# ROBUST DECISION MAKING WITH PARTIALLY CALIBRATED FORECASTS

#### **Anonymous authors**

000

001

002

004

005

010

011

012

013

014

015

017

018

019

023

024

027 028 029

031

032

033

037

038

039

040 041

042

043

044

045

046

047

048

049

050

051

052

054

056

057

Paper under double-blind review

#### **ABSTRACT**

Calibration has emerged as a foundational goal in "trustworthy machine learning", in part because of its strong decision theoretic semantics. Independent of the underlying distribution, and independent of the decision maker's utility function, calibration promises that amongst all policies mapping predictions to actions, the uniformly best policy is the one that "trusts the predictions" and acts as if they were correct. But this is true only of *fully calibrated* forecasts, which are tractable to guarantee only for very low dimensional prediction problems. For higher dimensional prediction problems (e.g. when outcomes are multiclass), weaker forms of calibration have been studied that lack these decision theoretic properties. In this paper we study how a conservative decision maker should map predictions endowed with these weaker ("partial") calibration guarantees to actions, in a way that is robust in a minimax sense: i.e. to maximize their expected utility in the worst case over distributions consistent with the calibration guarantees. We characterize their minimax optimal decision rule via a duality argument, and show that surprisingly, "trusting the predictions and acting accordingly" is recovered in this minimax sense by decision calibration (and any strictly stronger notion of calibration), a substantially weaker and more tractable condition than full calibration. For calibration guarantees that fall short of decision calibration, the minimax optimal decision rule is still efficiently computable, and we provide an empirical evaluation of a natural one that applies to any regression model solved to optimize squared error.

# 1 Introduction

Machine learning systems are increasingly deployed in high-stakes decision making domains such as healthcare, finance, and law. The predictive power of these models can be extraordinary, but scoring well on predictive error metrics does not directly guarantee that decisions downstream of those predictions will be correct. For predictions to be operationally useful, a decision-maker must be able to treat them as reliable inputs into a downstream decision making policy. This raises two fundamental questions:

On the Model Side: What does it mean for machine learning predictions to be trustworthy in decision-making contexts?

**On the Decision Making Side:** Given predictions that satisfy a particular type of "trustworthiness", how should the decision maker adapt its actions to the promised guarantees?

On the Model Side: A natural answer is that trustworthy predictions should directly support good decisions as they are. In other words, the decision-maker should be able to reliably best respond to the forecaster's predictions as if they were correct. Formally, let (X,Y) be a pair of random variables drawn from a joint distribution  $\mathcal{D}$ , where  $X \in \mathcal{X}$  represents the observed features and  $Y \in [0,1]^d$  is the outcome of interest. Let  $\mathcal{A}$  denote the action set, and suppose the decision-maker follows a policy  $a(\cdot):[0,1]^d \to \mathcal{A}$  mapping predictions to actions. Given a predictor f, the decision maker's performance when using a policy a is measured by its expected utility on the underlying distribution:

$$\mathbb{E}_{(X,Y)\sim\mathcal{D}}[u(a(f(X)),Y)],$$

where  $u(a,y) \in \mathbb{R}$  is a utility function. Given a forecaster  $f: \mathcal{X} \to [0,1]^d$ , the *plug-in best response* to a forecast is defined as

$$a_{\rm BR}(f(x)) = \underset{a \in \mathcal{A}}{\arg \max} \ u(a, f(x)).$$
 (1)

Thus, a forecaster f is trustworthy if the decision-maker's best-response policy  $a_{BR}(f(x))$  achieves higher utility than any other policy. When is this the case?

The classical answer lies in the notion of *calibration*. Intuitively, a forecaster is calibrated if, whenever it predicts a vector  $f(x) = v \in [0,1]^d$ , the empirical outcomes are consistent with that prediction. More formally, a forecaster f is said to be *fully calibrated* if for every  $v \in [0,1]^d$ ,

$$\mathbb{E}[Y \mid f(X) = v] = v.$$

It is well known that best responding to calibrated forecasts is the optimal decision policy among all policies that map forecasts to actions (Foster & Vohra, 1997; Kleinberg et al., 2023; Noarov et al., 2023; Roth, 2022).

However, achieving full calibration is extremely difficult, both in theory—the sample complexity of calibrating an existing forecaster without harming its accuracy grows exponentially with the outcome dimension d (Gopalan et al., 2024)—and in practice, where empirical evidence shows systematic deviations from calibration, ranging from neural networks to large language models (Guo et al., 2017; Kull et al., 2019; Gupta & Ramdas, 2022; Plaut et al., 2024). Thus, despite the appealing link between calibration and trustworthy ML-powered decision-making, this connection quickly breaks down in real-world applications.

On the Decision Making Side: Decision making from predictions admits two canonical extremes. At one end, the decision maker aggressively best responds to the forecasts, acting as if they were fully correct. At the other end, the decision maker conservatively plays a minimax-safety strategy,  $\arg\max_{a\in\mathcal{A}}\min_{y\in\mathcal{Y}}u(a,y)$ , treating the forecasts as if they carried no information about the instance.

Departing from these extremes, we treat a model f and it's forecast f(x) as information that constrains what the true, instance-conditional outcome distribution could be. In other words, after observing f(x), the decision maker considers the set of *candidate realities*—outcome distributions consistent with the forecast and the available calibration guarantees. Intuitively, the "volume" of this set is governed by the strength of calibration: under full calibration, the set collapses to the forecast itself (the prediction can be treated as reality, at least in expectation), whereas as calibration weakens, the set enlarges. A principled decision rule should therefore *tune its conservatism to what the reality could be*, consistent with the provided guarantees. This idea, together with the fragility of full calibration in practice, leads to the central question of this paper: *can we derive optimal decision-making policies under weaker and more practical conditions than full calibration?* 

We answer this question affirmatively. We introduce a framework based on *conservative* decision making that nevertheless fully exploits *partially* calibrated forecasts. This viewpoint echoes ideas in robust optimization and control, but it has not been systematically developed for post hoc decision making with partially calibrated machine-learning forecasts.

#### 1.1 OUR RESULTS

We consider a parameterized family of weighted calibration guarantees that have recently become a popular object of study (Hébert-Johnson et al., 2018; Gopalan et al., 2022). Informally speaking this family of guarantees constrains the residuals of a predictor f to be uncorrelated with a collection of "test functions"  $h \in \mathcal{H}$  mapping the range of f to the reals. When  $\mathcal{H}$  consists of all such test functions, we recover full calibration, but many popular variants of calibration (e.g. top label calibration, decision calibration, etc) can be expressed as instances of  $\mathcal{H}$ -calibration under much smaller/more tractable sets  $\mathcal{H}$ . Our contributions are as follows:

- 1. In Section 2 we formalize the following question: given a set of test functions  $\mathcal{H}$  and a predictor f(x) that is promised to satisfy  $\mathcal{H}$ -calibration, what decision rule  $a:[0,1]^d\to\mathcal{A}$ , mapping predictions to actions, will maximize a decision maker's expected utility in the worst case over all joint distributions over  $X\times Y$  that are consistent with the promise that f is  $\mathcal{H}$ -calibrated?
- 2. In Section 3 we answer this question by giving a closed-form for the decision maker's optimal decision rule, in terms of the dual variables of a convex program that can be efficiently computed for any finite  $\mathcal{H}$ .
- 3. In Section 4 we instantiate this decision rule for various calibration guarantees of interest. Of particular note, we find that when \$\mathcal{H}\$ corresponds to the tractable notion of decision calibration (Zhao et al., 2021; Noarov et al., 2023), then the optimal decision rule is the best response decision rule  $a_{\rm BR}$ , just as it is for (the intractable notion of) full calibration. In fact, it suffices that \$\mathcal{H}\$ contains the decision calibration constraints any larger set also makes best response the optimal decision rule. Thus what could have been a very large hierarchy of minimax optimal decision rules "collapses" to best response at the level of decision calibration. An upshot of this is that a predictor can be simultaneously decision calibrated for many downstream decision makers, and for each of them, best response will be their optimal decision policy in this minimax sense. We also derive the minimax optimal decision rule for a simple "self-orthogonality" calibration condition that will hold for any regression model with a linear final layer trained to optimize squared loss, and hence will be commonly satisfied without any algorithmic intervention.
- 4. In Section 5 we train a two-layer MLP to minimize squared loss on two regression datasets, and evaluate both the best-response decision rule and the robust decision rule that results from the self-orthogonality condition of squared error regression. We find that, as predicted by our theory, the robust decision rule outperforms the best-response decision rule under calibration-preserving distribution shift, and that the cost of this robustness is mild even under ideal conditions.

