

YOU'VE GOT TO BE EFFICIENT: AMBIGUITY, MISSPECIFICATION AND VARIATIONAL PREFERENCES

KARUN ADUSUMILLI[†]

ABSTRACT. This article introduces a framework for evaluating statistical decisions under both prior ambiguity and likelihood misspecification. We begin with an ambiguity set — a frequentist model that pairs a possibly misspecified likelihood with every possible prior — and uniformly expand it by a Kullback–Leibler radius to accommodate likelihood misspecification. We show that optimal decisions under this framework are equivalent to minimax decisions with an exponentially tilted loss function. Misspecification manifests as an exponential tilting of the loss, while ambiguity corresponds to a search for the least favorable prior. This separation between ambiguity and misspecification enables local asymptotic analysis under global misspecification, achieved by localizing the priors alone. Remarkably, for both estimation and treatment assignment, we show that optimal decisions coincide with those under correct specification, regardless of the degree of misspecification. These results extend to semi-parametric models. As a practical consequence, our findings imply that practitioners should prefer maximum likelihood over the simulated method of moments, and efficient GMM estimators — such as two-step GMM — over diagonally weighted alternatives.

This version: April 10, 2026

I would like to thank Xu Cheng, Frank Diebold, Wayne Gao, George Mailath, and seminar participants at the University of Pennsylvania for valuable discussions and comments that substantially improved this article.

[†]Department of Economics, University of Pennsylvania.

1. INTRODUCTION

Box (1976) famously observed that all models are wrong, since they are necessarily approximations of reality. Any researcher or decision-maker who relies on a statistical model to learn about a parameter of interest must therefore contend with the possibility that the likelihood is misspecified. At the same time, researchers are often unable or unwilling to commit to a single prior over the parameter. In practice, then, decision-makers confront both prior ambiguity and likelihood misspecification.

This article introduces a framework for evaluating statistical decisions under both sources of concern. Following Bayesian practice, we define a statistical model as a joint distribution comprising a prior and a likelihood. We argue that both components are necessary because they capture fundamentally different types of uncertainty. The prior encodes epistemic uncertainty — subjective uncertainty arising from incomplete knowledge about the parameter of interest — while the likelihood captures aleatoric uncertainty — the objective randomness inherent in any statistical experiment.

To account for prior ambiguity, we define an ambiguity set: a frequentist model that pairs a possibly misspecified likelihood with every possible prior. Following Cerreia-Vioglio et al. (2026), we then uniformly expand this set by a Kullback–Leibler radius to accommodate likelihood misspecification. The optimal decision rule is defined as the one that achieves the lowest expected loss under the worst-case model from this expanded set.

We show that optimal decisions under this formulation are equivalent to minimax decisions with an exponentially tilted loss function. Likelihood misspecification manifests as an exponential tilting of the loss, while prior ambiguity corresponds to a search for the least favorable prior. Our framework thus enables a clean separation between ambiguity and misspecification. Furthermore, when there is no fear misspecification, the optimal decisions reduce to the standard Wald (1950) formulation of minimax decisions under ambiguity alone — the formulation underlying most of frequentist analysis.

This separation between ambiguity and misspecification also enables us to develop a local asymptotic theory under global misspecification, achieved by localizing the priors around a reference parameter. Under mild conditions, the finite-sample likelihoods, which may themselves be misspecified, can be replaced by a limit experiment involving the Gaussian family as the reference likelihood. While a substantial literature studies local asymptotics under local misspecification, where misspecification typically manifests as an added bias in the Gaussian limit, our framework permits the Gaussian family itself to be globally misspecified in the limit experiment, thereby accommodating much richer classes of misspecification.

Local asymptotics also simplifies the search for optimal decisions, as these are considerably easier to characterize in the limit experiment. Quite remarkably, we find that for estimation and treatment assignment problems, optimal decisions coincide with those under correct specification, regardless of the degree of global misspecification. For these problems, it is therefore always optimal for the decision-maker to proceed as if the likelihood were correctly specified and select the resulting optimal decision rule. Intuitively, these results arise because misspecification under our formulation is symmetric around the reference Gaussian likelihood. Since the estimation and treatment-assignment loss functions are also symmetric around the parameter of interest, any estimator that is not efficient under correct specification would break this symmetry. Because nature chooses the least favorable likelihood specification given the decision-maker's choice of estimator, departures from symmetry necessarily incur higher decision risk.

We extend our local asymptotic theory to semi-parametric models, and show that the above results on optimal decisions apply to that setting as well.

Our findings have a number of practical consequences. In applications, researchers often employ inefficient estimators over efficient ones, a practice frequently justified on the grounds that under misspecification, no estimand recovers the precise parameter of economic interest (Andrews et al., 2025). Our results, however, suggest that this reasoning is incomplete. While the parameter of interest cannot be recovered with certainty under misspecification, our decision-theoretic analysis shows that efficient estimators under correct specification also deliver the lowest decision risk under arbitrary misspecification. In the case of parametric

models, these results suggest that practitioners should prefer maximum likelihood over the simulated method of moments, irrespective of the degree of misspecification. Similarly, in the context of GMM, practitioners should prefer efficient estimation methods, such as two-step GMM, over diagonally weighted or inefficient alternatives. Misspecification concerns alone cannot justify the use of the inefficient estimators over parametrically or semi-parametrically efficient alternatives.

1.1. Related literature. This article relates to an extensive literature on ambiguity and misspecification spanning economics, statistics, and computer science. A detailed comparison of our approach with alternative decision-theoretic frameworks is deferred to Section 2.4. Here, we restrict ourselves to a broad survey of the literature on ambiguity and misspecification.

The analysis of optimal decisions under prior ambiguity originates with Wald (1950). A substantial body of work in statistics has extended this framework to local asymptotics; we refer to Ibragimov and Hasminskii (1981); Le Cam (1986); van der Vaart and Wellner (1996); Van der Vaart (2000) for textbook treatments. A central result from this literature is that semi-parametrically efficient estimators are asymptotically minimax optimal under prior ambiguity.

The literature on misspecification is equally extensive. Huber (1964) proposes a contamination model to address likelihood misspecification. Hansen and Sargent (2011) develop an approach that involves selecting a worst-case likelihood from an ambiguity set defined by surrounding a reference or approximate likelihood with a Kullback–Leibler divergence ball of finite radius. The related field of Distributionally Robust Optimization (DRO) takes the reference distribution to be the empirical distribution $\hat{\mathbb{P}}$ of the data in the experiment, and employs more general measures of distance from $\hat{\mathbb{P}}$ to define ambiguity sets — including ϕ -divergence measures (e.g., reverse KL divergence and total-variation distance), Wasserstein distances and Levy-Prokhorov distances. We refer to Ben-Tal et al. (2009) for a textbook treatment of DRO, and to Rahimian and Mehrotra (2019), Kuhn et al. (2025) for recent surveys. These methods do not account for prior ambiguity and, consequently, do not reduce to the standard minimax formulation that underpins

frequentist analysis in the absence of misspecification concerns. Our approach instead follows the recent work of Cerreia-Vioglio et al. (2026) by first defining an ambiguity set to address prior ambiguity and then uniformly expanding this entire set by a KL divergence radius to accommodate likelihood misspecification. In contrast to the results from DRO, we find that optimal estimation and treatment assignment are invariant to the degree of misspecification.

This article adopts a decision-theoretic approach to ambiguity and misspecification, which is closely related to the literature in economic theory on variational preferences, as expounded in Maccheroni et al. (2006) and Cerreia-Vioglio et al. (2026). The econometrics literature has also studied alternative, non-decision-theoretic approaches to misspecification, and we refer to Armstrong (2025) for a recent survey. For instance, White (1982) defines pseudo-parameters as the probability limits of estimators, and views them as suitably defined approximations to the underlying parameter of interest. The partial identification approach of Manski (2003) proposes set-identifying the parameter under misspecification, while Masten and Poirier (2021) develop methods for sensitivity analysis. A limitation of these approaches relative to the decision-theoretic framework is that they do not directly identify the optimal statistical decision that a decision-maker should employ.

2. DECISION-MAKING UNDER AMBIGUITY AND MISSPECIFICATION

2.1. An illustrative example. To illustrate our formalism, we introduce the following running example. A decision-maker, Alice, is tasked with determining whether a drug should be approved for use in the US population. She is therefore interested in learning about the parameter $\theta \in \Theta$, defined as the average population treatment effect. We assume binary outcomes, so that the population outcome distribution is Bernoulli(θ).

To assess the drug's efficacy, the pharmaceutical company has conducted a randomized controlled trial with n observations. Given the observed data $\mathbf{x} \in \mathcal{X}$ from the trial, Alice seeks to choose a decision $\delta : \mathcal{X} \rightarrow \mathcal{A}$ so as to maximize her utility $u_n(\theta, \delta)$, or equivalently, minimize the loss function $l_n(\theta, \delta) = -u_n(\theta, \delta)$.

Examples of loss functions include the estimation loss,

$$l_n(\theta, \delta) = \ell(\sqrt{n}(\theta - \delta)), \quad \delta \in \Theta,$$

for some bowl-shaped function $\ell(\cdot)$, e.g., $\ell(z) = z^2$ for mean squared error, and the treatment-assignment loss,

$$l_n(\theta, \delta) = \sqrt{n}(\theta \mathbf{1}\{\theta \geq 0\} - \theta \delta), \quad \delta \in \{0, 1\}.$$

Under the estimation loss, the goal is to learn directly about the parameter θ , whereas under the treatment-assignment loss, the goal is to either approve ($\delta = 1$) or reject ($\delta = 0$) the drug for use in the entire population.

Unfortunately for Alice, the trial was conducted exclusively in the state of Pennsylvania. Because the drug is novel, she has no formal basis for judging whether, or to what degree, treatment effects observed in Pennsylvania are representative of those in the broader US population. This gives rise to model misspecification concerns. At the same time, Alice also faces ambiguity concerns, as she is unable to form an initial prior over θ . We now describe a formalism that accommodates both.

2.2. Bayesian, Frequentist and Misspecified models.

2.2.1. *(Bayesian) Statistical models.* We begin by formally defining the notion of a statistical model in the absence of ambiguity or misspecification concerns.

Since the loss function takes the form $l_n(\theta, \delta(\mathbf{x}))$, the payoff-relevant state of the world is given by $\omega = (\theta, \mathbf{x})$: an oracle who knows ω would recover Alice's loss with certainty. Following the framework of Savage or Anscombe–Aumann (Anscombe and Aumann, 1963), we define a model $m \equiv m(\theta, \mathbf{x})$ as a probability distribution over the payoff-relevant state $\omega = (\theta, \mathbf{x})$. This distribution admits a natural decomposition into a prior and a likelihood:

$$m(\theta, \mathbf{x}) = \pi(\theta) \otimes p_\theta(\mathbf{x}). \tag{2.1}$$

Here, $\pi(\theta)$ denotes the posited prior, the marginal distribution over θ , while $p_\theta(\mathbf{x}) = p(\mathbf{x} \mid \theta)$ denotes the posited likelihood, the conditional distribution of \mathbf{x}

given θ . Equation (2.1) is nothing more than the definition of a Bayesian statistical model; see Robert (2007, Definition 1.2.1).

The decomposition of a model into a prior and a likelihood is a canonical feature of Bayesian decision-making. Given the importance of this decomposition for what follows, it is worth understanding why both components are necessary. As we argue below, they capture fundamentally different sources of uncertainty: epistemic and aleatoric. In the terminology of Anscombe and Aumann (1963), these correspond to the uncertainties involved in horse gambles and roulette wheels.

Epistemic uncertainty refers to uncertainty arising from a lack of knowledge — uncertainty that can, in principle, be reduced through the acquisition of additional data or evidence. In the Anscombe–Aumann framework, this is the uncertainty of a horse gamble. Because the parameter θ enters Alice’s loss function directly, it is natural to regard it as a quantity that exists in principle but that Alice does not know. The prior $\pi(\theta)$ thus encodes Alice’s epistemic uncertainty due to her imperfect knowledge of θ . Crucially, Alice can conceptualize θ independently of the likelihood. She may, for instance, have access to prior information — such as data from related studies — that enables her to form a prior π without reference to whatever experiment the pharmaceutical company may have conducted.

Aleatoric uncertainty, by contrast, refers to inherent randomness, which is implicit in the design of any statistical experiment. The likelihood $p_\theta(\mathbf{x})$ captures precisely this source of uncertainty. In conducting the trial, the pharmaceutical company presumably drew a random sample of n observations from the population of Pennsylvania. This sampling procedure requires the use of an implicit or explicit random number generator and therefore introduces genuine randomness; in the Anscombe–Aumann framework, this is uncertainty generated by roulette wheels. The likelihood thus describes the distribution of the data \mathbf{x} induced by this randomness, for any given value of θ .

Importantly, in our framework, the likelihood does not rise to the status of a model. It provides only a mapping from the parameter θ to the distribution of \mathbf{x} . Because θ enters Alice’s loss function directly, knowledge of the correct likelihood would not enable Alice to obtain a probabilistic forecast of her loss, as she would still face epistemic uncertainty over θ .

2.2.2. *Models with prior ambiguity, aka Frequentist models.* We now incorporate prior ambiguity into our framework. Suppose that Alice is unable to form a single prior, perhaps because she is ambiguity-averse in the sense of Maccheroni et al. (2006). Instead, she posits a structured set of models,

$$\mathcal{Q} := \left\{ \pi(\theta) \otimes p_\theta(\mathbf{x}) : \pi \in \Delta(\Theta) \right\},$$

where $\Delta(\Theta)$ denotes the set of all probability distributions over θ , while continuing to treat the likelihood $p_\theta(\mathbf{x})$ as correctly specified. In the spirit of Wald (1950), Alice could then choose the decision rule that performs best against the worst-case model in \mathcal{Q} — effectively, the one associated with the least favorable prior — thereby guarding against prior ambiguity:

$$\delta_{n,f}^* := \arg \min_{\delta} \left[\sup_{m \in \mathcal{Q}} \mathbb{E}_m [l_n(\theta, \delta)] \right].$$

Because the Wald approach underpins much of frequentist analysis, we refer to \mathcal{Q} as a frequentist model and to $\delta_{n,f}^*$ as a frequentist (or minimax) decision rule.

2.2.3. *Models with prior ambiguity and likelihood misspecification.* Now suppose that Alice entertains the possibility that the likelihood $p_\theta(x)$ employed in her frequentist model \mathcal{Q} may not be correctly specified. Likelihood misspecification can arise in two distinct ways. The first is misspecification of functional form, e.g., specifying a Gaussian likelihood when the true data-generating process follows a t -distribution. In our running example, however, the outcomes are Bernoulli, so the likelihood is necessarily a product of Bernoulli distributions and functional-form misspecification is not a concern. In fact, as we show in Section 2.3, functional-form misspecification can always be ameliorated by making the likelihood sufficiently flexible. The misspecification that Alice confronts here is of the second kind: the link between the parameter of interest θ and the distribution over x may be incorrectly specified. In the running example, the true outcome distribution is $Y \sim \text{Bernoulli}(\theta_P)$, where θ_P is the average treatment effect in Pennsylvania. Likelihood misspecification arises because, in general, $\theta \neq \theta_P$.

To accommodate misspecification concerns, we follow Cerreia-Vioglio et al. (2026) and place a protective belt around the frequentist model \mathcal{Q} . Formally,

we posit a set of unstructured models of the form

$$\mathcal{M} = \left\{ m \in \Delta(\Theta, \mathcal{X}) : \min_{q \in \mathcal{Q}} R_q(m) \leq K \right\},$$

where $R_q(m) = \text{KL}(m \parallel q)$ denotes the Kullback–Leibler divergence and K is a constant reflecting Alice’s degree of ambiguity aversion. Alice could then choose the decision rule that performs best against the worst-case model in \mathcal{M} , thereby guarding against both prior ambiguity and misspecification:

$$\delta_n^* := \arg \min_{\delta} \left[\sup_{m \in \mathcal{M}} \mathbb{E}_m [l_n(\theta, \delta)] \right]. \quad (2.2)$$

For the remainder of this article, we decompose any generic model m as $m(\theta, \mathbf{x}) = \pi(\theta) \otimes m_{\theta}(\mathbf{x})$, and reserve the notation $p_{\theta}(\mathbf{x})$ for a reference likelihood specification, which may itself be misspecified. Define $R_{\mathcal{Q}}(m) := \min_{q \in \mathcal{Q}} R_q(m)$. Straightforward algebra yields

$$R_{\mathcal{Q}}(m) = \int \text{KL}(m_{\theta}(\cdot) \parallel p_{\theta}(\cdot)) d\pi(\theta), \quad (2.3)$$

so that the set of unstructured models can be equivalently written as

$$\mathcal{M} = \left\{ \pi(\theta) \otimes m_{\theta}(x) : \int \text{KL}(m_{\theta}(\cdot) \parallel p_{\theta}(\cdot)) d\pi(\theta) \leq K, \pi \in \Delta(\Theta) \right\}.$$

The class of unstructured models thus comprises every possible prior π , paired with all likelihoods $m_{\theta}(\cdot)$ satisfying the integrated KL constraint.

Note that \mathcal{M} expands \mathcal{Q} by adding a protective radius of KL divergence of size K . Because \mathcal{Q} already accommodates unrestricted prior ambiguity, this protective belt serves entirely to guard against likelihood misspecification. Within the class of alternative models, however, the prior and likelihood may interact in subtle ways: the constraint permits $m_{\theta}(\cdot)$ to deviate substantially from $p_{\theta}(\cdot)$ for certain values of θ , provided the associated prior π places low weight on those values.

To further understand this interaction between the prior and likelihood, it is instructive to see how likelihood misspecification would be modeled in the absence of prior ambiguity. If Alice were able to commit to a single prior π , then \mathcal{Q} would be a singleton and \mathcal{M} would consist of all models $\pi(\theta) \otimes m_{\theta}(\mathbf{x})$ such that

$m_\theta(\cdot) \in \mathcal{M}_{x|\theta}(\pi)$, where

$$\mathcal{M}_{x|\theta}(\pi) := \left\{ m_\theta(\cdot) : \int \text{KL}(m_\theta(\cdot) \parallel p_\theta(\cdot)) d\pi(\theta) \leq K \right\}.$$

When the prior is fixed, the class of candidate likelihoods thus depends on π ; the prior shapes which deviations from the reference likelihood are admissible. In the general case with unrestricted priors, \mathcal{M} can be interpreted as the union of $\pi(\theta) \otimes \mathcal{M}_{x|\theta}(\pi)$ over all $\pi \in \Delta(\Theta)$.

