THE PRICE OF COMPRESSING PREFERENCES: SHARP REGRET BOUNDS AND CYCLE CERTIFICATES FOR MECHANISM-READY REPRESENTATIONS

Anonymous authors
Paper under double-blind review

ABSTRACT

Reusable representations of private preferences are central to decision systems that must operate across changing prices, menus, and mechanisms, so compression should preserve the comparisons that determine downstream choices rather than raw utility coordinates. This requires a task-specific theory because outcome-independent utility shifts are behaviorally irrelevant, while merging several valuations can create multi-outcome worst-case regret that is invisible to coordinatewise reconstruction error, unpriced rankings, or pairwise diameter alone. We develop the *Price of Compressing Preferences* (PCP) framework for finite-outcome quasilinear economies, which identifies the exact information a menu-independent representation must retain to control regret uniformly over future priced menus. PCP first quotients valuations by additive constants and derives an exact strategic distortion, then converts each encoder cell into a system of pairwise difference constraints whose optimal surrogate is obtained by a linear program and whose obstruction is a maximum directed cycle mean, and finally lifts these certificates to message complexity, mechanism transfer, rate–distortion, and a Strategic Minimax quantizer. The theory shows that arbitrary deterministic response rules cannot outperform a single optimal surrogate and yields sharp diameter and covering bounds; experiments verify the LP–cycle identity to 1.1×10^{-15} on 240 random cells and, at $K = 16$, reduce held-out maximum strategic distortion relative to Anchored K-means by 13.7% on public projects and 24.3% on bundle selection.

1 INTRODUCTION

Reusable representations are useful only to the extent that they preserve the decisions for which the represented information will later be used. This requirement is especially demanding for private preferences: the same valuation may be queried under different prices, feasible sets, contracts, or mechanisms, so the downstream decision problem is not fixed when the representation is constructed. Related work on utility elicitation has long used decision loss or minimax regret, rather than raw parameter error, to determine which preference information is decision-relevant (Wang & Boutilier, 2003; Boutilier et al., 2006), while information-theoretic representation methods similarly emphasize preserving task-relevant information under compression (Tishby et al., 2000; Arumugam & Van Roy, 2021). Compressing a valuation by ordinary coordinate error can therefore spend capacity on economically irrelevant variation while missing small relative-value errors that are amplified by a future priced comparison. The relevant question is not how faithfully one reconstructs a raw utility vector, but what information must be preserved so that one finite representation remains reliable across all downstream priced choices.

We study this question in a finite-outcome quasilinear model, a canonical setting in mechanism design and quasilinear rationalizability (Myerson, 1981; Rochet, 1987). An encoder observes a valuation and maps it to a finite label *before* a downstream menu of feasible outcomes and payments is revealed; after observing the menu, a deterministic response rule chooses an outcome using only the label and the menu. Our target is uniform menu regret: the worst loss, in monetary units, relative to the action that the true valuation would choose. The timing is essential, because a menu-aware encoder could simply communicate an optimal action and would solve a different communication problem. This separation between private information, message-space restrictions, and downstream implementation

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.

is closely related to classical work on limited communication in mechanisms (Green & Laffont, 1986; 1982; Nisan & Segal, 2006; Van Zandt, 2007). It is also essential to distinguish this task from incentive compatibility alone: a constant mechanism is exactly incentive compatible without learning anything about valuations, so nontrivial representation lower bounds require a decision-preservation objective beyond incentive compatibility itself.

Several neighboring literatures address parts of this problem but not the reusable representation question. Work on communication-constrained mechanism design studies how restricted message spaces interact with incentives and implementation (Green & Laffont, 1986; 1982; Nisan & Segal, 2006; Van Zandt, 2007; Dobzinski, 2016; Dobzinski et al., 2022); related work measures mechanism complexity through menus or exploits structured low-dimensional models to reduce representation and communication burdens (Hart & Nisan, 2019; Cai & Daskalakis, 2022). Automated and learning-based mechanism design instead optimizes a mechanism for a specified environment or valuation distribution (Conitzer & Sandholm, 2002; Dütting et al., 2019), whereas regret-based preference elicitation studies how incomplete utility information can be refined for a decision problem (Wang & Boutilier, 2003; Boutilier et al., 2006). Information theory provides the complementary language of lossy compression and task-relevant representation, from classical rate–distortion theory (Shannon, 1993; Davisson, 1972; Cover et al., 1991) to information bottlenecks and decision-oriented rate–distortion (Tishby et al., 2000; Arumugam & Van Roy, 2021). These approaches leave three structural issues unresolved for a representation chosen before the menu: raw utility space contains an outcome-independent direction that no priced choice can observe; the difficulty of representing an entire encoder cell is a simultaneous system of relative-value constraints and cannot in general be read from pairwise diameter alone; and an upper bound for a chosen surrogate decoder does not establish what is unavoidable for an arbitrary deterministic downstream rule. We do not claim the maximum-cycle-mean machinery (Karp, 1978), classical rate–distortion theory, or cycle-based ideas from revealed preference and quasilinear rationalizability (Afriat, 1967; Varian, 1982; Rochet, 1987; Allen & Rehbeck, 2020) as new; the missing step is an operational identification that turns these ingredients into an exact representation theorem for all future priced menus.

We address these issues through the *Price of Compressing Preferences* (PCP) framework. PCP first quotients valuations by outcome-independent additive constants and measures the remaining error by the oscillation of valuation differences; this strategic distortion is exactly, rather than merely upper-bounding, the worst menu regret caused by replacing one valuation with a surrogate. The central difficulty appears when one label merges a set of valuations C: for each ordered pair of outcomes, the cell induces the largest required relative-value difference, and representing the cell with regret at most r becomes a system of difference inequalities for a surrogate potential. Feasibility of these inequalities is governed by directed cycles. Consequently, Theorem 3.1 shows that the optimal regret of *any* deterministic menu-response rule, the Chebyshev radius of C in strategic distortion, and the maximum directed cycle mean of the cell all coincide. A maximizing cycle constructs an adversarial priced menu for the converse, while feasible vertex potentials construct an optimal surrogate through a linear program. This equivalence resolves the null-direction, whole-cell, and decoder-optimality problems simultaneously; it then lifts to exact covering/message complexity, feasibility-induced dimension, mechanism-transfer guarantees, finite-alphabet rate–distortion, and the Strategic Minimax quantizer used in our experiments.

Contributions. Our contributions are fourfold. First, we give a sharp operational geometry for reusable preference representations: strategic distortion equals worst-case surrogate-induced menu regret, and cell-wise deterministic minimax regret equals both an ambient Chebyshev radius and a maximum directed cycle mean, yielding sharp diameter bounds and constructive primal–adversarial certificates. Second, we turn this characterization into information statements: the minimum label alphabet for uniform regret ε is exactly an ambient covering number in the quotient space, with communication scaling determined by the strategically visible dimension rather than the ambient parameter dimension; we also state the corresponding finite-alphabet rate–distortion formulation for expected certificates. Third, we show how representation error transfers to downstream economics by bounding approximate dominant-strategy deviations and welfare loss, and we analyze a Strategic Minimax quantizer with deterministic stability bounds under valuation-estimation error and an independent validation certificate. Fourth, controlled experiments verify the theoretical identities rather than only their upper bounds: cyclic cells achieve radii from $1/2$ to $15/16$ while the two-cycle bound stays at $1/2$, 240 random cells match the LP and cycle solutions to within 1.1×10^{-15}, and at $K = 16$ Strategic Minimax reduces held-out maximum strategic distortion relative to Anchored

K-means by 13.7% on public projects and 24.3% on bundle selection, while the average-risk results confirm that minimax compression need not dominate objectives it was not designed to optimize.

2 ECONOMIC TASK AND THE INDUCED GEOMETRY

This section derives the geometry relevant to a representation chosen before the downstream priced menu is known. Because priced choices depend only on pairwise valuation differences, valuations are identified up to an additive constant. The resulting strategic distortion is then shown to equal the sharp worst-case regret from acting on a surrogate valuation, which Section 3 lifts from valuation pairs to entire encoder cells.

