{\pi \in \Delta([-m, m]) : W(\pi, \pi^*) > \xi\}.$$

Since π^* is the unique least-favorable distribution (Proposition 4.19 in [Johnstone \[2019\]](#)), it is the unique maximizer of $f(\cdot)$. Therefore

$$\sup_{\pi \in S(\xi)} f(\pi) < f(\pi^*) = \underline{v}^*(m).$$

The inequality above implies the existence of a scalar $\nu(\xi) > 0$ such that

$$\sup_{\pi \in S(\xi)} f(\pi) < \underline{v}^*(m) - \nu(\xi).$$

Therefore, for any $\xi > 0$,

$$\begin{aligned} P_n(W(\pi_n, \pi^*) > \xi) &\leq P_n(f(\pi_n) < \underline{v}^*(m) - \nu(\xi)) \\ &= 1 - P_n(f(\pi_n) \geq \underline{v}^*(m) - \nu(\xi)) \\ &= 1 - P_n\left(\underline{v}^*(m) \geq \inf_{d \in \mathcal{D}} r(d, \pi_n) \geq \underline{v}^*(m) - \nu(\xi)\right). \end{aligned}$$

It follows from Theorem 1 and the definition of π_n that for any $\alpha \in (0, 1)$

$$P_n\left(\underline{v}^*(m) \geq \inf_{d \in \mathcal{D}} r(d, \pi_n) \geq \underline{v}^*(m) - \epsilon_n \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon_n))}}\right)\right) \geq 1 - \alpha.$$

Since $\epsilon_n \rightarrow 0$ and $I(\epsilon_n) \rightarrow \infty$, we have that there exists $n(\xi, \alpha)$ large enough such that for $n \geq n(\xi, \alpha)$:

$$\epsilon_n \left(1 + \sqrt{\frac{\ln(1/\alpha)}{\ln(I(\epsilon_n))}} \right) \leq \nu(\xi).$$

Therefore, for $n \geq n(\xi, \alpha)$

$$P_n \left(\underline{v}^*(m) \geq \inf_{d \in \mathcal{D}} r(d, \pi_n) \geq \underline{v}^*(m) - \nu(\xi) \right) \geq 1 - \alpha,$$

and consequently,

$$P_n(W(\pi_n, \pi^*) > \xi) \leq \alpha.$$

A.5 Proof of Theorem 3

Let π^* be the least-favorable distribution of the bounded normal mean problem. Algebra shows that for any $\theta \in [-m, m]$.

$$\begin{aligned} |R(d_n, \theta) - R(d_{\pi^*}, \theta)| &= \left| \int_{-\infty}^{\infty} \left((d_n(y) - \theta)^2 - (d_{\pi^*}(y) - \theta)^2 \right) \phi(y; \theta) dy \right| \\ &\leq \int_{-\infty}^{\infty} |(d_n(y) - \theta)^2 - (d_{\pi^*}(y) - \theta)^2| \phi(y; \theta) dy \\ &= \int_{-\infty}^{\infty} |(d_n(y)^2 - 2d_n(y)\theta - d_{\pi^*}(y)^2 + 2d_{\pi^*}(y)\theta)| \phi(y; \theta) dy \\ &= \int_{-\infty}^{\infty} |d_n(y) - d_{\pi^*}(y)| \cdot \underbrace{|d_n(y) + d_{\pi^*}(y) - 2\theta|}_{\leq 4m} \phi(y; \theta) dy \\ &\leq 4m \int_{-\infty}^{\infty} |d_n(y) - d_{\pi^*}(y)| \phi(y; \theta) dy. \end{aligned}$$

Theorem 2 showed that $W(\pi_n, \pi^*) \xrightarrow{P} 0$. Using Lemma 4 and the fact that $d_n := d_{\pi_n}$, we have that

$$\sup_{\theta \in [-m, m]} \int_{-\infty}^{\infty} |d_n(y) - d_{\pi^*}(y)| \phi(y; \theta) dy \xrightarrow{P} 0,$$

which implies that

$$\sup_{\theta \in [-m, m]} |R(d_n, \theta) - R(d_{\pi^*}, \theta)| \xrightarrow{P} 0. \quad (28)$$

To finalize the proof note that

$$\begin{aligned} \sup_{\theta \in [-m, m]} R(d_n, \theta) &= \sup_{\theta \in [-m, m]} R(d_n, \theta) - R(d_{\pi^*}, \theta) + R(d_{\pi^*}, \theta) \\ &\leq \sup_{\theta \in [-m, m]} |R(d_n, \theta) - R(d_{\pi^*}, \theta)| + \sup_{\theta \in [-m, m]} R(d_{\pi^*}, \theta) \\ &= \sup_{\theta \in [-m, m]} |R(d_n, \theta) - R(d_{\pi^*}, \theta)| + \bar{v}^*(m), \end{aligned}$$

where the last line follows from the fact that d_{π^*} is minimax, by Proposition 4.19 in [Johnstone \[2019\]](#). Therefore, for any $\varepsilon > 0$

$$P_n \left(\sup_{\theta \in [-m, m]} R(d_n, \theta) \leq \bar{v}^*(m) + \varepsilon \right) \geq P_n \left(\sup_{\theta \in [-m, m]} |R(d_n, \theta) - R(d_{\pi^*}, \theta)| \leq \varepsilon \right).$$

Equation (28) implies that $\{d_n\}_{n=1}^{\infty}$ is a stochastic minimax sequence.

