Bayesian Persuasion with a Risk-Conscious Receiver

author names withheld

Under Review for NExT-Game 2026

Abstract

We study Bayesian persuasion with a receiver who evaluates actions by Conditional Value-at-Risk. CVaR captures settings in which rare adverse outcomes matter for acceptance, including automated alerts, financial advice, and safety-critical AI-assisted decisions. The receiver's value is nonlinear in posterior beliefs, so direct recommendation by action fails: merging two signals that recommend the same action can change the relevant tail event and violate incentive compatibility. Our main result shows that this failure does not make the explicit finite-state problem hard. CVaR has a finite active-threshold representation, and signals can be refined by both the recommended action and the active CVaR facet. This refined revelation principle turns the sender's problem into an exact polynomial-size linear program. We then develop a posterior discretization scheme for finite-precision implementation and large-scale extensions, including a margin condition under which approximate incentive compatibility becomes strict incentive compatibility.

1. Introduction

Bayesian persuasion studies how an informed sender commits to an information policy in order to influence a receiver's action (Kamenica and Gentzkow, 2011). In the classical model, the receiver is risk-neutral. Her payoff from an action is linear in the posterior belief, and this linearity supports the usual revelation principle, concavification, and linear-programming formulations (Bergemann and Morris, 2019; Xu et al., 2015; Dughmi and Xu, 2021).

Many current information-design problems do not fit this premise. A clinician who sees an automated triage signal, a regulator who evaluates an AI release, or an operator who receives a security alert may care more about a small posterior tail of failures than about average performance. In such settings, an action can look good in expectation while remaining unacceptable because the disclosed information leaves too much downside risk. This makes tail-sensitive persuasion a natural topic for game-theoretic learning and AI-mediated decision systems.

We model this behavior through Conditional Value-at-Risk (CVaR), a standard tail-risk measure (Artzner et al., 1999; Rockafellar and Uryasev, 2000, 2002; Acerbi and Tasche, 2002). Under CVaR, the receiver's utility is nonlinear in the posterior. This breaks the classical action-based direct-recommendation argument. If two signals separately induce the same action, their merged posterior may have a different worst-tail distribution and may no longer induce that action.

At first sight, this suggests a difficult nonconvex persuasion problem. The main message of this paper is more precise. Action-based merging fails, but CVaR has a finite active-threshold structure. On a finite state space, the CVaR value of each action is the maximum of finitely many affine functions of the posterior, one for each possible shortfall threshold. If a signal is labeled not only by the recommended action but also by the active CVaR facet that supports that recommendation, incentive compatibility becomes a family of linear inequalities.

This gives an active-facet revelation principle for CVaR persuasion. The sender can refine signals by action-facet pairs without changing value. Each refined incentive region is a polytope, and the sender's objective is linear in the joint distribution of states and refined recommendations. Consequently, the explicit finite-state CVaR persuasion problem admits an exact polynomial-size linear program. If $m = |\Omega|$, $A = |\mathcal{A}|$, and each action has at most m CVaR facets, the LP has $O(Am^2)$ variables and $O(A^2m^2)$ incentive constraints.

This distinction matters for both theory and modeling. The nonlinearity of CVaR does not disappear, and the receiver's best-response regions are not the same as in the expected-utility model. What changes is the correct message space. Once the signal label records the active tail threshold, the sender no longer needs to reason over arbitrary nonconvex unions of posterior regions. The model therefore preserves the main computational appeal of Bayesian persuasion while making room for receivers who reject actions because of tail losses rather than average payoffs.

We then develop posterior discretization from this perspective. In the explicit finite-state CVaR model, discretization is not needed for computational tractability, since the active-facet LP is exact. Its role is finite-precision implementation and preparation for settings where the refined LP cannot be enumerated directly, such as oracle-described risk functionals or very large state and action spaces. The same approximation argument also identifies the stability condition needed for exact implementation: if the benchmark recommendations stay a positive margin away from receiver indifference, a margin-filtered discretization yields strict incentive compatibility.