#### 1.2 RELATED WORK

Rothblum & Yona (2023) consider a setting in which both the outcome and decision maker's action set are binary, and study how a decision maker should act to minimize their worst case regret over distributions such that the predictor has maximum calibration error bounded by  $\alpha$ : informally that  $|\mathbb{E}[Y|f(x)=v]-v| \leq \alpha$  for all v. The models f they study are (approximately) fully calibrated, which is a reasonable assumption in their setting, since they limit their study to 1-dimensional outcomes. In contrast, our interest is not (just) in quantitative measures of full calibration error, but rather qualitatively weaker calibration guarantees, as even approximate full calibration becomes intractable in high dimensions.

A line of recent work (Zhao et al., 2021; Kleinberg et al., 2023; Noarov et al., 2023; Roth & Shi, 2024; Hu & Wu, 2024; Okoroafor et al., 2025) has studied the guarantees that can be given to downstream decision makers who best respond to predictions that have weaker guarantees than full calibration (and which in the cases of Zhao et al. (2021); Noarov et al. (2023); Roth & Shi (2024) can be tractably guaranteed in higher dimensional outcome settings). These guarantees take the form of (external and swap) regret bounds, which are qualitatively weaker than the kind of "trustworthiness" promised by full calibration. Informally, regret bounds promise that the decision maker could not have done better by consistently playing a fixed action (or a fixed function remapping their actions to other actions), not that they could not have done better by using a different policy from predictions to actions. A notable exception is Hu & Wu (2024) who do give a weaker notion than full calibration that can recover this trustworthiness guarantee for binary outcome spaces — but their construction takes advantage of the structure of 1-dimensional proper scoring rules, and so does not extend beyond 1-dimensional outcomes. We show that even in high dimensions, the tractable "decision calibration" condition given by Zhao et al. (2021) recovers the same "trustworthiness" semantics of full calibration when viewed through our minimax decision making lens.

Analyzing minimax optimal decision policies is a common way of analyzing *robust* or *risk-averse* decision making guarantees, with deep roots in economics (Gilboa & Schmeidler, 1989; Hansen & Sargent, 2001; Manski, 2000; 2004; Manski & Tetenov, 2007; Manski, 2011), statistics (Wald, 1950), and robust optimization (Ben-Tal & Nemirovski, 2002; Kuhn et al., 2019; Duchi & Namkoong, 2021). For example, Carroll (2015) adopts this lens this in the context of contract theory and Kiyani et al. (2025) and Andrews & Chen (2025) do so in the context of conformal prediction. To the best of our knowledge, we are the first to apply this "robust" minimax lens to the problem of partially calibrated high dimensional forecasts.

# 2 ROBUST DECISION MAKING AND $\mathcal{H}$ -CALIBRATION

In this Section, we define  $\mathcal{H}$ -calibration as a flexible relaxation of full calibration and then introduce a framework to derive minimax optimal decision making policies that are designed to act on forecasters guaranteed to satisfy  $\mathcal{H}$ -calibration. This family of calibration guarantees has been studied extensively in the recent literature on multicalibration and its extensions (Hébert-Johnson et al., 2018; Dwork et al., 2021; Gopalan et al., 2022; Deng et al., 2023) — in particular,  $\mathcal{H}$ -calibration is a special case of what Gopalan et al. (2022) call weighted multicalibration.

 $\mathcal{H}$ -Calibration. Let  $\mathcal{H}$  be a set of functions  $h:[0,1]^d\to\mathbb{R}$ . A forecaster f is said to be  $\mathcal{H}$ -calibrated if for every  $h\in\mathcal{H}$ ,

$$\mathbb{E}[h(f(X))\cdot(Y-f(X))] = 0. \tag{2}$$

Equivalently, writing  $q(v) := \mathbb{E}[Y \mid f(X) = v]$  for the true conditional expectation,  $\mathcal{H}$ -calibration requires

$$\mathbb{E}[h(f(X)) \cdot (q(f(X)) - f(X))] = 0, \quad \forall h \in \mathcal{H}.$$
(3)

This definition captures a spectrum of guarantees. When  $\mathcal{H}$  contains all bounded measurable functions,  $\mathcal{H}$ -calibration reduces to full calibration — i.e. it requires that  $f(v) = q(v) := \mathbb{E}[Y \mid f(X) = v]$  almost surely. For smaller classes  $\mathcal{H}$ , the requirement is weaker and can be seen as a relaxation of calibration, enforcing consistency only with respect to a restricted set of tests.

**Robust Decision Making.** Fix an  $\mathcal{H}$ -calibrated forecaster f. Define the set

$$Q = \left\{ q : [0,1]^d \to [0,1]^d \mid \mathbb{E} \left[ h(f(X)) \cdot (q(f(X)) - f(X)) \right] = 0, \ \forall h \in \mathcal{H} \right\}. \tag{4}$$

In words,  $\mathcal Q$  consists of all candidate conditional expectations consistent with f satisfying  $\mathcal H$ -calibration. Because the perfect predictor  $f(X) = \mathbb E[Y|X]$  satisfies  $\mathcal H$ -calibration for every  $\mathcal H$ , the identity map q(v) = v is always in  $\mathcal Q$ —but in general the set may contain many maps. From the perspective of the decision-maker who knows f and the promised calibration guarantee  $\mathcal H$ , but does not know the underlying distribution, given a forecast f(x), the true expectation  $\mathbb E[Y\mid f(x)]$  is uncertain but must lie within  $\mathcal Q$ . As  $\mathcal H$  grows richer,  $\mathcal Q$  shrinks, eventually reducing to  $\{q(v)=v\}$  in the case of full calibration.

Faced with this uncertainty, a natural strategy is to adopt a robust policy that guards against the worst-case admissible reality. Formally, the robust decision rule is

$$a_{\text{robust}}(\cdot) = \underset{a(\cdot):[0,1]^d \to \mathcal{A}}{\arg \max} \quad \underset{q \in \mathcal{Q}}{\min} \ \mathbb{E}\big[u(a(f(X)), q(f(X)))\big]. \tag{5}$$

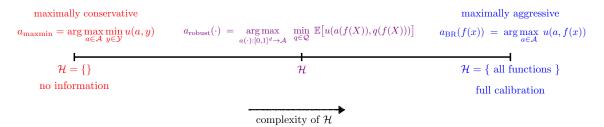


Figure 1: Schematic of interpolating property

That is, the decision-maker chooses an action policy that maximizes utility under the worst-case conditional expectation consistent with calibration guarantees.

Interpolating Property. The robust policy defined in Equation 5 interpolates between two classical extremes (look at Figure 1). When  $\mathcal{H}$  contains all functions, the set  $\mathcal{Q}$  reduces to the singleton  $\{q(v)=v\}$ , and  $a_{\text{robust}}$  coincides with the best-response policy  $a_{\text{BR}}(\cdot)$  (look at equation 1), aggressively using the prediction as if it were correct. At the opposite extreme, when  $\mathcal{H}$  is empty, the set  $\mathcal{Q}$  consists of all functions and the robust policy collapses to the constant *minimax safety strategy*, conservatively optimizing for the worst case, assuming that the predictor f has no relationship to the outcome.

$$a_{\text{Minimax}}(x) = \underset{a \in \mathcal{A}}{\arg \max} \underset{y \in [0,1]^d}{\min} u(a,y), \quad \forall x \in \mathcal{X},$$

Thus, the policy defined by Equation 5 provides a principled bridge between best-responding to calibrated forecasts and adopting fully conservative strategies. More broadly,  $\mathcal{H}$ -calibration offers a flexible language to describe varying levels of conservatism in decision-making, determined by the richness of the calibration guarantees available to the decision-maker.

The central theme of the remainder of this paper is to investigate the interaction between different levels of  $\mathcal{H}$ -calibration and the resulting optimal robust policies. Our focus is not on developing methods for achieving  $\mathcal{H}$ -calibration itself (for which we refer the reader to a rich line of recent work showing how to accomplish this in both the batch and online adversarial setting (Hébert-Johnson et al., 2018; Gopalan et al., 2022; Deng et al., 2023; Noarov et al., 2023; Globus-Harris et al., 2023)), but rather on understanding the decision-making consequences once such guarantees are in place. In the next section, we begin by analyzing the general problem of deriving optimal robust decision rules for arbitrary classes  $\mathcal{H}$ . We then specialize to the important case of decision calibration, showing that this weaker and more practical notion identifies large classes of partially calibrated forecasters for which best responding remains optimal. Beyond its theoretical appeal, this result is also practically useful: when a decision-maker can influence the design or post-processing of the forecaster, they can request a decision-calibrated forecaster, to which they can then simply, reliably, and optimally best respond.

