

Improving Robust Decisions with Data*

Xiaoyu Cheng[†]

February 10, 2026

Abstract

A decision-maker faces uncertainty governed by a data-generating process (DGP), which is only known to belong to a set of sequences of independent but possibly non-identical distributions. A robust decision maximizes the expected payoff against the worst possible DGP in this set. This paper characterizes when and how such robust decisions can be *objectively* improved with data — that is, yield higher expected payoffs under the true DGP regardless of which DGP is the truth. It further develops simple and novel inference procedures that achieve such improvement, while common methods (e.g., maximum likelihood) may fail to do so.

JEL: D81, D83, C44

Keywords: Maxmin expected utility, Γ -minimax, learning under ambiguity, statistical decisions

*This paper is a revised version of a chapter in my Northwestern University doctoral dissertation. The extended abstract of an earlier version appeared in EC'22. I am deeply indebted to Peter Klibanoff and Marciano Siniscalchi for their guidance and support. Comments from the editor and four anonymous referees have greatly improved the paper. I also thank Federico Bugni, Ivan Canay, Tristan Gagnon-Bartsch, Shaowei Ke, Yonggyun Kim, R.Vijay Krishna, Yijun Liu, Edwin Munoz Rodriguez, Isaias N.Chaves, Eran Shmaya, and Chen Zhao for insightful comments.

[†]Department of Economics, Florida State University, Tallahassee, FL, USA. E-mail: xcheng@fsu.edu.

1 Introduction

When a decision-maker (DM) lacks sufficient knowledge about the probability law governing an uncertain environment, they may want their decisions to be robust. This concern for robustness is often modeled by ranking choices by their worst-case expected payoffs among a set of possible laws (Blum and Rosenblatt, 1967; Gilboa and Schmeidler, 1989; Hansen and Sargent, 2007; Carroll, 2019). With insufficient knowledge, gathering data is a classic means of learning about the environment to help guide decisions. Here, learning involves using data to revise the set of possible laws. A revision with data is said to provide an *objective improvement* if the data-revised robust decision yields a higher expected payoff under the *true* law than the robust decision without any data-revision, the benchmark robust decision. This paper studies when and how data may be used to achieve such an objective improvement regardless of which possible law is the true one.

The analysis is conducted in a decision environment where it is a priori unclear whether and how data can generate objective improvement — one in which the true law cannot be uniquely identified even asymptotically. Specifically, the DM faces a countable sequence of random experiments (e.g., coin flips) that share the same set of outcomes. The probability law governing these realizations, referred to as the data-generating process (DGP), is a sequence of independent but possibly non-identical distributions over the outcomes. The DM initially knows there is a set of possible DGPs that contains the true one, observes sample data given by realizations of some experiments, and then makes a decision whose payoff depends only on future outcomes.

This non-identical decision environment captures settings with unobserved heterogeneity. For example, consider an online platform, say Netflix, making content recommendations to users based on feedback from other users with similar profiles. Users' preferences, however, are usually also determined by their various offline activities that cannot be observed, which often further affect how they interact with the platform. As a result, the sequential user feedback data collected by Netflix can be viewed as generated by a sequence of independent but possibly non-identical distributions.¹ In this case, Netflix is motivated to use its data in a robust manner to guard against the concern that the data could be more from one type of user, whereas future users that determine payoffs could be more of a different type.

Regardless of what data Netflix/the DM observes, they can always rely on their initial

¹See Cao (2014) for a discussion on the issue of non-IIDness in online recommender systems.

belief to make the benchmark decision, which is robust against all initially possible DGPs. Notice, because the experiments are independent, the benchmark decision does not depend on the data at all.² As a result, the benchmark decision ensures robustness, but at the cost of ignoring any information that might lead to better decisions.

The goal of this paper is to develop a strategy for using data that addresses both the concern for robustness and the desire for improvement. To preview the final output, the proposed strategy ensures that, in a broad class of decision problems, no matter which DGP in the initial set is true, a sufficient amount of data will almost surely lead the data-revised decision to objectively improve upon the benchmark decision, and sometimes strictly so. With finite samples, the same holds with a pre-specified probability. Hence, this strategy enables the DM to achieve higher expected payoffs under the truth while preserving robustness, i.e., failing to do so only with either zero or a small pre-specified probability.

What must be true for such a strategy? The following example illustrates a key observation of this paper.

Example (Introductory Example). *Suppose Netflix is deciding how to recommend a movie to a population of users who might either like it (thumbs-Up) or dislike it (thumbs-Down). Let $\{U, D\}$ denote these two possible outcomes observed after a recommendation. Let full state space be $\{U, D\}^\infty$, i.e., the infinite Cartesian product of these outcomes. Denote any probability distribution over $\{U, D\}$ by the probability of U . A data-generating process, P , can be written as $P = P_1 \times P_2 \times \dots$ with $P_i \in [0, 1]$. Suppose Netflix's initial knowledge about the users is given by the set $\{(1/3)^\infty\} \cup \{3/5, 1\}^\infty$, where $(1/3)^\infty$ denotes an i.i.d. sequence with marginal probability $1/3$ and the other set is defined as*

$$\{3/5, 1\}^\infty \equiv \{P : P_i \in \{3/5, 1\}\}.$$

Intuitively, users share some similarities: they either all like the movie with the same low probability $1/3$, or all like it with high but possibly heterogeneous probabilities, $3/5$ or 1 . Suppose Netflix observes feedback from having recommended the movie to N users and must decide how to recommend it to future users.

²With independence, the set of conditional distributions over future outcomes is solely determined by the set of DGPs regardless of the realized outcomes. Independence is assumed because this fact eases the exposition. In more general environments, the benchmark should be understood as the robust decision under full Bayesian or prior-by-prior updating (Pires, 2002). Namely, the DM applies Bayes' rule to update every DGP in the initial set conditioning on the sample data to obtain a set of posteriors over future outcomes. In this case, all results either directly apply or generalize under standard conditions, see Appendix B.

Decision Problem I. Suppose Netflix can recommend the movie either aggressively (a) or mildly (m).³ The aggressive recommendation yields a higher payoff than a mild one if users like the movie, but is worse otherwise. Specifically, let $a(U) = 2$, $a(D) = -1$, $m(U) = 1/2$, and $m(D) = 0$ be the state-contingent payoffs. Within the initial set of DGPs, the worst-case DGP for both types of recommendation is $(1/3)^\infty$, since it gives the lowest probability of getting a U . Thus, the benchmark decision is to choose m .

With the sample data, Netflix can make a data-revised decision by revising the initial set of DGPs. A commonly used method is maximum likelihood updating, which revises the initial set to the subset of DGPs that maximize the likelihood of observing the data.⁴ When the true DGP is $(1/3)^\infty$, with N sufficiently large, Netflix will almost surely observe sample data with an empirical frequency of U close to $1/3$. The likelihood of any such sample, however, is never maximized by the true DGP in the initial set. Indeed, given any such sequence of outcomes, one can always find a DGP in the set $\{3/5, 1\}^\infty$ that assigns probability 1 to outcome U whenever U is observed and probability $3/5$ to outcome U whenever D is observed. In this case, notice that for all N ,

$$(1/3)^{N/3} \times (2/3)^{2N/3} < 1^{N/3} \times (2/5)^{2N/3},$$

i.e., the likelihood of observing the given sample under this heterogeneous DGP is always strictly higher than under the true one. Thus, with maximum likelihood updating, Netflix will asymptotically almost surely rule out the true DGP from their data-revised set.⁵ As the worst-case probability of U in the data-revised set increases to $3/5$, the aggressive recommendation becomes Netflix's data-revised decision. However, this decision gives Netflix an expected payoff of 0 under the true DGP, which is strictly lower than the benchmark $(1/6)$. Therefore, with maximum likelihood updating, Netflix will almost surely be worse off than under their benchmark decision when the true DGP is $(1/3)^\infty$.

The preceding example shows that ruling out the true DGP can lead to objectively

³For instance, the movie can be recommended either on top of the front page or further down the list. All payoffs can be interpreted as cardinal measures of how much the recommendation affects users' likelihood of interacting with Netflix.

⁴Proposed and axiomatized in [Gilboa and Schmeidler \(1993\)](#) and [Cheng \(2022\)](#), this updating rule is applied in [Epstein and Schneider \(2007\)](#) for studying problems in a similar setting.

⁵While the present paper, to my knowledge, is the first to make such an observation in the context of maximum likelihood updating, it is intrinsically related to the infamous incidental parameter problem in making estimates using maximum likelihood, discovered by [Neyman and Scott \(1948\)](#) (see [Lancaster \(2000\)](#) for a review), and more broadly, the problem of overfitting. The same observation continues to hold for likelihood-based updating rules, such as those proposed in [Epstein and Schneider \(2007\)](#) and [Cheng \(2022\)](#).

worse decisions. Theorem 1 formalizes this observation by showing that accommodating (i.e., containing, with the necessary technical generalizations) the true DGP when revising the initial set is necessary for guaranteeing objective improvements. This necessity holds even when restricting attention to *basic* decision problems, those consist of two alternatives, one of which yields a constant payoff. Since likelihood-based rules fail this criterion, this paper develops a simple **empirical distribution method** that accommodates the truth, as illustrated below. In the introductory example, this also proves sufficient for objective improvement.

Example (Introductory Example Continued). *Under the empirical distribution method, Netflix revises the initial set to include DGPs whose average mixture of sample marginals, i.e., $1/N \sum_{i=1}^N P_i$ in this example, is close enough to the empirical distribution of observed outcomes. Kolmogorov's strong law of large numbers ensures that the true DGP will be retained in such data-revised sets asymptotically almost surely. More importantly, whenever the true DGP is retained, the data-revised decision is m if the true DGP is $(1/3)^\infty$, and a if the true DGP belongs to the set $\{3/5, 1\}^\infty$. In all cases, it is either the same or strictly better than the benchmark decision, thereby achieving objective improvement.*

Notice that in Decision Problem I, both alternatives rank the two states U and D the same in terms of payoffs. Hence, a *higher* probability of U leads to a *higher* expected payoff under each alternative. I define a decision problem (i.e., a set of alternatives) that possesses this property as a *monotone decision problem*.⁶ More specifically, all alternatives in a monotone decision problem are positive affine transformations of a common payoff function. In this class of problems, accommodating the truth is both necessary and sufficient for achieving objective improvement (Theorem 2). Monotone decision problems arise naturally in economic contexts; Section 3 further illustrates how this result may be applied to canonical principal-agent models.

However, this sufficiency does not extend beyond monotone decision problems (Theorem 3). In other words, such monotonicity is *the* property that characterizes when accommodating the truth suffices for objective improvement.

In searching for objective improvements in all decision problems, Theorem 4 and Corollary 3 together provide an impossibility result: It is achieved if and only if the true DGP is uniquely identified from the data. Such identification, however, is generally in-

⁶This terminology stems from the definition of a “monotone decision problem” in Athey and Levin (2018).

feasible in non-identical environments.⁷ Nevertheless, accommodating the truth still provides a weaker yet meaningful guarantee of objective improvement in all decision problems (Proposition 1). In particular, it ensures that the decision maker would never prefer to forgo the opportunity to revise decisions using data in exchange for the certainty-equivalent payoff of the benchmark decision.

To summarize, all these results jointly provide three decision-theoretic justifications for accommodating the truth in data revision: (1) it is necessary for objective improvements, (2) it is sufficient in monotone decision problems; and (3) it provides a weaker but still useful guarantee in all decision problems. The paper then proceeds to develop practical methods that accommodate the truth in non-identical environments.

First, I show that the simple empirical distribution method accommodates the truth asymptotically almost surely (Theorem 5).

Additional concerns arise when considering data samples of bounded size. The standard finite-sample approach is to develop methods that accommodate the truth with a pre-specified probability. This entails making statistical inferences from independent but non-identical samples, which can be computationally challenging. To address this difficulty, I develop an easy-to-implement method, the *augmented i.i.d. test*, which operates as a simple augmentation to the standard practice of constructing confidence regions from i.i.d. data. Theorem 6 establishes that this method accommodates the truth with at least the required pre-specified probability. For a concrete illustration, Section 5.1 presents a Bernoulli model with ambiguous nuisance parameters and shows that the data-revised sets given by the augmented i.i.d. test are simple extensions of the *Wilson Confidence Intervals* (Wilson, 1927).

Similar non-identical decision environments have been adopted in the literature to study, for instance, social learning with heterogeneous individuals (Reshidi et al., 2025; Chen, 2026), dynamic portfolio choice under ambiguous idiosyncratic shocks (Epstein and Schneider, 2007), and estimating parameters of incomplete models (Epstein et al., 2016). In all these models, the revision rules developed in this paper can provide new perspectives and often distinct predictions. As an illustration, in Section 5.2, I consider the model in Reshidi et al. (2025). They assume the DM applies prior-by-prior updating and show there can be non-vanishing ambiguity. Proposition 2 here shows all that ambiguity vanishes asymptotically if the DM instead applies the revision rules proposed here. In fact, learning under

⁷As in the example, for any sample size N , there always exist multiple DGPs in the set $\{3/5, 1\}^\infty$ such that they are the same up to the N -th experiment, but have different marginals over future outcomes.

the new revision rules is more effective compared to commonly used updating rules and is guaranteed to be correct.

Finally, I briefly summarize the main contributions of this paper to the literature on robust statistical decisions (Wald, 1950; Watson and Holmes, 2016; Hansen and Marinacci, 2016; Manski, 2021). A more detailed discussion is provided in Section 6. Most papers in this literature focus on only the “data-revised decision” in their decision context. The main innovation here is to explicitly compare the data-revised decision with the benchmark decision that could have been made without using data. This comparison turns out to be not always obvious. Indeed, I show that it is impossible to guarantee the data-revised decision to be always better whenever the true DGP is not uniquely identified. In light of this impossibility, the present paper contributes to the literature by (1) characterizing when the simple criterion of accommodating the truth suffices for objective improvement and (2) providing practical inference methods that achieve it in both asymptotic and finite-sample settings.

Outline. Section 2 introduces the decision environment. Section 3 characterizes when and how objective improvements can be achieved. Section 4 defines and studies the proposed revision rules. Section 5 presents two applications with parametric models. Section 6 provides a further discussion on related literature. All proofs are collected in the Appendix.

2 The Decision Environment

There is a countably infinite sequence of random experiments that all have the same finite set of outcomes S with generic element s .⁸ The experiments are ordered and indexed by the set $\mathbb{N} = \{1, 2, \dots\}$. The *full state space* is denoted by $\Omega = S^\infty \equiv \prod_i S_i$ with generic element ω . Let Σ denote the discrete σ -algebra on S and Σ^∞ the product σ -algebra on Ω .

For any given sample size $N \in \mathbb{N}$, let $S^N \equiv \prod_{i=1}^N S_i$ denote the *sample states*. A decision-maker (DM) observes *sample data* $\omega^N \in S^N$ and makes a decision whose payoff depends only on future unrealized experiments. Let $\Omega_N \equiv \prod_{i=N+1}^\infty S_i \equiv S_N^\infty$ denote the *future states* with generic element $\tilde{\omega}_N$. Define Σ_N and Σ_N^∞ similarly.

The DM’s decision is the choice from a set of alternatives, named *acts*. An act is defined as a bounded Σ_N^∞ -measurable function, $f : \Omega_N \rightarrow \mathbb{R}$, that maps future states to payoffs (measured in utilities). Let \mathcal{F} be the space of all acts endowed with the product

⁸The finiteness assumption is made for simplicity. Section 3.2 and Appendix B provide discussions on how to extend the results in this paper to infinite state spaces.

topology.⁹ An act is *finitely-based* if it depends only on finitely many experiments (i.e., not on tail events). Let $x \in \mathbb{R}$ denote a constant act that gives the same payoff x in all future states. A *decision problem* D is a nonempty and compact subset of \mathcal{F} . Let \mathcal{D} denote the collection of all decision problems. Say that a decision problem D is *binary* if $|D| = 2$.