2.3. Nuisance and structural parameters. A nuisance parameter is an unknown quantity that enters the likelihood but does not affect the loss function. In our running example, if the outcomes are distributed as $\mathcal{N}(\mu, \sigma^2)$ but Alice is solely interested in learning about the mean treatment effect μ , then σ^2 is a nuisance parameter. Nuisance parameters make the likelihood more flexible and can be used to ameliorate functional-form misspecification; indeed, one can allow for nonparametric specifications by making the nuisance parameters infinite-dimensional. As noted earlier, however, nuisance parameters cannot address the second type of misspecification concerning the link between the parameter of interest and the data. Even if Alice were to adopt a fully nonparametric specification of the outcome distribution, she would still face misspecification concerns, as treatment effects in Pennsylvania may differ fundamentally from those in the broader US population.

In contrast, we refer to the parameters that enter the utility function directly as structural parameters. In what follows, θ denotes the full collection of unknown parameters, which may include both structural and nuisance components. The structural parameters are modeled as known functions $\mu(\theta)$ of θ . With this notation, the loss functions introduced earlier take the more general form

$$l_n(\theta, \delta) = \ell(\sqrt{n}(\mu(\theta) - \delta)), \quad \delta \in \mu(\Theta),$$

for estimation loss, and

$$l_n(\theta, \delta) = \sqrt{n}(\mu(\theta) \mathbf{1}\{\mu(\theta) \geq 0\} - \mu(\theta) \delta), \quad \delta \in \{0, 1\},$$

for treatment-assignment loss.

The definitions of Bayesian, frequentist, and misspecified models remain unchanged; the introduction of nuisance parameters affects only the form of the loss functions.

2.4. Alternative approaches to misspecification: A comparison. Andrews et al. (2025) define an econometric model $(\theta, p_\theta(\cdot))$ as a combination of the parameter and the likelihood.¹ Apart from the prior over θ , this coincides with the definition of a Bayesian statistical model. Introducing the prior allows us to account for prior ambiguity, which, as we have seen, plays a key role even in the standard frequentist approach.

A growing recent literature has considered accounting for misspecification through partial identification; see, e.g., Ishihara and Kitagawa (2021); Yata (2021); Christensen et al. (2022); Montiel Olea et al. (2026). This literature supposes that the parameter of interest θ lies within a bounded distance, $d(\theta, \theta_P) \leq L$, of an identifiable parameter θ_P . In our running example, this would require Alice to assume that the population treatment effect θ differs from the treatment effect in Pennsylvania θ_P by at most L . While bounds of this form can arise naturally in a number of applications, the approach falls short as a general framework for misspecification for several reasons.

First, unlike our formalism, bounding the parameters directly lacks an axiomatic justification. Second, the approach is sensitive to the choice of the distance measure $d(\cdot)$, which in turn makes L difficult to calibrate. In the Bernoulli setting, for instance, it would not be reasonable to use Euclidean distance $d(\theta, \theta_P) = |\theta - \theta_P|$, as it does not respect the constraint $\theta \in [0, 1]$. Third, and most importantly, imposing a uniform bound $d(\theta, \theta_P) \leq L$ for all θ_P implies comparisons across different values of θ_P that may be at odds with the decision-maker’s actual preferences over ambiguity and misspecification. To see why, note that there is always epistemic uncertainty over the value of θ_P . There is no a priori reason to believe that the bound $d(\theta, \theta_P) \leq L$ provides the same degree of protection against misspecification when $\theta_P = 0.9$ as when $\theta_P = 0.1$. Depending on Alice’s preferences, e.g., her loss function, she may be less concerned about misspecification at high values of

¹In the terminology of Andrews et al. (2025), the likelihood is referred to as a data-generating process.

θ_P (which suggests the treatment is highly effective) than at low values. The constraint $d(\theta, \theta_P) \leq L$ is not directly linked to her attitudes toward misspecification; it is a constraint on parameters, not on payoff-relevant quantities.

In more closely related work, Andrews et al. (2020), Bonhomme and Weidner (2022) and Christensen and Connault (2023) characterize misspecification through a statistical distance $d(m_\theta(\cdot), p_\theta(\cdot))$, e.g., KL divergence, over likelihoods. The specific setups and goals of these works differ substantially from our own.² Quite apart from this, however, a uniform bound on the KL divergence over likelihoods, of the form $\sup_\theta d(m_\theta(\cdot), p_\theta(\cdot)) \leq L$, is subject to the same criticism as a bound $d(\theta, \theta_P) \leq L$ over parameters: there is no a priori reason to believe that it provides the same degree of protection against misspecification across different values of θ . As before, Alice may be less concerned about misspecification at high values of θ than at low values. Furthermore, quantities such as KL divergence are sensitive to the choice of the reference measure $p_\theta(\cdot)$, and KL divergences evaluated at different parameter values such as $R_{p_{\theta_1}}(m_{\theta_1})$ and $R_{p_{\theta_2}}(m_{\theta_2})$ are not directly comparable.

Our formulation avoids this problem because we first postulate an infinite-dimensional ambiguity set \mathcal{Q} and then uniformly expand it by a KL radius K . As in Cerreia-Vioglio et al. (2026), the value of K can be tied to the decision-maker’s underlying preferences over ambiguity and misspecification. But rather remarkably, as it turns out, our optimal decisions are asymptotically independent of the choice of K .

3. CHARACTERIZING OPTIMAL DECISIONS

We now characterize optimal decisions under ambiguity and misspecification. For reasons that will become apparent shortly, it is convenient to start with utility maximization rather than loss minimization. Following the framework described in Section 2.2.3, the optimal decision rule takes the form:

$$\delta_n^* := \arg \max_{\delta} \left[\inf_{m \in \mathcal{M}} \mathbb{E}_m [u_n(\theta, \delta)] \right] = \arg \max_{\delta} \inf_m \left\{ \mathbb{E}_m [u_n(\theta, \delta)] : R_{\mathcal{Q}}(m) \leq K \right\}.$$

²Andrews et al. (2020) study the relationship between descriptive statistics and structural parameters. Bonhomme and Weidner (2022) analyze estimation under local misspecification, i.e., when $d(m_\theta(\cdot), p_\theta(\cdot)) \rightarrow 0$. Christensen and Connault (2023) study partial identification of θ . These goals are all distinct from ours: devising optimal decisions under global ambiguity and misspecification.

The function $R_{\mathcal{Q}}(\cdot) : \Delta(\Theta \times \mathcal{X}) \rightarrow \mathbb{R}$ is strictly convex, so we may apply a minimax theorem to show that for each K there exists a multiplier λ such that³

$$\begin{aligned} \delta_n^* &= \arg \max_{\delta} V_n(\delta), \text{ where} \\ V_n(\delta) &:= \inf_m \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda R_{\mathcal{Q}}(m) \right\}. \end{aligned}$$

Following Cerreia-Vioglio et al. (2026), we refer to $V_n(\cdot)$ as the variational decision criterion. In fact, the variational criterion has a direct axiomatic justification through preferences, as in Cerreia-Vioglio et al. (2026), more so than the constraint version (2.2), which is only used as a motivation for the optimal decision.

Recalling the definition $R_{\mathcal{Q}}(m) := \min_{q \in \mathcal{Q}} R_q(m)$ and interchanging the the order of the $\min_{q \in \mathcal{Q}}$ and \inf_m operations, we can write

$$V_n(\delta) = \min_{q \in \mathcal{Q}} \inf_m \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda R_q(m) \right\}.$$

The Donsker-Varadhan variational formula yields

$$\inf_m \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda R_q(m) \right\} = -\lambda \ln \mathbb{E}_q[e^{-u_n(\theta, \delta)/\lambda}]. \quad (3.1)$$

Converting back to the loss function via $l_n(\theta, \delta) = -u_n(\theta, \delta)$, we obtain

$$V_n(\delta) = -\lambda \ln \left\{ \max_{q \in \mathcal{Q}} \mathbb{E}_q[e^{l_n(\theta, \delta)/\lambda}] \right\}, \quad (3.2)$$

and consequently, the optimal decision can be characterized as

$$\delta_n^* = \arg \min_{\delta} \max_{q \in \mathcal{Q}} \mathbb{E}_q[e^{l_n(\theta, \delta)/\lambda}]. \quad (3.3)$$

So far, the calculations above follow Cerreia-Vioglio et al. (2026). However, due to the special structure of \mathcal{Q} in our setting—comprising all possible priors paired

³Formally, we suppose that $\Theta \times \mathcal{X}$ is compact. Then, an application of Sion's minimax theorem for each δ shows that

$$\begin{aligned} & \max_{\delta} \inf_{m \in \Delta(\Theta \times \mathcal{X})} \left\{ \mathbb{E}_m[u_n(\theta, \delta)] : R_{\mathcal{Q}}(m) \leq K \right\} \\ &= \max_{\delta} \inf_{m \in \Delta(\Theta \times \mathcal{X})} \max_{\lambda \geq 0} \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda(R_{\mathcal{Q}}(m) - K) \right\} \\ &= \max_{\delta} \max_{\lambda \geq 0} \inf_{m \in \Delta(\Theta \times \mathcal{X})} \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda(R_{\mathcal{Q}}(m) - K) \right\}. \end{aligned}$$

The claim then follows by interchanging the max operations.

with the reference likelihood $p_\theta(x)$ —we can simplify (3.3) further:

$$\delta_n^* = \arg \min_{\delta} \max_{\pi \in \Delta(\Theta)} \int \mathbb{E}_{p(x|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right] d\pi(\theta). \quad (3.4)$$

In this expression, the quantity $R_n(\theta, \delta) := \mathbb{E}_{p(x|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right]$ admits a natural interpretation as the frequentist risk of the decision rule δ under the reference likelihood $p(x|\theta)$, evaluated with respect to the exponentiated loss $e^{l_n(\theta, \delta)/\lambda}$.

Notice that (3.4) corresponds to a standard minimax decision framework under exponentiated loss: we can interpret the optimal decision as the result of a two-player game in which nature chooses the least favorable prior while the decision-maker chooses the optimal rule. Equation (3.4) is therefore a key result of this article. It establishes that optimal decisions under both ambiguity and misspecification are equivalent to optimal decisions under ambiguity alone, but with an exponentiated loss function. The result also reveals how the two sources of concern separate naturally. The effect of misspecification is to transform the loss function into an exponentiated version; intuitively, the decision-maker magnifies the impact of large losses while attenuating the impact of small ones. The effect of ambiguity, as in the setting without misspecification, manifests in the search for the least favorable prior.

4. LOCAL ASYMPTOTICS WITH GLOBAL MISSPECIFICATION

It is rarely feasible to solve the minimax problem (3.4) exactly in finite samples. Instead, as is standard even in classical frequentist (i.e., minimax) settings, we turn to local asymptotic approximations.

Following the usual approach, we fix a reference parameter θ_0 and consider local perturbations of the form $\theta_0 + h/\sqrt{n}$. Priors $\pi(\theta)$ over θ are then mapped to local priors $\pi(h)$ over h .

In the case of treatment-assignment loss, $l_n(\theta, \delta) = n(\mu(\theta) \mathbf{1}\{\mu(\theta) \geq 0\} - \mu(\theta) \delta)$, local asymptotics arise naturally from the global minimax problem (3.4). The key observation is that the least favorable prior concentrates its mass on regions where the treatment effect $\mu(\theta)$ is of order $1/\sqrt{n}$. When $\mu(\theta)$ is of a higher order of magnitude than $1/\sqrt{n}$, determining the optimal assignment is asymptotically trivial. Conversely, when $\mu(\theta)$ is of a lower order of magnitude than $1/\sqrt{n}$,

the difference between treatment and status quo is negligible, so the loss is close to zero regardless of the choice of δ . It is therefore natural to choose a reference parameter θ_0 satisfying $\mu(\theta_0) = 0$, since the least favorable prior would concentrate around this value in any case. As Hirano and Porter (2009) showed, these same considerations apply to treatment-assignment problems in the absence of misspecification as well.

But how should one choose a reference θ_0 for estimation loss, and what is the meaning of local asymptotics in this setting? We offer two interpretations.

The first is that local asymptotics amounts to localizing prior ambiguity around a reference parameter θ_0 that the decision-maker believes is close to the true value. In our running example, Alice may have a priori reason to believe that the true treatment effect lies near θ_0 , even if she is uncertain about its exact value. It would then be natural to restrict her ambiguity set to a $1/\sqrt{n}$ neighborhood of θ_0 . Because this prior information is obtained independently of the data, likelihood misspecification does not affect the choice of θ_0 . Consequently, we can localize the priors, even as we can — and do — allow for global misspecification of the likelihood.

Under the second interpretation, the choice of the reference θ_0 is itself subject to adversarial optimization. The basic idea, following Ibragimov and Hasminskii (1981), is to decompose the global minimax problem into two stages: first, fix a reference θ_0 and evaluate the local minimax performance of a decision rule δ against local alternatives of the form $\theta_0 + h/\sqrt{n}$; then, in an outer step, select the least favorable reference θ_0 .

Here, we focus primarily on a theoretical development of the first interpretation. The second interpretation is detailed in Section 4.5, while the theory is developed in Appendix B.

4.1. Parametric models: Setup. We assume that the data consists of an i.i.d collection of outcomes $\mathbf{x} := \{Y_i\}_{i=1}^n$. Under the reference likelihood, Y_i is distributed as P_θ . Let ν denote a dominating measure for $\{P_\theta : \theta \in \mathbb{R}^d\}$, and set $p_\theta := dP_\theta/d\nu$. We require the reference class of likelihoods, $\{P_\theta\}_\theta$, to be quadratic mean differentiable (qmd):

Assumption 1. The class $\{P_\theta : \theta \in \mathbb{R}^d\}$ is qmd around θ_0 , i.e., there exists a score function $\psi(\cdot)$ such that for each $h \in \mathbb{R}^d$,

$$\int \left[\sqrt{p_{\theta_0+h}} - \sqrt{p_{\theta_0}} - \frac{1}{2} h^\top \psi \sqrt{p_{\theta_0}} \right]^2 d\nu = o(|h|^2).$$

Furthermore, the information matrix $I_0 := \mathbb{E}_{\theta_0}[\psi\psi^\top]$ is invertible.

In the illustrative example, the outcomes are modeled as Bernoulli, so Assumption 1 holds with $\psi(y) = (\theta_0(1 - \theta_0))^{-1} (y - \theta_0)$. More broadly, this assumption is rather mild and satisfied for almost all commonly used likelihood models, including the Normal, Cauchy, Exponential, and Poisson distributions. It is important to bear in mind that Assumption 1 constrains the reference class of likelihoods, p_θ , not the actual likelihoods, which are unknown.

We also assume that the function $\mu(\theta)$, which maps θ to structural parameters, satisfies a mild differentiability condition:

Assumption 2. There exists $\dot{\mu}_0 \in \mathbb{R}^d$ and $\epsilon_n \rightarrow 0$ independent of h such that $\sqrt{n}(\mu(\theta_0 + h/\sqrt{n}) - \mu(\theta_0)) = \dot{\mu}_0^\top h + \epsilon_n |h|^2$ for all bounded h .

Let $P_{n,h}$ denote the joint probability measure over the iid Y_1, \dots, Y_n when each $Y_i \sim P_{\theta_0+h/\sqrt{n}}$, and let $\mathbb{E}_{n,h}[\cdot]$ denote the corresponding expectation. Under local asymptotics, the minimal risk attained under the minimax problem (3.4) can be written as:

$$V_n^* = \min_{\delta} \max_{\pi(h)} \int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0+h/\sqrt{n},\delta)/\lambda} \right] d\pi(h). \quad (4.1)$$

4.2. Limit approximations and the Gaussian limit experiment. Define the standardized score statistic as

$$x_n = \frac{I_0^{-1/2}}{\sqrt{n}} \sum_{i=1}^n \psi(Y_i).$$

It is well known, see e.g., Van der Vaart (2000, Chapter 7), that quadratic mean differentiability (Assumption 1) implies $\mathbb{E}_{n,0}[\psi(Y_i)] = 0$. Then, by the central limit theorem,

$$x_n \xrightarrow[P_{n,0}]{d} x \sim \mathcal{N}(0, I). \quad (4.2)$$

Assumption 1 also implies the important property of Local Asymptotic Normality (LAN; Van der Vaart, 2000, Chapter 7):

$$\ln \frac{dP_{n,\theta_0+h/\sqrt{n}}}{dP_{n,\theta_0}} = h^\top I_0^{1/2} x_n - \frac{1}{2} h^\top I_0 h + o_{P_{n,0}}(1), \text{ uniformly over bounded } h. \quad (4.3)$$

Consider now a limit experiment in which the decision-maker observes a d -dimensional signal x , posited to be drawn from a reference Gaussian likelihood, $P_h(x) \sim \mathcal{N}(I_0^{-1/2}h, I)$. By the properties of the Gaussian distribution,

$$\ln \frac{dP_h}{dP_0} = h^\top I_0^{1/2} x - \frac{1}{2} h^\top I_0 h.$$

It follows from (4.2) and (4.3) that the reference likelihood ratios in the finite-sample experiment converge to their counterparts in the limit experiment:

$$\ln \frac{dP_{n,\theta_0+h/\sqrt{n}}}{dP_{n,\theta_0}} \xrightarrow{P_{n,0}} \ln \frac{dP_h}{dP_0}, \text{ for each } h.$$

Furthermore, Assumption 2 implies that the loss functions admit asymptotic approximations. For estimation-loss, defining $\tilde{\delta}_n = \sqrt{n}\dot{\mu}_0^\top(\delta_n - \theta_0)$ and assuming $\tilde{\delta}_n$ has a weak limit $\tilde{\delta}$, we have

$$l_n(\theta_0 + h/\sqrt{n}, \delta_n) \equiv \ell\left(\sqrt{n}\left(\mu(\theta_0 + h/\sqrt{n}) - \delta_n\right)\right) \rightsquigarrow \ell(\dot{\mu}_0^\top h - \tilde{\delta}), \quad (4.4)$$

where ‘ \rightsquigarrow ’ represents weak convergence. For treatment-assignment loss, since the reference parameter satisfies $\mu(\theta_0) = 0$, Assumption 2 implies

$$l_n(\theta_0 + h/\sqrt{n}, a) \rightarrow \dot{\mu}_0^\top h \mathbf{1}\{\dot{\mu}_0^\top h \geq 0\} - (\dot{\mu}_0^\top h)a, \text{ uniformly over } a \in \{0, 1\} \text{ and bounded } h. \quad (4.5)$$

Convergence of likelihood ratios implies asymptotic equivalence between the actual and limit experiments in the sense of Le Cam (1986). Combined with the loss function approximations above, this suggests that the minimax value V_n^* in (4.1) should converge to the minimax value V^* in the limit experiment, where

$$V^* := \min_{\tilde{\delta}} \max_{\pi(h)} \int \mathbb{E}_h \left[e^{l(h,\tilde{\delta})/\lambda} \right] d\pi(h), \text{ with} \quad (4.6)$$

$$l(h, \tilde{\delta}) = \begin{cases} \ell(\dot{\mu}_0^\top h - \tilde{\delta}) & \text{for estimation loss,} \\ \dot{\mu}_0^\top h \{ \mathbf{1}\{\dot{\mu}_0^\top h \geq 0\} - \tilde{\delta} \} & \text{for treatment-assignment loss.} \end{cases}$$

Formal statements to this effect are provided in Section 4.4.