1.1. Related Work

Our work connects Bayesian persuasion (Kamenica and Gentzkow, 2011; Bergemann and Morris, 2019) with algorithmic information design and risk-sensitive decision-making. Computational results for persuasion often rely on expected-utility linearity (Xu et al., 2015; Dughmi and Xu, 2021). Work on risk-conscious persuasion develops geometric tools for nonlinear receiver preferences (Anunrojwong et al., 2024). We show that for CVaR, the finite active-threshold representation yields a refined revelation principle and an exact LP. Posterior discretization uses the same sampling idea behind sparse approximations in large games (Althöfer, 1994; Lipton et al., 2003), but here it approximates the affine pieces of the CVaR envelope rather than expected payoff statistics.

2. Model

There is a finite state space $\Omega = \{\omega_1, \dots, \omega_m\}$, a common prior $\mu_0 \in \mathrm{int}(\Delta(\Omega))$, and a finite action set $\mathcal{A} = \{a_1, \dots, a_A\}$. The sender observes the state and commits to a signaling scheme $\pi : \Omega \to \Delta(\mathcal{S})$. After signal s, the receiver forms posterior

$$\mu_s(\omega) = \frac{\pi(s \mid \omega)\mu_0(\omega)}{\sum_{\omega' \in \Omega} \pi(s \mid \omega')\mu_0(\omega')}. \tag{1}$$

Equivalently, a signaling scheme induces a distribution over posteriors whose mean is μ_0.

The receiver evaluates action a at posterior μ through a risk functional $\rho(\mu, a)$ and chooses an action in $\arg\max_{a\in\mathcal{A}} \rho(\mu, a)$. The sender has payoff $v : \Omega\times\mathcal{A} \to \mathbb{R}$ and maximizes expected payoff subject to the receiver's incentive constraint. In this paper,

$$\rho(\mu, a) = \mathrm{CVaR}_r^\mu(u(\cdot, a)) = \sup_{b\in\mathbb{R}} \left\{ b - \frac{1}{r}\mathbb{E}_{\omega\sim\mu}\big[(b - u(\omega, a))^+\big] \right\}, \tag{2}$$

where $r \in (0, 1)$ is fixed. This reward-side CVaR evaluates the average utility in the worst r-tail. A smaller r describes a more cautious receiver.

For arbitrary nonlinear risk functionals, direct recommendations are not without loss. The same action may be optimal in two posterior regions but not at their convex combination. We therefore work with posterior-based signaling. The following support bound is standard and is used only to justify finite-support benchmarks.

Lemma 1 (Finite posterior support) *For any finite-state persuasion problem with any receiver risk functional $\rho(\mu, a)$, there exists an optimal signaling scheme with at most $m = |\Omega|$ supported posteriors.*

The proof is the usual Carathéodory support reduction for Bayes-plausible distributions over posteriors and appears in Appendix A. The sharper structure for CVaR comes from active facets, not from this support bound.

3. Exact Optimization Under CVaR

For each action a, let $f_a(\mu) = \mathrm{CVaR}_r^\mu(u(\cdot, a))$. CVaR has a finite polyhedral representation on a finite state space. If $M_R = \max_{\omega,a} |u(\omega, a)|$, then there is a set $\mathcal{L}_a$ with $|\mathcal{L}_a| \le m$ and vectors $c_{a,\ell} \in \mathbb{R}^m$ such that

$$f_a(\mu) = \max_{\ell\in\mathcal{L}_a}\langle c_{a,\ell}, \mu\rangle, \qquad \|c_{a,\ell}\|_\infty \le M_R\left(1 + \frac{2}{r}\right). \tag{3}$$

The proof is in Appendix B. Each facet corresponds to one threshold in the variational formula for CVaR.

For every action-facet pair (a, ℓ), define

$$P_{a,\ell} = \left\{\mu \in \Delta(\Omega) : \langle c_{a,\ell}, \mu\rangle \ge \langle c_{a',\ell'}, \mu\rangle \quad \forall a' \in \mathcal{A},\ \forall \ell' \in \mathcal{L}_{a'}\right\}. \tag{4}$$

If $\mu \in P_{a,\ell}$, then facet ℓ supports the CVaR value of action a, and action a is a receiver best response. Since all comparisons are linear, $P_{a,\ell}$ is a polytope.