I call a sequence of independent but possibly non-identical distributions over outcomes a *data-generating process (DGP)*. Formally, let $\Delta(\Omega)$ denote the set of all countably additive probability measures on (Ω, Σ^∞) , endowed with the topology of setwise convergence. The collection of all DGPs is the subset $\Delta_{\text{indep}}(\Omega) = \prod_{i=1}^\infty \Delta(S_i) \subseteq \Delta(\Omega)$. On $\Delta_{\text{indep}}(\Omega)$, the subspace topology inherited from $\Delta(\Omega)$ under setwise convergence coincides with the product topology on $\prod_{i=1}^\infty \Delta(S_i)$. For each $P \in \Delta_{\text{indep}}(\Omega)$, write P_i for its marginal distribution on S_i , P^N for its joint marginal on S^N , and P_N^∞ for its marginal on $\Omega_N = S_N^\infty$.

Suppose the DM knows there is an *initial set* \mathcal{P} of possible DGPs. When fixing some sample data ω^N , let P^* denote the *true DGP* that generates the data and governs the future states. The DM's initial knowledge is “correctly specified”, meaning that it always contains the true DGP. Using the terminology from Bayesian learning literature (Kalai and Lehrer, 1993), the DM's initial knowledge contains a “grain of truth”:

Assumption 1. $P^* \in \mathcal{P}$.

What this assumption also means is that every DGP in the initial set could be the true one governing the experiments. In addition, assume the set \mathcal{P} is compact and all $P \in \mathcal{P}$ have full support.¹⁰ Let \mathcal{P}_i , \mathcal{P}^N and \mathcal{P}_N^∞ denote the sets of their corresponding marginals.

Given a decision problem D , the DM can make a robust decision by using the initial set of DGPs. Formally, let the *benchmark decision*, denoted by $c(D) \in D$, be given by

$$c(D) \equiv \arg \max_{f \in D} \min_{P \in \mathcal{P}} \int_{\Omega_N} f(\tilde{\omega}_N) dP_N^\infty(\tilde{\omega}_N) = \arg \max_{f \in D} \min_{P_N^\infty \in \overline{\text{co}}(\mathcal{P}_N^\infty)} \int_{\Omega_N} f(\tilde{\omega}_N) dP_N^\infty(\tilde{\omega}_N),$$

where $\overline{\text{co}}(\mathcal{P}_N^\infty)$ denotes the closed and convex hull of \mathcal{P}_N^∞ . The “min” is well-defined as \mathcal{P} is compact. The equality follows since the minimum can be achieved at an extreme point and thus belongs to \mathcal{P} . Moreover, $c(D)$ is defined as a singleton by imposing an arbitrary (but fixed) tie-breaking rule.

⁹For simplicity, I suppress the dependence of \mathcal{F} on N . This is without loss of generality since any act on Ω_N can be isomorphically defined on Ω_M for all $M \in \mathbb{N}$.

¹⁰This is a standard assumption for studying the update of a set of probability measures. It is also used when establishing a Central Limit Theorem in this environment. See the proof of Lemma A.1 for details.

With sample data ω^N , the DM can also choose to revise their initial set to a *data-revised set* of DGPs. Let $\mathcal{P}(\omega^N) \subseteq \Delta_{indep}(\Omega)$ denote the revised set. Given Assumption 1, further let $\mathcal{P}(\omega^N) \subseteq \mathcal{P}$ as there is no need to consider DGPs outside \mathcal{P} . Importantly, the data-revised set is assumed to depend only on the initial set \mathcal{P} and sample data ω^N but not on D , i.e., the specific decision problem considered. In other words, the DM's revision rule is purely "data-based". In the literature on belief updating, this is known as consequentialism, see [Hanany and Klibanoff \(2007, 2009\)](#) and [Siniscalchi \(2009\)](#) for discussions.

Let $c(D, \omega^N) \in D$ denote the DM's *data-revised decision* given (the closed and convex hull of) their data-revised set $\mathcal{P}(\omega^N)$, i.e.,

$$c(D, \omega^N) \equiv \arg \max_{f \in D} \min_{P_N^\infty \in \overline{\text{co}}(\mathcal{P}(\omega^N)_N^\infty)} \int_{\Omega_N} f(\tilde{\omega}_N) dP_N^\infty(\tilde{\omega}_N),$$

where the "min" is well defined as $\overline{\text{co}}(\mathcal{P}(\omega^N)_N^\infty)$ is compact and $c(D, \omega^N)$ is also defined to be a singleton by imposing a tie-breaking rule, subject to the following consistency requirement:

Assumption 2. *If $c(D)$ is among the maximizers in the data-revised program, then $c(D, \omega^N) = c(D)$. Otherwise, the tie-breaking can be arbitrary.*

Assumption 2 makes sure that if the DM's data-revised set coincides with their initial set, then their data-revised decision is the same as the benchmark. This rules out uninteresting complications when the DM's decisions are different only because of different tie-breakings.¹¹

Importantly, the DM's *objective payoff* from an act is determined by the true DGP P^* that actually governs the future. To make this explicit, let

$$W(f, P^*) = \int_{\Omega_N} f(\tilde{\omega}_N) dP_N^{*\infty}(\tilde{\omega}_N)$$

denote the expected payoff of act f evaluated under the true DGP P^* . For any decision problem D , the DM's objective payoffs from their benchmark and data-revised decisions are then $W(c(D), P^*)$ and $W(c(D, \omega^N), P^*)$, respectively.

¹¹ Also notice both decisions are defined as deterministic choices from D , which seems to rule out randomizations from the DM's decisions. This is, in fact, more general as one can explicitly add those randomizations as acts and it will just be another well-defined decision problem. The current definition allows for richer decision patterns: By imposing different tie-breaking rules, it allows the DM to have different preferences in terms of whether randomization hedges against ambiguity. See [Saito \(2015\)](#) and [Ke and Zhang \(2020\)](#) for discussions and characterizations of such preferences.

Because the DM's decisions depend only on future states and are made using the robust criterion, it is useful to define the following notions that suitably generalize set inclusions:

Definition 1. *The data-revised set $\mathcal{P}(\omega^N)$ **accommodates** a DGP P if*

$$P_N^\infty \in \overline{\text{co}}(\mathcal{P}(\omega^N)_N^\infty).$$

*The data-revised set $\mathcal{P}(\omega^N)$ **refines** the initial set \mathcal{P} if*

$$\overline{\text{co}}(\mathcal{P}(\omega^N)_N^\infty) \subsetneq \overline{\text{co}}(P_N^\infty).$$

*Say that $\mathcal{P}(\omega^N)$ is a **truth-accommodating refinement** if it accommodates the true DGP and refines the initial set.*

Notice that if $P^* \in \mathcal{P}(\omega^N)$, then $\mathcal{P}(\omega^N)$ accommodates P^* . Moreover, $\mathcal{P}(\omega^N)$ refines \mathcal{P} only if $\mathcal{P}(\omega^N)$ is a proper subset of \mathcal{P} . In practice, since a convex combination of DGPs in $\Delta_{\text{indep}}(\Omega)$ remains in $\Delta_{\text{indep}}(\Omega)$ only when the DGPs share identical marginals in all but one experiment, a truth-accommodating refinement is thus essentially $(P^*)_N^\infty \in \mathcal{P}(\omega^N)_N^\infty \subsetneq \mathcal{P}_N^\infty$.

3 Objective Improvements

Fix an initial set \mathcal{P} , some sample data ω^N generated by the true DGP $P^* \in \mathcal{P}$. This section investigates what guarantees objective improvements across different decision problems. To this end, consider the following definition.

Definition 2. *Let $\mathcal{C} \subseteq \mathcal{D}$ denote a **class** of decision problems. A data revision is said to provide **objective improvement** in \mathcal{C} if, for all $D \in \mathcal{C}$,*

$$W(c(D, \omega^N), P^*) \geq W(c(D), P^*),$$

and the inequality is strict for some $D \in \mathcal{C}$.

Among all decision problems, the simplest possible form is the choice between an uncertain and a constant act. Such a canonical form is often used to model, for example, the decision of whether or not to approve a new drug, implement a new policy, convict a defendant, and invest in an asset. I call such decision problems, i.e., binary with a constant

act, *basic decision problems*. Arguably, basic decision problems are the building blocks of more complex decision problems, thus any general enough class of decision problems should include them as special cases. Guaranteeing objective improvements at least in all basic decision problems is a reasonable minimal requirement for any data revision. The following result identifies a necessary condition for this guarantee: The data-revised set must accommodate the true DGP.

Theorem 1. *If the data-revised set $\mathcal{P}(\omega^N)$ does not accommodate the true DGP P^* , then there exists a basic decision problem for which the data-revised decision is objectively worse than the benchmark decision, i.e., $W(c(D, \omega^N), P^*) < W(c(D), P^*)$.*

The proof of Theorem 1 relies on a standard separating hyperplane argument by noticing that $P^* \in \mathcal{P}$ (Assumption 1), but $(P^*)_N^\infty \notin \overline{co}(\mathcal{P}(\omega^N)_N^\infty)$ (Definition 1). The key takeaway is that if the data-revised set fails to accommodate the true DGP, then one can easily construct a basic decision problem for which the data-revised decision is objectively worse. In this sense, Theorem 1 highlights that accommodating the true DGP is a necessary condition for achieving objective improvement in any class of decision problems including basic decision problems as special cases.

A truth-accommodating refinement, a slight strengthening of this necessary condition, can be shown to be also sufficient for objective improvement in basic decision problems. The next section establishes this result by characterizing the essentially largest class of decision problems for which this sufficiency holds, a class that includes all basic decision problems.

3.1 Monotone Decision Problems

What structural feature of a decision problem makes a truth-accommodating refinement sufficient for objective improvement? To build intuition, consider a *betting decision problem*, where every available act is a bet on the *same* event. Formally, in a betting decision problem, there exists $A \subseteq \Omega_N$ such that for each f , for all $\tilde{\omega} \in A$ and $\tilde{\omega}' \in A^c$,

$$f(\tilde{\omega}) = f(A) \geq f(A^c) = f(\tilde{\omega}').$$

In words, all acts rank outcomes in the same way across the two events, A and A^c , but differ in their payoff tradeoffs. As a result, a higher belief in A (i.e., a greater probability assigned to A) leads to a higher expected payoff of every act. Moreover, all acts can

be ordered so that higher acts are optimal under higher beliefs. This is in line with the definition of monotone decision problems in Athey and Levin (2018),¹² which highlights that both payoffs and choices move in the same direction as beliefs.

This form of monotonicity is precisely what makes a truth-accommodating refinement sufficient for objective improvement. When choosing how much to bet on the event A , the worst-case DGP is always the one assigning the lowest probability to A . A truth-accommodating refinement raises this worst-case probability while keeping it below the true probability. Monotonicity then guarantees that the data-revised decision moves closer to the exact optimal act under P^* , thereby delivering a higher objective payoff than the benchmark.

Motivated by this intuition, I now generalize the betting problem to higher dimensions by defining a class of *monotone decision problems*.

Definition 3. A decision problem D is called a **monotone decision problem** if there exists an act $g \in \mathcal{F}$ (not necessarily contained in D) such that, for every $f \in D$,

$$f = \lambda_f g + c_f,$$

for some $\lambda_f \geq 0$ and $c_f \in \mathbb{R}$. Let \mathcal{D}_m denote the class of monotone decision problems.

Intuitively, all acts in a monotone decision problem are non-negative affine transformations of a common reference act g . Hence they induce the same ranking over future states and satisfy

$$W(f, P) = \lambda_f W(g, P) + c_f, \quad \forall f \in D.$$

Thus the DM’s evaluation reduces to the single statistic $W(g, P)$: a higher belief in g (i.e., a greater expected payoff of g) raises the expected payoff of every act. This generalizes the monotonicity observed in betting problems. Importantly, the “monotone” label encodes a *one-dimensional* structure: although the underlying state space may be high-dimensional, all payoff-relevant variation collapses to the single expectation of g . From a geometric perspective, after subtracting the constant, all acts are aligned in the same direction. Thus, choosing among these acts is essentially trading off between their sensitivity to beliefs (λ_f) against their guaranteed payoffs (c_f), just as in betting decision problems one trades off payoffs across two events.

¹²Their definition says that the optimal act is monotone in signal x , which corresponds to a posterior belief over states. The additional property here is that it also leads to a higher expected payoff.

Basic decision problems are trivially monotone, since the only non-constant act can be taken as the reference act. It thus follows from Theorem 1 that accommodating the true DGP is necessary for objective improvement in monotone decision problems. The next result establishes sufficiency.

Theorem 2. *A data revision provides objective improvement in monotone decision problems if and only if the data-revised set is a truth-accommodating refinement.*

Theorem 2 delivers a clear and powerful message: a truth-accommodating refinement is sufficient to guarantee objective improvement in monotone decision problems, regardless of the true DGP. Its proof builds on the intuition developed in betting decision problems, but the extension is substantial: monotone decision problems constitute a broader class that encompasses a wide range of economically relevant environments. To illustrate this concretely, I next show that the problem of choosing among linear contracts in a canonical principal-agent model is a monotone decision problem. Consequently, Theorem 2 applies directly and yields new insights.

Example. *[Improving Linear Contracts with Data]* A principal hires an agent to work on a project whose output is given by $y = \theta e + \epsilon$, where $\theta \in \mathbb{R}_+$ denotes the agent's productivity, $e \in \mathbb{R}_+$ the agent's effort, and ϵ a random noise with $\mathbb{E}[\epsilon] = 0$. The agent's effort cost is $c(e) = ke^2/2$ for some $k > 0$.

The principal faces uncertainty about the agent's productivity. Let $\Theta \subseteq \mathbb{R}_+$ denote the set of possible productivity levels and $\mathcal{P} \subseteq \Delta(\Theta)$ the principal's initial belief set. The principal's objective is to choose a robustly optimal **linear contract** to offer the agent, given \mathcal{P} and a possible revision using data from past projects.

A linear contract specifies a base wage $\alpha \in \mathbb{R}$ and a share $\beta \in [0, 1]$ of output. Given (α, β) , the agent chooses effort e to maximize $\mathbb{E}_\epsilon[\alpha + \beta(\theta e + \epsilon) - ke^2/2]$, yielding the optimal effort $e^*(\theta, \alpha, \beta) = \beta\theta/k \geq 0$. Thus, given (α, β) and θ , the agent's expected payoff is $U_A(\theta; \alpha, \beta) = \alpha + \beta^2\theta^2/2k$ and the principal's expected payoff is

$$\pi_{(\alpha, \beta)}(\theta) = \mathbb{E}_\epsilon[(1 - \beta)(\theta e^*(\theta, \alpha, \beta) + \epsilon) - \alpha] = \frac{\beta(1 - \beta)}{k}\theta^2 - \alpha.$$

Let $u_0 : \Theta \rightarrow \mathbb{R}$ denote the agent's outside-option payoff. Under robustness considerations, the principal restricts attention to contracts that satisfy individual rationality (IR)

uniformly across all types.¹³ This pins down contracts to those satisfying

$$\alpha \geq \sup_{\theta \in \Theta} \left(u_0(\theta) - \frac{\beta^2 \theta^2}{2k} \right) \equiv \alpha_{\min}(\beta).$$

Because the principal's payoff is decreasing in α , it is without loss to focus on $\alpha = \alpha_{\min}(\beta)$. Thus, the principal's expected payoff from a linear contract with share β becomes

$$\pi_\beta(\theta) = \frac{\beta(1 - \beta)}{k} \theta^2 - \alpha_{\min}(\beta) \equiv \lambda_\beta g(\theta) + c_\beta.$$

This implies that choosing among feasible contracts (i.e., among $\beta \in [0, 1]$) is to choose among acts π_β that are non-negative affine transformations of the common act $g(\theta) = \theta^2$, i.e., a monotone decision problem.

By Theorem 2, if the principal revises the initial belief set \mathcal{P} about the agent's productivity to a truth-accommodating refinement using data, the resulting data-revised contract is guaranteed to yield a higher expected payoff than the benchmark contract regardless of the true productivity distribution.

The linear-contract example illustrates how monotone decision problems can arise in economic settings. The key is to identify a decision problem where all available options can be represented as non-negative affine transformations of a common act. In this example, the assumption of linear contracts is essential for representing the principal's decision problem as monotone. This assumption may be partially justified by the prominent role of linear contracts under robustness concerns (Carroll, 2015; Carroll and Meng, 2016). Nevertheless, to extend the observations here to more contracting environments, the generalizations identified in the next section are useful.

