NExT-Game 2026: New Frontiers in Game-Theoretic Learning – ICML 2026 Workshop 1–15

Bayesian Persuasion with a Risk-Conscious Receiver

Yujing Chen
School of Mathematical Sciences, Peking University

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 information-design settings are tail-sensitive: a clinician interpreting an automated triage alert, a regulator evaluating an AI release, or a security operator may reject an action that is favorable on average but leaves an unacceptable posterior tail of failures. We capture this behavior with 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 value is nonlinear in the posterior, so two signals that separately induce the same action may, once merged, change the relevant tail event and 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

© Y. Chen.

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). Anunrojwong et al. (Anunrojwong et al., 2024) develop geometric tools for general nonlinear receiver preferences; our result is complementary, specializing to finite-state CVaR and exploiting its enumerable active thresholds to obtain an LP for arbitrary finite action sets. More generally, the argument applies to explicitly enumerable finite max-affine risk functionals, including fixed finite mixtures of CVaR levels, but not directly to continuous spectral or non-polyhedral risk measures. 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 $\Omega = \{\omega_1, \ldots, \omega_m\}$, a common prior $\mu_0 \in \operatorname{int}(\Delta(\Omega))$, and a finite action set $\mathcal{A} = \{a_1, \ldots, a_A\}$. 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 $\arg\max_{a \in \mathcal{A}} \rho(\mu, a)$. Unless stated otherwise, ties are broken in the

sender's favor: among receiver-optimal actions, the receiver selects one that maximizes the sender's conditional payoff. 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,

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

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 directly with posterior-based signaling. A standard finite-support reduction implies that an optimal finite-state scheme can be chosen with at most $m = |\Omega|$ posteriors; we use this fact only to justify the finite-support benchmark in the discretization result. The reduction is recalled in Appendix A. The sharper CVaR structure developed below comes instead from active facets.

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.

Example 1 (Action-only merging fails) *Let $r = 1/2$, $\Omega = \{\omega_1, \omega_2, \omega_3\}$, and consider two receiver actions $a, \tilde{a}$ with $u(\cdot, a) = \left(\frac{1}{5}, 1, \frac{1}{5}\right), u(\cdot, \tilde{a}) = \left(0, \frac{4}{5}, \frac{4}{5}\right)$. Let the prior be $\mu_0 = \left(\frac{3}{10}, \frac{1}{2}, \frac{1}{5}\right)$, and let the sender use two signals s_1, s_2 with*

$$\pi(s_1 \mid \cdot) = \left(\frac{1}{6}, \frac{9}{10}, 0\right), \qquad \pi(s_2 \mid \cdot) = \left(\frac{5}{6}, \frac{1}{10}, 1\right).$$

Each signal is sent with probability $1/2$. Bayes' rule gives $\mu_{s_1} = \left(\frac{1}{10}, \frac{9}{10}, 0\right), \mu_{s_2} = \left(\frac{1}{2}, \frac{1}{10}, \frac{2}{5}\right)$.

At these posteriors,

$$\left(f_a(\mu_{s_1}), f_{\tilde{a}}(\mu_{s_1})\right) = \left(\frac{21}{25}, \frac{16}{25}\right), \quad \left(f_a(\mu_{s_2}), f_{\tilde{a}}(\mu_{s_2})\right) = \left(\frac{1}{5}, 0\right).$$

Thus, a is the unique best response after either signal. However, merging the two signals yields posterior $\frac{1}{2}\mu_{s_1} + \frac{1}{2}\mu_{s_2} = \mu_0$, at which $\big(f_a(\mu_0), f_{\tilde{a}}(\mu_0)\big) = \big(\frac{1}{5}, \frac{8}{25}\big)$. Hence a is no longer optimal after merging.

The maximizing CVaR threshold for a is 1 at μ_{s_1} and 1/5 at μ_{s_2}. Thus the two signals have distinct refined labels, (a, ℓ_1) and $(a, \ell_{1/5})$, although both recommend a. Action-based merging fails, whereas the action–facet representation retains the two linear incentive regions separately.

Proposition 1 (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 2 (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 $L \leq Am$, 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

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

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: the refined signal type chooses a supporting facet as part of the recommendation, so no big-M encoding or binary selection variables are

needed. The weak inequalities in (6) implement the sender-preferred tie-breaking convention stated in Section 2.