Remark 2.1. Throughout this Paper, we assume the utility function u(a,v) for each  $a \in \mathcal{A}$ , u(a,v) is linear in its second argument  $v \in [0,1]^d$ . This captures, for example, settings in which v represents a distribution over d outcomes, and the decision maker has arbitrary utilities for each action/outcome pair, and wishes to maximize their expected utility given uncertainty (here, the linearity in v follows from the linearity of expectation).

#### 3 OPTIMAL DECISION POLICIES FOR FINITE DIMENSIONAL H-CALIBRATION

In this Section, we characterize the optimal robust decision making policies, i.e., solutions to Equation 5. Throughout this Section, we assume the function class  $\mathcal H$  is a finite dimensional space, i.e. it can be described as span of finitely many functions. Formally, let  $\mathcal H = \operatorname{span}\{h_1,\dots,h_k\}$  be the linear class generated by measurable  $h_i:[0,1]^d\to\mathbb R$ . Then the  $\mathcal H$ -calibration condition equation 3 is equivalent to the k linear moment equalities

$$\mathbb{E}[h_i(f(X)) \cdot (q(f(X)) - f(X))] = 0, \qquad i = 1, \dots, k,$$

so that the ambiguity set in equation 4 may be written as

$$Q = \left\{ q : [0,1]^d \to [0,1]^d \mid \mathbb{E}[h_i(f(X)) \cdot (q(f(X)) - f(X))] = 0 \text{ for } i = 1,\dots,k \right\}.$$

Intuitively, each equality enforces that, conditional on the forecast, the forecast error has zero correlation with the corresponding test  $h_i$ ; taken together, these constraints exhaust the information provided by  $\mathcal{H}$ -calibration criteria and hence precisely describe the admissible reality faced by the robust decision-maker in equation 5.

**Theorem 3.1** (Characterization of the Optimal Robust Policy). Suppose  $\mathcal{H} = \operatorname{span}\{h_1, \ldots, h_k\}$  with each  $h_i$ :  $[0,1]^d \to \mathbb{R}$ , and let  $\mathcal{Q}$  be defined as above. Then the minimax problem in Equation 5 admits a saddle point  $(a_{\text{robust}}, q^*)$  with the following structure:

233

234

235 236 237

238

239 240

241

242

243

244

245 246 247

248

249

250

251

252

253

254

255

256

257

259

260

261

262

263

264

265

266 267

268 269

270

271

272

273 274

275

276

277

278

279

280 281

282

283

284

286

287

289

There exist multipliers  $\lambda^* = (\lambda_1^*, \dots, \lambda_k^*)$  with each  $\lambda_i^* \in \mathbb{R}^d$  such that for almost every forecast v = f(x) the worst-case map  $q^*(v)$  solves

$$q^{\star}(v) \in \arg\min_{p \in [0,1]^d} \Big\{ \operatorname{val}(p) + p \cdot \sum_{i=1}^k h_i(v) \lambda_i^{\star} \Big\}, \quad \textit{where } \operatorname{val}(p) = \max_{a \in \mathcal{A}} u(a,p).$$

Given  $q^*$ , the optimal robust action at v is the best response to  $q^*(v)$ :

$$a_{\text{robust}}(v) \in \arg\max_{a \in \mathcal{A}} u(a, q^{\star}(v)).$$

Interpretation. Theorem 3.1 provides a transparent description of both the worst-case distribution the decisionmaker may face (up to the information encoded by H-calibration) and the corresponding optimal response. Operationally, for any realized forecast  $\nu = f(x)$  the theorem prescribes a two-step procedure: first compute the adversarial belief

$$q^{\star}(\nu) \in \arg\min_{p \in [0,1]^d} \left\{ \operatorname{val}(p) + p \cdot s^{\star}(\nu) \right\}, \qquad s^{\star}(\nu) := \sum_{i=1}^k h_i(\nu) \, \lambda_i^{\star},$$

and then best respond to  $q^*(\nu)$ , i.e. take  $a_{\text{robust}}(\nu) \in \arg\max_{a \in \mathcal{A}} u(a, q^*(\nu))$ . Thus, the optimal policy is *always* a best response, but, in general, not to the raw forecast f(x); rather, it best responds to the adversarially tilted distribution  $q^*(\nu)$  that is chosen to be most challenging under the calibration constraints. A useful byproduct of the theorem is pointwise computability: although  $a_{\text{robust}}$  is a a-priori a policy on  $\mathcal{X}$ , the characterization reduces its evaluation at a given  $\nu = f(x)$  to solving two low-dimensional problems, without constructing the entire mapping  $x \mapsto a_{\text{robust}}(x)$ .

From an optimization perspective, the multipliers  $\lambda^{\star}$  solve a finite-dimensional concave maximization problem (look at the proof of Theorem 3.1), and  $q^*(\nu)$  is obtained by a pointwise convex minimization in  $p \in [0,1]^d$ . Both stages can be carried out by standard, fast methods with provable guarantees (e.g., projected subgradient for the dual, or a simple primal-dual scheme), after which one evaluates  $q^*(\nu)$  via the pointwise minimization and takes the best response  $a_{\text{robust}}(\nu) = \arg \max_{a} u(a, q^{\star}(\nu))$ .

In the next section, we analyze the behavior of the resulting decision rules by specializing to concrete  $\mathcal{H}$ -classes. One might expect that Theorem 3.1 induces a vast and intricate hierarchy of policies whose form depends sensitively on H. Perhaps surprisingly, this is not the case. In particular, we show a sharp transition: for each decision maker, there exists a specific test class, precisely the one associated with decision calibration, such that as soon as  $\mathcal{H}$  contains this class, the adversarial tilt collapses  $(q^*(\nu) = \nu)$  for a.e.  $\nu$ ) and the optimal robust rule reduces to the plug-in best response to the forecaster. This places the spotlight on decision-calibrated forecasters as a practical, strictly weaker alternative to full calibration that nonetheless confers decision-theoretic trustworthiness: the decision-maker can simply, reliably, and optimally best respond to the forecast.

# INSTANTIATIONS OF THE ROBUST POLICY: DECISION CALIBRATION AND BEYOND

In this section, we specialize the general characterization derived in Theorem 3.1 to concrete test classes  $\mathcal{H}$  and derive the associated action policies. Our core result concerns decision calibration: a practically tractable guarantee under which the minimax-optimal robust policy collapses to the plug-in (best-response) rule. This identifies a simple and operational path to decision-theoretic trustworthiness that does not require full calibration.

#### 4.1 DECISION CALIBRATION AND PLUG-IN BEST RESPONSE OPTIMALITY

Here we define the variant of decision calibration given by Noarov et al. (2023), a slight strengthening of the definition originally given by Zhao et al. (2021). Fix a single decision problem with action set A and utility function u(a, v). For each action  $a \in \mathcal{A}$ , let

$$R_a = \left\{ v \in [0,1]^d : u(a,v) \ge u(a',v) \text{ for all } a' \in \mathcal{A} \right\}$$

be the (closed, convex) decision region on which a is a plug-in best response. The *decision-calibration class* is

$$\mathcal{H}_{\mathrm{dec}} = \{ \mathbf{1}_{R_a} : a \in \mathcal{A} \}.$$

Here, we denote  $\mathbf{1}_A(x) := \mathbf{1}\{x \in A\}$ . A forecaster f is decision calibrated if it is  $\mathcal{H}_{\mathrm{dec}}$ -calibrated, i.e.,

$$\mathbb{E}\left[\mathbf{1}_{R_a}(f(X))\left(Y-f(X)\right)\right]=0$$
 for all  $a\in\mathcal{A}$ .

**Theorem 4.1** (Decision calibration  $\Rightarrow$  plug-in best response optimality). If f is  $\mathcal{H}_{dec}$ -calibrated, then the minimax-optimal robust rule in equation 5 coincides with the plug-in best response:

$$a_{\text{robust}}(v) \in \arg\max_{a \in A} u(a, v)$$
 for almost every  $v = f(x)$ .

 $a_{\mathrm{robust}}(v) \in \arg\max_{a \in \mathcal{A}} u(a,v)$  for almost every v = f(x). Equivalently, under decision calibration, best responding to the forecaster is minimax optimal among all forecastbased policies.

Figure 2: Schematic of sharp transition

Put differently, upon observing a forecast v=f(x), the decision-maker need only best respond to v; no adversarial "tilt" survives the decision-calibration constraints. Conceptually, this upgrades the perviously known guarantees of decision calibration, phrased in terms of swap-regret (see Section 1.2 for discussion and references), to *minimax* optimality.