3.2 Monotone-Like Decision Problems

Theorem 2 identifies monotonicity as a sufficient condition for a truth-accommodating refinement to guarantee objective improvement when \mathcal{P} is arbitrary and acts are functions from states to payoffs. Its core intuition, in fact, extends more broadly as additional structure is imposed on either the initial set or the acts. Notice the key force behind Theorem 2 is two-fold: (i) all acts share a common worst-case DGP, and (ii) a truth-accommodating

¹³The uniform IR condition provides one natural way to ensure that the linear-contract problem is monotone. Nevertheless, the monotone structure can also be preserved under alternative IR formulations, provided that the set of participating types remains invariant across contracts.

refinement shifts the robust decision towards higher payoffs under every DGP in the revised set, thereby ensuring an objective improvement regardless of which DGP is the truth. In this section, I illustrate that these two forces can be found in several alternative formulations that are relevant in applications.

3.2.1 Monotone Decision Problems Under FOSD

Fix an order on Ω . Say that \mathcal{P} is *FOSD-comparable* if all DGPs in \mathcal{P} are totally ordered by first-order stochastic dominance with respect to this order. A decision problem D is **monotone under FOSD** if, under the same order on Ω , every act in D is increasing and for any $f, g \in D$, the difference $f - g$ is monotone (either increasing or decreasing). This requirement is weaker than Definition 3 as it does not restrict acts to be affine transformation of one another. Nevertheless, it preserves the two forces identified above and thus guarantees objective improvement under truth-accommodating refinements.

Corollary 1. *Fix an order on Ω and suppose \mathcal{P} is FOSD-comparable. If the data-revised set is a truth-accommodating refinement, then the data revision provides objective improvement in all decision problems that are monotone under FOSD.*

To see why, notice when \mathcal{P} is FOSD-comparable, all increasing acts share the same worst-case DGP. If $\mathcal{P}(\omega^N)$ is a truth-accommodating refinement, then the true DGP P^* dominates the common worst-case DGP P_1 in $\mathcal{P}(\omega^N)$, which in turn dominates the common worst-case DGP P_2 in \mathcal{P} . Let $f = c(D)$ and $g = c(D, \omega^N)$. It follows that $W(f, P_2) \leq W(g, P_2)$ and $W(g, P_1) > W(f, P_1)$. Let $h = g - f$ and by monotonicity, h must be increasing. Hence $W(h, P^*) \geq W(h, P_1) > 0$, which implies the desired objective improvement.

Example (Improving (Non-Linear) Contracts with Data). *Continuing the previous example, but additionally suppose the principal's initial knowledge \mathcal{P} is FOSD-comparable with respect to $\Theta \subseteq \mathbb{R}_+$. This additional structure allows the principal to consider more general contracts beyond linear ones while still being able to guarantee objective improvement under truth-accommodating refinements.*

Concretely, let $\Theta = [\underline{\theta}, \bar{\theta}] \subseteq \mathbb{R}_+$ and suppose the principal expands the contract space to include the following quadratic form of contracts:

$$w_{(\beta, \gamma)}(y) = \alpha_{\min}(\beta, \gamma) + \beta y + \frac{\gamma}{2} y^2,$$

with $\beta \in [0, 1]$, $\gamma < (1 - \beta)k/\bar{\theta}^2$, and $\alpha_{\min}(\beta, \gamma)$ the minimal base wage satisfying uniform IR. For simplicity, let $\epsilon \equiv 0$. The upper bound on γ ensures that $e^*(\theta, \beta, \gamma) = \beta\theta/(k - \gamma\theta^2)$ is always well-defined and the principal's payoff,

$$\pi_{(\beta, \gamma)}(\theta) = \frac{(1 - \beta)\beta\theta^2}{k - \gamma\theta^2} - \frac{\gamma\beta^2\theta^4}{2(k - \gamma\theta^2)^2} - \alpha_{\min}(\beta, \gamma),$$

is increasing in θ . However, it is not necessarily true that for any two feasible contracts (β, γ) and (β', γ') , the difference $\pi_{(\beta, \gamma)}(\theta) - \pi_{(\beta', \gamma')}(\theta)$ is monotone in θ . Nevertheless, if the principal's optimal contracts before and after data revision are such that this difference is monotone in θ , then objective improvement can be established by examining only these two contracts: by Corollary 1, objective improvement holds in the decision problem restricted to this pair. Since the principal in fact selects these same two contracts in the full problem, enlarging the contract space to include additional contracts that are not chosen does not affect the objective-improvement conclusion.

3.2.2 Monotone Moment Decision Problems

In fields such as information design (Gentzkow and Kamenica, 2016; Kolotilin et al., 2017; Dworczak and Martini, 2019), among others, decision problems are sometimes modeled as choosing among options whose payoffs depend on the distribution over states only through a scalar *moment* (e.g., the mean). Formally, fix a bounded measurable map $m : \Omega_N \rightarrow \mathbb{R}$ and write $m(P) \equiv \mathbb{E}_P[m(\omega)]$. A *moment act* is a function $\bar{f} : \mathbb{R} \rightarrow \mathbb{R}$ that assigns payoff $\bar{f}(m(P))$ under distribution P . Note that while every (state-contingent) act induces an expectation under each P , not every moment act corresponds to such a state-contingent act, especially when \bar{f} is nonlinear in the moment (e.g., $\bar{f}(x) = x^2$).

A decision problem \bar{D} is a **monotone moment decision problem** if all acts are increasing moment acts and *single crossing* holds: for all $\bar{f}, \bar{g} \in \bar{D}$, if $\bar{f}(x) \geq \bar{g}(x)$ at some x , then $\bar{f}(x') \geq \bar{g}(x')$ for all $x' \geq x$. As before, these conditions ensure the same two forces identified above. The following corollary summarizes the conclusion.

Corollary 2. *If the data-revised set is a truth-accommodating refinement, then the data revision provides objective improvement in all monotone moment decision problems.*

Example (Improving Linear Contracts with Data Under Generalized Preferences). *Continuing the linear-contract example, another key assumption that makes the principal's problem monotone is the principal's payoff form. Allowing the principal's preference to*

depend non-linearly on the expected output but still linear in the expected payment would generally break the monotonicity as in Definition 3. Specifically, let $u : \mathbb{R} \rightarrow \mathbb{R}$ be an increasing function such that the principal's payoff under a linear contract $(\alpha, \beta) \in \mathbb{R} \times [0, 1]$ and a distribution $P \in \Delta(\Theta)$ is

$$\bar{\pi}_{(\alpha, \beta)}(P) = u\left(\frac{\beta}{k}\mathbb{E}_P[\theta^2]\right) - \frac{\beta^2}{k}\mathbb{E}_P[\theta^2] - \alpha.$$

Let $m(P) = \mathbb{E}_P[\theta^2]$, then $\bar{\pi}_{(\alpha, \beta)}$ is in the form of a moment act. When $u(\cdot)$ is twice differentiable, $u'(z) \geq 1$ and $u'(z) + zu''(z) \geq 2$ on the relevant range of moments is sufficient for all $\bar{\pi}_{(\alpha, \beta)}$ to be increasing and pairwise single-crossing in $m(P)$. Then by Corollary 2, again, a truth-accommodating refinement guarantees objective improvement in choosing among linear contracts.

Monotone decision problems under FOSD and monotone moment decision problems illustrate two distinct yet complementary directions for generalizing the notion of monotonicity while preserving the two key forces underlying Theorem 2. As emphasized at the beginning, within any specific decision context, whenever these two forces are present, a truth-accommodating refinement guarantees objective improvement.

3.3 Beyond Monotonicity

Outside the monotone and monotone-like classes identified above, where a common worst case and a directional monotonicity of payoff differences obtain, the guarantee that truth-accommodating refinements yield objective improvement can break down. The next result shows that this failure is robust: it could arise for arbitrary non-singleton refinements, or arbitrary non-monotone binary decision problems.

Theorem 3. *The following statements are true.*

- (i) *For any non-singleton data-revised set $\mathcal{P}(\omega^N)$ that refines an initial set \mathcal{P} , there exists some $P^* \in \mathcal{P}(\omega^N)$ and a non-monotone decision problem D such that*

$$W(c(D, \omega^N), P^*) < W(c(D), P^*).$$

- (ii) *For all non-monotone binary decision problem D with finitely-based acts f_1 and f_2 , if there exists $P \neq P' \in \Delta_{indep}(\Omega)$ such that $W(f_1, P) > W(f_2, P)$, and $W(f_1, P') <$*

$W(f_2, P')$, then there exists an initial set \mathcal{P} , a data-revised set $\mathcal{P}(\omega^N)$, and a true DGP P^* such that $\mathcal{P}(\omega^N)$ is a truth-accommodating refinement, yet

$$W(c(D, \omega^N), P^*) < W(c(D), P^*).$$

Theorem 3 identifies two different impossibility directions for extending sufficiency beyond monotone decision problems. Part (i) rules out any guarantee based solely on the refinement itself: for every non-singleton refinement $\mathcal{P}(\omega^N)$ of \mathcal{P} there is a true DGP and a non-monotone problem for which the data-revised choice underperforms the benchmark. Part (ii) emphasizes that there is no particular way of deviating from monotonicity that maintains sufficiency.¹⁴ In other words, monotonicity is not merely sufficient; it is essentially the boundary for when a truth-accommodating refinement guarantees objective improvement.

To gain some intuition of why without monotonicity, a truth-accommodating refinement could lead to a strictly worse decision, consider the following example.

Example (Introductory Example continued). **Decision Problem II.** Netflix decides whether to include (i) or remove (r) the movie from their recommendations. The key difference from Decision Problem I is that the two actions now rank the states in opposite ways: removing the movie yields a higher payoff when users dislike it, whereas including it yields a higher payoff when users like it. Numerically, let the payoffs be $i(U) = 1$, $i(D) = -1$, $r(U) = 0$, and $r(D) = 1$. Because the two alternatives rank the two states differently, this decision problem is non-monotone.

In this case, Netflix's benchmark decision is to choose r. If Netflix revises the initial belief using the empirical distribution method described in the introduction, then their data-revised decision will be r when the true DGP is $(1/3)^\infty$ and i when the true DGP belongs to the set $\{3/5, 1\}^\infty$. However, if the true DGP is $(3/5)^\infty \in \{3/5, 1\}^\infty$, the expected payoff of i is $1/5$, strictly lower than the expected payoff of r, equal to $2/5$.

In Decision Problem II, the two acts rank the two states differently. While a higher probability of U implies a “higher” optimal act, it does not necessarily lead to a higher expected payoff. This non-monotonicity invalidates the previous intuition. In this case, the

¹⁴The statement restricts to finitely-based acts to avoid a small caveat involving tail events. See Remark A.1 in the proof for details. In addition, similar negative results can be stated for non-binary decision problems as the argument involves only two acts, those that are chosen by the benchmark and data-revised decisions. The presence of other acts only complicates the construction.

worst-case DGPs for the two acts are those assigning the lowest probability to U and D , respectively. While a truth-accommodating refinement still ensures that the lowest probabilities of U and D in the data-revised set are greater than those in the initial set but less than the true probabilities, the different levels of probability increase may lead the data-revised decision to be further away from the exact optimal act.

Decision Problem II illustrates one direction where a decision problem can deviate from monotonicity. For all the other directions, see the example in the proof of Theorem 3 for an illustration.

3.4 Impossibility of Objective Improvement in All Decisions

Are there stronger conditions than truth-accommodating refinements that could guarantee objective improvement beyond monotone decision problems? This section provides a negative answer when considering the set of all decision problems, i.e., when $\mathcal{C} = \mathcal{D}$.

Obviously, a sufficient condition is when P^* is uniquely identified from ω^N and \mathcal{P} is not a singleton. In this case, by letting $\mathcal{P}(\omega^N) = \{P^*\}$, the DM's data-revised decision is exactly optimal against P^* , thus always improves. However, revising a non-singleton initial set to a singleton set containing the true DGP is not always feasible, especially when the possible DGPs can be non-identical. When the data-revised set is not a singleton, the following theorem shows that objective improvement in all decision problems requires the true DGP to be effectively uniquely identified.

Theorem 4. *A data revision provides objective improvement in all decision problems if and only if there exists $\alpha \in (0, 1]$ such that*

$$\overline{co}(\mathcal{P}(\omega^N)_N^\infty) = \alpha P_N^{*\infty} + (1 - \alpha) \overline{co}(\mathcal{P}_N^\infty).$$

The condition $\overline{co}(\mathcal{P}(\omega^N)_N^\infty) = \alpha P_N^{*\infty} + (1 - \alpha) \overline{co}(\mathcal{P}_N^\infty)$ says that, after taking the closed convex hull of future marginals, the data-revised set needs to be a convex combination of the initial set and the true DGP. Observe that the only way to form such a data-revised set requires knowing exactly what P^* is. But if it is the case, the DM should just let $\{P^*\}$ be their data-revised set. On the other hand, if $\overline{co}(\mathcal{P}(\omega^N)_N^\infty) = \alpha P_N^{*\infty} + (1 - \alpha) \overline{co}(\mathcal{P}_N^\infty)$ for some $\alpha \in (0, 1]$, then the same cannot hold for any other P and α whenever $P_N^\infty \neq P_N^{*\infty}$. This can be seen by considering the probability of any arbitrary event. This impossibility thus leads to the following corollary.

Corollary 3. *Any given data-revised set can provide objective improvement in all decision problems under at most one DGP (up to the same marginal over future states).*

Given Corollary 3, if there are multiple DGPs the DM can by no means distinguish using the sample data (like the ones in the introductory example), then no data-revised set would be able to guarantee objective improvement in all decision problems under *all of them*.

Crucially, a truth-accommodating refinement does not have this issue: A data-revised set continues to be a truth-accommodating refinement no matter which DGP accommodated by it turns out to be the truth. Therefore, Theorem 2 indeed further implies that the objective improvement can be guaranteed simultaneously under multiple DGPs, contrasting to the conclusion in Corollary 3.

In addition, a truth-accommodating refinement can also provide a weaker improvement guarantee in all decision problems:

Proposition 1. *If the data-revised set $\mathcal{P}(\omega^N)$ is a truth-accommodating refinement, then for all $D \in \mathcal{D}$,*

$$W(c(D, \omega^N), P^*) \geq \min_{P \in \mathcal{P}} W(c(D), P),$$

and the inequality is strict for some $D \in \mathcal{D}$.

In words, the expected payoff from the data-revised decision is never lower than the guaranteed payoff of the benchmark decision. Hence, whenever the revision rule constitutes a truth-accommodating refinement, the DM would never prefer to forgo the opportunity to revise their decision using data in exchange for the benchmark's certainty-equivalent payoff. This improvement guarantee is relatively weak but not trivial: a misled data-revised decision could be objectively worse than this certainty equivalent. Truth accommodation ensures that it cannot happen. In other words, a truth-accommodating refinement guarantees that learning from data is always valuable relative to receiving the ex-ante certainty equivalent payoff, providing an additional rationale for adopting revision rules that accommodate the truth.

As a final remark, all results in this section also apply when comparing two data-revised sets, say $\mathcal{P}_1(\omega^N)$ and $\mathcal{P}_2(\omega^N)$. When both sets accommodate the truth and $\mathcal{P}_1(\omega^N)$ is a refinement of $\mathcal{P}_2(\omega^N)$, all conclusions about objective improvement holds by viewing $\mathcal{P}_1(\omega^N)$ as a truth-accommodating refinement of $\mathcal{P}_2(\omega^N)$.

4 Revision Rules that Accommodate the Truth

As established in the previous section, for improving robust decisions, it is necessary and sometimes sufficient to accommodate the true DGP in revising the initial set (the refinement part only guarantees it to be sometimes strict). Define a *revision rule* as the mapping from an initial set \mathcal{P} and sample data ω^N to a data-revised set $\mathcal{P}(\omega^N)$. When the sample size is unbounded, the following definition formalizes a notion of accommodating the truth in the asymptotic sense.

Definition 4. *A revision rule accommodates the truth asymptotically almost surely if, for any initial set \mathcal{P} , for every $P^* \in \mathcal{P}$, and for P^* -almost every $\omega \in \Omega$, there exists $\bar{N}(\omega)$ such that, for all $N \geq \bar{N}(\omega)$, the data-revised set $\mathcal{P}(\omega^N)$ accommodates the DGP P^* .*

This definition says that, regardless of which possible DGP governs the uncertainty, the revision rule ensures that, with a sufficient amount of data, the data-revised set accommodates the true DGP almost surely. Therefore, whenever accommodating the truth is sufficient for objective improvement, such improvements are also achieved asymptotically almost surely.