Proposition 2 (Facet revelation) *For every IC signaling scheme, there is another IC scheme with the same sender value whose signals are indexed by pairs (a, ℓ). Here a is the recommended action and $\ell \in \mathcal{L}_a$ is an active CVaR facet at the induced posterior.*

The proof is in Appendix C. The point is not that signals with the same action can be merged. They may not be mergeable. The point is that once the active facet is included in the signal label, the relevant incentive region is convex and linear.

Theorem 3 (Exact LP) *Consider a finite-state persuasion problem with one CVaR receiver, finite action set $\mathcal{A}$, explicit payoffs, prior μ_0, and rational risk level $r \in (0, 1)$. Let*

$$L = \sum_{a \in \mathcal{A}} |\mathcal{L}_a|.$$

An optimal incentive-compatible signaling scheme can be computed by a linear program with mL nonnegative joint-mass variables and at most $m + L^2$ linear constraints, apart from nonnegativity. Since $L \le Am$, this is $O(Am^2)$ variables and $O(A^2m^2)$ incentive constraints. In particular, the problem is solvable in time polynomial in the explicit input size.

The LP uses variables $q_{a,\ell,\omega}$, interpreted as the joint probability of state ω and a refined recommendation (a, ℓ). Bayes plausibility is

$$\sum_{a \in \mathcal{A}} \sum_{\ell \in \mathcal{L}_a} q_{a,\ell,\omega} = \mu_0(\omega), \qquad \forall \omega \in \Omega. \tag{5}$$

The refined incentive constraints are

$$\sum_{\omega \in \Omega} q_{a,\ell,\omega} \big(c_{a,\ell}(\omega) - c_{a',\ell'}(\omega) \big) \ge 0 \qquad \forall (a, \ell), (a', \ell'). \tag{6}$$

The objective is

$$\max_q \sum_{a \in \mathcal{A}} \sum_{\ell \in \mathcal{L}_a} \sum_{\omega \in \Omega} q_{a,\ell,\omega} v(\omega, a). \tag{7}$$

The proof in Appendix D shows that every feasible LP solution induces a valid signaling scheme and every incentive-compatible signaling scheme maps to a feasible LP solution of the same value.

The formulation is deliberately linear. It does not use a big-M encoding, binary variables, or a choice of one active facet after the posterior is known. Instead, the refined signal type chooses the facet as part of the recommendation, and the LP assigns joint probability mass directly to state-type pairs. This is useful when ties occur. If several CVaR facets or several actions are active at the same posterior, the sender may choose any payoff-maximizing refined label among them, and the weak incentive constraints remain valid. Thus the LP captures the same tie-breaking convention as the original persuasion problem without adding a separate selection rule.

4. Finite-Precision Discretization

Theorem 3 solves the unrestricted explicit finite-state model exactly. We now give a separate finite-precision result. Its role is to approximate posterior beliefs by a finite grid while preserving CVaR comparisons. This is useful for implementation and for large or implicit variants where one does not want to enumerate every refined action-facet type.

Let $\mathcal{D}_k$ be the set of k-uniform posteriors,

$$\mathcal{D}_k = \left\{ \mu \in \Delta(\Omega) : \mu(\omega) \in \left\{ 0, \frac{1}{k}, \dots, 1 \right\} \text{ for all } \omega \right\}. \tag{8}$$

Lemma 4 (Uniform approximation by k-uniform posteriors) *Let $C_R = M_R(1 + 2/r)$. For any posterior $\mu \in \Delta(\Omega)$ and tolerance $\epsilon_R > 0$, if*

$$k \geq \frac{2C_R^2}{\epsilon_R^2} \log(2mA), \tag{9}$$

then there exists $\bar{\mu} \in \mathcal{D}_k$ such that

$$\max_{a \in \mathcal{A}} \max_{\ell \in \mathcal{L}_a} |\langle c_{a,\ell}, \mu - \bar{\mu} \rangle| \leq \epsilon_R. \tag{10}$$

Consequently, $\max_{a \in \mathcal{A}} |\rho(\mu, a) - \rho(\bar{\mu}, a)| \leq \epsilon_R$.