The preceding result assumes that the information conveyed by the forecaster to the decision-maker is exhausted by the decision-calibration tests  $\{\mathbf{1}_{R_a}\}_{a\in\mathcal{A}}$ . In practice, a forecaster might satisfy additional calibration equalities,

$$\mathbb{E}[h(f(X)) \cdot \{Y - f(X)\}] = 0,$$

for functions h beyond the indicators  $\mathbf{1}_{R_a}$ . The next theorem shows that the plug-in optimality conclusion is stable under such enrichments. This is intuitive: if a forecaster is trustworthy, then making it more calibrated (i.e., adding information) should not diminish that trustworthiness.

**Theorem 4.2** (Decision calibration is sufficient, and remains sufficient under richer tests). Let  $\mathcal{H}$  be any test class that contains the decision-calibration indicators,  $\mathcal{H}_{dec} = \{\mathbf{1}_{R_a} : a \in \mathcal{A}\}$ . If f is perfectly  $\mathcal{H}$ -calibrated, then the minimax-optimal robust rule in equation 5 coincides (a.e.) with the plug-in best response:

$$a_{\text{robust}}(v) \in \arg\max_{a \in \mathcal{A}} u(a, v)$$
 for a.e.  $v = f(x)$ .

**Sharp transition.** At first glance, one might expect a *gradual* deconservatization: as  $\mathcal{H}$  is enriched with more tests, the robust policy equation 5 should steadily move from fully conservative toward plug-in best-response. Theorems 4.1–4.2 reveal a sharper picture (look at Figure 2). Once  $\mathcal{H}$  contains the  $|\mathcal{A}|$  decision tests  $\{\mathbf{1}_{R_a}\}_{a\in\mathcal{A}}$ , the adversarial tilt vanishes  $(q^{\star}(\nu) = \nu$  a.e.) and the robust rule *collapses* to the plug-in best response (given in equation 1). Enlarging  $\mathcal{H}$  beyond these indicators does not add conservatism: the minimax-optimal policy remains "trust the forecast and best respond."

Decision calibration is a minimal, task-specific threshold at which robust decision making and plug-in best-response coincide, providing a crisp target for forecaster design and a clear requirement for downstream decision makers.

Note that this result is a qualitatively stronger guarantee than the kinds of *regret* guarantees that have already been shown to be satisfied by decision calibration Zhao et al. (2021); Noarov et al. (2023) — those bound the performance of best responding to decision calibrated forecasts with respect to various benchmarks, but do not establish that best responding to decision calibrated forecasts are the optimal decision rule. Their optimality depends on the minimax view that we introduce in this work.

As a byproduct, this leads to another practical advantage of decision calibration: a single forecaster can be made simultaneously reliable for a *collection* of downstream decision problems. Intuitively, if the forecast passes the decision calibration tests of each problem, then none of the decision makers needs additional robustness, the plug-in best-response is minimax-optimal for all of them.

**Corollary 4.3** (Simultaneous plug-in optimality across multiple decisions). Let  $u_1, \ldots, u_m$  be m decision problems, with respective action sets  $A_j$  and linear utilities  $u_j(a, v)$  in  $v \in [0, 1]^d$ . For each j and  $a \in A_j$ , let

$$R_{a,j} = \{ v \in [0,1]^d : u_j(a,v) \ge u_j(a',v) \text{ for all } a' \in \mathcal{A}_j \}$$

be the plug-in decision region of action a in problem j, and define the combined test class

$$\mathcal{H}^{ ext{all}}_{ ext{dec}} \ = \ igcup_{j=1}^m ig\{ \, \mathbf{1}_{R_{a,j}} : \ a \in \mathcal{A}_j \, ig\}.$$

If f is  $\mathcal{H}$ -calibrated for some  $\mathcal{H}$  satisfying  $\mathcal{H}^{\mathrm{all}}_{\mathrm{dec}} \subseteq \mathcal{H}$ , then for every  $j \in \{1, \ldots, m\}$  the minimax-optimal robust policy for problem j coincides (a.e.) with the plug-in best response:

$$a_{\mathrm{robust},j}(v) \in \arg\max_{a \in \mathcal{A}_j} u_j(a,v)$$
 for a.e.  $v = f(x)$ .

*Proof.* For each problem j, the included indicators  $\{\mathbf{1}_{R_{a,j}}\}_{a\in\mathcal{A}_j}$  ensure that  $\mathcal{H}$  contains the decision-calibration tests of problem j. Theorem 4.2 then applies verbatim to each j, yielding plug-in optimality problem by problem.

#### 4.2 Beyond Decision Calibration: Generic $\mathcal{H}$ -Specifications from Training Pipelines

Thus far we have focused on *decision calibration*, which, when attainable, collapses  $a_{\rm robust}$  to the plug-in best response. In practice, two regimes arise. (i) If one can influence the forecaster's training pipeline, decision calibration is the natural target: it is minimal and task-specific, and our sharp-transition results guarantee plug-in minimax optimality. (ii) If one *cannot* control training, the forecaster might not be decision calibrated for the downstream task. Identifying its precise partial-calibration profile may be difficult, yet certain moment conditions arise *structurally* from standard training procedures. Here we give two examples of how to leverage such "free" structure to specify usable  $\mathcal{H}$ 's and derive the associated robust policies.

**Self-orthogonality from squared-loss training.** A ubiquitous example is *self-orthogonality* (a form of self-calibration) that follows from first-order optimality when a model with a linear last layer is trained to minimize mean squared error. This includes the universally adopted cases of regression with either a linear model or a neural network with a linear head, trained by mean squared error.

**Proposition 4.4** (Self-orthogonality under squared loss). Let  $X \mapsto z_{\phi}(X) \in \mathbb{R}^k$  be a representation and  $f_{\theta}(X) = Wz_{\phi}(X) \in \mathbb{R}^d$  a linear last layer. Suppose  $\theta = (\phi, W)$  is trained to a first-order stationary point of the expected squared loss

$$\mathcal{L}(\theta) = \frac{1}{2} \mathbb{E} \left[ \left\| f_{\theta}(X) - Y \right\|_{2}^{2} \right].$$

Then the following calibration moments hold:

$$\mathbb{E}[z_{\phi}(X)\left(Y - f_{\theta}(X)\right)^{\top}] = 0 \quad and \quad \mathbb{E}[f_{\theta}(X)\left(Y - f_{\theta}(X)\right)^{\top}] = 0.$$

In particular,  $f_{\theta}$  is  $\mathcal{H}$ -calibrated for the test class  $\mathcal{H} = \{h_j(v) = e_j^{\top}v : j = 1, ..., d\}$  (and for any linear combination thereof).

**Implications.** Proposition 4.4 supplies a generic, pipeline-induced  $\mathcal{H}$ -calibration guarantee whenever a linear head is trained to stationarity under squared loss. Specializing Theorem 3.1 to this case yields a particularly simple dual. For d=1 (for example the case of 1-dimensional regression) and  $\mathcal{H}=\{h(v)=v\}$ , the multiplier is a scalar  $\lambda$ , and for each forecast  $\nu=f(x)$  the worst-case distribution becomes

$$q^{\star}(\nu) \;\in\; \arg\min_{p\in[0,1]} \Big\{ \mathrm{val}(p) + \lambda\,\nu\,p \Big\}, \qquad \mathrm{val}(p) = \max_{a\in\mathcal{A}} \big\{\,u(a,p)\,\big\}.$$

The robust action is the best response to  $q^*(\nu)$ :

$$a_{\text{robust}}(\nu) \in \arg\max_{a \in \mathcal{A}} u(a, q^{\star}(\nu)).$$

When u(a, p) is linear in p and  $\mathcal{A}$  is finite, val is convex piecewise linear in p, so the inner minimization reduces to evaluating a finite set of candidate points (the endpoints and pairwise breakpoints of val). The dual objective

$$G(\lambda) \ = \ \mathbb{E}\big[\min_{p \in [0,1]} \{ \operatorname{val}(p) + \lambda \, f(X) \, p \} \big] \ - \ \lambda \, \mathbb{E}\big[f(X)^2\big]$$

is concave in  $\lambda$  and can be maximized by standard one-dimensional methods with provable guarantees (e.g., bisection on a monotone subgradient). In higher dimensions (d>1) the correction term,  $\lambda \nu p$  becomes  $\Lambda \nu p$  for a matrix of multipliers  $\Lambda$ , and the pointwise inner problem remains a small convex program over  $p \in [0,1]^d$ ; for finite  $\mathcal{A}$  and linear utilities, it is again efficiently solvable.