Likelihood-based rules have been shown in the introductory example to violate this property. Here, I propose a revision rule based on empirical distributions: For any sample data ω^N , let $\Phi(\omega^N) \in \Delta(S)$ denote the *empirical distribution*, i.e., for any outcome $s \in S$, $\Phi(\omega^N)(s) \equiv N^{-1} \sum_{i=1}^N I\{\omega_i = s\}$. For any $P \in \Delta_{indep}(\Omega)$ and for any N , the *average of sample marginals*, $\bar{P}^N \in \Delta(S)$, is defined to be the distribution over S given by the average mixture of the marginal distributions over each component of the sample states, i.e., $\bar{P}^N \equiv N^{-1} \sum_{i=1}^N P_i$. For any $p, q \in \Delta(S)$, let $\rho(p, q)$ denote the sup-norm distance.

Definition 5. *The data-revised sets are obtained by the **empirical distribution method** if, for some pre-specified $\epsilon > 0$ and for every ω^N ,*

$$\mathcal{P}(\omega^N) = \{P \in \mathcal{P} : \rho(\bar{P}^N, \Phi(\omega^N)) < \epsilon\}.$$

The empirical distribution method is a formalization of the simple heuristic of retaining a DGP if its average of sample marginals is close enough to the empirical distribution. Such a heuristic can be used when DGPs are i.i.d., for it happens to coincide with retaining DGPs that maximize the likelihood in this special case. With possible non-identical DGPs, while maximum likelihood is no longer useful, this heuristic remains valid.

Theorem 5. *For all $\epsilon > 0$, the empirical distribution method accommodates the truth asymptotically almost surely.*

The proof of Theorem 5 is to verify that Kolmogorov's strong law of large numbers holds. This also suggests that the independence assumption is not crucial for this result. As long as there is a version of the strong law of large numbers, one can obtain the same conclusion. Notable cases include when the DGPs satisfy Markov property or are weakly dependent (De Jong, 1995).

In addition, the empirical distribution method is not the only revision rule that accommodates the truth asymptotically almost surely. For instance, when the odd and even experiments are known to have different characteristics, one may apply the empirical distribution method separately to the odd and even experiments to obtain a potentially more refined data-revised set. Such revision rules, however, typically need to be tailored to the specific structure in a case-by-case manner. In contrast, the empirical distribution method stands out for its simplicity and general applicability.

4.1 Accommodating the Truth with Finite Sample

As Theorem 5 holds for all $\epsilon > 0$, letting $\epsilon \rightarrow 0$ obtains the theoretic limit of the data-revised sets under the empirical distribution method. For applications with finite samples, the standard approach is to derive the ϵ -bound as a function of the sample size that ensures a pre-specified asymptotic probability of accommodating the truth.¹⁵ This section develops a novel and simple method to achieve this. First, define the following finite-sample notion of accommodating the truth.

Definition 6. *A revision rule **accommodates the truth with an asymptotic level** $1 - \alpha$ if, for any initial set \mathcal{P} and for every $P^* \in \mathcal{P}$,*

$$\liminf_{N \rightarrow \infty} P^*(\{\omega^N : \mathcal{P}(\omega^N) \text{ accommodates } P^*\}) \geq 1 - \alpha.$$

A revision rule satisfying Definition 6 ensures that, regardless of which DGP governs the data, the data-revised set accommodates the true DGP with asymptotic probability at least $1 - \alpha$. In particular, the level $1 - \alpha$ is understood uniformly — the probability bound

¹⁵With a finite sample, the only revision rule that accommodates the truth almost surely is to keep using the initial set, because of the full-support assumption. Thus, the only meaningful notion of accommodating the truth with a finite sample is the probabilistic approach, which is also a standard practice in statistical inferences.

holds uniformly over all possible DGPs in \mathcal{P} . Consequently, objective improvement is guaranteed with at least the same asymptotic probability whenever accommodating the truth is sufficient.

Definition 6 effectively requires the data-revised sets to serve as consistent *confidence regions* for the true DGP.¹⁶ One way to ensure this property is to construct the confidence regions directly and use them as the data-revised sets.

Constructing confidence regions is theoretically straightforward exploiting the well-known duality with hypothesis tests. For any $P \in \mathcal{P}$, consider testing the null hypothesis $P^* = P$, against the unrestricted alternative $P^* \neq P$. Let $A_{N,\alpha}(P) \subseteq S^N$ denote the *region of acceptance*: the set of sample data under which the null cannot be rejected. Then find regions that satisfy

$$\liminf_{N \rightarrow \infty} P(A_{\alpha,N}(P)) \geq 1 - \alpha,$$

that is, the probability of accepting the null when it is true is at least $1 - \alpha$ asymptotically. Given any data ω^N , construct the data-revised set as

$$\mathcal{P}(\omega^N) = \{P \in \mathcal{P} : \omega^N \in A_{\alpha,N}(P)\}.$$

Such data-revised sets contain the true DGP with asymptotic probability $1 - \alpha$, uniformly across all P in \mathcal{P} .

When all possible DGPs are i.i.d., this construction is standard and tractable for two convenient features. First, by the central limit theorem, the regions of acceptance can be obtained from probability contours of the corresponding multivariate Gaussian approximations. The mean vectors and covariance matrices depend only on the unique marginal distribution and thus remain fixed as the sample size grows. Second, because each i.i.d. distribution is uniquely determined by its marginal, the number of tests remains fixed regardless of the sample size.

Both features do not carry over when DGPs may be non-identical. First, in this case, both the mean vectors and covariance matrices depend on all marginals across sample states. As a result, for every sample size, both need to be recomputed even for the same DGP. Second, a non-identical DGP is determined by all its marginals, so each additional observation increases the number of tests.

To overcome these difficulties, I present in the following a novel method for constructing confidence regions for non-identical DGPs. This method retains the simplicity of the

¹⁶The coverage is in the weaker sense of accommodating but the difference is unimportant.

i.i.d. case while guaranteeing that the resulting confidence regions cover the true non-identical DGP with at least the required asymptotic level.

Formally, for any $p \in \Delta(S)$, let p^∞ denote the i.i.d. distribution over Ω with marginal p . Let $A_{N,\alpha}^*(p^\infty)$ denote its region of acceptance with asymptotic level $1 - \alpha$, constructed using the corresponding Gaussian approximation.¹⁷ For any $P \in \Delta_{indep}(\Omega)$, recall $\bar{P}^N \in \Delta(S)$ denotes the average of sample marginals. Let $A_{N,\alpha}^*((\bar{P}^N)^\infty)$ denote the region of acceptance constructed using the Gaussian approximation of the i.i.d. distribution $(\bar{P}^N)^\infty$. Consider the following revision rule:

Definition 7. *The data-revised sets are obtained using the **augmented i.i.d. test** with asymptotic level $1 - \alpha$ if, for every ω^N ,*

$$\mathcal{P}(\omega^N) = \{P \in \mathcal{P} : \omega^N \in A_{N,\alpha}^*((\bar{P}^N)^\infty)\}.$$

Intuitively, the augmented i.i.d. test follows a two-step procedure:

- (i) For any sample data ω^N , construct a confidence region as if the initial set consists of all i.i.d. DGPs.
- (ii) For each DGP in the initial set, retain it in the data-revised set if its average of sample marginals coincides with the marginal of some i.i.d. DGP in the previous confidence region.

Notice the first step is the standard procedure for constructing confidence regions from i.i.d. sample data. The essential departure is the second step, which *augments* the i.i.d. confidence region by also including the possible non-identical DGPs. Importantly, this augmentation is achieved through a straightforward comparison, adding virtually no computational difficulty. Therefore, implementing the augmented i.i.d. test is as tractable as conventional statistical inferences based on i.i.d. samples.

Theorem 6. *The augmented i.i.d. test with asymptotic level $1 - \alpha$ accommodates the truth with the same asymptotic level.*

The proof of this theorem relies on a key observation: For all $P \in \Delta_{indep}(\Omega)$, for all N and α ,

$$A_{N,\alpha}^*(P) \subseteq A_{N,\alpha}^*((\bar{P}^N)^\infty),$$

¹⁷The exact form is standard and is given by equation (A.5) in the appendix, with additional notations.

which further implies that

$$P(A_{N,\alpha}^*(P)) \leq P(A_{N,\alpha}^*((\bar{P}^N)^\infty)).$$

Therefore, when testing the null hypothesis $P^* = P$, using the region of acceptance $A_{N,\alpha}^*((\bar{P}^N)^\infty)$, the probability of accepting P when it is true is at least greater than the probability when using $A_{N,\alpha}^*(P)$. The latter probability is, by construction, asymptotically greater than $1 - \alpha$.

The key relation, $A_{N,\alpha}^*(P) \subseteq A_{N,\alpha}^*((\bar{P}^N)^\infty)$, is shown in Lemma A.2 by deriving a result relating the average covariance matrices of the two distributions, P and $(\bar{P}^N)^\infty$. Specifically, subtracting the average covariance matrix of P_N from that of $((\bar{P}^N)^\infty)_N$ yields a positive semi-definite matrix. This result generalizes a well-known variance inequality for mixtures of binomial distributions (Wang, 1993) to the multinomial case.¹⁸

Remark 1. *One potential concern is that computing the probability contours of multivariate Gaussian distributions may be difficult when $|S|$ is large. A practical alternative is to construct a Bonferroni-type confidence region by forming confidence intervals for the probability of each outcome s , each with confidence level $1 - \alpha/(|S| - 1)$. The intersection of all such confidence intervals then yields a confidence region with overall confidence level $1 - \alpha$. However, this region is generally more conservative than that obtained directly from the multivariate Gaussian distribution. Moreover, using the corresponding result in Wang (1993) for binomial distributions, one can show that the Bonferroni-type confidence region constructed from i.i.d. distributions guarantees at least the same coverage probability for non-identical distributions.*

5 Applications with Parametric Models

This section illustrates how the theoretical results translate into familiar statistical and economic settings. In particular, it studies two applications where the initial sets are given by specific parametric models. The first application showcases a setting where the data-revised sets obtained under the augmented i.i.d. test have a closed-form solution, and are closely related to the standard Wilson confidence interval. The second application highlights new

¹⁸Specifically, the average variance of an i.i.d. binomial distribution is always weakly greater than the average variance of a non-identical binomial distribution whose average mean is the same as the i.i.d. distribution.

findings in a commonly studied model of learning under ambiguity. These applications confirm the practical relevance of the proposed revision rules and the associated theoretical results.

5.1 Bernoulli Model with Ambiguous Nuisance Parameters

This model is a generalization of the one studied in [Walley \(1991\)](#). Suppose the DM faces a sequence of coin flips, with outcome space $\Omega = \{H, T\}^\infty$, representing heads and tails. The probability of getting a head from the i -th coin flip is determined by both a *structural parameter* $\theta \in [0, 1]$ and a *nuisance parameter* $\psi_i \in [0, 1]$:

$$(1 - \delta)\theta + \delta\psi_i \in [0, 1],$$

for some fixed $\delta \in [0, 1]$. The structural parameter is common across all flips and can be interpreted as a systematic characteristic of the coins. But each coin flip is also affected by its idiosyncratic feature captured by ψ_i . Throughout, I use the probability of H to represent a probability distribution over $\{H, T\}$. Because each ψ_i is only known to lie in $[0, 1]$, each structural parameter θ corresponds to a set of possible DGPs:

$$\mathcal{P}_\theta = \{P \in \Delta_{indep}(\Omega) : P_i \in [(1 - \delta)\theta, (1 - \delta)\theta + \delta], \forall i\}.$$

Let $\Theta = [0, 1]$ denote the set of structural parameters. The initial set of DGPs is therefore $\mathcal{P} = \cup_{\theta \in \Theta} \mathcal{P}_\theta$. The DM observes outcomes of N coin flips. Their goal is to estimate the true structural parameter and make a set-valued prediction for the probability of getting a head in the next flip. The benchmark estimate is simply the initial set of parameters $\Theta = [0, 1]$, so the benchmark prediction is $\mathcal{P}_{N+1} = [0, 1]$.

Consider the DM's asymptotic prediction using the empirical distribution method. For simplicity, I ignore the pre-specified ϵ by taking it to be arbitrarily small. Then the data-revised set is given by

$$\mathcal{P}(\omega^N) = \left\{ P \in \mathcal{P} : N^{-1} \sum_{i=1}^N P_i = \Phi(\omega^N) \right\}.$$

Let the DM's data-revised estimate of the structural parameter be given by

$$\Theta(\omega^N) \equiv \{\theta \in \Theta : \mathcal{P}_\theta \cap \mathcal{P}(\omega^N) \neq \emptyset\},$$

i.e., a structural parameter is considered possible whenever there is a corresponding DGP retained in $\mathcal{P}(\omega^N)$. For any $\theta \in \Theta$, there exists $P \in \mathcal{P}_\theta$ satisfying the above condition if and only if $\Phi(\omega^N) \in [(1 - \delta)\theta, (1 - \delta)\theta + \delta]$.

As a result, it follows that

$$\Theta(\omega^N) = \left[\max \left\{ \frac{\Phi(\omega^N) - \delta}{1 - \delta}, 0 \right\}, \min \left\{ \frac{\Phi(\omega^N)}{1 - \delta}, 1 \right\} \right].$$

Notice the data-revised prediction is also completely determined by the data-revised estimate of the structural parameter and is given by

$$\mathcal{P}(\omega^N)_{N+1} = \left[\max \left\{ \Phi(\omega^N) - \delta, 0 \right\}, \min \left\{ \Phi(\omega^N) + \delta, 1 \right\} \right].$$

Intuitively, the asymptotic prediction under the empirical distribution method is a “ δ -fattening” of the observed empirical frequency of heads.

Next, consider the finite-sample estimate and prediction at asymptotic level $1 - \alpha$ using the augmented i.i.d. test. For any sample data ω^N , the first step is to construct the confidence interval as if the underlying DGPs were i.i.d. binomial distributions. Specifically, the corresponding confidence interval is the *Wilson Interval*.¹⁹ Let $z_{\alpha/2}$ denote the upper $100(\alpha/2)\%$ quantile of the standard normal distribution. Let $[\underline{W}(\omega^N), \overline{W}(\omega^N)]$ denote the Wilson Interval which has the following closed-form expressions:

$$\begin{aligned} \overline{W}(\omega^N) &= \frac{N\Phi(\omega^N) + z_{\alpha/2}^2/2}{N + z_{\alpha/2}^2} + \frac{z_{\alpha/2}N^{1/2}}{N + z_{\alpha/2}^2} \left(\Phi(\omega^N)(1 - \Phi(\omega^N)) + z_{\alpha/2}^2/(4N) \right)^{1/2}; \\ \underline{W}(\omega^N) &= \frac{N\Phi(\omega^N) + z_{\alpha/2}^2/2}{N + z_{\alpha/2}^2} - \frac{z_{\alpha/2}N^{1/2}}{N + z_{\alpha/2}^2} \left(\Phi(\omega^N)(1 - \Phi(\omega^N)) + z_{\alpha/2}^2/(4N) \right)^{1/2}. \end{aligned}$$

Given the i.i.d. confidence interval, the second step is to consider non-identical DGPs whose average of sample marginals falls into this confidence interval: A structural parameter θ is retained in the data-revised estimate if and only if $[(1 - \delta)\theta, (1 - \delta)\theta + \delta] \cap [\underline{W}(\omega^N), \overline{W}(\omega^N)] \neq \emptyset$. Hence, the data-revised estimate is

$$\Theta(\omega^N) = \left[\max \left\{ \frac{\underline{W}(\omega^N) - \delta}{1 - \delta}, 0 \right\}, \min \left\{ \frac{\overline{W}(\omega^N)}{1 - \delta}, 1 \right\} \right].$$

¹⁹Different from the probably more famous Wald Interval which uses the sample variance, Wilson Interval is constructed by directly inverting the statistical tests, thus using the null variance. Wilson Interval has considerably better asymptotic performance than the Wald Interval. See [Brown et al. \(2001\)](#) for a discussion.

Similarly, the data-revised prediction in this case is

$$\mathcal{P}(\omega^N)_{N+1} = [\max \{\underline{W}(\omega^N) - \delta, 0\}, \min \{\overline{W}(\omega^N) + \delta, 1\}].$$

Notice the prediction is again a δ -fattening of the Wilson Interval. Therefore, The resulting expressions are not only analytically tractable but also intuitively interpretable.