It is instructive to compare our asymptotic approach with the more traditional analysis of local asymptotics under local misspecification. In locally misspecified models, the KL divergence between the true and reference likelihoods is assumed to decline at a $1/n$ rate. Consequently, as highlighted in Andrews et al. (2020), Bonhomme and Weidner (2022) and Müller and Norets (2024), local misspecification manifests as asymptotic bias in the Gaussian limit experiment. Our framework differs fundamentally since it permits the Gaussian likelihood approximation itself to be globally misspecified. This is possible because our asymptotic theory approximates only the reference finite-sample likelihood ratios with Gaussian likelihoods; it makes no claim about the convergence of the true likelihood ratios. Global misspecification consequently manifests not as bias in the Gaussian limit, but as an exponential tilting of the loss function.

Equation (4.6) suggests that asymptotically optimal decision rules can be derived by solving the minimax problem in the limit experiment and mapping the solutions back to the finite-sample setting. Since optimal decisions are considerably easier to characterize under Gaussian likelihoods, this reduction illustrates the key benefit of the local asymptotic approach.

4.3. Characterization of optimal decisions in the limit experiment. We begin with estimation-loss. Since $\ell(\cdot)$ is bowl-shaped, so is $e^{\ell(\cdot)/\lambda}$. It then follows from Anderson’s lemma, see e.g., Van der Vaart (2000, Proposition 8.6), that the minimax-optimal estimator in the limit experiment — the solution to (4.6) — is simply

$$\tilde{\delta}^* = \mu_0^\top I_0^{-1/2} x.$$

Remarkably, $\tilde{\delta}^*$ is independent of λ , which governs the degree of misspecification. In fact, $\tilde{\delta}^*$ coincides with most efficient estimator, the maximum likelihood estimator, under correct misspecification, which corresponds to $\lambda = \infty$. In other words, the optimal estimator under ambiguity and misspecification is identical to the optimal estimator under ambiguity alone.

For treatment-assignment loss, Anderson’s lemma does not apply. Nevertheless, as the following proposition shows, the optimal decision rule again takes a simple

form: it recommends treatment whenever the MLE of the treatment effect under correct specification is positive.

Proposition 1. *The minimax-optimal decision rule in the limit experiment under the treatment assignment loss is $\tilde{\delta}^* = \mathbf{1}\{\hat{\mu}_0^\top I_0^{-1/2} x \geq 0\}$. The corresponding least-favorable prior is a symmetric two-point prior supported on $(-h^*, h^*)$, with $h^* := \frac{\Delta^*}{\hat{\mu}^\top I_0^{-1} \hat{\mu}} I_0^{-1} \hat{\mu}$, and*

$$\Delta^* = \arg \max_{\Delta \geq 0} \left\{ \left(e^{\frac{\Delta}{\lambda}} - 1 \right) \Phi(-\Delta) \right\}.$$

As with the optimal estimator, the optimal treatment-assignment rule is independent of the degree of misspecification.

Intuitively, these results arise because misspecification under our formulation is unstructured and therefore symmetric around the reference Gaussian likelihood. Since the loss functions are also symmetric around the reference θ_0 , any estimator that is not efficient under correct specification would break this symmetry. Because nature chooses the least favorable likelihood specification given the decision-maker's choice of estimator, departures from symmetry would necessarily incur higher decision risk. It is therefore always optimal for the decision-maker to proceed as if the likelihood were correctly specified and select the resulting optimal decision rule.

4.4. Formal results. We now formally establish the asymptotic equivalence of experiments through two results. First, we show that the minimax value V^* in the limit experiment forms an asymptotic lower bound on the sequence of optimal decision risks in the finite-sample experiments. Second, we show that plug-in versions of the optimal limit-experiment decision rules $\tilde{\delta}^*$ – obtained by replacing $I_0^{-1/2} x$ with the finite sample MLE $\hat{\theta}_{\text{mle}}$ – are asymptotically optimal, in the sense that their decision-risks converge to V^* . Specifically, we argue that the asymptotically decisions are given by

$$\hat{\delta}_n^* = \begin{cases} \mu(\hat{\theta}_{\text{mle}}) & \text{for estimation,} \\ \mathbf{1}\{\mu(\hat{\theta}_{\text{mle}}) \geq 0\} & \text{for treatment assignment.} \end{cases} \quad (4.7)$$

As discussed at the beginning of this section our formal results require localization of ambiguity. This involves restricting attention to the set of compactly

supported priors $\Delta_M(\mathcal{H}) \equiv \{\pi(h) : \text{supp}(\pi) \in [-M, M]\}$ for some $M < \infty$. This restriction is not needed for the lower bound, but plays a role in establishing that the bound is attained by the plug-in rules.

Theorem 1. *(Lower bound) Suppose that Assumptions 1 and 2 hold. Then, under both the estimation and treatment-assignment loss functions,*

$$\liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h/\sqrt{n}, \delta)/\lambda} \right] d\pi(h) \geq V^*.$$

We place a mild regularity condition on the MLE:

Assumption 3. *The maximum-likelihood estimator $\hat{\theta}_{mle}$ admits a locally linear score-function approximation:*

$$\hat{\theta}_{mle} - \theta_0 = I_0^{-1/2} x_n + o_{P_{n,0}}(1).$$

To avoid technical issues relating to the existence of moments for the estimation problem, our theory also requires that $\ell(\cdot)$ be bounded. We state this as an additional assumption.

Assumption 4. *The function $\ell(\cdot)$ is bounded.*

Assumption 4 implies the estimation loss $l_n(\theta, \delta)$ is bounded. As for the treatment-assignment loss, since we work with compact priors, it is automatically bounded.

Theorem 2. *(Asymptotic optimality of plug-in rules) Suppose that Assumptions 1-4 hold. Then, under both the estimation and treatment-assignment loss functions,*

$$\lim_{M \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] d\pi(h) = V^*.$$

In fact, the MLE can be replaced by any asymptotically efficient estimator satisfying Assumption 3. All such estimators attain the same minimax lower bound V^* ; our theory therefore does not distinguish among them.

The statements of Theorems 1 and 2 are new even in the absence of misspecification. Standard local asymptotic minimax theorems are typically stated in terms of a maximum over discrete sets of h values, rather than the $\max_{\pi(h) \in \Delta_M(\mathcal{H})}$ operation employed here, see, e.g., Van der Vaart (2000, Proposition 8.11). This distinction is consequential for our setting, as the discrete formulation does not lend itself

to a natural interpretation of prior ambiguity. Our results avoid this limitation through a different method of proof that directly accommodates optimization over a compact space of local priors.

4.5. Non-local priors. The formal results above were derived under priors localized around a reference parameter θ_0 . As noted earlier, with a slight strengthening of the assumptions, we can allow for global priors over a compact set Θ ; essentially, we require the assumptions to be valid uniformly over $\theta_0 \in \Theta$. In this case, the lower bounds on decision risk correspond to the least favorable choice of the reference.

Formally, we show that under both loss functions,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_n(\theta,\delta)/\lambda} \right] d\pi(\theta) \\ & \geq \begin{cases} \sup_{\theta_0 \in \Theta} V_{\theta_0}^* & \text{for estimation loss,} \\ \sup_{\{\theta_0 \in \Theta: \mu(\theta_0)=0\}} V_{\theta_0}^* & \text{for treatment-assignment loss,} \end{cases} \end{aligned}$$

where $V_{\theta_0}^*$ represents the minimal decision-risk in the limit experiment (as defined in Section 4.2) evaluated at a given reference parameter θ_0 . The formal statement, together with the required assumptions, is provided in Appendix B.

In the same appendix, we show that the decisions in (4.7) are asymptotically optimal under global priors as well, in that they attain this bound:

$$\begin{aligned} & \lim_{K \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta,\delta)/\lambda} \right] d\pi(\theta) \\ & = \begin{cases} \sup_{\theta_0 \in \Theta} V_{\theta_0}^* & \text{for estimation loss,} \\ \sup_{\{\theta_0 \in \Theta: \mu(\theta_0)=0\}} V_{\theta_0}^* & \text{for treatment-assignment loss,} \end{cases} \end{aligned}$$

where $l_{n,K}(\cdot) := K \wedge l_n(\cdot)$ denotes the loss function truncated at level K .

To the best of our knowledge, these results appear new even in the context of no misspecification concerns. A number of authors, e.g., Hirano (2025, Section 3.2), have raised concerns about the interpretation of asymptotic analysis under local reparameterizations. Our results address these concerns by showing that global minimax analysis is equivalent to local asymptotic analysis with a least favorable choice of reference.

4.6. Application: Maximum likelihood vs Simulated method of moments. As noted by Andrews et al. (2025), researchers in applied work often prefer the simulated method of moments — which involves subjective selection of moment functions — over maximum likelihood, despite the latter’s superior efficiency. This preference is frequently justified by the argument that “with misspecification concerns, moment estimators are often more reliable” (Bordalo et al., 2020).

Our results suggest, however, that this reasoning is incomplete. Under our framework, there is no tradeoff between misspecification robustness and efficiency. Efficient estimators are misspecification-robust to an arbitrary degree, and inefficient estimators remain suboptimal even in the presence of potential misspecification. Misspecification concerns alone therefore cannot justify the use of the simulated method of moments over maximum likelihood.

This is not to say that the simulated method of moments should never be used. The selection of moment functions can often be understood as a form of model selection. For instance, if the parameter of interest is the treatment effect on a subgroup of the population, it may be reasonable to selectively overweight the relevant portion of the sample. This, however, is a problem of model selection rather than model misspecification per se.

5. SEMI-PARAMETRIC MODELS

The preceding sections focused on parametric models for the likelihood. In many applications, however, the exact distributional form is left unspecified, motivating the use of semi-parametric models. As discussed in 2.3, our methodology extends naturally to the semi-parametric and nonparametric settings by allowing for infinite-dimensional nuisance parameters.

In semi-parametric models, the structural parameter is typically a regular functional $\mu := \mu(P)$, of the unknown population distribution P . Common examples of regular functionals include the mean, median, and quantiles. For simplicity, we assume that μ is scalar-valued. The loss functions are then

$$l_n(P, \delta) = \ell\left(\sqrt{n}(\mu(P) - \delta)\right),$$

for estimation loss, and

$$l_n(P, \delta) = \sqrt{n}(\mu(P) \mathbf{1}\{\mu(P) \geq 0\} - \mu(P) \delta), \quad \delta \in \{0, 1\},$$

for treatment-assignment loss. The population distribution P plays the same role as θ in the parametric analysis. While P is unknown, we are only interested in its scalar functional $\mu(P)$; the remaining features of P are treated as an infinite dimensional nuisance parameter.

As in Section 2.2, we consider a framework in which the decision-maker confronts both prior ambiguity — an inability to form a single prior over the parameter P — and model misspecification — the concern that the population distribution P may not coincide with the outcome distribution \hat{P} in the experimental sample. In our running example, Alice may adopt a fully nonparametric specification of the outcome distribution yet still face misspecification concerns, as treatment effects in Pennsylvania could differ fundamentally from those in the broader US population. More generally, misspecification also arises when the semi-parametric model imposes restrictions that are not satisfied by the combination of the true μ and the sample outcome distribution \hat{P} . For instance, in the GMM framework, the researcher specifies a moment condition $\mathbb{E}_P[m(Y_i, \mu)] = 0$, where $m(\cdot) \in \mathbb{R}^p$ is a known vector of moment functions, $\mu \in \mathbb{R}^d$ is the structural parameter of interest, and P is the unknown population distribution. Misspecification arises if this restriction does not hold in the sample, i.e., $\mathbb{E}_{\hat{P}}[m(Y_i, \mu)] \neq 0$ for the true μ . In our running example, which fits within the GMM framework with $m(Y_i, \mu) = Y_i - \mu$, Alice may be concerned that $\mathbb{E}_{\hat{P}}[Y_i - \mu] \neq 0$.

The central message of our semi-parametric results is that the sample analogs of the optimal decisions characterized in Section 4.3 remain optimal in the semi-parametric setting. In essence, one replaces the score statistic from the parametric setting with an efficient influence function process associated with the functional of interest. For example, if the goal is to conduct inference on the mean, one replaces x_n with the cumulative sum of outcomes $n^{-1/2} \sum_{i=1}^n (Y_i/\sigma)$, where $\sigma^2 := \text{Var}[Y_i]$.

5.1. Local asymptotics for semi-parametric models. It is easiest to discuss ambiguity and misspecification in semi-parametric settings using local asymptotics. Our local asymptotic analysis employs the formalism of Van der Vaart (2000, Section 25.3). Let \mathcal{P} denote the class of candidate population distributions with bounded variance, dominated by some measure ν . We fix a reference distribution $P_0 \in \mathcal{P}$, and surround it with various smooth one-dimensional parametric sub-models, $\{P_{s,h} : s \leq \zeta\}$ for some $\zeta > 0$, whose score function is $h(\cdot)$, and that pass through P_0 at $s = 0$ (i.e., $P_{0,h} = P_0$). Formally, these sub-models satisfy

$$\int \left[\frac{dP_{s,h}^{1/2} - dP_0^{1/2}}{s} - \frac{1}{2} h dP_0^{1/2} \right]^2 d\nu \rightarrow 0 \text{ as } s \rightarrow 0. \quad (5.1)$$

As in the parametric setting, for the treatment-assignment problem, we choose P_0 such that $\mu(P_0) = 0$.

By Van der Vaart (2000, Lemma 25.14), condition (5.1) implies $\int h dP_0 = 0$ and $\int h^2 dP_0 < \infty$. The set of all such functions h is termed the tangent space $T(P_0)$, which is a subset of the Hilbert space $L^2(P_0)$ endowed with the inner product $\langle f, g \rangle = \mathbb{E}_{P_0}[fg]$ and norm $\|f\| = \mathbb{E}_{P_0}[f^2]^{1/2}$. For any $h \in T(P_0)$, let $P_{n,h}$ denote the joint probability measure over Y_1, \dots, Y_n , when each Y_i is an iid draw from $P_{1/\sqrt{n},h}$, and let $\mathbb{E}_{n,h}[\cdot]$ denote the corresponding expectation. An important consequence of (5.1) is the LAN property:

$$\sum_{i=1}^n \ln \frac{dP_{1/\sqrt{n},h}}{dP_0}(Y_i) = \frac{1}{\sqrt{n}} \sum_{i=1}^n h(Y_i) - \frac{1}{2} \|h\|^2 + o_{P_{n,0}}(1), \text{ uniformly over bounded } \|h\|. \quad (5.2)$$

Let $\psi \in T(P_0)$ denote the efficient influence function corresponding to μ , defined by the property that for any $h \in T(P_0)$,

$$\frac{\mu(P_{s,h}) - \mu(P_0)}{s} - \langle \psi, h \rangle = o(s). \quad (5.3)$$

Set $\sigma^2 = \mathbb{E}_{P_0}[\psi^2]$. The semi-parametric analogue of the score statistic in the semi-parametric setting is the standardized efficient influence function process

$$x_n := \frac{\sigma^{-1}}{\sqrt{n}} \sum_{i=1}^n \psi(Y_i).$$

Any element $h \in T(P_0)$ admits the orthogonal decomposition $h = \langle \psi/\sigma, h \rangle \psi/\sigma + \tilde{h}$, where \tilde{h} is orthogonal to ψ (i.e., $\langle \psi, \tilde{h} \rangle = 0$). The component $\mu = \langle \psi, h \rangle$ represents the structural parameter, while \tilde{h} represents an infinite dimensional nuisance parameter. Although the full perturbation direction h is unknown, only the projection onto the efficient influence function is relevant for learning about μ .

For each $h \in T(P_0)$, define $\mu_n(h) := \mu(P_{1/\sqrt{n},h})$, and let $\mathbf{x} := (Y_1, \dots, Y_n)$ denote the collection of outcomes. We can rewrite the loss functions in terms of h as

$$l_n(h, \delta) = \begin{cases} \ell(\sqrt{n}(\mu_n(h) - \delta)) & \text{for estimation loss,} \\ \sqrt{n}(\mu_n(h) \mathbf{1}\{\mu_n(h) \geq 0\} - \mu_n(h) \delta) & \text{for treatment-assignment loss.} \end{cases}$$

In the semi-parametric setting, a Bayesian statistical model, $m(h, \mathbf{x})$, is defined as a joint probability distribution over both h, \mathbf{x} . As in Section 2.2, it admits the decomposition

$$m(h, \mathbf{x}) = \pi(h) \otimes p_{n,h}(\mathbf{x}),$$

where $\pi \in \Delta(T(P_0))$ denotes a prior over the tangent space $T(P_0)$, and $p_{n,h}(\mathbf{x}) := \prod_i dP_{1/\sqrt{n},h}(Y_i)/d\nu$ represents the likelihood, a parametric sub-model. Prior ambiguity is incorporated by defining a frequentist model, a structured set of models, as

$$\mathcal{Q} := \left\{ \pi(h) \otimes p_{n,h}(\mathbf{x}) : \pi \in \Delta(T(P_0)) \right\}.$$

Finally, misspecification concerns are addressed by placing a protective belt around \mathcal{Q} , yielding the set of unstructured models

$$\mathcal{M} = \left\{ m \in \Delta(T(P_0), \mathcal{X}) : \min_{q \in \mathcal{Q}} R_q(m) \leq K \right\}.$$

Expanding the set of structured models allows us to account for the possibility that the true distribution of the experimental data \mathbf{x} is not captured by $p_{n,h}(\mathbf{x})$ for any h .

The decision-maker chooses the decision rule that performs best against the worst-case model in \mathcal{M} , thereby guarding against both prior ambiguity and misspecification:

$$\delta_n^* := \arg \min_{\delta} \left[\sup_{m \in \mathcal{M}} \mathbb{E}_m[l_n(h, \delta)] \right].$$

As in Section 3, applying a Lagrangian formulation and the same sequence of calculations yields the following characterization of minimal decision risk:

$$V_n^* = \min_{\delta} \max_{\pi \in \Delta(T(P_0))} \int \mathbb{E}_{n,h} \left[e^{l_n(h,\delta)/\lambda} \right] d\pi(h). \quad (5.4)$$

5.2. Formal results: Semi-parametric models. We impose the following regularity conditions throughout this section:

Assumption 5. *The sub-models $\{P_{s,h}; h \in T(P_0)\}$ satisfy (5.1). Furthermore, they admit an efficient influence function, $\psi(\cdot)$, for $\mu(P)$ such that*

$$\sqrt{n} \left(\mu(P_{1/\sqrt{n},h}) - \mu_0 \right) = \langle \psi, h \rangle + \epsilon_n \|h\|^2,$$

where $\mu_0 := \mu(P_0)$, and ϵ_n is independent of h for bounded $\|h\|$.