Let $\mathcal{X} = \{0, \dots, m-1\}$, $m \geq 2$. A type θ induces a valuation $v_\theta \in \mathbb{R}^m$; since only the induced valuation matters, we write simply v. Utilities are quasilinear, $u(x, p) = v(x) - p$, with values and payments in the same monetary unit.

A priced menu is $M = (S, p)$ with $\emptyset \neq S \subseteq \mathcal{X}$ and $p \in \mathbb{R}^S$. Let $\mathcal{M}$ denote the set of all such menus, and define $\mathrm{BR}_v(M) = \arg\max_{x \in S}\{v(x) - p_x\}$ and $\mathrm{Reg}_v(x; M) = \max_{y \in S}\{v(y) - p_y\} - \{v(x) - p_x\}$. Thus x is a δ-best response iff $\mathrm{Reg}_v(x; M) \leq \delta$.

We study deterministic representations chosen before the menu is revealed; the downstream rule may then observe the full menu. Payments are unrestricted real numbers, which is needed for the exact converse statements. Randomized source coding is treated separately in Section 5.

Adding an outcome-independent constant leaves every priced comparison unchanged, so only relative values are strategically relevant.

Definition 2.1 (Strategic distortion)**.** *For $v, w \in \mathbb{R}^m$, define $d_{\mathrm{str}}(v, w) := \max_{x,y \in \mathcal{X}} |[v(y) - v(x)] - [w(y) - w(x)]| = \mathrm{osc}(v - w)$, $\mathrm{osc}(a) := \max_x a_x - \min_x a_x$. Accordingly, the strategic type space is $\mathcal{Q} = \mathbb{R}^m / \mathrm{span}\{\mathbf{1}\}$ with quotient norm $\|[a]\|_{\mathcal{Q}} = \mathrm{osc}(a)$. We use the unique anchored representative $\widetilde{v} = v - v(0)\mathbf{1}$, for which $\widetilde{v}(0) = 0$.*

The quotient geometry also has an exact operational interpretation.

Proposition 2.2 (Exact comparison geometry)**.** *The function d_{str} is a pseudometric on $\mathbb{R}^m$ with $d_{\mathrm{str}}(v, w) = 0$ iff $v - w \in \mathrm{span}\{\mathbf{1}\}$, and hence induces a metric on $\mathcal{Q}$. Moreover,*

$$d_{\mathrm{str}}(v, w) = \sup_{\substack{x,y \in \mathcal{X} \\ p_x, p_y \in \mathbb{R}}} \left| \big([v(y) - p_y] - [v(x) - p_x]\big) - \big([w(y) - p_y] - [w(x) - p_x]\big) \right|.$$

Thus $d_{\mathrm{str}}(v, w)$ is exactly the smallest uniform additive error over all priced pairwise comparisons.

Two further equivalent characterizations of this geometry are collected in Appendix B.1.

Proposition 2.2 identifies the relevant geometry; the next result shows that its distance is exactly the regret penalty from replacing the true valuation by a surrogate.

Theorem 2.3 (Sharp menu-regret transfer)**.** *Let $v, w \in \mathbb{R}^m$ and $\delta \geq 0$. For every menu $M = (S, p) \in \mathcal{M}$ and $x \in S$ with $\mathrm{Reg}_w(x; M) \leq \delta$, $\mathrm{Reg}_v(x; M) \leq \delta + d_{\mathrm{str}}(v, w)$. The bound is sharp: $\sup_{M=(S,p)\in\mathcal{M}} \sup_{\substack{x \in S \\ \mathrm{Reg}_w(x;M) \leq \delta}} \mathrm{Reg}_v(x; M) = \delta + d_{\mathrm{str}}(v, w)$. In particular, for $\delta = 0$, $\sup_{M \in \mathcal{M}} \sup_{x \in \mathrm{BR}_w(M)} \mathrm{Reg}_v(x; M) = d_{\mathrm{str}}(v, w)$.*

Thus $d_{\mathrm{str}}(v, w)$ is the exact universal additive certificate for replacing v by w: absolute utility levels are irrelevant, whereas preserving only unpriced rankings is insufficient. A finite representation, however, merges multiple valuations into one message, so Section 3 turns this pointwise result into a cell-wise minimax characterization.

Proofs of Proposition 2.2 and Theorem 2.3 are deferred to Appendix B.1.

3 A CELL'S OPTIMAL REGRET IS A MAXIMUM CYCLE MEAN

The previous section treats one valuation–surrogate pair. A finite representation instead maps many valuations to one message, so the decoder knows only that the true valuation lies in an encoder cell

C. We now solve this cell-wise deterministic minimax problem exactly and show that arbitrary menu-response rules, optimal surrogate decoding, and a maximum directed cycle mean all have the same value.

Let C be a nonempty bounded set of anchored valuations, so $v(0) = 0$ for every $v \in C$. Let $\mathcal{A} = \{a : \mathcal{M} \to \mathcal{X} : a(S,p) \in S \text{ for every } (S,p) \in \mathcal{M}\}$. Define $\mathfrak{R}(C) := \inf_{a\in\mathcal{A}} \sup_{v\in C} \sup_{M\in\mathcal{M}} \mathrm{Reg}_v(a(M);M)$, $\rho(C) := \inf_{w\in\mathbb{R}^m} \sup_{v\in C} d_{\mathrm{str}}(v,w)$. The center in $\rho(C)$ is an ambient surrogate and need not belong to C; by shift invariance it may be anchored at $w(0) = 0$.

For $x, y \in \mathcal{X}$, let $A_{xy}(C) := \sup_{v\in C}\{v(y) - v(x)\}$, and let Γ be the set of simple directed cycles on $\mathcal{X}$, including zero-weight self-loops.

Theorem 3.1 (Cell minimax–radius–cycle equivalence)**.** *For every nonempty bounded cell C,*

$$\mathfrak{R}(C) = \rho(C) = \max_{\gamma\in\Gamma} \frac{1}{|\gamma|} \sum_{(x,y)\in\gamma} A_{xy}(C). \tag{1}$$

The infimum defining $\rho(C)$ is attained, and an optimal anchored surrogate is obtained from

$$\begin{aligned} \min_{r,w} \quad & r \\ \text{s.t.} \quad & w_y - w_x \geq A_{xy}(C) - r, \quad \forall x, y \in \mathcal{X}, \\ & w_0 = 0. \end{aligned} \tag{2}$$

For any optimizer $w^\star$, deterministic tie-breaking in $\mathrm{BR}_{w^\star}(M)$ attains $\mathfrak{R}(C) = \rho(C)$.

Moreover, if $\gamma^\star$ maximizes equation 1, there exists a priced menu $M_{\gamma^\star}$ such that $\sup_{v\in C} \mathrm{Reg}_v(a(M_{\gamma^\star}); M_{\gamma^\star}) \geq \rho(C)$ for every $a \in \mathcal{A}$. Thus arbitrary deterministic menu-response rules cannot improve on an optimal surrogate decoder.

Theorem 3.1 is both geometric and operational: the maximum cycle mean is the smallest strategic-distortion radius containing the cell and, simultaneously, the unavoidable all-menu regret of any deterministic decoder that observes only the cell identity.

For a target radius r, each ordered pair (x, y) requires the surrogate potential to satisfy $w_y - w_x \geq A_{xy}(C) - r$. These inequalities cannot be checked independently: summing them around a directed cycle cancels the surrogate potentials and leaves a consistency condition on the cell itself. The maximum cycle mean is therefore exactly the smallest common slack that removes every positive cyclic inconsistency. This is why a genuinely multi-outcome obstruction appears even when every pair of valuations in the cell is close.

To compare this characterization with ordinary pairwise geometry, define the strategic diameter $\Delta(C) := \sup_{u,v\in C} d_{\mathrm{str}}(u, v)$. The quantity $\Delta(C)/2$ is exactly the largest two-cycle mean, whereas longer cycles may impose a strictly larger radius.

Corollary 3.2 (Sharp diameter bounds)**.** *For every nonempty bounded cell C, $\frac{1}{2}\Delta(C) \leq \rho(C) \leq \frac{m-1}{m}\Delta(C)$. Both constants are sharp.*

The lower bound can be strict already with three outcomes; an explicit cycle-obstruction example is given in Appendix B.2.3.