5.2 Gaussian Signals with Ambiguous Variances

Prior-by-prior or full-Bayesian updating is the most commonly used update rule in models of learning under ambiguity in the literature. However, its asymptotic result is often hard to derive and is known only in specific parametric models. This section revisits one such model from [Reshidi et al. \(2025\)](#) and shows that applying the empirical distribution method yields a simpler analysis and entirely different conclusions.

In this model, the DM aims to learn the state of the world $\theta \in \Theta \equiv \mathbb{R}$ by observing a countably infinite sequence of signals denoted by $\{x_i\}_{i=1}^\infty$. Each x_i is a Gaussian random variable with mean θ and variance σ_i^2 , and let $g_i(\theta, \sigma_i)$ denote its probability density function. The signals are mutually independent, but each σ_i is only known to lie in $[\underline{\sigma}, \bar{\sigma}]$. Thus, the DM observes a sequence of independent but possibly heterogeneous Gaussian random variables.

Formally, for each state θ , let

$$\mathcal{P}_\theta = \left\{ \prod_{i=1}^\infty g_i(\theta, \sigma_i) : \sigma_i \in [\underline{\sigma}, \bar{\sigma}], \forall i \right\}$$

denote the set of possible DGPs over the signal sequence. The initial set is $\mathcal{P} = \cup_{\theta \in \Theta} \mathcal{P}_\theta$. For every $N \in \mathbb{N}$, let $\hat{x}^N \equiv (\hat{x}_1, \hat{x}_2, \dots, \hat{x}_N)$ be a sequence of signal realizations and let $\mathcal{P}(\hat{x}^N)$ denote the data-revised set of DGPs. Because the DM's goal is to learn the true state, define

$$\Theta(\hat{x}^N) \equiv \{\theta \in \Theta : \mathcal{P}_\theta \cap \mathcal{P}(\hat{x}^N) \neq \emptyset\}$$

as the set of states consistent with the data-revised set of DGPs.

Directly applying full Bayesian updating here would amount to retain all possible DGPs, implying $\mathcal{P}(\hat{x}^N) \equiv \mathcal{P}$ and hence $\Theta(\hat{x}^N) \equiv \Theta$ for all \hat{x}^N . In [Reshidi et al. \(2025\)](#), they apply full Bayesian updating differently by assuming a prior μ over Θ . Then applying full Bayesian updating is to apply Bayes' rule to update μ under each possible DGP.

Specifically, for each θ , let $P_\theta \in \mathcal{P}_\theta$ denote a specific DGP, the posterior probability

$$\mu(\theta|\hat{x}^N) = \frac{\mu(\theta)P_\theta(\hat{x}^N)}{\int_\Theta \mu(\theta')P_{\theta'}(\hat{x}^N)d\mu(\theta')}.$$

This yields a set of posterior beliefs over Θ .²⁰ Their main result (Theorem 1) shows that in any state θ and some possible DGP, the set of posteriors converges almost surely to a set of degenerate distributions over a non-singleton set of states. That is, ambiguity does not vanish asymptotically.

Consider revising the initial set using the empirical distribution method. Because only the mean matters, it suffices to consider their sample mean.

Definition 8. *The revision of states is obtained by the **sample mean method** if, for some pre-specified $\epsilon > 0$, and for every \hat{x}^N ,*

$$\Theta(\hat{x}^N) = \left\{ \theta \in \Theta : \left| N^{-1} \sum_{i=1}^N \hat{x}_i - \theta \right| < \epsilon \right\}.$$

As in the empirical distribution method, $\Theta(\hat{x}^N)$ retains a state if it is close enough to the sample mean. Because the mean of each marginal distribution equal to θ , applying Kolmogorov's strong law of large numbers yields the following result.

Proposition 2. *For any $\theta^* \in \Theta$ and any $P^* \in \mathcal{P}_{\theta^*}$, the revision of states obtained by the sample mean method with any $\epsilon > 0$ contains the true state asymptotically almost surely, i.e., for any $\epsilon > 0$,*

$$P^* \left(\hat{x} : \lim_{N \rightarrow \infty} \theta^* \in \Theta(\hat{x}^N) \right) = 1.$$

As the conclusion holds for any $\epsilon > 0$, taking ϵ arbitrarily small makes the revised set of states arbitrarily precise. Thus, even with Gaussian signals that have unknown and possibly heterogeneous variances, the true state can still be identified asymptotically almost surely. This stands in sharp contrast to full Bayesian updating, under which ambiguity persists.

The key factor enabling asymptotic identification here is that all Gaussian signals share the same mean. When the means themselves are also ambiguous, the sample-mean method may still yield asymptotic ambiguity over a non-vanishing set of states. Nonetheless, the

²⁰This way of applying full Bayesian updating can be incorporated into the present framework by letting the state space be $\Theta \times S^\infty$ and allowing dependence in DGPs. Applying (B.1) in Appendix B yields exactly the same posterior distribution over Θ .

key here is that this asymptotic prediction is obtained straightforwardly with the sample-mean method, whereas the corresponding asymptotic result for full-Bayesian updating in this setting remains unknown.

6 Related Literature

The decision environment formulated in this paper is closely related to some in the literature on decisions under ambiguity. In particular, it directly generalizes the setting studied by [Epstein and Schneider \(2007\)](#). They assume the DM applies maximum likelihood updating to revise the initial set. The present paper highlights possible concerns with this approach. [Epstein et al. \(2016\)](#) develop robust confidence regions when the possible data-generating processes are belief functions. Belief functions impose restrictions on the possible marginal distributions. In contrast, the environment studied here allows for arbitrary marginals. But the main conceptual difference from their paper and other papers on asymptotic learning under ambiguity, such as [Marinacci \(2002\)](#) and [Marinacci and Massari \(2019\)](#), is that the present paper emphasizes implications for decision making in addition to asymptotic learning.

This paper also contributes to the literature on dynamic decisions under ambiguity by proposing new rules for revising or updating sets of distributions. See [Gilboa and Marinacci \(2013\)](#) and [Cheng \(2022\)](#) for recent developments. The essential departure of the present paper from this literature is that it evaluates decisions using an objective criterion. The objective criterion leads to a characterization of accommodating-the-truth property, conceptually analogous to statistical consistency. In this way, this paper draws a connection between a classical concept from statistics and the theory of decisions under ambiguity, following the line of research by [Cerreia-Vioglio et al. \(2013\)](#) and [Denti and Pomatto \(2022\)](#). This objective approach also resonates with some recent studies of misspecified learning that evaluate performance according to an objective measure, such as [Frick et al. \(2024\)](#) and [He and Libgober \(2025\)](#).

Finally, this paper develops a useful augmenting technique for making inferences in the presence of independent but non-identical distributions. In essence, this technique can be applied on top of standard statistical procedures developed under the i.i.d. assumption. The data-revised set naturally serves as a set-valued identification object and is therefore related to the partial identification literature (see surveys by [Tamer \(2010\)](#), [Canay and Shaikh \(2017\)](#), and [Molinari \(2020\)](#)). In that literature, the decision environment is typically for-

mulated so that the data-generating distribution is point identified, while the payoff-relevant parameters are only partially identified, see, for example, [Christensen et al. \(2023\)](#). By contrast, the present paper focuses on situations where the distribution itself is partially identified and develops relevant inference methods for such settings.

A Proofs of Results

For ease of exposition, unless otherwise specified, I use the notations \mathcal{P} , $\mathcal{P}(\omega^N)$, and P to represent $\overline{\text{co}}(\mathcal{P})N^\infty$, $\overline{\text{co}}(\mathcal{P}(\omega^N))N^\infty$, and P_N^∞ , respectively. Accordingly, a truth-accommodating refinement is one that satisfies $P^* \in \mathcal{P}(\omega^N) \subsetneq \mathcal{P}$.

A.1 Proof of Theorem 1

Proof. Suppose $\mathcal{P}(\omega^N)$ does not accommodate P^* , i.e., $P^* \notin \mathcal{P}(\omega^N)$. By the strong separating hyperplane theorem (See, for example, Corollary 5.80 in [Aliprantis and Border \(2006\)](#)), there exists a Σ_N^∞ -measurable bounded function $f : \Omega_N \rightarrow \mathbb{R}$ and a real number $x \in \mathbb{R}$ such that

$$\int_{\Omega_N} f(\tilde{\omega}_N) dP^*(\tilde{\omega}_N) < x, \text{ and } \min_{P \in \mathcal{P}(\omega^N)} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N) > x.$$

Consider the basic decision problem $D = \{f, x\}$, the above inequalities imply that $c(D) = x$, $c(D, \omega^N) = f$, but $W(f, P^*) < x = W(x, P^*)$. \square

A.2 Proof of Theorem 2

Proof. The necessity is a direct consequence of Theorem 1 as basic decision problems are monotone. The sufficiency is proved as follows.

Fix any initial set \mathcal{P} . By definition, there exists $g \in \mathcal{F}$ such that for all $f \in D$, $f = \lambda_f g + c_f$ for some $\lambda_f \geq 0$ and $c_f \in \mathbb{R}$. Since for any $f \in D$ and any $P \in \Delta(\Omega)$,

$$W(f, P) = \lambda_f W(g, P) + c_f,$$

it follows that all $f \in D$ and g share the same worst-case DGP that delivers the lowest

$W(g, P)$. Based on this observation, further denote

$$\begin{aligned}[a, b] &\equiv \left[\min_{P \in \mathcal{P}} W(g, P), \max_{P \in \mathcal{P}} W(g, P) \right] \subseteq \mathbb{R}, \\ [a', b'] &\equiv \left[\min_{P \in \mathcal{P}(\omega^N)} W(g, P), \max_{P \in \mathcal{P}(\omega^N)} W(g, P) \right] \subseteq \mathbb{R}, \\ p^* &\equiv W(g, P^*) \in \mathbb{R}.\end{aligned}$$

Then $\mathcal{P}(\omega^N)$ is a truth-accommodating refinement of \mathcal{P} implies that $p^* \in [a', b'] \subseteq [a, b]$.

Denote $c(D) = f_1$ and $c(D, \omega^N) = f_2$. If $f_1 = f_2$, then the conclusion holds trivially. Otherwise, it must hold that

$$\begin{aligned}\lambda_{f_1} \cdot a + c_{f_1} &\geq \lambda_{f_2} \cdot a + c_{f_2}, \\ \lambda_{f_1} \cdot a' + c_{f_1} &< \lambda_{f_2} \cdot a' + c_{f_2}.\end{aligned}$$

The second inequality must be strict, otherwise, since f_1 is a solution under the revised set, the tie-breaking assumption forces $c(D, \omega^N) = f_1$, contradicting $c(D, \omega^N) = f_2 \neq f_1$. Strict inequality further implies that $a' > a$. Then, as λ_{f_1} and λ_{f_2} are both non-negative, combining both inequalities yields $\lambda_{f_2} > \lambda_{f_1}$. Thus, for any $p \geq a'$, it must hold that

$$\lambda_{f_1} \cdot p + c_{f_1} < \lambda_{f_2} \cdot p + c_{f_2}.$$

Therefore, as $p^* \geq a'$, it follows that $W(c(D, \omega^N), P^*) > W(c(D), P^*)$.

To show the existence of a monotone decision problem where strict improvement occurs, one can apply a separating hyperplane argument to construct a basic decision problem with $f_1 \neq f_2$. \square

A.3 Proof of Theorem 3

Proof. (i) This follows from Theorem 4. If $\mathcal{P}(\omega^N)$ refines \mathcal{P} , there always exists $P^* \in \mathcal{P}(\omega^N)$ (for example, those on the boundary) such that the condition in Theorem 4 fails. Then by Theorem 2, the decision problem where objective improvement fails must be a non-monotone decision problem.

(ii) Let $D = \{f_1, f_2\} \in \mathcal{D}$ be a non-monotone decision problem with finitely-based acts. For some finite $M > N$, let S_N^M denote the finite set of experiments where both acts

are measurable and let $\omega_N^M = (\omega_{N+1}, \dots, \omega_M)$. For each $P \in \Delta_{indep}(S_N^M)$, identify it as a vector $(P_{N+1}, \dots, P_M) \in (\Delta(S))^{M-N} \subset \mathbb{R}^{(M-N) \times |S|}$. Notice the set $(\Delta(S))^{M-N}$ is convex and compact.²¹ Next, for $j \in \{1, 2\}$, the function

$$W(f_j, P) = \sum_{\omega_N^M \in S_N^M} \prod_{i=N+1}^M P_i(\omega_i) f_j(\omega_N^M),$$

is continuous over $(\Delta(S))^{M-N}$. Therefore, for each $x \in (\min_{\omega_N^M} f_j(\omega_N^M), \max_{\omega_N^M} f_j(\omega_N^M))$, there exists a full support $P \in \Delta_{indep}(S_N^M)$ such that $W(f_j, P) = x$.

D is non-monotone implies both f_1 and f_2 are non-constant. Pick any $\hat{\omega}_N \in \Omega_N$ and define

$$\hat{f}_j(\tilde{\omega}_N) = f_j(\tilde{\omega}_N) - f_j(\hat{\omega}_N), \text{ for } j \in \{1, 2\}.$$

Observe that D is monotone if and only if \hat{f}_1 and \hat{f}_2 are in the same ray, i.e., there exists $\alpha > 0$ such that $\hat{f}_2 = \alpha \hat{f}_1$. Thus, when D is non-monotone, there are two possible cases:

Case 1: \hat{f}_1 and \hat{f}_2 are in the opposite ray, i.e., there exists $\alpha < 0$ such that $\hat{f}_2 = \alpha \hat{f}_1$.

Since S_N^M is finite, all relevant acts and DGPs can be identified as vectors in a finite-dimensional Euclidean space. Henceforth, I use the standard inner product $\langle P, f \rangle$ to denote the expected payoff $W(f, P)$. First observe that, fix an f , any $P \in \Delta(S_N^M)$ can be uniquely decomposed into \hat{P} and \hat{P}^\perp such that

$$\langle P, \hat{f} \rangle = \langle \hat{P}, \hat{f} \rangle + \langle \hat{P}^\perp, \hat{f} \rangle = \langle \hat{P}, \hat{f} \rangle + 0.$$

That is, \hat{P} is parallel to \hat{f} and \hat{P}^\perp is orthogonal to \hat{f} . When \hat{f}_2 and \hat{f}_1 are in the opposite ray, it follows that the same decomposition works for both \hat{f}_1 and \hat{f}_2 . Thus, for any $P \in \Delta_{indep}(S_N^M)$, let \hat{P} denote its component parallel to both \hat{f}_1 and \hat{f}_2 , then

$$\langle P, \hat{f}_j \rangle = \langle \hat{P}, \hat{f}_j \rangle, \text{ for } j \in \{1, 2\}.$$

And thus,

$$\langle P, f_j \rangle = f_j(\hat{\omega}_N) + \text{sgn}(\langle \hat{P}, \hat{f}_j \rangle) \|\hat{P}\| \cdot \|\hat{f}_j\|, \text{ for } j \in \{1, 2\},$$

where $\text{sgn}(\langle \hat{P}, \hat{f}_j \rangle)$ is always opposite for f_1 and f_2 .

²¹A convex combination in the space $(\Delta(S))^{M-N}$ is different from that in $\Delta_{indep}(S_N^M)$. The former is a convex combination of the marginal distributions, while the latter is a convex combination of the joint distributions. Thus, this convexity does not contradict the fact that $\Delta_{indep}(S_N^M)$ is not convex.

Fix any initial set \mathcal{P} that includes both P and P' . Define the following set using f_1 :

$$\{\text{sgn}(\langle \hat{P}, \hat{f}_1 \rangle) \|\hat{P}\| : P \in \overline{\text{co}}(\mathcal{P})\},$$

which is closed and convex, and thus can be represented by a closed interval $[a, b] \subseteq \mathbb{R}$. Then for any $P \in \mathcal{P}$, let $p \in [a, b]$ denote its corresponding element in $[a, b]$ and notice

$$\begin{aligned}\langle P, f_1 \rangle &= f_1(\hat{\omega}_N) + \|\hat{f}_1\|p, \\ \langle P, f_2 \rangle &= f_2(\hat{\omega}_N) - \|\hat{f}_2\|p.\end{aligned}$$

It further implies that a is the minimizer for f_1 and b is the minimizer for f_2 . Let P_1 and P_2 be the two minimizers in \mathcal{P} corresponding to a and b , respectively. Suppose $c(D) = f_1$, the other case is symmetric.