A.6 Proof of Theorem 4

Take any constant $c > 0$ such that $S \subseteq [-c, c]$. By Lemma 4, for any $y \in \mathbb{R}$

$$|d_n(y) - d_{\pi^*}(y)| \leq \frac{C(y)}{b(y, \pi_n)} W(\pi_n, \pi^*),$$

where $C(y) := C_1(y) + mC_2(y)$ and $b(y, \pi_n) := \int_{[-m, m]} \phi(y; \theta) d\pi_n$. Since $C(\cdot)$ is continuous and S is compact, $\sup_{y \in S} C(y) =: A < \infty$.

Using triangle inequality, for any $y \in [-c, c]$ and for any $\theta \in [-m, m]$

$$\phi(y; \theta) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(y - \theta)^2}{2}\right) \geq \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{(c + m)^2}{2}\right) =: B > 0,$$

Since this holds for every $\theta \in [-m, m]$, integrating over any $\pi \in \Delta([-m, m])$ yields

$$b(y, \pi_n) = \int_{[-m, m]} \phi(y; \theta) d\pi_n \geq B,$$

uniformly over all $y \in S$.

Combining this together,

$$|d_n(y) - d_{\pi^*}(y)| \leq \frac{A}{B} W(\pi_n, \pi^*).$$

Therefore, by Theorem 2, for any $\xi > 0$

$$P_n \left(\sup_{y \in S} |d_n(y) - d_{\pi^*}(y)| > \xi \right) \leq P_n \left(W(\pi_n, \pi^*) > \frac{B}{A} \xi \right) \rightarrow 0.$$

Appendix B. Main Auxiliary Lemmas

B.1 Proof of Lemma 3

The proof of Theorem 1 uses the following auxiliary Lemma.

Lemma 3 (Probability bound) *Let \tilde{g}_j and g_j be defined as in the proof of Theorem 1.*

For any $\delta > 0$ and any $\pi \in \Delta^{I-1}$,

$$\mathbb{P} \left(\frac{1}{J} \sum_{j=1}^J (\tilde{g}_j - g_j)^\top (\pi_j - \pi) < \delta \right) \geq 1 - \exp \left(-\frac{J\delta^2}{128m^4} \right).$$

Proof. For each $j = 1, \dots, J$, define $x_j := \pi_j - \pi$ and

$$\Delta_j := \tilde{g}_j - g_j = \left((d_{\pi_j}(Y_{j,i}) - \theta_i)^2 - \mathbb{E}_{\theta_i} [(d_{\pi_j}(Y_{j,i}) - \theta_i)^2] \right)_{i=1}^I,$$

where $Y_{j,i} \sim \mathcal{N}(\theta_i, 1)$ (independently of π_j). Since $d_{\pi_j}(Y_{j,i}) \in [-m, m]$ and $\theta_i \in [-m, m]$, we have $(d_{\pi_j}(Y_{j,i}) - \theta_i)^2 \in [0, 4m^2]$ for every $i \in \{1, \dots, I\}$. Thus, each component of $\Delta_j = \tilde{g}_j - g_j$ lies in $[-4m^2, 4m^2]$, and so $\|\Delta_j\|_\infty \leq 4m^2$. Let $L := 8m^2$ and let $\{\mathcal{F}_j\}_{j \geq 0}$ be the natural filtration generated by $\{(Y_{j,1}, \dots, Y_{j,I})\}_{j \geq 0}$. Note that under the natural filtration, $\mathbb{E}[\Delta_j | \mathcal{F}_{j-1}] = 0$. Moreover, since $\|x_j\|_1 \leq 2$ and $\|\Delta_j\|_\infty \leq 4m^2$, we have

$$|\langle \Delta_j, x_j \rangle| \leq \|x_j\|_1 \|\Delta_j\|_\infty \leq 8m^2.$$

Hence, conditional on \mathcal{F}_{j-1} , the random variable $\langle \Delta_j, x_j \rangle$ is mean zero and takes values on a subset of the interval $[-8m^2, 8m^2]$. Hoeffding's lemma (Lemma 2.2 in [Boucheron, Lugosi, and Massart \[2013\]](#)) therefore yields, for every $\gamma > 0$,

$$\mathbb{E}[e^{\gamma \langle \Delta_j, x_j \rangle} | \mathcal{F}_{j-1}] \leq \exp \left(\frac{\gamma^2 (16m^2)^2}{8} \right) = \exp \left(\frac{\gamma^2 L^2}{2} \right). \quad (29)$$