**Zero-bias and bin-wise calibration.** A widely available source of partial calibration comes from *post-hoc recalibration* that many practitioners already apply (mean correction, histogram binning, isotonic-style step fits on a held-out split). These procedures enforce generic (not task-specific) moment constraints that are directly usable in our framework. We focus on *bin-wise* calibration: take a partition of the forecast range into bins  $\{B_1, \ldots, B_J\}$  and enforce, for each bin,

$$\mathbb{E}\Big[\mathbf{1}_{\{f(X)\in B_j\}}\left(Y-f(X)\right)\Big]=0, \qquad j=1,\ldots,J.$$

This corresponds to the test class  $\mathcal{H}_{\text{bin}} = \{\mathbf{1}_{B_j} : j = 1, \dots, J\}$ , and reduces to zero-bias when J=1 with  $B_1 = [0,1]^d$ .

**Proposition 4.5** (Robust policy under bin-wise calibration). *Let the utility be linear in the outcome and the action set* A *be finite. If* f *is*  $\mathcal{H}_{bin}$ -calibrated, then with

$$m_i := \mathbb{E}[f(X) \mid f(X) \in B_i] = \mathbb{E}[Y \mid f(X) \in B_i],$$

the worst-case belief is piecewise constant

$$q^{\star}(v) = m_j \quad \text{for } v \in B_j \text{ (a.e.)},$$

and the robust action best-responds to the bin mean:

$$a_{\text{robust}}(v) \in \arg\max_{a \in \mathcal{A}} \{u(a, m_j)\}$$
 for  $v \in B_j$  (a.e.).

**Implications.** Bin-wise calibration  $\mathcal{H}_{\mathrm{bin}}$  can be obtained cheaply via standard post-hoc methods (histogram binning or isotonic regression), and Proposition 4.5 yields an especially simple, closed-form characterization of the robust policy. Computing  $a_{\mathrm{robust}}$  reduces to: (i) estimating  $m_j$  on a calibration split, and (ii) at test time, mapping v to its bin  $B_j$  and best-responding to  $m_j$ . No additional optimization is needed to compute actions. As a special case, when J=1 we recover the global-mean constraint  $\mathbb{E}[Y-f(X)]=0$ . Then  $q^*$  is constant,  $q^*(v)\equiv \bar{m}$ , with  $\bar{m}=\mathbb{E}[f(X)]=\mathbb{E}[Y]$ , and the robust rule ignores v and plays  $\arg\max_{a\in\mathcal{A}}u(a,\bar{m})$ . As the partition is refined, the robust rule moves from a single global plug-in best response at  $\bar{m}$  to a piecewise plug-in best response at  $m_j$ , yielding a richer, finer-grained decision policy.

# 5 EXPERIMENTS

In this section, we empirically evaluate the validity and practical consequences of our framework by implementing our methods on two real-world datasets. We compare the *plug-in best response* ( $a_{\rm BR}$ ), a standard and widely used baseline, against the *robust policy* ( $a_{\rm robust}$ ), which enjoys minimax optimality guarantees under  $\mathcal{H}$ -calibration.

We focus on two classes of metrics. *Nominal performance* measures average utility when the test data are i.i.d. from the same distribution as the training and calibration splits; this reflects an optimistic regime that often degrades in practice. *Adversarial performance* probes the other extreme by altering the test-time outcome distribution in two ways: (i) a worst case tailored to the plug-in policy, and (ii) a worst case induced by the robust dual, tailored to the robust policy. In both cases, the adversarial distributions respect the  $\mathcal{H}$ -calibration constraints and are therefore indistinguishable, from the decision-maker's perspective, from i.i.d. test draws given an  $\mathcal{H}$ -calibrated forecaster.

Our theory predicts that two patterns should emerge. First, by minimax optimality, the robust policy should dominate the plug-in rule when each is evaluated against its own worst-case distribution (and typically dominate the plug-in rule under the adversary tuned to hurt the plug-in). Second, because  $(a_{\text{robust}}, q^*)$  forms a saddle point of equation 5, when both policies are evaluated under the robust-tuned adversary, the robust policy should not underperform the plug-in rule. Under nominal i.i.d. evaluation, the plug-in rule may achieve higher utility, reflecting the lack of need for conservatism in that regime.

#### 5.1 CASE STUDIES: BIKE SHARING AND CALIFORNIA HOUSING

We evaluate our framework on two real-world regression datasets with distinct decision-making interpretations.

**Bike Sharing (UCI).** The UCI *Bike Sharing* (daily) dataset Fanaee-T & Gama (2014) records daily rider counts alongside calendar and weather covariates (season, month, weekday, holiday, working day, weather state, temperature, humidity, wind). The outcome  $Y \in [0,1]$  is the rescaled total rider count, and the decision-maker chooses a staffing/capacity multiplier from  $\mathcal{A} = \{0.8, 1.0, 1.2\}$ , interpretable as conservative, nominal, and aggressive provisioning.

California Housing. The California Housing dataset Pace & Barry (1997) records median house values (rescaled to [0,1]) with demographic and geographic covariates (median income, housing age, population, latitude/longitude, etc.). Here the decision-maker chooses an investment multiplier from  $\mathcal{A} = \{0.6, 0.75, 0.90\}$ , interpretable as conservative, nominal, and aggressive investment.

Utility specification. In both settings we adopt the utility

$$u(a, y) = \alpha \, a \, y - C(a),$$

which is linear in y. The benefit term  $\alpha$  a y captures service or return proportional to realized outcome y, scaled by  $\alpha > 0$ . The cost term C(a) grows in a, penalizing aggressive choices via over-provisioning costs or investment risk. This form tunes the under/over-trade-off without departing from linearity. For Bike Sharing we use  $(\alpha, C) = (0.9, \{0.02, 0.05, 0.1\})$ , while for California Housing we use  $(\alpha, C) = (0.9, \{0.02, 0.05, 0.20\})$ . The qualitative conclusions of this Section remain the same under other reasonable parameter choices.

Table 1: Mean utility on the test set under natural i.i.d. evaluation and two adversarial evaluations. Adversaries respect  $\mathcal{H}$ -calibration ( $\mathcal{H} = \{h(v) = v\}$ ).

Dataset	i.i.d.		Worst-case for robust		Worst-case for plug-in	
	Plug-in	Robust	Plug-in	Robust	Plug-in	Robust
Bike Sharing (UCI)	0.474	0.463	0.402	0.410	0.393	0.412
California Housing	0.216	0.207	0.160	0.164	0.155	0.166

Forecasting model. In both datasets, the forecaster f is a two-layer MLP trained by MSE. By the self-orthogonality property of linear heads under squared loss (see Section §??), the learned forecaster approximately satisfies  $\mathcal{H}$ -calibration with  $\mathcal{H} = \{h(v) = v\}$ , which is the calibration constraint used to derive the robust policy  $a_{\text{robust}}$ . All experiments use an i.i.d. train/calibration/test split (60/20/20). We use the calibration data to substitute any population level expectation that is needed to be computed to derive  $a_{\text{robust}}$ .

**Results.** Table 1 reports the mean utilities. The results match theory: under adversaries tailored to the robust policy, the robust rule achieves at least the plug-in performance; under adversaries tuned to harm the plug-in rule, the robust policy secures noticeably higher utility, reflecting its minimax protection. Moreover, the robust policy outperforms the plug-in best response when each is evaluated against its own worst-case distribution.

#### 6 CONCLUSION AND LIMITATIONS

We developed a decision-theoretic framework for acting on partially calibrated forecasts via a minimax-optimal robust policy over  $\mathcal{H}$ -calibrated forecasters. We then identified a sharp transition in the behavior of these policies: for any decision problem with m actions, there exist m decision tests (the decision-calibration class) such that, once they are included in  $\mathcal{H}$ , the robust policy collapses to the plug-in best response. This spotlights decision calibration as a natural requirement whenever the decision-maker can influence the training pipeline. Moreover, even when decision calibration is unavailable, we showed that generic properties induced by standard training and post hoc procedures (e.g., self-orthogonality under squared loss and bin-wise calibration) yield usable test classes  $\mathcal{H}$  and tractable robust policies within our framework.

Our guarantees and saddle-point characterizations are stated at the population level, assuming exact  $\mathcal{H}$ -calibration. In practice,  $\mathcal{H}$ -constraints are estimated from finite calibration/validation splits, introducing finite-sample error. Quantifying the resulting finite-sample degradation is an important direction for future work. We also assumed u(a,v) is linear in v and that  $\mathcal{A}$  is finite; these cover many expected-utility settings. However, extending to nonlinear (e.g., risk-averse) utilities could broaden the scope of our results.

#### REFERENCES

Isaiah Andrews and Jiafeng Chen. Certified decisions. arXiv preprint arXiv:2502.17830, 2025.

Aharon Ben-Tal and Arkadi Nemirovski. Robust optimization–methodology and applications. *Mathematical programming*, 92(3):453–480, 2002.