As $P, P' \in \mathcal{P}$, both acts can be the unique maximizer of $W(f, P)$ among $P \in \mathcal{P}$, thus

$$\begin{aligned}f_1(\hat{\omega}_N) + \|\hat{f}_1\|a &< f_2(\hat{\omega}_N) - \|\hat{f}_2\|a, \\ f_1(\hat{\omega}_N) + \|\hat{f}_1\|b &> f_2(\hat{\omega}_N) - \|\hat{f}_2\|b.\end{aligned}$$

Therefore, there then exists $c \in (a, b)$ such that

$$f_1(\hat{\omega}_N) + \|\hat{f}_1\|c = f_2(\hat{\omega}_N) - \|\hat{f}_2\|c.$$

Since $f_1(\hat{\omega}_N) + \|\hat{f}_1\|c > f_1(\hat{\omega}_N) + \|\hat{f}_1\|a$, there further exists $d \in (c, b)$ such that

$$f_1(\hat{\omega}_N) + \|\hat{f}_1\|a < f_2(\hat{\omega}_N) - \|\hat{f}_2\|d.$$

By continuity of $\langle P, f_2 \rangle$ in P , there exists $\tilde{P} \in \Delta_{\text{indep}}(S_N^M)$ such that $\langle \tilde{P}, f_2 \rangle - f_2(\hat{\omega}_N) = -\|\hat{f}_2\|d$, i.e., the data revised set $\mathcal{P}(\omega^N) = \{P_1, \tilde{P}\}$, after taking the closed and convex hull, would correspond to $[a, d] \subsetneq [a, b]$. The last inequality implies $c(D, \omega^N) = f_2$.

Again by continuity, there exists $P^* \in \Delta_{\text{indep}}(S_N^M)$ that corresponds to some $p^* \in (c, d]$. Further let the data-revised set be $\mathcal{P}(\omega^N) = \{P_1, \tilde{P}, P^*\}$, notice this does not change the data-revised decision and accommodates P^* . However,

$$f_1(\hat{\omega}_N) + \|\hat{f}_1\|p^* > f_2(\hat{\omega}_N) - \|\hat{f}_2\|p^*,$$

i.e., $W(c(D), P^*) > W(c(D, \omega^N), P^*)$.

Case 2: \hat{f}_2 is not in the affine hull of \hat{f}_1 . As $W(f_1, P) > W(f_2, P)$ and $W(f_1, P') < W(f_2, P')$, by continuity, there exists $P_0 \in \Delta_{indep}(S_N^M)$ such that $W(f_1, P_0) = W(f_2, P_0)$.

For any $P \in \Delta_{indep}(S_N^M)$, decompose it into \hat{P} and \hat{P}^\perp with respect to \hat{f}_2 so that

$$\langle P, \hat{f}_2 \rangle = \langle \hat{P}, \hat{f}_2 \rangle.$$

Then for f_1 , one has

$$\langle P, \hat{f}_1 \rangle = \langle \hat{P}, \hat{f}_1 \rangle + \langle \hat{P}^\perp, \hat{f}_1 \rangle,$$

where $\langle \hat{P}^\perp, f_1 \rangle \neq 0$ whenever \hat{P}^\perp is non-zero.

Since P_0 has full support, it lies in the interior of $(\Delta_{indep}(S_N^M))$ (identified as the set $(\Delta(S))^{M-N}$). Thus there always exists $\epsilon > 0$ such that the open ball (in $\mathbb{R}^{(M-N) \times |S|}$) centered at P_0 with radius ϵ is entirely contained in $\Delta_{indep}(S_N^M)$. Let \mathcal{P} be a superset of this open ball and suppose $c(D) = f_1$, the other case is symmetric.

From the open ball, one can also find P_ϵ such that

$$\hat{P}_\epsilon = \hat{P}_0, \text{ and } \langle \hat{P}_\epsilon^\perp, \hat{f}_1 \rangle > \langle \hat{P}_0^\perp, \hat{f}_1 \rangle.$$

This further implies

$$\begin{aligned} \langle P_\epsilon, f_1 \rangle &= f_1(\hat{\omega}_N) + \langle \hat{P}_\epsilon, \hat{f}_1 \rangle + \langle \hat{P}_\epsilon^\perp, \hat{f}_1 \rangle \\ &= f_1(\hat{\omega}_N) + \langle \hat{P}_0, \hat{f}_1 \rangle + \langle \hat{P}_\epsilon^\perp, \hat{f}_1 \rangle \\ &> f_1(\hat{\omega}_N) + \langle \hat{P}_0, \hat{f}_1 \rangle + \langle \hat{P}_0^\perp, \hat{f}_1 \rangle \\ &= \langle P_0, f_1 \rangle = \langle P_0, f_2 \rangle = f_1(\hat{\omega}_N) + \langle \hat{P}_0, \hat{f}_2 \rangle \\ &= f_2(\hat{\omega}_N) + \langle \hat{P}_\epsilon, \hat{f}_2 \rangle = \langle P_\epsilon, f_2 \rangle. \end{aligned}$$

Similarly, one can find $P_{-\epsilon}$ in the open ball with the property that

$$\langle P_{-\epsilon}, f_1 \rangle < \langle P_0, f_1 \rangle = \langle P_0, f_2 \rangle = \langle P_{-\epsilon}, f_2 \rangle.$$

Then let the data-revised set be $\{P_{-\epsilon}, P_\epsilon\}$ and let $P^* = P_\epsilon$. In this case, the data-revised decision is $c(D, \omega^N) = f_2$, but

$$W(f_2, P^*) = \langle P_\epsilon, f_2 \rangle < \langle P_\epsilon, f_1 \rangle = W(f_1, P^*),$$

i.e., the conclusion.

The following example provides a simpler illustration of the main intuition in Case 2:

Example. Suppose $D = \{f_1, f_2\}$ and all acts depend on a single experiment with three possible outcomes, i.e., $S = \{s_1, s_2, s_3\}$. Let $f_1 = (0, 2, 0)$ and $f_2 = (-1, 1, 1)$. \hat{f}_2 is not in the affine hull of \hat{f}_1 . In terms of the marginal distribution over S , let $\mathcal{P} = \Delta(S)$, $\mathcal{P}(\omega^N) = \Delta(\{s_2, s_3\})$, and $P^* = \delta_{s_2}$. Notice $\mathcal{P}(\omega^N)$ is a truth-accommodating refinement. Then $c(D) = f_1$, $c(D, \omega^N) = f_2$, but $W(f_1, P^*) = 2 > 1 = W(f_2, P^*)$.

Remark A.1. The assumption of finitely-based acts is used to embed $\Delta_{indep}(S_N^M)$ into Euclidean space and show that $W(f, P)$ is continuous over $\Delta_{indep}(S_N^M)$. This continuity breaks down when the acts are not finitely based, i.e., depend on tail events. By Kolmogorov's zero-one law, all tail events have a probability of either 0 or 1, and thus $W(f, P)$ is no longer continuous in P .

In particular, consider a non-monotone decision problem where the two acts are bets on and against a tail event A . This corresponds to Case 1 in the proof. For this decision problem, however, all DGPs in \mathcal{P} correspond to the two endpoints a and b and nothing else. Therefore, if the data-revised set is a truth-accommodating refinement, then the data-revised decision is either the same as the benchmark decision or exactly optimal under the true DGP. In other words, a truth-accommodating refinement, “inconveniently”, leads to an objective improvement in such a non-monotone decision problem. This fact, however, does not generalize to all non-monotone problems with tail-event acts. Notice the previous example is still true when replacing s_1, s_2 and s_3 by some tail events. The assumption of finitely-based acts is needed to simplify the statement in light of this subtlety.

□

A.4 Proof of Theorem 4

Proof. IF. The inequality $W(c(D, \omega^N), P^*) \geq W(c(D), P^*)$ is trivially true when the data-revised and benchmark decisions are the same. Consider any $D \in \mathcal{D}$ where these two are different, denote $f = c(D)$ and $f' = c(D, \omega^N)$. The condition in the theorem says that, for some $\alpha \in (0, 1]$,

$$\mathcal{P}(\omega^N) = \alpha P^* + (1 - \alpha) \mathcal{P}. \quad (\text{A.1})$$

By optimality of f and f' under \mathcal{P} and $\mathcal{P}(\omega^N)$ respectively, one has

$$\min_{P \in \mathcal{P}} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N) \geq \min_{P \in \mathcal{P}} \int_{\Omega_N} f'(\tilde{\omega}_N) dP(\tilde{\omega}_N), \quad (\text{A.2})$$

and

$$\min_{P \in \mathcal{P}(\omega^N)} \int_{\Omega_N} f'(\tilde{\omega}_N) dP(\tilde{\omega}_N) > \min_{P \in \mathcal{P}(\omega^N)} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N), \quad (\text{A.3})$$

where the strict inequality is required by the tie-breaking assumption (Assumption 2) and that $f \neq f'$. Substituting (A.1) into (A.3) yields,

$$\begin{aligned} & \alpha \int_{\Omega_N} f'(\tilde{\omega}_N) dP^*(\tilde{\omega}_N) + (1 - \alpha) \min_{P \in \mathcal{P}} \int_{\Omega_N} f'(\tilde{\omega}_N) dP(\tilde{\omega}_N) \\ & > \alpha \int_{\Omega_N} f(\tilde{\omega}_N) dP^*(\tilde{\omega}_N) + (1 - \alpha) \min_{P \in \mathcal{P}_N} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N). \end{aligned}$$

Rearranging terms to get,

$$\begin{aligned} & \alpha \left[\int_{\Omega_N} f'(\tilde{\omega}_N) dP^*(\tilde{\omega}_N) - \int_{\Omega_N} f(\tilde{\omega}_N) dP^*(\tilde{\omega}_N) \right] \\ & > (1 - \alpha) \left[\min_{P \in \mathcal{P}_N^\infty} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N) - \min_{P \in \mathcal{P}} \int_{\Omega_N} f'(\tilde{\omega}_N) dP(\tilde{\omega}_N) \right] \geq 0, \end{aligned}$$

where the last inequality follows from inequality (A.2). As $\alpha > 0$, one has,

$$W(c(D, \omega^N), P^*) - W(c(D), P^*) > 0.$$

ONLY IF. Consider the contra-positive statement: If there does not exist any $\alpha \in (0, 1]$ such that the condition holds, then there must exist $D \in \mathcal{D}$ with $W(c(D, \omega^N), P^*) < W(c(D), P^*)$. If $\mathcal{P}(\omega^N)$ does not accommodate P^* , then the conclusion follows from Theorem 1. Consider the case where $\mathcal{P}(\omega^N)$ accommodates P^* . The presumption implies

$$\mathcal{P}(\omega^N) \neq \mathcal{P}_\alpha \equiv \alpha P^* + (1 - \alpha) \mathcal{P}, \forall \alpha \in (0, 1],$$

It thus implies the existence of $\hat{\alpha} \in [0, 1]$ such that

$$\mathcal{P}(\omega^N) \subsetneq \mathcal{P}_\alpha, \forall \alpha \leq \hat{\alpha},$$

and for all $\alpha > \hat{\alpha}$, there always exists $P' \in \mathcal{P}(\omega^N)$ such that $P' \notin \mathcal{P}_\alpha$. In other words, P' is on the boundary of both $\mathcal{P}(\omega^N)$ and $\mathcal{P}_{\hat{\alpha}}$. Since both sets are closed and convex, they admit the same supporting hyperplane at P' . In other words, there exists an act f , whose minimum expectation among both $\mathcal{P}(\omega^N)$ and $\mathcal{P}_{\hat{\alpha}}$ are achieved at P' . It therefore implies

that,

$$\min_{P \in \mathcal{P}(\omega^N)} W(f, P) = \min_{P \in \hat{\alpha}P^* + (1-\hat{\alpha})\mathcal{P}} W(f, P) = \hat{\alpha}W(f, P^*) + (1-\hat{\alpha}) \min_{P \in \mathcal{P}} W(f, P).$$

Next, because $\mathcal{P}(\omega^N) \subsetneq \mathcal{P}_{\hat{\alpha}}$, by strong separating hyperplane theorem, there also must exist an act g such that for some $\beta > \hat{\alpha}$,

$$\hat{\alpha}W(g, P^*) + (1-\hat{\alpha}) \min_{P \in \mathcal{P}} W(g, P) < \min_{P \in \mathcal{P}(\omega^N)} W(g, P) = \beta W(g, P^*) + (1-\beta) \min_{P \in \mathcal{P}} W(g, P),$$

where $\beta > \hat{\alpha}$ follows from the fact that for such g , $W(g, P^*) > \min_{P \in \mathcal{P}} W(g, P)$. One can normalize g by taking mixtures with constant acts to get

$$W(g, P^*) = 1/2, \text{ and } \min_{P \in \mathcal{P}} W(g, P) = 0.$$

Then normalize f similarly to get for some $\epsilon_1, \epsilon_2 \in (0, (\beta - \hat{\alpha})/2)$ such that

$$W(f, P^*) = 1/2 + \epsilon_1 \text{ and } \min_{P \in \mathcal{P}} W(f, P) = \epsilon_2.$$

Then for the decision problem $D = \{f, g\}$, it is the case that $c(D) = f$ and $c(D, \omega^N) = g$. The second claim follows from

$$\begin{aligned} \min_{P \in \mathcal{P}(\omega^N)} W(f, P) &= \hat{\alpha}W(f, P^*) + (1-\hat{\alpha}) \min_{P \in \mathcal{P}} W(f, P) \\ &= \hat{\alpha}(1/2 + \epsilon_1) + (1-\hat{\alpha})\epsilon_2 < \beta/2 \\ &= \beta W(g, P^*) + (1-\beta) \min_{P \in \mathcal{P}} W(g, P) = \min_{P \in \mathcal{P}(\omega^N)} W(g, P). \end{aligned}$$

But $W(f, P^*) > W(g, P^*)$, i.e. the conclusion. □

A.5 Proof of Proposition 1

Proof. For any $D \in \mathcal{D}$, if the data-revised set accommodates the true DGP P^* and refines the initial set, then

$$\begin{aligned} W(c(D, \omega^N), P^*) &\geq \min_{P \in \mathcal{P}(\omega^N)} W(c(D, \omega^N), P) \\ &\geq \min_{P \in \mathcal{P}(\omega^N)} W(c(D), P) \\ &\geq \min_{P \in \mathcal{P}} W(c(D), P), \end{aligned}$$

where the first inequality follows from $P^* \in \mathcal{P}(\omega^N)$, the second inequality follows from the fact that $c(D, \omega^N)$ is optimal with respect to $\mathcal{P}(\omega^N)$, the third inequality follows from $\mathcal{P}(\omega^N) \subsetneq \mathcal{P}$. Notice the last inequality can be strict for some $D \in \mathcal{D}$ as $\mathcal{P}(\omega^N) \subsetneq \mathcal{P}$. \square

A.6 Proof of Theorem 5

Proof. The only argument is to show all possible DGPs satisfy the condition of Kolmogorov's strong law of large numbers:

Kolmogorov's SLLN. Let $\{X_i\}$ be a sequence of independent random variables. Define $Y_N = N^{-1} \sum_{i=1}^N X_i$ and $\bar{\mu}_N = N^{-1} \sum_{i=1}^N \mu_i$. If $E[X_i] = \mu_i$ and $\text{var}(X_i) = \sigma_i^2$,

$$\lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{\sigma_i^2}{i^2} < \infty,$$

then $Y_N - \bar{\mu}_N \rightarrow 0$ almost surely as $N \rightarrow \infty$.