We now bound the moment generating function of the average $\frac{1}{J} \sum_{j=1}^J \langle \Delta_j, x_j \rangle$ by the law of iterated expectations. Conditioning on \mathcal{F}_{J-1} and applying (29),

$$\begin{aligned} \mathbb{E} \left[\exp \left(\frac{\gamma}{J} \sum_{j=1}^J \langle \Delta_j, x_j \rangle \right) \right] &= \mathbb{E} \left[\exp \left(\frac{\gamma}{J} \sum_{j=1}^{J-1} \langle \Delta_j, x_j \rangle \right) \cdot \mathbb{E} \left[\exp \left(\frac{\gamma}{J} \langle \Delta_J, x_J \rangle \right) \mid \mathcal{F}_{J-1} \right] \right] \\ &\leq \exp \left(\frac{\gamma^2 L^2}{2J^2} \right) \cdot \mathbb{E} \left[\exp \left(\frac{\gamma}{J} \sum_{j=1}^{J-1} \langle \Delta_j, x_j \rangle \right) \right]. \end{aligned}$$

Repeating the argument sequentially from $j = J$ to $j = 1$ gives

$$\mathbb{E} \left[\exp \left(\frac{\gamma}{J} \sum_{j=1}^J \langle \Delta_j, x_j \rangle \right) \right] \leq \prod_{j=1}^J \exp \left(\frac{\gamma^2 L^2}{2J^2} \right) = \exp \left(\frac{\gamma^2 L^2}{2J} \right).$$

By Markov's inequality, for any $\gamma > 0$ and $\delta > 0$,

$$\begin{aligned} &\mathbb{P} \left(\frac{1}{J} \sum_{j=1}^J \langle \Delta_j, x_j \rangle \geq \delta \right) \leq \exp \left(\frac{\gamma^2 L^2}{2J} - \gamma \delta \right) \\ \implies &\mathbb{P} \left(\frac{1}{J} \sum_{j=1}^J (\tilde{g}_j - g_j)^\top (\pi_j - \pi) < \delta \right) \geq 1 - \exp \left(\frac{\gamma^2 L^2}{2J} - \gamma \delta \right). \end{aligned}$$

Minimizing the exponent over $\gamma > 0$ yields $\gamma^* = \delta J / L^2$, and substituting into the expression above gives the resulting probability bound

$$\exp \left(\frac{\delta^2 J}{2L^2} - \frac{\delta^2 J}{L^2} \right) = \exp \left(-\frac{J\delta^2}{2L^2} \right) = \exp \left(-\frac{J\delta^2}{128m^4} \right).$$

■

B.2 Proof of Lemma 4

The proof of Theorem 3 uses the following auxiliary Lemma.

Lemma 4 *Let $\{\pi_n\}_{n=1}^\infty$ be a stochastic sequence of probability distributions in $\Delta([-m, m])$ such that $W(\pi_n, \pi) \xrightarrow{p} 0$. Let $d_n := d_{\pi_n}$ denote the Bayes rule associated with π_n . Then, for*

any $\xi > 0$

$$P_n \left(\sup_{\theta \in [-m, m]} \int_{-\infty}^{\infty} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy > \xi \right) \rightarrow 0$$

as $n \rightarrow \infty$.

Proof. The proof of Lemma 8 showed that for any $y \in \mathbb{R}$ and any sequence $\{\pi_n\}_{n=1}^{\infty}$

$$|d_{\pi_n}(y) - d_\pi(y)| \leq \frac{1}{b(y, \pi_n)} (C_1(y) + mC_2(y)) W(\pi_n, \pi),$$

where

$$C_1(y) := \frac{1 + m(|y| + m)}{\sqrt{2\pi}}, \quad C_2(y) := \frac{|y| + m}{\sqrt{2\pi}}, \quad b(y, \pi_n) := \int_{[-m, m]} \phi(y; \theta) d\pi_n.$$

Let $C(y) := C_1(y) + mC_2(y)$. Then, we have

$$|d_{\pi_n}(y) - d_\pi(y)| \leq \frac{C(y)}{b(y, \pi_n)} W(\pi_n, \pi). \quad (30)$$

Fix $\alpha \in (0, 1)$ and let $z_{1-\alpha/2}$ be the $1 - (\alpha/2)$ quantile of a standard normal. For any $\theta \in [-m, m]$ we have

$$\begin{aligned} \int_{-\infty}^{\infty} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy &= \int_{-\infty}^{\theta - z_{1-\alpha/2}} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy \\ &+ \int_{\theta - z_{1-\alpha/2}}^{\theta + z_{1-\alpha/2}} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy \\ &+ \int_{\theta + z_{1-\alpha/2}}^{\infty} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy \\ &\leq \int_{\theta - z_{1-\alpha/2}}^{\theta + z_{1-\alpha/2}} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy + 2m\alpha \end{aligned} \quad (31)$$

where the last inequality follows from the fact that $|d_n(y) - d_\pi(y)|$ is bounded by $2m$. Note

also that for any $y \in [\theta - z_{1-\alpha/2}, \theta + z_{1-\alpha/2}]$ and any $\pi \in \Delta([-m, m])$ we have

$$b(y, \pi) \geq C_3(\alpha) := \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}(2m + z_{1-\alpha/2})^2\right).$$

