# LEARNING UNANIMOUSLY ACCEPTABLE LOTTERIES FROM NOISY APPROVAL QUERIES

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Randomization can make collective agreement possible even when no deterministic choice satisfies every agent. We study how to find a unanimously acceptable lottery using only noisy binary approval queries, when agents have unknown utilities and acceptance thresholds. Our model allows feedback to become uninformative near acceptance thresholds, making agreement at these boundaries difficult to resolve. We develop algorithms for exact consensus under finite precision and for a robust formulation that allows either answer on borderline instances while preserving unanimous acceptance guarantees. Complementary lower bounds establish the optimal dependence on precision and margin parameters, up to logarithmic factors. We also distinguish finding one acceptable lottery from learning conservative approximations to every agent's acceptance set. Learning sets that exclude rejected lotteries and retain all lotteries with sufficient margin requires quadratically many queries in the number of alternatives, even when a common lottery with substantial margin is supplied. Checking a proposed lottery with sufficient margin, by comparison, has a query cost independent of that number. We exploit this distinction to develop learning-augmented algorithms that save queries when a predicted lottery is useful while preserving worst-case guarantees for arbitrary predictions. Finally, we establish when exact decisions remain possible without known noise parameters and show how knowledge of the response rule can further reduce query complexity.

## 1 INTRODUCTION

Suppose a decision maker must choose between two options, and two groups disagree about which is better. The first group accepts a randomized choice only if the first option is chosen with probability at least $1/2$; the second group requires the same probability for the second option. Neither option alone is acceptable to both groups. The equal mixture is, but it is the only such choice. If either group raises its requirement by even a small amount, agreement becomes impossible.

This simple example becomes a learning problem when the groups' requirements are unknown. The decision maker can propose a randomized choice, called a *lottery*, and ask each group whether it is acceptable. Such queries may be noisy: repeated human assessments can disagree, and evaluations on finite samples can give different answers (Dawid & Skene, 1979; Snow et al., 2008; Cotter et al., 2019). At the equal mixture in our example, a group whose requirement is met exactly may answer accept or reject with equal probability. Repeating that query therefore need not reveal whether agreement is possible. An error about this single lottery changes the answer to the entire problem.

Our model captures AI deployment decisions that must meet several minimum standards, as in model evaluation and auditing (Mitchell et al., 2019; Raji et al., 2020). Consider a choice among model variants or policies with different rules for seeking human review. An affected group may require minimum accuracy on its requests, a safety team may impose an upper bound on the rate of harmful responses, and an operator may limit expected review cost. We represent each requirement by an *agent*: unanimity means meeting every requirement separately, so better accuracy cannot compensate for exceeding a harm limit. When a lottery selects a policy with fixed probabilities, independently of the request, each expected score is the probability-weighted average of the policies' scores. A mixture can therefore meet all requirements even when no individual policy does, as in fair and constrained classification (Agarwal et al., 2018; Cotter et al., 2019; Hardt et al., 2016).



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

The decision maker proposes such mixtures to evaluators and receives only noisy binary (accept/reject) answers, without observing their scores or thresholds. Appendix B gives the precise interpretation of a safety requirement and a query.

Choo et al. (2026) introduced the problem with truthful, noiseless answers. Their algorithms ask about mixtures of pairs of alternatives and use the answers to determine whether all agents accept a common lottery. This is a form of preference elicitation: learning enough about unknown preferences to make a decision (Angluin, 1988; Boutilier & Rosenschein, 2016; Conitzer & Sandholm, 2002; Lu & Boutilier, 2020).

We ask how many queries are needed when answers are *noisy*. The central issue is whether the algorithm must resolve agreement exactly, as in our equal mixture example, or may leave some borderline cases undecided. We also distinguish finding one acceptable lottery from learning all acceptable lotteries for each agent. These objectives require different amounts of information.

### 1.1 Contributions

We consider $n$ agents and $m$ alternatives. The *margin* of a lottery for an agent is the difference between the agent's expected utility and her threshold. A positive margin means that the requirement is exceeded; zero means it is met exactly. Under *uniform noise*, every answer is correct with probability at least $1 - \eta > 1/2$. Under *near-threshold noise*, an answer at nonzero margin $g$ is correct with probability at least $1/2 + \lambda|g|^\beta$. Thus, $\lambda > 0$ controls the guaranteed accuracy and $\beta > 0$ describes how quickly this guarantee approaches random guessing near a threshold. Section 2 gives the full definitions; Table 1 summarizes the bounds.

**Exact agreement.** Uniform noise can be handled by repeating each query of a noiseless algorithm (Section 2.2). Near-threshold noise requires more care. Suppose utilities and thresholds are multiples of $\varepsilon$. Although these numbers are spaced by $\varepsilon$, the mixture probabilities at which acceptance changes (turning points) can be only order $\varepsilon^2$ apart. We query between possible turning points and obtain precision dependence $\varepsilon^{-4\beta}$. A lower bound matches this dependence up to logarithmic factors, already with two alternatives (Section 3). With arbitrary real utilities and thresholds, no finite worst-case expected query bound is possible.

**Agreement with a margin gap.** For $\rho > 0$, we require a lottery to be found whenever some lottery exceeds every threshold by at least $\rho$. We require a negative answer whenever agreement remains impossible even after every threshold is lowered by $\rho$. In the intervening cases a negative answer is allowed, but every returned lottery must still satisfy all original thresholds. This *robust consensus* problem has margin dependence $\rho^{-2\beta}$, up to logarithmic factors, without a finite-precision assumption (Section 4).

**Learning acceptable choices and checking a proposal.** Our robust algorithm learns a set for each agent that contains every lottery of margin at least $\rho$ and contains no rejected lottery. Learning all these sets requires quadratically many queries in $m$, even when a common margin-$1/4$ lottery is given (Theorem 4.4). This lower bound applies to learning the sets; whether consensus search needs the same dependence remains open. In comparison, checking a proposed lottery with common margin at least $\rho$ costs $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ queries, where $\delta$ is the allowed error probability. This cost is independent of $m$. We combine such checks with a general search and retain its worst-case bound for any proposal (Section 5).

**What the algorithm knows about the noise.** Without known noise parameters or a margin bound, an exact procedure terminates with finite expected query count on every infeasible instance and every instance with a common positive margin. Termination cannot be guaranteed on all remaining feasible instances (Appendix H). If the full response rule is known, responses can reveal numerical margins. For $\beta = 1$, the exact precision cost improves from $\varepsilon^{-4}$ to $\varepsilon^{-2}$, and learning the sets becomes linear in $m$ (Appendix I).

### 1.2 Related Work

**Learning acceptable lotteries.** Our starting point is the noiseless model of Choo et al. (2026), who learn one agent's acceptance set using $O(m\log(2/\varepsilon))$ queries under finite precision. We use their reconstruction from two-alternative mixtures and their $\varepsilon^2$ lower bound on the separation of possible

Table 1: Query bounds, with logarithmic factors omitted. Here $n$ counts agents, $m$ counts alternatives, $\varepsilon$ is utility/threshold spacing, and $\rho$ is the margin gap. The notation $\widetilde{O}_\beta$ and $\widetilde{\Omega}_\beta$ hides logarithmic factors in the displayed parameters and $1/\delta$; constants may depend on $\beta$. Precise bounds appear in the stated results. The final row concerns learning a set $\widehat{A}_i$ with $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i$ for every agent, even when a common margin-$1/4$ lottery is supplied.

| **Task and noise model** | **Upper bound** | **Lower bound** |
|---|---|---|
| Uniform noise: simulate $D$ queries | $\widetilde{O}(D(1-2\eta)^{-2})$ | — |
| Near-threshold: exact consensus | $\widetilde{O}_\beta(nm\lambda^{-2}\varepsilon^{-4\beta})$ | $\widetilde{\Omega}_\beta(\lambda^{-2}[(n-1)\varepsilon^{-4\beta} + m\varepsilon^{-2\beta}])$ |
| Near-threshold: robust consensus | $\widetilde{O}_\beta(nm^2\lambda^{-2}\rho^{-2\beta})$ | $\widetilde{\Omega}_\beta((n+m-1)\lambda^{-2}\rho^{-2\beta})$ |
| Near-threshold: learn all inner sets | $\widetilde{O}_\beta(nm^2\lambda^{-2}\rho^{-2\beta})$ | $\widetilde{\Omega}_\beta(nm^2\lambda^{-2}\rho^{-2\beta})$ |

Near-threshold upper bounds assume known $\lambda, \beta$. Lower bounds are for worst-case expected query counts and $\delta \in (0, 1/4)$. The exact row assumes $\varepsilon$-precision; its lower bound has $n, m \geq 2$ and $\varepsilon \leq 1/4$. The robust lower bound has $n, m \geq 2$ and $\rho \leq 1/10$. The last lower bound has even $m \geq 4$ and $\rho \leq 1/16$. Under near-threshold noise, exact consensus without finite precision or a margin gap has no finite worst-case expected query bound.

turning points. Their searches may query a zero-margin lottery, where near-threshold noise makes repetition inconclusive. Our additional step is to choose queries between possible turning points and bound their margins away from zero, giving exact consensus with precision dependence $\varepsilon^{-4\beta}$ and a matching lower bound up to logarithmic factors. Without finite precision, our robust formulation instead learns inner approximations that retain every lottery with sufficient margin. We establish the corresponding $\rho^{-2\beta}$ dependence and a quadratic lower bound in $m$ for learning these sets, even when a common acceptable lottery is supplied; this is not a lower bound for consensus search.

Choo et al. (2026) also use predictions and can check an acceptable proposed lottery with one noiseless query per agent. We extend this approach to noisy answers: a common positive margin permits reliable checking with a query cost independent of $m$, whereas acceptance with zero margin need not. Combining these checks with general search preserves the worst-case bound for arbitrary predictions.

**Noisy learning and search.** Blum et al. (1998) study concept learning when queries near a decision boundary may be unreliable. Castro & Nowak (2008) analyze active learning under conditions on the noise and the boundary. A central distinction is the objective: small classification error under a distribution can tolerate errors at individual boundary points, whereas our exact decision may depend on just one such point. Noisy decision tree computation (Feige et al., 1994) gives the repetition argument for uniform noise. Near-threshold noise requires either choosing queries away from a boundary or allowing tests to leave some cases unresolved. Appendix A discusses further connections to social choice, geometric learning, and stochastic response models.

## 2 Model and Preliminaries

We separate what makes a lottery acceptable from what a noisy answer reveals about its acceptance. For a positive integer $z$, let $[z] := \{1, \ldots, z\}$. There are $n$ *agents*, indexed by $N := [n]$, and $m$ *alternatives*, $S := \{s_1, \ldots, s_m\}$. A *lottery* $\mathbf{x} = (x_1, \ldots, x_m)$ assigns probability $x_j$ to alternative $s_j$. The set of lotteries is the *probability simplex* $\Delta(S) := \{\mathbf{x} \in \mathbb{R}^m_{\geq 0} : \sum_{j=1}^m x_j = 1\}$. The *pure lottery* $\mathbf{e}_j$ chooses $s_j$ with probability one.

Agent $i$ has unknown utilities $u_i = (u_{i1}, \ldots, u_{im}) \in [0,1]^m$ and an unknown threshold $\tau_i \in (0, 1]$. The expected utility of $\mathbf{x}$ is $\langle u_i, \mathbf{x}\rangle = \sum_j u_{ij}x_j$. Its *margin* is $g_i(\mathbf{x}) := \langle u_i, \mathbf{x}\rangle - \tau_i$, and agent $i$'s *acceptance set* is $A_i := \{\mathbf{x} \in \Delta(S) : g_i(\mathbf{x}) \geq 0\}$. Positive, zero, and negative margins mean that utility exceeds, equals, and falls below the threshold. Equality counts as acceptance. Since $\sum_j x_j = 1$, we have $g_i(\mathbf{x}) = \sum_j x_j g_i(\mathbf{e}_j)$: a lottery's margin is the weighted average of its pure lottery margins. The true answer, or *noiseless label*, is $q_i(\mathbf{x}) := \mathbf{1}\{g_i(\mathbf{x}) \geq 0\}$, with 1 for acceptance and 0 for rejection. The acceptance condition is a linear inequality, so $A_i$ is a *halfspace restricted to*

*the simplex*. Unless every lottery is accepted or every lottery is rejected, the lotteries with $g_i(\mathbf{x}) = 0$ form the boundary of $A_i$ within the simplex.

An *edge* of the simplex consists of mixtures of two alternatives. For distinct $k, k' \in [m]$, let $\mathbf{x}_{k,k'}(\alpha) := (1-\alpha)\mathbf{e}_k + \alpha\mathbf{e}_{k'}$, where $\alpha \in [0,1]$ is the probability of $s_{k'}$. For agent $i$, AcceptAll and RejectAll mean $A_i = \Delta(S)$ and $A_i = \varnothing$, respectively.

The following example shows how changing the thresholds can reduce the common acceptance set from an interval to a single boundary lottery, or make it empty.

**Example: agreement between two agents.** Let $u_1 = (1, 0)$ and $u_2 = (0, 1)$. With $\tau_1 = \tau_2 = 0.4$, neither pure lottery is accepted by both agents, but every $\mathbf{x}_{1,2}(\alpha)$ with $0.4 \le \alpha \le 0.6$ is. In particular, $g_1(\mathbf{x}_{1,2}(1/2)) = g_2(\mathbf{x}_{1,2}(1/2)) = 0.1$. Raising both thresholds to $0.5$ leaves only the equal mixture, with zero margin for both agents. Raising them to $0.6$ makes agreement impossible. Figure 1 shows these three cases.

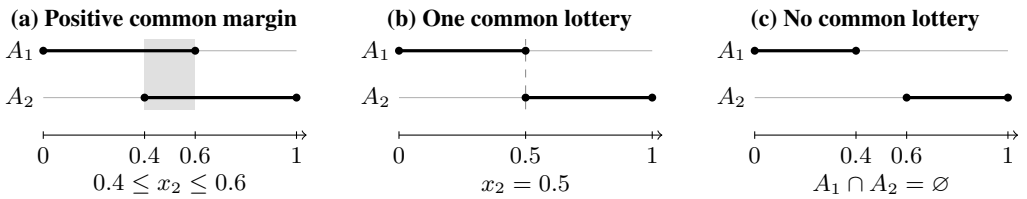


Figure 1: Two alternatives to illustrate the difficulty. Horizontal axes give $x_2 = \alpha$ in $\mathbf{x}_{1,2}(\alpha)$. Each thick segment shows one agent's accepted values, including its filled endpoints. Thresholds are $0.4$, $0.5$, and $0.6$ for both agents in panels (a), (b), and (c). In (b), the only common lottery is on both boundaries, where near-threshold responses may be unbiased.

**Consensus objectives.** A lottery is *acceptable* to agent $i$ if it belongs to $A_i$, and it is *unanimously acceptable* if it belongs to $\bigcap_{i \in N} A_i$.

**Definition 2.1** (Exact consensus)**.** In the *exact consensus* problem, an algorithm outputs either a lottery or NULL; a lottery is correct if it lies in $\bigcap_{i \in N} A_i$, and NULL is correct if and only if $\bigcap_{i \in N} A_i = \varnothing$. For $\delta \in (0, 1)$, the algorithm is $(1-\delta)$-correct if, for every instance and every response rule satisfying the stated noise model, it returns a correct output with probability at least $1 - \delta$.

The robust formulation asks for a decision when the instance is sufficiently far from the borderline cases. Fix $\rho \in (0, 1]$ and define $A_i^{+\rho} := \{\mathbf{x} \in \Delta(S) : g_i(\mathbf{x}) \ge \rho\}$, $A_i^{-\rho} := \{\mathbf{x} \in \Delta(S) : g_i(\mathbf{x}) \ge -\rho\}$. Thus $A_i^{+\rho} \subseteq A_i \subseteq A_i^{-\rho}$: the first set raises the acceptance threshold by $\rho$, while the last lowers it by $\rho$.

**Definition 2.2** ($\rho$-robust consensus)**.** The $\rho$-robust consensus problem distinguishes the following two cases: YES : $\bigcap_{i \in N} A_i^{+\rho} \neq \varnothing$ and NO : $\bigcap_{i \in N} A_i^{-\rho} = \varnothing$. An instance may satisfy neither condition. The algorithm must output either a lottery or NULL. For $\delta \in (0, 1)$, an algorithm is $(1-\delta)$-correct if, for every instance and every response rule satisfying the stated noise model, the following requirements hold with probability at least $1 - \delta$: it outputs a lottery in $\bigcap_{i \in N} A_i$ in the YES case, it outputs NULL in the NO case, and whenever it outputs a lottery, including when neither condition holds, that lottery belongs to $\bigcap_{i \in N} A_i$.

The parameter $\rho$ separates the two required cases; it is not a guaranteed margin of the returned lottery. In Section 4, we will see why the output margin cannot in general be required to be a fixed positive fraction of $\rho$, even in the YES case. For example, with $\rho = 0.05$, Figure 1(a) is a robust YES instance and 1(c) is a robust NO instance. Figure 1(b) satisfies neither condition: a robust algorithm may return the equal mixture or NULL, whereas an exact algorithm must return the equal mixture. On the correctness event, a robust NULL output rules out common margin $\rho$, not necessarily every unanimously acceptable lottery.

**Finite precision for exact consensus.** Fix $\varepsilon \in (0, 1]$ with $Q := 1/\varepsilon \in \mathbb{N}$. An instance has $\varepsilon$-*precision* if each $u_{ij}$ and $\tau_i$ is an integer multiple of $\varepsilon$. For example, $\varepsilon = 0.01$ allows utilities and thresholds specified to two decimal places. The algorithm may still query arbitrary lotteries.

On an edge where acceptance changes, the *turning point* is the value of $\alpha$ at which $g_i(\mathbf{x}_{k,k'}(\alpha)) = 0$. It is a ratio of differences of utilities and thresholds. Under $\varepsilon$-precision it therefore lies in

$$R_Q := \{p/q \in [0,1] : 1 \le q \le Q,\ 0 \le p \le q,\ \gcd(p,q) = 1\},$$

the rational numbers whose reduced denominator is at most $Q$. Exact edge searches will distinguish these finitely many possibilities. The robust results and the procedure for unknown parameters in Appendix H do not assume finite precision.

### 2.1 Noisy Queries

At round $t$, the algorithm chooses an agent-lottery pair $(i_t, \mathbf{x}_t) \in N \times \Delta(S)$ and receives a bit $Y_t \in \{0,1\}$: 1 denotes acceptance and 0 rejection. A randomized rule producing these responses is an *oracle*. Query complexity counts these responses, not the computation between queries. Unless stated otherwise, upper bounds hold on every execution; lower bounds concern the worst-case expected count over instances and *admissible* oracles, namely those satisfying the noise model.

Let $\mathcal{F}_{t-1}$ contain the *query history* (earlier queries and answers) and the algorithm's private random seed. The next query and stopping decision may depend on this information. The *response mean* is the conditional probability of an accept answer; its difference from $1/2$ is the *bias*. Responses need not be independent: the following bounds must hold after every history, apart from histories of probability zero.

Under *uniform noise*, the probability of a correct answer is bounded away from $1/2$ at every lottery, as in bounded classification noise (Massart & Nédélec, 2006).

**Definition 2.3** (Uniform noise)**.** The oracle satisfies $\mathsf{UniformNoise}(\eta)$, with $\eta \in [0, 1/2)$, if, at every round $t$, $\Pr(Y_t = q_{i_t}(\mathbf{x}_t) \mid \mathcal{F}_{t-1}) \ge 1 - \eta$.

Under *near-threshold noise*, answers may become less reliable as utility approaches the acceptance threshold, as in stochastic choice and boundary-noise models (Train, 2009; Blum et al., 1998; Locatelli et al., 2018).

**Definition 2.4** (Near-threshold noise)**.** Fix $\lambda \in (0, 1/4]$ and $\beta > 0$. The oracle satisfies $\mathsf{NearThresholdNoise}(\lambda, \beta)$ if, at every round $t$,

- $g_{i_t}(\mathbf{x}_t) > 0 \implies \mathbb{E}[Y_t \mid \mathcal{F}_{t-1}] \ge \frac{1}{2} + \lambda |g_{i_t}(\mathbf{x}_t)|^\beta$,
- $g_{i_t}(\mathbf{x}_t) < 0 \implies \mathbb{E}[Y_t \mid \mathcal{F}_{t-1}] \le \frac{1}{2} - \lambda |g_{i_t}(\mathbf{x}_t)|^\beta$,
- $g_{i_t}(\mathbf{x}_t) = 0 \implies \mathbb{E}[Y_t \mid \mathcal{F}_{t-1}] \ge \frac{1}{2}$.

At a boundary point, acceptance is the true answer, but the responses may be unbiased. Away from the boundary, $\lambda|g|^\beta$ lower-bounds the advantage over random guessing; larger $\beta$ allows this advantage to decrease faster near zero. The model gives inequalities, not a specified response function. In particular, a more positive margin need not give a larger response mean, and the mean may depend on the preceding history.

**What repeated answers reveal.** If the conditional response mean is always at least $1/2 + b$, or always at most $1/2 - b$, for a known $b > 0$, then $O(b^{-2}\log(1/\delta))$ repeated answers determine its side of $1/2$ with error probability at most $\delta$. This follows from the conditional concentration argument of Hoeffding (1963); Azuma (1967) and does not require independent replies (Appendix C.1). Near-threshold noise gives such a positive bound away from a boundary, but not at the boundary itself. In Figure 1(b), repeating queries at the only common lottery may therefore produce nothing more than fair coin flips.

Without a known positive lower bound on the bias, no test can distinguish every positive bias from every negative bias and also stop in finite expected time at zero bias (Proposition G.1). Appendix H explains what exact decisions remain possible when the noise parameters are unknown.

### 2.2 Uniform Noise: Simulating Noiseless Queries

Under uniform noise, repeat each query and take a majority vote. This works even at an acceptance boundary. A noiseless algorithm making at most $D$ queries on any answer sequence can be simulated

with probability at least $1-\delta$ using at most $D\lceil 2\log(D/\delta)/(1-2\eta)^2\rceil$ noisy queries, preserving its adaptive choices and early stopping (Feige et al., 1994). Applying this to the $O(nm\log(2/\varepsilon))$-query finite-precision routine of Choo et al. (2026) gives the uniform-noise bound in Table 1; Appendix D gives both proofs. The remaining sections address near-threshold noise, where repeating a query at zero margin need not reveal its true label.

## 3 Exact Consensus under Near-Threshold Noise

Exact consensus requires finding a common lottery even when it lies on every agent's acceptance boundary. Asking about that lottery repeatedly may yield only fair coin flips. Our solution is to learn the acceptance sets by querying other lotteries, then find a point in their intersection. We use the noiseless reconstruction of Choo et al. (2026), but change its queries: finite precision gives a finite list of possible turning points, and we query between them.

On an edge from a rejected $\mathbf{e}_k$ to an accepted $\mathbf{e}_{k'}$, the margin increases with $\alpha$. Its turning point $\alpha^*$ belongs to $R_Q$, as defined in the preliminaries. A midpoint between consecutive elements of $R_Q$ avoids every possible turning point. We need more than nonzero margin, however: the following margin bound determines how many repetitions make each comparison reliable.

**Lemma 3.1.** *Under $\varepsilon$-precision, suppose $q_i(\mathbf{e}_k)=0$ and $q_i(\mathbf{e}_{k'})=1$. Then the turning point $\alpha^*$ of $g_i(\mathbf{x}_{k,k'}(\alpha))$ lies in $R_Q$. If $r_h<r_{h+1}$ are consecutive elements of $R_Q$ and $z=(r_h+r_{h+1})/2$, then $z\neq\alpha^*$ and $|g_i(\mathbf{x}_{k,k'}(z))|\geq\varepsilon^2/2$.*

A bound on the distance between turning points alone would not suffice: utility may change slowly along an edge. Finite precision also relates that rate of change to the turning point's denominator, giving the stated margin bound. The example in Appendix E.1 attains this $\varepsilon^2$ scale. Thus the guaranteed bias may be only of order $\lambda\varepsilon^{2\beta}$, giving repetition cost $\lambda^{-2}\varepsilon^{-4\beta}$.

To use these edge comparisons, we first identify which pure lotteries are accepted. A rejected pure lottery has margin at most $-\varepsilon$, whereas an accepted one has response mean at least $1/2$, including at equality. Comparing the fraction of accept answers with a cutoff slightly below $1/2$ therefore distinguishes the two cases.

If all pure lotteries are accepted or all are rejected, then every lottery is accepted or every lottery is rejected; we return the description AcceptAll or RejectAll for that agent. Otherwise, the known pure lottery labels and at most $m-1$ selected turning points determine an inequality $\langle\mathbf{c}_i,\mathbf{x}\rangle\geq 1$ describing exactly the agent's acceptance set (Choo et al., 2026); see Lemma C.5. This inequality suffices; we do not need to estimate $u_i$ and $\tau_i$ separately. We then find a lottery satisfying every agent's inequality, or conclude that none exists. Algorithm 3 in Appendix E.5 gives the details.

**Theorem 3.2.** *Under $\varepsilon$-precision and near-threshold noise with known parameters $(\lambda,\beta)$, for every $\delta\in(0,1)$ there is a $(1-\delta)$-correct exact-consensus algorithm with query complexity $O_\beta(nm\lambda^{-2}\varepsilon^{-4\beta}\log(2/\varepsilon)\log(2nm\log(2/\varepsilon)/\delta))$.*

The dependence on $\varepsilon$ comes from distinguishing nearby turning points. Next, we show that this cost is necessary even with two alternatives, and it can be incurred separately for each of $n-1$ agents.