The first part of Assumption 5 simply states the definition of parametric sub-models. The second part of Assumption 5 slightly strengthens (5.3).

Since $L^2(P_0)$ is a Hilbert space, it is possible to select $\{\phi_1, \phi_2, \dots\} \in L^2(P_0)$ in such a manner that $\{\psi/\sigma, \phi_1, \phi_2, \dots\}$ is a set of orthonormal basis functions for the closure of $T(P_0)$; the division by σ in the first component ensures $\|\psi/\sigma\|^2 = 1$. We can also choose these bases so they lie in $T(P_0)$, i.e., $\mathbb{E}_{P_0}[\phi_j] = 0$ for all j . By the Hilbert space isometry, each $h \in T(P_0)$ is then associated with an element from the l_2 space of square integrable sequences, $(\mu/\sigma, \gamma_1, \gamma_2, \dots)$, where $\mu = \langle \psi, h \rangle$ and $\gamma_k = \langle \phi_k, h \rangle$ for all $k \neq 0$. Consequently, any prior $\pi(h)$ over $T(P_0)$ can be represented as a prior over l_2 .

As in Section 4.4, our formal results require localization of ambiguity. This involves restricting attention to priors that are supported on a compact subset, K_M , of l_2 , defined as

$$K_M \equiv \left\{ h = (\mu/\sigma, \gamma_1, \dots) : \|h\| \leq M, \lim_{J \rightarrow \infty} \sup_{(\gamma_1, \gamma_2, \dots)} \sum_{j=J}^{\infty} |\gamma_j|^2 = 0 \right\}.$$

The compactness condition essentially requires the set of candidate h to be sufficiently smooth.

5.2.1. Lower bounds. As in the parametric setting, we show that the minimax value V^* in the limit experiment also forms an asymptotic lower bound on the

sequence of optimal decision risks, V_n^* , in the finite sample semi-parametric experiments. However, since the previous definition of the limit experiment used a different interpretation of h , we will need to modify the construction slightly.

Specifically, we now consider a limit experiment where we observe a one dimensional signal x , posited to be drawn from a reference Gaussian likelihood, $P_\mu(x) \sim \mathcal{N}(\mu/\sigma, 1)$. Let $\mathbb{E}_\mu[\cdot]$ denote the expectation corresponding to P_μ . The minimax value V^* in this limit experiment is then defined as

$$V^* := \min_{\tilde{\delta}} \max_{\rho(\mu) \in \Delta(\mathbb{R})} \int \mathbb{E}_\mu \left[e^{l(\mu, \tilde{\delta})/\lambda} \right] d\rho(\mu), \text{ with} \quad (5.5)$$

$$l(\mu, \tilde{\delta}) = \begin{cases} \ell(\mu - \tilde{\delta}) & \text{for estimation loss,} \\ \mu \{ \mathbf{1}\{\mu \geq 0\} - \tilde{\delta} \} & \text{for treatment-assignment loss.} \end{cases}$$

It is straightforward to verify that the value of V^* in (5.5) is, in fact, the same as that in (4.6) when $\theta = \mu$ and $I_0 = 1/\sigma^2$.

Theorem 3. *Suppose that Assumption 5 holds. Then, under both the estimation and treatment-assignment loss functions,*

$$\liminf_{n \rightarrow \infty} \min_{\tilde{\delta}} \max_{\pi(h) \in \Delta(K_M)} \int \mathbb{E}_{n,h} \left[e^{l_n(h, \tilde{\delta})/\lambda} \right] d\pi(h) \geq V^*.$$

5.2.2. *Asymptotically optimal decisions.* As in the parametric setting, asymptotically optimal decisions under ambiguity and misspecification are the same as those under ambiguity alone. Let $\hat{\mu}_n$ denote any semi-parametrically efficient estimator for μ , understood as satisfying the following assumption:

Assumption 6. *The estimator $\hat{\mu}$ attains the semi-parametric efficiency bound, in that it admits a locally linear influence-function approximation:*

$$\hat{\mu}_n - \mu_0 = \sigma x_n + o_{P_{n,0}}(1).$$

As we show below, asymptotically optimal decisions are then given by

$$\hat{\delta}_n^* = \begin{cases} \hat{\mu}_n & \text{for estimation,} \\ \mathbf{1}\{\hat{\mu}_n \geq 0\} & \text{for treatment assignment.} \end{cases}$$

Theorem 4. *Suppose that Assumptions 4-6 hold. Then, under both the estimation and treatment-assignment loss functions,*

$$\lim_{M \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta(K_M)} \int \mathbb{E}_{n,h} \left[e^{l_n(h, \hat{\delta}_n^*)/\lambda} \right] d\pi(h) = V^*.$$

5.3. Application: Optimal GMM estimators under misspecification. The Generalized Method of Moments (GMM) is an example of a semi-parametric model that is widely used in economic applications. Recall that in the GMM framework, the researcher specifies a moment condition $\mathbb{E}_P[m(Y_i, \mu)] = 0$, where $m(\cdot) \in \mathbb{R}^p$ is a known vector of moments, $\mu \in \mathbb{R}^d$ is the structural parameter, and P is the population distribution. When $p > d$, the GMM model is said to be over-identified. In this setting, the efficient influence function is given by

$$\psi(Y_i) = G_0^\top \Omega_0^{-1} m(Y_i, \mu_0),$$

where $\mu_0 := \mu(P_0)$ is the unique solution to $\mathbb{E}_{P_0}[m(Y_i, \mu)] = 0$ under the reference distribution P_0 , $G_0 := \mathbb{E}_{P_0}[\nabla_\mu m(Y_i, \mu)]$ and $\Omega_0 := \mathbb{E}_{P_0}[m(Y_i, \mu_0)m(Y_i, \mu_0)^\top]$.

In the absence of misspecification concerns, it is well known that several estimators that are asymptotically efficient, including 2-step GMM and continuously updated GMM (CUGMM), among others. In practice, however, researchers are often concerned that the model may be misspecified, i.e., $\mathbb{E}_{\hat{P}}[m(Y_i, \mu)] \neq 0$ under the true structural parameter μ and the sample outcome distribution \hat{P} . Under misspecification, different estimators converge to different limits under the distribution of sample outcomes \hat{P} , and researchers often resort to inefficient estimators employing a weighing matrix other than the optimal $W^* = G_0^\top \Omega_0^{-1}$. This practice is frequently justified on the grounds that “...under misspecification the two-step GMM estimator is no more efficient than any other estimator... each weighting leads us to recover a different parameter” (Andrews et al., 2025).

Our results suggest, however, that this reasoning is incomplete. When the decision-maker confronts both prior ambiguity and model misspecification as in our framework, the optimal estimator coincides with that under prior ambiguity alone. Consequently, 2-step GMM remains superior to diagonally-weighted GMM, even under an arbitrary degree of misspecification. Researchers who wish to employ inefficient weighting should therefore provide explicit justification for doing so. If,

for instance, the researcher believes that misspecification is not unstructured but that certain forms or directions of misspecification are more likely than others, this could in principle lead to different estimation strategies. Even so, it would be difficult to justify use of identify or diagonal weighting on the basis of directional misspecification alone, since such weighting typically preserves symmetry across directions.

Apart from requiring efficiency, our results do not distinguish among different efficient estimators. Under local asymptotics, all efficient estimators attain the same decision risk V^* , so additional criteria must be employed to select among them. For instance, imposing invariance would lead one to prefer CUGMM over two-step GMM.

6. EXTENSIONS

In this section, we discuss several variations and extensions of our framework.

6.1. Alternative discrepancy measures. So far, we have used relative entropy $R_q(m)$ to measure the discrepancy between statistical models. Hansen and Sargent (2011, Chapter 1.8) provide two important reasons for using relative entropy. First, it leads to a tractable characterization of minimal decision risk and optimal decisions, as demonstrated in Section 3. Second, as discussed in Hansen and Sargent (2011, Chapter 9), it can be linked to risk-sensitivity adjustment through the theory of large deviations.

It is, however, possible to employ alternative discrepancy measures. Relative entropy, is but a special case of the ϕ -divergence class of discrepancies, which take the general form

$$D_\phi(m||q) = \int \phi \left(\frac{dm}{dq} \right) dq,$$

where $\phi : [0, \infty) \rightarrow (-\infty, \infty]$ is a convex function satisfying $\phi(x) < \infty$ for all $x > 0$, $\phi(1) = 0$ and $\phi(0) = \lim_{x \rightarrow 0} \phi(x)$. Setting $\phi(x) = x \ln x$ recovers KL divergence.

Under this more general class of discrepancies, the decision-risk of a rule δ can be characterized in variational form as

$$V_{n,\phi}(\delta) = \min_{q \in \mathcal{Q}} \inf_m \left\{ \mathbb{E}_m \left[u_n(\theta, \delta) \right] + \lambda D_\phi(m \| q) \right\}.$$

Let $\phi^*(\cdot)$ denote the convex conjugate of $\phi(\cdot)$. The decision risk $V_{n,\phi}(\delta)$ then admits the dual representation

$$\begin{aligned} V_{n,\phi}(\delta) &= \lambda \min_{q \in \mathcal{Q}} \sup_{\eta \in \mathbb{R}} \left\{ \eta - \int \phi^* \left(\eta - \frac{u_n(\theta, \delta)}{\lambda} \right) dq \right\} \\ &= \lambda \min_{\pi \in \Delta(\Theta)} \sup_{\eta \in \mathbb{R}} \left\{ \eta - \int \mathbb{E}_{p(x|\theta)} \left[\phi^* \left(\eta + \frac{l_n(\theta, \delta)}{\lambda} \right) \right] d\pi(\theta) \right\}, \end{aligned}$$

where the first equality is due to Cerreia-Vioglio et al. (2026, Proposition 1) and the second equality exploits the specific structure of \mathcal{Q} in our setting.

By standard properties of convex conjugates, $\phi^*(x) < \infty$ for all bounded x if and only if $\lim_{x \rightarrow \infty} \phi(x)/x = \infty$. Since our loss functions $l_n(\theta, \delta)$ are generally unbounded, $V_{n,\phi}(\delta)$ is therefore ∞ for any ϕ -divergence measure satisfying $\lim_{x \rightarrow \infty} \phi(x)/x < \infty$, unless sharp support restrictions are imposed on the class of priors. In other words, the decision risk is trivially infinite for these divergence measures whenever the class of priors is sufficiently rich. We therefore argue that such discrepancies are not well suited for a framework that accommodates both ambiguity and misspecification. Notable examples in this category include total variation ($\phi(x) = |x - 1|/2$), squared Hellinger distance ($\phi(x) = (\sqrt{x} - 1)^2/2$) and Pearson χ^2 divergence ($\phi(x) = (x - 1)^2/x$).

The underlying issue with these divergence measures is that the misspecification they permit is too broad. Consider, for instance, the total-variation metric: it is possible to have $dm/dq = \infty$, meaning $q \in \mathcal{Q}$ may not share the same support as the true model, even as total variation remains finite. This would imply that \mathcal{Q} is blatantly misspecified, in the sense that any specification test would surely reject it almost surely. In contrast, as Hansen and Sargent (2011, Chapter 9) argue, when $R_{\mathcal{Q}}(m) < \infty$, the approximating model can still be regarded as plausibly correct, since it would not be rejected with probability one.

Among the ϕ -divergence measures satisfying $\lim_{x \rightarrow \infty} \phi(x)/x = \infty$, the only commonly used divergence apart from KL divergence is the Neyman χ^2 divergence ($\phi(x) = (x - 1)^2$). The Neyman χ^2 divergence is a stronger discrepancy measure than KL divergence: it is possible to have $\chi^2(m||q) < \infty$ while $R_q(m) = \infty$. Consequently, a KL-based misspecification set $\mathcal{M} = \left\{ m \in \Delta(\Theta, \mathcal{X}) : \min_{q \in \mathcal{Q}} D_q(m) \leq K \right\}$ is strictly larger than the corresponding χ^2 -based set $\mathcal{M}_{\chi^2} = \left\{ m \in \Delta(\Theta, \mathcal{X}) : \min_{q \in \mathcal{Q}} D_{\chi^2}(m||q) \leq K \right\}$. Based on this insight, we conjecture that our results on the optimality of efficient decisions continue to hold for the Neyman χ^2 divergence as well, though we leave the formal analysis to future work.

6.2. Asymmetric loss functions. The estimation and treatment assignment loss functions considered thus far share a crucial property: they are symmetric in the sense that overestimating θ by a given amount incurs the same loss as underestimating it by that amount. This symmetry is crucial to the result that optimal decisions do not depend on the degree of misspecification.

There are, however, many loss functions that lack this symmetry. A prominent example is the linex loss $l(\theta, \delta) = e^{\mu(\theta) - \delta} - (\mu(\theta) - \delta) - 1$ which penalizes positive errors much more than negative errors. Under such losses, the optimal estimator is biased even in the absence of misspecification. Incorporating misspecification concerns introduces additional bias whose magnitude depends on the misspecification parameter λ ; see Appendix C for details in the context of linex loss.⁴ Intuitively, misspecification entails an exponential tilting of the loss function, which further exacerbates any asymmetry already present in the loss. Consequently, for asymmetric losses, optimal decisions under ambiguity and misspecification may not coincide with those under ambiguity alone.

⁴Under misspecification, the linex loss function must be truncated to keep the minimax risk finite. The optimal estimator therefore also depends on the level of truncation. However, for any level of truncation, the bias of the optimal estimator decreases in λ (note that $\lambda = 0$ corresponds to no misspecification risk).

6.3. Relaxing caution: Smooth ambiguity aversion and hierarchical Bayes.

Our framework employs the Waldian approach to ambiguity by selecting the worst-case model within the ambiguity set – the structured class of models \mathcal{Q} . Cerreia-Vioglio et al. (2026) term the preference axiom underlying this approach as ‘caution’. An alternative is to adopt smooth ambiguity aversion, as in Klibanoff et al. (2005), by introducing a probability distribution $\varrho_{\mathcal{Q}}(\cdot)$ over \mathcal{Q} . Intuitively, smooth ambiguity aversion corresponds to making the decision-maker less cautious: rather than guarding against the worst case, she averages over models according to $\varrho_{\mathcal{Q}}(\cdot)$. Since \mathcal{Q} pairs a candidate likelihood with every possible prior, a distribution over \mathcal{Q} is equivalent to a distribution $\varrho(\pi)$ over the space of priors $\Delta(\Theta)$: the latter can simply be interpreted as a hyperprior in a Bayesian hierarchical model.

Cerreia-Vioglio et al. (2026, Section 6.1) show that under smooth ambiguity aversion, the decision risk takes the form

$$\begin{aligned}\tilde{V}_n(\delta) &= \int_{\mathcal{Q}} \phi_{\mathcal{Q}} \left(\inf_m \left\{ \mathbb{E}_m [u_n(\theta, \delta)] + \lambda R_q(m) \right\} \right) d\varrho_{\mathcal{Q}}(q) \\ &= \int_{\mathcal{Q}} \phi_{\mathcal{Q}} \left(-\lambda \ln \mathbb{E}_q \left[e^{-u_n(\theta, \delta)/\lambda} \right] \right) d\varrho_{\mathcal{Q}}(q),\end{aligned}$$

where $\phi_{\mathcal{Q}}(\cdot)$ is a monotone function and the second equality uses (3.1). Setting $\phi_{\mathcal{Q}}(t) = -e^{-t/\lambda}$ and converting $\varrho_{\mathcal{Q}}(q)$ to a prior $\varrho(\pi)$ over $\Delta(\Theta)$ yields

$$\begin{aligned}\tilde{V}_n(\delta) &= - \int_{\Delta(\Theta)} \left(\int \mathbb{E}_{p(x|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right] d\pi(\theta) \right) d\varrho(\pi) \\ &= - \int \mathbb{E}_{p(x|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right] d\bar{\pi}(\theta),\end{aligned}$$

where $\bar{\pi}(\theta)$ is the effective prior induced by the hyperprior $\varrho(\cdot)$ over $\Delta(\Theta)$. Hence, under this specification of $\phi_{\mathcal{Q}}(\cdot)$, misspecification combined with smooth ambiguity aversion with $\phi_{\mathcal{Q}}(t) = -e^{-t/\lambda}$ is equivalent to misspecification with a single hierarchical prior.

The choice of $\phi_{\mathcal{Q}}(t) = -e^{-t/\lambda}$, however, uses the same parameter λ to govern both aversion to prior ambiguity and sensitivity to model misspecification. It may therefore be more natural to set $\phi_{\mathcal{Q}}(t) = -e^{-t/\xi}$, where ξ captures aversion to prior uncertainty separately from the misspecification parameter λ . With this choice,

the decision risk becomes

$$\begin{aligned}\tilde{V}_n(\delta) &= - \int_{\mathcal{Q}} \left(\mathbb{E}_q \left[e^{l_n(\theta, \delta) / \lambda} \right] \right)^{\lambda / \xi} d\varrho_{\mathcal{Q}}(q) \\ &= - \int_{\Delta(\Theta)} \left(\int \mathbb{E}_{p(\mathbf{x}|\theta)} \left[e^{l_n(\theta, \delta) / \lambda} \right] d\pi \right)^{\lambda / \xi} d\varrho(\pi).\end{aligned}\tag{6.1}$$

As $\xi \rightarrow 0$, aversion to prior ambiguity grows without bound and the objective $\tilde{V}_n(\delta)$ reduces to the decision risk $V_n(\delta)$ from (3.2) — the Waldian formulation involving the least favorable prior, and the primary focus of this article. For $\xi \in (0, \infty)$, the decision-maker exhibits less aversion to prior ambiguity, but this comes at the cost of a nonlinear objective when $\lambda \neq \xi$, which complicates the analysis. A formal treatment of the more general criterion (6.1) when $\xi \notin \{\lambda, \infty\}$ is therefore left for future research.

6.4. Model averaging. While we have so far focused on misspecification of a single likelihood, practitioners are often interested in selecting or averaging among multiple competing likelihood specifications, each potentially subject to varying degrees of misspecification concern.