Therefore, (30) and (31) imply

$$\int_{-\infty}^{\infty} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy \leq 2m\alpha + \frac{1}{C_3(\alpha)} W(\pi_n, \pi) \int_{-\infty}^{\infty} C(y) \phi(y; \theta) dy.$$

Since α is arbitrary, $W(\pi_n, \pi) \xrightarrow{P} 0$, and

$$C := \sup_{\theta \in [-m, m]} \int_{-\infty}^{\infty} C(y) \phi(y; \theta) dy < \infty,$$

we conclude that for any $\xi > 0$

$$P_n \left(\sup_{\theta \in [-m, m]} \int_{-\infty}^{\infty} |d_n(y) - d_\pi(y)| \phi(y; \theta) dy > \xi \right) \rightarrow 0$$

as $n \rightarrow \infty$. ■

B.3 Proof of Lemma 5

The following lemma provides a formal statement of the stochastic sequence of invariant estimators result discussed in Section 5.2.

Lemma 5 *Fix $m > 0$ and horizon h . Let $\{d_n\}_{n=1}^{\infty}$ be a stochastic minimax sequence for the bounded normal mean problem:*

$$Y \sim \mathcal{N}(\theta, 1), \quad \theta \in [-m, m].$$

For each n , define an invariant estimator as in Equation (23)

$$\delta_{h,n}(\hat{\theta}_{LP}, \Delta) = \hat{\theta}_{VAR} - \sigma_\Delta d_n(\Delta).$$

Then $\{\delta_{h,n}\}_{n=1}^{\infty}$ is a stochastic minimax sequence for the LP-VAR problem over the class of invariant estimators. That is, for any $\xi > 0$,

$$P_n \left(\sup_{\theta_h, |b| \leq m \cdot \sigma_{\Delta}} \mathcal{R}(\delta_{h,n}, \theta_h, b) > \bar{v}_h^*(m; \Sigma) + \xi \right) \rightarrow 0, \text{ as } n \rightarrow \infty,$$

where

$$\bar{v}_h^*(m; \Sigma) := \sigma_{\text{VAR}}^2 + \sigma_{\Delta}^2 \bar{v}^*(m)$$

Proof. Let $\theta := \frac{b}{\sigma_{\Delta}}$. By the maintained bound $|b| \leq m \cdot \sigma_{\Delta}$, we have $\theta \in [-m, m]$.

By definition of Δ ,

$$\hat{\theta}_{\text{VAR}} = \hat{\theta}_{\text{LP}} + \sigma_{\Delta} \Delta.$$

Therefore,

$$\delta_{h,n}(\hat{\theta}_{\text{LP}}, \Delta) = \hat{\theta}_{\text{LP}} + \sigma_{\Delta} \Delta - \sigma_{\Delta} d_n(\Delta).$$

Thus, in the notation of invariant estimator,

$$\delta_{h,n}(0, \Delta) = \sigma_{\Delta} \Delta - \sigma_{\Delta} d_n(\Delta).$$

Using Equation (21), we obtain

$$\begin{aligned} \mathcal{R}(\delta_{h,n}, \theta_h, b) &= \sigma_{\text{VAR}}^2 + \sigma_{\Delta}^2 \mathbb{E}_{b/\sigma_{\Delta}} \left[\left(\left[\Delta - \frac{\delta_{h,n}(0, \Delta)}{\sigma_{\Delta}} \right] - \frac{b}{\sigma_{\Delta}} \right)^2 \right] \\ &= \sigma_{\text{VAR}}^2 + \sigma_{\Delta}^2 \mathbb{E}_{\theta} [(d_n(\Delta) - \theta)^2]. \end{aligned}$$

The second equality follows because

$$\Delta - \frac{\delta_{h,n}(0, \Delta)}{\sigma_{\Delta}} = \Delta - \frac{\sigma_{\Delta} \Delta - \sigma_{\Delta} d_n(\Delta)}{\sigma_{\Delta}} = d_n(\Delta).$$

Since $\Delta \sim \mathcal{N}(\theta, 1)$, the last expectation is exactly the risk of d_n in the bounded normal

mean problem:

$$\mathbb{E}_\theta [(d_n(\Delta) - \theta)^2] = R(d_n, \theta).$$

Hence,

$$\mathcal{R}(\delta_{h,n}, \theta_h, b) = \sigma_{\text{VAR}}^2 + \sigma_\Delta^2 R(d_n, \theta).$$

Taking supremum over θ_h and $|b| \leq m\sigma_\Delta$ gives

$$\sup_{\theta_h, |b| \leq m\sigma_\Delta} \mathcal{R}(\delta_{h,n}, \theta_h, b) = \sigma_{\text{VAR}}^2 + \sigma_\Delta^2 \sup_{\theta \in [-m, m]} R(d_n, \theta),$$

using the risk of invariant estimators does not depend on θ_h after the reduction.