For robustness to arbitrary tie-breaking, fix a margin $\gamma > 0$. For every competing action $a' \neq a$, replace the corresponding constraints in (6) by

$$\sum_{\omega\in\Omega} q_{a,\ell,\omega}\big(c_{a,\ell}(\omega) - c_{a',\ell'}(\omega)\big) \geq \gamma \sum_{\omega\in\Omega} q_{a,\ell,\omega}, \qquad \forall \ell' \in \mathcal{L}_{a'}.$$

The within-action facet constraints remain weak. The resulting formulation is still an LP and computes the optimal γ-strictly incentive-compatible scheme, whose implementation is independent of tie-breaking. Without a positive margin, a sender-preferred optimum may rely on exact indifference and need not remain implementable under adversarial tie-breaking.

4. Finite-Precision Discretization

Theorem 2 solves the explicit finite-state CVaR persuasion problem exactly. We therefore use discretization only as a finite-precision construction. It also describes how a supplied finite set of candidate posteriors can be handled by a state-contingent LP.

Let $C_R = M_R(1 + 2/r)$, and let

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

For any $\epsilon_R > 0$, choosing $k \geq \frac{2C_R^2}{\epsilon_R^2}\log(2mA)$ ensures that every posterior μ has a proxy $\bar{\mu} \in \mathcal{D}_k$ satisfying

$$\max_{a\in\mathcal{A},\ \ell\in\mathcal{L}_a} |\langle c_{a,\ell}, \mu - \bar{\mu}\rangle| \leq \epsilon_R; \tag{9}$$

the proof is given in Appendix E.

Given a candidate grid set $\mathcal{G} \subseteq \mathcal{D}_k$, define

$$\widehat{\Sigma}(\mathcal{G}) = \left\{(\bar{\mu}, a) \in \mathcal{G} \times \mathcal{A} : \rho(\bar{\mu}, a) \geq \max_{a'\in\mathcal{A}} \rho(\bar{\mu}, a') - 2\epsilon_R\right\}. \tag{10}$$

Thus, given $\mathcal{G}$, the reduced alphabet is constructed by evaluating the A CVaR values at each $\bar{\mu} \in \mathcal{G}$ and retaining its $2\epsilon_R$-optimal actions.

The state-contingent LP uses joint-mass variables $\phi(\omega, \sigma)$, where $\sigma = (\bar{\mu}_\sigma, a_\sigma) \in \widehat{\Sigma}(\mathcal{G})$, and $\lambda_\sigma = \sum_{\omega\in\Omega} \phi(\omega, \sigma)$. Besides the state-wise probability constraints $\sum_{\sigma\in\widehat{\Sigma}(\mathcal{G})} \phi(\omega, \sigma) = \mu_0(\omega), \forall \omega \in \Omega$, the LP imposes the linear cell constraints

$$-\epsilon_R\lambda_\sigma \leq \sum_{\omega\in\Omega} \phi(\omega, \sigma)\left(c_{a,\ell}(\omega) - \langle c_{a,\ell}, \bar{\mu}_\sigma\rangle\right) \leq \epsilon_R\lambda_\sigma, \tag{11}$$

for every $\sigma \in \widehat{\Sigma}(\mathcal{G})$, $a \in \mathcal{A}$, and $\ell \in \mathcal{L}_a$. Thus the posterior induced by each label remains uniformly close to its grid proxy on all CVaR affine pieces.

Theorem 3 (Finite-precision discretization) *Assume that an optimal exact-IC signaling scheme has finite support. For any $\epsilon > 0$, set $\epsilon_R = \epsilon/4$, choose*

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

and run the state-contingent LP with $\mathcal{G} = \mathcal{D}_k$. The resulting signaling scheme has IC regret at most ϵ at every supported posterior-action pair, and its sender utility is at least the exact-IC optimum. The LP has $m|\widehat{\Sigma}(\mathcal{D}_k)|$ variables and $O\left(m + L|\widehat{\Sigma}(\mathcal{D}_k)|\right)$ linear constraints, where $L = \sum_{a\in\mathcal{A}} |\mathcal{L}_a| \leq Am$. In particular, its size is polynomial in Am^k.