The proof samples k independent states from μ and applies Hoeffding's inequality to all CVaR facets; see Appendix E. Based on this grid, define approximate signal labels

$$\widehat{\Sigma} = \left\{ (\bar{\mu}, a) \in \mathcal{D}_k \times \mathcal{A} : \rho(\bar{\mu}, a) \geq \max_{a' \in \mathcal{A}} \rho(\bar{\mu}, a') - 2\epsilon_R \right\}. \tag{11}$$

Solving the state-contingent LP over labels in $\widehat{\Sigma}$ gives the following guarantee.

Theorem 5 (Finite-precision discretization) *Assume that an optimal exact-IC signaling scheme has finite support. For any $\epsilon > 0$, set $\epsilon_R = \epsilon/4$ and choose*

$$k = O\left(\frac{C_R^2 \log(mA)}{\epsilon^2} \right).$$

The discretized state-contingent LP returns a signaling scheme whose every supported posterior-action pair has IC regret at most ϵ, and whose sender utility is at least the exact-IC optimum. Its running time is polynomial in Am^k.

This theorem should be read as a finite-precision and extension result. In the baseline finite-state model, Theorem 3 gives the exact polynomial-time solution.

The reason to keep this result is practical rather than foundational. In high-dimensional instances, the exact LP may still be large because it enumerates all states, actions, and CVaR facets. In extensions with implicit action sets, sampled state spaces, or local access to payoff vectors, a grid method can work with a smaller set of candidate posteriors and can be combined with local active-facet screening. The guarantee above states what must be preserved: all affine CVaR pieces that matter for incentives must be approximated uniformly. This keeps the discretized scheme close in receiver regret while leaving the exact tractability statement intact.

Approximate IC is the natural worst-case guarantee near indifference. If two receiver actions have almost the same CVaR value, a small posterior perturbation can change the exact best response. Exact implementation is recovered under a margin condition. For a posterior-action pair (μ, a), define

$$\Gamma(\mu, a) = \rho(\mu, a) - \max_{a' \neq a} \rho(\mu, a'). \tag{12}$$

A positive value of $\Gamma(\mu, a)$ means that a is a strict best response. For $\gamma > 0$, define the margin-filtered alphabet

$$\widehat{\Sigma}_\gamma = \left\{ (\bar{\mu}, a) \in \mathcal{D}_k \times \mathcal{A} : \Gamma(\bar{\mu}, a) \geq \frac{\gamma}{2} \right\}. \tag{13}$$

Proposition 6 (Margin stability) *If* $\max_{a\in\mathcal{A}} |\rho(\mu, a) - \rho(\bar{\mu}, a)| \le \epsilon$*, then* $\Gamma(\bar{\mu}, a) \ge \Gamma(\mu, a) - 2\epsilon$ *for every* $a \in \mathcal{A}$*.*

Theorem 7 (Strict IC under a margin condition) *Suppose there is a finite-support benchmark scheme with sender value* OPT_γ *such that every supported recommendation* (μ, a) *satisfies* $\Gamma(\mu, a) \ge \gamma > 0$*. Run the discretized LP with the margin-filtered alphabet* $\widehat{\Sigma}_\gamma$ *and choose* $\epsilon_R < \gamma/4$*. Then every supported recommendation in the induced scheme is strictly incentive compatible. More precisely, every supported signal* σ *satisfies*

$$\Gamma(\mu_\sigma, a_\sigma) \ge \frac{\gamma}{2} - 2\epsilon_R > 0. \tag{14}$$

The sender utility is at least OPT_γ*, and the LP size is polynomial in* Am^k *with* $k = O(C_R^2 \log(mA)/\gamma^2)$*.*

The proofs are in Appendix E. The margin condition is not a new modeling assumption for the exact LP. It only states when a finite-precision approximation preserves exact receiver incentives signal by signal.

5. Conclusion

CVaR preferences change the revelation structure of Bayesian persuasion. Action-only recommendations are no longer sufficient because the receiver's relevant tail event can change when signals are merged. The finite-state CVaR problem nevertheless remains tractable: after refining recommendations by active CVaR facets, incentive compatibility is linear and the sender's problem is an ordinary LP.