Therefore, for any $\xi > 0$,

$$\begin{aligned} & P_n \left(\sup_{\theta_h, |b| \leq m\sigma_\Delta} \mathcal{R}(\delta_{h,n}, \theta_h, b) > \bar{v}_h^*(m; \Sigma) + \xi \right) \\ &= P_n \left(\sigma_{\text{VAR}}^2 + \sigma_\Delta^2 \sup_{\theta \in [-m, m]} R(d_n, \theta) > \sigma_{\text{VAR}}^2 + \sigma_\Delta^2 \bar{v}^*(m) + \xi \right) \\ &= P_n \left(\sup_{\theta \in [-m, m]} R(d_n, \theta) > \bar{v}^*(m) + \frac{\xi}{\sigma_\Delta^2} \right). \end{aligned}$$

Since $\{d_n\}_{n=1}^\infty$ is a stochastic minimax sequence for the bounded normal mean problem, the last probability converges to zero. Hence

$$P_n \left(\sup_{\theta_h, |b| \leq m\sigma_\Delta} \mathcal{R}(\delta_{h,n}, \theta_h, b) > \bar{v}_h^*(m; \Sigma) + \xi \right) \rightarrow 0.$$

■

Appendix C. Mathematical Properties of the Maximin Problem

In this section, we focus on the problem of finding a least-favorable distribution π^* ; that is, a distribution π^* that solves (6). To this purpose, it will be convenient to define the function $f : \Delta([-m, m]) \rightarrow \mathbb{R}_+$ given by

$$f(\pi) := \inf_{d \in \mathcal{D}} r(d, \pi) = \inf_{d \in \mathcal{D}} \int_{[-m, m]} R(d, \theta) d\pi. \quad (32)$$

Lemma 6 below shows that the function $f(\cdot)$ is Lipschitz with respect to the 1-Wasserstein metric over $\Delta([-m, m])$.

Lemma 6 (The optimized Bayes Risk is Lipschitz in π) *Suppose that for any $\pi \in \Delta([-m, m])$ there exists a Bayes rule d_π as defined in (5). Then, for any $\pi, \pi' \in \Delta([-m, m])$*

$$|f(\pi) - f(\pi')| \leq (3m^2 + 4m) \cdot W(\pi, \pi').$$

Proof. Let π, π' be two arbitrary elements of $\Delta([-m, m])$. Without loss of generality, let us assume that $f(\pi) \geq f(\pi')$. Let $L(m) := 3m^2 + 4m$ be the Lipschitz constant of Lemma 1. Then

$$\begin{aligned} 0 \leq f(\pi) - f(\pi') &= \int_{[-m, m]} R(d_\pi, \theta) d\pi - \int_{[-m, m]} R(d_{\pi'}, \theta) d\pi' \\ &\leq \int_{[-m, m]} R(d_{\pi'}, \theta) d\pi - \int_{[-m, m]} R(d_{\pi'}, \theta) d\pi' \\ &\quad \text{(by definition of } d_\pi\text{)} \\ &= L(m) \left(\int_{[-m, m]} \frac{R(d_{\pi'}, \theta)}{L(m)} d\pi - \int_{[-m, m]} \frac{R(d_{\pi'}, \theta)}{L(m)} d\pi' \right) \\ &\leq L(m)W(\pi, \pi'), \end{aligned}$$

where the last inequality follows from (13) and the fact that, according to Lemma 1,

$$\left\| \frac{R(d_{\pi'}, \theta)}{L(m)} \right\|_{\text{Lip}} = \frac{1}{L(m)} \sup_{\theta \neq \theta'} \frac{|R(d_{\pi'}, \theta) - R(d_{\pi'}, \theta')|}{|\theta - \theta'|} \leq 1.$$

■

Lemma 7 below further shows that the optimized Bayes Risk is a concave function of π and shows that the risk function can be viewed as a supergradient.

Lemma 7 (The optimized Bayes Risk is Concave in π) *Let π, π' be two arbitrary elements of $\Delta([-m, m])$. For any $\lambda \in [0, 1]$, $\lambda\pi + (1 - \lambda)\pi'$*

$$f(\lambda\pi + (1 - \lambda)\pi') \geq \lambda f(\pi) + (1 - \lambda)f(\pi').$$

Moreover, if the Bayes rule $d_{\pi'}$ exists:

$$f(\pi') \geq f(\pi) + \left(\int_{[-m, m]} R(d_{\pi'}, \theta) d\pi' - \int_{[-m, m]} R(d_{\pi'}, \theta) d\pi \right)$$

for any $\pi \in \Delta([-m, m])$.