The full guarantee in Theorem 3 uses the complete grid $\mathcal{D}_k$, whose cardinality

$$|\mathcal{D}_k| = \binom{k+m-1}{m-1}$$

can be large. Thus discretization is a finite-precision guarantee rather than a scalable replacement for the exact LP. In practice, the same formulation can be run on a screened candidate set $\mathcal{G} \subseteq \mathcal{D}_k$, but such screening does not retain the worst-case guarantee without an additional coverage or separation argument.

A margin-filtered variant recovers strict incentive compatibility. Specifically, if a benchmark scheme has a uniform action gap $\gamma > 0$, then retaining only grid labels whose recommended action has gap at least $\gamma/2$, and taking $\epsilon_R < \gamma/4$, yields a scheme that is strictly IC at every supported signal. The formal statement and proof appear in Appendix E.

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

Lemma 4 (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.*

Proof 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,

$$\mathrm{CVaR}_r^{\mu}(X_a) = \sup_{b \in \mathbb{R}} \left\{ b - \frac{1}{r} \sum_{\omega \in \Omega} \mu(\omega)(b - X_a(\omega))^+ \right\}. \tag{12}$$

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

$$c_{a,j}(\omega) = z_j - \frac{1}{r}(z_j - u(\omega, a))^+. \tag{13}$$

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 1] 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 $\ell' \in \mathcal{L}_{a'}$,

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

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 2] 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 1, $\mu_s \in P_{a_s,\ell_s}$. 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

Lemma 5 (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*

$$k \geq \frac{2C_R^2}{\epsilon_R^2} \log(2mA), \tag{14}$$

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

$$\max_{a \in \mathcal{A}} \max_{\ell \in \mathcal{L}_a} |\langle c_{a,\ell}, \mu - \bar{\mu} \rangle| \leq \epsilon_R. \tag{15}$$

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

Proof 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 $|Y_i| \le C_R$, Hoeffding's inequality gives

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

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. ■

Theorem 6 (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.

Proof 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

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

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 5, 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. ■

Proposition 7 (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 $a \in \mathcal{A}$.*

Proof For any competing action $a' \ne a$, 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 $a' \ne a$ proves the claim. ■

Theorem 8 (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 $\Gamma(\mu, a) \geq \gamma > 0$. 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*

$$\Gamma(\mu_\sigma, a_\sigma) \geq \frac{\gamma}{2} - 2\epsilon_R > 0. \tag{16}$$

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

Proof Let the benchmark scheme be $\{(p_j, \mu_j, a_j)\}_{j \in J}$, with $\Gamma(\mu_j, a_j) \geq \gamma$ for every supported signal. For each j, choose a grid posterior $\bar{\mu}_j$ using Lemma 5. Since $\epsilon_R < \gamma/4$, Proposition 7 gives

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

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 6 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

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

Because $\sigma \in \widehat{\Sigma}_\gamma$, we have $\Gamma(\bar{\mu}_\sigma, a_\sigma) \geq \gamma/2$. Applying Proposition 7 again gives

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

Thus every supported recommendation is strictly incentive compatible. The size bound follows from Lemma 5 with $\epsilon_R = \Theta(\gamma)$ and $|\mathcal{D}_k| \leq m^k$. ■

Appendix F. Synthetic Computational Checks

Purpose. We include a small set of synthetic computational checks to clarify the practical meaning of the active-facet LP. The goal is not to claim scalability for very large or combinatorial state spaces. Rather, the checks show that the exact formulation can be solved by an off-the-shelf LP solver on explicit finite instances, that repeated utility values and empty refined types can make the realized LP much smaller than the worst-case bound, and that unstructured instances still grow quickly.

The action-only construction is not used as an optimization baseline. Example 1 shows that action-only merging can be infeasible under CVaR. We therefore compare the full active-facet LP with a constraint-generated implementation of the same LP.

Implementation. For each action, the candidate CVaR facets are generated from the distinct receiver utility thresholds. We then screen out affine pieces that never attain the action's CVaR envelope on the simplex, and screen out refined types whose global active-facet region is empty. Empty refined types are removed only as variables; their facets remain available as comparison facets for the incentive constraints. Let L_{env} denote the number of retained comparison facets and let T denote the number of nonempty refined types. The resulting LP has mT variables, and the number of instantiated refined IC constraints is at most $T(L_{\text{env}} - 1)$.