Gabriel Carroll. Robustness and linear contracts. American Economic Review, 105(2):536–563, 2015.

Zhun Deng, Cynthia Dwork, and Linjun Zhang. Happymap: A generalized multicalibration method. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, pp. 41–1. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2023.

John C Duchi and Hongseok Namkoong. Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics*, 49(3):1378–1406, 2021.

Cynthia Dwork, Michael P Kim, Omer Reingold, Guy N Rothblum, and Gal Yona. Outcome indistinguishability. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 1095–1108, 2021.

Hadi Fanaee-T and Joao Gama. Event labeling combining ensemble detectors and background knowledge. *Progress in Artificial Intelligence*, 2(2):113–127, 2014.

Dean P Foster and Rakesh V Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, 1997.

Itzhak Gilboa and David Schmeidler. Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, 18(2):141–153, 1989.

Ira Globus-Harris, Declan Harrison, Michael Kearns, Aaron Roth, and Jessica Sorrell. Multicalibration as boosting for regression. In *International Conference on Machine Learning*, pp. 11459–11492. PMLR, 2023.

Parikshit Gopalan, Michael P Kim, Mihir A Singhal, and Shengjia Zhao. Low-degree multicalibration. In *Conference on Learning Theory*, pp. 3193–3234. PMLR, 2022.

- Parikshit Gopalan, Lunjia Hu, and Guy N Rothblum. On computationally efficient multi-class calibration. In *The Thirty Seventh Annual Conference on Learning Theory*, pp. 1983–2026. PMLR, 2024.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pp. 1321–1330. PMLR, 2017.
- Chirag Gupta and Aaditya Ramdas. Top-label calibration and multiclass-to-binary reductions. In *International Conference on Learning Representations*. OpenReview, 2022.
- Lars Peter Hansen and Thomas J Sargent. Robust control and model uncertainty. *American Economic Review*, 91 (2):60–66, 2001.
- Ursula Hébert-Johnson, Michael Kim, Omer Reingold, and Guy Rothblum. Multicalibration: Calibration for the (computationally-identifiable) masses. In *International Conference on Machine Learning*, pp. 1939–1948. PMLR, 2018.
- Lunjia Hu and Yifan Wu. Predict to minimize swap regret for all payoff-bounded tasks. In 65th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2024, Chicago, IL, USA, October 27-30, 2024, pp. 244–263. IEEE, 2024.
- Shayan Kiyani, George Pappas, Aaron Roth, and Hamed Hassani. Decision theoretic foundations for conformal prediction: Optimal uncertainty quantification for risk-averse agents, 2025. URL https://arxiv.org/abs/2502.02561.
- Bobby Kleinberg, Renato Paes Leme, Jon Schneider, and Yifeng Teng. U-calibration: Forecasting for an unknown agent. In Gergely Neu and Lorenzo Rosasco (eds.), *The Thirty Sixth Annual Conference on Learning Theory, COLT 2023, 12-15 July 2023, Bangalore, India*, volume 195 of *Proceedings of Machine Learning Research*, pp. 5143–5145. PMLR, 2023.
- Daniel Kuhn, Peyman Mohajerin Esfahani, Viet Anh Nguyen, and Soroosh Shafieezadeh-Abadeh. Wasserstein distributionally robust optimization: Theory and applications in machine learning. In *Operations research & management science in the age of analytics*, pp. 130–166. Informs, 2019.
- Meelis Kull, Miquel Perello Nieto, Markus Kängsepp, Telmo Silva Filho, Hao Song, and Peter Flach. Beyond temperature scaling: Obtaining well-calibrated multi-class probabilities with dirichlet calibration. *Advances in neural information processing systems*, 32, 2019.
- Charles F Manski. Identification problems and decisions under ambiguity. *Journal of Econometrics*, 95(2):415–442, 2000.
- Charles F Manski. Statistical treatment rules for heterogeneous populations. *Econometrica*, 72(4):1221–1246, 2004.
- Charles F Manski. Choosing treatment policies under ambiguity. Annual Review of Economics, 3:25–49, 2011.
- Charles F Manski and Aleksey Tetenov. Admissible treatment rules for a risk-averse planner. *Econometrica*, 75 (3):715–752, 2007.
- Georgy Noarov, Ramya Ramalingam, Aaron Roth, and Stephan Xie. High-dimensional prediction for sequential decision making. *arXiv preprint arXiv:2310.17651*, 2023.
- Princewill Okoroafor, Robert Kleinberg, and Michael P Kim. Near-optimal algorithms for omniprediction. *arXiv* preprint arXiv:2501.17205, 2025.
- R Kelley Pace and Ronald Barry. Sparse spatial autoregressions. *Statistics & Probability Letters*, 33(3):291–297, 1997.
- Benjamin Plaut, Nguyen X Khanh, and Tu Trinh. Probabilities of chat Ilms are miscalibrated but still predict correctness on multiple-choice q&a. *arXiv preprint arXiv:2402.13213*, 2024.
- Aaron Roth. Uncertain: Modern topics in uncertainty estimation. Lecture Notes, 11:30–31, 2022.

Aaron Roth and Mirah Shi. Forecasting for swap regret for all downstream agents. In Dirk Bergemann, Robert Kleinberg, and Daniela Sabán (eds.), *Proceedings of the 25th ACM Conference on Economics and Computation, EC 2024, New Haven, CT, USA, July 8-11, 2024*, pp. 466–488. ACM, 2024.

- Guy N Rothblum and Gal Yona. Decision-making under miscalibration. In *14th Innovations in Theoretical Computer Science Conference*, *ITCS 2023*, pp. 92. Schloss Dagstuhl-Leibniz-Zentrum fur Informatik GmbH, Dagstuhl Publishing, 2023.
- Abraham Wald. Statistical decision functions. In *Breakthroughs in Statistics: Foundations and Basic Theory*, pp. 342–357. Springer, 1950.
- Shengjia Zhao, Michael Kim, Roshni Sahoo, Tengyu Ma, and Stefano Ermon. Calibrating predictions to decisions: A novel approach to multi-class calibration. *Advances in Neural Information Processing Systems*, 34:22313–22324, 2021.

# **Appendix**

#### A PROOFS

#### **Proof of Theorem 3.1**

*Proof.* We begin from the robust formulation

$$\max_{a(\cdot):\mathcal{X}\to\mathcal{A}} \min_{q\in\mathcal{Q}} \mathbb{E}[u(a(f(X)), q(f(X)))], \tag{6}$$

where  $\mathcal{A} \subset \mathbb{R}^m$  is compact,  $u(\cdot,\cdot)$  is linear in its second component,  $\mathcal{Q}$  is the nonempty, convex, and compact set of measurable maps  $q:[0,1]^d \to [0,1]^d$  satisfying the linear moment equalities in equation 4, and  $a(\cdot)$  ranges over measurable policies with values in  $\mathcal{A}$ . The mapping  $(a,q) \mapsto \mathbb{E}[u(a(f(X)),q(f(X)))]$  is convex in q (since  $u(a,\cdot)$  is linear, hence convex, in q and expectation preserves convexity), concave in q (as a pointwise maximum over linear functionals in q on the compact set q). Hence, by Sion's minimax theorem,

$$\max_{a(\cdot)} \ \min_{q \in \mathcal{Q}} \ \mathbb{E} \big[ u(a(f(X)), q(f(X))) \big] \ = \ \min_{q \in \mathcal{Q}} \ \max_{a(\cdot)} \mathbb{E} \big[ u(a(f(X)), q(f(X))) \big].$$

Fix any  $q \in \mathcal{Q}$ . The inner maximization over policies separates pointwise in v = f(x), yielding the value function

$$\operatorname{val}(p) \; \triangleq \; \max_{a \in \mathcal{A}} u(a,p) \quad \text{and} \quad \max_{a(\cdot)} \; \mathbb{E}\big[u\big(a(f(X)),q(f(X))\big)\big] = \mathbb{E}\big[\operatorname{val}\big(q(f(X))\big)\big] \; .$$

Therefore the robust value equals the convex adversarial problem

$$\min_{q \in \mathcal{Q}} \mathbb{E}[\operatorname{val}(q(f(X)))], \tag{7}$$

which will be analyzed via Lagrangian duality below.