In particular, diameter captures only the strongest two-outcome conflict. Longer cycles can combine several individually moderate relative-value demands into a larger global inconsistency, so a clustering objective based only on pairwise diameter can systematically underestimate the regret required to represent the entire cell.

Theorem 3.1 concerns deterministic response rules; randomization can strictly reduce worst-case expected regret, as shown in Appendix B.2.4. Proofs of Theorem 3.1 and Corollary 3.2 are deferred to Appendix B.2.

4 What a certificate implies for mechanisms

We now translate a strategic-distortion certificate into downstream mechanism guarantees. The implication is deliberately one-way: small distortion is sufficient to transfer incentive and welfare guarantees from surrogate to true valuations, but need not be necessary for a fixed mechanism.

Consider n agents with common finite outcome set $\mathcal{X}$. Agent i has type $\theta_i \in \Theta_i$, private valuation $v_i(\theta_i, \cdot) : \mathcal{X} \to \mathbb{R}$, and quasilinear utility $u_i(\theta_i; x, p_i) = v_i(\theta_i, x) - p_i$.

For each agent, let Z_i be a finite label set, $E_i : \Theta_i \to Z_i$ an encoder, and $w_{i,z_i} : \mathcal{X} \to \mathbb{R}$ the surrogate associated with label z_i. Write $Z = \prod_i Z_i$. A latent direct mechanism consists of $x : Z \to \mathcal{X}$ and payments $p_i : Z \to \mathbb{R}$.

We say that the latent mechanism is δ_i-DSIC for agent i on the surrogate label space if, for all $z_i, z_i' \in Z_i$ and $z_{-i} \in Z_{-i}$,

$$w_{i,z_i}(x(z_i', z_{-i})) - p_i(z_i', z_{-i}) \le w_{i,z_i}(x(z_i, z_{-i})) - p_i(z_i, z_{-i}) + \delta_i. \tag{3}$$

The quantifier is over the full report alphabet, since every available label is a feasible deviation.

For a true type θ_i, write $z_i = E_i(\theta_i)$ and define its pointwise certificate by $D_i(\theta_i) := d_{\mathrm{str}}(v_i(\theta_i, \cdot), w_{i,E_i(\theta_i)})$. It bounds every discrepancy between true and surrogate utility gains across two outcomes.

Theorem 4.1 (Incentive and welfare transfer)**.** *Suppose the latent mechanism is δ_i-DSIC in the sense of equation 3. Then, for every θ_i, z_{-i}, and alternative label z_i',*

$$\begin{aligned} &u_i(\theta_i; x(z_i', z_{-i}), p_i(z_i', z_{-i})) \\ &\quad - u_i(\theta_i; x(E_i(\theta_i), z_{-i}), p_i(E_i(\theta_i), z_{-i})) \le \delta_i + D_i(\theta_i). \end{aligned} \tag{4}$$

Hence, if $\sup_{\theta_i} D_i(\theta_i) \le \varepsilon_i$, reporting $E_i(\theta_i)$ is $(\delta_i + \varepsilon_i)$-approximately dominant.

If additionally $x(z) \in \arg\max_{x \in \mathcal{X}} \sum_i w_{i,z_i}(x)$, and $x^\star(\theta) \in \arg\max_{x \in \mathcal{X}} \sum_i v_i(\theta_i, x)$, then

$$\sum_{i=1}^n v_i(\theta_i, x^\star(\theta)) - \sum_{i=1}^n v_i(\theta_i, x(E(\theta))) \le \sum_{i=1}^n D_i(\theta_i). \tag{5}$$

Thus uniform bounds $D_i(\theta_i) \le \varepsilon_i$ imply welfare loss at most $\sum_i \varepsilon_i$.

Theorem 4.1 requires only one distortion term because a unilateral deviation compares exactly two outcomes and the payment difference is identical under the true and surrogate valuations. The welfare bound follows from the same pairwise comparison principle applied agent by agent.

Importantly, this transfer is modular: once a representation is certified in d_{str}, the same certificate can be reused across any downstream latent mechanism satisfying the stated surrogate-side properties. The representation therefore need not be retrained for each allocation rule, which is precisely the advantage of controlling all priced pairwise comparisons before the final mechanism is selected.

A concrete surrogate Groves implementation is stated in Appendix B.3.

These guarantees are sufficient rather than converse statements for arbitrary fixed mechanisms. Likewise, average distortion yields only corresponding average ex-post guarantees and does not imply uniform DSIC or typewise interim BIC. Individual rationality requires an explicit outside-option condition, and because d_{str} removes outcome-independent utility shifts, it does not by itself control absolute payment levels, revenue, budget balance, or fairness.

The proof of Theorem 4.1 and the surrogate Groves specialization are given in Appendix B.3.

5 Information required by a reusable representation

The preceding sections identify the distortion relevant to reusable representations and the exact minimax regret of an encoder cell. We now turn these geometric results into information requirements. We first characterize uniform message complexity by strategic covering numbers, then identify the effective dimension for structured linear valuations, and finally consider the distributional counterpart through finite-alphabet rate–distortion. The first two guarantees are uniform and single-shot; the last is expected and asymptotic.

Throughout, we identify each class in $\mathcal{Q} = \mathbb{R}^m / \operatorname{span}\{\mathbf{1}\}$ with its anchored representative from Section 2. Let $\mathcal{V} \subset \mathcal{Q}$ be nonempty and bounded. For $\varepsilon > 0$, let K_ε be the smallest $K \ge 1$ for

which there exist a deterministic encoder $E : \mathcal{V} \to \{1, \dots, K\}$ and deterministic response rules $a_z : \mathcal{M} \to \mathcal{X}$, with $a_z(S, p) \in S$, such that

$$\sup_{[v]\in\mathcal{V}} \sup_{M\in\mathcal{M}} \mathrm{Reg}_v\big(a_{E([v])}(M); M\big) \le \varepsilon. \tag{6}$$

This regret is invariant to the representative of $[v]$.

For $c \in \mathcal{Q}$, write $\overline{B}_{d_{\mathrm{str}}}(c, \varepsilon) := \{q \in \mathcal{Q} : d_{\mathrm{str}}(q, c) \le \varepsilon\}$. Define the ambient covering number

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) := \min\left\{ K : \mathcal{V} \subseteq \bigcup_{z=1}^{K} \overline{B}_{d_{\mathrm{str}}}(c_z, \varepsilon) \text{ for some } c_1, \dots, c_K \in \mathcal{Q} \right\}. \tag{7}$$

The centers are ambient: they need not belong to $\mathcal{V}$ or to the original parametric valuation family. Boundedness of $\mathcal{V}$ ensures finiteness for every $\varepsilon > 0$.

Theorem 5.1 (Exact message complexity)**.** *For every $\varepsilon > 0$,*

$$K_\varepsilon = N_{\mathrm{amb}}(\mathcal{V}, \varepsilon). \tag{8}$$

Hence the minimum fixed-length binary description has $B_\varepsilon = \lceil \log_2 K_\varepsilon \rceil$ bits. The converse holds even when each label invokes an arbitrary deterministic menu-response rule rather than a surrogate valuation.

Theorem 5.1 lifts the cell-wise characterization of Theorem 3.1 to an exact representation theorem: surrogate decoding is message-optimal among deterministic downstream rules, and the relevant compression problem is covering the quotient valuation space in d_{str}. This accounting concerns only the one-way label sent before the menu is known; it does not include codebook storage, valuation acquisition, or menu-dependent interaction, and therefore is not a lower bound for arbitrary interactive implementations.

The two directions have complementary meanings. A strategic ε-cover directly gives an encoder by assigning each valuation to a nearby center and then best responding to that center. Conversely, any label whose downstream rule achieves uniform regret ε must, by Theorem 3.1, correspond to a cell of strategic radius at most ε. Hence no more elaborate deterministic decoder can pack a strictly larger cell into one message.

Feasibility-induced dimension. The covering formulation also reveals that ambient parameter dimension can overstate communication complexity. Consider $v_\theta(x) = \langle \theta, \psi(x) \rangle$ with $\theta \in \Theta \subseteq \mathbb{R}^d$, and define $L := \mathrm{span}\{\psi(x) - \psi(y) : x, y \in \mathcal{X}\}$, $\quad r := \dim L$. Then $d_{\mathrm{str}}(v_\theta, v_{\theta'}) = \max_{x,y\in\mathcal{X}} |\langle \theta - \theta', \psi(x) - \psi(y) \rangle|$. Thus the strategically null directions are exactly $L^\perp$, while the restriction to L is a norm. Hence r is the largest dimension visible through feasible priced comparisons.