We solve two sparse LP implementations. The full LP materializes all refined IC constraints. The constraint-generated LP starts from a smaller set and iteratively adds violated refined IC constraints. After termination, we explicitly evaluate every omitted refined IC inequality and report the maximum omitted-constraint violation. This is a numerical feasibility check, not a separate exact certificate.

Instance families. All runs use $r = 0.25$, priors sampled from Dirichlet($\mathbf{1}$), and sender payoffs sampled independently from $[0, 1]$. In the structured / repeated-utility family, one action is constant and the remaining receiver utilities are sampled from q discrete levels. In the continuous stress family, receiver utilities are sampled independently from $[0, 1]$, so each action typically has close to m distinct thresholds.

Results.

Results. Tables 1 and 2 report medians and interquartile ranges over seeds for which both formulations completed and passed the numerical agreement checks described above.

The tables report median [IQR] over successful seeds; the Solved column counts seeds that solved both the full LP and the constraint-generated LP, with |OPT_full-OPT_CG| <= 1e-07 and maximum omitted-constraint violation at most 1e-08. Unsuccessful seeds remain in the raw CSV. Experiments used Python 3.9.13, NumPy 1.26.4, SciPy 1.13.1 (HiGHS backend bundled with SciPy 1.13.1), CPU macOS-26.4.1-x86_64-i386-64bit (x86_64), RAM 32.0 GB, and 10 seeds per configuration. HiGHS primal/dual feasibility tolerances were set to 1e-09. Full and CG times include the common facet/type screening; CG time also includes all constraint scans. Rows with zero successful seeds show – for median [IQR] columns because these summaries are defined over successful Full+CG checks only.

Table 1: Structured / repeated-utility instances. Entries are median [IQR] over solved seeds.

Solved	m	A	q	L_{env}	T	Variables	Full rows	Full time	CG rows	CG time
10/10	10	5	5	20 [19, 20]	12 [12, 13]	120 [120, 130]	224 [216, 242]	0.07 [0.07, 0.07]	124 [103, 136]	0.08 [0.08, 0.08]
10/10	25	10	5	46 [45, 46]	28 [27, 28]	700 [681, 700]	1260 [1199, 1260]	0.23 [0.22, 0.23]	722 [636, 772]	0.32 [0.28, 0.35]
10/10	50	20	5	96 [96, 96]	58 [58, 58]	2900 [2900, 2900]	5510 [5510, 5510]	1.43 [1.41, 1.45]	3174 [2990, 3638]	3.03 [2.64, 3.23]
10/10	100	20	5	96 [96, 96]	58 [58, 58]	5800 [5800, 5800]	5510 [5510, 5510]	3.62 [3.59, 3.75]	3728 [3452, 3860]	9.08 [7.21, 9.62]
10/10	100	20	10	191 [191, 191]	96 [96, 96]	9600 [9600, 9600]	18240 [18240, 18240]	19.93 [19.42, 20.11]	13134 [12846, 13833]	80.92 [60.86, 95.12]

Table 2: Continuous-utility stress instances. Entries are median [IQR] over solved seeds.

Solved	m	A	L_{env}	T	Full rows	Full time	CG rows	CG time	Obj. gap	Max omitted viol.
10/10	10	5	50 [50, 50]	40 [37, 42]	1936 [1825, 2046]	0.21 [0.21, 0.21]	1036 [856, 1223]	0.42 [0.34, 0.47]	0.0e+00 [0.0e+00, 1.9e-16]	0.0e+00 [0.0e+00, 0.0e+00]
10/10	20	10	200 [200, 200]	173 [164, 175]	34427 [32736, 34775]	3.63 [3.28, 3.92]	18403 [17446, 19614]	23.73 [17.94, 26.43]	2.2e-16 [1.1e-16, 3.1e-16]	0.0e+00 [0.0e+00, 6.5e-19]
8/10	30	10	300 [300, 300]	268 [265, 271]	79982 [79310, 81029]	16.97 [16.34, 17.66]	39538 [35708, 44361]	85.37 [64.27, 101.41]	2.2e-16 [1.9e-16, 2.8e-16]	0.0e+00 [0.0e+00, 2.2e-18]