For large or oracle-described instances, finite-precision discretization provides a controlled approximation that preserves receiver incentives up to a uniform regret bound. When the benchmark recommendations have a positive margin, the same discretization can be filtered to recover strict incentive compatibility. This separation is useful for risk-aware information design: the nonlinearity of CVaR changes the right representation of signals, but it does not by itself make the explicit finite-state persuasion problem computationally intractable.

References

Carlo Acerbi and Dirk Tasche. On the coherence of expected shortfall. *Journal of Banking & Finance*, 26(7):1487–1503, July 2002.

Ingo Althöfer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355, March 1994.

Jerry Anunrojwong, Krishnamurthy Iyer, and David Lingenbrink. Persuading risk-conscious agents: A geometric approach. *Operations Research*, 72(1):151–166, January 2024.

Philippe Artzner, Freddy Delbaen, Jean-Marc Eber, and David Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, July 1999.

Dirk Bergemann and Stephen Morris. Information design: A unified perspective. *Journal of Economic Literature*, 57(1):44–95, March 2019.

Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian persuasion. *SIAM Journal on Computing*, 50(3):16–68, January 2021. Num Pages: STOC16-97.

Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, October 2011.

Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM conference on Electronic commerce*, EC '03, pages 36–41, New York, NY, USA, June 2003. ISBN 978-1-58113-679-1.

R. Tyrrell Rockafellar and Stanislav Uryasev. Optimization of conditional value-at-risk. *The Journal of Risk*, 2(3):21–41, 2000.

R. Tyrrell Rockafellar and Stanislav Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, July 2002.

Haifeng Xu, Zinovi Rabinovich, Shaddin Dughmi, and Milind Tambe. Exploring information asymmetry in two-stage security games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1057–1063, Austin, Texas, January 2015. ISBN 978-0-262-51129-2.

Appendix A. Finite Support

Proof [Proof of Lemma 1] Let $m = |\Omega|$. Any signaling scheme induces a Bayes-plausible distribution over posteriors, namely a probability measure τ on $\Delta(\Omega)$ with mean μ_0. For each posterior μ, fix a receiver best response that maximizes the sender value among tied actions, and let $\Phi(\mu)$ be the sender's induced payoff at that posterior. The sender's value is $\int \Phi(\mu) d\tau(\mu)$, while Bayes plausibility imposes the moment equations $\int \mu(\omega) d\tau(\mu) = \mu_0(\omega)$.

An optimal value is attained by an extreme point of the feasible set of probability measures with these moment constraints. By the standard Carathéodory argument for such measures, an extreme feasible distribution has support size at most m, because posterior vectors lie in an $(m-1)$-dimensional simplex and the total mass constraint supplies the remaining affine coordinate. The support points are the posteriors induced by signals, and their weights are the signal probabilities. ■

Appendix B. Finite-Facet Representation of CVaR

Proof Fix action a and write $X_a(\omega) = u(\omega, a)$. For posterior μ,

$$\mathrm{CVaR}_r^{\mu}(X_a) = \sup_{b \in \mathbb{R}} \left\{ b - \frac{1}{r} \sum_{\omega \in \Omega} \mu(\omega)(b - X_a(\omega))^+ \right\}. \tag{15}$$

Let the distinct realized utility values of X_a be $z_1 < \cdots < z_q$, with $q \le m$. For fixed μ, the displayed objective is concave and piecewise linear in b, with breakpoints among these values. Outside $[z_1, z_q]$, the objective is weakly improved by moving to the nearest endpoint. On each interval between adjacent breakpoints, an affine function attains a maximum at an endpoint unless it is constant. Thus the supremum can be restricted to $b \in \{z_1, \ldots, z_q\}$.