**Theorem 3.3.** *Fix $n\geq 2$, $m=2$, $\varepsilon\in(0,1/4]$, and $\delta\in(0,1/4)$. Under $\varepsilon$-precision and* NearThresholdNoise$(\lambda,\beta)$*, every $(1-\delta)$-correct exact-consensus algorithm has worst-case expected query complexity $\Omega_\beta((n-1)\lambda^{-2}\varepsilon^{-4\beta}\log(1/\delta))$.*

The lower bound allows arbitrary adaptive queries, not just edge searches: the difficulty is deciding whether agreement exists, rather than reconstructing every acceptance set. Appendix E.6 gives the proof.

For one agent, an accepted lottery exists if and only if some pure lottery is accepted. Using the construction of Choo et al. (2026), in which any pure lottery may be the only accepted choice, Theorem G.2 gives a lower bound of $\Omega(m\lambda^{-2}\varepsilon^{-2\beta}\log(1/\delta))$. Checking the pure lotteries matches it up to logarithmic factors. Thus the larger $\varepsilon^{-4\beta}$ cost arises from agreement between agents. Appendix G.3 combines these bounds for all $n,m\geq 2$ as in Table 1; they do not show that every alternative incurs the $\varepsilon^{-4\beta}$ cost.

Finite precision makes a finite worst-case query bound possible. Without it, two turning points can be separated by an arbitrarily small distance $\Delta > 0$, and the response distributions can be correspondingly difficult to distinguish, as illustrated by the following result.

**Proposition 3.4.** *Fix $\delta \in (0, 1/4)$ and assume* NearThresholdNoise$(\lambda, \beta)$. *If utilities and thresholds may be arbitrary real numbers, then for every $\Delta \in (0, 1/4]$ there are two instances with $n = m = 2$ such that any algorithm returning a correct exact output with probability at least $1 - \delta$ on each of them needs $\Omega(\lambda^{-2}\Delta^{-2\beta}\log(1/\delta))$ expected queries on one of them. Since $\Delta$ can be arbitrarily small, exact consensus has no finite worst-case expected query bound.*

## 4 Robust Consensus without Finite Precision

Without finite precision, we cannot enumerate the possible turning points. The robust objective changes what must be learned: we may leave the boundary unresolved, provided we find agreement whenever a common margin-$\rho$ lottery exists. This permits tests whose cost depends on $\rho$ rather than parameter precision. Learning correctly labelled regions while leaving boundary points undecided is also studied by Locatelli et al. (2018); here each learned region must also contain every lottery with margin at least $\rho$.

Checking a few lotteries for each agent is insufficient: agreement may require a mixture we have not checked. We therefore learn a set $\widehat{A}_i$ containing only lotteries that agent $i$ accepts, while retaining every lottery with margin at least $\rho$. Such a set is called an *inner approximation*, and its two requirements are $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i$. The first inclusion keeps every possible common choice with margin $\rho$; the second ensures acceptance. Hence the learned sets have a common point in every robust YES instance, and any common point is unanimously acceptable. They have no common point in a robust NO instance. Figure 2 illustrates one agent's set.

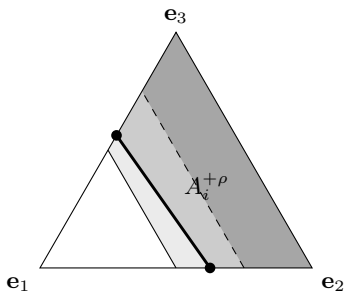


For $u_i = (0, 1, 1)$, $\tau_i = 1/2$, and $\rho = 1/4$:

$$A_i = \{\mathbf{x} : x_1 \leq 1/2\}, \quad A_i^{+\rho} = \{\mathbf{x} : x_1 \leq 1/4\},$$

where both sets are restricted to $\Delta(S)$. The two marked edge points have margins $1/8$ and $1/16$. All mixtures of these points with $\mathbf{e}_2$ and $\mathbf{e}_3$ form $\widehat{A}_i$.
All its points are accepted, and it contains the entire darkest region $A_i^{+\rho}$.

Figure 2: Learning an inner approximation does not require learning the true boundary. The thin solid line is $g_i(\mathbf{x}) = 0$, the thick line bounds $\widehat{A}_i$, and the dashed line is $g_i(\mathbf{x}) = \rho$. The accepted edge points may have different margins; the inclusions $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i$ still hold.

The test below makes the distinction we need with high probability. On its correctness event, test output 1 implies acceptance, and test output 0 implies margin below a chosen value $\gamma$. A lottery with $0 \leq g_i(\mathbf{x}) < \gamma$ may give either test output. This test can therefore finish even at a boundary point, where the exact label may be impossible to learn by repetition.

**Lemma 4.1.** *Under* NearThresholdNoise$(\lambda, \beta)$ *with known $\lambda, \beta$, for any $i \in N$, $\mathbf{x} \in \Delta(S)$, $\gamma \in (0, 1]$, and $\delta_0 \in (0, 1)$, $O(\lambda^{-2}\gamma^{-2\beta}\log(2/\delta_0))$ repeated answers at $(i, \mathbf{x})$ give a bit $a$ such that, with probability at least $1 - \delta_0$, $a = 1 \Longrightarrow g_i(\mathbf{x}) \geq 0$ and $a = 0 \Longrightarrow g_i(\mathbf{x}) < \gamma$.*

*The guarantee also holds conditionally on the history before the test.*

For each agent, test every pure lottery with $\gamma = \rho/2$. For each pair of pure lotteries with different test outputs, repeatedly halve the interval of mixtures between them to find an accepted point with margin below $\rho$. Let $\widehat{A}_i$ consist of all mixtures of the points so obtained and the pure lotteries whose tests returned 1; this is their *convex hull*. These points suffice because every corner of $A_i^{+\rho}$ is either a pure lottery or a mixture of two pure lotteries (Motzkin et al., 1953). Finally, find a point in $\bigcap_i \widehat{A}_i$, or return NULL if it is empty. Algorithm 7 in Appendix F gives the details.

**Theorem 4.2.** *For every $\rho \in (0, 1]$ and $\delta \in (0, 1)$, there is a $(1 - \delta)$-correct algorithm for $\rho$-robust consensus under near-threshold noise with known parameters $(\lambda, \beta)$, without finite precision, using $O_\beta(nm^2\lambda^{-2}\rho^{-2\beta}\log(2/\rho)\log\frac{4nm^2\log(2/\rho)}{\delta})$ queries. With probability at least $1 - \delta$, every output lottery is accepted by every agent on every instance, including when neither required condition holds.*

A linear objective can also be maximized over $\bigcap_i \widehat{A}_i$ at no additional query cost. In a robust YES instance, the same correctness event guarantees a value at least that of any lottery with common margin $\rho$ (Appendix F.5).

The returned lottery need not have margin comparable to $\rho$. For any fixed $c > 0$ independent of $m$, no algorithm can guarantee margin $c\rho$ on every robust YES instance. The noise inequalities need not reveal which accepted lotteries have larger margins, even after arbitrarily many queries (Appendix F.6). A known response rule can remove this limitation (Appendix I).

The robust algorithm's repetition cost $\lambda^{-2}\rho^{-2\beta}$ comes from distinguishing negative margins from margins at least $\rho/2$. Queried lotteries may have margins between 0 and $\rho/2$, where the test allows either output. The next result shows that this dependence on $\rho$, together with the linear dependence on $n$, is necessary even with two alternatives.

**Theorem 4.3.** *Fix $n \geq 2$, $m = 2$, $\rho \in (0, 1/10]$, and $\delta \in (0, 1/4)$. Under* NearThresholdNoise$(\lambda, \beta)$*, every $(1 - \delta)$-correct $\rho$-robust consensus algorithm has worst-case expected query complexity $\Omega_\beta((n - 1)\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$.*

A separate construction gives an $\Omega_\beta(m\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$ lower bound with one agent (Theorem G.3). Together, the two constructions give the bound for all $n, m \geq 2$ in Table 1; Appendices F.7 and G.4 give the proofs.

Our consensus upper bound is quadratic in $m$, whereas the lower bound is linear. Whether consensus itself requires the quadratic cost remains open. We can establish that cost for the stronger task performed by our algorithm: learning all the inner approximations. The next theorem allows any description of the learned sets.

**Theorem 4.4.** *Let $n \geq 1$, $m = 2r \geq 4$, $\rho \in (0, 1/16]$, and $\delta \in (0, 1/4)$. Suppose a randomized procedure returns sets $\widehat{A}_1, \ldots, \widehat{A}_n$ such that, on every instance under* NearThresholdNoise$(\lambda, \beta)$*, $\Pr(\forall i \in N : A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i) \geq 1 - \delta$. The sets may have any description, provided that membership is a measurable function of the completed queries and responses and requires no further queries. Then the worst-case expected query complexity is $\Omega_\beta\left(nm^2\lambda^{-2}\rho^{-2\beta}\log(1/\delta)\right)$. This holds even with independent responses, known thresholds $\tau_i = 1/2$, and a supplied lottery having margin at least $1/4$ for every agent.*

The bound matches our set learning algorithm up to logarithmic factors; Appendix G.5 gives the proof. It quantifies a distinction central to preference elicitation (Blum et al., 2004): knowing one good choice does not make every agent's acceptable choices cheap to learn. For consensus, however, that one choice may be all we need. This suggests checking a promising lottery before learning any sets.

## 5 Robust Consensus with a Predicted Lottery

A previous decision or preliminary evaluation may suggest a promising lottery $\mathbf{x}^{(0)}$. Its quality is unknown, so trusting it without checking would sacrifice the acceptance guarantee. *Learning-augmented algorithms*, also called algorithms with predictions, address this tension: useful predictions improve performance, while arbitrary predictions retain a worst-case guarantee (Lykouris & Vassilvitskii, 2021; van den Brand et al., 2024).

Choo et al. (2026) apply this approach to noiseless approval, where one query per agent checks a proposed lottery. With noisy answers, acceptance with zero margin is not enough to make such a check reliable. Here a useful prediction has $g_i(\mathbf{x}^{(0)}) \geq \rho$ for every agent, as the equal mixture in Figure 1(a) does for $\rho \leq 0.1$. We make no assumption about how the prediction was obtained.

The first step is to check one lottery rather than learn entire sets. Applying Lemma 4.1 with $\gamma = \rho$ gives the following guarantee, whose query cost is independent of $m$.

**Proposition 5.1.** *Under* NearThresholdNoise$(\lambda, \beta)$ *with known* $\lambda, \beta$*, given a lottery* $\mathbf{x}$*, there is a test using* $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ *queries such that, with probability at least* $1-\delta$*, it returns* YES *when* $\min_{i\in N} g_i(\mathbf{x}) \geq \rho$ *and returns* NO *when* $\min_{i\in N} g_i(\mathbf{x}) < 0$*. On the same event, every* YES *output implies* $\mathbf{x} \in \bigcap_{i\in N} A_i$*, including when neither condition holds.*

Appendix F.8 gives the proof and the matching lower bound in the main statistical parameters.

When an agent's test returns 0, acceptance remains unresolved: the margin is below $\rho$, but need not be negative. We learn that agent's inner approximation and keep the current lottery if it lies in all sets learned so far; otherwise we choose a new point in their intersection. An empty intersection leads to NULL. We repeat the checks for the remaining agents until they all return 1. Each restart learns one new set, so there are at most $n$ restarts. Algorithm 8 in Appendix G.6 gives the details; its cost depends on the number $K$ of sets learned.

**Theorem 5.2.** *For every* $\rho \in (0, 1]$*,* $\delta \in (0, 1)$*, and initial lottery* $\mathbf{x}^{(0)}$*, there is an adaptive algorithm that is* $(1-\delta)$*-correct for* $\rho$*-robust consensus under near-threshold noise with known parameters* $(\lambda, \beta)$*. With the same probability, every lottery it outputs lies in* $\bigcap_{i\in N} A_i$*, including when neither required condition holds. On every execution that constructs inner approximations for* $K$ *agents, its query count is*

$$O_\beta\left(\lambda^{-2}\rho^{-2\beta}\left[n(K+1) + Km^2\log(2/\rho)\right]\log\frac{8n^2m^2\log(2/\rho)}{\delta}\right).$$

*In particular, if the initial lottery has margin at least* $\rho$ *for every agent, then, with probability at least* $1-\delta$*,* $K = 0$ *and the cost is* $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$*.*

**Worst-case guarantee for any prediction.** The adaptive procedure may repeat checks after changing its lottery, introducing the $n(K+1)$ term. To keep these checks from exceeding the general search budget, alternate one noisy query from Algorithm 7 and one from Algorithm 8, each with error allowance $\delta/2$, and return the first completed output (van den Brand et al., 2024). The combined procedure is $(1-\delta)$-correct for any $\mathbf{x}^{(0)}$ and uses at most twice the worst-case query budget of Algorithm 7 with error allowance $\delta/2$. If $g_i(\mathbf{x}^{(0)}) \geq \rho$ for every agent, it still stops within $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ queries with probability at least $1-\delta$. Appendix G.7 proves this for history-dependent responses as well.

For $\rho \leq 1/10$ and $\delta < 1/4$, even an algorithm allowed to return a different lottery needs $\Omega_\beta(n\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$ expected queries on some instances where the prediction has common margin $\rho$ (Appendix G.7). Thus the fast rate is optimal in these parameters, up to logarithmic factors, for algorithms correct under arbitrary predictions. The fast-rate guarantee uses the prediction's own margins; closeness to an acceptable lottery alone is not enough.

## 6 CONCLUSION

Noisy approval changes not only how often to ask a question, but which questions can resolve agreement. Repetition alone may be inconclusive at the only common lottery. Finite precision permits exact decisions by allowing queries between possible turning points; a margin gap instead permits boundary uncertainty while preserving unanimous acceptance of every returned lottery. The query cost also depends on the required output. Learning every inner approximation requires quadratically many queries in the number of alternatives, even when a common acceptable lottery is given. Checking a proposed lottery with sufficient margin has no such dependence; our prediction-based algorithm exploits this without losing the general worst-case guarantee. Unknown noise parameters add only logarithmic costs whenever the instance is infeasible or admits a common positive margin; a known response rule can reveal numerical margins and reduce the cost of learning.

Several questions remain open. The exact and robust consensus bounds do not match jointly in $n$ and $m$; the quadratic lower bound concerns learning the sets. Our lower bounds for near-threshold noise use worst-case response rules allowed by Definition 2.4; they need not hold for a fixed rule. It would be useful to study an unknown monotone response function with both upper and lower bounds on its bias. Another direction is to allow nonlinear acceptance rules or feedback whose accuracy changes through deliberation. The current results require linear acceptance and the stated conditional response bounds after every query history.

### AI USE STATEMENT

We used generative AI tools to revise the prose and presentation of technical proofs, to assist in checking mathematical arguments and proposing corrections, and to help identify relevant literature. We did not use these tools to formulate the model or to generate or analyze experimental data. The authors are responsible for independently verifying the mathematical arguments, references, and all AI-assisted revisions, and take full responsibility for the final content.

## REFERENCES


Alekh Agarwal, Alina Beygelzimer, Miroslav Dudík, John Langford, and Hanna Wallach. A reductions approach to fair classification. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, pp. 60–69, 2018.

Eitan Altman. *Constrained Markov Decision Processes*. Chapman and Hall, 1999.

Dana Angluin. Queries and concept learning. *Machine Learning*, 2:319–342, 1988.

Dana Angluin and Philip Laird. Learning from noisy examples. *Machine Learning*, 2:343–370, 1988.

Francis J. Anscombe and Robert J. Aumann. A definition of subjective probability. *The Annals of Mathematical Statistics*, 34(1):199–205, 1963.

Kazuoki Azuma. Weighted sums of certain dependent random variables. *Tohoku Mathematical Journal, Second Series*, 19(3):357–367, 1967.

Maria-Florina Balcan, Alina Beygelzimer, and John Langford. Agnostic active learning. *Journal of Computer and System Sciences*, 75(1):78–89, 2009.

Xiaohui Bei, Ning Chen, and Shengyu Zhang. Solving linear programming with constraints unknown. In *Proceedings of the 42nd International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 129–142, 2015.

Avrim Blum, Prasad Chalasani, Sally A. Goldman, and Donna K. Slonim. Learning with unreliable boundary queries. *Journal of Computer and System Sciences*, 56(2):209–222, 1998.

Avrim Blum, Jeffrey Jackson, Tuomas Sandholm, and Martin Zinkevich. Preference elicitation and query learning. *Journal of Machine Learning Research*, 5:649–667, 2004.

Craig Boutilier and Jeffrey S. Rosenschein. Incomplete information and communication in voting. In Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (eds.), *Handbook of Computational Social Choice*, pp. 223–258. Cambridge University Press, 2016.

Ralph Allan Bradley and Milton E. Terry. Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika*, 39(3/4):324–345, 1952.

Steven J. Brams and Peter C. Fishburn. Approval voting. *American Political Science Review*, 72(3): 831–847, 1978.

Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (eds.). *Handbook of Computational Social Choice*. Cambridge University Press, 2016.

M. V. Burnashev and K. Sh. Zigangirov. An interval estimation problem for controlled observations. *Problemy Peredachi Informatsii*, 10:51–61, 1974.

Rui M. Castro and Robert D. Nowak. Minimax bounds for active learning. *IEEE Transactions on Information Theory*, 54(5):2339–2353, 2008.

Lin Chen, Hamed Hassani, and Amin Karbasi. Near-optimal active learning of halfspaces via query synthesis in the noisy setting. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 1798–1804, 2017.

Davin Choo, Paul W. Goldberg, and Nicholas Teh. Learning unanimously acceptable lotteries via queries. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, 2026. Extended version available as arXiv:2604.17505.

Vincent Conitzer and Tuomas Sandholm. Vote elicitation: Complexity and strategy-proofness. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI)*, pp. 392–397, 2002.

Andrew Cotter, Maya Gupta, Heinrich Jiang, Nathan Srebro, Karthik Sridharan, Serena Wang, Blake Woodworth, and Seungil You. Training well-generalizing classifiers for fairness metrics and other data-dependent constraints. In *Proceedings of the 36th International Conference on Machine Learning (ICML)*, pp. 1397–1405, 2019.

Trung Dang, Daniel Halpern, Anuran Makur, Alexandros Psomas, Japneet Singh, and Paritosh Verma. Fair division under inaccurate preferences. *arXiv preprint arXiv:2602.24169*, 2026.

Alexander P. Dawid and Allan M. Skene. Maximum likelihood estimation of observer error-rates using the EM algorithm. *Journal of the Royal Statistical Society: Series C*, 28(1):20–28, 1979.

Ilias Diakonikolas, Vasilis Kontonis, Christos Tzamos, and Nikos Zarifis. Learning halfspaces with massart noise under structured distributions. In *Proceedings of 33rd Conference on Learning Theory (COLT)*, pp. 1486–1513, 2020.

Ilias Diakonikolas, Daniel M. Kane, Vasilis Kontonis, Christos Tzamos, and Nikos Zarifis. Efficiently learning halfspaces with tsybakov noise. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pp. 88–101, 2021.

Ilias Diakonikolas, Daniel Kane, and Mingchen Ma. Active learning of general halfspaces: Label queries vs membership queries. In *Proceedings of the 38th International Conference on Neural Information Processing Systems (NeurIPS)*, pp. 49180–49213, 2024.

Uriel Feige, Prabhakar Raghavan, David Peleg, and Eli Upfal. Computing with noisy information. *SIAM Journal on Computing*, 23(5):1001–1018, 1994.

Peter C. Fishburn. *Utility Theory for Decision Making*. Wiley, 1970.

Peter C. Fishburn. Lotteries and social choices. *Journal of Economic Theory*, 5(2):189–207, 1972.

Peter C. Fishburn. Acceptable social choice lotteries. In Hans W. Gottinger and Werner Leinfellner (eds.), *Decision Theory and Social Ethics*, volume 17 of *Theory and Decision Library*, pp. 133–152. Springer, 1978.

Peter I. Frazier, Shane G. Henderson, and Rolf Waeber. Probabilistic bisection converges almost as quickly as stochastic approximation. *Mathematics of Operations Research*, 44(2):651–667, 2019.

Aditya Gangrade, Aditya Gopalan, Venkatesh Saligrama, and Clayton Scott. Testing the feasibility of linear programs with bandit feedback. In *Proceedings of the 41st International Conference on Machine Learning (ICML)*, pp. 14513–14539, 2024.

Javier García and Fernando Fernández. A comprehensive survey on safe reinforcement learning. *Journal of Machine Learning Research*, 16:1437–1480, 2015.

Allan Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45(3):665–681, 1977.

Paul W. Goldberg and Stephen Kwek. The precision of query points as a resource for learning convex polytopes with membership queries. In *Proceedings of the 13th Annual Conference on Computational Learning Theory (COLT)*, pp. 225–235, 2000.

Navin Goyal and Luis Rademacher. Learning convex bodies is hard. In *Proceedings of the 22nd Annual Conference on Learning Theory (COLT)*, 2009.

Lucas Gretta and Eric Price. Sharp noisy binary search with monotonic probabilities. In *Proceedings of the 51st International Colloquium on Automata, Languages, and Programming (ICALP)*, pp. 75:1–75:19, 2024.

Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*. Springer, 1993.

Daniel Halpern, Safwan Hossain, and Jamie Tucker-Foltz. Computing voting rules with elicited incomplete votes. In *Proceedings of the 25th ACM Conference on Economics and Computation (EC)*, pp. 941–963, 2024.

Steve Hanneke. A bound on the label complexity of agnostic active learning. In *Proceedings of the 24th International Conference on Machine Learning*, pp. 353–360, 2007.

Steve Hanneke and Liu Yang. Minimax analysis of active learning. *Journal of Machine Learning Research*, 16:3487–3602, 2015.

Sariel Har-Peled, Mitchell Jones, and Saladi Rahul. Active-learning a convex body in low dimensions. *Algorithmica*, 83:1885–1917, 2021.

Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. In *Proceedings of the 30th International Conference on Neural Information Processing Systems (NeurIPS)*, pp. 3323–3331, 2016.

Israel N. Herstein and John Milnor. An axiomatic approach to measurable utility. *Econometrica*, 21 (2):291–297, 1953.

Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 58(301):13–30, 1963.

Max Hopkins, Daniel Kane, Shachar Lovett, and Gaurav Mahajan. Noise-tolerant, reliable active classification with comparison queries. In *Proceedings of the 33rd Conference on Learning Theory (COLT)*, pp. 1957–2006, 2020.

Michael Horstein. Sequential transmission using noiseless feedback. *IEEE Transactions on Information Theory*, 9(3):136–143, 1963.

Richard M. Karp and Robert Kleinberg. Noisy binary search and its applications. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 881–890, 2007.

Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best-arm identification in multi-armed bandit models. *Journal of Machine Learning Research*, 17(1):1–42, 2016.

Ori Kelner. Learning halfspaces with membership queries. *arXiv preprint arXiv:2012.10985*, 2020.

Jean-François Laslier and M. Remzi Sanver (eds.). *Handbook on Approval Voting*. Springer, 2010.

Yin Tat Lee, Aaron Sidford, and Santosh S. Vempala. Efficient convex optimization with membership oracles. In *Proceedings of the 31st Conference On Learning Theory (COLT)*, pp. 1292–1294, 2018.

Yinan Li and Chicheng Zhang. Efficient active learning halfspaces with tsybakov noise: A non-convex optimization approach. In *Proceedings of The 27th International Conference on Artificial Intelligence and Statistics (AISTATS)*, pp. 4744–4752, 2024.

Zihan Li, Pasin Manurangsi, Jonathan Scarlett, and Warut Suksompong. Complexity of round-robin allocation with potentially noisy queries. *Information and Computation*, 306:105332, 2025.

Zihan Li, Yan Hao Ling, Jonathan Scarlett, and Warut Suksompong. Envy-free allocation of indivisible goods via noisy queries. In *Proceedings of the 43rd International Conference on Machine Learning (ICML)*, 2026. Extended version available as arXiv:2602.06361.

Andrea Locatelli, Alexandra Carpentier, and Samory Kpotufe. An adaptive strategy for active learning with smooth decision boundary. In *Proceedings of Algorithmic Learning Theory (ALT)*, pp. 547–571, 2018.

Tyler Lu and Craig Boutilier. Preference elicitation and robust winner determination for single- and multi-winner social choice. *Artificial Intelligence*, 279:103203, 2020.

R. Duncan Luce. *Individual Choice Behavior: A Theoretical Analysis*. Wiley, 1959.

Thodoris Lykouris and Sergei Vassilvitskii. Competitive Caching with Machine Learned Advice. *Journal of the ACM*, 68(4):1–25, 2021.

Wolfgang Maass and György Turán. Lower bound methods and separation results for on-line learning models. *Machine Learning*, 9(2–3):107–145, 1992. doi: 10.1023/A:1022637031594.

Enno Mammen and Alexandre B. Tsybakov. Smooth discrimination analysis. *The Annals of Statistics*, 27(6):1808–1829, 1999.

Pascal Massart and Élodie Nédélec. Risk bounds for statistical learning. *The Annals of Statistics*, 34(5):2326–2366, 2006.

Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *Proceedings of the 2019 Conference on Fairness, Accountability, and Transparency (FAT*)*, pp. 220–229, 2019.

Theodore S. Motzkin, Howard Raiffa, Gerald L. Thompson, and Robert M. Thrall. The double description method. In Harold W. Kuhn and Albert W. Tucker (eds.), *Contributions to the Theory of Games*, volume 2, pp. 51–74. Princeton University Press, 1953.

Harikrishna Narasimhan, Harish Ramaswamy, Aadirupa Saha, and Shivani Agarwal. Consistent multiclass algorithms for complex performance measures. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 2398–2407, 2015.

Harikrishna Narasimhan, Harish G. Ramaswamy, Shiv Kumar Tavker, Drona Khurana, Praneeth Netrapalli, and Shivani Agarwal. Consistent multiclass algorithms for complex metrics and constraints. *Journal of Machine Learning Research*, 25(367):1–81, 2024.