For the next result, let $\mathcal{V}_\Theta = \{[v_\theta] : \theta \in \Theta\}$, and let K_ε refer to the message complexity of this valuation family. Write P_L for orthogonal projection onto L.

Proposition 5.2 (Dimensional scaling under full-dimensional support)**.** *Let $\mathcal{V}_\Theta = \{[v_\theta] : \theta \in \Theta\}$ and let P_L denote orthogonal projection onto L. Suppose $r > 0$, $P_L\Theta$ is bounded, and contains a nonempty r-dimensional relative open ball. Then $\log_2 K_\varepsilon = r \log_2(1/\varepsilon) + O(1)$, $\quad \varepsilon \downarrow 0$. The $O(1)$ term is independent of ε. Without the interior assumption, r alone need not determine the covering exponent.*

An explicit family attaining the r-dimensional scaling, together with its exact ambient covering number, is given in Appendix B.4.

Thus feasibility determines which parameter directions are strategically observable, while type support determines how many such directions must actually be resolved.

In particular, a high-dimensional parameterization need not imply a high communication burden: directions in $L^\perp$ may change the parameter vector substantially while leaving every feasible priced comparison unchanged. The exponent r therefore measures decision-relevant, rather than merely representational, dimension.

Expected certificates and rate–distortion. Uniform covering is appropriate for worst-case guarantees. When valuations instead follow a known distribution and only expected pointwise distortion is required, the corresponding object is a finite-alphabet rate–distortion problem under d_{str}.

Let V be a finite-valued random valuation and $\mathcal{W} \subset \mathbb{R}^m$ a fixed finite reproduction alphabet. Define

$$R_{\mathcal{W}}(D) := \min_{q(w|v):\, \mathbb{E}[d_{\mathrm{str}}(V,W)] \le D} I(V;W), \tag{9}$$

with mutual information measured in bits and value $+\infty$ when the constraint is infeasible. The reproduction alphabet is fixed; optimizing its locations is a different problem.

By Theorem 2.3, $d_{\mathrm{str}}(V,W)$ is the realized worst-case menu loss of using W as a surrogate. Thus equation 9 averages a pointwise all-menu certificate, $\mathbb{E}[\sup_M \mathrm{Reg}(M)]$, which generally differs from $\sup_M \mathbb{E}[\mathrm{Reg}(M)]$. The formulation is therefore an expected certificate, not a Bayesian-mechanism converse.

Proposition 5.3 (Quotient sufficiency for rate–distortion)**.** *For finite V, $R_{\mathcal{W}}(D)$ depends on the source only through the distribution of $[V]$ and is convex and nonincreasing on its effective domain. If $\mathcal{W}$ contains a representative of every quotient class in the support of V, then $R_{\mathcal{W}}(0) = H([V])$. Moreover, the classical finite-alphabet rate–distortion coding theorem applies to equation 9; the precise finite-block converse and asymptotic achievability statement are given in Appendix B.4.*

Proposition 5.3 specializes classical finite-alphabet rate–distortion theory (Shannon, 1993; Davisson, 1972; Cover et al., 1991) to the strategic quotient metric; it is not a new source-coding theorem. The test channel is information-theoretic, and the operational result concerns ordinary long-block codes rather than strategic truthful reporting. Proofs and coding details are in Appendix B.4.

6 Learning and certifying a finite codebook

The preceding results characterize the population problem when the valuation family is known. We now consider finite data. We first define an empirical strategic quantizer, then bound the effect of valuation-estimation error, and finally state what an independent validation sample can certify. These are distinct issues: training optimization and mean generalization alone do not provide worst-case guarantees on unseen types.

Given training valuations $v^1, \dots, v^N \in \mathcal{Q}$ and codebook size K, consider

$$\widehat{\varepsilon}_K := \min_{w_1,\dots,w_K \in \mathcal{Q}} \max_{1 \le j \le N} \min_{1 \le z \le K} d_{\mathrm{str}}(v^j, w_z). \tag{10}$$

By Theorem 5.1, this is the empirical strategic covering radius with K ambient centers; it does not imply coverage of unseen valuations.

Our *strategic minimax quantizer* alternates nearest-center assignments in d_{str} with exact cell-center updates using the linear program equation 2. This assignment–update template is analogous to classical Lloyd-style quantization (Lloyd, 1982), while the minimax objective is closer in spirit to metric k-center formulations (Gonzalez, 1985); here, however, the center update is the exact strategic-radius problem characterized by Theorem 3.1. Each step weakly decreases equation 10; empty-cell centers are unchanged, and iteration stops when assignments stabilize or a prescribed limit is reached. This procedure is not globally optimal in general; for small training sets we use an exponential subset dynamic program as an independent global oracle.

Each alternating iteration is polynomial-time; a direct implementation costs $O(NKm + Nm^2 + Km^3)$ up to lower-order terms. The derivation and numerical cross-checks are given in Appendix B.5.

As throughout the paper, this optimization assumes that the valuation vectors available to the encoder have already been obtained; elicitation or estimation cost is not part of the communication budget.

Stability under valuation estimation. Let $C = \{v^j : j \in I\}$ and $\widetilde{C} = \{\widetilde{v}^j : j \in I\}$ be matched finite cells satisfying $d_{\mathrm{str}}(v^j, \widetilde{v}^j) \le \eta$ for every $j \in I$.

Proposition 6.1 (Stability under valuation estimation)**.** *Under the preceding matched η-perturbation,* $\left|\rho(C) - \rho(\widetilde{C})\right| \le \eta, \qquad \sup_{v \in C} d_{\mathrm{str}}(v, \widetilde{w}^\star) \le \rho(C) + 2\eta$, *where $\widetilde{w}^\star$ is any optimal center of $\widetilde{C}$.*

Moreover, for any fixed codebook $\{w_z\}_{z=1}^K$*, if* $\widehat{z}(\widetilde{v}) \in \arg\min_z d_{\mathrm{str}}(\widetilde{v}, w_z)$ *and* $d_{\mathrm{str}}(v, \widetilde{v}) \le \eta$*, then* $d_{\mathrm{str}}(v, w_{\widehat{z}(\widetilde{v})}) \le \min_z d_{\mathrm{str}}(v, w_z) + 2\eta$.

Thus matched valuation error perturbs the optimal cell radius by at most η, while learning or selecting through perturbed valuations can incur the larger 2η penalty. These guarantees are deterministic and require no sampling assumptions.

What an independent validation sample certifies. Fix an encoder and codebook independently of the validation sample, and let $D(v) := d_{\mathrm{str}}(v, w_{E(v)})$. For i.i.d. validation draws $V_1^{\mathrm{val}}, \dots, V_{N_{\mathrm{val}}}^{\mathrm{val}}$, define $\widehat{D}_{\mathrm{val}} := N_{\mathrm{val}}^{-1} \sum_j D(V_j^{\mathrm{val}})$.

Proposition 6.2 (Independent validation certificate)**.** *Suppose* $0 \le D(V) \le L_0$ *almost surely. Then, for every* $\alpha \in (0,1)$*, with probability at least* $1-\alpha$*,* $\mathbb{E}[D(V)] \le \widehat{D}_{\mathrm{val}} + L_0 \sqrt{\frac{\log(1/\alpha)}{2N_{\mathrm{val}}}}$.

Proposition 6.2 controls population mean distortion, not a uniform worst-case guarantee over unseen types. This distinction separates three levels of certification: the training objective controls only the observed sample, an independent validation set can certify population mean risk, and a uniform guarantee requires geometric coverage of the deployment support. Uniform deployment guarantees therefore need additional coverage assumptions; for example, an h-net extends an empirical covering radius by at most h. The precise statement is given in Appendix B.5.4.