For each z_j, define

$$c_{a,j}(\omega) = z_j - \frac{1}{r}(z_j - u(\omega, a))^+. \tag{16}$$

Since $\sum_\omega \mu(\omega) = 1$, the objective at $b = z_j$ equals $\langle c_{a,j}, \mu \rangle$. Hence $f_a(\mu) = \max_j \langle c_{a,j}, \mu \rangle$, with $q \le m$. If $|u(\omega, a)| \le M_R$, then $|z_j| \le M_R$ and $(z_j - u(\omega, a))^+ \le 2M_R$, so $\|c_{a,j}\|_\infty \le M_R(1 + 2/r)$. ■

Appendix C. Active-Facet Revelation

Proof [Proof of Proposition 2] Consider any incentive-compatible signaling scheme. Let s be a signal with probability λ_s, posterior μ_s, and recommended action a_s. Since a_s is incentive compatible, $f_{a_s}(\mu_s) \ge f_{a'}(\mu_s)$ for every $a' \in \mathcal{A}$. Choose one active facet $\ell_s \in \mathcal{L}_{a_s}$ satisfying

$$f_{a_s}(\mu_s) = \langle c_{a_s, \ell_s}, \mu_s \rangle.$$

Then for every $a' \in \mathcal{A}$ and every $\ell' \in \mathcal{L}_{a'}$,

$$\langle c_{a_s, \ell_s}, \mu_s \rangle = f_{a_s}(\mu_s) \ge f_{a'}(\mu_s) \ge \langle c_{a', \ell'}, \mu_s \rangle.$$

Thus $\mu_s \in P_{a_s,\ell_s}$. Relabeling signal s by (a_s, ℓ_s) does not change its probability, posterior, recommended action, or sender payoff. Applying this to every signal gives the refined scheme. ■

Appendix D. Exact Active-Facet LP

Proof [Proof of Theorem 3] For each refined type (a, ℓ) and state ω, introduce $q_{a,\ell,\omega} \geq 0$, the joint probability of state ω and refined recommendation (a, ℓ). The sender's objective and Bayes-plausibility constraints are those stated in Section 3. The posterior induced by a positive-mass type (a, ℓ) is

$$\mu_{a,\ell}(\omega) = \frac{q_{a,\ell,\omega}}{\lambda_{a,\ell}}, \qquad \lambda_{a,\ell} = \sum_{\omega} q_{a,\ell,\omega}.$$

The condition $\mu_{a,\ell} \in P_{a,\ell}$ is equivalent, after multiplying by $\lambda_{a,\ell}$, to

$$\sum_{\omega} q_{a,\ell,\omega}\big(c_{a,\ell}(\omega) - c_{a',\ell'}(\omega)\big) \geq 0 \qquad \forall (a', \ell').$$

Thus every feasible LP solution induces an incentive-compatible signaling scheme.

Conversely, take any incentive-compatible signaling scheme. For each signal s, choose an active facet ℓ_s for its recommended action a_s. By the argument in Proposition 2, $\mu_s \in P_{a_s,\ell_s}$. Assign joint mass $\lambda_s \mu_s(\omega)$ to $q_{a_s,\ell_s,\omega}$, aggregating signals with the same refined type. The constraints hold because each $P_{a,\ell}$ is defined by linear inequalities and is convex. The objective value is unchanged, so the LP optimum equals the optimal persuasion value. The size bound follows from $L = \sum_a |\mathcal{L}_a| \leq Am$. ■

Appendix E. Discretization Proofs

Proof [Proof of Lemma 4] Sample $X_1, \ldots, X_k$ independently from μ and let $\bar{\mu}$ be their empirical distribution. Fix action a and facet ℓ. The variables $Y_i = c_{a,\ell}(X_i)$ have mean $\langle c_{a,\ell}, \mu \rangle$, and their empirical average equals $\langle c_{a,\ell}, \bar{\mu} \rangle$. Since $|Y_i| \leq C_R$, Hoeffding's inequality gives

$$\Pr\left(|\langle c_{a,\ell}, \bar{\mu} - \mu \rangle| \geq \epsilon_R\right) \leq 2 \exp\left(-\frac{k\epsilon_R^2}{2C_R^2}\right).$$

There are at most mA action-facet pairs. The stated lower bound on k makes the union bound smaller than one, so some empirical posterior satisfies all facet inequalities. Since each CVaR value is the maximum over its facets, the same ϵ_R bound holds for every action value. ■

Proof [Proof of Theorem 5] The discretized LP is the standard state-contingent LP over labels $\widehat{\Sigma}$. Its variables $\varphi(\omega, \sigma)$ specify the probability of sending label $\sigma = (\bar{\mu}_\sigma, a_\sigma)$ in state ω. The LP maximizes sender payoff subject to probability constraints and cell constraints

$$|\langle c_{a,\ell}, \mu_\sigma - \bar{\mu}_\sigma \rangle| \leq \epsilon_R \qquad \forall a, \ell$$

for every supported induced posterior μ_σ. These inequalities are written linearly after multiplying by the signal probability.