Fix any possible DGP $P \in \Delta_{\text{indep}}(\Omega)$. For any $s \in S$, let $X_i = I\{\omega_i = s\}$. Then X_i 's are independent random variables with $E[X_i] = P_i(s)$ and $\text{var}(X_i) = P_i(s)(1 - P_i(s))$. To see the condition of SLLN holds,

$$\begin{aligned} \lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{\sigma_i^2}{i^2} &= \lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{P_i(s_j)(1 - P_i(s_j))}{i^2} \\ &\leq \lim_{N \rightarrow \infty} \sum_{i=1}^N \frac{1}{i^2} = \frac{\pi^2}{6} < \infty. \end{aligned}$$

Thus, $\Phi(\omega^N)(s) - \bar{P}^N(s) \rightarrow 0$ almost surely for all $s \in S$. SLLN implies that for P -almost every ω and for any $\epsilon > 0$, there exists $\bar{N}(\omega, \epsilon, s)$ such that for all $N \geq \bar{N}(\omega, \epsilon, s)$,

$|\Phi(\omega^N)(s) - \bar{P}^N(s)| < \epsilon$. By definition, for any $\epsilon > 0$, the data-revised set obtained using the empirical distribution method can be written as

$$\mathcal{P}(\omega^N) = \{P \in \mathcal{P} : \cap_{s \in S} |\bar{P}^N(s) - \Phi(\omega^N)(s)| < \epsilon\}.$$

Let $\bar{N}(\omega, \epsilon) = \max\{\bar{N}(\omega, \epsilon, s)\}$, then SLLN implies that for all $N \geq \bar{N}(\omega, \epsilon)$, $P \in \mathcal{P}(\omega^N)$. Thus, the conclusion holds for all $\epsilon > 0$. \square

A.7 Proof of Theorem 6

Proof of Theorem 6. Some additional notations are required to state this proof. Let $S = \{s_1, \dots, s_d\}$. For $s \in \{s_1, \dots, s_{d-1}\}$, let e_s denote the corresponding standard basis in \mathbb{R}^{d-1} . Let $\mathbf{X}_i : \Omega \rightarrow \mathbb{R}^{d-1}$ denote a $(d-1)$ -dimensional random vector such that $\mathbf{X}_i = e_{\omega_i}$ if $\omega_i \in \{s_1, \dots, s_{d-1}\}$ and $\mathbf{X}_i = 0$ if $\omega_i = s_d$. For each P , let \mathbf{P}_i denote the $(d-1)$ -dimensional vector whose coordinates are $P_i(s)$ for $s \in \{s_1, \dots, s_{d-1}\}$. If P is the underlying probability measure, then the mean vector of \mathbf{X}_i is $\boldsymbol{\mu}_i = \mathbf{P}_i$. Moreover, the covariance matrix $\boldsymbol{\Sigma}_i$ of the random vector \mathbf{X}_i is

$$[\boldsymbol{\Sigma}_i]_{kl} = \begin{cases} P_i(s_k)(1 - P_i(s_k)) & \text{if } k = l; \\ -P_i(s_k)P_i(s_l) & \text{if } k \neq l. \end{cases}$$

Given the full-support assumption, the covariance matrix is always positive definite. Let $\bar{\mathbf{X}}_N \equiv N^{-1} \sum_{i=1}^N \mathbf{X}_i$ be the random vector given by the average. Let $\bar{\boldsymbol{\mu}}_N \equiv N^{-1} \sum_{i=1}^N \boldsymbol{\mu}_i$ and $\bar{\boldsymbol{\Sigma}}_N \equiv N^{-1} \sum_{i=1}^N \boldsymbol{\Sigma}_i$.

Lemma A.1. *Let \mathcal{P} be any compact subset of $\Delta_{indep}(\Omega)$ such that all $P \in \mathcal{P}$ have full support. For any $P \in \mathcal{P}$, the Central Limit Theorem holds, i.e., $\sqrt{N}(\bar{\mathbf{X}}_N - \bar{\boldsymbol{\mu}}_N) \xrightarrow{D} \mathcal{N}(0, \bar{\boldsymbol{\Sigma}}_N)$.*

The proof is standard by combining Liapounov's CLT for Triangular arrays with the Cramer-Wold device. Lemma A.1 implies one can use Gaussian approximations to construct acceptance regions. Specifically, for a $(d-1)$ -dimensional random vector $\mathbf{X} \sim \mathcal{N}_{d-1}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, the ellipsoidal region given by the set of all vectors \mathbf{x} satisfying the following has a probability of $1 - \alpha$:

$$(\mathbf{x} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu}) \leq \chi_{d-1}^2(\alpha),$$

where $\chi_{d-1}^2(\alpha)$ is the upper $100\alpha\%$ quantile for the chi-square distribution with $d - 1$ degrees of freedom. By definition, the set of all such vectors is the probability contour of the multivariate Gaussian distribution $\mathcal{N}_{d-1}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ with probability $1 - \alpha$. For any $P \in \Delta_{indep}(\Omega)$, define

$$A_{N,\alpha}^*(P) = \{\mathbf{x} \in \mathbb{R}^{d-1} : (\mathbf{x} - \bar{\boldsymbol{\mu}}_N)^\top \bar{\boldsymbol{\Sigma}}_N^{-1} (\mathbf{x} - \bar{\boldsymbol{\mu}}_N) \leq \chi_{d-1}^2(\alpha)\}, \quad (\text{A.4})$$

i.e., the probability contour of the multivariate Gaussian distribution $\mathcal{N}_{d-1}(\bar{\boldsymbol{\mu}}_N, \bar{\boldsymbol{\Sigma}}_N)$ with probability $1 - \alpha$. Then Lemma A.1 guarantees that

$$\liminf_{N \rightarrow \infty} P(A_{N,\alpha}^*(P)) \geq 1 - \alpha.$$

Let $\hat{\boldsymbol{\Sigma}}_N$ denote the covariance matrix of the i.i.d. distribution $(\bar{P}^N)^\infty$ and define

$$A_{N,\alpha}^*((\bar{P}^N)^\infty) = \{\mathbf{x} \in \mathbb{R}^{d-1} : (\mathbf{x} - \bar{\boldsymbol{\mu}}_N)^\top \hat{\boldsymbol{\Sigma}}_N^{-1} (\mathbf{x} - \bar{\boldsymbol{\mu}}_N) \leq \chi_{d-1}^2(\alpha)\}, \quad (\text{A.5})$$

i.e., the corresponding probability contour for the multivariate Gaussian distribution $\mathbf{N}_{d-1}(\bar{\boldsymbol{\mu}}_N, \hat{\boldsymbol{\Sigma}}_N)$.

Lemma A.2. *For any $P \in \Delta_{indep}(\Omega)$, N and α , $A_{N,\alpha}^*(P) \subseteq A_{N,\alpha}^*((\bar{P}^N)^\infty)$.*

Lemma A.2 further implies $P(A_{N,\alpha}^*(P)) \leq P(A_{N,\alpha}^*((\bar{P}^N)^\infty))$. Therefore, it follows that

$$\liminf_{N \rightarrow \infty} P(A_{N,\alpha}^*((\bar{P}^N)^\infty)) \geq \liminf_{N \rightarrow \infty} P(A_{N,\alpha}^*(P)) \geq 1 - \alpha.$$

Finally, notice that

$$\begin{aligned} \liminf_{N \rightarrow \infty} P(\{\omega^N : P_N^\infty \in \overline{co}(\mathcal{P}(\omega^N)_N^\infty)\}) &\geq \liminf_{N \rightarrow \infty} P(\{\omega^N : P \in \mathcal{P}(\omega^N)\}) \\ &= \liminf_{N \rightarrow \infty} P(A_{N,\alpha}^*((\bar{P}^N)^\infty)) \geq 1 - \alpha. \end{aligned}$$

Thus, the conclusion. □

Proof of Lemma A.1. The following definitions are taken from White (1984):

Liapounov's CLT for Triangular Arrays. Let $\{Z_{Ni}\}$ be a sequence of independent random scalars with $\mu_{Ni} \equiv E[Z_{Ni}]$, $\sigma_{Ni}^2 = var(Z_{Ni})$, and $E|Z_{Ni} - \mu_{Ni}|^{2+\delta} < \Delta < \infty$ for some $\delta > 0$ and all N and i . Define $\bar{Z}_N \equiv N^{-1} \sum_{i=1}^N Z_{Ni}$, $\bar{\mu}_N \equiv N^{-1} \sum_{i=1}^N \mu_{Ni}$ and

$\bar{\sigma}_N^2 \equiv \text{var}(\sqrt{N}\bar{Z}_N) = N^{-1} \sum_{i=1}^N \sigma_{Ni}^2$. If $\bar{\sigma}_N^2 > \delta' > 0$ for all N sufficiently large, then

$$\sqrt{N}(\bar{Z}_N - \bar{\mu}_N)/\bar{\sigma}_N \xrightarrow{D} \mathcal{N}(0, 1).$$

Cramer-Wold device. Let $\{\mathbf{b}_N\}$ be a sequence of random $k \times 1$ vectors and suppose that for any real $k \times 1$ vector $\boldsymbol{\lambda}$ such that $\boldsymbol{\lambda}^\top \boldsymbol{\lambda} = 1$, $\boldsymbol{\lambda}^\top \mathbf{b}_N \xrightarrow{D} \boldsymbol{\lambda}^\top \mathbf{Z}$ where \mathbf{Z} is a $k \times 1$ vector with joint distribution function F . Then the limiting distribution of \mathbf{b}_N exists and equals F .

For any $\boldsymbol{\lambda}$ with $\boldsymbol{\lambda}^\top \boldsymbol{\lambda} = 1$, consider the sequence of random variables given by $Y_i = \boldsymbol{\lambda}^\top \mathbf{X}_i$. Then one has $E[Y_i] = \boldsymbol{\lambda}^\top \boldsymbol{\mu}_i$ and $\text{var}(Y_i) = \boldsymbol{\lambda}^\top \boldsymbol{\Sigma}_i \boldsymbol{\lambda}$. For $\bar{Y}_N = N^{-1} \sum_{i=1}^N Y_i$, one has $E[\bar{Y}_N] = \boldsymbol{\lambda}^\top \bar{\boldsymbol{\mu}}_N$ and

$$\text{var}(\sqrt{N}\bar{Y}_N) = \boldsymbol{\lambda}^\top \bar{\boldsymbol{\Sigma}}_N \boldsymbol{\lambda} = \boldsymbol{\lambda}^\top (\bar{\boldsymbol{\Sigma}}_N^{1/2})^\top (\bar{\boldsymbol{\Sigma}}_N^{1/2}) \boldsymbol{\lambda}.$$

The existence of $\bar{\boldsymbol{\Sigma}}_N^{1/2}$ is guaranteed by the fact that $\bar{\boldsymbol{\Sigma}}_N$ is positive definite. Moreover, let it be the positive square root of $\bar{\boldsymbol{\Sigma}}_N$ so that it is unique, symmetric, and positive definite (Theorem 7.2.6 (a) in [Horn and Johnson \(2012\)](#)). Let $Z_{Ni} = \boldsymbol{\lambda}^\top \bar{\boldsymbol{\Sigma}}_N^{-1/2} \mathbf{X}_i$. Then one has $E[\bar{Z}_N] = \boldsymbol{\lambda}^\top \bar{\boldsymbol{\Sigma}}_N^{-1/2} \bar{\boldsymbol{\mu}}$ and

$$\begin{aligned} \text{var}(\sqrt{N}\bar{Z}_N) &= \boldsymbol{\lambda}^\top (\bar{\boldsymbol{\Sigma}}_N^{-1/2})^\top \text{var}(\sqrt{N}\bar{\mathbf{X}}_N) \bar{\boldsymbol{\Sigma}}_N^{-1/2} \boldsymbol{\lambda} \\ &= \boldsymbol{\lambda}^\top (\bar{\boldsymbol{\Sigma}}_N^{-1/2})^\top (\bar{\boldsymbol{\Sigma}}_N^{1/2})^\top (\bar{\boldsymbol{\Sigma}}_N^{1/2}) (\bar{\boldsymbol{\Sigma}}_N^{-1/2}) \boldsymbol{\lambda} = 1. \end{aligned}$$

If Liapounov's CLT holds, then $\sqrt{N}(\bar{Z}_N - E[\bar{Z}_N]) \xrightarrow{D} \mathcal{N}(0, 1)$. Applying the Cramer-Wold device will imply the Lemma. Thus, it remains to show that the Liapounov condition holds for any $P \in \mathcal{P}$ and for any $\boldsymbol{\lambda}$. That is, for some $\delta > 0$, there exists $\Delta < \infty$ such that for any N and i , $E|Z_{Ni}|^{2+\delta} < \Delta$. By Minkowski's inequality,

$$E[|Z_{Ni}|^{2+\delta}] \leq \sum_j \lambda_j E|(\bar{\boldsymbol{\Sigma}}_N^{-1/2} \mathbf{X}_i)_j|^{2+\delta}.$$

Notice that \mathbf{X}_i is either 0 or e_k . The RHS is less than the following

$$(\max_{k,l} [\bar{\boldsymbol{\Sigma}}_N^{-1/2}]_{kl})^{2+\delta}.$$

Since $\bar{\boldsymbol{\Sigma}}_N^{-1/2}$ is symmetric and positive semi-definite, the largest entry must be on the diagonal. The trace of the matrix is equal to the sums of the eigenvalues. Both the diagonal

elements and eigenvalues are non-negative due to positive definiteness (Corollary 7.1.5 in [Horn and Johnson \(2012\)](#)). Thus, it suffices to show the eigenvalues of $\bar{\Sigma}_N^{-1/2}$ are bounded from above. Moreover, as each eigenvalue of $\bar{\Sigma}_N^{-1/2}$ is the square root of the eigenvalue of $\bar{\Sigma}_N^{-1}$. Thus, one only needs to show the eigenvalues of $\bar{\Sigma}_N^{-1}$ are bounded from above, which is equivalent to showing the eigenvalues of $\bar{\Sigma}_N$ are bounded away from zero.

For every Σ_i , the smallest eigenvalue is bounded below by $\min_j P_i(s_j)$ ([Watson, 1996](#)). Furthermore, as each Σ_i is symmetric and positive definite, the minimum eigenvalue of $\sum_{i=1}^n \Sigma_i$ (n times the minimum eigenvalue of $\bar{\Sigma}_N$) is greater than the sum of the minimum eigenvalues of every Σ_i (Corollary 4.3.15 in [Horn and Johnson \(2012\)](#)). Therefore, it suffices to show that $P_i(s_j)$ is bounded away from zero for every i and j . By assumption, \mathcal{P} is compact thus $\min_{j,i} P_i(s_j)$ exists. By full-support assumption, the minimum is always positive. \square

Proof of Lemma A.2. It suffices to show that for any $P \in \Delta_{indep}(\Omega)$, N , constant $c \in \mathbb{R}$, and for all \mathbf{x} ,

$$(\mathbf{x} - \bar{\mu}_N)^\top \bar{\Sigma}_N^{-1} (\mathbf{x} - \bar{\mu}_N) \leq c \Rightarrow (\mathbf{x} - \bar{\mu}_N)^\top \hat{\Sigma}_N^{-1} (\mathbf{x} - \bar{\mu}_N) \leq c,$$

which is equivalent to

$$(\mathbf{x} - \bar{\mu}_N)^\top \bar{\Sigma}_N^{-1} (\mathbf{x} - \bar{\mu}_N) - (\mathbf{x} - \bar{\mu}_N)^\top \hat{\Sigma}_N^{-1} (\mathbf{x} - \bar{\mu}_N) \geq 0, \forall \mathbf{x}.$$

In other words, for any \mathbf{x} ,

$$\mathbf{x}^\top (\bar{\Sigma}_N^{-1} - \hat{\Sigma}_N^{-1}) \mathbf{x} \geq 0$$

That is, the lemma is true if $(\bar{\Sigma}_N^{-1} - \hat{\Sigma}_N^{-1})$ is a positive semi-definite matrix. It is further equivalent to, according to Corollary 7.7.4 in [Horn and Johnson \(2012\)](#), $(\hat{\Sigma}_N - \bar{\Sigma}_N)$ being positive semi-definite. The two covariance matrices are given by, respectively,

$$[\hat{\Sigma}_N]_{kl} = \begin{cases} \bar{P}^N(s_k)(1 - \bar{P}^N(s_k)) & \text{if } k = l, \\ -\bar{P}^N(s_k)\bar{P}^N(s_l) & \text{if } k \neq l, \end{cases}$$

and

$$[\bar{\Sigma}_N]_{kl} = \begin{cases} N^{-1} \sum_{i=1}^N P_i(s_k)(1 - P_i(s_k)) & \text{if } k = l, \\ -N^{-1} \sum_{i=1}^N P_i(s_k)P_i(s_l) & \text{if } k \neq l. \end{cases}$$

By algebra, one can show,

$$N[\hat{\Sigma}_N - \bar{\Sigma}_N]_{kl} = \begin{cases} \sum_{i=1}^N (P_i(s_k) - \bar{P}^N(s_k))^2 & \text{if } k = l; \\ \sum_{i=1}^N (P_i(s_k) - \bar{P}^N(s_k))(P_i(s_l) - \bar{P}^N(s_l)) & \text{if } k \neq l. \end{cases}$$

Notice that $N[\hat{\Sigma}_N - \bar{\Sigma}_N]_{kl}$ is a Gram Matrix $G_{kl} = \langle v_k, v_l \rangle$, where the set of vectors are given by:

$$[v_k]_i = P_i(s_k) - \bar{P}^N(s_k).$$

A Gram Matrix is always positive semi-definite (Theorem 7.2.10 in [Horn and Johnson \(2012\)](#)), therefore $\hat{\Sigma}_N - \bar{\Sigma}_N$ is positive semi-definite as desired. \square

B When Experiments can be Dependent

Throughout the main text, the possible data-generating processes are assumed to be independent across experiments. In this section, I show that all results in Section 3 continue to hold after a slight modification of the definitions to take dependence into account. I then discuss how results in Section 4 can also be generalized accordingly.