Introduce vector Lagrange multipliers  $\lambda_i \in \mathbb{R}^d$  for the d-dimensional equalities in equation 4, and let  $\lambda = (\lambda_1, \dots, \lambda_k)$ . Define

$$s(v) \triangleq \sum_{i=1}^{k} h_i(v) \lambda_i \in \mathbb{R}^d, \quad v \in [0, 1]^d.$$

The Lagrangian of equation 7 is

$$L(q,\lambda) = \mathbb{E}\big[\mathrm{val}\big(q(f(X))\big)\big] + \sum_{i=1}^k \lambda_i \cdot \mathbb{E}\big[h_i\big(f(X)\big)\big(q(f(X)) - f(X)\big)\big].$$

By linearity of expectation,

$$L(q,\lambda) = \mathbb{E}\Big[\mathrm{val}\big(q(f(X))\big) + q(f(X)) \cdot s\big(f(X)\big) - f(X) \cdot s\big(f(X)\big)\Big].$$

The dual function is obtained by minimizing  $L(q, \lambda)$  over measurable  $q : [0, 1]^d \to [0, 1]^d$ . Since the integrand depends on q only through q(f(X)), the infimum can be taken *pointwise* in the forecast value v = f(X):

$$G(\lambda) \ = \ \inf_{q} L(q,\lambda) \ = \ \mathbb{E} \Big[ \ \inf_{p \in [0,1]^d} \big\{ \mathrm{val}(p) + p \cdot \mathit{s} \big( f(X) \big) \big\} \ \Big] \ - \ \mathbb{E} \big[ f(X) \cdot \mathit{s} \big( f(X) \big) \big].$$

The primal problem equation 7 is convex (convex objective, affine constraints) and feasible (e.g., q(v) = v), thereby strong duality holds. Hence,

$$\min_{q \in \mathcal{Q}} \mathbb{E} \big[ \mathrm{val} \big( q(f(X)) \big) \big] \; = \; \max_{\lambda \in (\mathbb{R}^d)^k} G(\lambda),$$

and there exists a maximizing multiplier  $\lambda^*$ . Define

$$s^{\star}(v) \triangleq \sum_{i=1}^{k} h_i(v) \lambda_i^{\star} \in \mathbb{R}^d.$$

By the definition of  $G(\lambda)$  and strong duality, any primal optimizer  $q^* \in \mathcal{Q}$  must minimize the Lagrangian at  $\lambda^*$ . Since the dependence on q is only through q(f(X)), this yields the pointwise characterization, for v = f(x) almost surely,

$$q^{\star}(v) \in \arg\min_{p \in [0,1]^d} \left\{ \operatorname{val}(p) + p \cdot s^{\star}(v) \right\}.$$

With  $q^*$  fixed, define the policy

$$a_{\text{robust}}(v) \in \arg\max_{a \in A} u(a, q^*(v)).$$

Then, by the definition of val and the construction of  $q^*$ ,

$$\max_{a(\cdot)} \mathbb{E}\big[u\big(a(f(X)),q^{\star}(f(X))\big)\big] = \mathbb{E}\big[\mathrm{val}\big(q^{\star}(f(X))\big)\big] = \min_{q \in \mathcal{Q}} \mathbb{E}\big[\mathrm{val}\big(q(f(X))\big)\big]\,,$$

which shows that  $(a_{\text{robust}}, q^*)$  is a saddle point of equation 6. In particular,  $a_{\text{robust}}$  is optimal for the outer maximization, and  $q^*$  is worst–case optimal for the inner minimization, with  $q^*$  characterized pointwise by the minimization problem above and determined by the dual multiplier  $\lambda^*$ . This matches the statement of Theorem 3.1 and completes the proof.

### **Proof of Theorem 4.1:**

*Proof.* We use the reduction

$$\max_{a(\cdot)} \min_{q \in \mathcal{Q}} \mathbb{E} \big[ u(a(f(X)), q(f(X))) \big] \ = \ \min_{q \in \mathcal{Q}} \mathbb{E} \big[ \mathrm{val}(q(f(X))) \big],$$

established in the proof of Theorem 3.1. Fix the decision regions

$$R_a = \{ v \in [0,1]^d : u(a,v) \ge u(a',v) \ \forall \ a' \in \mathcal{A} \},$$

each convex. Under  $\mathcal{H}_{dec} = \{\mathbf{1}_{R_a} : a \in \mathcal{A}\}$ , admissible q satisfy

$$\mathbb{E}\big[\mathbf{1}_{R_a}(f(X))\{q(f(X)) - f(X)\}\big] = 0 \quad \forall a,$$

equivalently (whenever  $\mathbb{P}(f(X) \in R_a) > 0$ ),

$$\mathbb{E}[q(f(X)) \mid f(X) \in R_a] = \mathbb{E}[f(X) \mid f(X) \in R_a] =: \mu_a \in R_a.$$

By Jensen's inequality (convexity of val), for any  $q \in \mathcal{Q}$  and any a,

$$\mathbb{E}[\operatorname{val}(q(f(X))) \mid f(X) \in R_a] \geq \operatorname{val}(\mu_a).$$

Define the piecewise-constant  $\bar{q}(v) = \sum_a \mu_a \mathbf{1}_{R_a}(v)$ . Then  $\bar{q} \in \mathcal{Q}$  and, conditionally on  $f(X) \in R_a$ , we have  $\bar{q}(f(X)) = \mu_a$  a.s., hence the bound is attained:

$$\mathbb{E}\big[\mathrm{val}\big(\bar{q}(f(X))\big)\big] = \sum_{a} \mathbb{P}(f(X) \in R_a) \, \mathrm{val}(\mu_a) \le \mathbb{E}\big[\mathrm{val}\big(q(f(X))\big)\big] \quad \forall q \in \mathcal{Q}.$$

Thus a worst-case belief is  $q^* = \bar{q}$ , region-wise constant with  $q^*(v) = \mu_a$  on  $R_a$ .

Finally, since  $\mu_a \in R_a$ , by definition of  $R_a$  we have  $u(a, \mu_a) \ge u(a', \mu_a)$  for all a', so a is a best response to  $\mu_a$ . Therefore the robust action at  $v \in R_a$  is

$$a_{\text{robust}}(v) \in \arg\max_{a'} u(a', q^*(v)) = \arg\max_{a'} u(a', \mu_a) \ni a,$$

which coincides (a.e.) with the plug-in best response to v. This proves Theorem 4.1.

#### **Proof of Theorem 4.2:**

Recall  $val(p) = \max_{a \in \mathcal{A}} u(a, p)$  and the reduction

$$\max_{a(\cdot)} \min_{q \in \mathcal{Q}_{\mathcal{H}}} \, \mathbb{E} \big[ u \big( a(f(X)), q(f(X)) \big) \big] \; = \; \min_{q \in \mathcal{Q}_{\mathcal{H}}} \, \mathbb{E} \big[ \mathrm{val} \big( q(f(X)) \big) \big] \, ,$$

established earlier in the proof of Theorem 3.1. Moreover, the identity map  $q_{id}(v) = v$  always lies in  $\mathcal{Q}_{\mathcal{H}}$  (the perfect forecaster is consistent with every  $\mathcal{H}$ -calibration constraint), so for any policy  $a(\cdot)$ ,

$$\min_{q \in \mathcal{Q}_{\mathcal{H}}} \mathbb{E}\left[u\left(a(f(X)), q(f(X))\right)\right] \leq \mathbb{E}\left[u\left(a(f(X)), f(X)\right)\right]. \tag{8}$$

Let  $a_{BR}(v) \in \arg \max_{a \in \mathcal{A}} u(a, v)$  be a plug-in best response. We show that, assuming  $\mathcal{H}$  contains the decision-calibration tests  $\{\mathbf{1}_{R_a}\}_{a \in \mathcal{A}}$ ,

$$\mathbb{E}\left[u\left(a_{\mathrm{BR}}(f(X)), q(f(X))\right)\right] = \mathbb{E}\left[u\left(a_{\mathrm{BR}}(f(X)), f(X)\right)\right] \quad \forall q \in \mathcal{Q}_{\mathcal{H}}.\tag{9}$$

<sup>&</sup>lt;sup>1</sup>Fix any deterministic tie-breaking so that  $a_{BR}$  and the regions  $R_a = \{v : a_{BR}(v) = a\}$  are measurable.