Robert D. Nowak. Noisy generalized binary search. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems (NeurIPS)*, pp. 1366–1374, 2009.

Robert D. Nowak. The geometry of generalized binary search. *IEEE Transactions on Information Theory*, 57(12):7893–7906, 2011.

Dominik Peters, Ariel D. Procaccia, and David Zhu. Robust rent division. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*, pp. 13864–13876, 2022.

Inioluwa Deborah Raji, Andrew Smart, Rebecca N. White, Margaret Mitchell, Timnit Gebru, Ben Hutchinson, Jamila Smith-Loud, Daniel Theron, and Parker Barnes. Closing the AI accountability gap: Defining an end-to-end framework for internal algorithmic auditing. In *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency (FAT*)*, pp. 33–44, 2020.

Jens Schreiter, Duy Nguyen-Tuong, Mona Eberts, Bastian Bischoff, Heiner Markert, and Marc Toussaint. Safe exploration for active learning with gaussian processes. In *Machine Learning and Knowledge Discovery in Databases*, pp. 133–149. Springer, 2015.

Clayton Scott and Robert Nowak. A neyman-pearson approach to statistical learning. *IEEE Transactions on Information Theory*, 51:3806–3819, 2005.

Rion Snow, Brendan O'Connor, Daniel Jurafsky, and Andrew Y. Ng. Cheap and fast—but is it good? evaluating non-expert annotations for natural language tasks. In *Proceedings of the 2008 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 254–263, 2008.

Yanan Sui, Alkis Gotovos, Joel W. Burdick, and Andreas Krause. Safe exploration for optimization with gaussian processes. In *Proceedings of the 32nd International Conference on Machine Learning (ICML)*, pp. 997–1005, 2015.

Louis L. Thurstone. A law of comparative judgment. *Psychological Review*, 34(4):273–286, 1927.

Kenneth E. Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2 edition, 2009.

Alexandre B. Tsybakov. Optimal aggregation of classifiers in statistical learning. *The Annals of Statistics*, 32(1):135–166, 2004.

Jan van den Brand, Sebastian Forster, Yasamin Nazari, and Adam Polak. On dynamic graph algorithms with predictions. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pp. 3534–3557, 2024.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Rolf Waeber, Peter I. Frazier, and Shane G. Henderson. Bisection search with noisy responses. *SIAM Journal on Control and Optimization*, 51(3):2261–2279, 2013.

Yining Wang and Aarti Singh. Noise-adaptive margin-based active learning and lower bounds under tsybakov noise condition. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 2180–2186, 2016.

Richard Zeckhauser. Majority rule with lotteries on alternatives. *The Quarterly Journal of Economics*, 83(4):696–703, 1969.

# A Additional Related Work

**Randomized social choice, approval, and elicitation.** Randomized social choice studies lotteries as social outcomes, including majority rule over lotteries and axiomatic treatments of acceptable lotteries (Fishburn, 1972; 1978; Zeckhauser, 1969). Gibbard (1977) studies strategic manipulation when voting mechanisms use randomization, and Brandt et al. (2016) surveys computational social choice more broadly. Approval voting (Brams & Fishburn, 1978; Laslier & Sanver, 2010) is related because agents give binary acceptability information. Its usual input, however, is approval of individual alternatives. In our model, the algorithm asks about mixtures, and each answer concerns whether an expected utility reaches an unknown threshold.

Work on vote elicitation, incomplete preferences, and robust winner determination asks how much information is needed to compute a winner or a set of possible winners under a voting rule (Boutilier & Rosenschein, 2016; Conitzer & Sandholm, 2002; Halpern et al., 2024; Lu & Boutilier, 2020). Our query type and objective differ: we propose lotteries, receive noisy accept/reject answers, and seek a point common to all acceptance sets rather than a winner under an aggregation rule.

**Fair division with noisy preferences.** Li et al. (2025) study round-robin allocation with noisy item comparisons and value queries. Li et al. (2026) use Gaussian observations of item values and a positive gap in the best achievable envy to find an envy-free allocation. Uncertain numerical valuations are also studied in rent division (Peters et al., 2022) and fair division under inaccurate preferences (Dang et al., 2026). These works concern allocations of items. Our questions instead concern approval of lotteries at unknown thresholds, with response accuracy that may approach random guessing at equality. We study the resulting costs of exact agreement and of retaining every lottery with sufficient margin.

Blum et al. (2004) connect preference elicitation to query learning and distinguish finding a desired outcome from learning each agent's preferences. Our distinction between consensus and learning acceptance sets follows this perspective. Theorem 4.4 gives a quantitative bound for the latter task under noisy approval, even when a common acceptable lottery is supplied.

**Membership query and geometric concept learning.** For one agent, each query asks whether a lottery belongs to an unknown halfspace. Such a question is called a *membership query*. Membership queries, and queries that ask whether a proposed set equals the unknown set, are classical learning models (Angluin, 1988), and halfspaces have been studied with exact and noisy query access (Chen et al., 2017; Kelner, 2020). A useful distinction is whether the learner may choose its query points. In active learning, the algorithm often chooses which points to label from points drawn from a distribution. Membership queries instead allow the algorithm to choose arbitrary points. These forms of access can have different computational and statistical costs for halfspaces (Diakonikolas et al., 2024). Our model allows arbitrary lottery queries, but requires an exact or robust feasibility decision for the agents' common intersection.

The unreliable boundary model of Blum et al. (1998) is particularly close to our setting: queries near a concept's boundary may be corrupted. Our model specifies a lower bound proportional to $|g_i(\mathbf{x})|^\beta$ on the response bias. This gives quantitative margin-dependent bounds, while still allowing an unbiased response at a boundary. Exact consensus is sensitive to such points because they may be the only common lotteries.

Goldberg & Kwek (2000) treat the precision of query points as a resource for learning convex polytopes, which are bounded sets described by finitely many linear inequalities. More broadly, learning and optimization with membership oracles have been studied in learning theory and convex optimization (Goyal & Rademacher, 2009; Har-Peled et al., 2021; Lee et al., 2018; Grötschel et al., 1993). Here finite precision gives a lower bound on the absolute margin at the midpoints used by our exact searches.

Dimension-dependent lower bounds for exact halfspace learning are also classical. Maass & Turán (1992) give an almost quadratic bound on a finite grid with noiseless membership queries and queries asking whether the proposed set equals the unknown set. Our inner-set lower bound allows arbitrary lottery queries and depends on noise, margin, and confidence; it also holds when a common acceptable lottery is supplied.

**Noise conditions in statistical learning theory.** Uniform noise is a conditional version of bounded, or Massart, noise (Massart & Nédélec, 2006; Diakonikolas et al., 2020): the error probability may vary between queries but is uniformly below $1/2$. Independent flips with a fixed probability, as in Angluin & Laird (1988), are a special case. Our version also permits history-dependent responses. Repetition reduces the error probability by the classical concentration bounds of Hoeffding (1963); Azuma (1967). Near-threshold noise is related to low-noise and margin-noise conditions, including Tsybakov's condition (Mammen & Tsybakov, 1999; Tsybakov, 2004), and halfspace learning under these conditions (Diakonikolas et al., 2021; Li & Zhang, 2024). Those problems usually measure classification error or excess risk under a distribution. We instead decide feasibility of a particular collection of unknown constraints.

Other active learning methods study the number of labels needed under general noise, including methods that query where candidate classifiers disagree (Balcan et al., 2009; Hanneke, 2007; Hanneke & Yang, 2015). Margin-based active learning examines the cost of querying near a decision boundary under Tsybakov-type noise (Castro & Nowak, 2008; Wang & Singh, 2016). The difference in objectives matters here: an error on a boundary of probability zero may not affect classification risk, but can change whether the agents have any common lottery.

Locatelli et al. (2018) go beyond average classification error: they estimate a decision boundary uniformly and learn regions whose labels are correct, leaving points near the boundary undecided. Their pointwise lower bound on response bias is closely related to Definition 2.4. Hopkins et al. (2020) also learn classifiers that may leave points unlabelled and, with high probability, make no errors, using additional queries that compare two points. Our guarantees concern utility margins and the common intersection of several acceptance sets on the simplex. Exact consensus must also resolve cases where equality is the only way to agree. The power $\rho^{-2\beta}$ comes from the familiar cost of distinguishing small response biases; our results establish its consequences for unanimous acceptance and the additional cost of learning all the inner sets.

**Searching with noisy comparisons.** Our edge searches are related to noisy binary search and probabilistic bisection: methods that locate a threshold using comparisons that may be incorrect (Burnashev & Zigangirov, 1974; Horstein, 1963; Karp & Kleinberg, 2007; Waeber et al., 2013). Noisy decision-tree computation (Feige et al., 1994) and generalized binary search (Nowak, 2009; 2011) also combine adaptive queries with repeated observations. In our exact problem, however, a boundary query may be unbiased. The search must therefore use points known to differ from every possible turning point.

Frazier et al. (2019) use repeated sign tests when answer accuracy depends on the query point and discuss why such a test need not terminate at the root. Karp & Kleinberg (2007); Gretta & Price (2024) study unknown monotone response probabilities and allow uncertainty around a target probability. Our response probabilities need not be monotone along an edge, although the true acceptance labels are; our accuracy requirements are stated in terms of margins rather than response probabilities.

**Finding decisions that satisfy constraints.** Convex optimization also studies decision problems that leave points close to a boundary unresolved, including weak membership and robust feasibility. Our robust objective separates an instance with a common positive margin from one that remains infeasible after all thresholds are lowered. The available information differs. Classical convex-optimization methods use reliable answers about whether a point is feasible, or a linear inequality explaining why it is infeasible (Grötschel et al., 1993; Lee et al., 2018), whereas our responses may become arbitrarily unreliable near a boundary. Bei et al. (2015) study linear feasibility when a query identifies a violated constraint without revealing its coefficients. Their answers are noiseless, and the algorithm cannot select which constraint is reported. We can query a chosen agent, but her answer may be uninformative near her threshold. Gangrade et al. (2024) test feasibility of unknown linear constraints from noisy numerical scores. They express feasibility through the sign of the largest common margin and adapt their stopping time to its unknown magnitude. Their model already permits history-dependent noise. Our feedback is a binary answer from one chosen agent, with only bounds relating its probability to the margin; numerical constraint values are not observed. Our results establish the query costs for exact consensus and for learning inner sets under this binary feedback model.

Safe Bayesian optimization and safe active learning seek to keep queried points acceptable while learning or optimizing, often using a probabilistic model of how nearby points are related (Schreiter et al., 2015; Sui et al., 2015). We do not require queried lotteries to be accepted; only the returned lottery must be unanimously acceptable. Our guarantees are worst-case rather than Bayesian.

**Stochastic choice and response models.** Thurstone, Bradley–Terry, Luce, and random utility models describe stochastic choices (Bradley & Terry, 1952; Luce, 1959; Thurstone, 1927; Train, 2009). They motivate less predictable responses near indifference. Our near-threshold condition captures this behavior without fixing a response function such as logit or probit. The general lower bounds range over all rules allowed by Definition 2.4. A known response function can reveal numerical margins as well as acceptance, leading to the faster rates in Appendix I.

## B When Is Linear Acceptance Appropriate?

Our acceptance condition $\mathbf{x} \in A_i \iff \langle u_i, \mathbf{x} \rangle \geq \tau_i$ applies when an agent requires an expected score to reach a minimum level. We first derive the condition from that requirement, then discuss its interpretation for machine learning decisions and its limitations.

**Expected metric approval.** The alternatives $s_1, \ldots, s_m$ may be complete model versions, rollout policies, monitoring plans, or other fully specified choices. Suppose agent $i$ evaluates $s_j$ using a bounded metric $Z_i(s_j) \in [0, 1]$, such as a pass indicator, a validation score, or one minus a normalized harm or cost. A lottery first draws $J \sim \mathbf{x}$ and then uses the selected alternative's outcome distribution. Taking expectation over both the draw and the outcome gives us

$$\mathbb{E}[Z_i(s_J)] = \sum_{j=1}^{m} x_j \mathbb{E}[Z_i(s_j)].$$

With $u_{ij} = \mathbb{E}[Z_i(s_j)]$, the requirement $\mathbb{E}[Z_i(s_J)] \geq \tau_i$ is exactly our acceptance condition. Linearity here follows directly from averaging over the selected alternative.

Bounded scores and a nontrivial threshold can be rescaled together to the ranges in Section 2. A requirement met by every lottery is described by AcceptAll, and one met by none is described by RejectAll. An upper bound on an expected cost can likewise be written as a lower bound on a rescaled benefit. Rescaling preserves acceptance, but changes numerical margins, so noise and margin parameters are interpreted after normalization. A two-sided tolerance on a linear rate difference can be written as two inequalities. Multiple requirements from one stakeholder can therefore be represented as multiple agents; unanimity means satisfying them all.

**Acceptance relative to a minimum utility.** There is also a classical decision-theoretic interpretation. Classical expected utility theory gives conditions under which a person evaluates a lottery by averaging utilities of its alternatives. These include consistent comparisons, continuity, and an independence condition: mixing two options with the same third option preserves their ranking. Acceptance relative to a status quo or a minimum level then has the form $\langle u_i, \mathbf{x} \rangle \geq \tau_i$ after rescaling. This interpretation follows from the von Neumann–Morgenstern and mixture-space representation results; related expected utility results also allow uncertainty about states of the world (Anscombe & Aumann, 1963; Fishburn, 1970; Herstein & Milnor, 1953; von Neumann & Morgenstern, 1944). Our results do not require this broader preference interpretation. They only require the particular queried acceptance rule to be a threshold on an expected metric.

**Why this model is natural in ML.** A system can randomize among finitely many classifiers or deployment policies. Error, accuracy, recall, coverage, false-positive rates, and group-conditional true-positive rates are expectations of losses or indicators. When the policy first selects a model according to the lottery, these rates are linear in the mixing probabilities.

This structure is used in fair and constrained classification. Equalized odds and equal opportunity constrain group-conditional false-positive or true-positive rates (Hardt et al., 2016). The reductions approach of Agarwal et al. (2018) optimizes over distributions of classifiers subject to linear constraints on group rates. More general constrained learning methods use fairness, recall, coverage, and churn constraints and may return stochastic classifiers or mixtures of deterministic models (Cotter et al., 2019). Work on complex metrics and constraints de-

scribes classifiers through their confusion matrices; loss-based objectives are linear in these matrices, and some ratio thresholds become linear inequalities when the denominator is positive (Narasimhan et al., 2015; 2024). Neyman–Pearson classification controls one expected error rate while optimizing another (Scott & Nowak, 2005). Constrained Markov decision processes and safe reinforcement learning similarly use expected returns and expected cost constraints (Altman, 1999; García & Fernández, 2015).

For a concrete safety requirement, suppose harmful responses must occur on at most $2\%$ of requests from a fixed population. Let $u_{ij}$ be the probability that alternative $s_j$ does not produce a harmful response and let $\tau_i = 0.98$. Then $\langle u_i, \mathbf{x} \rangle \geq \tau_i$ is exactly the required bound when the policy is selected independently of the request. Separate agents can impose the same kind of bound within particular groups, so an overall average does not replace their requirements.

A query presents the lottery $\mathbf{x}$ to the evaluator for requirement $i$. The evaluator may assess a fresh sample of requests and return only whether the proposed mixture meets that requirement. Different samples or human assessments may yield different answers, especially when the expected score is close to the threshold. Our noise assumptions describe how reliably these answers reflect the underlying expected score; finite-sample evaluation alone does not imply the stated probability bounds. Thus the application requires both the linear acceptance rule and the conditional response bounds in Section 2.1.

**Scope and limitations.** The model does not cover arbitrary preferences over the distribution of outcomes. Variance penalties, nonlinear interactions, and many tail-risk or worst-case requirements need not be linear in the lottery. Some additional requirements can be expressed through separate expected metrics, such as expected squared loss, subgroup rates, costs, or benefits. Others require nonlinear or nonconvex acceptance sets.

The query guarantees rely on this restriction: each agent's acceptable lotteries satisfy a linear inequality within the simplex. They need not extend to more general preferences over randomized outcomes.

# C Statistical and Geometric Facts Used in the Proofs

The algorithms repeatedly ask the same question to reduce the chance of an incorrect answer. The lower bounds take the opposite view: if two instances give very similar answers, distinguishing them requires many queries. We give the probability tools for both arguments, then show how a small number of edge turning points determine an agent's acceptance set. Throughout, we use the margins $g_i$, edge lotteries $\mathbf{x}_{k,k'}(\alpha)$, and finite candidate set $R_Q$ defined in Section 2.

## C.1 Statistical tools

The repeated-answer bounds below apply the classical concentration argument of Hoeffding (1963); Azuma (1967) conditionally after each answer. Repeated answers need not be independent. We require only that the probability of a 1 response, conditional on the answers seen so far, stays on the stated side of $1/2$. In part (i), the bias toward the true label is at least $b$ under either label, so a majority vote suffices. In part (ii), label 1 may have mean exactly $1/2$; moving the decision threshold below $1/2$ makes it possible to identify this label as well. This second case will let us classify accepted pure lotteries that lie on the boundary.

**Lemma C.1.** *Let $T \geq 1$ be an integer and $b > 0$. Let $\mathcal{F}_0 \subseteq \cdots \subseteq \mathcal{F}_T$ describe the information available through each response. The response $Y_t \in \{0,1\}$ is $\mathcal{F}_t$-measurable, meaning that it is known after round $t$. For each label $\ell \in \{0,1\}$, write $\Pr_\ell$ and $\mathbb{E}_\ell$ for probability and expectation when $\ell$ is the true label. Set $\overline{Y} = T^{-1}\sum_{t=1}^{T} Y_t$, the fraction of responses equal to 1.*

*(i) If, for every $t$, $\mathbb{E}_1[Y_t \mid \mathcal{F}_{t-1}] \geq 1/2 + b$ and $\mathbb{E}_0[Y_t \mid \mathcal{F}_{t-1}] \leq 1/2 - b$, then $\widehat{L} = \mathbf{1}\{\overline{Y} > 1/2\}$ satisfies $\Pr_\ell(\widehat{L} \neq \ell) \leq e^{-2Tb^2}$ for each $\ell$.*

*(ii) If, for every $t$, $\mathbb{E}_1[Y_t \mid \mathcal{F}_{t-1}] \geq 1/2$ and $\mathbb{E}_0[Y_t \mid \mathcal{F}_{t-1}] \leq 1/2 - b$, then $\widehat{L}_{\mathrm{os}} = \mathbf{1}\{\overline{Y} \geq 1/2 - b/2\}$ satisfies $\Pr_\ell(\widehat{L}_{\mathrm{os}} \neq \ell) \leq e^{-Tb^2/2}$ for each $\ell$.*

*These bounds also hold conditionally on the information before the repeated test. No independence assumption is needed.*

*Proof.* Fix either label and write $\mu_t = \mathbb{E}[Y_t \mid \mathcal{F}_{t-1}]$ under its law. Conditional on $\mathcal{F}_{t-1}$, $Y_t - \mu_t$ has mean zero and lies in an interval of length one. Hoeffding's exponential bound (Hoeffding, 1963) gives, for every real $\theta$, $\mathbb{E}[e^{\theta(Y_t-\mu_t)} \mid \mathcal{F}_{t-1}] \le e^{\theta^2/8}$. Successive conditioning gives us

$$\mathbb{E}\left[\exp\left(\theta \sum_{t=1}^{T}(Y_t - \mu_t)\right)\middle|\mathcal{F}_0\right] \le e^{T\theta^2/8}.$$

For $a > 0$, Markov's inequality with $\theta = 4a/T$, and then with $\theta = -4a/T$, gives us

$$\Pr\left(\sum_{t=1}^{T}(Y_t - \mu_t) \ge a\middle|\mathcal{F}_0\right) \le e^{-2a^2/T}, \quad \text{and} \quad \Pr\left(\sum_{t=1}^{T}(Y_t - \mu_t) \le -a\middle|\mathcal{F}_0\right) \le e^{-2a^2/T}.$$

These bounds allow the conditional means to change after every response.

In part (i), an incorrect majority requires a deviation of at least $Tb$ from $\sum_t \mu_t$: downwards for label 1 and upwards for label 0. In part (ii), the threshold lies halfway between the two bounds on the means, so an incorrect decision requires a deviation of at least $Tb/2$. Substitution proves both error bounds, conditionally on $\mathcal{F}_0$ and hence also unconditionally. □

The next lemma is used for lower bounds. A correct algorithm must distinguish instances that require different outputs. We measure how different their answer distributions are using Kullback–Leibler (KL) divergence. This quantity is nonnegative and equals zero for identical distributions. Part (i) gives the divergence needed for a reliable decision. Part (ii) bounds the divergence obtained from the queries. Together, the two parts turn a bound on how much one query reveals into a lower bound on the number of queries. This is the standard statistical comparison for adaptive sampling; see Kaufmann et al. (2016, Lemma 1). We give the conditional argument and justify its use for unbounded query counts.

**Lemma C.2.** *Let $P_0$ and $P_1$ be probability laws of complete* query histories *on two instances. A complete history includes all queries and responses, the algorithm's private randomness, its stopping time, and its output. Write $\mathrm{KL}(P_0\|P_1)$ for their Kullback–Leibler divergence.*

*(i) If an event $E$ determined by the complete history satisfies $P_0(E) \ge 1 - \delta$ and $P_1(E) \le \delta$, where $\delta \in (0, 1/4)$, then $\mathrm{KL}(P_0\|P_1) \ge c_0 \log(1/\delta)$ with $c_0 = 1/4$.*

*(ii) Suppose the algorithm uses the same decision rule and private-randomness distribution on the two instances. Let $\mathcal{H}$ be a set of agent-lottery queries. After almost every earlier history under $P_0$, suppose that queries outside $\mathcal{H}$ have identical conditional response distributions, and queries in $\mathcal{H}$ have conditional response KL divergence at most $\kappa$. If $T_{\mathcal{H}}$ counts queries in $\mathcal{H}$, then $\mathrm{KL}(P_0\|P_1) \le \kappa \mathbb{E}_{P_0}[T_{\mathcal{H}}]$.*

*The query count may be random and unbounded.*

*Proof.* For part (i), let $p = P_0(E)$ and $q = P_1(E)$. Considering only whether $E$ occurs cannot increase KL divergence, so

$$\mathrm{KL}(P_0\|P_1) \ge p \log \frac{p}{q} + (1-p) \log \frac{1-p}{1-q}.$$

If $q = 0$, this bound is infinite. Otherwise, the right-hand side increases with $p$ and decreases with $q$ when $p > q$; values at $p = 1$ are understood by continuity. Consequently,

$$\mathrm{KL}(P_0\|P_1) \ge (1-2\delta) \log \frac{1-\delta}{\delta} \ge \frac{1}{4} \log(1/\delta).$$

The last inequality uses $1 - 2\delta \ge 1/2$ and $1 - \delta \ge \sqrt{\delta}$ for $0 < \delta < 1/4$.

For part (ii), include the private random seed at the start of the history. It has the same law in both instances. Given that seed and the earlier answers, the algorithm chooses the same query and makes the same stopping and output decisions. Thus only answers contribute to the divergence.

Let $T$ be the total number of queries. First truncate the history after $M$ queries, keeping the output if the algorithm has stopped by then, and append a common terminal symbol after stopping. Write $P_0^{(M)}, P_1^{(M)}$ for these truncated laws. The chain rule gives

$$\mathrm{KL}(P_0^{(M)} \| P_1^{(M)}) = \mathbb{E}_{P_0}\left[\sum_{t=1}^{T\wedge M} \mathrm{KL}(P_0(Y_t \mid \mathcal{F}_{t-1}) \| P_1(Y_t \mid \mathcal{F}_{t-1}))\right]$$
$$\leq \kappa\, \mathbb{E}_{P_0}\left[\sum_{t=1}^{T\wedge M} \mathbf{1}\{(i_t, \mathbf{x}_t) \in \mathcal{H}\}\right].$$

All conditional distributions here are evaluated at the same preceding history, not at separately generated histories from the two instances.

As $M$ increases, these finite histories generate the complete history, including the infinite sequence on a nonterminating run. KL divergence for the complete history is the limit of its divergences for these finite histories. Hence

$$\mathrm{KL}(P_0 \| P_1) = \lim_{M\to\infty} \mathrm{KL}(P_0^{(M)} \| P_1^{(M)}).$$

The query counts on the right increase to $T_{\mathcal{H}}$. Monotone convergence now proves the claim; no bound on the stopping time is needed, and the expected count may be infinite. □

### C.2 Finite-precision turning points and acceptance inequalities

Two facts make exact learning possible with few edge searches. First, finite precision leaves only finitely many possible turning points on each edge. Second, after the pure lottery labels are known, at most $m-1$ selected turning points determine the entire acceptance set. To see why the second fact is plausible, recall that a lottery's margin is the weighted average of its pure lottery margins. The turning points reveal the ratios between these margins, which suffice to decide acceptance even without knowing the utilities or threshold. This reconstruction is due to Choo et al. (2026); we give its formulas and proof below for completeness. It does not require finite precision.

**Lemma C.3.** *The set $R_Q$ satisfies $|R_Q| \leq (Q+1)^2$.*

*Proof.* For denominator $q \in [Q]$, there are at most $q+1$ numerators. Counting all pairs, including those that are not in lowest terms, gives

$$|R_Q| \leq \sum_{q=1}^{Q}(q+1) = \frac{Q(Q+3)}{2} \leq (Q+1)^2.$$ □

*Remark* C.4. Enumerate the pairs $1 \leq q \leq Q$, $0 \leq p \leq q$, retain those with $\gcd(p,q) = 1$, and sort the resulting rationals. The same list $R_Q = \{r_1 < \cdots < r_M\}$ is used for every agent and edge. This preprocessing uses no queries; the query bound does not include the time or memory needed to form the list.