Expected coding. For the distributional objective of Section 5, we solve the fixed-library finite-alphabet rate–distortion problem with Blahut–Arimoto updates (Blahut, 1972; Arimoto, 1972). This optimizes the stochastic encoder for fixed reproduction points, not the decoder locations themselves; implementation and optimality-gap details are given in Appendix A.7.

Proofs of Propositions 6.1 and 6.2, the quantizer complexity derivation, and the finite-net extension are collected in Appendix B.5.

7 Experiments

Our evaluation prioritizes the operational properties of the representation: sharp menu regret, cell certificates, and their stability. A smaller comparison study then tests what these guarantees imply for learned codebooks. All environments are synthetic and provide exact valuation oracles. Appendix A specifies the complete protocol and reports every setting, including cases where minimax compression does not improve average performance.

Why whole-cell geometry matters. We construct cyclic cells with unit diameter and $m = 2, \dots, 16$ outcomes. Their radii, computed independently by a linear program and the cycle algorithm, increase from $1/2$ to $15/16$, while the two-cycle bound remains $1/2$ (Figure 1a, left). Thus the gap is a multi-outcome obstruction, rather than an optimization failure. On 240 random cells, exhaustive cycle enumeration provides a third check; the maximum LP–cycle discrepancy is 1.1×10^{-15}. We also construct the cycle-derived price menu and evaluate *every* possible deterministic action: the best action's worst-type regret equals the LP radius.

Support, frequencies, and sharpness. Holding a two-type cell fixed while increasing one type's frequency leaves the minimax radius unchanged. The mean-center worst-type distortion instead increases from 1 to 1.98 (Figure 1a, right). Separately, 1000 valuation pairs verify the sharp pointwise transfer with explicitly constructed prices; random menu sampling can miss the worst case. Constructive covering/packing checks and perturbed-valuation experiments recover the dimension and stability predictions, including a near-tight 2η encoding example. These checks evaluate decisions or independent optimization problems, rather than only plotting the bounds. Across 20 small codebook problems with an exact subset-DP oracle, Strategic Minimax attains mean and maximum radius-to-optimum ratios of 1.0066 and 1.0409, respectively; remaining local optimization gaps are retained in the appendix.

Matched-budget compression and mechanism transfer. Four agents use learned representations in public-project selection (37 outcomes) and bundle selection with pair interactions (32 outcomes).

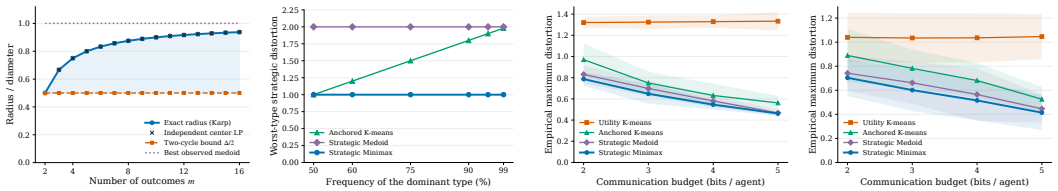

(a) Two distinct properties of universal regret. Left: longer cycles make the exact cell radius strictly larger than half the diameter. Right: changing source frequencies does not change a cell's minimax radius, but can degrade a mean center's worst-type regret.

(b) Held-out maximum strategic distortion under concentrated preferences: public projects (left) and bundle selection (right). Lines show means and bands show one sample standard deviation over five seeds. The same data, label budgets, and monetary scales are used for all methods.

Figure 1: Property verification and matched-budget compression results.

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
Public projects (37 outcomes)					
Utility K-means	1.329 ± 0.059	0.164 ± 0.016	1.045 ± 0.190	0.184 ± 0.030	0.041 ± 0.003
Anchored K-means	0.633 ± 0.111	$\mathbf{0.085 \pm 0.008}$	0.273 ± 0.070	$\mathbf{0.027 \pm 0.005}$	0.047 ± 0.003
Strategic Medoid	0.581 ± 0.055	0.175 ± 0.012	0.237 ± 0.022	0.092 ± 0.012	$\mathbf{0.039 \pm 0.007}$
Strategic Minimax	$\mathbf{0.546 \pm 0.046}$	0.150 ± 0.037	$\mathbf{0.225 \pm 0.077}$	0.074 ± 0.012	0.522 ± 0.102
Bundle selection (32 outcomes)					
Utility K-means	1.036 ± 0.203	0.129 ± 0.017	0.701 ± 0.043	0.146 ± 0.018	0.041 ± 0.005
Anchored K-means	0.681 ± 0.145	$\mathbf{0.075 \pm 0.012}$	0.320 ± 0.053	$\mathbf{0.032 \pm 0.006}$	0.045 ± 0.004
Strategic Medoid	0.565 ± 0.218	0.149 ± 0.019	$\mathbf{0.197 \pm 0.018}$	0.072 ± 0.011	$\mathbf{0.038 \pm 0.009}$
Strategic Minimax	$\mathbf{0.515 \pm 0.167}$	0.120 ± 0.034	0.247 ± 0.096	0.070 ± 0.023	0.429 ± 0.064

Table 1: Concentrated preferences at four bits ($K = 16$): mean $\pm$ sample standard deviation over five seeds. All losses are in monetary units. Maximum distortion and deviation are empirical maxima; welfare loss is a profile average. Bold marks the lowest mean in each environment and column. Fit time includes the warm-start computation for Strategic Minimax.

We compare Strategic Minimax with Utility K-means, Anchored K-means, and Strategic Medoid at $K \in \{4, 8, 16, 32\}$, using matched data and five seeds. The concentrated distribution makes rare preferences relevant to the worst-type objective; a uniform distribution is retained as a control. At $K = 16$, Strategic Minimax reduces mean-over-seed test maxima relative to Anchored K-means by 13.7% and 24.3%, respectively (Figure 1b; Table 1). The improvements over Strategic Medoid are smaller (5.9% and 8.7%). Anchored K-means has lower mean distortion and mean welfare loss in these concentrated settings: the minimax objective does not promise dominance on average-risk metrics. Enumerating all unilateral labels for each sampled profile verifies both incentive and welfare transfer inequalities. The reported test maxima remain finite-sample statistics, not uniform certificates for unseen valuations.

8 Limitations and Conclusion

Our exact converse requires finite feasible outcomes, quasilinear private values, unrestricted relative prices, deterministic response rules, and a representation chosen before the menu. The theory does not infer valuations from behavior, solve strategic elicitation, certify arbitrary neural misreport optimization, or prove communication lower bounds for a single mechanism. Continuous outcome sets, bounded payments, randomized response rules, and distribution-shift guarantees require separate analyses. The present experiments are deliberately small and oracle-based.

Within this scope, a cell's information loss has an exact economic meaning: it is simultaneously the unavoidable deterministic menu regret, a surrogate-reconstruction radius, and a maximum directed cycle mean. This connects reusable economic representations to constructive primal and adversarial certificates, while keeping downstream incentive guarantees distinct from the stronger task that makes nontrivial information lower bounds possible.

9 STATEMENTS

AI USE STATEMENT

In this work, we used generative AI tools primarily for language editing, including grammar correction, improving clarity and fluency, and refining the presentation of the manuscript. We also used generative AI tools to assist with running, debugging, and checking experimental code and computational workflows, as well as to help check mathematical derivations and improve the rigor, consistency, and clarity of proof presentation. Generative AI was not used to determine the core research ideas, theoretical results, experimental design, or scientific conclusions of this work. All AI-assisted mathematical arguments, code execution, outputs, and manuscript edits were independently inspected and verified by the authors. The authors take full responsibility for the accuracy, originality, reproducibility, and final content of the manuscript.

REFERENCES

Sydney N Afriat. The construction of utility functions from expenditure data. *International economic review*, 8(1):67–77, 1967.

Roy Allen and John Rehbeck. Satisficing, aggregation, and quasilinear utility. *Available at SSRN 3180302*, 2020.

Suguru Arimoto. An algorithm for computing the capacity of arbitrary discrete memoryless channels. *IEEE Transactions on Information Theory*, 18(1):14–20, 1972.

Dilip Arumugam and Benjamin Van Roy. Deciding what to learn: A rate-distortion approach. In *International conference on machine learning*, pp. 373–382. PMLR, 2021.

Richard Blahut. Computation of channel capacity and rate-distortion functions. *IEEE transactions on Information Theory*, 18(4):460–473, 1972.

Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization and utility elicitation using the minimax decision criterion. *Artificial Intelligence*, 170(8-9): 686–713, 2006.

Yang Cai and Constantinos Daskalakis. Recommender systems meet mechanism design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pp. 897–914, 2022.

Vincent Conitzer and Tuomas Sandholm. Complexity of mechanism design. *arXiv preprint cs/0205075*, 2002.

Thomas M Cover, Joy A Thomas, and John Kieffer. *Elements of information theory*, volume 2. wiley New York, 1991.

L Davisson. Rate distortion theory: A mathematical basis for data compression. *IEEE Transactions on Communications*, 20(6):1202–1202, 1972.

Shahar Dobzinski. Computational efficiency requires simple taxation. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 209–218. IEEE, 2016.

Shahar Dobzinski, Shiri Ron, and Jan Vondrák. On the hardness of dominant strategy mechanism design. In *Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing*, pp. 690–703, 2022.

Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David Parkes, and Sai Srivatsa Ravindranath. Optimal auctions through deep learning. In *International Conference on Machine Learning*, pp. 1706–1715. PMLR, 2019.

Teofilo F Gonzalez. Clustering to minimize the maximum intercluster distance. *Theoretical computer science*, 38:293–306, 1985.

Jerry Green and Jean-Jacques Laffont. *Limited communication and incentive-compatibility*. Harvard Institute of Economic Research, 1982.

Jerry R Green and Jean-Jacques Laffont. Incentive theory with data compression. *Uncertainty, Information and Communication: Essays in Honor of Kenneth Arrow, edited by Walter P. Heller, Ross M. Starr, and David A. Starrett*, pp. 239–53, 1986.

Sergiu Hart and Noam Nisan. Selling multiple correlated goods: Revenue maximization and menu-size complexity. *Journal of Economic Theory*, 183:991–1029, 2019.

Richard M Karp. A characterization of the minimum cycle mean in a digraph. *Discrete mathematics*, 23(3):309–311, 1978.

Stuart Lloyd. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2): 129–137, 1982.

Roger B Myerson. Optimal auction design. *Mathematics of operations research*, 6(1):58–73, 1981.

Noam Nisan and Ilya Segal. The communication requirements of efficient allocations and supporting prices. *Journal of Economic Theory*, 129(1):192–224, 2006.

Jean-Charles Rochet. A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of mathematical Economics*, 16(2):191–200, 1987.

Claude E Shannon. Coding theorems for a discrete source with a fidelity criterioninstitute of radio engineers, international convention record, vol. 7, 1959. 1993.

Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.

Timothy Van Zandt. Communication complexity and mechanism design. *Journal of the European Economic Association*, 5(2-3):543–553, 2007.

Hal R Varian. The nonparametric approach to demand analysis. *Econometrica: Journal of the Econometric Society*, pp. 945–973, 1982.

Tianhan Wang and Craig Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *Ijcai*, volume 3, pp. 309–316, 2003.

A Experimental details and additional results

A.1 Reproducibility and evaluation scope

The paper configuration uses seeds $\{0, 1, 2, 3, 4\}$; each seed changes the environment features, data draws, and initialization randomness. Unless stated otherwise, plots show the mean and one sample standard deviation across these five complete instances, rather than a standard error or confidence interval. Arithmetic is double precision and BLAS uses one thread. Ties are resolved by the smallest outcome or label index.

These are controlled experiments with known valuation vectors. They assess the geometry and its numerical implementation, not the statistical cost of inferring private values from behavior. A finite-test-set maximum is not equated with a population worst-case guarantee. The validation split used for model selection is not reused to claim the independent-validation certificate of Proposition 6.2.

A.2 Operational regret and independent cell certificates

Pointwise transfer. For each seed and $m \in \{3, 5, 9, 17\}$, we generate 50 anchored pairs. We draw an anchored Gaussian surrogate w and anchored Gaussian error e, then rescale e so that its oscillation is uniform on $[0.05, 2]$ and set $v = w + e$. Thus there are 1000 pairs in total. Let x and y minimize and maximize e, respectively, and write $d = \mathrm{osc}(e)$. For the two-outcome menu, set $p_x = 0$ and $p_y = w_y - w_x + 10^{-6}d$. We evaluate the actual surrogate choice and its true regret. The surrogate has a unique best response and the true regret is $d - 10^{-6}d$. For approximate choice, use $\delta = d/4$, set $p_y = w_y - w_x - \delta$, and explicitly evaluate choosing x; its surrogate regret is δ and its true regret is $d + \delta$. As a diagnostic of finite menu sampling, we also evaluate 64 full menus with independent standard Gaussian prices per pair. These random-menu maxima are lower bounds on the all-menu supremum, not an alternative certificate.

Sharp cyclic cells. For $m = 2, \dots, 16$, take every cyclic rotation of $(0, m-1, m-2, \dots, 1)/m$ and anchor each resulting valuation. We compute the diameter from $A_{xy} + A_{yx}$, solve the anchored center LP, run the maximum-cycle-mean algorithm, evaluate its recovered center, and minimize the maximum distance over all observed medoids. The cycle $0 \to 1 \to \cdots \to m-1 \to 0$ gives a price-menu witness. These cells have unit diameter and radius $(m-1)/m$; all numbers in the figure are evaluated by the implemented routines.

Random cells and adversarial prices. For each seed and $m \in \{3, 4, 5, 6\}$, we generate 12 cells, each containing seven valuations obtained by anchoring independent uniform entries on $[-1, 1]$. This gives 240 cells. Each is solved using three independent routes: the center LP with HiGHS, Karp's dynamic program, and exhaustive enumeration of all simple directed cycles. For a maximizing cycle $(x_0, \dots, x_{\ell-1})$ of mean r, define prices on its vertices by

$$p_{x_0} = 0, \qquad p_{x_{j+1}} - p_{x_j} = A_{x_j x_{j+1}} - r.$$

We then form the entire valuation-by-action regret matrix for this fixed menu and compute $\min_x \max_{v \in C} \mathrm{Reg}_v(x; M)$. The witness therefore checks all deterministic actions on the same menu; it does not choose a different menu after seeing the decoder's action.

Cell support versus empirical frequencies. We hold the cell support fixed at $\{(0, -1), (0, 1)\}$ and use 200 observations. The fraction of $(0, 1)$ is $0.50, 0.60, 0.75, 0.90, 0.95$, or 0.99. All methods receive both types and one label. We compare the Euclidean mean center, the exact observed medoid, and the exact ambient minimax center. The latter retains radius 1; the mean center's radius reaches 1.98. This controlled example isolates a support-based objective from a frequency-weighted objective and is not presented as a broad empirical advantage on unknown real-world data.

A.3 Covering, global optimality, and perturbations

Constructive dimension check. For $r \in \{1, 2, 3, 4\}$ and $q \in \{1, 2, 4, 8\}$, partition $[0, 1]^r$ into q^r axis-aligned cubes and use their midpoints as representatives of $v_a = (0, a_1, -a_1, \dots, a_r, -a_r)$. We evaluate the oscillation error at every vertex of every cube, obtaining radius $1/q$. For $q > 1$, a separate q-point endpoint grid per coordinate gives q^r packing points. A nearest-neighbor computation verifies that their minimum strategic separation is $2/(q-1) > 2/q$. The matching covering and packing counts check both directions at these resolutions. The bit-count plot is backed by these constructions, rather than solely by evaluating the formula for K_ε.

Small-instance global oracle. For each seed, generate eight anchored Gaussian valuations over five outcomes. Evaluate the exact cell radius of every nonempty subset and solve the partition dynamic program for $K = 1, 2, 3, 4$:

$$F_k(S) = \min_{\emptyset \neq T \subseteq S} \max\{\rho(T), F_{k-1}(S \setminus T)\}, \qquad F_1(S) = \rho(S), \quad F_k(\emptyset) = 0.$$

Requiring T to contain the least indexed member of S removes duplicate partitions. We compare learned training radii with this global oracle. For this optimization-only diagnostic the same eight points are used to select restarts; it is not a held-out generalization experiment. Non-unit ratios are retained and explicitly reveal local optimization gaps.