Write  $\mu_a := \mathbb{E}[f(X) \mid f(X) \in R_a]$  whenever  $\mathbb{P}(f(X) \in R_a) > 0$  (if  $\mathbb{P}(f(X) \in R_a) = 0$ , any choice of  $\mu_a$  is harmless since the corresponding terms vanish). Then

$$\mathbb{E}\left[u\left(a_{\mathrm{BR}}(f(X)),q(f(X))\right)\right] = \sum_{a\in\mathcal{A}} \mathbb{E}\left[u\left(a,q(f(X))\right)\mathbf{1}_{\{f(X)\in R_a\}}\right]$$

$$\stackrel{(i)}{=} \sum_{a\in\mathcal{A}} \mathbb{P}(f(X)\in R_a)\ u(a,\ \mathbb{E}[q(f(X))\mid f(X)\in R_a])$$

$$\stackrel{(ii)}{=} \sum_{a\in\mathcal{A}} \mathbb{P}(f(X)\in R_a)\ u(a,\ \mathbb{E}[f(X)\mid f(X)\in R_a])$$

$$= \sum_{a\in\mathcal{A}} \mathbb{P}(f(X)\in R_a)\ u(a,\mu_a)$$

$$\stackrel{(iii)}{=} \sum_{a\in\mathcal{A}} \mathbb{E}\left[u\left(a,f(X)\right)\mathbf{1}_{\{f(X)\in R_a\}}\right]$$

$$= \mathbb{E}\left[u\left(a_{\mathrm{BR}}(f(X)),f(X)\right)\right].$$

Here: (i) uses that  $u(a,\cdot)$  is linear in its second argument, so  $\mathbb{E}[u(a,q(f(X)))\mid f(X)\in R_a]=u(a,\mathbb{E}[q(f(X))\mid f(X)\in R_a]);$  (ii) uses the decision-calibration equalities  $\mathbb{E}[\mathbf{1}_{R_a}(f(X))\{q(f(X))-f(X)\}]=0$ , equivalently  $\mathbb{E}[q(f(X))\mid f(X)\in R_a]=\mathbb{E}[f(X)\mid f(X)\in R_a]=\mu_a$  whenever  $\mathbb{P}(f(X)\in R_a)>0$ ; and (iii) again uses linearity:  $u(a,\mu_a)=u(a,\mathbb{E}[f(X)\mid f(X)\in R_a])=\mathbb{E}[u(a,f(X))\mid f(X)\in R_a].$ 

Combining equation 8, the optimality of best response on the *perceived* outcomes,

$$\mathbb{E}[u(a(f(X)), f(X))] \leq \mathbb{E}[u(a_{BR}(f(X)), f(X))]$$
 for all policies  $a(\cdot)$ ,

and the invariance equation 9, we obtain the minimax dominance

$$\min_{q \in \mathcal{Q}_{\mathcal{H}}} \mathbb{E}\big[u\big(a_{\mathrm{BR}}(f(X)),q(f(X))\big)\big] \ = \ \mathbb{E}\big[u\big(a_{\mathrm{BR}}(f(X)),f(X)\big)\big] \ \geq \ \min_{q \in \mathcal{Q}_{\mathcal{H}}} \ \mathbb{E}\big[u\big(a(f(X)),q(f(X))\big)\big] \ ,$$

for every forecast-based policy  $a(\cdot)$ . Hence the plug-in best response is minimax optimal under any  $\mathcal H$  that contains the decision-calibration tests, as claimed.

# **Proof of Proposition 4.4**:

*Proof.* Assume  $\mathbb{E}||z_{\phi}(X)||_{2}^{2} < \infty$  and  $\mathbb{E}||Y||_{2}^{2} < \infty$  so that all derivatives and expectations below are well-defined and we may interchange expectation and differentiation by dominated convergence. Write  $z := z_{\phi}(X) \in \mathbb{R}^{k}$  and  $f := f_{\theta}(X) = Wz \in \mathbb{R}^{d}$ . The squared-loss risk is

$$\mathcal{L}(\theta) = \frac{1}{2} \mathbb{E}[\|f - Y\|_2^2] = \frac{1}{2} \mathbb{E}[(Wz - Y)^{\top}(Wz - Y)].$$

For the linear head  $W \in \mathbb{R}^{d \times k}$ , the gradient with respect to W satisfies the standard identity

$$\nabla_W \left( \frac{1}{2} \|Wz - Y\|_2^2 \right) = (Wz - Y) z^\top \in \mathbb{R}^{d \times k}.$$

Taking expectation and interchanging  $\nabla$  with  $\mathbb{E}$  yields

$$\nabla_W \mathcal{L}(\theta) = \mathbb{E}[(f - Y) z^\top].$$

At a first-order stationary point (in particular, when the gradient with respect to W vanishes) we have

$$\mathbb{E}\big[(f-Y)\,z^{\top}\big] = 0_{d\times k}.$$

Transposing gives

$$\mathbb{E}[z(f-Y)^{\top}] = 0_{k \times d} \iff \mathbb{E}[z(Y-f)^{\top}] = 0_{k \times d},$$

which is the first claimed moment identity.

For the second identity, observe that f = Wz, hence

$$\mathbb{E}\big[f\left(Y-f\right)^{\top}\big] \;=\; \mathbb{E}\big[Wz\left(Y-f\right)^{\top}\big] \;=\; W\,\mathbb{E}\big[z\left(Y-f\right)^{\top}\big] \;=\; W\,\mathbf{0}_{k\times d} \;=\; \mathbf{0}_{d\times d}.$$

Therefore both  $\mathbb{E}[z_{\phi}(X)(Y-f_{\theta}(X))^{\top}]=0$  and  $\mathbb{E}[f_{\theta}(X)(Y-f_{\theta}(X))^{\top}]=0$  hold. In particular, for each coordinate  $j=1,\ldots,d$ ,  $\mathbb{E}[e_{j}^{\top}f_{\theta}(X)(Y-f_{\theta}(X))^{\top}]=0$  and  $\mathbb{E}[z_{\phi}(X)e_{j}^{\top}(Y-f_{\theta}(X))]=0$ , so  $f_{\theta}$  is  $\mathcal{H}$ -calibrated for  $\mathcal{H}=\{h_{j}(v)=e_{j}^{\top}v:\ j=1,\ldots,d\}$  and for any linear combination thereof. This proves the proposition.

#### **Proof of Proposition 4.5:**

*Proof.* By the reduction established earlier (see the proof of Theorem 3.1), the robust problem

$$\max_{a(\cdot)} \min_{q \in \mathcal{Q}} \mathbb{E}[u(a(f(X)), q(f(X)))]$$

with linear utilities and finite A is equivalent to the convex program

$$\min_{q \in \mathcal{Q}} \mathbb{E}\left[\operatorname{val}(q(f(X)))\right], \qquad \operatorname{val}(p) := \max_{a \in \mathcal{A}} u(a, p),$$

subject to the  $\mathcal{H}_{\mathrm{bin}}$ -calibration constraints

$$\mathbb{E}[\mathbf{1}_{\{f(X)\in B_i\}}(q(f(X)) - f(X))] = 0, \quad j = 1,\dots, J.$$

Write  $E_j := \{f(X) \in B_j\}$  and assume  $\mathbb{P}(E_j) > 0$  (bins with zero probability are immaterial). Then the constraints are equivalent to

$$\mathbb{E}[q(f(X)) \mid E_i] = \mathbb{E}[f(X) \mid E_i] =: m_i, \qquad j = 1, \dots, J.$$

Because  $u(a, \cdot)$  is linear in the outcome, val is the pointwise maximum of linear maps and hence convex. Decomposing by bins and applying Jensen's inequality gives, for any feasible q,

$$\mathbb{E}\left[\operatorname{val}(q(f(X)))\right] = \sum_{j=1}^{J} \mathbb{P}(E_j) \,\mathbb{E}\left[\operatorname{val}(q(f(X))) \mid E_j\right]$$

$$\geq \sum_{j=1}^{J} \mathbb{P}(E_j) \,\operatorname{val}(\mathbb{E}[q(f(X)) \mid E_j])$$

$$= \sum_{j=1}^{J} \mathbb{P}(E_j) \,\operatorname{val}(m_j).$$

Define the piecewise-constant candidate

$$\bar{q}(v) := \sum_{i=1}^{J} m_j \, \mathbf{1}_{B_j}(v).$$

Then  $\bar{q}$  is feasible, since for each i,

$$\mathbb{E}\left[\mathbf{1}_{E_i}\left(\bar{q}(f(X)) - f(X)\right)\right] = \mathbb{P}(E_j)\left(m_j - \mathbb{E}[f(X) \mid E_j]\right) = 0,$$

and it attains the Jensen lower bound because  $\bar{q}(f(X)) = m_j$  almost surely on  $E_j$ :

$$\mathbb{E}\left[\operatorname{val}(\bar{q}(f(X))) \mid E_j\right] = \operatorname{val}(m_j).$$

Therefore  $\bar{q}$  is an optimizer, and any minimizer  $q^*$  can be chosen (a.e.) piecewise constant with  $q^*(v) = m_j$  for  $v \in B_j$ .

Finally, fixing such a  $q^*$ , the robust action at forecast  $v \in B_i$  solves

$$a_{\text{robust}}(v) \in \arg\max_{a \in \mathcal{A}} u(a, q^{\star}(v)) = \arg\max_{a \in \mathcal{A}} u(a, m_j),$$

which depends only on the bin index, i.e., it is the best response to the bin mean. This proves the claim.  $\Box$