Binary search on this list uses at most $\lceil \log_2 M \rceil + 1 = O(\log(2/\varepsilon))$ comparisons. With noiseless answers, it may query the candidates themselves. Under near-threshold noise, Algorithm 1 instead queries between consecutive candidates. Lemma 3.1 gives absolute margin at least $\varepsilon^2/2$ at each such midpoint. The proof of that margin bound is in Appendix E.

**Lemma C.5.** *Fix an agent $i$. Suppose the pure lottery labels are known, and let $L_{\mathrm{Acc}} = \{j \in [m] : q_i(\mathbf{e}_j) = 1\}$ and $L_{\mathrm{Rej}} = \{j \in [m] : q_i(\mathbf{e}_j) = 0\}$, both nonempty. For $k \in L_{\mathrm{Rej}}$ and $j \in L_{\mathrm{Acc}}$, let $\alpha_{k,j} \in (0,1]$ be the turning point satisfying $g_i(\mathbf{x}_{k,j}(\alpha_{k,j})) = 0$. Choose $r \in L_{\mathrm{Rej}}$.*

*If $\alpha_{r,j} = 1$ for every $j \in L_{\mathrm{Acc}}$, set $(\mathbf{c}_i)_j = 1$ for $j \in L_{\mathrm{Acc}}$ and $(\mathbf{c}_i)_j = 0$ for $j \in L_{\mathrm{Rej}}$. Otherwise, choose $a \in L_{\mathrm{Acc}}$ with $\alpha_{r,a} < 1$ and set*

$$(\mathbf{c}_i)_r = 0, \quad (\mathbf{c}_i)_j = \frac{1}{\alpha_{r,j}} \quad \text{for all } j \in L_{\mathrm{Acc}}, \quad (\mathbf{c}_i)_k = \frac{1-\alpha_{k,a}(\mathbf{c}_i)_a}{1-\alpha_{k,a}} \quad \text{for all } k \in L_{\mathrm{Rej}} \setminus \{r\}.$$

*In the second case $g_i(\mathbf{e}_a) > 0$, so $\alpha_{k,a} \in (0,1)$ and every denominator is nonzero. In either case, $A_i = \{\mathbf{x} \in \Delta(S) : \langle \mathbf{c}_i, \mathbf{x}\rangle \geq 1\}$. Thus at most $m-1$ edge turning points suffice once the pure lottery labels are known.*

*Proof.* A rejected pure lottery has negative margin, whereas an accepted one has nonnegative margin. The margin is affine along their edge, so it has positive slope and a unique zero in $(0, 1]$.

If $\alpha_{r,j} = 1$ for every $j \in L_{\mathrm{Acc}}$, then every accepted pure lottery has margin zero. Thus

$$g_i(\mathbf{x}) = \sum_{j=1}^{m} x_j g_i(\mathbf{e}_j) = \sum_{k \in L_{\mathrm{Rej}}} x_k g_i(\mathbf{e}_k)$$

is nonnegative exactly when $x_k = 0$ for every $k \in L_{\mathrm{Rej}}$. This is precisely the condition $\langle \mathbf{c}_i, \mathbf{x} \rangle \geq 1$ for the proposed coefficients.

Otherwise, $\alpha_{r,a} < 1$ implies $g_i(\mathbf{e}_a) > 0$, so $0 < \alpha_{k,a} < 1$ for every rejected $k$. We show that the formulas in the lemma give

$$(\mathbf{c}_i)_j = 1 - \frac{g_i(\mathbf{e}_j)}{g_i(\mathbf{e}_r)} \quad \text{for all } j \in [m].$$

The denominator is strictly negative. The identity holds at $j = r$. For an accepted $j$, the equation

$$(1 - \alpha_{r,j}) g_i(\mathbf{e}_r) + \alpha_{r,j} g_i(\mathbf{e}_j) = 0$$

gives $1 - g_i(\mathbf{e}_j)/g_i(\mathbf{e}_r) = 1/\alpha_{r,j}$, as required. For a rejected $k \neq r$, use the boundary equation on its edge to $\mathbf{e}_a$ to obtain

$$(1 - \alpha_{k,a}) \left(1 - \frac{g_i(\mathbf{e}_k)}{g_i(\mathbf{e}_r)}\right) + \alpha_{k,a} (\mathbf{c}_i)_a = 1.$$

Solving this equation gives exactly the stated formula for $(\mathbf{c}_i)_k$. Therefore, for every $\mathbf{x} \in \Delta(S)$,

$$\langle \mathbf{c}_i, \mathbf{x} \rangle - 1 = -\frac{\sum_j x_j g_i(\mathbf{e}_j)}{g_i(\mathbf{e}_r)} = -\frac{g_i(\mathbf{x})}{g_i(\mathbf{e}_r)}.$$

Since $g_i(\mathbf{e}_r) < 0$, the inequality $\langle \mathbf{c}_i, \mathbf{x} \rangle \geq 1$ is equivalent to $g_i(\mathbf{x}) \geq 0$.

The first case uses $|L_{\mathrm{Acc}}| \leq m - 1$ turning points. The second uses these same turning points and $|L_{\mathrm{Rej}}| - 1$ further ones, for a total of $m - 1$. □

### C.3 Learning acceptance sets without noise

We restate the noiseless elicitation method of Choo et al. (2026) using the finite candidate list and reconstruction lemma. This is the algorithm used by the uniform-noise simulation.

**Lemma C.6** (Noiseless finite-precision consensus elicitation)**.** *Under $\varepsilon$-precision, one agent's acceptance set can be recovered using at most $m + (m-1)(\lceil \log_2 |R_Q| \rceil + 1)$ noiseless queries. The result is* AcceptAll*,* RejectAll*, or a vector $\mathbf{c}_i$ describing $A_i = \{\mathbf{x} \in \Delta(S) : \langle \mathbf{c}_i, \mathbf{x} \rangle \geq 1\}$. Applying this to all agents solves exact consensus using at most $n(m + (m-1)(\lceil \log_2 |R_Q| \rceil + 1)) = O(nm \log(2/\varepsilon))$ noiseless queries.*

*Proof.* Query all $m$ pure lotteries for agent $i$ and form $L_{\mathrm{Acc}}$ and $L_{\mathrm{Rej}}$ as in Lemma C.5. Since $g_i(\mathbf{x}) = \sum_j x_j g_i(\mathbf{e}_j)$, return AcceptAll if all pure lotteries are accepted, and RejectAll if all are rejected.

Otherwise choose $r \in L_{\mathrm{Rej}}$. On every edge from $\mathbf{e}_r$ to $\mathbf{e}_j$, $j \in L_{\mathrm{Acc}}$, the labels satisfy

$$q_i(\mathbf{x}_{r,j}(\alpha)) = \mathbf{1}\{\alpha \geq \alpha_{r,j}\}.$$

By Lemma 3.1, $\alpha_{r,j} \in R_Q$. Because boundary points are accepted, binary search for the first accepted value in the ordered list $R_Q = \{r_1 < \cdots < r_M\}$ recovers this turning point exactly. Each search takes at most $\lceil \log_2 M \rceil + 1$ queries.

If all recovered turning points equal 1, use the first case of Lemma C.5. Otherwise choose $a \in L_{\mathrm{Acc}}$ with $\alpha_{r,a} < 1$, recover $\alpha_{k,a}$ for every $k \in L_{\mathrm{Rej}} \setminus \{r\}$ by the same search, and use the second case. The reconstruction lemma gives a correct $\mathbf{c}_i$ from at most $m - 1$ searches. Lemma C.3 gives the stated per-agent query count.

Apply this procedure to every agent. If any result is RejectAll, return NULL. Otherwise let $I$ be the agents for which a vector $\mathbf{c}_i$ was returned, and solve

$$\text{find } \mathbf{x} \in \Delta(S) \quad \text{such that} \quad \langle \mathbf{c}_i, \mathbf{x} \rangle \geq 1 \quad \text{for all } i \in I.$$

Agents outside $I$ accept every lottery, so the feasible region is exactly $\bigcap_i A_i$, including when $I$ is empty. Return a feasible lottery if one exists, and NULL otherwise. This last step uses no queries. □

# D Proofs and Additional Results for Section 2.2

## D.1 Simulation under uniform noise

The simulation must handle adaptive query choices as well as repeated answers. A uniform bound on the number of noiseless queries, including after incorrectly supplied labels, lets us allocate an error allowance before the execution begins.

**Theorem D.1.** *Let a noiseless adaptive query algorithm make at most $D \geq 1$ queries on every sequence of supplied labels, and fix $\delta \in (0,1)$. Under* $\mathsf{UniformNoise}(\eta)$*, there is a noisy query simulation such that, with probability at least $1-\delta$, every majority response supplied to the noiseless algorithm equals the corresponding noiseless label. On this event, if the noisy and noiseless executions use the same private random seed, they ask the same queries, stop after the same number of simulated queries, and return the same output. The simulation uses at most $D\lceil 2\log(D/\delta)/(1-2\eta)^2\rceil$ noisy queries.*

*Proof.* Fix the instance and an admissible oracle, and use the same private random seed in the simulated and noiseless executions. Replace each requested query by $T = \lceil \frac{2}{(1-2\eta)^2} \log \frac{D}{\delta} \rceil$ repetitions, supplying the majority label and resolving ties as 0. Conditional on the history before this test, every answer has bias at least $(1-2\eta)/2$ towards the true label. Lemma C.1(i) bounds the probability of an incorrect majority by $\exp(-T(1-2\eta)^2/2) \leq \delta/D$. The depth bound holds for every supplied label sequence, including incorrect ones, so at most $D$ tests are made. Averaging each conditional error bound and taking a union bound shows that all supplied labels are correct with probability at least $1-\delta$.

On this event, induction on the number of simulated queries shows that both executions agree: the same seed and preceding labels give the same next query, and the next supplied label is again the noiseless one. Their stopping decisions and outputs therefore agree as well. The noisy query count is at most $DT$ on every response sequence. □

## D.2 Exact consensus under uniform noise

Applying the simulation to Lemma C.6 gives exact consensus. The only extra precaution is to keep the reconstruction defined on answer sequences that could not occur without noise.

**Corollary D.2.** *For every $\delta \in (0,1)$, under $\varepsilon$-precision, exact consensus with* $\mathsf{UniformNoise}(\eta)$ *is $(1-\delta)$-correct using $O((nm\log(2/\varepsilon))/(1-2\eta)^2 \cdot \log(e\,nm\log(2/\varepsilon)/\delta))$ noisy queries.*

*Proof.* Apply Theorem D.1 to the algorithm in Lemma C.6. To make it well defined for arbitrary supplied labels, stop and return NULL whenever the reconstruction would divide by zero. Lemma C.5 shows that this never happens on a genuine noiseless execution.

Each binary search halves a finite candidate list regardless of the supplied labels. For each agent there are at most $m$ pure lottery queries and $m-1$ such searches, each taking at most $\lceil \log_2 |R_Q| \rceil + 1$ queries. Hence every supplied label sequence has depth at most

$$D = n(m + (m-1)(\lceil \log_2 |R_Q| \rceil + 1)) = O(nm\log(2/\varepsilon)).$$

The simulation is therefore correct with probability at least $1-\delta$ and uses at most

$$D\left\lceil \frac{2}{(1-2\eta)^2} \log \frac{D}{\delta} \right\rceil = O\left( \frac{nm\log(2/\varepsilon)}{(1-2\eta)^2} \log \frac{e\,nm\log(2/\varepsilon)}{\delta} \right)$$

noisy queries. The final feasibility computation uses no queries. □

# E Proofs and Additional Results for Section 3

## E.1 A worked example of midpoint separation

We now explain why the margin involves $\varepsilon^2$. With $\varepsilon = 1/4$, utilities $(0,1)$ or $(0,3/4)$ and threshold $1/4$ give turning points $1/4$ or $1/3$. At the midpoint $7/24$, the respective margins are $1/24$ and $-1/32$ (Figure 3); the smaller absolute margin is exactly $\varepsilon^2/2$.

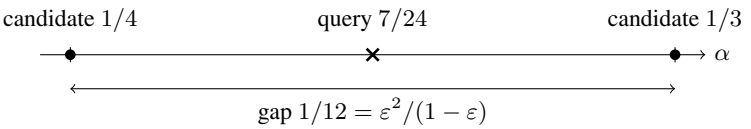


Figure 3: Querying between two possible boundary positions. At $7/24$, the agent accepts if the turning point is $1/4$ and rejects if it is $1/3$. The midpoint has nonzero margin in both instances, so repeated answers distinguish them.

### E.2 Proof of Lemma 3.1

*Proof.* The key point is that a small utility difference also limits the denominator of the turning point. Since $u_{ik} < \tau_i \le u_{ik'}$, the unique turning point is

$$\alpha^* = \frac{\tau_i - u_{ik}}{u_{ik'} - u_{ik}} \in (0, 1], \quad g_i(\mathbf{x}_{k,k'}(\alpha)) = (u_{ik'} - u_{ik})(\alpha - \alpha^*).$$

Both differences in the ratio are integer multiples of $\varepsilon$. If $\alpha^* = p/q$ is in lowest terms, its denominator therefore satisfies

$$q \le \frac{u_{ik'} - u_{ik}}{\varepsilon} \le Q.$$

In particular, $\alpha^* \in R_Q$ and $u_{ik'} - u_{ik} \ge q\varepsilon$.

Every different $p'/q' \in R_Q$ satisfies

$$\left|\frac{p'}{q'} - \frac{p}{q}\right| = \frac{|p'q - pq'|}{qq'} \ge \frac{1}{qQ},$$

since the numerator is a nonzero integer. Now let $z = (r_h + r_{h+1})/2$ for consecutive candidates. The turning point cannot lie strictly between them. If it is an endpoint, the distance to their midpoint is at least $1/(2qQ)$. If it is outside the interval, even its distance to the nearer endpoint is at least $1/(qQ)$. Thus, in either case,

$$|g_i(\mathbf{x}_{k,k'}(z))| = (u_{ik'} - u_{ik})|z - \alpha^*| \ge \frac{q\varepsilon}{2qQ} = \frac{\varepsilon^2}{2}.$$

In particular, $z \ne \alpha^*$. □

### E.3 Searching between candidate turning points

Algorithm 1, FindTurningPoint, uses binary search to find the turning point in the ordered list $R_Q$. It asks queries halfway between consecutive candidates, where Lemma 3.1 guarantees a nonzero margin. An accepted answer keeps the candidates to the left; a rejected answer keeps those to the right. Repeating each query then makes these comparisons reliable. The queried lottery may be accepted or rejected, and its coordinates need not be multiples of $\varepsilon$: finite precision restricts the unknown parameters only.

**Theorem E.1.** *Assume $\varepsilon$-precision and near-threshold noise with parameters $(\lambda, \beta)$. Let $\alpha^*$ be the turning point on the edge from a rejected $\mathbf{e}_k$ to an accepted $\mathbf{e}_{k'}$. With a supplied bound $0 < b_e \le \lambda(\varepsilon^2/2)^\beta$, Algorithm 1 returns $\alpha^*$ with probability at least $1 - \delta_e$ and uses at most $BT_e$ queries. The probability bound holds conditionally on the history before the call, apart from histories of probability zero. For $b_e = \lambda(\varepsilon^2/2)^\beta$, the query count is $O_\beta(\lambda^{-2}\varepsilon^{-4\beta}\log(2/\varepsilon)\log\frac{2\log(2/\varepsilon)}{\delta_e})$.*

*Proof.* First suppose every majority label is correct. Initially the candidate interval contains $\alpha^*$. At any step, $z = (r_h + r_{h+1})/2$ is not in $R_Q$, whereas $\alpha^*$ is. The margin increases along the edge from the rejected endpoint to the accepted endpoint. Thus an accepted midpoint means $\alpha^* < z$, equivalently $\alpha^* \le r_h$. A rejected midpoint means $\alpha^* > z$, equivalently $\alpha^* \ge r_{h+1}$. Either update retains the true turning point and reduces the candidate count by at least half, up to rounding. After at most $B$ comparisons, the remaining candidate is $\alpha^*$.

**Algorithm 1** FINDTURNINGPOINT

**Require:** Agent $i \in N$; distinct $k, k' \in [m]$; error allowance $\delta_e \in (0,1)$; $\varepsilon$-precision; supplied $0 < b_e \le \lambda(\varepsilon^2/2)^\beta$.
*Assumption:* $q_i(\mathbf{e}_k) = 0$ and $q_i(\mathbf{e}_{k'}) = 1$.
Use the precomputed list $R_Q = \{r_1 < \cdots < r_M\}$; set $B \leftarrow \lceil \log_2 M \rceil + 1$.
$T_e \leftarrow \lceil (2b_e^2)^{-1} \log(B/\delta_e) \rceil$.
$\ell \leftarrow 1, u \leftarrow M$.
**while** $\ell < u$ **do**
  $h \leftarrow \lfloor (\ell + u)/2 \rfloor$; $z \leftarrow (r_h + r_{h+1})/2$.
  Query $(i, \mathbf{x}_{k,k'}(z))$ $T_e$ times; let $\overline{Y}$ be the empirical mean.
  **if** $\overline{Y} > 1/2$ **then**
    $u \leftarrow h$.
  **else**
    $\ell \leftarrow h + 1$.
**return** $r_\ell$.

We now bound the chance of an incorrect comparison, conditional on the history before the call. Every queried point is a midpoint between consecutive candidates, even if an earlier comparison was incorrect. Lemma 3.1 therefore gives absolute margin at least $\varepsilon^2/2$ at each query. Conditional on the history before its repeated test, the response bias toward the true label is at least $b_e$. Lemma C.1(i) gives error probability at most $e^{-2T_e b_e^2} \le \delta_e/B$. A union bound over at most $B$ comparisons proves correctness. Finally, $B = O(\log(2/\varepsilon))$ and $b_e^{-2} = 2^{2\beta}\lambda^{-2}\varepsilon^{-4\beta}$ for the stated choice of $b_e$, which gives the query bound. □

### E.4 SINGLE-AGENT LEARNING

To use the edge search, we first need to know which pure lotteries are accepted. An accepted pure lottery may lie on the boundary and have response mean $1/2$, so ordinary majority voting is not sufficient. Finite precision nevertheless separates rejection from acceptance: every rejected pure lottery has margin at most $-\varepsilon$, while every accepted one has margin at least zero. We place the decision threshold slightly below $1/2$ to distinguish these two cases.

**Lemma E.2.** *Under $\varepsilon$-precision and* NearThresholdNoise$(\lambda, \beta)$*, fix an agent $i$, pure lottery $\mathbf{e}_j$, error allowance $\delta_v \in (0,1)$, and supplied bound $0 < b_v \le \lambda\varepsilon^\beta$. Query $(i, \mathbf{e}_j)$ $T_v = \lceil 2b_v^{-2} \log(1/\delta_v) \rceil$ times and return $\widehat{q}_j = \mathbf{1}\{\overline{Y} \ge 1/2 - b_v/2\}$. The probability that $\widehat{q}_j \ne q_i(\mathbf{e}_j)$ is at most $\delta_v$, also conditionally on the preceding history. For $b_v = \lambda\varepsilon^\beta$, the test uses $O(\lambda^{-2}\varepsilon^{-2\beta}\log(2/\delta_v))$ queries.*

*Proof.* If $\mathbf{e}_j$ is rejected, $\varepsilon$-precision gives $g_i(\mathbf{e}_j) \le -\varepsilon$, so every response has conditional mean at most $1/2 - b_v$. If it is accepted, the conditional mean is at least $1/2$, including when its margin is zero. Lemma C.1(ii) gives error probability at most $e^{-T_v b_v^2/2} \le \delta_v$. □

Algorithm 2, LEARNACCEPTANCESET, first classifies the pure lotteries, then searches at most $m-1$ edges and applies Lemma C.5. It outputs either an acceptance inequality or the conclusion that every lottery, or no lottery, is accepted. The checks before division keep the algorithm defined even when noisy tests are incorrect.

**Theorem E.3.** *Under $\varepsilon$-precision and* NearThresholdNoise$(\lambda, \beta)$*, let Algorithm 2 use supplied bounds $0 < b_v \le \lambda\varepsilon^\beta$ and $0 < b_e \le \lambda(\varepsilon^2/2)^\beta$. For every agent $i$ and $\delta_i \in (0,1)$, with probability at least $1 - \delta_i$ it returns a correct description of $A_i$:* AcceptAll *if $A_i = \Delta(S)$,* RejectAll *if $A_i = \varnothing$, or a vector $\mathbf{c}_i$ with $A_i = \{\mathbf{x} \in \Delta(S) : \langle \mathbf{c}_i, \mathbf{x} \rangle \ge 1\}$. The guarantee also holds conditionally on the preceding history. Its query count is*

$$O\left( m b_v^{-2} \log \frac{4m}{\delta_i} + m b_e^{-2} \log(2/\varepsilon) \log \frac{2m \log(2/\varepsilon)}{\delta_i} \right).$$

*For $b_v = \lambda\varepsilon^\beta$ and $b_e = \lambda(\varepsilon^2/2)^\beta$, this is $O_\beta(m\lambda^{-2}\varepsilon^{-4\beta}\log(2/\varepsilon)\log(2m\log(2/\varepsilon)/\delta_i))$.*

**Algorithm 2** LEARNACCEPTANCESET

**Require:** Agent $i \in N$; $\delta_i \in (0,1)$; $\varepsilon$-precision; supplied bounds $0 < b_v \le \lambda\varepsilon^{\beta}$ and $0 < b_e \le \lambda(\varepsilon^2/2)^{\beta}$.

Test every pure lottery by Lemma E.2, using $b_v$ and error allowance $\delta_i/(4m)$.
$L_{\mathrm{Acc}} \leftarrow \{j : \widehat{q}_j = 1\}$; $L_{\mathrm{Rej}} \leftarrow [m] \setminus L_{\mathrm{Acc}}$.
**if** $L_{\mathrm{Rej}} = \varnothing$ **then**
    **return** AcceptAll.
**if** $L_{\mathrm{Acc}} = \varnothing$ **then**
    **return** RejectAll.
Choose $r \in L_{\mathrm{Rej}}$.
For every $j \in L_{\mathrm{Acc}}$, set $\alpha_{r,j} \leftarrow$ FINDTURNINGPOINT$(i, r, j, 3\delta_i/(4m))$, using $b_e$.
**if** $\alpha_{r,j} = 1$ for every $j \in L_{\mathrm{Acc}}$ **then**
    Set $(\mathbf{c}_i)_j = 1$ for $j \in L_{\mathrm{Acc}}$ and $(\mathbf{c}_i)_j = 0$ for $j \in L_{\mathrm{Rej}}$.
    **return** $\mathbf{c}_i$.
Choose $a \in L_{\mathrm{Acc}}$ with $\alpha_{r,a} < 1$.
For every $k \in L_{\mathrm{Rej}} \setminus \{r\}$, set $\alpha_{k,a} \leftarrow$ FINDTURNINGPOINT$(i, k, a, 3\delta_i/(4m))$, using $b_e$.
**if** $\alpha_{r,j} = 0$ for some $j \in L_{\mathrm{Acc}}$, or $\alpha_{k,a} = 1$ for some $k \in L_{\mathrm{Rej}} \setminus \{r\}$ **then**
    **return** RejectAll.
Set $(\mathbf{c}_i)_r = 0$ and $(\mathbf{c}_i)_j = 1/\alpha_{r,j}$ for $j \in L_{\mathrm{Acc}}$.
Set $(\mathbf{c}_i)_k = (1 - \alpha_{k,a}(\mathbf{c}_i)_a)/(1 - \alpha_{k,a})$ for $k \in L_{\mathrm{Rej}} \setminus \{r\}$.
**return** $\mathbf{c}_i$.

*Proof.* Fix the history before the call. Lemma E.2 and a union bound give probability at most $\delta_i/4$ that a pure lottery label is incorrect. If all labels agree and are correct, the returned AcceptAll or RejectAll is correct because $g_i(\mathbf{x}) = \sum_j x_j g_i(\mathbf{e}_j)$.

Otherwise, condition on any pure-test history with all labels correct. Every edge subsequently searched has a rejected first endpoint and an accepted second endpoint. This remains true even if an earlier search was incorrect, since the endpoint indices are chosen from the fixed sets $L_{\mathrm{Rej}}$ and $L_{\mathrm{Acc}}$. Each of at most $m-1$ calls to FINDTURNINGPOINT therefore has conditional error probability at most $3\delta_i/(4m)$. A further union bound gives probability at most $3\delta_i/4$ that any turning point is incorrect.

When all labels and turning points are correct, Lemma C.5 gives the exact acceptance set and ensures that every denominator used is nonzero. The total error probability is at most $\delta_i/4 + 3\delta_i/4 = \delta_i$.

There are $m$ pure tests and at most $m-1$ edge searches on every response sequence. Summing their bounds gives the two query terms. Substituting the stated $b_v, b_e$, and using $\varepsilon^{-2\beta} \le \varepsilon^{-4\beta}$, gives the final bound. □

### E.5 PROOF OF THEOREM 3.2

Algorithm 3 learns each agent's acceptance inequality and then asks whether these inequalities have a common solution in the simplex. This last step is a linear feasibility computation and uses no queries. Each learning call has error allowance $\delta_i = \delta/n$; the repetition counts are specified in Lemma E.2 and Algorithm 1.

*Proof.* By Theorem E.3, each call returns an incorrect description with conditional probability at most $\delta/n$. There are at most $n$ calls, so with probability at least $1 - \delta$ every returned description is correct. On this event, any RejectAll output proves that the common intersection is empty. Otherwise, agents outside $I$ accept every lottery, and $\{\mathbf{x} \in \Delta(S) : \langle \mathbf{c}_i, \mathbf{x} \rangle \ge 1 \text{ for every } i \in I\} = \bigcap_{i \in N} A_i$. Thus the final feasibility problem gives a correct exact output, including when $I$ is empty.