Proof. Let π, π' be two arbitrary elements of $\Delta([-m, m])$. For any $\lambda \in [0, 1]$, $\lambda\pi + (1 - \lambda)\pi'$ is also an element of $\Delta([-m, m])$. By definition

$$\begin{aligned} f(\lambda\pi + (1 - \lambda)\pi') &= \inf_{d \in \mathcal{D}} \int_{[-m, m]} R(d, \theta) d(\lambda\pi + (1 - \lambda)\pi') \\ &= \inf_{d \in \mathcal{D}} \left(\lambda \int_{[-m, m]} R(d, \theta) d\pi + (1 - \lambda) \int_{[-m, m]} R(d, \theta) d\pi' \right) \\ &\geq \lambda \inf_{d \in \mathcal{D}} \int_{[-m, m]} R(d, \theta) d\pi + (1 - \lambda) \inf_{d \in \mathcal{D}} \int_{[-m, m]} R(d, \theta) d\pi' \\ &= \lambda f(\pi) + (1 - \lambda)f(\pi'). \end{aligned}$$

To show the last part of the lemma, note that

$$\begin{aligned}
f(\pi') &= \int_{[-m,m]} R(d_{\pi'}, \theta) d\pi' \\
&= \left(\int_{[-m,m]} R(d_{\pi'}, \theta) d\pi' - \int_{[-m,m]} R(d_{\pi'}, \theta) d\pi \right) + \int_{[-m,m]} R(d_{\pi'}, \theta) d\pi \\
&\geq \left(\int_{[-m,m]} R(d_{\pi'}, \theta) d\pi' - \int_{[-m,m]} R(d_{\pi'}, \theta) d\pi \right) + \inf_{d \in \mathcal{D}} \int_{[-m,m]} R(d, \theta) d\pi \\
&= \left(\int_{[-m,m]} R(d_{\pi'}, \theta) d\pi' - \int_{[-m,m]} R(d_{\pi'}, \theta) d\pi \right) + f(\pi).
\end{aligned}$$

■

Lemma 6 and Lemma 7 show that the problem of finding a least-favorable distribution is equivalent to the problem of maximizing a concave, Lipschitz function over the space of all Borel probability measures over $[-m, m]$. This is an infinite-dimensional concave maximization problem over $\Delta([-m, m])$.

Lemma 8 (Sequential continuity of $d_\pi(y)$ at π w.r.t. to the 1-Wasserstein Metric)

Let $\pi \in \Delta([-m, m])$. Take any sequence π_n such that

$$W(\pi_n, \pi) \xrightarrow{p} 0.$$

Then,

$$|d_{\pi_n}(y) - d_\pi(y)| \xrightarrow{p} 0$$

for almost every $y \in \mathbb{R}$.

Proof. The Bayes rule $d_\pi(\cdot)$ for the bounded normal mean problem is the posterior mean function:

$$d_\pi(y) := \frac{a(y, \pi)}{b(y, \pi)},$$

where

$$a(y, \pi) := \int_{[-m,m]} \theta \phi(y; \theta) d\pi, \quad b(y, \pi) := \int_{[-m,m]} \phi(y; \theta) d\pi,$$

and $\phi(y; \theta)$ is the p.d.f. of a $\mathcal{N}(\theta, 1)$ evaluated at y . Algebra shows that for any $y \in \mathbb{R}$, and any priors $\pi_n, \pi \in \Delta([-m, m])$

$$\begin{aligned}
d_{\pi_n}(y) - d_{\pi}(y) &= \frac{a(y, \pi_n)}{b(y, \pi_n)} - \frac{a(y, \pi)}{b(y, \pi)} \\
&= \frac{a(y, \pi_n)b(y, \pi) - a(y, \pi)b(y, \pi_n)}{b(y, \pi_n)b(y, \pi)} \\
&= \frac{a(y, \pi_n)b(y, \pi) - a(y, \pi)b(y, \pi) + a(y, \pi)b(y, \pi) - a(y, \pi)b(y, \pi_n)}{b(y, \pi_n)b(y, \pi)} \\
&= \frac{[a(y, \pi_n) - a(y, \pi)]b(y, \pi) + [b(y, \pi) - b(y, \pi_n)]a(y, \pi)}{b(y, \pi_n)b(y, \pi)} \\
&= \frac{1}{b(y, \pi_n)}[a(y, \pi_n) - a(y, \pi)] + d_{\pi}(y) \frac{1}{b(y, \pi_n)}[b(y, \pi) - b(y, \pi_n)].
\end{aligned}$$

Note that the functions

$$\varphi_1(\theta; y) := \theta\phi(y; \theta) \quad \text{and} \quad \varphi_2(\theta; y) := \phi(y; \theta),$$

are both Lipschitz continuous over $\theta \in [-m, m]$ with constants

$$C_1 := \frac{1 + m(|y| + m)}{\sqrt{2\pi}}, \quad C_2 := \frac{|y| + m}{\sqrt{2\pi}},$$

respectively. This means that, using the definition of the 1-Wasserstein metric in (13),

$$\begin{aligned}
|d_{\pi_n}(y) - d_{\pi}(y)| &\leq \frac{1}{b(y, \pi_n)} (C_1 + mC_2) W(\pi_n, \pi) \\
&= \frac{1}{b(y, \pi_n) - b(y, \pi) + b(y, \pi)} (C_1 + mC_2) W(\pi_n, \pi) \\
&\leq \frac{1}{-C_2W(\pi_n, \pi) + b(y, \pi)} (C_1 + mC_2) W(\pi_n, \pi).
\end{aligned}$$