Let $p_{1,\theta}(\mathbf{x}) \equiv p_1(\mathbf{x}|\theta)$ and $p_{2,\theta}(\mathbf{x}) \equiv p_2(\mathbf{x}|\theta)$ denote two candidate likelihoods, and let $\mathcal{Q}_1, \mathcal{Q}_2$ denote the corresponding frequentist models, where, as in Section 2.2.2,

$$\mathcal{Q}_a := \left\{ \pi(\theta) \otimes p_{a,\theta}(\mathbf{x}) : \pi \in \Delta(\Theta) \right\}, \quad a \in \{1, 2\}.$$

Suppose that, treating each likelihood in isolation, the decision maker contemplates a misspecification set for each of the form

$$\mathcal{M}_a = \left\{ m \in \Delta(\Theta, \mathcal{X}) : \min_{q \in \mathcal{Q}_a} R_{q_a}(m) \leq K_a \right\}, \quad a \in \{1, 2\}.$$

Here, K_a quantifies the decision-maker's misspecification concern for each model: $K_2 \gg K_1$ implies that the decision maker has greater concern about likelihood 2 being misspecified than about likelihood 1. Our framework permits the two likelihoods to be nested, i.e., $\{p_{2,\theta}(\mathbf{x})\}_\theta \subseteq \{p_{1,\theta}(\mathbf{x})\}_\theta$, which can be achieved by fixing some components of θ in $p_{2,\theta}(\mathbf{x})$. For instance, in the running example, Alice may be choosing between two specifications $Y \sim \mathcal{N}(\mu, \sigma^2)$ and $Y \sim \mathcal{N}(\mu, 1)$; here, $\theta := (\mu, \sigma^2)$, but $\sigma^2 = 1$ is fixed in the second specification.

Rather than treating the models in isolation, however, the decision maker may wish to combine them. The overall set of misspecified models under consideration can then be taken to be the intersection

$$\mathcal{M} = \mathcal{M}_1 \cap \mathcal{M}_2,$$

which we assume to be non-empty; otherwise, the two frequentist models would be incompatible. With this choice of \mathcal{M} , the optimal decision becomes

$$\begin{aligned} \delta_n^* &:= \inf_{\delta} \inf_{m \in \mathcal{M}} \mathbb{E}_m[l_n(\theta, \delta)] \\ &= \inf_{\delta} \inf_m \left\{ \mathbb{E}_m[u_n(\theta, \delta)] : R_{\mathcal{Q}_1}(m) \leq K_1, R_{\mathcal{Q}_2}(m) \leq K_2 \right\}. \end{aligned}$$

As $R_{\mathcal{Q}_a}(\cdot)$ is strictly convex, standard duality arguments show that we can alternatively characterize δ_n^* through the more convenient variational form:

$$\begin{aligned} \delta_n^* &= \inf_{\delta} V_n(\delta); \quad \text{where} \\ V_n(\delta) &:= \inf_m \left\{ \mathbb{E}_m[u_n(\theta, \delta)] + \lambda (\alpha R_{\mathcal{Q}_1}(m) + (1 - \alpha) R_{\mathcal{Q}_2}(m)) \right\}, \end{aligned}$$

for some $\alpha \in [0, 1]$ and $\lambda \geq 0$ that depend on K_1, K_2 . In particular, $\alpha > 1/2$ whenever $K_2 \gg K_1$: the decision risk places greater weight on the likelihood that is less likely to be misspecified.

Let

$$p_{\alpha, \theta}(\mathbf{x}) \equiv p_{\alpha}(\mathbf{x}|\theta) := \frac{p_1^{\alpha}(\mathbf{x}|\theta) \cdot p_2^{1-\alpha}(\mathbf{x}|\theta)}{\int p_1^{\alpha}(\mathbf{x}|\theta) \cdot p_2^{1-\alpha}(\mathbf{x}|\theta) d\mathbf{x}}$$

denote the (normalized) geometric mixture likelihood that combines $p_1(\mathbf{x}|\theta)$ and $p_2(\mathbf{x}|\theta)$ with mixing weights α and $1 - \alpha$, respectively. Recalling the decomposition $m = \pi(\theta) \otimes m_{\theta}(\mathbf{x})$ and applying (2.3) yields

$$\begin{aligned} V_n(\delta) &= \inf_{\pi(\theta) \otimes m_{\theta}(\mathbf{x})} \int \left\{ \mathbb{E}_{m_{\theta}(\mathbf{x})}[u_n(\theta, \delta)] + \lambda (\alpha \text{KL}(m_{\theta} \| p_{1, \theta}) + (1 - \alpha) \text{KL}(m_{\theta} \| p_{2, \theta})) \right\} d\pi(\theta) \\ &= \inf_{\pi(\theta) \otimes m_{\theta}(\mathbf{x})} \int \left\{ \mathbb{E}_{m_{\theta}(\mathbf{x})}[u_n(\theta, \delta)] + \lambda \cdot \text{KL}(m_{\theta} \| p_{\alpha, \theta}) \right\} d\pi(\theta), \end{aligned}$$

where the last step in makes use of the fact that the weighted sum of KL divergences can be expressed—up to a constant that does not depend on \mathbf{x} or δ —as a single KL divergence against a geometric mixture. Applying the Donsker–Varadhan variational formula again then gives

$$\begin{aligned} V_n(\delta) &= \inf_{\pi(\theta)} \int \inf_{m_\theta(\mathbf{x})} \left\{ \mathbb{E}_{m_\theta(\mathbf{x})} [u_n(\theta, \delta)] + \lambda \cdot \text{KL}(m_\theta \| p_{\alpha, \theta}) \right\} d\pi(\theta) \\ &= \inf_{\pi(\theta)} -\lambda \ln \left\{ \int \mathbb{E}_{p_\alpha(\mathbf{x}|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right] d\pi(\theta) \right\}. \end{aligned}$$

The optimal decision therefore solves

$$\delta_n^* = \arg \min_{\delta} \max_{\pi \in \Delta(\Theta)} \int \mathbb{E}_{p_\alpha(\mathbf{x}|\theta)} \left[e^{l_n(\theta, \delta)/\lambda} \right] d\pi(\theta).$$

Optimal decisions under multiple possibly misspecified candidate likelihoods are thus equivalent to minimax decisions with an exponentiated loss function and a single geometric mixture likelihood $p_\alpha(\mathbf{x}|\theta)$. All of our theoretical results therefore continue to apply upon reinterpreting $p_\alpha(\mathbf{x}|\theta)$ as the relevant reference likelihood.

7. CONCLUSION

In this article, we have introduced a framework for evaluating statistical decisions under both prior ambiguity and likelihood misspecification. Misspecification manifests as an exponential tilting of the loss function, while ambiguity corresponds to a search for the least favorable prior. We also develop a theory of local asymptotics under global misspecification, achieved by localizing the priors around a reference parameter, and use this theory to characterize optimal estimation and treatment-assignment decisions. Remarkably, in both cases, optimal decisions coincide with those under correct likelihood specification.

The proposed framework opens several avenues for further research. While we discuss some examples of asymmetric loss functions in Appendix C, a general theory for characterizing optimal decisions under such losses remains to be developed. As noted there, optimal decisions under asymmetric loss may depend on the degree of misspecification. On model selection, while Section 6.4 provides an initial treatment, a richer characterization would require taking the misspecification risks of competing likelihoods to converge to each other, so that meaningful

tradeoffs between efficiency and misspecification robustness may emerge. A further extension would be to separate ambiguity concerns over the prior from those over candidate likelihoods, employing smooth ambiguity aversion for the latter, as discussed in Section 6.3. This would bring the framework closer to the literature on Bayesian model averaging and model selection.

REFERENCES

- ANDREWS, I., J. CHEN, AND O. TECCHIO (2025): “The Purpose of an Estimator is What it Does: Misspecification, Estimands, and Over-Identification,” *arXiv preprint arXiv:2508.13076*.
- ANDREWS, I., M. GENTZKOW, AND J. M. SHAPIRO (2020): “On the Informativeness of Descriptive Statistics for Structural Estimates,” *Econometrica*, 88, 2231–2258.
- ANSCOMBE, F. J., AND R. J. AUMANN (1963): “A Definition of Subjective Probability,” *The Annals of Mathematical Statistics*, 34, 199–205.
- ARMSTRONG, T. B. (2025): “Misspecification in Econometrics: A Selective Review.”
- BEN-TAL, A., L. E. GHAOUI, AND A. NEMIROVSKI (2009): *Robust Optimization*: Princeton University Press.
- BONHOMME, S., AND M. WEIDNER (2022): “Minimizing Sensitivity to Model Misspecification,” *Quantitative Economics*, 13, 907–954.
- BORDALO, P., N. GENNAIOLI, Y. MA, AND A. SHLEIFER (2020): “Overreaction in Macroeconomic Expectations,” *American Economic Review*, 110, 2748–2782.
- BOX, G. E. (1976): “Science and Statistics,” *Journal of the American Statistical Association*, 71, 791–799.
- CERREIA-VIOGLIO, S., L. P. HANSEN, F. MACCHERONI, AND M. MARINACCI (2026): “Making Decisions under Model Misspecification,” *Review of Economic Studies*, 93, 892–925.
- CHRISTENSEN, T., AND B. CONNAULT (2023): “Counterfactual Sensitivity and Robustness,” *Econometrica*, 91, 263–298.
- CHRISTENSEN, T., H. R. MOON, AND F. SCHORFHEIDE (2022): “Optimal Discrete Decisions when Payoffs are Partially Identified,” *arXiv preprint*

- arXiv:2204.11748*.
- DONTCHEV, A. L., AND R. T. ROCKAFELLAR (2009): *Implicit Functions and Solution Mappings* Volume 543: Springer.
- HANSEN, L. P., AND T. J. SARGENT (2011): *Robustness*: Princeton university press.
- HIRANO, K. (2025): “Wald’s Statistical Decision Theory for Policy Analysis and Adaptive Experiments.”
- HIRANO, K., AND J. R. PORTER (2009): “Asymptotics for Statistical Treatment Rules,” *Econometrica*, 77, 1683–1701.
- HUBER, P. J. (1964): “Robust Estimation of a Location Parameter,” *Annals of Mathematical Statistics*, 35, 73–101.
- IBRAGIMOV, I., AND R. HASMINSKII (1981): *Statistical Estimation: Asymptotic Theory*: Springer.
- ISHIHARA, T., AND T. KITAGAWA (2021): “Evidence Aggregation for Treatment Choice,” *arXiv preprint arXiv:2108.06473*.
- KLIBANOFF, P., M. MARINACCI, AND S. MUKERJI (2005): “A Smooth Model of Decision Making under Ambiguity,” *Econometrica*, 73, 1849–1892.
- KUHN, D., S. SHAFIEE, AND W. WIESEMANN (2025): “Distributionally Robust Optimization,” *Acta Numerica*, 34, 579–804.
- LE CAM, L. M. (1986): *Asymptotic Methods in Statistical Decision Theory*: Springer-Verlag.
- MACCHERONI, F., M. MARINACCI, AND A. RUSTICHINI (2006): “Ambiguity Aversion, Robustness, and the Variational Representation of Preferences,” *Econometrica*, 74, 1447–1498.
- MANSKI, C. F. (2003): *Partial Identification of Probability Distributions*: Springer.
- MASTEN, M. A., AND A. POIRIER (2021): “Salvaging Falsified Instrumental Variable Models,” *Econometrica*, 89, 1449–1469.
- MONTIEL OLEA, J. L., C. QIU, AND J. STOYE (2026): “Decision Theory for Treatment Choice Problems with Partial Identification,” *Review of Economic Studies*, rdag015.

- MÜLLER, U. K., AND A. NORETS (2024): “Locally Robust Semiparametrically Efficient Bayesian Inference,” *Working paper*.
- RAHIMIAN, H., AND S. MEHROTRA (2019): “Distributionally Robust Optimization: A Review,” *arXiv preprint arXiv:1908.05659*.
- ROBERT, C. P. (2007): *The Bayesian Choice: From Decision-Theoretic Foundations to Computational Implementation*: Springer.
- VAN DER VAART, A. W. (2000): *Asymptotic Statistics*: Cambridge university press.
- VAN DER VAART, A. W., AND J. WELLNER (1996): *Weak Convergence and Empirical Processes: With Applications to Statistics*: Springer Science & Business Media.
- WALD, A. (1950): “Statistical Decision Functions,” in *Breakthroughs in Statistics: Foundations and Basic Theory*: Springer, 342–357.
- WHITE, H. (1982): “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica: Journal of the econometric society*, 1–25.
- YATA, K. (2021): “Optimal Decision Rules under Partial Identification,” *arXiv preprint arXiv:2111.04926*.

APPENDIX A. PROOFS

A.1. Proof of Proposition 1. The claim follows if we show that the decision maker's choice of $\tilde{\delta}^* = \mathbf{1}\{\dot{\mu}_0^\top I_0^{-1/2} x \geq 0\}$ and nature's choice of a two point prior supported on $(-h^*, h^*)$ constitute a Nash equilibrium. Note that the proof applies even if we allow for randomized decisions, where $\tilde{\delta}$ is interpreted as the probability of treatment assignment and $\mathbb{E}_h[\cdot]$ represents the expectation over both the likelihood and the policy randomization. The optimal decision, however, does not involve randomization.

Let $\sigma = \sqrt{\dot{\mu}_0^\top I_0^{-1} \dot{\mu}_0}$. Consider the best response of the decision-maker to any symmetric two-point prior, π_Δ , supported on

$$\left(-\frac{\Delta}{\sigma^2} I_0^{-1} \dot{\mu}_0, \frac{\Delta}{\sigma^2} I_0^{-1} \dot{\mu}_0 \right),$$

where $\Delta \in [0, \infty)$ is arbitrary. Let $\mathbb{P}_{h|x}$ denote the posterior probability over h given this prior. Some straightforward algebra shows that

$$q(x) := \mathbb{P}_{h|x} \left(h = -\frac{\Delta}{\sigma^2} I_0^{-1} \dot{\mu}_0 \right) = \frac{1}{1 + \exp \left\{ 2\Delta \left(\dot{\mu}_0^\top I_0^{-1/2} x \right) / \sigma^2 \right\}}.$$

The Bayes-optimal response to π_Δ is therefore

$$\begin{aligned} \tilde{\delta}_\Delta &= \mathbf{1} \left\{ \mathbb{E}_{h|x} \left[\exp \left\{ -\frac{\dot{\mu}_0^\top h}{\lambda} \mathbf{1}\{\dot{\mu}_0^\top h < 0\} \right\} \right] \leq \mathbb{E}_{h|x} \left[\exp \left\{ \frac{\dot{\mu}_0^\top h}{\lambda} \mathbf{1}\{\dot{\mu}_0^\top h \geq 0\} \right\} \right] \right\} \\ &= \mathbf{1} \left\{ e^{\frac{\sigma\Delta}{2\lambda}} q(x) + (1 - q(x)) \leq q(x) + e^{\frac{\sigma\Delta}{2\lambda}} (1 - q(x)) \right\} \\ &= \mathbf{1} \{q(x) \leq 1/2\} = \mathbf{1} \left\{ \dot{\mu}_0^\top I_0^{-1/2} x \geq 0 \right\}. \end{aligned}$$

Hence, $\tilde{\delta}^* = \mathbf{1}\{\dot{\mu}_0^\top I_0^{-1/2} x \geq 0\}$ is a Bayes optimal response to π_Δ for any $\Delta > 0$.

We now consider nature's best response to $\tilde{\delta}^*$. Observe that for any $\dot{\mu}_0^\top h \geq 0$,

$$\begin{aligned} \mathbb{E}_h \left[e^{l(h, \tilde{\delta}^*)/\lambda} \right] &= \mathbb{E}_h \left[e^{\frac{\dot{\mu}_0^\top h}{\lambda} \mathbf{1}\{\dot{\mu}_0^\top I_0^{-1/2} x \leq 0\}} \right] \\ &= \mathbb{E}_0 \left[e^{\frac{\dot{\mu}_0^\top h}{\lambda} \mathbf{1}\{x \leq -\dot{\mu}_0^\top h\}} \right] = e^{\frac{|\dot{\mu}_0^\top h|}{\lambda}} \Phi(-|\dot{\mu}_0^\top h|) + 1 - \Phi(-|\dot{\mu}_0^\top h|). \end{aligned}$$

Similarly, for any $\dot{\mu}_0^\top h < 0$,

$$\begin{aligned}\mathbb{E}_h \left[e^{l(h, \tilde{\delta}^*)/\lambda} \right] &= \mathbb{E}_0 \left[e^{-\frac{\dot{\mu}_0^\top h}{\lambda} \mathbf{1}\{x \geq -\dot{\mu}_0^\top h\}} \right]. \\ &= e^{\frac{|\dot{\mu}_0^\top h|}{\lambda}} \Phi(-|\dot{\mu}_0^\top h|) + 1 - \Phi(-|\dot{\mu}_0^\top h|).\end{aligned}$$

Thus, the expected loss is the constant in $|\dot{\mu}_0^\top h|$, implying that nature is indifferent between choosing $-h$ and h for any h . Nature's best response to $\tilde{\delta}^*$ is therefore any prior supported on $\{h : |\dot{\mu}_0^\top h| = \Delta^*\}$, where

$$\Delta^* := \arg \max_{\Delta \geq 0} \left\{ \left(e^{\frac{\Delta}{\lambda}} - 1 \right) \Phi(-\Delta) \right\}.$$

It is straightforward to verify that Δ^* exists and is unique. Thus, the symmetric two-point prior, π_{Δ^*} , supported on $(-h^*, h^*)$ is a best response to $\tilde{\delta}^*$.

A.2. Proof of Theorem 1. For estimation-loss, the theorem is a direct consequence of Van der Vaart (2000, Proposition 8.11); see also van der Vaart and Wellner (1996, Theorem 3.11.5).

We therefore focus on proving the theorem for treatment-assignment loss. Let

$$\begin{aligned}R_n(h, \delta) &:= \mathbb{E}_{n,h} \left[e^{l_n(\theta+h/\sqrt{n}, \delta)/\lambda} \right] \text{ and} \\ R(h, \tilde{\delta}) &:= \mathbb{E}_h \left[e^{l(h, \tilde{\delta})/\lambda} \right]\end{aligned}$$

denote the frequentist-risks of δ and $\tilde{\delta}$ in the finite-sample and limit experiments, respectively. As with the proof of Proposition 1, we allow for randomized decisions. For randomized decisions, $\delta, \tilde{\delta}_n$ should be understood as the probability that the treatment is assigned, and $\mathbb{E}_{n,h}[\cdot], \mathbb{E}_h[\cdot]$ should be understood as expectations over both the likelihood and the policy randomization.