On the structured repeated-utility instances, the realized LP dimensions are substantially smaller than the worst-case bounds because each non-safe action has only q distinct utility levels. For example, when $A = 20$ and $q = 5$, the preprocessing retains $L_{\text{env}} = 96$ comparison facets and $T = 58$ nonempty refined types for both $m = 50$ and $m = 100$. Hence the number of full refined IC rows remains 5,510, while the number of variables increases from 2,900 to 5,800. The full LP solves the tested $m = 100, A = 20, q = 5$ instances in $3.62\,[3.59, 3.75]$ seconds and the $q = 10$ instances in $19.93\,[19.42, 20.11]$ seconds. Thus, for explicit instances with repeated utility values, the exact formulation is directly solvable at the tested scales.

Constraint generation materializes fewer refined IC rows, but it is not faster on these instances. For example, at $m = 100, A = 20, q = 5$, it uses 3,728 [3,452, 3,860] rows instead of 5,510, but requires $9.08\,[7.21, 9.62]$ seconds versus $3.62\,[3.59, 3.75]$ seconds for the full LP. At $q = 10$, it uses 13,134 [12,846, 13,833] rows instead of 18,240, while its total runtime is $80.92\,[60.86, 95.12]$ seconds versus $19.93\,[19.42, 20.11]$ seconds for the full LP. The repeated LP solves and global violation scans therefore outweigh the reduction in instantiated rows in the tested structured instances.

The continuous-utility stress instances are closer to the worst case: all tested configurations retain $L_{\text{env}} = mA$ comparison facets. As m increases from 10 to 30 with $A = 10$, the median number of full IC rows grows from 1,936 to 79,982, and the full-LP runtime grows from 0.21 to 16.97 seconds. Constraint generation again adds fewer rows, but its total runtime is larger because of repeated solves and complete constraint scans. For all solved seeds, the full and constraint-generated formulations agree in objective value up to numerical precision, and the maximum omitted-constraint violation is below the prescribed tolerance.

For $m = 40$ and $A = 10$, no seed completed both formulations and passed the numerical agreement checks under the imposed solver limits. We therefore omit median summaries for this configuration. This outcome is consistent with the rapid growth of the unstructured formulation and reinforces that constraint generation is a sparse numerical implementation of the same LP, rather than a worst-case scalability improvement.

Appendix G. 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.

Camera-Ready Edits Summary

Relative to the original submission, we made the following substantive revisions in response to the reviewers' comments.

- **Illustrative example.** We added a three-state, two-action signaling example showing that two signals that separately recommend the same action cannot be merged under CVaR. The example identifies the distinct active thresholds and shows how action–facet labels retain the relevant linear incentive regions.
- **Constructive discretization outline.** We added an explicit description of how to construct the approximate label set from candidate grid points, formulate the state-contingent LP with state–label joint-mass variables, and impose the linear cell constraints. We now distinguish the full-grid theoretical guarantee from screened candidate sets used in practice, for which no worst-case coverage guarantee is claimed.
- **Computational scope.** We clarified that the active-facet LP is an exact method for explicit finite-state instances, rather than a general scalable method for high-dimensional or combinatorial state spaces. The discussion now makes explicit that the LP size depends on the realized facet family as well as the worst-case bound.
- **Synthetic computational checks.** We added appendix experiments comparing the full active-facet LP with a constraint-generated implementation of the same formulation. The tables report realized facet counts, empty-type screening, runtimes, objective agreement, omitted-constraint violations, and completion counts across fixed random seeds.
- **Relation to prior work and scope beyond CVaR.** We clarified that our result complements geometric approaches for general nonlinear receiver preferences: we specialize to finite-state CVaR and exploit its enumerable active-threshold representation to obtain an LP for arbitrary finite action sets. We also state that the same argument applies to explicitly enumerable finite max-affine risk functionals, including fixed finite mixtures of CVaR levels, but not directly to continuous spectral or non-polyhedral risk measures.
- **Tie-breaking and exposition.** We made the sender-preferred tie-breaking convention explicit. We also explain that a γ-strict version of the LP is robust to arbitrary tie-breaking, and relate this modification to the margin-filtered discretization result. To accommodate these additions within the main-paper limit, we moved the standard finite-support reduction and detailed discretization and margin-stability statements to the appendix.