Since $b(y, \pi) > 0$ for any $y \in \mathbb{R}$, this means that for any sequence π_n such that $W(\pi_n, \pi) \xrightarrow{P} 0$, we have that

$$|d_{\pi_n}(y) - d_{\pi}(y)| \xrightarrow{P} 0$$

for almost every $y \in \mathbb{R}$. ■

Appendix D. Notation

Table 2: Notation

Symbol	Description
Statistical model and estimation	
m	Bound on the parameter space, $m > 0$.
θ	Unknown mean, $\theta \in [-m, m]$.
Y	Observed data, $Y \sim \mathcal{N}(\theta, 1)$.
$\mathcal{L}(a, \theta)$	Loss function, $(a - \theta)^2$.
\mathcal{D}	Class of all measurable decision rules $d: \mathcal{Y} \rightarrow [-m, m]$.
$R(d, \theta)$	Risk of decision rule d at θ : $\mathbb{E}_\theta[(d(Y) - \theta)^2]$.
$\phi(y; \theta)$	Density of $\mathcal{N}(\theta, 1)$ evaluated at y .
Minimax and Bayes framework	
$\Delta([-m, m])$	Set of all Borel probability measures on $[-m, m]$.
π	A prior distribution, $\pi \in \Delta([-m, m])$.
$r(d, \pi)$	Bayes risk of d with respect to π : $\int R(d, \theta) d\pi$.
d_π	Bayes rule with respect to prior π (posterior mean).
$f(\pi)$	Optimized Bayes Risk: $\inf_{d \in \mathcal{D}} r(d, \pi)$.
π^*	Least-favorable distribution for the original problem.
$\bar{v}^*(m)$	Minimax value: $\inf_{d \in \mathcal{D}} \sup_\theta R(d, \theta)$.
$v^*(m)$	Maximin value: $\sup_\pi \inf_{d \in \mathcal{D}} r(d, \pi)$.
$\bar{v}_{\text{linear}}(m)$	Minimax risk over linear estimators: $m^2/(1 + m^2)$.
Discretization	
I	Number of grid points in $[-m, m]$.
θ_i	i -th grid point: $-m + 2m(i - 1)/(I - 1)$.
$\Delta(m, I)$	Grid spacing: $2m/(I - 1)$.
Δ^{I-1}	$(I-1)$ -dimensional probability simplex.
$v(m; I)$	Value of the discretized maximin problem.
$I(\epsilon)$	Grid size: $1 + (3m^3 + 4m^2)/(\epsilon/2)$.
Stochastic mirror ascent algorithm	
ϵ	Tolerance (target approximation accuracy).
M	Upper bound on the risk function: $4m^2$.
J	Number of epochs: $2M^2 \log(I)/(\epsilon/2)^2$.
η	Step size (learning rate): $(\epsilon/2)/M^2$.
$w_{i,j}$	Unnormalized weight on grid point θ_i at epoch j .
π_j	Prior at epoch j : normalized weights $w_{i,j-1}/\sum_k w_{k,j-1}$.
g_j	Supergradient at epoch j : $(R(d_{\pi_j}, \theta_1), \dots, R(d_{\pi_j}, \theta_I))'$.
\tilde{g}_j	Stochastic (Monte Carlo) estimate of g_j .
$y_{i,j}$	Monte Carlo draw: $y_{i,j} \sim \mathcal{N}(\theta_i, 1)$.
π^ϵ	Averaged prior over J iterations: $J^{-1} \sum_{j=1}^J \pi_j$.
LP-VAR application	
h	Horizon of the impulse response.
θ_h	Impulse response coefficient of interest at horizon h .
$\hat{\theta}_{\text{LP}}$	Local projection estimator of θ_h .
$\hat{\theta}_{\text{VAR}}$	VAR estimator of θ_h .
b	Bias of the VAR estimator.
Σ	Covariance matrix of $(\hat{\theta}_{\text{LP}}, \hat{\theta}_{\text{VAR}})$.
σ_{LP}^2	Variance of the LP estimator.
σ_{VAR}^2	Variance of the VAR estimator.
Δ	Standardized Difference between the VAR and LP estimators.
σ_Δ^2	Variance of $(\hat{\theta}_{\text{VAR}} - \hat{\theta}_{\text{LP}})$ (equal to $\sigma_{\text{LP}}^2 - \sigma_{\text{VAR}}^2$).
$\mathcal{R}(d_h, \theta_h, b)$	Risk of an invariant LP-VAR estimator.
$\bar{v}_h^*(m; \Sigma)$	Minimax value in the LP-VAR problem: $\sigma_{\text{VAR}}^2 + \sigma_\Delta^2 \bar{v}^*(m)$.

Appendix E. Additional Simulation Results

E.1 Theoretical vs empirical Lipschitz constants

Our theoretical results establish that for any decision rule d taking values in $[-m, m]$, the risk function is Lipschitz on $[-m, m]$ with constant $L(m) = 3m^2 + 4m$, that is,

$$|R(d, \theta) - R(d, \theta')| \leq (3m^2 + 4m)|\theta - \theta'| \quad \forall \theta, \theta' \in [-m, m].$$

This bound is uniform over all admissible decision rules and is used to control the difference between the worst-case risk over the full interval and the maximum risk over a discretized m grid.