For this section only, let the initial set \mathcal{P} be a compact subset of $\Delta(\Omega)$. As before, assume that every $P \in \mathcal{P}$ has full support. For any $P \in \mathcal{P}$ and sample data $\omega^N \in S^N$, define the marginal distribution over future states conditional on ω^N by

$$P(\tilde{\omega}_N \mid \omega^N) = \frac{P(\omega^N, \tilde{\omega}_N)}{P(\omega^N)}, \quad \forall \tilde{\omega}_N \in \Omega_N. \quad (\text{B.1})$$

Let $\mathcal{P}(\cdot \mid \omega^N)$ denote the set of such conditional distributions.

The *benchmark decision* continues to use the initial set \mathcal{P} but conditions on the observed sample ω^N through the set of marginals $\mathcal{P}(\cdot \mid \omega^N)$. This corresponds to the decision of a

DM who applies prior-by-prior, or full-Bayesian, updating to each $P \in \mathcal{P}$. Formally,

$$c(D | \omega^N) \equiv \arg \max_{f \in D} \min_{P \in \mathcal{P}(\cdot | \omega^N)} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N) = \arg \max_{f \in D} \min_{P \in \overline{\text{co}}(\mathcal{P}(\cdot | \omega^N))} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N).$$

Let $\mathcal{P}(\omega^N)$ denote the data-revised set, and $\mathcal{P}(\omega^N)(\cdot | \omega^N)$ the corresponding set of conditional distributions. The *data-revised decision* is then

$$c(D, \omega^N | \omega^N) \equiv \arg \max_{f \in D} \min_{P \in \overline{\text{co}}(\mathcal{P}(\omega^N)(\cdot | \omega^N))} \int_{\Omega_N} f(\tilde{\omega}_N) dP(\tilde{\omega}_N).$$

In words, the data-revised decision is obtained through a two-step procedure: first revising the initial set to $\mathcal{P}(\omega^N)$, and then conditioning each DGP in the revised set on the observed data.

The data-revised set *accommodates* a DGP P if

$$P(\cdot | \omega^N) \in \overline{\text{co}}(\mathcal{P}(\omega^N)(\cdot | \omega^N)).$$

This condition holds whenever $P \in \overline{\text{co}}(\mathcal{P}(\omega^N))$, that is, whenever the data-revised set contains the true DGP. The data-revised set *refines* the initial set if

$$\overline{\text{co}}(\mathcal{P}(\omega^N)(\cdot | \omega^N)) \subsetneq \overline{\text{co}}(\mathcal{P}(\cdot | \omega^N)).$$

A *truth-accommodating refinement* is a data-revised set that both accommodates the true DGP P^* and refines the initial set.

Notice that all results in Section 3 rely on the independence assumption only through the definitions of the benchmark decision, the data-revised decision, and the truth-accommodating refinement. Once these definitions are modified to incorporate dependence, all proofs continue to hold without change. More specifically, one can simply let $\mathcal{P}(\cdot | \omega^N)$ be the initial set, $\mathcal{P}(\omega^N)(\cdot | \omega^N)$ the data-revised set, and $P^*(\cdot | \omega^N)$ the true DGP, and directly apply the results.

Theorem B.1. *Suppose $\mathcal{P} \subset \Delta(\Omega)$. If the benchmark decision, data-revised decision, and truth-accommodating refinement are defined as in this section, then all results in Section 3 continue to hold.*

For results in Section 4, as already noted in the main text, Theorem 5 continues to hold as long as a version of the strong law of large numbers applies. Theorem 6 is more delicate,

but by the same token, it is natural to conjecture that it can also be generalized under similar conditions.

Finally, Theorem 5 can be further extended to the case where S is infinite. In that case, the proof would invoke the generalized Glivenko-Cantelli theorem for independent but non-identically distributed random variables, as presented in Wellner (1981).

References

Aliprantis, Charalambos D. and Kim C. Border (2006) *Infinite Dimensional Analysis*, Berlin/Heidelberg: Springer-Verlag, [10.1007/3-540-29587-9](https://doi.org/10.1007/3-540-29587-9).

Athey, Susan and Jonathan Levin (2018) “The value of information in monotone decision problems,” *Research in Economics*, 72 (1), 101–116, <https://doi.org/10.1016/j.rie.2017.01.001>.

Blum, Jeffrey R and Judah Rosenblatt (1967) “On Partial a Priori Information in Statistical Inference,” *The Annals of Mathematical Statistics*, 38 (6), 1671 – 1678, [10.1214/aoms/1177698602](https://doi.org/10.1214/aoms/1177698602).

Brown, Lawrence D., T. Tony Cai, and Anirban DasGupta (2001) “Interval Estimation for a Binomial Proportion,” *Statistical Science*, 16 (2), 101 – 133, [10.1214/ss/1009213286](https://doi.org/10.1214/ss/1009213286).

Canay, Ivan A. and Azeem M. Shaikh (2017) *Practical and Theoretical Advances in Inference for Partially Identified Models*, 2 of Econometric Society Monographs, 271–306: Cambridge University Press, [10.1017/9781108227223.009](https://doi.org/10.1017/9781108227223.009).

Cao, Longbing (2014) “Non-IIDness Learning in Behavioral and Social Data,” *The Computer Journal*, 57 (9), 1358–1370, [10.1093/comjnl/bxt084](https://doi.org/10.1093/comjnl/bxt084).

Carroll, Gabriel (2015) “Robustness and Linear Contracts,” *American Economic Review*, 105 (2), 536–63, [10.1257/aer.20131159](https://doi.org/10.1257/aer.20131159).

——— (2019) “Robustness in Mechanism Design and Contracting,” *Annual Review of Economics*, 11 (1), 139–166, [10.1146/annurev-economics-080218-025616](https://doi.org/10.1146/annurev-economics-080218-025616).

Carroll, Gabriel and Delong Meng (2016) “Robust Contracting with Additive Noise,” *Journal of Economic Theory*, 166, 586–604, [10.1016/j.jet.2016.10.002](https://doi.org/10.1016/j.jet.2016.10.002).

Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio (2013) “Ambiguity and robust statistics,” *Journal of Economic Theory*, 148 (3), 974–1049, [10.1016/j.jet.2012.10.003](https://doi.org/10.1016/j.jet.2012.10.003).

Chen, Jaden Yang (2026) “Sequential Learning under Informational Ambiguity,” *American Economic Review*, 116 (1), 209–45, [10.1257/aer.20231394](https://doi.org/10.1257/aer.20231394).

Cheng, Xiaoyu (2022) “Relative Maximum Likelihood updating of ambiguous beliefs,” *Journal of Mathematical Economics*, 99, 102587, <https://doi.org/10.1016/j.jmateco.2021.102587>.

Christensen, Timothy, Hyungsik Roger Moon, and Frank Schorfheide (2023) “Optimal Decision Rules when Payoffs are Partially Identified,” Unpublished Manuscript, University College London.

De Jong, Robert M (1995) “Laws of Large Numbers for Dependent Heterogeneous Processes,” *Econometric Theory*, 11 (2), 347–358, [10.1017/S0266466600009208](https://doi.org/10.1017/S0266466600009208).

Denti, Tommaso and Luciano Pomatto (2022) “Model and Predictive Uncertainty: A Foundation for Smooth Ambiguity Preferences,” *Econometrica*, 90 (2), 551–584, [10.3982/ECTA18009](https://doi.org/10.3982/ECTA18009).

Dworczak, Piotr and Giorgio Martini (2019) “The Simple Economics of Optimal Persuasion,” *Journal of Political Economy*, 127 (5), 1993–2048, [10.1086/701813](https://doi.org/10.1086/701813).

Epstein, Larry G., Hiroaki Kaido, and Kyoungwon Seo (2016) “Robust Confidence Regions for Incomplete Models,” *Econometrica*, 84 (5), 1799–1838, <https://doi.org/10.3982/ECTA13394>.

Epstein, Larry G. and Martin Schneider (2007) “Learning under Ambiguity,” *The Review of Economic Studies*, 74 (4), 1275–1303, <http://www.jstor.org/stable/4626181>.

Frick, Mira, Ryota Iijima, and Yuhta Ishii (2024) “Welfare Comparisons for Biased Learning,” *American Economic Review*, 114 (6), 1612–49, [10.1257/aer.20210410](https://doi.org/10.1257/aer.20210410).

Gentzkow, Matthew and Emir Kamenica (2016) “A Rothschild-Stiglitz Approach to Bayesian Persuasion,” *American Economic Review*, 106 (5), 597–601, [10.1257/aer.p20161049](https://doi.org/10.1257/aer.p20161049).

Gilboa, Itzhak and Massimo Marinacci (2013) “Ambiguity and the Bayesian Paradigm,” in Acemoglu, Daron, Manuel Arellano, and Eddie Dekel eds. *Advances in Economics and Econometrics: Tenth World Congress*, 1 of Econometric Society Monographs, 179–242: Cambridge University Press, [10.1017/CBO9781139060011.008](https://doi.org/10.1017/CBO9781139060011.008).

Gilboa, Itzhak and David Schmeidler (1989) “Maxmin expected utility with non-unique prior,” *Journal of Mathematical Economics*, 18 (2), 141–153, [https://doi.org/10.1016/0304-4068\(89\)90018-9](https://doi.org/10.1016/0304-4068(89)90018-9).

——— (1993) “Updating Ambiguous Beliefs,” *Journal of Economic Theory*, 59 (1), 33–49, <https://doi.org/10.1006/jeth.1993.1003>.

Hanany, Eran and Peter Klibanoff (2007) “Updating preferences with multiple priors,” *Theoretical Economics*, 2, 261–298.

——— (2009) “Updating Ambiguity Averse Preferences,” *The B.E. Journal of Theoretical Economics*, 9 (1), 0000102202193517041547, [doi:10.2202/1935-1704.1547](https://doi.org/10.2202/1935-1704.1547).

Hansen, Lars Peter and Massimo Marinacci (2016) “Ambiguity Aversion and Model Misspecification: An Economic Perspective,” *Statistical Science*, 31 (4), 511 – 515, [10.1214/16-STS570](https://doi.org/10.1214/16-STS570).

Hansen, Lars Peter and Thomas J. Sargent (2007) *Robustness*, Princeton: Princeton University Press, [doi:10.1515/9781400829385](https://doi.org/10.1515/9781400829385).

He, Kevin and Jonathan Libgober (2025) “Misspecified Learning and Evolutionary Stability,” *Journal of Economic Theory*, 230, 106082, [10.1016/j.jet.2025.106082](https://doi.org/10.1016/j.jet.2025.106082).

Horn, Roger A and Charles R Johnson (2012) *Matrix analysis*: Cambridge university press.

Kalai, Ehud and Ehud Lehrer (1993) “Rational Learning Leads to Nash Equilibrium,” *Econometrica*, 61 (5), 1019–1045, <http://www.jstor.org/stable/2951492>.

Ke, Shaowei and Qi Zhang (2020) “Randomization and Ambiguity Aversion,” *Econometrica*, 88 (3), 1159–1195, <https://doi.org/10.3982/ECTA15182>.

Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li (2017) “Persuasion of a privately informed receiver,” *Econometrica*, 85 (6), 1949–1964, <http://www.jstor.org/stable/44955187>.

Lancaster, Tony (2000) “The incidental parameter problem since 1948,” *Journal of Econometrics*, 95 (2), 391–413, [https://doi.org/10.1016/S0304-4076\(99\)00044-5](https://doi.org/10.1016/S0304-4076(99)00044-5).

Manski, Charles F. (2021) “Econometrics for Decision Making: Building Foundations Sketched by Haavelmo and Wald,” *Econometrica*, 89 (6), 2827–2853, <https://doi.org/10.3982/ECTA17985>.

Marinacci, Massimo (2002) “Learning from ambiguous urns,” *Statistical Papers*, 43 (1), 143–151, [10.1007/s00362-001-0092-5](https://doi.org/10.1007/s00362-001-0092-5).

Marinacci, Massimo and Filippo Massari (2019) “Learning from ambiguous and misspecified models,” *Journal of Mathematical Economics*, 84, 144–149, <https://doi.org/10.1016/j.jmateco.2019.07.012>.

Molinari, Francesca (2020) *Microeconomics with partial identification*, 7 of Handbook of Econometrics, 355–486: Elsevier, <https://doi.org/10.1016/bs.hoe.2020.05.002>.

Neyman, J. and Elizabeth L. Scott (1948) “Consistent Estimates Based on Partially Consistent Observations,” *Econometrica*, 16 (1), 1–32, <http://www.jstor.org/stable/1914288>.

Pires, Cesaltina Pacheco (2002) “A rule for updating ambiguous beliefs,” *Theory and Decision*, 53 (2), 137–152.

Reshidi, Pëllumb, João Thereze, and Mu Zhang (2025) “Asymptotic Learning with Ambiguous Information,” *American Economic Journal: Microeconomics*, 17 (3), 244–88, [10.1257/mic.20230142](https://doi.org/10.1257/mic.20230142).

Saito, Kota (2015) “Preferences for Flexibility and Randomization under Uncertainty,” *American Economic Review*, 105 (3), 1246–71, [10.1257/aer.20131030](https://doi.org/10.1257/aer.20131030).

Siniscalchi, Marciano (2009) “Two Out of Three Ain’t Bad: A Comment on “the Ambiguity Aversion Literature: A Critical Assessment”,” *Economics and Philosophy*, 25 (3), 335–356, [10.1017/S0266267109990277](https://doi.org/10.1017/S0266267109990277).

Tamer, Elie (2010) “Partial Identification in Econometrics,” *Annual Review of Economics*, 2 (1), 167–195, [10.1146/annurev.economics.050708.143401](https://doi.org/10.1146/annurev.economics.050708.143401).

Wald, Abraham (1950) *Statistical Decision Functions*, Wiley Publications in Statistics: Mathematical statistics: Wiley.

Walley, Peter (1991) *Statistical Reasoning with Imprecise Probabilities*: Chapman & Hall.

Wang, Yuan H (1993) “On the number of success in independent trials,” *Statistica Sinica*, 3 (2), 295–312, <http://www.jstor.org/stable/24304959>.

Watson, Geoffrey S. (1996) “Spectral Decomposition of the Covariance Matrix of a Multinomial,” *Journal of the Royal Statistical Society. Series B (Methodological)*, 58 (1), 289–291, <http://www.jstor.org/stable/2346179>.

Watson, James and Chris Holmes (2016) “Approximate Models and Robust Decisions,” *Statistical Science*, 31 (4), 465 – 489, [10.1214/16-STS592](https://doi.org/10.1214/16-STS592).

Wellner, Jon A. (1981) “A Glivenko-Cantelli theorem for empirical measures of independent but non-identically distributed random variables,” *Stochastic Processes and their Applications*, 11 (3), 309–312, [https://doi.org/10.1016/0304-4149\(81\)90033-8](https://doi.org/10.1016/0304-4149(81)90033-8).

White, Halbert (1984) “CHAPTER V - Central Limit Theory,” in White, Halbert ed. *Asymptotic Theory for Econometricians*, 107–131, San Diego: Academic Press, <https://doi.org/10.1016/B978-0-12-746650-7.50009-8>.

Wilson, Edwin B. (1927) “Probable Inference, the Law of Succession, and Statistical Inference,” *Journal of the American Statistical Association*, 22 (158), 209–212, [10.1080/01621459.1927.10502953](https://doi.org/10.1080/01621459.1927.10502953).