We start by proving the following lemma:

Lemma 1. *Suppose Assumptions 1 and 2 hold. Let $\{\delta_n\}_n$ be a sequence of treatment decisions with associated frequentist regret $R_n(h, \delta_n)$. Then, there exists a subsequence $\{n_k\}_k$ and a (possibly randomized) treatment decision $\tilde{\delta}$ in the limit experiment such that $\lim_{k \rightarrow \infty} R_{n_k}(h, \delta_{n_k}) = R(h, \tilde{\delta})$.*

Proof. As $\delta_n \in [0, 1]$ is uniformly bounded, it is tight. Combined with (4.3), it follows that the joint

$$\left(\delta_n, \ln \frac{dP_{n, \theta_0 + h/\sqrt{n}}}{dP_{n, \theta_0}} \right)$$

is also tight. Hence, by Prohorov's theorem, given any sequence $\{n\}$, there exists a further sub-sequence $\{n_k\}$ such that

$$\begin{aligned} \left(\begin{array}{c} \delta_{n_k} \\ \ln \frac{dP_{n_k, \theta_0 + h/\sqrt{n_k}}}{dP_{n_k, \theta_0}} \end{array} \right) &\xrightarrow{P_{n,0}} \left(\begin{array}{c} \bar{\delta} \\ \ln V \end{array} \right); \\ V &\sim \exp \left\{ h^\top I_0^{1/2} x - \frac{1}{2} h^\top I_0 h \right\}, \end{aligned} \quad (\text{A.1})$$

where $x \sim \mathcal{N}(0, I)$ and $\bar{\delta} \in [0, 1]$ is some tight limit of δ_{n_k} . Observe that $V \geq 0$ and $E[V] = 1$. Therefore, by an application of Le Cam's third lemma,

$$\delta_{n_k} \xrightarrow{P_{n,h}} \mathcal{L}; \text{ where } \mathcal{L}(B) := E \left[\mathbb{I}\{\bar{\delta} \in B\} V \right] \quad \forall B \in \mathcal{B}(\mathbb{R}). \quad (\text{A.2})$$

Define $\tilde{\delta} = E[\bar{\delta}|x]$. By construction, $\tilde{\delta}$ is a valid treatment policy in the limit experiment. Furthermore, by (A.2),

$$\begin{aligned} \lim_{k \rightarrow \infty} \mathbb{E}_{n_k, h}[\delta_{n_k}] &= E \left[\bar{\delta} e^{h^\top I_0^{1/2} x - \frac{1}{2} h^\top I_0 h} \right] \\ &= E \left[\tilde{\delta} e^{h^\top I_0^{1/2} x - \frac{1}{2} h^\top I_0 h} \right] = \mathbb{E}_h[\tilde{\delta}], \end{aligned}$$

where the second equality follows by the law of iterated expectations, and the last equality follows by standard change-of-measure arguments.

Observe that

$$\begin{aligned} R_{n_k}(h, \delta) &= e^{l_{n_k}(\theta_0 + h/\sqrt{n_k}, 1)/\lambda} \mathbb{E}_{n_k, h}[\delta_{n_k}] + e^{l_{n_k}(\theta_0 + h/\sqrt{n_k}, 0)/\lambda} \mathbb{E}_{n_k, h}[1 - \delta_{n_k}] \text{ and} \\ R(h, \tilde{\delta}) &= e^{l(h, 1)/\lambda} \mathbb{E}_h[\tilde{\delta}] + e^{l(h, 0)/\lambda} \mathbb{E}_h[1 - \tilde{\delta}]. \end{aligned}$$

Now, (4.5) implies

$$l_{n_k}(\theta_0 + h/\sqrt{n_k}, a) \rightarrow l(h, a), \text{ for each } a \in \{0, 1\} \text{ and } h.$$

At the same time, we have shown that $\mathbb{E}_{n_k, h}[\delta_{n_k}] \rightarrow \mathbb{E}_h[\tilde{\delta}]$ for each h . Taken together, these results prove $\lim_{k \rightarrow \infty} R_{n_k}(j, \delta_{n_k}) = R(h, \tilde{\delta})$. \square

Returning to the proof of Theorem 1, let π_{Δ^*} denote the least-favorable prior, i.e., the symmetric two-point prior supported on

$$\left(-\frac{\Delta^*}{\sigma^2} I_0^{-1} \dot{\mu}_0, \frac{\Delta^*}{\sigma^2} I_0^{-1} \dot{\mu}_0 \right).$$

Clearly, for all M sufficiently large,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h/\sqrt{n}, \delta)/\lambda} \right] d\pi(h) \\ & \geq \liminf_{n \rightarrow \infty} \min_{\delta} \int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h/\sqrt{n}, \delta)/\lambda} \right] d\pi_{\Delta^*}(h). \end{aligned}$$

Let $\{\delta_n\}_n$ denote any sequence along which the $\liminf_{n \rightarrow \infty} \min_{\delta}$ on the right hand side of the above expression is attained. By the definition of π_{Δ^*} ,

$$\int \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h/\sqrt{n}, \delta_n)/\lambda} \right] d\pi_{\Delta^*}(h) = \frac{1}{2} R_n(-h^*, \delta_n) + \frac{1}{2} R_n(h^*, \delta_n).$$

Lemma 1 then implies the existence of further subsequence, $\{n_k\}_k$, and a treatment decision $\tilde{\delta}$ in the limit experiment such that

$$\begin{aligned} \frac{1}{2} R_{n_k}(-h^*, \delta_{n_k}) + \frac{1}{2} R_{n_k}(h^*, \delta_{n_k}) & \rightarrow \frac{1}{2} R(-h^*, \tilde{\delta}) + \frac{1}{2} R(h^*, \tilde{\delta}) \\ & = \int \mathbb{E}_h \left[e^{l(h, \tilde{\delta})/\lambda} \right] d\pi_{\Delta^*}(h) \\ & \geq \int \mathbb{E}_h \left[e^{l(h, \tilde{\delta}^*)/\lambda} \right] d\pi_{\Delta^*}(h) = V^*, \end{aligned}$$

where the inequality follows from the fact that $\tilde{\delta}^*$ is the best-response to the least-favorable prior π_{Δ^*} in the limit experiment, and the final equality is just the definition of minimax-risk. This completes the proof of the theorem.

A.3. Proof of Theorem 2.

Estimation. We start with the case of estimation. Consider any sequence $h_n \rightarrow h$. By (4.2), (4.3) and Assumption 3,

$$\left(\begin{array}{c} \sqrt{n}(\hat{\theta}_{\text{mle}} - \theta_0) \\ \ln \frac{dP_{n, \theta_0 + h_n/\sqrt{n}}}{dP_{n, \theta_0}} \end{array} \right) \xrightarrow{P_{n,0}} \left(\begin{array}{c} I_0^{-1/2} x \\ h^\top I_0^{1/2} x - \frac{1}{2} h^\top I_0 h \end{array} \right), \text{ where } x \sim \mathcal{N}(0, I). \quad (\text{A.3})$$

By Le Cam's third lemma,

$$\sqrt{n}(\hat{\theta}_{\text{mle}} - \theta_0) \xrightarrow{P_{n, h_n}} \mathcal{N}(h, I_0^{-1}), \quad (\text{A.4})$$

and therefore, in view of Assumption 2, it follows that for each $h_n \rightarrow h$,

$$l_n \left(\theta_0 + h_n / \sqrt{n}, \hat{\delta}_n^* \right) = \ell \left(\sqrt{n} \left(\mu(\theta_0 + h_n / \sqrt{n}) - \mu(\hat{\theta}_{\text{MLE}}) \right) \right) \\ \xrightarrow{P_{n, h_n}} \ell \left(\dot{\mu}_0^\top I_0^{-1/2} x \right), \text{ where } x \sim \mathcal{N}(0, I).$$

Since $l_n(\cdot)$ is uniformly bounded by Assumption 4, standard properties of weak convergence imply

$$\mathbb{E}_{n, h_n} \left[e^{l_n(\theta_0 + h_n / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] \rightarrow \mathbb{E}_0 \left[e^{\ell(x) / \lambda} \right] = \mathbb{E}_h \left[e^{\ell(\dot{\mu}_0^\top h - \bar{\delta}^*) / \lambda} \right] \quad (\text{A.5})$$

for every sequence $h_n \rightarrow h$. Define

$$f_n(h) := \mathbb{E}_{n, h} \left[e^{l_n(\theta_0 + h / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] \text{ and} \\ f(h) := \mathbb{E}_h \left[e^{\ell(\dot{\mu}_0^\top h - \bar{\delta}^*) / \lambda} \right] = \mathbb{E}_h \left[e^{l(h, \bar{\delta}^*) / \lambda} \right].$$

Equation (A.5) implies continuous convergence of $f_n(\cdot)$ to $f(\cdot)$, i.e., $f_n(h_n) \rightarrow f(h)$ for every $h_n \rightarrow h$. But continuous convergence on compact sets implies uniform convergence, so

$$\sup_{|h| \leq M} |f_n(h) - f(h)| \rightarrow 0. \quad (\text{A.6})$$

Now, consider a sequence of priors $\{\pi_n(h)\}_n$ along which

$$\limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n, h} \left[e^{l_n(\theta_0 + h / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] d\pi(h)$$

is attained. Since $\Delta_M(\mathcal{H})$ represents the space of compactly supported priors, it is compact under the metric of weak convergence. Hence, there exists a further sub-sequence $\{\pi_{n_j}(h)\}_j$ such that $\pi_{n_j}(h)$ converges weakly to some $\tilde{\pi}(h) \in \Delta_M(\mathcal{H})$. Furthermore, as $e^{\ell(\cdot) / \lambda}$ is uniformly bounded, so is $f(\cdot)$. Combining these observations with (A.6), standard properties of weak convergence yield

$$\limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n, h} \left[e^{l_n(\theta_0 + h / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] d\pi(h) \\ = \lim_{j \rightarrow \infty} \int f_{n_j}(h) d\pi_{n_j}(h) \\ \leq \lim_{j \rightarrow \infty} \sup_{|h| \leq M} |f_{n_j}(h) - f(h)| + \lim_{j \rightarrow \infty} \left| \int f(h) d\pi_{n_j}(h) - \int f(h) d\tilde{\pi}(h) \right| + \int f(h) d\tilde{\pi}(h) \\ = \int f(h) d\tilde{\pi}(h) \leq \max_{\pi(h) \in \Delta(\mathcal{H})} \int \mathbb{E}_h \left[e^{l(h, \bar{\delta}^*) / \lambda} \right] d\pi(h) := V^*.$$

Since the above is valid for any $M < \infty$, this completes the proof of Theorem 2 for the estimation problem.

Treatment assignment. We now turn to the case of treatment assignment. Recall that here we choose θ_0 so that $\mu(\theta_0) = 0$. Then, (A.4) and Assumption 2 imply

$$\begin{aligned} \mathbf{1} \left\{ \mu(\hat{\theta}_{\text{mle}}) \geq 0 \right\} &= \mathbf{1} \left\{ \sqrt{n} \left(\mu(\hat{\theta}_{\text{mle}}) - \mu(\theta_0) \right) \geq 0 \right\} \\ &\xrightarrow{P_{n,h_n}} \mathbf{1} \left\{ \dot{\mu}_0^\top I_0^{-1/2} x \geq 0 \right\}, \text{ where } x \sim \mathcal{N}(I_0^{-1/2} h, I). \end{aligned}$$

Hence, by standard properties of weak convergence, for each $h_n \rightarrow h$,

$$\mathbb{E}_{n,h_n} \left[\hat{\delta}_n^* \right] = P_{n,h_n} \left(\mu(\hat{\theta}_{\text{mle}}) \geq 0 \right) \rightarrow P_h \left(\dot{\mu}_0^\top I_0^{-1/2} x \geq 0 \right) = \mathbb{E}_h \left[\tilde{\delta}^* \right]. \quad (\text{A.7})$$

Note that under the treatment assignment loss,

$$\mathbb{E}_{n,h_n} \left[e^{l_n(\theta_0 + h_n / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] = e^{l_n(\theta_0 + h_n / \sqrt{n}, 1) / \lambda} \mathbb{E}_{n,h_n} \left[\hat{\delta}_n^* \right] + e^{l_n(\theta_0 + h_n / \sqrt{n}, 0) / \lambda} \mathbb{E}_{n,h_n} \left[1 - \hat{\delta}_n^* \right].$$

Now, (4.5) yields

$$l_n(\theta_0 + h_n / \sqrt{n}, a) \rightarrow l(h, a), \text{ for each } a \in \{0, 1\} \text{ and } h_n \rightarrow h.$$

Combined with (A.7), this proves

$$\begin{aligned} \mathbb{E}_{n,h_n} \left[e^{l_n(\theta_0 + h_n / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] &\rightarrow e^{l(h, 1) / \lambda} \mathbb{E}_h \left[\tilde{\delta}^* \right] + e^{l(h, 0) / \lambda} \mathbb{E}_h \left[1 - \tilde{\delta}^* \right] \\ &= \mathbb{E}_h \left[e^{l(h, \tilde{\delta}^*) / \lambda} \right], \text{ for each } h_n \rightarrow h. \end{aligned}$$

As before, define

$$\begin{aligned} f_n(h) &:= \mathbb{E}_{n,h} \left[e^{l_n(\theta_0 + h / \sqrt{n}, \hat{\delta}_n^*) / \lambda} \right] \text{ and} \\ f(h) &:= \mathbb{E}_h \left[e^{l(h, \tilde{\delta}^*) / \lambda} \right], \end{aligned}$$

and observe that $f(h)$ is bounded under treatment-assignment loss whenever $|h| \leq M$. Consequently, the remainder of the proof follows by applying the same arguments as in the case of estimation.

A.4. Proof of Theorem 3. For estimation-loss, the theorem is a direct consequence of Van der Vaart (2000, Theorem 25.21). We therefore focus on proving the theorem for treatment-assignment loss.

Recall that, via the Hilbert space isometry, we can construct an orthonormal basis $(\psi/\sigma, \phi_1, \phi_2, \dots)$ such that each $h \in T(P_0)$ can be identified with square integrable sequence of the form $(\mu/\sigma, \gamma_1, \gamma_2, \dots) \in l_2$, where $\mu = \langle \psi, h \rangle$ and $\gamma_k = \langle \phi_k, h \rangle$ for all $k \neq 0$. Consider the class of priors $\tilde{\Pi}$ on l_2 that assign a prior $\rho(\mu) \in \Delta_M(\mathbb{R})$ to μ , supported on $\{\mu : |\mu| \leq M\}$, while placing a point-mass at the origin $(\gamma_1, \gamma_2, \dots) = (0, 0, \dots)$ for the remaining components. Any $\pi \in \tilde{\Pi}$ is then equivalent to a probability distribution over $T(P_0)$ such that h takes the form $(\mu/\sigma^2)\psi$, where $\mu \sim \rho(\cdot)$.

Clearly,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(h) \in \Delta(K_M)} \int \mathbb{E}_{n,h} \left[e^{l_n(h, \delta)/\lambda} \right] d\pi(h) \\ & \geq \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\rho(\mu) \in \Delta_M(\mathbb{R})} \int \mathbb{E}_{n,h} \left[\exp \left\{ l_n \left(\frac{\mu}{\sigma^2} \psi, \delta \right) / \lambda \right\} \right] d\rho(\mu). \end{aligned}$$

Now, consider sub-models of the form $P_{1/\sqrt{n}, (\mu/\sigma)\psi/\sigma}$ for $\mu \in \mathbb{R}$. By (5.2),

$$\sum_{i=1}^n \ln \frac{P_{1/\sqrt{n}, (\mu/\sigma)\psi/\sigma}(Y_i)}{dP_0} = \frac{\mu}{\sigma\sqrt{n}} \sum_{i=1}^n \frac{\psi}{\sigma}(Y_i) - \frac{\mu^2}{2\sigma^2} + o_{P_{1/\sqrt{n}, 0}}(1). \quad (\text{A.8})$$

Comparing with (4.3), we observe that the family

$$\left\{ P_{1/\sqrt{n}, (\mu/\sigma)\psi/\sigma} : \mu \in \mathbb{R} \right\}$$

is equivalent to a parametric model with (normalized) score $\psi(\cdot)/\sigma$ and local parameter μ (observe that $\mathbb{E}_{P_0}[(\psi/\sigma)^2] = 1$). Furthermore, the second part of Assumption 5 implies

$$l_n \left(\frac{\mu}{\sigma^2} \psi, a \right) \rightarrow \mu \mathbf{1}\{\mu \geq 0\} - \mu a, \quad \text{uniformly over } a \in \{0, 1\} \text{ and bounded } \mu. \quad (\text{A.9})$$

Consequently, we can apply the same arguments as in the proof of Theorem 1 to show that

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\rho(\mu) \in \Delta_M(\mathbb{R})} \int \mathbb{E}_{n,h} \left[\exp \left\{ l_n \left(\frac{\mu}{\sigma^2} \psi, \delta \right) / \lambda \right\} \right] d\rho(\mu) \\ & \geq \min_{\delta} \max_{\rho(\mu) \in \Delta(\mathbb{R})} \int \mathbb{E}_{\mu} \left[e^{\{\mu \mathbf{1}\{\mu \geq 0\} - \mu \delta\} / \lambda} \right] d\rho(\mu). \end{aligned}$$

But by (5.5), the term on the right is just the definition of V^* .

A.5. Proof of Theorem 4.

Estimation. We start with the case of estimation. Denote $\mu_0 = \mu(P_0)$. By Assumption 6 and the central limit theorem,

$$\sqrt{n}(\hat{\mu}_n - \mu_0) \xrightarrow[P_{n,0}]{d} \mathcal{N}(0, \sigma^2). \quad (\text{A.10})$$

Consider any sequence $h_n \rightarrow h$, where convergence is in terms of the $L^2(P_0)$ norm. Recall that any $h \in T(P_0)$ admits the orthogonal decomposition $h = (\mu/\sigma)\psi/\sigma + \tilde{h}$, where $\mu := \langle \psi, h \rangle$ and \tilde{h} is orthogonal to ψ . Applying the central limit theorem again,

$$\begin{pmatrix} \frac{1}{\sqrt{n}} \sum_{i=1}^n \psi(Y_i)/\sigma \\ \frac{1}{\sqrt{n}} \sum_{i=1}^n \tilde{h}(Y_i) \end{pmatrix} \xrightarrow[P_{n,0}]{d} \begin{pmatrix} x \\ Z \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & \|\tilde{h}\|^2 \end{pmatrix} \right). \quad (\text{A.11})$$

Combining (A.10), (A.11), (5.2) and the fact $\|h\|^2 = (\mu^2/\sigma^2) + \|\tilde{h}\|^2$, we obtain

$$\begin{pmatrix} \sqrt{n}(\hat{\mu}_n - \mu_0) \\ \ln \frac{dP_{n,h_n}}{dP_{n,0}} \end{pmatrix} \xrightarrow[P_{n,0}]{d} \begin{pmatrix} \sigma x \\ \ln V \end{pmatrix}, \text{ where} \quad (\text{A.12})$$

$$V \sim \exp \left\{ \left(\frac{\mu}{\sigma} x - \frac{\mu^2}{2\sigma^2} \right) + \left(Z - \frac{1}{2} \|\tilde{h}\|^2 \right) \right\}.$$

Observe that $V \geq 0$ and $E[V] = 1$. Therefore, by an application of Le Cam's third lemma,

$$\sqrt{n}(\hat{\mu}_n - \mu_0) \xrightarrow[P_{n,h_n}]{d} \mathcal{L}; \text{ where } \mathcal{L}(B) := E[\mathbb{I}\{\sigma x \in B\}V] \quad \forall B \in \mathcal{B}(\mathbb{R}). \quad (\text{A.13})$$

In other words, for every bounded and continuous $g(\cdot)$,

$$\begin{aligned} \mathbb{E}_{n,h_n} \left[g \left(\sqrt{n}(\hat{\mu}_n - \mu_0) \right) \right] &\rightarrow E \left[g(\sigma x) \exp \left\{ \left(\frac{\mu}{\sigma} x - \frac{\mu^2}{2\sigma^2} \right) + \left(Z - \frac{1}{2} \|\tilde{h}\|^2 \right) \right\} \right] \\ &= E \left[g(\sigma x) \exp \left\{ \left(\frac{\mu}{\sigma} x - \frac{\mu^2}{2\sigma^2} \right) \right\} \right], \end{aligned}$$

where the equality is due to the fact x and Z are independent, and $Z \sim \mathcal{N}(0, \|\tilde{h}\|^2)$. Hence, applying the portmanteau lemma and standard change of measure arguments yields

$$\sqrt{n}(\hat{\mu}_n - \mu_0) \xrightarrow[P_{n,h_n}]{d} \mathcal{N}(\langle \psi, h \rangle, \sigma^2), \quad (\text{A.14})$$

for every $h_n \rightarrow h$.