For arbitrary supplied bounds, summing the single-agent query counts gives

$$O\left( nmb_v^{-2} \log \frac{4mn}{\delta} + nmb_e^{-2} \log(2/\varepsilon) \log \frac{2nm \log(2/\varepsilon)}{\delta} \right).$$

**Algorithm 3** EXACTCONSENSUS

**Require:** $\delta \in (0,1)$; $\varepsilon$-precision; the positive bias bounds $b_v, b_e$ from Algorithm 2.

```
I ← ∅.
for i ∈ N do
    Run LEARNACCEPTANCESET(i, δ/n) with these bounds.
    if the result is RejectAll then
        return NULL.
    else if the result is a vector c_i then
        I ← I ∪ {i}.
Find x ∈ Δ(S) with ⟨c_i, x⟩ ≥ 1 for every i ∈ I.
Return such a lottery if one exists, and NULL otherwise.
```

Substituting $b_v = \lambda\varepsilon^\beta$ and $b_e = \lambda(\varepsilon^2/2)^\beta$ proves the theorem. The final feasibility computation uses no additional queries. □

### E.6 EXACT LOWER BOUND CONSTRUCTION

We construct an instance with exactly one common lottery. Moving one agent's turning point from $\varepsilon$ to $\varepsilon/(1-\varepsilon)$ makes this lottery unacceptable and leaves no common lottery. The move is only of order $\varepsilon^2$. We choose response probabilities that are identical outside the interval between these two turning points. Within the interval, their difference is only of order $\lambda\varepsilon^{2\beta}$. Thus every query gives little information about whether consensus is possible.

#### E.6.1 INSTANCES TO BE DISTINGUISHED

Let $m = 2$ and use the lotteries $\mathbf{x}_{1,2}(\alpha)$ from the preliminaries. Since $\varepsilon \le 1/4$, $Q = 1/\varepsilon \ge 4$. Define

$$\alpha_0 = \varepsilon = \frac{1}{Q}, \quad \alpha_1 = \frac{\varepsilon}{1-\varepsilon} = \frac{1}{Q-1}, \quad \Delta = \alpha_1 - \alpha_0 = \frac{\varepsilon^2}{1-\varepsilon}.$$

Then $\varepsilon^2 \le \Delta \le (4/3)\varepsilon^2$. The two turning points are consecutive in $R_Q$: if $1/Q < p/q < 1/(Q-1)$ with $q \le Q$, then $(Q-1)p < q \le Q$ forces $p = 1$, leaving the impossible condition $Q-1 < q < Q$.

Agent 1 has utilities $(1, 0)$ and threshold $1-\varepsilon$, so it accepts exactly $\alpha \le \alpha_0$. Each agent $h \in \{2, \dots, n\}$ has one of two types. Superscripts denote the type, while the subscript continues to denote the agent:

| type | $u_h$ | $\tau_h$ | $g_h^b(\mathbf{x}_{1,2}(\alpha))$ |
|---|---|---|---|
| 0 | $(0,1)$ | $\varepsilon$ | $\alpha - \alpha_0$ |
| 1 | $(0, 1-\varepsilon)$ | $\varepsilon$ | $(1-\varepsilon)(\alpha - \alpha_1)$. |

All utilities and thresholds satisfy $\varepsilon$-precision. When every agent $h \ge 2$ has type 0, the only common lottery is $\mathbf{x}_{1,2}(\alpha_0)$. Changing any one of them to type 1 requires $\alpha \ge \alpha_1 > \alpha_0$, making consensus infeasible.

The noise model bounds the response bias without specifying it exactly. We use this freedom to make the two types give identical answer distributions wherever their labels agree. A Bernoulli response with mean $p$ returns 1 with probability $p$ and 0 otherwise.

**Lemma E.4.** *Fix an agent $h \ge 2$ in this construction. Its two types admit near-threshold response rules with parameters $(\lambda, \beta)$ and independent randomness at each query, whose Bernoulli means $p_0(\alpha), p_1(\alpha)$ satisfy $\mathrm{KL}(\mathrm{Bernoulli}(p_0(\alpha)) \| \mathrm{Bernoulli}(p_1(\alpha))) \le C\lambda^2\Delta^{2\beta}$ for all $\alpha \in [0,1]$, for a universal constant $C$. The bound holds in the reverse direction as well. The means depend only on the current query, not on its history, and all other agents may use the same response rules in both instances.*

*Proof.* When the two types agree on acceptance, the noise model allows us to give them identical response probabilities. Choose

$$p_0(\alpha) = p_1(\alpha) = \begin{cases} \frac{1}{2} - \lambda|g_h^1(\mathbf{x}_{1,2}(\alpha))|^\beta, & \alpha < \alpha_0, \\ \frac{1}{2} + \lambda|g_h^0(\mathbf{x}_{1,2}(\alpha))|^\beta, & \alpha \ge \alpha_1. \end{cases}$$

Below $\alpha_0$, both margins are negative and the type-1 absolute margin is larger. At or above $\alpha_1$, both margins are nonnegative and the type-0 margin is larger. Choosing the larger required bias gives a response probability allowed for both types.

On $[\alpha_0, \alpha_1)$, type 0 accepts and type 1 rejects, so their response probabilities must differ. Both absolute margins are at most $\Delta$, allowing us to set

$$p_0(\alpha) = \frac{1}{2} + \lambda\Delta^\beta, \quad p_1(\alpha) = \frac{1}{2} - \lambda\Delta^\beta.$$

At $\alpha_0$, the positive probability is also allowed for the zero type-0 margin. All chosen means lie in $[1/4, 3/4]$, because the margin magnitudes and $\Delta$ are at most one and $\lambda \le 1/4$.

The distributions are identical outside this interval. Inside it, use the elementary bound, valid for $0 \le b \le 1/4$,

$$\mathrm{KL}(\mathrm{Bernoulli}(1/2 + b) \| \mathrm{Bernoulli}(1/2 - b)) = 2b \log \frac{1 + 2b}{1 - 2b} \le \frac{8b^2}{1 - 2b} \le 16b^2.$$

The first inequality follows from $\log(1 + x) \le x$. The KL divergence is unchanged when these two Bernoulli laws are reversed. Taking $b = \lambda\Delta^\beta$ proves the claim. □

**Corollary E.5.** *Let $\delta \in (0, 1/4)$. With the response rules of Lemma E.4, every adaptive randomized test that identifies agent $h$'s type with error at most $\delta$ under either type satisfies $\mathbb{E}_b[T_h] \ge c_\beta \lambda^{-2} \varepsilon^{-4\beta} \log(1/\delta)$ for all $b \in \{0, 1\}$, where $T_h$ counts queries to that agent and $c_\beta > 0$ depends only on $\beta$.*

*Proof.* We compare the divergence needed for a correct decision with the divergence obtained from the queries. Let $P_0, P_1$ be the laws of the complete query histories and let $E$ be the event that the test returns type 0. Correctness and Lemma C.2(i) give us

$$\mathrm{KL}(P_0 \| P_1) \ge c_0 \log(1/\delta). \tag{1}$$

Only queries to $h$ distinguish the instances, so Lemmas E.4 and C.2(ii) give us

$$\mathrm{KL}(P_0 \| P_1) \le C\lambda^2 \Delta^{2\beta} \mathbb{E}_0[T_h]. \tag{2}$$

Combining the bounds and using $\Delta^{-2\beta} = (1 - \varepsilon)^{2\beta} \varepsilon^{-4\beta} \ge (3/4)^{2\beta} \varepsilon^{-4\beta}$ proves the result under type 0. Reverse the two laws and use $E^c$ to obtain the result under type 1. An infinite expected count already satisfies the lower bound. □

### E.7 Proof of Theorem 3.3

*Proof.* Each of the $n - 1$ agents can independently make the common lottery unacceptable. The algorithm must therefore query every one of these agents enough to rule out its type-1 instance. Use the instances in Appendix E.6.1. Let $P_0$ be the law of the complete query history when every agent $h \ge 2$ has type 0, and let $P_h$ be the law when only agent $h$ changes to type 1. Use the response rules of Lemma E.4, with the same rule for agent 1 in every instance.

Let $E$ be the event that the algorithm stops and returns a lottery. The reference instance is feasible, whereas each comparison instance is infeasible. Correctness gives $P_0(E) \ge 1 - \delta$ and $P_h(E) \le \delta$. If $T_h$ counts queries to agent $h$, Lemma C.2 yields

$$c_0 \log(1/\delta) \le \mathrm{KL}(P_0 \| P_h) \le C\lambda^2 \Delta^{2\beta} \mathbb{E}_{P_0}[T_h].$$

All these expected counts are taken under the same reference law $P_0$, so they can be added. If $T$ is the total query count,

$$\mathbb{E}_{P_0}[T] \ge \sum_{h=2}^{n} \mathbb{E}_{P_0}[T_h] \ge c_\beta (n - 1) \lambda^{-2} \varepsilon^{-4\beta} \log(1/\delta).$$

Here we again used $\Delta \le (4/3)\varepsilon^2$. This is an admissible instance, so its expected count lower-bounds the worst case. The argument also applies if some expectation is infinite. □

### E.8 Proof of Proposition 3.4: No uniform finite query bound

Without finite precision, the distance between turning points can be arbitrarily small. Choo et al. (2026) already establish that exact consensus has no finite worst-case query bound with noiseless answers and unrestricted real parameters for algorithms required to be always correct. The following construction quantifies the cost in terms of noise, separation $\Delta$, and the allowed error probability.

**Lemma E.6.** *Fix $\Delta \in (0, 1/4]$, $\lambda \in (0, 1/4]$, and $\beta > 0$. With two alternatives, consider an agent $h$ whose utilities are $(0, 1)$ and whose threshold is either $1/2$ (type 0) or $1/2 + \Delta$ (type 1). There are admissible Bernoulli response rules for the two types such that every query $\mathbf{x}_{1,2}(\alpha)$ has KL divergence at most $C\lambda^2\Delta^{2\beta}$ in either direction. The means depend only on the query, and responses use independent randomness at each query.*

*Proof.* Set

$$p_0(\alpha) = p_1(\alpha) = \begin{cases} \frac{1}{2} - \lambda(1/2 + \Delta - \alpha)^\beta, & \alpha < 1/2, \\ \frac{1}{2} + \lambda(\alpha - 1/2)^\beta, & \alpha \geq 1/2 + \Delta, \end{cases}$$

and set $p_0(\alpha) = 1/2 + \lambda\Delta^\beta$, $p_1(\alpha) = 1/2 - \lambda\Delta^\beta$ on $[1/2, 1/2 + \Delta)$. Outside this interval both types have the same label, and the chosen common probability uses the larger absolute margin. Inside it the labels differ, but both absolute margins are at most $\Delta$. Thus all means satisfy Definition 2.4, including its boundary condition, and lie in $[1/4, 3/4]$. The Bernoulli calculation in Lemma E.4 gives the bound in both directions. □

*Proof of Proposition 3.4.* Let agent 1 have utilities $(1, 0)$ and threshold $1/2$, so it accepts $\mathbf{x}_{1,2}(\alpha)$ exactly when $\alpha \leq 1/2$. Let agent 2 have one of the types in Lemma E.6. With type 0, the unique common lottery is $\mathbf{x}_{1,2}(1/2)$. With type 1, no common lottery exists.

Let $P_0, P_1$ be the laws of the complete query histories, using the lemma's response rules for agent 2 and the same admissible rule for agent 1 in both instances. For the event $E$ that the algorithm stops and returns a lottery, correctness gives $P_0(E) \geq 1 - \delta$ and $P_1(E) \leq \delta$. Therefore

$$\mathrm{KL}(P_0 \| P_1) \geq c_0 \log(1/\delta). \tag{3}$$

If $T_2$ counts queries to agent 2, the per-query bound gives

$$\mathrm{KL}(P_0 \| P_1) \leq C\lambda^2\Delta^{2\beta}\mathbb{E}_{P_0}[T_2]. \tag{4}$$

Hence the expected total query count on the feasible instance is at least $c\lambda^{-2}\Delta^{-2\beta}\log(1/\delta)$. For every finite proposed worst-case bound, choosing $\Delta$ sufficiently small makes this lower bound larger. Thus no finite worst-case expected query bound is possible without a restriction on the separation. □

## F Proofs and Additional Results for Section 4

For robust consensus, we learn a set between $A_i^{+\rho}$ and $A_i$: it includes every lottery with margin at least $\rho$, and every lottery in it is accepted. Recovering the exact boundary is unnecessary. We first give a test that can establish acceptance. We then use this test to select pure lotteries and points on edges, and return all mixtures of these selected points. The procedures use $\lambda$ and $\beta$ to set their repetition counts.

### F.1 One-sided margin tests

The test distinguishes negative margin from margin at least $\gamma$. It may return either value when $0 \leq g_i(\mathbf{x}) < \gamma$. Its decision threshold is slightly above $1/2$, so a return value of 1 establishes acceptance with high probability. A return value of 0 means only that the margin is below $\gamma$; the lottery may still be accepted.

**Lemma F.1.** *Under* $\mathsf{NearThresholdNoise}(\lambda, \beta)$, *let* $a = \textsc{PositiveTest}_\gamma(i, \mathbf{x}, \delta_0)$. *With probability at least $1 - \delta_0$, $a = 1 \implies g_i(\mathbf{x}) \geq 0$, $a = 0 \implies g_i(\mathbf{x}) < \gamma$. The same guarantee holds conditionally on the history before the test. It uses $O(\lambda^{-2}\gamma^{-2\beta}\log(2/\delta_0))$ queries.*

**Algorithm 4** POSITIVETEST$_\gamma$

**Require:** Agent $i \in N$; $\mathbf{x} \in \Delta(S)$; $\gamma \in (0, 1]$; $\delta_0 \in (0, 1)$; known $\lambda, \beta$.

1: $b \leftarrow \lambda\gamma^\beta$; $T \leftarrow \lceil 2b^{-2} \log(1/\delta_0) \rceil$.
2: Query $(i, \mathbf{x})$ $T$ times; let $\overline{Y}$ be the empirical mean.
3: Return 1 if $\overline{Y} \geq 1/2 + b/2$, and 0 otherwise.

*Proof.* Condition on the history before the test. If $g_i(\mathbf{x}) \geq \gamma$, every response has conditional mean at least $1/2 + b$. Returning 0 then requires a downward deviation of more than $b/2$ in the empirical mean. If $g_i(\mathbf{x}) < 0$, every conditional mean is at most $1/2$, and returning 1 requires an upward deviation of at least $b/2$. The concentration bound proved in Lemma C.1 bounds either probability by $e^{-Tb^2/2} \leq \delta_0$.

Only one of these cases can apply to the fixed queried point, so there is no need to add their error probabilities. In the remaining case $0 \leq g_i(\mathbf{x}) < \gamma$, either returned bit satisfies both implications. This also proves Lemma 4.1. □

### F.2 FINDING ACCEPTED EDGE POINTS BELOW MARGIN $\rho$

Suppose an edge starts at a pure lottery with margin below $\rho/2$ and ends at an accepted pure lottery. We seek an accepted point on this edge whose margin is below $\rho$. Repeatedly bisecting the edge keeps one endpoint below margin $\rho/2$ and the other accepted. When the distance between their edge positions is at most $\rho/2$, their margins differ by at most $\rho/2$, since $|u_{ik} - u_{id}| \leq 1$. The accepted endpoint then has the required margin.

**Algorithm 5** INNEREDGE

**Require:** Agent $i \in N$; distinct $d, k \in [m]$; $\rho \in (0, 1]$; $\delta_e \in (0, 1)$; known $\lambda, \beta$.
*Assumption:* $g_i(\mathbf{e}_d) < \rho/2$ and $g_i(\mathbf{e}_k) \geq 0$.

$\gamma \leftarrow \rho/2$; $L \leftarrow \lceil \log_2(2/\rho) \rceil + 1$.
$\alpha^- \leftarrow 0$; $\alpha^+ \leftarrow 1$.
**for** $t = 1, \ldots, L$ **do**
    $z \leftarrow (\alpha^- + \alpha^+)/2$.
    $a \leftarrow$ POSITIVETEST$_\gamma(i, \mathbf{x}_{d,k}(z), \delta_e/L)$.
    **if** $a = 1$ **then**
        $\alpha^+ \leftarrow z$.
    **else**
        $\alpha^- \leftarrow z$.
**return** $\mathbf{y}_{d,k} := \mathbf{x}_{d,k}(\alpha^+)$.

**Lemma F.2.** *Under* NearThresholdNoise$(\lambda, \beta)$*, suppose* $g_i(\mathbf{e}_d) < \rho/2$ *and* $g_i(\mathbf{e}_k) \geq 0$*. With probability at least* $1 - \delta_e$*, Algorithm 5 returns* $\mathbf{y}_{d,k}$ *with* $0 \leq g_i(\mathbf{y}_{d,k}) < \rho$*. The guarantee holds conditionally on any preceding history for which the input conditions hold. For* $L = \lceil \log_2(2/\rho) \rceil + 1$*, the query count is*

$$O_\beta\left(\lambda^{-2}\rho^{-2\beta} L \log \frac{L}{\delta_e}\right) = O_\beta\left(\lambda^{-2}\rho^{-2\beta} \log(2/\rho) \log\left(2 + \frac{\log(2/\rho)}{\delta_e}\right)\right).$$

*Proof.* Condition on the history before the call. Each of the $L$ tests has conditional error probability at most $\delta_e/L$, even when its midpoint depends on earlier answers. A union bound gives probability at least $1 - \delta_e$ that every test satisfies Lemma F.1. On this event, the maintained conditions are

$$g_i(\mathbf{x}_{d,k}(\alpha^-)) < \rho/2, \quad g_i(\mathbf{x}_{d,k}(\alpha^+)) \geq 0.$$

They hold at the initial endpoints. A test returning 1 preserves the condition at $\alpha^+$; a test returning 0 preserves the condition at $\alpha^-$. These statements do not require a rejected left endpoint or an increasing margin: both endpoint conditions may hold at the same point.

After $L$ bisections, $\alpha^+ - \alpha^- = 2^{-L} \leq \rho/2$. Since utilities lie in $[0, 1]$,

$$0 \leq g_i(\mathbf{x}_{d,k}(\alpha^+)) \leq g_i(\mathbf{x}_{d,k}(\alpha^-)) + |u_{ik} - u_{id}|(\alpha^+ - \alpha^-) < \rho/2 + \rho/2 = \rho.$$

This proves the required margin bound. Each test uses $O_\beta(\lambda^{-2}\rho^{-2\beta}\log(L/\delta_e))$ queries. Multiplying by $L = O(\log(2/\rho))$ gives the stated cost. □

### F.3 INNER APPROXIMATION OF ONE ACCEPTANCE SET

We retain the pure lotteries whose tests return 1. For every edge from a remaining pure lottery to a retained one, we also retain the point returned by INNEREDGE. With high probability, all retained points are accepted. The *convex hull* of the set $V_i$ of retained points, denoted $\mathrm{conv}(V_i)$, consists of all mixtures of these points: each receives a nonnegative weight, and the weights sum to one. Such mixtures remain accepted because their margins are weighted averages of the points' margins.

The less immediate fact is that these mixtures include every lottery with margin at least $\rho$. We prove this by checking the vertices, or corner points, of $A_i^{+\rho}$. Each vertex is either a pure lottery or lies on an edge, so the pure lottery tests and edge searches suffice.

**Algorithm 6** INNERAPPROX

**Require:** Agent $i \in N$; $\rho \in (0,1]$; $\delta_i \in (0,1)$; known $\lambda, \beta$.
$\gamma \leftarrow \rho/2$.
**for** $j \in [m]$ **do**
  $a_j \leftarrow \mathrm{POSITIVETEST}_\gamma(i, \mathbf{e}_j, \delta_i/(4m))$.
$K_i \leftarrow \{j : a_j = 1\}$; $D_i \leftarrow [m] \setminus K_i$.
$V_i \leftarrow \{\mathbf{e}_k : k \in K_i\}$.
**for** $d \in D_i$ **do**
  **for** $k \in K_i$ **do**
    $\mathbf{y}_{d,k} \leftarrow \mathrm{INNEREDGE}(i, d, k, \rho, \delta_i/(4m^2))$.
    $V_i \leftarrow V_i \cup \{\mathbf{y}_{d,k}\}$.
**return** $V_i$.

**Lemma F.3.** *Under* NearThresholdNoise$(\lambda, \beta)$*, Algorithm 6 returns a finite set* $V_i \subseteq \Delta(S)$ *such that, with probability at least* $1 - \delta_i$, $A_i^{+\rho} \subseteq \widehat{A}_i := \mathrm{conv}(V_i) \subseteq A_i$, *where* $\mathrm{conv}(\varnothing) = \varnothing$. *The guarantee also holds conditionally on the preceding history. For the sets* $D_i, K_i$ *obtained from the pure tests, the query count is*

$$O_\beta\left(\lambda^{-2}\rho^{-2\beta}[m + |D_i||K_i|\log(2/\rho)]\log\frac{4m^2\log(2/\rho)}{\delta_i}\right).$$

*In particular, it is* $O_\beta(m^2\lambda^{-2}\rho^{-2\beta}\log(2/\rho)\log(4m^2\log(2/\rho)/\delta_i))$ *in the worst case.*

*Proof.* **The selected points are accepted.** The pure tests satisfy all implications in Lemma F.1 except with probability at most $\delta_i/4$. Conditional on any pure-test history where they hold,

$$k \in K_i \implies g_i(\mathbf{e}_k) \geq 0, \quad d \in D_i \implies g_i(\mathbf{e}_d) < \rho/2.$$

Every edge call then satisfies the assumptions of Lemma F.2. There are at most $|D_i||K_i| \leq m^2$ calls, each with conditional error probability at most $\delta_i/(4m^2)$. Thus, with probability at least $1-\delta_i/2$, all pure-test implications hold and every selected edge point has margin in $[0, \rho)$. Work on this event. Every point in $V_i$ is accepted, so affinity of $g_i$ gives $\widehat{A}_i \subseteq A_i$.

If $K_i = \varnothing$, every pure margin is below $\rho/2$ and $A_i^{+\rho} = \varnothing$. If $D_i = \varnothing$, then $\widehat{A}_i = \Delta(S)$. The desired other inclusion is immediate in either case, and also when $A_i^{+\rho} = \varnothing$.

**Which points must the inner set contain?** In the remaining cases, $A_i^{+\rho}$ is a nonempty bounded intersection of finitely many linear inequalities, hence a polytope. Such a set is the convex hull of its vertices (Grötschel et al., 1993), so it suffices to show that every vertex belongs to $\widehat{A}_i$. We first justify the description of these vertices directly; the argument also covers lower-dimensional sets. If a point $\mathbf{z}$ has margin strictly above $\rho$ and two positive coordinates, shifting a sufficiently small amount of probability between those coordinates in either direction keeps it in $A_i^{+\rho}$. It is then strictly between two distinct feasible points and is not a vertex. If $g_i(\mathbf{z}) = \rho$ and at least three coordinates are positive, the two linear equations $\sum_j h_j = 0$ and $\sum_j h_j g_i(\mathbf{e}_j) = 0$ have a nonzero solution $\mathbf{h}$

with $h_j = 0$ whenever $z_j = 0$. For sufficiently small $t > 0$, both $\mathbf{z} + t\mathbf{h}$ and $\mathbf{z} - t\mathbf{h}$ are nonnegative lotteries of margin $\rho$. Again $\mathbf{z}$ is not a vertex.

It follows that a vertex is either a pure lottery with margin at least $\rho$, or a non-pure point of margin $\rho$ on an edge. In the latter case, the endpoint margins lie strictly on opposite sides of $\rho$. Indeed, the value $\rho$ is a strict convex combination of the two margins; if both equalled $\rho$, the entire edge would be feasible and its interior points would not be vertices. The same argument applies at any real margin level, including zero.

**Every required vertex belongs to $\widehat{A}_i$.** A pure vertex has margin at least $\rho$, so its test returns 1 and it is retained. Orient any non-pure vertex's edge from $\mathbf{e}_j$ to $\mathbf{e}_k$ with $g_i(\mathbf{e}_j) < \rho < g_i(\mathbf{e}_k)$. Necessarily $k \in K_i$. If $j \in K_i$, both endpoints are retained. Otherwise, the algorithm retains $\mathbf{y}_{j,k}$ of margin below $\rho$. Along this edge the margin is strictly increasing, so the level-$\rho$ vertex lies between $\mathbf{y}_{j,k}$ and $\mathbf{e}_k$. In both cases the vertex belongs to $\widehat{A}_i$.

Since $\widehat{A}_i$ is convex and contains all these vertices, it contains $A_i^{+\rho}$. The event used has probability at least $1 - \delta_i/2 \geq 1 - \delta_i$, conditionally on the history before the call.

Finally, there are $m$ pure tests and $|D_i||K_i|$ edge calls on every response sequence. Their costs from Lemmas F.1 and F.2 give the displayed bounds. □

### F.4 Robust consensus algorithm and proof of Theorem 4.2

Once the sets $V_i$ are known, finding a point common to their convex hulls requires no more queries. A common lottery may have a different representation as a mixture of each agent's retained points. If every $V_i$ is nonempty, introduce one nonnegative weight $\theta_{i,\mathbf{v}}$ for each $\mathbf{v} \in V_i$. For each agent, these weights must sum to one and give the same lottery $\mathbf{x}$. The final linear feasibility problem is

$$\begin{aligned} &\text{find } \mathbf{x},\ \{\theta_{i,\mathbf{v}} : i \in N,\ \mathbf{v} \in V_i\} \\ &\text{subject to } \mathbf{x} \in \Delta(S), \\ &\qquad \theta_{i,\mathbf{v}} \geq 0 && (i \in N,\ \mathbf{v} \in V_i), \\ &\qquad \sum_{\mathbf{v} \in V_i} \theta_{i,\mathbf{v}} = 1 && (i \in N), \\ &\qquad \mathbf{x} = \sum_{\mathbf{v} \in V_i} \theta_{i,\mathbf{v}} \mathbf{v} && (i \in N). \end{aligned} \tag{5}$$

Since $|V_i| \leq m + m^2$, this program has $O(nm^2)$ variables, $O(nm)$ equality constraints, and $O(nm^2)$ nonnegativity constraints.