Estimated valuations. For each seed, generate 40 anchored Gaussian valuations over seven outcomes and matched anchored Gaussian perturbations normalized to unit oscillation. Multiply the perturbations by $\eta \in \{0, 0.025, 0.05, 0.1, 0.2, 0.4, 0.8\}$. We solve the original and perturbed cells and evaluate radius change, true-cell error of the estimated center, and encoding error with a fixed eight-center Gaussian codebook. The checks use the bounds $\eta, 2\eta$, and 2η, respectively. For a near-tight encoder construction, use centers $(0, -1), (0, 1)$, true valuation $(0, -\eta + \epsilon)$, and estimate $(0, \epsilon)$, with a small positive ϵ. Direct evaluation approaches excess 2η.

Deterministic scope. For the two valuations $(0, -1), (0, 1)$ and 121 relative prices in $[-1.5, 1.5]$, we enumerate the two deterministic choices and independently solve the two-action lottery LP. The worst-menu values are 1 and $1/2$, respectively. This verifies the importance of the deterministic-response restriction; randomized decoding is not included in the deterministic codebook comparison.

A.4 Comparison environments and data splits

Public projects. There are eight projects and feasible bundles have cardinality at most two, including the empty bundle. Thus there are 37 outcomes. Draw independent $\phi_j \sim \mathcal{N}(0, I_3)$ and let $\psi_{\mathrm{raw}}(S) = \sum_{j \in S} \phi_j$.

Bundle selection with interactions. There are five items and all 32 subsets are feasible. Draw independent $\phi_j, \gamma_{jk} \sim \mathcal{N}(0, I_4)$ and define

$$\psi_{\mathrm{raw}}(S) = \sum_{j \in S} \phi_j + 0.35 \sum_{\{j,k\} \subseteq S} \gamma_{jk}.$$

The pair features introduce complementarity/substitution at the outcome level while retaining a known low-dimensional private type.

For both environments, use $\psi(S) = \psi_{\mathrm{raw}}(S) / \max_T \|\psi_{\mathrm{raw}}(T)\|_1$ and

$$v_{\theta,c}(S) = \langle \theta, \psi(S) \rangle + c.$$

This common normalization is fixed before fitting any method. It bounds the outcome-dependent values by one for $\theta \in [-1, 1]^r$. The empty bundle is outcome zero. Outcome-independent offsets are $c \sim \mathrm{Unif}[-4, 4]$ and are independent of the relevant type. We use two distributions on θ:

$$\begin{aligned} \text{Uniform:} \quad & \mathrm{Unif}([-1,1]^r), \\ \text{Concentrated:} \quad & 0.9\,\mathrm{Unif}([-0.15, 0.15]^r) + 0.1\,\mathrm{Unif}([-1,1]^r). \end{aligned}$$

Both have the same full-dimensional support. Concentration changes how strongly average reconstruction emphasizes the central region. It is a declared stress condition, not a distribution chosen separately for each method. The uniform controls are reported in full.

Each environment/distribution/seed has 768 training, 256 validation, and 1536 test types. The feature map is shared across these splits, but the type draws are independent. Four i.i.d. agents share a population codebook; their private types need not coincide. For mechanism evaluation, we sample 384 four-agent index profiles with replacement from the held-out pool. Every method uses exactly the same profile indices. Profile reuse is respected by aggregating the reported uncertainty across seeds, not by treating all sampled agents as independent experimental replications. The label budgets are $K \in \{4, 8, 16, 32\}$, or two to five bits per agent.

A.5 Methods, optimization, and controls

Utility K-means and Anchored K-means. Both use standard Lloyd iterations, K-means++ initialization, ten restarts, a 300-iteration limit, and tolerance 10^{-4}. These settings are fixed across all instances. Utility K-means receives raw valuations; Anchored K-means receives $v - v(0)\mathbf{1}$. Each uses its standard squared-Euclidean nearest-center encoder.

Strategic Medoid. Starting from the first training valuation, farthest-first traversal in d_{str} selects K centers. Assignments use nearest strategic distance. For a cell, the new center is the observed member minimizing maximum strategic distance to the cell. Its scores are computed exactly from $A_{xy}(C) + u_x - u_y$, avoiding a quadratic table over observations. There is one initialization, at most 12 alternating iterations, and unchanged assignments terminate optimization. Empty centers are retained.

Strategic Minimax. Assignments use d_{str}, and every nonempty cell update has the exact ambient minimax radius. The three initializations are: the same farthest-first codebook as Strategic Medoid; the anchored K-means codebook; and farthest-first traversal starting at a seeded random training valuation. Each receives at most 12 iterations. Among the final candidates, we select the smallest validation maximum, using validation mean distortion only to break ties. Neither test valuations nor mechanism outcomes enter fitting or selection. This extra computation is reported, and a one-initialization control isolates it from the effect of exact cell centers.

The cycle solver computes $r = \rho(C)$ and sets $B_{xy} = A_{xy}(C) - r$. To specify the choice among optimal centers, initialize both potentials at the cell mean $\bar{v}$ and apply $m-1$ simultaneous relaxations:

$$\ell_y \leftarrow \max\{\ell_y, \max_x(\ell_x + B_{xy})\},$$
$$u_x \leftarrow \min\{u_x, \min_y(u_y - B_{xy})\}.$$

Both resulting potentials satisfy the difference constraints. We anchor their midpoint and check that its radius is at most $r + 2 \times 10^{-8}$. This mean-based tie rule does not relax the primary minimax objective. Every restart's training maximum is checked for nonincrease. Independent LP and exhaustive-cycle checks are performed in the diagnostic suite, rather than repeatedly invoking an LP during all codebook fits.

Ablations and offset sensitivity. At $K = 16$, we also evaluate the anchored K-means centers using strategic nearest-center encoding and run Strategic Minimax with only its common farthest-first initialization. These controls distinguish a change of decoder, unrestricted cell centers, and initialization selection. For the offset sweep, reuse public-project concentrated base valuations and multiply their saved uniform offsets by $s \in \{0, 0.5, 1, 2, 4, 8\}$. Utility K-means is refitted at every scale. The anchored and strategic methods reuse fixed codebooks because the quotient input is identical; we directly verify their distortion invariance. These are different, explicitly specified protocols, not equally expensive refits presented as a runtime comparison.

A.6 Exact mechanism evaluation and reported metrics

For each profile, decode the prescribed labels, maximize the sum of surrogate valuations over the complete finite outcome set, and use Groves payments with $h_i = 0$:

$$x(z) \in \arg\max_x \sum_i w_{z_i}(x), \qquad p_i(z) = -\sum_{j \neq i} w_{z_j}(x(z)).$$

For each agent and each of its K reports, recompute the allocation and payment. The maximum gain, including the prescribed report, is therefore nonnegative and exact over the available unilateral alphabet. We also enumerate the true first-best allocation. The tests check

$$G_i \leq D_i, \qquad L_{\text{wel}} \leq \sum_i D_i,$$

as well as zero surrogate deviation gain, up to 10^{-7} numerical tolerance. These inequalities are evaluated for all reports at each sampled profile; we do not enumerate every possible opponents' profile. The ratios in the certificate figures are diagnostic utilization ratios, not confidence levels. A zero-over-zero ratio is recorded as zero.

Representation metrics are mean, 95th-percentile, and maximum distortion over the test pool. Mechanism metrics include mean and maximum deviation, and mean, 95th-percentile, and maximum welfare loss over the profile bank. The primary comparison concerns maximum distortion; the average metrics are retained to show the objective tradeoff. Wall-clock fit time includes all three Strategic Minimax starts and charges it for the anchored K-means warm start, even when that computation is cached in the pipeline. Evaluation and I/O times are saved separately. No significance claims are made from the standard-deviation bands alone.