Observe that, by the second part of Assumption 5,

$$\sqrt{n} \left(\mu(P_{1/\sqrt{n}, h_n}) - \mu_0 \right) = \langle \psi, h \rangle + o_{P_n, 0}(1) \quad (\text{A.15})$$

for every $h_n \rightarrow h$. Combining (A.14), (A.15) and the fact $\ell(\cdot)$ is bounded by Assumption 4, we obtain

$$\begin{aligned} l_n(h_n, \hat{\delta}_n^*) &= \ell \left(\sqrt{n} \left(\mu(P_{1/\sqrt{n}, h_n}) - \mu_0 \right) + \sqrt{n} (\mu_0 - \hat{\mu}_n) \right) \\ &\xrightarrow{P_{n, h_n}} \ell(\sigma x), \text{ where } x \sim \mathcal{N}(0, 1). \end{aligned}$$

Since $l_n(\cdot)$ is uniformly bounded, standard properties of weak convergence imply that for every $h_n \rightarrow h$,

$$\mathbb{E}_{n, h_n} \left[e^{l_n(h_n, \hat{\delta}_n^*)/\lambda} \right] \rightarrow \mathbb{E}_0 \left[e^{\ell(\sigma x)/\lambda} \right] = \mathbb{E}_{\langle \psi, h \rangle} \left[e^{\ell(\langle \psi, h \rangle - \tilde{\delta}^*)/\lambda} \right], \quad (\text{A.16})$$

where for any $\mu \in \mathbb{R}$, $\mathbb{E}_\mu[\cdot]$ represents the expectation under the limit experiment described in Section 5.2.1, and $\tilde{\delta}^* := \sigma x$, with $x \sim \mathcal{N}(\mu/\sigma, 1)$ under P_μ , is the optimal decision rule under that limit experiment.

Define

$$\begin{aligned} f_n(h) &:= \mathbb{E}_{n, h} \left[e^{l_n(h, \hat{\delta}_n^*)/\lambda} \right] \text{ and} \\ f(h) &:= \mathbb{E}_{\langle \psi, h \rangle} \left[e^{\ell(\langle \psi, h \rangle - \tilde{\delta}^*)/\lambda} \right]. \end{aligned}$$

Equation (A.16) implies continuous convergence of $f_n(\cdot)$ to $f(\cdot)$ as functions from l_2 to \mathbb{R} , i.e., $f_n(h_n) \rightarrow f(h)$ for every $h_n \rightarrow h$. But continuous convergence on compact sets implies uniform convergence, so

$$\sup_{h \in K_M} |f_n(h) - f(h)| \rightarrow 0. \quad (\text{A.17})$$

Now, consider a sequence of priors $\{\pi_n(h)\}_n$ along which

$$\limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta(K_M)} \int \mathbb{E}_{n, h} \left[e^{l_n(h, \hat{\delta}_n^*)/\lambda} \right] d\pi(h)$$

is attained. Since $K_M \subset l_2$ is a compact set, $\Delta(K_M)$ is compact under the metric of weak convergence. Hence, there exists a further sub-sequence $\{\pi_{n_j}(h)\}_j$ such that $\pi_{n_j}(h)$ converges weakly to some $\tilde{\pi}(h) \in \Delta(K_M)$. Furthermore, since $e^{\ell(\cdot)/\lambda}$

is uniformly bounded, so is $f(\cdot)$. Combined with (A.17), standard properties of weak convergence then imply

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\pi(h) \in \Delta(K_M)} \int \mathbb{E}_{n,h} \left[e^{l_n(h, \hat{\delta}_n^*)/\lambda} \right] d\pi(h) \\
&= \lim_{j \rightarrow \infty} \int f_{n_j}(h) d\pi_{n_j}(h) \\
&\leq \lim_{j \rightarrow \infty} \sup_{h \in K_M} |f_{n_j}(h) - f(h)| + \lim_{j \rightarrow \infty} \left| \int f(h) d\pi_{n_j}(h) - \int f(h) d\tilde{\pi}(h) \right| + \int f(h) d\tilde{\pi}(h) \\
&= \int f(h) d\tilde{\pi}(h) = \int \mathbb{E}_{\langle \psi, h \rangle} \left[e^{\ell(\langle \psi, h \rangle - \delta^*)/\lambda} \right] d\tilde{\pi}(h) \\
&\leq \max_{\rho(\mu) \in \Delta(\mathbb{R})} \int \mathbb{E}_{\mu} \left[e^{\ell(\mu - \delta^*)/\lambda} \right] d\rho(\mu) := V^*.
\end{aligned}$$

Since the above is valid for any $M < \infty$, this completes the proof of Theorem 4 for the estimation problem.

Treatment assignment. We now turn to the case of treatment assignment. Recall that here we choose P_0 so that $\mu_0 := \mu(P_0) = 0$. Then, (A.14), (A.15) and the second part of Assumption 5 imply

$$\begin{aligned}
\mathbf{1} \{ \hat{\mu}_n \geq 0 \} &= \mathbf{1} \{ \sqrt{n} (\hat{\mu}_n - \mu_0) \geq 0 \} \\
&\xrightarrow{P_{n, h_n}} \mathbf{1} \{ \sigma x \geq 0 \}, \text{ where } x \sim \mathcal{N}(\langle \psi, h \rangle / \sigma, 1).
\end{aligned}$$

Hence, by standard properties of weak convergence, for each $h_n \rightarrow h$, we have

$$\mathbb{E}_{n, h_n} [\hat{\delta}_n^*] = P_{n, h_n} (\hat{\mu}_n \geq 0) \rightarrow P_{\langle \psi, h \rangle} (\sigma x \geq 0) = \mathbb{E}_{\langle \psi, h \rangle} [\tilde{\delta}^*], \quad (\text{A.18})$$

where for any $\mu \in \mathbb{R}$, P_{μ} and $\mathbb{E}_{\mu}[\cdot]$ represent the probabilities and expectations under the limit experiment described in Section 5.2.1, and $\tilde{\delta}^* := \mathbf{1} \{ \sigma x \geq 0 \}$ — with $x \sim \mathcal{N}(\mu/\sigma, 1)$ under P_{μ} — is the optimal decision rule under that limit experiment.

Note that under the treatment assignment loss,

$$\mathbb{E}_{n, h_n} \left[e^{l_n(h_n, \hat{\delta}_n^*)/\lambda} \right] = e^{l_n(h_n, 1)/\lambda} \mathbb{E}_{n, h_n} [\hat{\delta}_n^*] + e^{l_n(h_n, 0)/\lambda} \mathbb{E}_{n, h_n} [1 - \hat{\delta}_n^*].$$

Now, (4.5) implies

$$l_n(h_n, a) \rightarrow l(\langle \psi, h \rangle, a), \text{ for each } a \in \{0, 1\} \text{ and } h_n \rightarrow h,$$

where $l(\cdot, \cdot)$ denotes the treatment assignment loss under the limit experiment, as defined in (5.5). Combined with (A.18), this proves

$$\begin{aligned} \mathbb{E}_{n, h_n} \left[e^{l_n(h_n, \hat{\delta}_n^*)/\lambda} \right] &\rightarrow e^{l(\langle \psi, h \rangle, 1)/\lambda} \mathbb{E}_{\langle \psi, h \rangle} \left[\tilde{\delta}^* \right] + e^{l(\langle \psi, h \rangle, 0)/\lambda} \mathbb{E}_{\langle \psi, h \rangle} \left[1 - \tilde{\delta}^* \right] \\ &= \mathbb{E}_{\langle \psi, h \rangle} \left[e^{l(\langle \psi, h \rangle, \tilde{\delta}^*)/\lambda} \right], \text{ for each } h_n \rightarrow h. \end{aligned}$$

As before, define

$$\begin{aligned} f_n(h) &:= \mathbb{E}_{n, h} \left[e^{l_n(h, \hat{\delta}_n^*)/\lambda} \right] \text{ and} \\ f(h) &:= \mathbb{E}_{\langle \psi, h \rangle} \left[e^{l(\langle \psi, h \rangle, \tilde{\delta}^*)/\lambda} \right], \end{aligned}$$

and observe that $f(h)$ is bounded under treatment-assignment loss whenever $h \in K_M$ (which implies $\langle \psi, h \rangle \leq M$). Consequently, the remainder of the proof follows by applying the same arguments as in the case of estimation.

APPENDIX B. LOCAL ASYMPTOTICS WITH GLOBAL PRIORS

We can allow unrestricted prior and reference parameter choice by employing uniform versions of local asymptotic normality and Assumption 2. In what follows, the parameter space, Θ , is assumed to be a compact set. Let $P_{n,\theta}$ represent the joint probability measure over the iid Y_1, \dots, Y_n when each $Y_i \sim P_\theta$, and let $\mathbb{E}_{n,\theta}[\cdot]$ denote the corresponding expectation.

Assumption A1. *The class $\{P_\theta : \theta \in \Theta\}$ satisfies a uniform LAN property, i.e., there exists a score function $\psi_\theta(\cdot)$ and information matrix $I_\theta := \mathbb{E}_\theta[\psi_\theta \psi_\theta^\top]$ such that for each $\theta_n \rightarrow \theta \in \Theta$ and $h_n \rightarrow h \in \mathbb{R}^d$,*

$$\ln \frac{dP_{n,\theta_n+h_n/\sqrt{n}}}{dP_{n,\theta_n}} = h^\top I_{\theta_n}^{1/2} x_{n,\theta_n} - \frac{1}{2} h^\top I_{\theta_n} h + o_{P_{n,\theta_n}}(1),$$

where

$$x_{n,\theta_n} := \frac{I_{\theta_n}^{-1/2}}{\sqrt{n}} \sum_{i=1}^n \psi_{\theta_n}(Y_i) \xrightarrow{P_{n,\theta_n}} \mathcal{N}(0, I).$$

Furthermore, the information matrix $I_\theta := \mathbb{E}_\theta[\psi_\theta \psi_\theta^\top]$ is invertible, continuous in θ and $0 < \inf_\theta \lambda_{\min}(I_\theta^{-1}) < \sup_\theta \lambda_{\max}(I_\theta^{-1}) < \infty$.⁵

Assumption A2. *The function $\mu(\cdot)$ is Lipschitz continuous over Θ . Specifically, there exists $\dot{\mu}_\theta \in \mathbb{R}^d$ and $\epsilon_n \rightarrow 0$ independent of θ such that $\sqrt{n}(\mu(\theta + h/\sqrt{n}) - \mu(\theta)) = \dot{\mu}_\theta^\top h + \epsilon_n |h|^2$ uniformly over all $\theta \in \Theta$ and bounded h . Furthermore, for the treatment-assignment problem, $\dot{\mu}_\theta$ is continuous in θ and $\inf_{\theta \in \tilde{\Theta}} \|\dot{\mu}_\theta\| > 0$, where $\tilde{\Theta} := \{\theta \in \Theta : \mu(\theta) = 0\}$.*

Assumption A1 follows Ibragimov and Hasminskii (1981, Definition 2.2). Primitive conditions for this assumption can be found in Ibragimov and Hasminskii (1981, Sections 2.6 & 2.7); essentially, one needs a uniform version of quadratic mean differentiability. Assumption A2 is a uniform version of Assumption 2. The main additional requirement is that we need the derivative $\dot{\mu}_\theta$ to be non-zero on the zero-set of $\mu(\cdot)$, i.e., whenever $\mu(\theta) = 0$.

We now define a reference parameter dependent limit experiment. Suppose that the decision-maker observes a d -dimensional signal x , posited to be drawn from

⁵Footnote, here $\lambda_{\min}(A)$ and $\lambda_{\max}(A)$ represent the minimum and maximum eigenvalues of a matrix A .

a reference Gaussian likelihood, $P_{\theta,h}(x) \sim \mathcal{N}(I_\theta^{-1/2}h, I)$. Let V_θ^* represent the parameter dependent minimal decision-risk in this experiment:

$$V_\theta^* := \min_{\tilde{\delta}} \max_{\pi(h) \in \Delta(\mathbb{R})} \int \mathbb{E}_{\theta,h} \left[e^{l_\theta(h,\tilde{\delta})/\lambda} \right] d\pi(h), \quad \text{with} \quad (\text{B.1})$$

$$l_\theta(h, \tilde{\delta}) = \begin{cases} \ell(\dot{\mu}_\theta^\top h - \tilde{\delta}) & \text{for estimation loss,} \\ \dot{\mu}_\theta^\top h \left\{ \mathbf{1}\{\dot{\mu}_\theta^\top h \geq 0\} - \tilde{\delta} \right\} & \text{for treatment-assignment loss.} \end{cases}$$

Observe that the corresponding optimal decisions are $\tilde{\delta}_\theta^* = \dot{\mu}_\theta^\top I_\theta^{-1/2} x$ for estimation and $\tilde{\delta}_\theta^* = \mathbf{1}\{\dot{\mu}_\theta^\top I_\theta^{-1/2} x \geq 0\}$ for treatment assignment.

We then obtain the following lower bound under global priors:

Theorem 5. *Suppose that Assumptions A1 and A2 hold. Then, under both the estimation and treatment-assignment loss functions,*

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\tilde{\delta}} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_n(\theta,\tilde{\delta})/\lambda} \right] d\pi(\theta) \\ & \geq \begin{cases} \sup_{\theta \in \Theta} V_\theta^* & \text{for estimation loss,} \\ \sup_{\{\theta \in \Theta: \mu(\theta)=0\}} V_\theta^* & \text{for treatment-assignment loss.} \end{cases} \end{aligned}$$

To show that the MLE based decisions in (4.7) also remain asymptotically optimal under global priors, we impose a stronger assumption on the properties of MLE:

Assumption A3. *The maximum-likelihood estimator $\hat{\theta}_{mle}$ admits a uniform locally linear score-function approximation, i.e., for any $\theta_n \rightarrow \theta \in \Theta$,*

$$\sqrt{n} I_\theta^{1/2} \left(\hat{\theta}_{mle} - \theta_n \right) \xrightarrow{P_{n,\theta_n}} \mathcal{N}(0, I).$$

Furthermore, for any $\epsilon > 0$, there exists $M < \infty$ such that

$$\sup_{\theta} P_{n,\theta} \left(\left| \sqrt{n} \left(\mu(\hat{\theta}_{mle}) - \mu(\theta) \right) \right| > M \right) \leq \epsilon.$$

Sufficient conditions for Assumption A3 can be found in Ibragimov and Hasminskii (1981, Theorem 3.1).

As in Section 4.4, we require that $\ell(\cdot)$ be bounded. Additionally, we also truncate the loss for the treatment-assignment problem to avoid issues relating to the non-existence of moments. We state these requirements as an additional assumption below.

Assumption A4. *The function $\ell(\cdot)$ is bounded. Additionally, for the treatment assignment problem, we replace $l_n(\theta, \delta)$ with the truncated loss $l_{n,K}(\theta, \delta) = K \wedge l_n(\theta, \delta)$.*

Theorem 6. *Suppose that Assumptions A1-A4 hold. Then, under both the estimation and treatment-assignment loss functions,*

$$\begin{aligned} & \lim_{K \rightarrow \infty} \limsup_{n \rightarrow \infty} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta,\delta)/\lambda} \right] d\pi(\theta) \\ &= \begin{cases} \sup_{\theta \in \Theta} V_\theta^* & \text{for estimation loss,} \\ \sup_{\{\theta \in \Theta: \mu(\theta)=0\}} V_\theta^* & \text{for treatment-assignment loss.} \end{cases} \end{aligned}$$

B.1. Proof of Theorem 5. Observe that for any reference parameter θ ,

$$\begin{aligned} & \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_n(\theta,\delta)/\lambda} \right] d\pi(\theta) \\ & \geq \liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,\theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n},\delta)/\lambda} \right] d\pi(h). \end{aligned}$$

For the case of treatment assignment, we choose $\theta \in \{\theta \in \Theta : \mu(\theta) = 0\}$. Then, making use of Assumptions A1 and A2, we can employ the same arguments as in the proof of Theorem 1 to show that

$$\liminf_{n \rightarrow \infty} \min_{\delta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,\theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n},\delta)/\lambda} \right] d\pi(h) \geq V_\theta^*.$$

The claim thus follows since the above holds for any $\theta \in \Theta$ under estimation loss, and any $\theta \in \{\theta \in \Theta : \mu(\theta) = 0\}$ under treatment-assignment loss.

B.2. Proof of Theorem 6.