**Algorithm 7** RobustConsensus

**Require:** $\rho \in (0, 1]$; $\delta \in (0, 1)$; known $\lambda, \beta$.
1: **for** $i \in N$ **do**
2: $V_i \leftarrow \text{InnerApprox}(i, \rho, \delta/n)$.
3: **if** $V_i = \varnothing$ **then**
4: **return** NULL.
5: Solve the linear feasibility problem (5).
6: Return its lottery $\mathbf{x}$ if feasible, and NULL otherwise.

*Proof of Theorem 4.2.* Each call to InnerApprox has conditional error probability at most $\delta/n$. A union bound over at most $n$ calls gives probability at least $1 - \delta$ that every constructed set satisfies $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i$. Work on this event.

In a Yes instance, a common margin-$\rho$ lottery belongs to every constructed $\widehat{A}_i$. No $V_i$ is empty, and the final program is feasible. Every returned lottery lies in all $\widehat{A}_i$, hence in all $A_i$. In a No instance, a point in all $\widehat{A}_i$ would also lie in all $A_i^{-\rho}$, contradicting the No condition. The algorithm therefore returns NULL. On an instance satisfying neither condition, any returned lottery still belongs to all $\widehat{A}_i \subseteq A_i$. This proves all requirements of Definition 2.2.

The query count follows by summing Lemma F.3 over the at most $n$ calls with $\delta_i = \delta/n$:

$$O_\beta\left(nm^2\lambda^{-2}\rho^{-2\beta}\log(2/\rho)\log\frac{4nm^2\log(2/\rho)}{\delta}\right).$$

The linear program adds no queries. □

### F.5 Optimizing a linear objective

If we have a linear objective for choosing among lotteries, we can maximize it over $\bigcap_{i\in N}\widehat{A}_i$ at no additional query cost. On the same event of probability at least $1-\delta$, $\bigcap_{i\in N}A_i^{+\rho}\subseteq\bigcap_{i\in N}\widehat{A}_i\subseteq\bigcap_{i\in N}A_i$. Thus, in a robust YES instance, every agent accepts the returned lottery, and its objective value is at least as large as that of any lottery with common margin $\rho$.

### F.6 Why the output margin need not be a fixed fraction of $\rho$

The returned lottery is accepted, but its margin need not be a fixed positive fraction of $\rho$. No algorithm can ensure margin $c\rho$ for a fixed $c>0$ independent of $m$ on every robust YES instance. To see why, fix $n=1$ and $0<\rho\le 1/2$. For each $j\in[m]$, take $\tau_1=1/2$, $u_{1j}=1/2+\rho$, and $u_{1k}=1/2$ for $k\neq j$. Then $g_1(\mathbf{x})=\rho x_j$: every lottery is accepted, and $\mathbf{e}_j$ has margin $\rho$. Achieving large margin requires assigning enough probability to alternative $j$. Independent replies with constant mean $1/2+\lambda\rho^\beta$ satisfy the noise model on every instance, including at zero margin. These replies give no information about $j$, so every adaptive algorithm has the same distribution of queries, replies, and outputs on all instances. For any $c\in(0,1]$, a lottery has at most $1/c$ coordinates at least $c$. Averaging over $j$, some instance has probability at most $1/(mc)$ of receiving an output with margin at least $c\rho$. For fixed $c>0$, this is less than $1-\delta$ for sufficiently large $m$. A known response rule can allow a positive output margin; see Appendix I.

### F.7 Robust lower bound construction and proof of Theorem 4.3

The exact lower bound changed feasibility by moving a turning point by order $\varepsilon^2$. Robust consensus does not require distinguishing all such nearly feasible instances. We instead move a threshold by $5\rho$, enough to change a robust YES instance into a robust NO instance. The response probabilities can again be identical wherever the two types agree on acceptance, and differ by only order $\lambda\rho^\beta$ elsewhere.

**Lemma F.4.** *Fix $\rho\in(0,1/10]$, $\lambda\in(0,1/4]$, and $\beta>0$. With two alternatives, let agent $h$ have utilities $(0,1)$ and threshold $1/2-2\rho$ in type 0, or $1/2+3\rho$ in type 1. Thus, $g_h^0(\mathbf{x}_{1,2}(\alpha))=\alpha-1/2+2\rho$ and $g_h^1(\mathbf{x}_{1,2}(\alpha))=\alpha-1/2-3\rho$. There are Bernoulli response rules satisfying* NearThresholdNoise$(\lambda,\beta)$ *for these types such that, at every $\alpha\in[0,1]$, $\mathrm{KL}(\mathrm{Bernoulli}(p_0(\alpha))\|\mathrm{Bernoulli}(p_1(\alpha)))\le 16\lambda^2(5\rho)^{2\beta}$. Their means depend only on the query, and responses use independent randomness at each query.*

*Proof.* The thresholds lie in $(0,1)$ because $\rho\le 1/10$. Set

$$p_0(\alpha)=p_1(\alpha)=\begin{cases}\frac{1}{2}-\lambda(1/2+3\rho-\alpha)^\beta, & \alpha<1/2-2\rho,\\ \frac{1}{2}+\lambda(\alpha-1/2+2\rho)^\beta, & \alpha\ge 1/2+3\rho.\end{cases}$$

In each case the two types agree on acceptance, and the common probability uses the larger absolute margin. On $[1/2-2\rho,1/2+3\rho)$, their labels differ and both absolute margins are at most $5\rho$. Set $p_0(\alpha)=1/2+\lambda(5\rho)^\beta$ and $p_1(\alpha)=1/2-\lambda(5\rho)^\beta$ there. These probabilities satisfy the required inequalities, including at the two boundary points, and lie in $[1/4,3/4]$. The distributions are identical outside this interval. Inside it, the Bernoulli bound from Lemma E.4 gives the stated divergence. □

*Proof of Theorem 4.3.* Agent 1 has utilities $(1,0)$ and threshold $1/2$, giving $g_1(\mathbf{x}_{1,2}(\alpha))=1/2-\alpha$. Every other agent has one of the types in Lemma F.4. When all have type 0, the lottery $\mathbf{x}_{1,2}(1/2-\rho)$ has margin exactly $\rho$ for every agent. This is a YES instance.

If only agent $h \geq 2$ changes to type 1, there is no common lottery even after every threshold is lowered by $\rho$. Indeed, membership in $A_1^{-\rho}$ requires $\alpha \leq 1/2 + \rho$, whereas membership in $A_h^{-\rho}$ requires $\alpha \geq 1/2 + 2\rho$. These inequalities are incompatible, so every such comparison instance is a NO instance.

Let $P_0$ be the law of the complete query history of the all-type-0 instance and $P_h$ that of the instance where only $h$ changes. Use the lemma's response rules and keep agent 1's response rule fixed. For the event $E$ that the algorithm stops and returns a lottery, correctness gives $P_0(E) \geq 1 - \delta$ and $P_h(E) \leq \delta$. Hence

$$\mathrm{KL}(P_0 \| P_h) \geq c_0 \log(1/\delta). \tag{6}$$

Let $T$ be the total query count and $T_h$ the number of queries to agent $h$. Only those queries distinguish this pair, so Lemma C.2(ii) gives

$$\mathrm{KL}(P_0 \| P_h) \leq 16\lambda^2 (5\rho)^{2\beta} \mathbb{E}_{P_0}[T_h]. \tag{7}$$

The algorithm must distinguish the YES instance from every one of these NO instances. Combining the bounds for each $h$ and adding them under the same law $P_0$ yields

$$\mathbb{E}_{P_0}[T] \geq \sum_{h=2}^{n} \mathbb{E}_{P_0}[T_h] \geq c_\beta (n-1) \lambda^{-2} \rho^{-2\beta} \log(1/\delta).$$

This proves the worst-case expected lower bound. Infinite expectations already satisfy it. □

### F.8 CHECKING A FIXED LOTTERY

The preceding algorithm searches for a common lottery. Here the lottery is supplied, and we only check it. A NO answer is required when some agent rejects it, but is also allowed when all agents accept it and some agent's margin is below $\rho$. Thus a NO answer need not establish rejection or rule out a common acceptable lottery.

**Definition F.5** (Fixed-lottery test)**.** Fix $\rho \in (0, 1]$ and a supplied lottery $\mathbf{x} \in \Delta(S)$. With probability at least $1 - \delta$, the test must return YES if $g_i(\mathbf{x}) \geq \rho$ for every $i$, and NO if $g_i(\mathbf{x}) < 0$ for some $i$. If neither condition holds, either answer is allowed. On every instance, a YES answer must imply $\mathbf{x} \in \bigcap_i A_i$ on the same event of probability at least $1 - \delta$.

**Proposition F.6.** *Under* $\mathsf{NearThresholdNoise}(\lambda, \beta)$ *with known* $\lambda, \beta$*, the fixed-lottery test can be solved using* $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ *queries for every* $\delta \in (0, 1)$*. More generally, a supplied lower bound* $0 < b_\rho \leq \lambda\rho^\beta$ *suffices, with query count* $O(nb_\rho^{-2}\log(2n/\delta))$*.*

*Proof.* For each agent, run $\textsc{PositiveTest}_\rho(i, \mathbf{x}, \delta/n)$. Return YES if all tests return 1, and NO otherwise. With probability at least $1 - \delta$, every test has the implications in Lemma F.1. On this event, common margin at least $\rho$ forces every test to return 1, while any negative margin forces a 0. A YES output always means every agent accepts $\mathbf{x}$. This proves Proposition 5.1 as well.

Each of the $n$ tests costs $O(\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ queries. When only $b_\rho$ is supplied, use cutoff $1/2 + b_\rho/2$ and $\lceil 2b_\rho^{-2}\log(n/\delta)\rceil$ repetitions. The same proof applies because $b_\rho$ is a lower bound on the required positive bias. □

The next result shows that the factors $n$, $\lambda^{-2}$, $\rho^{-2\beta}$, and $\log(1/\delta)$ are necessary for checking a supplied lottery. This remains true even if the test can query other lotteries to learn about the agents.

**Proposition F.7.** *Fix* $n \geq 1$*,* $\rho \in (0, 1/10]$*,* $\lambda \in (0, 1/4]$*,* $\beta > 0$*, and* $\delta \in (0, 1/4)$*. There is a fixed two-alternative lottery* $\mathbf{x}^*$ *such that every adaptive randomized test that is* $(1 - \delta)$*-correct on every instance and admissible oracle has worst-case expected query count* $\Omega_\beta(n\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$*. This holds even when the test may query any lottery, not only* $\mathbf{x}^*$*.*

*Proof.* Let $\mathbf{x}^* = \mathbf{x}_{1,2}(1/2 - \rho)$, and let each of the $n$ agents have one of the types in Lemma F.4. At the supplied lottery, $g_h^0(\mathbf{x}^*) = \rho$ and $g_h^1(\mathbf{x}^*) = -4\rho$. Let $P_0$ be the law of the complete query history when all agents have type 0 and $P_h$ the law when only agent $h$ changes to type 1, using the lemma's response rules. The supplied lottery meets the test's YES condition under $P_0$ and its NO condition under $P_h$. For the event $E$ that the test returns YES, correctness gives

$$\mathrm{KL}(P_0 \| P_h) \geq c_0 \log(1/\delta). \tag{8}$$

Let $T$ be the total query count and $T_h$ the number of queries to agent $h$. Each such query, wherever it is made, has divergence at most $16\lambda^2(5\rho)^{2\beta}$. Other agents have identical response distributions in this pair of instances. Therefore

$$\mathrm{KL}(P_0 \| P_h) \le 16\lambda^2(5\rho)^{2\beta}\mathbb{E}_{P_0}[T_h]. \tag{9}$$

Adding the resulting lower bounds over all $h \in N$ gives us

$$\mathbb{E}_{P_0}[T] \ge c_\beta n \lambda^{-2}\rho^{-2\beta}\log(1/\delta).$$

This proves the proposition. □

Thus checking a supplied lottery already requires the statistical factor $\lambda^{-2}\rho^{-2\beta}$. The additional factors in our robust algorithm come from constructing inner approximations, not from checking one point. This comparison does not show that the same additional factors are necessary for consensus search.

# G Proofs for the Additional Results

## G.1 Deciding at an unbiased response

Without a known positive lower bound on the bias, a test cannot both distinguish arbitrarily small positive and negative biases and stop in finite expected time at zero bias. The same difficulty arises in the sign tests of probabilistic bisection (Frazier et al., 2019). Here $\mathrm{Bernoulli}(p)$ means independent binary answers, each equal to 1 with probability $p$.

**Proposition G.1.** *Fix $\delta \in (0, 1/4)$. No sequential test with outputs "positive" and "negative" that is given no positive lower bound on the bias can simultaneously satisfy all three properties below:*

*(i) it has finite expected stopping time under* $\mathrm{Bernoulli}(1/2)$ *replies;*

*(ii) for every $b \in (0, 1/4]$, under* $\mathrm{Bernoulli}(1/2 + b)$ *replies it outputs "positive" with probability at least $1 - \delta$;*

*(iii) for every $b \in (0, 1/4]$, under* $\mathrm{Bernoulli}(1/2 - b)$ *replies it outputs "negative" with probability at least $1 - \delta$.*

Appendix H avoids requiring a decision at zero bias: it permits continued sampling when the observations do not yet justify an output.

*Proof of Proposition G.1.* The idea is that a decision made after finitely many fair coin flips cannot change abruptly when we give the coin a very small bias. Suppose such a test exists, and let $T$ be the number of queries it makes before stopping. Under unbiased replies, finite expected stopping time implies $T < \infty$ almost surely. Since the test must then return one of its two outputs, at least one output has probability at least $1/2$.

Write $\Pr_p$ for probabilities when each reply is an independent Bernoulli bit with mean $p$. Suppose the output with probability at least $1/2$ is "positive." Choose $M$ large enough that

$$\Pr_{1/2}(T \le M \text{ and the output is positive}) > 1/2 - \delta/2.$$

The event above depends only on the first $M$ bits and the algorithm's private random seed. Its probability varies continuously with $p$: sum the probabilities of the finitely many length-$M$ binary sequences, weighting each by the probability that the algorithm makes this decision on that sequence. Thus, for sufficiently small $b > 0$,

$$\Pr_{1/2-b}(T \le M \text{ and the output is positive}) > 1/2 - \delta > \delta,$$

contradicting property (iii). If "negative" has probability at least $1/2$ under unbiased replies, use mean $1/2 + b$ instead to contradict property (ii). □

### G.2 Exact lower bound in the number of alternatives

For one agent, an accepted lottery exists exactly when a pure lottery is accepted. The algorithm may nevertheless have to check every alternative, because any one of them can be the only accepted choice.

**Theorem G.2.** *Fix $m \geq 2$, $n = 1$, $\varepsilon \in (0, 1]$ with $1/\varepsilon \in \mathbb{N}$, and $\delta \in (0, 1/4)$. Under $\varepsilon$-precision and $\mathsf{NearThresholdNoise}(\lambda, \beta)$, every $(1-\delta)$-correct exact-consensus algorithm has worst-case expected query complexity $\Omega(m\lambda^{-2}\varepsilon^{-2\beta}\log(1/\delta))$.*

*Proof of Theorem G.2.* Use the construction of Choo et al. (2026) in which any pure lottery may be the only accepted one, scaled to $\varepsilon$-precision. In the reference instance, set $u_1 = 0$ and $\tau_1 = \varepsilon$, so all lotteries are rejected. In comparison instance $j \in [m]$, set $u_1 = \varepsilon \mathbf{e}_j$ and keep $\tau_1 = \varepsilon$. Its margin is $g_1^j(\mathbf{x}) = \varepsilon(x_j - 1)$, so $\mathbf{e}_j$ is its only accepted lottery. All parameters satisfy $\varepsilon$-precision.

In the reference instance, use independent Bernoulli replies of mean $1/2 - \lambda\varepsilon^\beta$ at every query. In comparison $j$, use this same mean except at $\mathbf{e}_j$, where the mean is $1/2$. These rules are admissible: away from $\mathbf{e}_j$, the comparison margin is negative with magnitude at most $\varepsilon$; at $\mathbf{e}_j$, it is zero. Hence only queries exactly at $\mathbf{e}_j$ distinguish comparison $j$ from the reference.

Let $P_0, P_j$ be the complete-history laws in these instances. Let $T$ be the total query count and $T_j$ count queries at $\mathbf{e}_j$. The event of returning NULL has probability at least $1 - \delta$ under $P_0$ and at most $\delta$ under $P_j$. Lemma C.2 therefore gives

$$c_0 \log(1/\delta) \leq \mathrm{KL}(P_0 \| P_j) \leq 4\lambda^2\varepsilon^{2\beta}\mathbb{E}_{P_0}[T_j].$$

For the last inequality, $\log z \leq z - 1$ gives

$$\mathrm{KL}(\mathrm{Bernoulli}(p) \| \mathrm{Bernoulli}(q)) \leq \frac{(p-q)^2}{q(1-q)},$$

and here $q = 1/2$.

A query can be at only one pure lottery, so $\sum_j T_j \leq T$. All expectations are under the same reference law; adding the $m$ bounds yields

$$\mathbb{E}_{P_0}[T] \geq \frac{c_0 m}{4\lambda^2\varepsilon^{2\beta}} \log(1/\delta).$$

An infinite expected count already satisfies this inequality. □

### G.3 One-agent exact consensus and extension to additional dimensions

For one agent, checking pure lotteries attains the bound in Theorem G.2 up to logarithmic factors. Since $g_1(\mathbf{x}) = \sum_j x_j g_1(\mathbf{e}_j)$, a mixture cannot have higher margin than every pure lottery it mixes. An accepted lottery therefore exists if and only if a pure lottery is accepted. Testing each pure lottery with error allowance $\delta/m$ (Lemma E.2) and returning an accepted one, or NULL if none exists, uses $O(m\lambda^{-2}\varepsilon^{-2\beta}\log(2m/\delta))$ queries. The larger $\varepsilon^{-4\beta}$ cost is a consequence of requiring agreement between agents, already with two alternatives.

Theorems 3.3 and G.2 extend to any fixed $n, m \geq 2$ when $\varepsilon \leq 1/4$. In the two-alternative construction, give each extra alternative the same utility as the first alternative for every agent. For each query, combine their probabilities with that of the first alternative and use the original response distribution. Thus the extra alternatives provide no additional information. In the one-agent construction, add agents who accept every lottery and always answer 1. Taking the larger of the resulting two lower bounds gives, up to a factor of two, $\Omega_\beta(\lambda^{-2}\log(1/\delta)[(n-1)\varepsilon^{-4\beta} + m\varepsilon^{-2\beta}])$. These bounds do not show that all $m$ alternatives incur the $\varepsilon^{-4\beta}$ cost.

### G.4 Robust lower bound in the number of alternatives

The corresponding robust construction makes the identity of the only margin-$\rho$ lottery unknown. A query concentrated on one alternative then gives no information about most of the other possibilities.

**Theorem G.3.** *Fix $m \geq 2$, $n = 1$, $\rho \in (0, 1/10]$, and $\delta \in (0, 1/4)$. Under* NearThresholdNoise$(\lambda, \beta)$*, every $(1-\delta)$-correct $\rho$-robust consensus algorithm has worst-case expected query complexity $\Omega_\beta(m\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$.*

*Proof of Theorem G.3.* In the reference instance, set $u_1 = 0$ and $\tau_1 = 2\rho$. The margin is $-2\rho$ everywhere, so $A_1^{-\rho} = \varnothing$: this is a robust NO instance. For each $j \in [m]$, define comparison instance $j$ by $u_1 = 3\rho\mathbf{e}_j$ and $\tau_1 = 2\rho$. Its margin is $g_1^j(\mathbf{x}) = 3\rho x_j - 2\rho$, and $g_1^j(\mathbf{e}_j) = \rho$, making it a robust YES instance. All parameters are permitted because $\rho \leq 1/10$.

Set $p_0 = 1/2 - \lambda(2\rho)^\beta$ and $p_+ = 1/2 + \lambda\rho^\beta$. Use independent replies with mean $p_0$ everywhere in the reference instance. In comparison $j$, use $p_0$ when $x_j < 2/3$ and $p_+$ when $x_j \geq 2/3$. In these regions the comparison margin is respectively negative of magnitude at most $2\rho$, or nonnegative and at most $\rho$. The rules are therefore admissible, including at the boundary $x_j = 2/3$.

Let $P_0, P_j$ be the complete-history laws. Let $T$ be the total query count and $T_j$ count queries with $x_j \geq 2/3$. The event of returning NULL has probability at least $1-\delta$ under $P_0$ and at most $\delta$ under $P_j$. Moreover, the replies differ only at these $T_j$ queries. At each one,

$$\mathrm{KL}(\mathrm{Bernoulli}(p_0)\|\mathrm{Bernoulli}(p_+)) \leq \frac{16}{3}(p_0 - p_+)^2 \leq C_\beta\lambda^2\rho^{2\beta},$$

since $p_+ \in [1/2, 3/4]$. Lemma C.2 gives

$$\mathbb{E}_{P_0}[T_j] \geq c_\beta\lambda^{-2}\rho^{-2\beta}\log(1/\delta).$$

No lottery has two coordinates at least $2/3$. Thus $\sum_j T_j \leq T$, and summing under $P_0$ proves the result, also when the expected count is infinite. □

Giving extra alternatives the same utilities as the first alternative, as in Appendix G.3, and adding agents with constant margin $1/2$ preserve both required cases. Taking the larger of the resulting two lower bounds gives, for every $n, m \geq 2$, worst-case expected query complexity $\Omega_\beta((n+m-1)\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$.

### G.5 Query complexity of inner approximations

*Proof of Theorem 4.4.* We choose $r^2$ edge lotteries that an inner set must contain in a reference instance. For each lottery, a comparison instance rejects it and therefore requires its exclusion. The algorithm must distinguish all these comparisons, but may query anywhere in the simplex. We will show that the sum of the divergences from any one query is only $O_\beta(\lambda^2\rho^{2\beta})$. Adding the information required by Lemma C.2 then gives the factor $r^2$ per agent.

**Instances and required set membership.** Put $r = m/2$ and partition $[m]$ into two sets $D, K$ of size $r$. In the reference instance, the pure lotteries in $D$ will be rejected and those in $K$ accepted. Set $a = \frac{\rho}{100r^{1/\beta}}$ and $b = \frac{a(1-4\rho)}{400\rho}$. The choice of $a$ gives $r^2a^{2\beta} = 100^{-2\beta}\rho^{2\beta}$, which will bound the information from queries rejected by the reference. The choice of $b$ will make comparison $(d,k)$ reject the selected point on edge $(d,k)$. Since $r \geq 2$ and $\rho \leq 1/16$, we have $0 < a \leq \rho/100$ and $0 < b \leq 1/40000$. For this one-agent construction, superscripts indicate the instance. Define the reference margin by

$$g_1^0(\mathbf{x}) = -a\sum_{d\in D} x_d + \frac{1}{2}\sum_{k\in K} x_k,$$

and, for each $(d,k) \in D \times K$, define the comparison margin by

$$g_1^{d,k}(\mathbf{x}) = -\frac{a}{100}x_d + a\sum_{d'\in D\setminus\{d\}} x_{d'} + bx_k + \frac{1}{2}\sum_{k'\in K\setminus\{k\}} x_{k'}.$$

All these margins are realized with threshold $\tau_1 = 1/2$: add $1/2$ to each displayed coefficient to obtain utilities in $[0,1]$. The supplied lottery $\mathbf{x}^{(0)} = \frac{1}{r}\sum_{k\in K}\mathbf{e}_k$ has reference margin $1/2$ and comparison margin $((r-1)/2 + b)/r \geq 1/4$, since $r \geq 2$. Thus the same known lottery remains acceptable with a large margin in every instance. Learning the entire inner set will still require many queries.

Let $\alpha = (\rho + a)/(1/2 + a)$ and $\mathbf{z}_{d,k} = \mathbf{x}_{d,k}(\alpha)$. Each of these points has reference margin $\rho$. Also $\alpha \le 2(\rho + a) \le 2.02\rho < 4\rho$, and the definition of $b$ gives $b + a/100 = a/(400\rho)$. Hence

$$g_1^{d,k}(\mathbf{z}_{d,k}) = \left(b + \frac{a}{100}\right)\alpha - \frac{a}{100} = \frac{a}{100}\left(\frac{\alpha}{4\rho} - 1\right) < 0.$$

A valid inner approximation must therefore contain $\mathbf{z}_{d,k}$ in the reference instance and omit it in comparison $(d, k)$. This is the difference between the outputs that we will use in the information bound.