Soundness follows because every supported label is $2\epsilon_R$-optimal at its grid center, while all CVaR values move by at most ϵ_R between the grid center and the induced posterior. Hence the recommended action has IC regret at most $4\epsilon_R = \epsilon$. Completeness follows by taking an optimal finite-support exact-IC scheme, replacing each posterior by the grid proxy given by Lemma 4, and keeping the same state-contingent probabilities. Exact optimality at the original posterior becomes $2\epsilon_R$-optimality at the grid center, so the label is available. The sender objective is preserved. Since $|\mathcal{D}_k| \le m^k$, the LP size is polynomial in Am^k. ■

Proof [Proof of Proposition 6] For any competing action $a' \neq a$, the proxy condition gives

$$\rho(\bar{\mu}, a) \ge \rho(\mu, a) - \epsilon, \qquad \rho(\bar{\mu}, a') \le \rho(\mu, a') + \epsilon.$$

Thus

$$\rho(\bar{\mu}, a) - \rho(\bar{\mu}, a') \ge \rho(\mu, a) - \rho(\mu, a') - 2\epsilon.$$

Taking the minimum over all $a' \neq a$ proves the claim. ■

Proof [Proof of Theorem 7] Let the benchmark scheme be $\{(p_j, \mu_j, a_j)\}_{j\in J}$, with $\Gamma(\mu_j, a_j) \ge \gamma$ for every supported signal. For each j, choose a grid posterior $\bar{\mu}_j$ using Lemma 4. Since $\epsilon_R < \gamma/4$, Proposition 6 gives

$$\Gamma(\bar{\mu}_j, a_j) \ge \gamma - 2\epsilon_R > \frac{\gamma}{2}.$$

Hence each $(\bar{\mu}_j, a_j)$ belongs to the margin-filtered alphabet. The same state-contingent construction used in the completeness proof of Theorem 5 maps the benchmark scheme into a feasible solution of the filtered LP and preserves sender value, so the filtered LP obtains value at least OPT_γ.

Now take any feasible solution of the filtered LP and any supported signal $\sigma = (\bar{\mu}_\sigma, a_\sigma)$. The LP cell constraints imply

$$\max_{a\in\mathcal{A}} |\rho(\mu_\sigma, a) - \rho(\bar{\mu}_\sigma, a)| \le \epsilon_R.$$

Because $\sigma \in \widehat{\Sigma}_\gamma$, we have $\Gamma(\bar{\mu}_\sigma, a_\sigma) \ge \gamma/2$. Applying Proposition 6 again gives

$$\Gamma(\mu_\sigma, a_\sigma) \ge \frac{\gamma}{2} - 2\epsilon_R > 0.$$

Thus every supported recommendation is strictly incentive compatible. The size bound follows from Lemma 4 with $\epsilon_R = \Theta(\gamma)$ and $|\mathcal{D}_k| \le m^k$. ■

Appendix F. Limitations and Broader Impacts

The results assume finite state and action spaces, a known prior, known utilities, and a fixed CVaR level. The exact LP also relies on explicit enumeration of states, actions, and CVaR facets. In very large, combinatorial, or oracle-described environments, explicit enumeration may be unavailable, which is why finite-precision and local discretization methods remain useful.

The framework is motivated by AI-assisted and safety-critical decision systems. A positive use is to audit whether a disclosure policy remains persuasive when receivers care about tail losses. The same tools can be misused by a sender who knows the receiver's risk preferences and designs information to steer behavior near indifference boundaries. The results should therefore be interpreted as a basis for analyzing and auditing risk-aware information policies.