Estimation. We start with the case of estimation. Consider any sequence $\theta_n \rightarrow \theta \in \Theta$ and $h_n \rightarrow h$. By (4.2), (4.3) and Assumption 3,

$$\left(\begin{array}{c} \sqrt{n}(\hat{\theta}_{\text{mle}} - \theta_n) \\ \ln \frac{dP_{n,\theta_n+h_n/\sqrt{n}}}{dP_{n,\theta_n}} \end{array} \right) \xrightarrow{P_{n,\theta_n}} \left(\begin{array}{c} I_\theta^{-1/2} x \\ h^\top I_\theta^{1/2} x - \frac{1}{2} h^\top I_\theta h \end{array} \right), \text{ where } x \sim \mathcal{N}(0, I). \quad (\text{B.2})$$

Le Cam's third lemma then yields

$$\sqrt{n}(\hat{\theta}_{\text{mle}} - \theta_n) \xrightarrow{P_{n,\theta_n+h_n/\sqrt{n}}} \mathcal{N}(h, I_\theta^{-1}). \quad (\text{B.3})$$

Therefore, in view of Assumption A2, it follows that for each $\theta_n \rightarrow \theta$ and $h_n \rightarrow h$,

$$\begin{aligned} l_n(\theta_n + h_n/\sqrt{n}, \hat{\delta}_n^*) &= \ell\left(\sqrt{n}\left(\mu(\theta_n + h_n/\sqrt{n}) - \mu(\hat{\theta}_{\text{mle}})\right)\right) \\ &= \left(\sqrt{n}\left(\mu(\theta_n + h_n/\sqrt{n}) - \mu(\theta_n)\right) + \sqrt{n}\left(\mu(\theta_n) - \mu(\hat{\theta}_{\text{mle}})\right)\right) \\ &\xrightarrow{P_{n,\theta_n+h_n/\sqrt{n}}} \ell\left(\dot{\mu}_\theta^\top I_\theta^{-1/2} x\right), \text{ where } x \sim \mathcal{N}(0, I). \end{aligned}$$

Since $l_n(\cdot)$ is uniformly bounded by Assumption A4, standard properties of weak convergence imply

$$\mathbb{E}_{n,\theta_n+h_n/\sqrt{n}} \left[e^{l_n(\theta_n+h_n/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] \rightarrow \mathbb{E}_{\theta,0} \left[e^{\ell(x)/\lambda} \right] = \mathbb{E}_{\theta,h} \left[e^{\ell(\dot{\mu}_\theta^\top h - \tilde{\delta}_\theta^*)/\lambda} \right] \quad (\text{B.4})$$

for every sequence $(\theta_n, h_n) \rightarrow (\theta, h)$. Define

$$\begin{aligned} f_n(\theta, h) &:= \mathbb{E}_{n,\theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] \text{ and} \\ f(\theta, h) &:= \mathbb{E}_{\theta,h} \left[e^{\ell(\dot{\mu}_\theta^\top h - \tilde{\delta}_\theta^*)/\lambda} \right] = \mathbb{E}_{\theta,h} \left[e^{l_\theta(h, \tilde{\delta}_\theta^*)/\lambda} \right]. \end{aligned}$$

Equation (B.4) implies continuous convergence of $f_n(\cdot)$ to $f(\cdot)$, i.e., $f_n(\theta_n, h_n) \rightarrow f(\theta, h)$ for every $(\theta_n, h_n) \rightarrow (\theta, h)$. But continuous convergence on compact sets implies uniform convergence, so

$$\sup_{(\theta,h) \in \Theta \times [-M,M]} |f_n(\theta, h) - f(\theta, h)| \rightarrow 0. \quad (\text{B.5})$$

Now, observe that for any $M < \infty$,

$$\begin{aligned} &\limsup_{n \rightarrow \infty} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_n(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta) \\ &\leq \limsup_{n \rightarrow \infty} \max_{\theta \in \Theta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n,\theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] d\pi(h). \end{aligned}$$

Consider a sequence $\{(\theta_n, \pi_n(h))\}_n$ along which the limsup on the right hand side is attained. Since $\Delta_M(\mathcal{H})$ represents the space of compactly supported priors, it is compact under the metric of weak convergence. Hence, there exists a further sub-sequence $\{(\theta_{n_j}, \pi_{n_j}(h))\}_j$ such that $(\theta_{n_j}, \pi_{n_j}(h))$ converges weakly to some $(\tilde{\theta}, \tilde{\pi}(h)) \in \Theta \times \Delta_M(\mathcal{H})$. Furthermore, as $e^{\ell(\cdot)/\lambda}$ is uniformly bounded, so is $f(\cdot)$. Combining these observations with (B.5), standard properties of weak convergence yield

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\theta \in \Theta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n, \theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] d\pi(h) \\
&= \lim_{j \rightarrow \infty} \int f_{n_j}(\theta_{n_j}, h) d\pi_{n_j}(h) \\
&\leq \lim_{j \rightarrow \infty} \sup_{\theta \in \Theta} \sup_{|h| \leq M} |f_{n_j}(\theta, h) - f(\theta, h)| \\
&\quad + \lim_{j \rightarrow \infty} \left| \int f(\theta_{n_j}, h) d\pi_{n_j}(h) - \int f(\theta_{n_j}, h) d\tilde{\pi}(h) \right| + \lim_{j \rightarrow \infty} \int f(\theta_{n_j}, h) d\tilde{\pi}(h) \\
&= \int \lim_{j \rightarrow \infty} f(\theta_{n_j}, h) d\tilde{\pi}(h).
\end{aligned}$$

Now, observe that by Assumption A1 (continuity of I_θ) and standard properties of the Gaussian distribution, $\lim_{j \rightarrow \infty} f(\theta_{n_j}, h) = f(\tilde{\theta}, h)$ for each h . Consequently,

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\theta \in \Theta} \max_{\pi(h) \in \Delta_M(\mathcal{H})} \int \mathbb{E}_{n, \theta+h/\sqrt{n}} \left[e^{l_n(\theta+h/\sqrt{n}, \hat{\delta}_n^*)/\lambda} \right] d\pi(h) \\
&= \int f(\tilde{\theta}, h) d\tilde{\pi}(h) \leq \sup_{\theta \in \Theta} \max_{\pi(h) \in \Delta(\mathcal{H})} \int \mathbb{E}_{\theta, h} \left[e^{l_\theta(h, \hat{\delta}_\theta^*)/\lambda} \right] d\pi(h) := \sup_{\theta \in \Theta} V_\theta^*.
\end{aligned}$$

This proves Theorem 6 for the estimation problem.

Treatment assignment. Recall the definition $\tilde{\Theta} := \{\theta \in \Theta : \mu(\theta) = 0\}$ from Assumption A2 and note that $\tilde{\Theta}$ a compact set due to the Lipschitz continuity of $\mu(\theta)$, as imposed in Assumption A2. In addition, denote

$$\begin{aligned}
\tilde{\Theta}_{n, M} &:= \{\theta \in \Theta : |\mu(\theta)| \leq M/\sqrt{n}\}, \\
\tilde{\Theta}_{n, M}^+ &:= \{\theta \in \Theta : \mu(\theta) > M/\sqrt{n}\}, \text{ and} \\
\tilde{\Theta}_{n, M}^- &:= \{\theta \in \Theta : \mu(\theta) < -M/\sqrt{n}\}.
\end{aligned}$$

We can decompose

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta) \\
& \leq \limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}^+} \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] \\
& \quad \vee \limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}^-} \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] \vee \limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}^-} \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right]. \quad (\text{B.6})
\end{aligned}$$

Note that $l_{n,K}(\theta, \hat{\delta}_n^*) \leq K$. Then, by the second part of Assumption A4, we may choose $M < \infty$ large enough such that

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}^+} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta) \\
& \leq \limsup_{n \rightarrow \infty} \left\{ e^{K/\lambda} \sup_{\theta \in \Theta} P_{n,\theta} \left(\sqrt{n} \left(\mu(\hat{\theta}_{\text{mle}}) - \mu(\theta) \right) < -M \right) + 1 \right\} \leq 1 + \eta,
\end{aligned}$$

for any η arbitrarily small. In a similar vein,

$$\limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}^-} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta) \leq 1 + \eta.$$

It remains to analyze the term

$$\limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}_{n,M}} \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] = \limsup_{n \rightarrow \infty} \max_{\pi \in \Delta(\tilde{\Theta}_{n,M})} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta).$$

By Dontchev and Rockafellar (2009), Lipschitz continuity of $\mu(\theta)$ and $\inf_{\theta \in \tilde{\Theta}} \|\dot{\mu}_\theta\| > 0$ (both required under Assumption A2) imply metric regularity of $\mu(\cdot)$ near its zero set, i.e., $|\mu(\theta)| > c \cdot d(\theta, \tilde{\Theta})$ for some $c > 0$. Consequently, there exists $L := M/c < \infty$ such that

$$\tilde{\Theta}_{n,M} \subseteq \left\{ \theta + h/\sqrt{n} : \theta \in \tilde{\Theta}, |h| \leq L \right\}.$$

Hence,

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \max_{\pi \in \Delta(\tilde{\Theta}_{n,M})} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(\theta) \\
& \leq \limsup_{n \rightarrow \infty} \max_{\theta \in \tilde{\Theta}} \max_{\pi \in \Delta([-L, L])} \int \mathbb{E}_{n,\theta} \left[e^{l_{n,K}(\theta, \hat{\delta}_n^*)/\lambda} \right] d\pi(h).
\end{aligned}$$

Consider any sequence $(\theta_n, h_n) \in \tilde{\Theta} \times [-L, L]$ such that $(\theta_n, h_n) \rightarrow (\theta, h)$. Since $\tilde{\Theta}$ is a compact set, $\theta \in \tilde{\Theta}$, i.e., $\mu(\theta) = 0$. Then, (B.2) and Assumption A2 imply

$$\mathbf{1}\{\mu(\hat{\theta}_{\text{mle}}) \geq 0\} = \mathbf{1}\{\sqrt{n}(\mu(\hat{\theta}_{\text{mle}}) - \mu(\theta_n + h_n/\sqrt{n})) + \sqrt{n}(\mu(\theta_n + h_n/\sqrt{n}) - \mu(\theta)) \geq 0\} \\ \xrightarrow{P_{n, \theta_n + h_n/\sqrt{n}}} \mathbf{1}\{\dot{\mu}_\theta^\top I_\theta^{-1/2} x \geq 0\}, \text{ where } x \sim \mathcal{N}(I_\theta^{-1/2} h, I).$$

Hence, by standard properties of weak convergence, for each $(\theta_n, h_n) \in \tilde{\Theta} \times [-L, L] \rightarrow (\theta, h)$,

$$\mathbb{E}_{n, \theta_n + h_n/\sqrt{n}} [\hat{\delta}_n^*] = P_{n, \theta_n + h_n/\sqrt{n}}(\mu(\hat{\theta}_{\text{mle}}) \geq 0) \rightarrow P_{\theta, h}(\dot{\mu}_\theta^\top I_\theta^{-1/2} x \geq 0) = \mathbb{E}_{\theta, h} [\tilde{\delta}_\theta^*]. \quad (\text{B.7})$$

Note that for the treatment assignment loss,

$$\mathbb{E}_{n, \theta_n + h_n/\sqrt{n}} [e^{l_{n, K}(\theta_0 + h_n/\sqrt{n}, \hat{\delta}_n^*)/\lambda}] \\ = e^{l_{n, K}(\theta_0 + h_n/\sqrt{n}, 1)/\lambda} \mathbb{E}_{n, \theta_n + h_n/\sqrt{n}} [\hat{\delta}_n^*] + e^{l_{n, K}(\theta_0 + h_n/\sqrt{n}, 0)/\lambda} \mathbb{E}_{n, \theta_n + h_n/\sqrt{n}} [1 - \hat{\delta}_n^*].$$

Now, Assumption A2 yields

$$l_n(\theta_n + h_n/\sqrt{n}, a) \rightarrow l_{\theta, K}(h, a), \text{ for each } a \in \{0, 1\} \text{ and } (\theta_n, h_n) \rightarrow (\theta, h),$$

where

$$l_{\theta, K}(h, a) := K \wedge \dot{\mu}_\theta^\top h \{\mathbf{1}\{\dot{\mu}_\theta^\top h \geq 0\} - a\}.$$

Combined with (B.7), this proves

$$\mathbb{E}_{n, \theta_n + h_n/\sqrt{n}} [e^{l_n(\theta_n + h_n/\sqrt{n}, \hat{\delta}_n^*)/\lambda}] \\ \rightarrow e^{l_{\theta, K}(h, 1)/\lambda} \mathbb{E}_{\theta, h} [\tilde{\delta}_\theta^*] + e^{l_{\theta, K}(h, 0)/\lambda} \mathbb{E}_{\theta, h} [1 - \tilde{\delta}_\theta^*] \\ = \mathbb{E}_{\theta, h} [e^{l_{\theta, K}(h, \tilde{\delta}_\theta^*)/\lambda}], \text{ for each } (\theta_n, h_n) \in \tilde{\Theta} \times [-L, L] \rightarrow (\theta, h).$$

As before, define $f_n, f : \tilde{\Theta} \times [-L, L] \rightarrow \mathbb{R}$ such that

$$f_n(\theta, h) := \mathbb{E}_{n, \theta + h/\sqrt{n}} [e^{l_{n, K}(\theta + h/\sqrt{n}, \hat{\delta}_n^*)/\lambda}] \text{ and} \\ f(\theta, h) := \mathbb{E}_h [e^{l_{\theta, K}(h, \tilde{\delta}_\theta^*)/\lambda}].$$

Observe that $f(\theta, h)$ is bounded under treatment-assignment loss by construction. Consequently, by similar arguments as in the case of estimation, it follows

$$\begin{aligned}
& \limsup_{n \rightarrow \infty} \sup_{\theta \in \tilde{\Theta}} \max_{\pi \in \Delta([-L, L])} \int \mathbb{E}_{n, \theta} \left[e^{l_{n, K}(\theta, \hat{\delta}_n^*) / \lambda} \right] d\pi(h) \\
& \leq \sup_{\theta \in \tilde{\Theta}} \max_{\pi(h) \in \Delta(\mathcal{H})} \int \mathbb{E}_{\theta, h} \left[e^{l_{\theta, K}(h, \tilde{\delta}_\theta^*) / \lambda} \right] d\pi(h) \\
& \leq \sup_{\theta \in \tilde{\Theta}} \max_{\pi(h) \in \Delta(\mathcal{H})} \int \mathbb{E}_{\theta, h} \left[e^{l_\theta(h, \tilde{\delta}_\theta^*) / \lambda} \right] d\pi(h) := \sup_{\theta \in \tilde{\Theta}} V_\theta^*.
\end{aligned}$$

Observe that, by definition, $V_\theta^* > 1$ for any $\theta \in \tilde{\Theta}$. Combined with (B.6) and the fact $\eta > 0$ can be set arbitrarily small, it follows that

$$\limsup_{n \rightarrow \infty} \max_{\pi(\theta) \in \Delta(\Theta)} \int \mathbb{E}_{n, \theta} \left[e^{l_{n, K}(\theta, \hat{\delta}_n^*) / \lambda} \right] d\pi(\theta) \leq \sup_{\theta \in \tilde{\Theta}} V_\theta^* \vee K,$$

as stated by the theorem.

APPENDIX C. ASYMMETRIC LOSS FUNCTIONS: THE CASE OF LINEX

In the limit experiment, the linex loss takes the form

$$l(h, \tilde{\delta}) = e^{(\dot{\mu}_0^\top h - \tilde{\delta})} - (\dot{\mu}_0^\top h - \tilde{\delta}) - 1.$$

We start by analyzing the minimax optimal decision under no misspecification. Observe that $\dot{\mu}_0^\top I_0^{-1/2} x$ is a sufficient statistic for $\dot{\mu}_0^\top h$. Also, the loss function is convex and location-invariant, in the sense that $l(h + z, \tilde{\delta} + \dot{\mu}_0^\top z)$ for all $z \in \mathbb{R}^d$. Consequently, the Hunt-Stein theorem implies that the minimax optimal estimator should be the minimum risk equivariant estimator. Any equivariant estimator in this setting must be of the form $\tilde{\delta}_z(x) = \dot{\mu}_0^\top (I_0^{-1/2} x + z)$. The frequentist risk of $\tilde{\delta}_z$ is

$$\begin{aligned} R(h, \tilde{\delta}_z) &= e^{-\dot{\mu}_0^\top z} \mathbb{E}_h \left[e^{(\dot{\mu}_0^\top h - \dot{\mu}_0^\top I_0^{-1/2} x)} \right] + \dot{\mu}_0^\top z - 1 \\ &= e^{-\dot{\mu}_0^\top z} e^{\dot{\mu}_0^\top I_0^{-1} \dot{\mu}_0} + \dot{\mu}_0^\top z - 1. \end{aligned}$$

Optimizing over the value of z — or equivalently, that of $\dot{\mu}_0^\top z$ — we find that the frequentist risk is minimized at $\dot{\mu}_0^\top z^* = \frac{1}{2} \dot{\mu}_0^\top I_0^{-1} \dot{\mu}_0$. Consequently, in the absence of misspecification, the minimax optimal estimator takes the form

$$\tilde{\delta}^* = \dot{\mu}_0^\top I_0^{-1/2} x + \frac{1}{2} \dot{\mu}_0^\top I_0^{-1} \dot{\mu}_0.$$

Recall that optimal decisions under misspecification are equivalent to minimax optimal decisions with an exponential tilt of the loss function. Since the resulting risk would be infinite under the Gaussian experiment, we truncate the linex loss at some large value M to obtain $l_M(h, \tilde{\delta}) = \min \{l(h, \tilde{\delta}), M\}$. The decision-risk under misspecification is then given by

$$V^* = \min_{\tilde{\delta}} \max_{\pi(h)} \int \mathbb{E}_h \left[e^{l_M(h, \tilde{\delta})/\lambda} \right] d\pi(h).$$

Observe that $\exp\{l_M(h, \tilde{\delta})/\lambda\}$ is still convex and location-invariant. Consequently, we can apply the Hunt-Stein theorem again to conclude that the minimax optimal estimator must be of the form $\tilde{\delta}_z(x) = \dot{\mu}_0^\top I_0^{-1/2} x + \dot{\mu}_0^\top z$. Since the frequentist risks, $R_\lambda(h, \tilde{\delta}_z) := \mathbb{E}_h \left[e^{l_M(h, \tilde{\delta}_z)/\lambda} \right]$, of equivariant estimators are independent of h , we

have

$$R_\lambda(h, \tilde{\delta}_z) = R_\lambda(0, \tilde{\delta}_z) = \mathbb{E}_0 \left[e^{l_M(0, \tilde{\delta}_z)/\lambda} \right].$$

The optimal value of $\dot{\mu}_0^\top z^*$ therefore solves

$$\begin{aligned} \dot{\mu}_0^\top z^* = \Delta_M^*(\lambda) &:= \arg \min_{\Delta} \mathbb{E}_0 \left[e^{l_M(0, \dot{\mu}_0^\top I_0^{-1/2} x + \Delta)/\lambda} \right] \\ &= \arg \min_{\Delta} \mathbb{E}_Y \left[e^{l_M(0, Y + \Delta)/\lambda} \right], \text{ where } Y \sim \mathcal{N}(0, \dot{\mu}_0^\top I_0^- \dot{\mu}_0). \end{aligned}$$

Applying the implicit function theorem, some tedious but straightforward algebra shows that $\partial_\lambda \Delta_M^*(\lambda) < 0$, i.e., the minimax optimal shift decreases in λ . Indeed, as $M \rightarrow \infty$ and $\lambda \rightarrow 0$ — i.e., misspecification risk decreases — the minimax optimal shift converges to that under no misspecification:

$$\lim_{M \rightarrow \infty} \lim_{\lambda \rightarrow 0} \Delta_M^*(\lambda) = \frac{1}{2} \dot{\mu}_0^\top I_0^{-1} \dot{\mu}_0.$$

On the other hand, when $\lambda \rightarrow \infty$ — i.e., misspecification risk explodes — we find that $\Delta_M^*(\lambda) \rightarrow \infty$ and the minimax optimal estimator also becomes ∞ . Hence, in the case of linear loss, the optimal estimator depends on the degree of misspecification.