**Choose similar response probabilities across instances.** For a query $\mathbf{x}$, let $q = \sum_{k \in K} x_k$, its total probability on $K$. Then

$$g_1^0(\mathbf{x}) = (a + 1/2)q - a, \quad -a/100 \le g_1^{d,k}(\mathbf{x}) \le q/2 + a.$$

Use Bernoulli responses with reference mean

$$p_0(\mathbf{x}) = \begin{cases} 1/2 - \lambda a^\beta, & g_1^0(\mathbf{x}) < 0, \\ 1/2 + \lambda(q/2 + a)^\beta, & g_1^0(\mathbf{x}) \ge 0. \end{cases}$$

In comparison $(d, k)$, use the same mean whenever the reference and comparison agree on acceptance. When they disagree, use

$$p_{d,k}(\mathbf{x}) = \begin{cases} 1/2 - \lambda(a/100)^\beta, & g_1^0(\mathbf{x}) \ge 0 > g_1^{d,k}(\mathbf{x}), \\ 1/2 + \lambda(2a)^\beta, & g_1^0(\mathbf{x}) < 0 \le g_1^{d,k}(\mathbf{x}). \end{cases}$$

Generate replies independently at each query. A negative reference margin has magnitude at most $a$, and a negative comparison margin has magnitude at most $a/100$. Hence the common negative mean is permitted whenever both reject. Every nonnegative margin is at most $q/2 + a$, so the common positive mean is permitted whenever both accept. If the reference rejects, then $q < a/(a + 1/2)$ and $q/2 + a < 2a$, justifying the positive mean used by any accepting comparison. The remaining differing negative mean uses the bound $a/100$. Zero margins are counted as accepted throughout. Finally, $q/2 + a < 1$ and $2a < 1$, so all means lie in $[1/4, 3/4]$.

**Bound the information obtained by one query.** There are two cases. If the reference accepts the queried lottery, at most one comparison can reject it. If the reference rejects it, many comparisons may accept it, but all their response means will be very close to the reference mean. We prove both statements.

Suppose first that $g_1^0(\mathbf{x}) \ge 0$ but $g_1^{d,k}(\mathbf{x}) < 0$. The only negative term in the comparison margin is $-(a/100)x_d$. Each positive term must be smaller in magnitude, giving

$$a(1 - q - x_d) < \frac{a}{100}x_d, \quad \frac{1}{2}(q - x_k) < \frac{a}{100}x_d.$$

The first inequality gives $x_d > 100(1 - q)/101 \ge (1 - q)/2$, so at most one $d$ can occur. Reference acceptance gives $q \ge a/(a + 1/2) \ge a > 0$. The second inequality therefore gives

$$x_k > q - \frac{a}{50}x_d \ge q - \frac{a}{50} \ge \frac{49}{50}q > \frac{q}{2}.$$

At most one $k$ can hold more than half of the probability on $K$. Thus at most one pair $(d, k)$ has a negative comparison margin.

We next bound the difference between the means for that one comparison. The inequality $bx_k < (a/100)x_d \le (a/100)(1 - x_k)$ gives $x_k < 4\rho$. The preceding inequality on $q - x_k$ then gives

$$q < x_k + a/50 < 4\rho + a/50 \le 5\rho.$$

Consequently,

$$|p_0(\mathbf{x}) - p_{d,k}(\mathbf{x})| = \lambda((q/2 + a)^\beta + (a/100)^\beta) \le C_\beta \lambda \rho^\beta.$$

The constants $C_\beta$ may depend on $\beta$ but not on $n, m, \lambda, \rho$, or $\delta$.

For Bernoulli means in $[1/4, 3/4]$,

$$\mathrm{KL}(\mathrm{Bernoulli}(p) \| \mathrm{Bernoulli}(p')) \le \frac{(p - p')^2}{p'(1 - p')} \le \frac{16}{3}(p - p')^2.$$

If $g_1^0(\mathbf{x}) \geq 0$, at most the one comparison just considered contributes nonzero divergence, so the total is at most $C_\beta \lambda^2 \rho^{2\beta}$. If $g_1^0(\mathbf{x}) < 0$, all $r^2$ comparisons may contribute, but each differing mean is only $\lambda(a^\beta + (2a)^\beta)$ away from $p_0(\mathbf{x})$. Their total divergence is at most $C_\beta r^2 \lambda^2 a^{2\beta} \leq C_\beta \lambda^2 \rho^{2\beta}$. Thus every possible query satisfies

$$\sum_{d \in D,\, k \in K} \mathrm{KL}(\mathrm{Bernoulli}(p_0(\mathbf{x})) \| \mathrm{Bernoulli}(p_{d,k}(\mathbf{x}))) \leq C_\beta \lambda^2 \rho^{2\beta}. \tag{10}$$

**Adding the requirements across agents.** For $n$ agents, let $P_0$ be the distribution of the complete query history when every agent has the reference margin and response rule. Let $P_{i,d,k}$ be the distribution when only agent $i$ changes to comparison $(d, k)$. The same supplied $\mathbf{x}^{(0)}$ has margin at least $1/4$ for every agent in every instance.

Let $E_{i,d,k}$ be the event that the algorithm returns sets with $\mathbf{z}_{d,k} \in \widehat{A}_i$. This is measurable by the theorem's assumption on the returned descriptions, and does not occur on a run that returns no sets. The required set inclusions give

$$P_0(E_{i,d,k}) \geq 1 - \delta, \quad P_{i,d,k}(E_{i,d,k}) \leq \delta.$$

Therefore $\mathrm{KL}(P_0 \| P_{i,d,k}) \geq c_0 \log(1/\delta)$ for each comparison. Let $T$ be the total query count. A query to agent $i$ contributes only to comparisons that change that agent. Sum the divergences in the chain rule of Lemma C.2 over all comparisons, always taking expectations under $P_0$. At each query, apply (10) to bound this sum. For an unbounded query count, first stop the calculation after a fixed number of queries and then let that number increase, as justified by monotone convergence in the lemma. We obtain

$$c_0 n r^2 \log(1/\delta) \leq \sum_{i \in N} \sum_{d \in D,\, k \in K} \mathrm{KL}(P_0 \| P_{i,d,k}) \leq C_\beta \lambda^2 \rho^{2\beta} \mathbb{E}_{P_0}[T].$$

If the reference expectation is infinite, the lower bound is immediate. Otherwise, rearrange and use $r = m/2$ to obtain the claimed result. □

### G.6 Robust consensus with an initial lottery

The algorithm follows the adaptive elicitation method of Choo et al. (2026): check the proposed lottery, and learn an agent's set only when needed. Here a noisy test may leave acceptance unresolved, so we learn an inner approximation when that test returns 0. Recall that the test in Lemma F.1 has two guarantees: output 1 implies acceptance, and margin at least $\rho$ implies output 1. Output 0 need not mean rejection; it means that the test did not establish acceptance at its chosen accuracy.

Once some sets have been learned, the algorithm chooses a lottery in their intersection and checks it with the remaining agents. Each restart learns the set of one more agent, which limits the number of restarts. With high probability, each learned set contains every lottery with margin at least $\rho$ for that agent, so the intersection of the learned sets never excludes a common lottery with margin at least $\rho$.

*Proof of Theorem 5.2.* **All tests and learned sets are correct together.** Use the error allowances stated in Algorithm 8. Each restart adds a new agent to $L$, so the algorithm stops after at most $n + 1$ passes through the loop, on every response sequence. There are at most $n$ tests before the first restart, fewer than $n(n + 1)$ later tests, and at most $n$ calls to construct inner approximations.

The next test or set construction may depend on earlier replies. Nevertheless, conditional on that earlier history, Lemmas F.1 and F.3 give the stated guarantee for each call. Adding the error probabilities gives at most $\delta/4 + \delta/4 + \delta/2 = \delta$. Thus, with probability at least $1 - \delta$, every test has its two implications and every constructed set satisfies $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i$. Work on this event for the rest of the proof.

**The output meets the robust objective.** In a robust YES instance, any common lottery with margin at least $\rho$ lies in every learned set. Their intersection is therefore never empty, so the algorithm cannot return NULL.

Whenever the algorithm returns $\mathbf{x}$, each agent in $L$ accepts it because $\mathbf{x} \in \widehat{A}_i \subseteq A_i$. Every remaining agent's test returned 1 at that same $\mathbf{x}$ in the final pass, which also implies acceptance. Hence

**Algorithm 8** Adaptive robust consensus from an initial lottery $\mathbf{x}^{(0)}$

**Require:** Initial lottery $\mathbf{x}^{(0)} \in \Delta(S)$; $\rho \in (0, 1]$; $\delta \in (0, 1)$; known $\lambda, \beta$.
Use error allowance $\delta/(4n)$ for each test before the first restart, $\delta/(4n(n+1))$ for each later test, and $\delta/(2n)$ for each inner approximation.

$L \leftarrow \varnothing$; $\mathbf{x} \leftarrow \mathbf{x}^{(0)}$. ▷ $L$ contains agents whose inner sets have been learned.
**loop**
  **if** $L \neq \varnothing$ **then**
    **if** $\bigcap_{i \in L} \widehat{A}_i = \varnothing$ **then**
      Return NULL.
    Keep $\mathbf{x}$ if $\mathbf{x} \in \bigcap_{i \in L} \widehat{A}_i$; otherwise choose a lottery in this intersection.
  **for** each $i \notin L$ **do**
    Run POSITIVETEST$_\rho$ on $(i, \mathbf{x})$.
    **if** the test returns 0 **then**
      $V_i \leftarrow$ INNERAPPROX$(i, \rho, \delta/(2n))$; $\widehat{A}_i \leftarrow \mathrm{conv}(V_i)$.
      Add $i$ to $L$ and restart the loop.
  Return $\mathbf{x}$.

every returned lottery is unanimously acceptable, including on instances satisfying neither robust condition. In a robust NO instance, no such lottery exists because $\bigcap_i A_i \subseteq \bigcap_i A_i^{-\rho} = \varnothing$. Since the algorithm always stops, it must then return NULL.

**The cost depends on how many sets are learned.** The following count holds on every response sequence, without conditioning on the correctness event. If the algorithm learns sets for $K$ agents, it has $K$ restarts and makes at most $n(K+1)$ margin tests. Each test uses $O(\lambda^{-2}\rho^{-2\beta}\log(2n(n+1)/\delta))$ queries. Each of the $K$ inner approximations costs

$$O_\beta\left(m^2\lambda^{-2}\rho^{-2\beta}\log(2/\rho)\log\frac{8nm^2\log(2/\rho)}{\delta}\right).$$

Adding these costs and using the larger logarithm in the theorem proves the $K$-dependent bound. If $g_i(\mathbf{x}^{(0)}) \geq \rho$ for every agent, all first-pass tests return 1 on the same event. Then $K = 0$, and the $n$ tests use $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ queries. □

### G.7 WORST-CASE GUARANTEE WITH A PREDICTED LOTTERY

A good prediction can save queries, but a poor one should not force us to spend more than the bound for the algorithm without a prediction. To obtain both guarantees, run Algorithms 7 and 8 with error allowance $\delta/2$ each. Keep their internal states separate, alternate one noisy query from each, and return the first completed output. After each response, perform the corresponding algorithm's computations until it either returns an output or requests its next query.

The replies need not be independent. At each query, the noise inequalities hold conditional on the full preceding history, including the other algorithm's queries and answers. Taking a further conditional expectation with respect to the information available to either algorithm preserves these inequalities. Thus each algorithm's repeated tests have the same concentration bounds even though queries from the other algorithm occur between their repetitions. Each proof's conditional error bounds and union bounds still apply.

For the analysis, imagine continuing the unfinished algorithm after the first output. With probability at least $1 - \delta$, both completed outputs meet their guarantees. On this event, whichever output was returned first is correct for robust consensus, and any returned lottery is unanimously acceptable.

By the time either algorithm has received $t$ of its own answers, the combined procedure has made at most $2t$ queries. Algorithm 7 has the deterministic worst-case budget in Theorem 4.2, with $\delta/2$ in place of $\delta$. Thus the combined procedure uses at most twice that budget for every prediction.

When $g_i(\mathbf{x}^{(0)}) \geq \rho$ for all agents, the algorithm using the initial lottery completes its first pass within $O(n\lambda^{-2}\rho^{-2\beta}\log(2n/\delta))$ of its own queries on the same event of probability at least $1 - \delta$. The combined procedure then stops within twice this many queries, or earlier.

**Why a useful prediction still requires queries.** The lower bound for a fixed-lottery test does not by itself prove optimality for consensus: a consensus algorithm may return a different lottery. We therefore compare instances in which changing one agent makes *every* lottery unacceptable. Fix $n, m \geq 1$, $\rho \in (0, 1/10]$, and $\delta \in (0, 1/4)$, and supply the same arbitrary lottery $\mathbf{x}^{(0)}$ in all instances. In the reference instance, set $\tau_i = 1/2$ and $u_{ij} = 1/2 + \rho$ for every $i, j$. Thus every lottery, including the prediction, has common margin $\rho$.

For each $h \in N$, change only agent $h$'s utilities to $u_{hj} = 1/2 - 2\rho$ for every $j$. This agent now has constant margin $-2\rho$, so the changed instance is a robust NO instance. All parameters remain in their permitted ranges. Use independent Bernoulli replies of mean $1/2 + \lambda\rho^\beta$ for every unchanged agent, and mean $1/2 - \lambda(2\rho)^\beta$ for the changed agent. These are admissible response rules. Their means lie in $[1/4, 3/4]$, and the bound $\mathrm{KL}(\mathrm{Bernoulli}(p) \| \mathrm{Bernoulli}(q)) \leq (p-q)^2/(q(1-q))$ gives one-query divergence at most $C_\beta \lambda^2 \rho^{2\beta}$.

Let $P_0$ and $P_h$ denote the complete-history laws in the reference instance and comparison $h$, and let $T_h$ count queries to agent $h$. An algorithm correct for every prediction must return a lottery with probability at least $1 - \delta$ under $P_0$ and with probability at most $\delta$ under $P_h$. Lemma C.2 therefore gives

$$c_0 \log(1/\delta) \leq \mathrm{KL}(P_0 \| P_h) \leq C_\beta \lambda^2 \rho^{2\beta} \mathbb{E}_{P_0}[T_h].$$

Summing over agents proves the lower bound $\Omega_\beta(n\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$ on the reference instance, where the prediction is useful. Querying or returning a different lottery cannot avoid this cost, because each agent's margin is constant across lotteries.

The same rate applies to a query budget by which the algorithm returns a lottery with probability at least $1 - \delta$ on the reference instance. Truncate both histories at that budget and consider the event that a lottery has been returned. Its probabilities are still at least $1 - \delta$ and at most $\delta$, respectively. The same KL comparison applies, while the sum of the truncated query counts is at most the budget. This gives the claimed lower bound for the high-probability fast rate as well as for expected cost.

## H Exact Consensus with Unknown Noise and Margin

So far, the algorithms have used known parameters to decide how many times to repeat a query and how far to search along an edge. In practice, neither the noise parameters nor a positive lower bound on the relevant margins may be available. Can the algorithm instead gather more replies until it can justify an exact decision? Adapting to unknown noise is also studied in active learning (Wang & Singh, 2016; Locatelli et al., 2018). Our problem adds a requirement: an exact algorithm must justify the absence of a common lottery, as well as find one when possible.

We give a procedure that receives none of $\lambda, \beta, \varepsilon$, or $\rho$. Its stopping behavior depends on the *largest common margin*, defined by $\nu := \max_{\mathbf{x} \in \Delta(S)} \min_{i \in N} g_i(\mathbf{x})$. For a fixed lottery, the minimum is its margin for the least satisfied agent; maximizing this minimum gives the best margin that all agents can achieve together. This same quantity is used to test feasibility from noisy numerical feedback by Gangrade et al. (2024). Here it determines when binary approval answers suffice for an exact decision.

The maximum is attained because the margins are continuous and the simplex is compact. Thus $\nu > 0$ means that some lottery has positive margin for every agent. If $\nu < 0$, every lottery is rejected by at least one agent, so consensus is infeasible. If $\nu = 0$, consensus exists, but every common lottery is on at least one agent's acceptance boundary. The three panels of Figure 1 have $\nu = 0.1, 0, -0.1$, respectively.

**Theorem H.1.** *For every $\delta \in (0, 1/2)$, there is a procedure given only $n, m, \delta$ with the following guarantees under* $\mathsf{NearThresholdNoise}(\lambda, \beta)$*, without finite precision. On every instance, with probability at least $1 - \delta$, any returned lottery belongs to $\bigcap_{i \in N} A_i$, and any* NULL *output means that this intersection is empty. If $\nu \neq 0$, the procedure terminates almost surely. Its expected query complexity is*

$$O_\beta\left(nm^2\lambda^{-2}|\nu|^{-2\beta}\log\frac{2}{|\nu|}\,\log^2\frac{enm}{\delta\lambda|\nu|}\right).$$

*With probability at least $1 - \delta$, it also returns a correct output within this bound.*

For $\nu \neq 0$, the power of $|\nu|$ is the same as the power of $\rho$ in Theorem 4.2: learning without the parameters adds only logarithmic factors. The procedure must, however, meet a stronger condition before returning NULL. Inner approximations may omit acceptable lotteries, so an empty intersection of inner sets does not prove that consensus is infeasible. We therefore construct an inner set $\widehat{A}_i$ containing only accepted lotteries and an outer set $\overline{A}_i$ containing every accepted lottery:

$$\widehat{A}_i \subseteq A_i \subseteq \overline{A}_i.$$

A common point of the inner sets is unanimously acceptable. An empty intersection of the outer sets proves infeasibility. If the inner sets have no common point but the outer sets do, the observations do not yet justify either decision. The algorithm then repeats the construction with a larger query budget.

The exception at $\nu = 0$ is necessary. Even when the noise parameters are known, no procedure can be required to stop reliably on every such instance.

**Proposition H.2.** *Fix $\delta \in (0, 1/2)$. Any procedure satisfying the correctness guarantee of Theorem H.1 on every instance has the following property: for every instance with $\nu = 0$, there is an admissible oracle under which the probability of returning a lottery after finitely many queries is at most $\delta$, and the probability of returning any output after finitely many queries is at most $2\delta$. Consequently, no procedure is $(1 - \delta)$-correct in the sense of Definition 2.1 on the entire class with arbitrary real utilities and thresholds.*

To see the reason, start with an instance whose largest common margin is zero. A small decrease in utilities and increase in thresholds makes it infeasible. Yet the response probabilities can remain arbitrarily close at every query. Any procedure that usually stops after finitely many queries in the original instance would then usually make the same decision in a sufficiently close infeasible instance. Appendix H.3 gives the formal argument.

### H.1 Constructing inner and outer approximations

We first describe one run with a fixed query budget. Fix integers $M, L \geq 1$ and an error allowance $\delta_0 \in (0, 1)$. Each test uses $M$ replies. Each edge search takes $L$ bisection steps, where one step tests the midpoint and keeps one half of the current interval. The run uses at most $B = 4nm^2(L + 1)$ tests in total. Increasing $M$ makes the tests more precise; increasing $L$ shortens the final interval of each edge search.

**A test with an undecided outcome.** Since no lower bound on response bias is known, a test must be allowed to leave the sign undecided. At each queried point, take $M$ replies, let $\overline{Y}$ be their average, and set

$$a_M = \sqrt{\frac{\log(2B/\delta_0)}{2M}}.$$

Return $+$ if $\overline{Y} > 1/2 + a_M$, return $-$ if $\overline{Y} < 1/2 - a_M$, and otherwise return "undecided." The choice of $a_M$ gives the simultaneous deviation bound below. The rule uses neither $\lambda$ nor $\beta$.

Let $\overline{\mu}$ be the average of the conditional response means in this test, each conditioned on all replies before it. The concentration argument in Lemma C.1 gives

$$\Pr(|\overline{Y} - \overline{\mu}| > a_M \mid \text{history before the test}) \leq 2e^{-2Ma_M^2} = \delta_0/B.$$

Adding the error probabilities of at most $B$ tests shows that every test has $|\overline{Y} - \overline{\mu}| \leq a_M$ with probability at least $1 - \delta_0$, also conditional on the history before the run. On this event,

$$+ \implies g_i(\mathbf{x}) \geq 0, \quad - \implies g_i(\mathbf{x}) < 0.$$

Indeed, accepted points have conditional means at least $1/2$, so they cannot receive $-$. Rejected points have conditional means below $1/2$, so they cannot receive $+$.

To describe which margins the test can reliably distinguish, define $\gamma_M = (2a_M/\lambda)^{1/\beta}$ for the proof only. If $g_i(\mathbf{x}) > \gamma_M$, every conditional mean exceeds $1/2 + 2a_M$. Subtracting the possible error $a_M$ still leaves the observed average above the cutoff for $+$, so the test returns $+$. Similarly, margin below $-\gamma_M$ implies output $-$. Therefore

$$\text{not} + \implies g_i(\mathbf{x}) \leq \gamma_M, \quad \text{not} - \implies g_i(\mathbf{x}) \geq -\gamma_M.$$

The procedure never evaluates $\gamma_M$; it is used only to bound the accuracy of the sets it constructs.

**Inner approximation: retain points labelled $+$.** Test every pure lottery. Let $K_i$ be the indices receiving $+$ and $D_i = [m] \setminus K_i$. For each $d \in D_i$ and $k \in K_i$, search the edge from $\mathbf{e}_d$ to $\mathbf{e}_k$ for $L$ bisection steps. At every step, the left endpoint has an outcome other than $+$ and the right endpoint has outcome $+$. A midpoint with outcome $+$ replaces the right endpoint; either other outcome replaces the left endpoint.

Define $\widehat{A}_i$ as the convex hull of the pure lotteries in $K_i$ and the final right endpoints. If no point is retained, set $\widehat{A}_i = \varnothing$.

All retained points are accepted. Since a mixture of accepted lotteries is also accepted, $\widehat{A}_i \subseteq A_i$. At the end of a searched edge, the left endpoint has margin at most $\gamma_M$. The endpoints differ by $2^{-L}$ in edge parameter, and the margin changes at rate $|u_{ik} - u_{id}| \leq 1$ along this parameter. The retained right endpoint therefore has margin at most $\gamma_M + 2^{-L}$.

**Outer approximation: discard only points labelled $-$.** For the outer set, keep every pure lottery that the test has not established to be rejected. Reuse the pure lottery test outcomes and search every edge from a pure lottery labelled $-$ to one not labelled $-$, again for $L$ bisection steps. Keep a left endpoint labelled $-$ and a right endpoint not labelled $-$. A midpoint labelled $-$ replaces the left endpoint; either other outcome replaces the right endpoint.

Define $\overline{A}_i$ as the convex hull of all pure lotteries not labelled $-$ and all final left endpoints. Here we retain the rejected endpoint of each final interval so that no accepted boundary point is excluded. If all pure lotteries are labelled $-$, set $\overline{A}_i = \varnothing$; if none is, set $\overline{A}_i = \Delta(S)$.

To prove $A_i \subseteq \overline{A}_i$, recall the geometric fact in Appendix F.3: every vertex of $A_i$ is either an accepted pure lottery or a point where the acceptance boundary meets a simplex edge. Every accepted pure lottery is retained because its label cannot be $-$.

For a non-pure boundary vertex, the two endpoint margins have opposite signs, by the vertex argument in Appendix F.3 applied at level zero. Orient the edge from its rejected endpoint to its accepted endpoint. If the rejected endpoint is not labelled $-$, both endpoints are retained. Otherwise, this edge is searched. Its margin is strictly increasing, and the final left endpoint is still strictly rejected; the boundary vertex therefore lies between this point and the original accepted endpoint. Both are retained. Thus every vertex of $A_i$ belongs to $\overline{A}_i$, and hence $A_i \subseteq \overline{A}_i$. If all pure lotteries are labelled $-$, all pure margins are strictly negative, so $A_i = \varnothing$ as required.

The outer set contains only lotteries whose margins are at least $-\gamma_M - 2^{-L}$. To see this, each retained pure lottery has margin at least $-\gamma_M$. On a searched edge, the final right endpoint also has margin at least $-\gamma_M$, even though it need not be accepted. The final left endpoint is at most $2^{-L}$ away in edge parameter, so its margin is at least $-\gamma_M - 2^{-L}$. The margin of a mixture is the corresponding average of the margins, which gives the same lower bound throughout the outer convex hull.

**Set inclusions and the stopping decision.** Whenever all tests satisfy their deviation bounds, we have $\widehat{A}_i \subseteq A_i \subseteq \overline{A}_i$, for every $M, L$. Put $s = \gamma_M + 2^{-L}$, adding the uncertainty from the tests and the remaining edge-interval length. When $s \leq 1$, the more precise inclusions are

$$A_i^{+s} \subseteq \widehat{A}_i \subseteq A_i \subseteq \overline{A}_i \subseteq A_i^{-s}. \tag{11}$$

For the first inclusion, apply the vertex argument of Appendix F.3 at level $s$. A pure vertex has margin at least $s > \gamma_M$, so it is labelled $+$ and retained. A non-pure vertex has margin $s$ and lies on an edge with endpoint margins strictly below and strictly above $s$. The higher-margin endpoint is labelled $+$. If both endpoints are labelled $+$, their segment is retained. Otherwise, the inner search returns an accepted point of margin at most $s$ on this increasing edge. The level-$s$ vertex lies between that point and the retained higher-margin endpoint, including when the returned point itself has margin $s$. This proves $A_i^{+s} \subseteq \widehat{A}_i$. If $K_i = \varnothing$, all pure margins are at most $\gamma_M < s$, so $A_i^{+s} = \varnothing$ and the inclusion is again valid.

The stopping rule now follows directly from the set inclusions. Return a point in $\bigcap_i \widehat{A}_i$ if one exists. Otherwise, return NULL only if $\bigcap_i \overline{A}_i$ is empty. If neither condition holds, make no decision and continue to another run. Either returned output is correct with probability at least $1 - \delta_0$,

conditionally on the history before the run. Both intersection checks are linear feasibility problems of the form in Appendix F.4.

There are at most $nm + 2nm^2L \leq B$ tests: one per pure lottery and at most two searches per agent and pair of alternatives. Each test uses $M$ replies, giving at most $BM = O(nm^2ML)$ queries in the run.