To assess the conservativeness of this bound, we also report an empirical estimate of this quantity evaluated as

$$\hat{L}_{\text{grid}}(m) = \max_{i < j} \frac{|\hat{R}(d_{\pi^\epsilon}, \theta_i) - \hat{R}(d_{\pi^\epsilon}, \theta_j)|}{|\theta_i - \theta_j|},$$

where $\hat{R}(d_{\pi^\epsilon}, \theta_i)$ is the Monte Carlo estimate of the risk at the grid point θ_i . Figure 14 illustrates this comparison.




Figure 14: Comparison of theoretical and empirical Lipschitz constants.

E.2 A certified upper bound on the worst-case risk

The worst-case risk evaluated on the finite grid is a priori lower than the true maximal value, since the supremum may be attained between grid points. To certify the gap, we use Lipschitz continuity of the risk function.

For any $\theta \in [\theta_i, \theta_{i+1}]$, applying Lemma 1 at each endpoint and averaging gives

$$R(d_{\pi^\epsilon}, \theta) \leq \frac{1}{2} \left[R(d_{\pi^\epsilon}, \theta_i) + (3m^2 + 4m)|\theta - \theta_i| \right] + \frac{1}{2} \left[R(d_{\pi^\epsilon}, \theta_{i+1}) + (3m^2 + 4m)|\theta - \theta_{i+1}| \right].$$

Since $|\theta - \theta_i| + |\theta - \theta_{i+1}| = \Delta(m, I)$, taking supremum over $\theta \in [-m, m]$

$$\sup_{\theta \in [-m, m]} R(d_{\pi^\epsilon}, \theta) \leq \min \left\{ \max_{i \in \{1, \dots, I-1\}} \frac{R(d_{\pi^\epsilon}, \theta_i) + R(d_{\pi^\epsilon}, \theta_{i+1})}{2} + \frac{(3m^2 + 4m)\Delta(m, I)}{2}, 4m^2 \right\},$$

where $\Delta(m, I) = 2m/(I - 1)$ is the grid spacing.

Figure 15 applies analogous Lipschitz-certification to truncated maximum likelihood estimator $d_{\text{trunc}}(y) = \max\{\min\{y, m\}, -m\}$ (green line) and *clipped* linear estimator (purple line) $d_{\text{linear, clipped}} = \max\{\min\{c \cdot y, m\}, -m\}$, where $c = m^2/(1 + m^2)$. All risks are evaluated with 1,000 Monte Carlo draws over the 200,000-point grid.




Figure 15: Worst-case risk of d_{π^ϵ} (certified bound) vs. alternative estimators and analytical bounds.

E.3 Comparison with Casella and Strawderman [1981] closed-form for $m \leq 1$

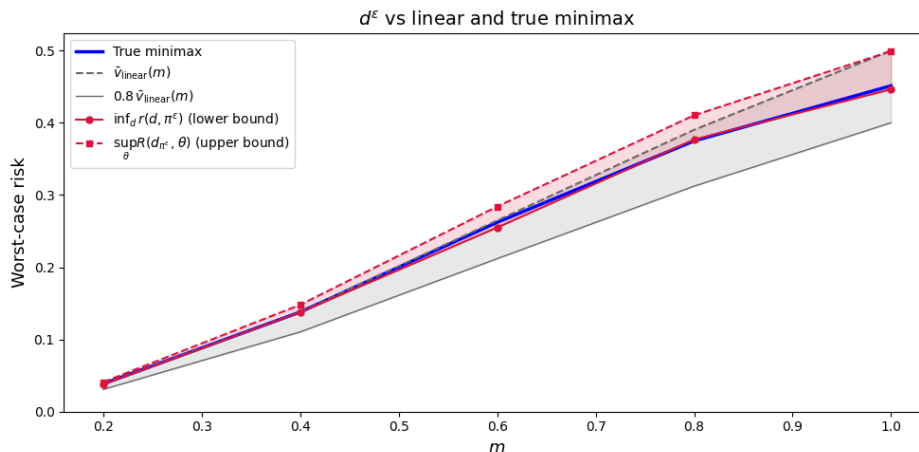


Figure 16: d_{π^ϵ} vs linear vs true minimax.

E.4 Stability of the risk evaluation under Monte Carlo randomness

One concern with the estimator d_{π^ϵ} is that—due to randomness introduced by the stochastic mirror ascent procedure—different runs of the algorithm could generate very different estimators. One way to assuage this concern is to report the behavior of the risk function across different seeds. Figure 17 displays the risk $R(d_{\pi^\epsilon}, \theta)$ for different seeds evaluated for $\theta \in [-2, 2]$. The shaded region in the graph represents the area between the 2.5% and the 97.5% quantiles across different seeds. The solid curve shows the across-seed mean. Each profile is estimated by Monte Carlo over 100 independent seeds with 1,000 draws per seed.




Figure 17: Risk of $d_{\pi^{\epsilon}}$ for $m = 2$.