### H.2 Increasing repetitions and search depth

*Proof of Theorem H.1.* **The schedule uses no unknown parameters.** The procedure must eventually increase both the number of replies per test and the depth of each edge search. It does not know which requires the larger share of the query budget. At stage $r = 0, 1, 2, \ldots$, run the preceding construction for each $j \in \{0, \ldots, r\}$ with $M = 2^j$, $L = 2^{r-j}$, and $\delta_{r,j} = \frac{\delta}{4(r+1)4^r}$. Each run uses new replies, and the procedure stops at the first output. Every choice has $ML = 2^r$, but a different balance between repetitions and search depth. Across the $r + 1$ runs, the number of replies per test ranges from 1 to $2^r$, while the depth ranges from $2^r$ to 1. Trying all these choices avoids needing to know the right balance in advance.

The sum of the error allowances is

$$\sum_{r\geq 0}\sum_{j=0}^{r}\delta_{r,j} = \frac{\delta}{4}\sum_{r\geq 0}4^{-r} = \frac{\delta}{3}.$$

Thus, with probability at least $1 - \delta/3$, all tests in every executed run meet their deviation bounds, and every returned output is correct. This conclusion also applies when $\nu = 0$: the procedure may continue indefinitely, but the probability of an incorrect output is still at most $\delta/3$.

Since $L \geq 1$, a run costs at most $BM = 4nm^2(L + 1)M \leq 8nm^2ML = 8nm^2 2^r$ queries. Summing the costs of all $r + 1$ runs and all preceding stages gives at most $16nm^2(r + 1)2^r$ queries through stage $r$.

**A sufficiently accurate run decides every instance with $\nu \neq 0$.** For the analysis only, fix $\nu \neq 0$ and let $L_* := \lceil \log_2 \frac{4}{|\nu|} \rceil$. This depth makes the remaining edge-interval length at most $|\nu|/4$. Once $r \geq \lceil \log_2 L_* \rceil$, stage $r$ includes the choice $j = r - \lceil \log_2 L_* \rceil$. Its depth lies in $[L_*, 2L_*)$, and it uses at least $2^r/(2L_*)$ replies per test. For this run,

$$\log(2B/\delta_{r,j}) \leq C(\log(enm/\delta) + r).$$

Consequently, for a sufficiently large constant depending only on $\beta$, the condition

$$2^r \geq C_\beta \lambda^{-2}|\nu|^{-2\beta}L_*(\log(enm/\delta) + r) \tag{12}$$

ensures $\gamma_M \leq |\nu|/4$. Combining this test accuracy with $2^{-L} \leq |\nu|/4$ gives $s \leq |\nu|/2$ in (11).

If $\nu > 0$, a lottery achieving the largest common margin belongs to every $A_i^{+s} \subseteq \widehat{A}_i$. The inner intersection is then nonempty, and the run returns a common lottery. If $\nu < 0$, the outer intersection is contained in $\bigcap_i A_i^{-s}$. This latter intersection is empty: any lottery in it would have common margin at least $-s > \nu$, contradicting the definition of $\nu$. The run then returns NULL. Thus either sign of $\nu$ leads to a decision whenever this run's tests meet their deviation bounds.

There is a stage $r_*$ satisfying (12) for which

$$2^{r_*} \leq C_\beta \lambda^{-2}|\nu|^{-2\beta}L_* \log\frac{enm}{\delta\lambda|\nu|}, \quad \text{and} \quad r_* + 1 \leq C_\beta \log\frac{enm}{\delta\lambda|\nu|}.$$

To see this, choose $r_*$ by rounding up the base-two logarithm of $\lambda^{-2}|\nu|^{-2\beta}L_* \log(enm/(\delta\lambda|\nu|))$ and adding a sufficiently large integer depending only on $\beta$. Since $L_* = O(\log(2/|\nu|))$, that logarithm is $O_\beta(\log(enm/(\delta\lambda|\nu|)))$. Thus $r_*$ has the second displayed bound and is at least $\lceil \log_2 L_* \rceil$. The added integer multiplies $2^{r_*}$ by a constant, whereas the right-hand side of (12) grows only linearly with $r_*$. A sufficiently large choice therefore gives that condition and the first displayed bound simultaneously.

Condition (12) continues to hold at larger stages because $2^r/(\log(enm/\delta) + r)$ increases with $r$. Substituting the bounds for $r_*$ into the cost through stage $r_*$, and using $L_* = O(\log(2/|\nu|))$, gives

the theorem's query bound whenever all tests meet their deviation bounds. That event has probability at least $1 - \delta/3$, and hence at least $1 - \delta$.

**Almost-sure termination and expected cost.** A high-probability query bound alone does not imply finite expected cost, since the remaining runs could be arbitrarily long. We must therefore also bound the probability of continuing past an accurate stage. Fix $r \geq r_*$ and its choice $j = r - \lceil \log_2 L_* \rceil$ above. If this run is reached, condition on the full history before it. With conditional probability at least $1 - \delta_{r,j}$, all its deviation bounds hold and it returns an output, regardless of the results of earlier runs. Therefore reaching the end of stage $r$ without an output has probability at most $\delta_{r,j}$, giving

$$\Pr(\text{not stopped by the end of stage } r) \leq \frac{\delta}{4(r+1)4^r}.$$

This probability tends to zero, proving almost-sure termination. It also decreases faster than the cost of the next stage increases. A stage $r > r_*$ is reached only if the preceding stage ended without an output. Multiplying the cost of stage $r$ by the displayed bound for reaching it shows that the total expected cost after stage $r_*$ is at most

$$\sum_{r>r_*} Cnm^2(r+1)2^r \frac{\delta}{4r4^{r-1}} \leq C'nm^2\delta \sum_{r>r_*} 2^{-r} \leq C''nm^2.$$

Adding the deterministic cost through stage $r_*$ proves the expected bound. None of these arguments requires independence between responses or runs. □

### H.3 Impossibility when the common margin is zero

*Proof of Proposition H.2.* We compare any instance with $\nu = 0$ to infeasible instances obtained by arbitrarily small changes in its parameters. The response distributions will approach those of the original instance, so any event determined by finitely many queries must have nearly the same probability.

Fix an instance with $\nu = 0$. Give it Bernoulli replies with means

$$p_i(\mathbf{x}) = \frac{1}{2} + \lambda \operatorname{sgn}(g_i(\mathbf{x}))|g_i(\mathbf{x})|^\beta, \quad \operatorname{sgn}(0) = 0,$$

using independent randomness at each query. Here $\operatorname{sgn}(z)$ is 1 for $z > 0$, $-1$ for $z < 0$, and 0 at zero. These means lie in $[1/4, 3/4]$ and satisfy Definition 2.4.

For $h \in (0, 1)$, change the parameters to

$$u_i^{(h)} = (1-h)u_i, \quad \tau_i^{(h)} = (1-h)\tau_i + h.$$

They remain in the permitted ranges. Their margins are $g_i^{(h)}(\mathbf{x}) = (1-h)g_i(\mathbf{x}) - h$, so the largest common margin becomes $(1-h)\nu - h = -h$. Thus every changed instance is infeasible. Use the same formula for its response means, now with $g_i^{(h)}$.

The function $z \mapsto \operatorname{sgn}(z)|z|^\beta$ is uniformly continuous on $[-1, 1]$, and

$$\sup_{i,\mathbf{x}} |g_i^{(h)}(\mathbf{x}) - g_i(\mathbf{x})| \leq 2h.$$

The response means therefore approach the original ones uniformly as $h \downarrow 0$. To transfer this to adaptive queries, fix $M$ and give the two executions the same private random seed. While their histories agree, they make the same next query. Generate their Bernoulli replies together using the same uniform random number, so that the probability of different replies is the absolute difference of their means. The probability of any difference in the first $M$ queries is then at most $M$ times the maximum difference of the means, which tends to zero. If both executions have already stopped, append identical terminal symbols.

Every changed instance is infeasible, so returning a lottery within $M$ queries is incorrect and has probability at most $\delta$. Letting $h \downarrow 0$ transfers this bound to the original instance. Then letting $M \to \infty$ shows that the probability of returning any lottery after finitely many queries is at most $\delta$.

The original instance is feasible because the maximum $\nu = 0$ is attained. Returning NULL there is incorrect and also has probability at most $\delta$. The probability of returning any output after finitely many queries is therefore at most $2\delta$. Exact $(1-\delta)$-correctness would require a correct finite lottery output with probability at least $1 - \delta > \delta$, which is impossible. □

# I KNOWN RESPONSE RULE

Under near-threshold noise, the same response mean may correspond to several different margins. Knowing $\lambda$ and $\beta$ gives bounds on the margin but need not determine its value. We now ask what becomes possible when the entire relation between the margin and the response mean is known. Specifically, assume that $\lambda, \beta$ are known and

$$\mathbb{E}[Y_t \mid \mathcal{F}_{t-1}] = \frac{1}{2} + \lambda \operatorname{sgn}(g_{i_t}(\mathbf{x}_t))|g_{i_t}(\mathbf{x}_t)|^{\beta}, \quad \operatorname{sgn}(0) = 0. \tag{13}$$

Recovering utility differences from a specified response function is standard in stochastic choice models (Train, 2009); Dang et al. (2026) also use a known comparison model to estimate preferences for fair division. Here the equality lets us estimate numerical margins and compare query rates with those under the general noise condition. The response mean at $\mathbf{e}_j$ determines $g_i(\mathbf{e}_j)$. Moreover, every lottery's margin is a weighted average of its pure lottery margins: $g_i(\mathbf{x}) = \sum_j x_j g_i(\mathbf{e}_j)$. Because the weights are nonnegative and sum to one, estimating all pure lottery margins to accuracy $\rho/4$ estimates every lottery's margin to that same accuracy. Only $nm$ different agent-lottery pairs need to be queried.

For example, with $\beta = 1$ and $\lambda = 1/4$, response mean $0.6$ at $\mathbf{e}_j$ implies $g_i(\mathbf{e}_j) = 0.4$ under (13). Under Definition 2.4 alone, the same mean is compatible with any margin between $0$ and $0.4$. The extra information comes from the equality, not merely from knowing the noise parameters.

**Theorem I.1.** *Under* (13)*, for every $\rho \in (0, 1]$ and $\delta \in (0, 1/4)$, $\rho$-robust consensus can be solved using $O_\beta(nm\lambda^{-2}\rho^{-2\max\{1,\beta\}}\log(nm/\delta))$ queries. The algorithm learns sets $\widehat{A}_i, \overline{A}_i$, each defined by one linear inequality on $\Delta(S)$, such that, with probability at least $1-\delta$, simultaneously for all $i$, $A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i \subseteq \overline{A}_i \subseteq A_i^{-\rho}$. On the same event, every returned lottery has margin at least $\rho/4$ for every agent. For $\beta \geq 1$ and $\rho \leq 1/4$, every algorithm that learns inner approximations satisfying $\Pr(\forall i : A_i^{+\rho} \subseteq \widehat{A}_i \subseteq A_i) \geq 1-\delta$ on every instance requires $\Omega_\beta(nm\lambda^{-2}\rho^{-2\beta}\log(1/\delta))$ expected queries in the worst case, regardless of the descriptions of its sets. For each fixed lottery, membership in a returned set must be measurable from the completed queries, responses, and algorithm randomness, without further oracle calls.*

For $\beta \geq 1$, learning the inner sets now has linear dependence on $m$, compared with the quadratic dependence required under the general noise model. The power of $\rho$ is unchanged. The returned lottery also has margin at least $\rho/4$ for every agent, whereas under the general model no fixed positive fraction of $\rho$ can be guaranteed (Section 4). As before, the lower bound here concerns learning every agent's inner set; it does not establish the cost of finding just one common lottery. This known-response lower bound does not assume that a common lottery with positive margin is supplied to the algorithm.

**Theorem I.2.** *Under* (13) *and $\varepsilon$-precision, for $\varepsilon \leq 1/4$ and $\delta \in (0, 1/4)$, there is a $(1-\delta)$-correct exact consensus algorithm using $O_\beta(nm\lambda^{-2}\varepsilon^{-2\max\{1,\beta\}}\log(nm/\delta))$ queries. Already for $n = m = 2$, every such algorithm has worst-case expected query complexity at least*

$$\Omega_\beta\left(\lambda^{-2}\log(1/\delta)\cdot\begin{cases}\varepsilon^{-4\beta}, & 0<\beta<1/2,\\ \varepsilon^{-2}, & 1/2\leq\beta\leq 1,\\ \varepsilon^{-2\beta}, & \beta>1.\end{cases}\right).$$

Finite precision turns approximate margin estimates into exact ones. Each pure lottery margin is a multiple of $\varepsilon$. Once its response mean is estimated accurately enough, the known response rule identifies which multiple it is. We then have all of the exact linear acceptance constraints and can check whether they admit a common lottery.

Since $\beta$ is known, we can choose in advance between this algorithm and the one in Theorem 3.2. Taking the smaller stated query bound matches all three powers of $\varepsilon$ in the lower bound. For fixed $n = m = 2$ and $\beta \geq 1/2$, the bounds also match in $\lambda$ and $\delta$; for $\beta < 1/2$, the difference is logarithmic in the precision and confidence parameters. In particular, at $\beta = 1$, the optimal precision dependence improves from $\varepsilon^{-4}$ under the general model to $\varepsilon^{-2}$ under the known response rule.

### I.1 Proof of Theorem I.1

*Proof.* **Estimate the pure lottery margins.** The first step is to translate an error in the estimated response mean into an error in the estimated margin. Subtracting $1/2$ from a response mean and dividing by $\lambda$ gives $\operatorname{sgn}(g_i(\mathbf{e}_j))|g_i(\mathbf{e}_j)|^\beta$. Its inverse is $z \mapsto \operatorname{sgn}(z)|z|^{1/\beta}$ on $[-1,1]$. For $0 < \beta \le 1$, its derivative on $[-1,1]$ has absolute value at most $1/\beta$, so the inverse changes by at most $|z-z'|/\beta$. For $\beta \ge 1$, it satisfies

$$|\operatorname{sgn}(z)|z|^{1/\beta} - \operatorname{sgn}(z')|z'|^{1/\beta}| \le 2^{1-1/\beta}|z-z'|^{1/\beta}.$$

For inputs of the same sign, the latter follows from $(a+b)^{1/\beta} \le a^{1/\beta} + b^{1/\beta}$ for $a, b \ge 0$. For opposite signs, use $a^{1/\beta} + b^{1/\beta} \le 2^{1-1/\beta}(a+b)^{1/\beta}$.

Take $\lceil C_\beta \lambda^{-2} \rho^{-2\max\{1,\beta\}} \log(2nm/\delta) \rceil$ replies for every agent at every pure lottery. For sufficiently large $C_\beta$, conditional concentration and a union bound make all $nm$ empirical means accurate to $c_\beta \lambda \rho^{\max\{1,\beta\}}$ with probability at least $1-\delta$. Here $c_\beta > 0$ is chosen small enough that the inverse bounds above give margin error at most $\rho/4$.

To form each estimate, subtract $1/2$ from the empirical mean, divide by $\lambda$, and replace any value outside $[-1,1]$ by the nearer endpoint. This replacement cannot increase its distance from the true rescaled mean. Apply the inverse function to obtain estimates satisfying

$$\max_{i,j} |\widehat{g}_i(\mathbf{e}_j) - g_i(\mathbf{e}_j)| \le \rho/4.$$

The concentration bound uses the known conditional mean in (13), so independence need not be assumed separately.

**Extend the estimates to all lotteries.** Define $\widehat{g}_i(\mathbf{x}) = \sum_j x_j \widehat{g}_i(\mathbf{e}_j)$. On the same event,

$$|\widehat{g}_i(\mathbf{x}) - g_i(\mathbf{x})| \le \sum_j x_j |\widehat{g}_i(\mathbf{e}_j) - g_i(\mathbf{e}_j)| \le \rho/4 \quad \text{for all } \mathbf{x} \in \Delta(S).$$

We use an offset of $\rho/2$ on either side of the estimated acceptance boundary to define

$$\widehat{A}_i = \{\mathbf{x} \in \Delta(S) : \widehat{g}_i(\mathbf{x}) \ge \rho/2\}, \quad \overline{A}_i = \{\mathbf{x} \in \Delta(S) : \widehat{g}_i(\mathbf{x}) \ge -\rho/2\}.$$

If $g_i(\mathbf{x}) \ge \rho$, then $\widehat{g}_i(\mathbf{x}) \ge 3\rho/4$, so $\mathbf{x} \in \widehat{A}_i$. Conversely, $\mathbf{x} \in \widehat{A}_i$ implies $g_i(\mathbf{x}) \ge \rho/4$. If $g_i(\mathbf{x}) \ge 0$, then $\widehat{g}_i(\mathbf{x}) \ge -\rho/4$, so $\mathbf{x} \in \overline{A}_i$. Finally, $\mathbf{x} \in \overline{A}_i$ implies $g_i(\mathbf{x}) \ge -3\rho/4 \ge -\rho$. These four implications prove all the stated inclusions. Return a point in $\bigcap_i \widehat{A}_i$ if it is nonempty, and NULL otherwise. A robust YES instance has a common lottery with margin at least $\rho$, which belongs to every inner set. Every returned lottery has common margin at least $\rho/4$, so no lottery can be returned on a robust NO instance. This proves robust correctness and the claimed output margin.

**Lower bound for learning the inner sets.** We use one comparison for each of the $nm$ agent-alternative pairs. The reference instance requires every pure lottery to belong to every returned set. Comparison $(i,j)$ requires $\mathbf{e}_j$ to be omitted from agent $i$'s set. We then bound how much a single query can reveal across these comparisons.

In the reference instance, give every agent constant margin $\rho$. For each $(i,j) \in N \times [m]$, change only agent $i$ to margin $\rho(1-2x_j)$. These instances have threshold $1/2$ and utilities $1/2+\rho$ in every coordinate, except for utility $1/2-\rho$ in coordinate $j$ of the changed agent. They are valid for $\rho \le 1/4$.

Let $P_0, P_{ij}$ be the distributions of the complete query histories, using independent Bernoulli replies with the means in (13), and let $T$ be the total query count. In the reference instance $A_i^{+\rho} = \Delta(S)$, whereas in comparison $(i,j)$, agent $i$ rejects $\mathbf{e}_j$. The measurable event of returning sets with $\mathbf{e}_j \in \widehat{A}_i$ therefore has probability at least $1-\delta$ under $P_0$ and at most $\delta$ under $P_{ij}$. Lemma C.2(i) gives

$$\mathrm{KL}(P_0 \| P_{ij}) \ge c_0 \log(1/\delta).$$

For $\beta \ge 1$, $z \mapsto \operatorname{sgn}(z)|z|^\beta$ changes by at most $\beta|z-z'|$ on $[-1,1]$. Factor $\rho^\beta$ out of the two response biases and apply this bound to 1 and $1-2x_j$. At a query $\mathbf{x}$ to agent $i$, the means differ

by at most $2\beta\lambda\rho^\beta x_j$. Thus a query that puts little probability on alternative $j$ reveals little about comparison $(i, j)$.

All means lie in $[1/4, 3/4]$, so the sum of the one-query divergences over $j$ is at most

$$C_\beta \lambda^2 \rho^{2\beta} \sum_j x_j^2 \le C_\beta \lambda^2 \rho^{2\beta}.$$

A query to another agent contributes nothing to these comparisons. Summing the adaptive KL chain rule over all $nm$ comparison instances gives

$$c_0 nm \log(1/\delta) \le \sum_{i,j} \mathrm{KL}(P_0 \| P_{ij}) \le C_\beta \lambda^2 \rho^{2\beta} \mathbb{E}_{P_0}[T].$$

If the number of queries is unbounded, first apply the calculation through a fixed query count and then let that count increase, as in Lemma C.2. Rearranging proves the lower bound. If the expected query count under the reference instance is infinite, the lower bound holds immediately. □

### I.2 Proof of Theorem I.2

*Proof.* **Recover the margins from their finite grid.** Under $\varepsilon$-precision, every $g_i(\mathbf{e}_j)$ belongs to $\varepsilon\mathbb{Z} \cap [-1, 1]$, the multiples of $\varepsilon$ between $-1$ and 1. Each candidate margin has a different known response mean. To identify the true margin, it is enough to estimate that mean more accurately than the gap between neighboring candidates. The known response means corresponding to successive grid values are separated by at least

$$c_\beta \lambda \varepsilon^{\max\{1,\beta\}}.$$

For $0 < \beta \le 1$, the derivative of $\mathrm{sgn}(z)|z|^\beta$ away from zero is $\beta|z|^{\beta-1} \ge \beta$ on $[-1, 1]$. A step from or to zero changes the value by $\varepsilon^\beta \ge \varepsilon$. For $\beta \ge 1$, the smallest successive difference on either side occurs next to zero and equals $\varepsilon^\beta$.

Estimate the response mean for each agent at every pure lottery to less than one third of this minimum separation, and choose the nearest candidate mean. The true candidate is then the unique nearest one. Conditional concentration and a union bound recover all pure margins exactly with probability at least $1 - \delta$, within the stated query count. The true acceptance constraints are then

$$\sum_j g_i(\mathbf{e}_j) x_j \ge 0 \quad \text{for all } i \in N.$$

Exact linear feasibility on $\Delta(S)$ solves consensus. When $\beta < 1/2$, Theorem 3.2 also applies under the known response rule and gives the smaller precision exponent $4\beta$.

**Lower bound for $0 < \beta \le 1$.** We compare an instance with exactly one common lottery to an infeasible instance obtained by slightly reducing one utility. Unlike in the general noise model, we must use the response means imposed by (13), so we bound their difference at every possible query.

Use two agents and two alternatives. Agent 1 has utilities $(1, 0)$ and threshold $1 - \varepsilon$, hence margin $\varepsilon - \alpha$ at $\mathbf{x}_{1,2}(\alpha)$. Agent 2 has threshold $\varepsilon$ and utilities either $(0, 1)$ or $(0, 1 - \varepsilon)$. Its two margins are

$$g_2^0(\mathbf{x}_{1,2}(\alpha)) = \alpha - \varepsilon, \quad g_2^1(\mathbf{x}_{1,2}(\alpha)) = (1 - \varepsilon)\alpha - \varepsilon.$$

All parameters satisfy $\varepsilon$-precision. Agent 1 accepts exactly when $\alpha \le \varepsilon$. In the first instance, agent 2 accepts exactly when $\alpha \ge \varepsilon$, leaving the unique common lottery $\mathbf{x}_{1,2}(\varepsilon)$. In the second, agent 2 requires $\alpha \ge \varepsilon/(1 - \varepsilon) > \varepsilon$, so consensus is infeasible. Let $p_0(\alpha), p_1(\alpha)$ be agent 2's response means under (13).

Let $T$ be the total query count and $P_0$ the complete-history law of the feasible instance. Use independent Bernoulli replies with the prescribed means in both instances. Split the queried edge into two intervals. Both turning points belong to $[0, 2\varepsilon]$ because $\varepsilon/(1 - \varepsilon) \le 2\varepsilon$. Within this interval, the margins differ by $\varepsilon\alpha \le 2\varepsilon^2$. For $0 < \beta \le 1$,

$$|\mathrm{sgn}(a)|a|^\beta - \mathrm{sgn}(b)|b|^\beta| \le 2^{1-\beta}|a - b|^\beta.$$

The same equal-sign and opposite-sign inequalities used in the preceding proof give this bound. Hence $|p_0(\alpha) - p_1(\alpha)| \le C_\beta \lambda \varepsilon^{2\beta}$ in this interval.

For $\alpha \geq 2\varepsilon$, both margins are positive and $g_2^1(\mathbf{x}_{1,2}(\alpha)) \geq (1/2 - \varepsilon)\alpha \geq \alpha/4$. The mean value theorem, using $\beta - 1 \leq 0$ and $\alpha \leq 1$, gives

$$|p_0(\alpha) - p_1(\alpha)| \leq \lambda\beta(\alpha/4)^{\beta-1}\varepsilon\alpha \leq C_\beta\lambda\varepsilon.$$

Thus, over all possible queries, the mean difference is at most $C_\beta\lambda\max\{\varepsilon, \varepsilon^{2\beta}\}$. The means lie in $[1/4, 3/4]$, so each query to agent 2 has KL divergence at most $C_\beta\lambda^2\varepsilon^{2\min\{1,2\beta\}}$. Queries to agent 1 contribute zero because its parameters and response means are unchanged. Returning a lottery after finitely many queries has probability at least $1 - \delta$ in the feasible instance and at most $\delta$ in the infeasible one. Lemma C.2 therefore gives

$$\mathbb{E}_{P_0}[T] \geq c_\beta\lambda^{-2}\varepsilon^{-2\min\{1,2\beta\}}\log(1/\delta).$$

If $\beta < 1/2$, the exponent $2\min\{1, 2\beta\}$ is $4\beta$; if $1/2 \leq \beta \leq 1$, it is 2. This proves the first two cases of the theorem.

**Lower bound for $\beta > 1$.** Here a simpler pair suffices. Set $u_{ij} = \varepsilon$ for both agents and both alternatives, and set $\tau_i = \varepsilon$. Every margin is zero, so all lotteries are accepted. In the comparison instance, change one agent's utilities to zero, keeping its threshold $\varepsilon$. That agent then has margin $-\varepsilon$ everywhere, making consensus infeasible. Both instances satisfy $\varepsilon$-precision.

For the changed agent, the two response means are $1/2$ and $1/2 - \lambda\varepsilon^\beta$, giving KL divergence at most $C\lambda^2\varepsilon^{2\beta}$ per query. The other agent's replies have the same distribution in both instances. Apply Lemma C.2 to the event of returning a lottery after finitely many queries, as above, to obtain $\Omega(\lambda^{-2}\varepsilon^{-2\beta}\log(1/\delta))$ expected queries. □