A.7 Finite-alphabet rate–distortion

The source and reproduction library are both $\{(0, t) : t \in \{-1, -3/4, \ldots, 3/4, 1\}\}$, with uniform source probabilities. The distortion matrix is computed as $|t - \widehat{t}|$. For 40 logarithmically spaced multipliers $\beta \in [0.1, 128]$, initialize the reproduction probabilities uniformly and run Blahut–Arimoto for at most 50,000 iterations, stopping at a feasible primal–dual gap of 10^{-9} nats. The library contains every source type, so zero distortion is representable; this experiment does not optimize a continuous library. For a reproduction marginal q, write $Z_i = \sum_j q_j e^{-\beta d_{ij}}$ and $s_{\max} = \max_j \sum_i p_i e^{-\beta d_{ij}} / Z_i$. The conditional distribution supplies the primal value $I + \beta \mathbb{E}[d]$, and $-\sum_i p_i \log Z_i - \log s_{\max}$ is a feasible dual value. Their gap is saved for every multiplier. Rate is plotted in bits, whereas the Lagrangian gap is measured in nats.

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
Public projects (37 outcomes)					
Utility K-means	1.558 ± 0.156	0.684 ± 0.066	1.328 ± 0.136	0.641 ± 0.044	$\mathbf{0.042 \pm 0.002}$
Anchored K-means	0.577 ± 0.080	0.265 ± 0.029	$\mathbf{0.236 \pm 0.021}$	0.055 ± 0.013	0.046 ± 0.003
Strategic Medoid	0.514 ± 0.048	0.266 ± 0.027	0.249 ± 0.027	0.058 ± 0.005	0.059 ± 0.016
Strategic Minimax	$\mathbf{0.496 \pm 0.065}$	$\mathbf{0.257 \pm 0.027}$	0.253 ± 0.038	$\mathbf{0.053 \pm 0.008}$	0.655 ± 0.115
Bundle selection (32 outcomes)					
Utility K-means	1.258 ± 0.226	0.551 ± 0.083	0.978 ± 0.226	0.510 ± 0.085	$\mathbf{0.045 \pm 0.009}$
Anchored K-means	0.604 ± 0.068	0.262 ± 0.041	0.268 ± 0.059	0.071 ± 0.010	0.047 ± 0.004
Strategic Medoid	$\mathbf{0.504 \pm 0.104}$	0.263 ± 0.044	0.268 ± 0.074	0.069 ± 0.021	0.049 ± 0.015
Strategic Minimax	0.517 ± 0.094	$\mathbf{0.249 \pm 0.040}$	$\mathbf{0.242 \pm 0.046}$	$\mathbf{0.066 \pm 0.018}$	0.578 ± 0.069

Table 2: Uniform-distribution controls at $K = 16$. The conventions match Table 1; all methods and all seeds are retained.

A.8 Complete comparisons and additional figures

The uniform controls and all communication budgets are reported below. In particular, Strategic Medoid slightly improves on Strategic Minimax's four-bit maximum distortion in the uniform bundle environment. The full table retains this result. The concentrated settings also retain Anchored K-means' advantage in mean distortion and welfare loss. The certificate bounds do not impose a ranking of average mechanism performance.

B Proof

B.1 Proofs for Section 2

B.1.1 Proof of Proposition 2.2

Proof. Let $a = v - w$. By Definition 2.1,

$$d_{\mathrm{str}}(v, w) = \mathrm{osc}(a) = \max_{x \in \mathcal{X}} a_x - \min_{x \in \mathcal{X}} a_x.$$

Nonnegativity is immediate. Moreover,

$$\begin{aligned} d_{\mathrm{str}}(w, v) &= \mathrm{osc}(-a) \\ &= \max_x(-a_x) - \min_x(-a_x) \\ &= -\min_x a_x + \max_x a_x \\ &= \max_x a_x - \min_x a_x \\ &= d_{\mathrm{str}}(v, w), \end{aligned}$$

so d_{str} is symmetric, and clearly $d_{\mathrm{str}}(v, v) = 0$. To verify the triangle inequality, let $u, v, w \in \mathbb{R}^m$ and write

$$a := u - v, \qquad b := v - w.$$

For every coordinate x,

$$a_x + b_x \le \max_z a_z + \max_z b_z,$$

and therefore

$$\max_x(a_x + b_x) \le \max_x a_x + \max_x b_x.$$

Likewise, for every x,

$$a_x + b_x \ge \min_z a_z + \min_z b_z,$$

so

$$\min_x(a_x + b_x) \ge \min_x a_x + \min_x b_x.$$

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
2 bits ($K = 4$)					
Utility K-means	1.469 ± 0.064	0.681 ± 0.065	0.984 ± 0.217	0.640 ± 0.061	0.026 ± 0.001
Anchored K-means	0.902 ± 0.073	0.422 ± 0.035	0.503 ± 0.111	0.176 ± 0.044	$\mathbf{0.023 \pm 0.003}$
Strategic Medoid	0.860 ± 0.085	0.428 ± 0.027	0.534 ± 0.097	0.161 ± 0.029	0.049 ± 0.013
Strategic Minimax	$\mathbf{0.778 \pm 0.041}$	$\mathbf{0.411 \pm 0.034}$	$\mathbf{0.456 \pm 0.076}$	$\mathbf{0.154 \pm 0.037}$	0.214 ± 0.029
3 bits ($K = 8$)					
Utility K-means	1.502 ± 0.111	0.682 ± 0.066	1.201 ± 0.115	0.648 ± 0.075	$\mathbf{0.030 \pm 0.003}$
Anchored K-means	0.711 ± 0.153	0.333 ± 0.040	0.374 ± 0.030	0.097 ± 0.016	0.037 ± 0.004
Strategic Medoid	0.664 ± 0.098	0.337 ± 0.034	$\mathbf{0.340 \pm 0.073}$	0.095 ± 0.015	0.060 ± 0.014
Strategic Minimax	$\mathbf{0.616 \pm 0.074}$	$\mathbf{0.326 \pm 0.032}$	0.372 ± 0.057	$\mathbf{0.090 \pm 0.017}$	0.398 ± 0.066
4 bits ($K = 16$)					
Utility K-means	1.558 ± 0.156	0.684 ± 0.066	1.328 ± 0.136	0.641 ± 0.044	$\mathbf{0.042 \pm 0.002}$
Anchored K-means	0.577 ± 0.080	0.265 ± 0.029	$\mathbf{0.236 \pm 0.021}$	0.055 ± 0.013	0.046 ± 0.003
Strategic Medoid	0.514 ± 0.048	0.266 ± 0.027	0.249 ± 0.027	0.058 ± 0.005	0.059 ± 0.016
Strategic Minimax	$\mathbf{0.496 \pm 0.065}$	$\mathbf{0.257 \pm 0.027}$	0.253 ± 0.038	$\mathbf{0.053 \pm 0.008}$	0.655 ± 0.115
5 bits ($K = 32$)					
Utility K-means	1.514 ± 0.160	0.612 ± 0.046	1.060 ± 0.121	0.415 ± 0.031	$\mathbf{0.069 \pm 0.007}$
Anchored K-means	0.484 ± 0.097	0.213 ± 0.023	$\mathbf{0.188 \pm 0.024}$	0.034 ± 0.004	0.084 ± 0.017
Strategic Medoid	0.423 ± 0.052	0.210 ± 0.021	0.194 ± 0.034	0.037 ± 0.004	0.081 ± 0.024
Strategic Minimax	$\mathbf{0.404 \pm 0.054}$	$\mathbf{0.203 \pm 0.022}$	0.200 ± 0.030	$\mathbf{0.034 \pm 0.004}$	1.194 ± 0.178

Table 3: Public projects, uniform preferences: every communication budget and method. Mean $\pm$ sample standard deviation over five seeds. Maximum distortion and deviation are empirical maxima; welfare loss is a profile mean.

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
2 bits ($K = 4$)					
Utility K-means	1.321 ± 0.053	0.160 ± 0.015	0.749 ± 0.046	0.183 ± 0.024	0.027 ± 0.006
Anchored K-means	0.970 ± 0.153	$\mathbf{0.126 \pm 0.022}$	0.375 ± 0.091	$\mathbf{0.087 \pm 0.029}$	$\mathbf{0.021 \pm 0.005}$
Strategic Medoid	0.832 ± 0.084	0.324 ± 0.093	0.409 ± 0.151	0.117 ± 0.026	0.042 ± 0.007
Strategic Minimax	$\mathbf{0.787 \pm 0.066}$	0.296 ± 0.155	$\mathbf{0.303 \pm 0.086}$	0.101 ± 0.013	0.219 ± 0.059
3 bits ($K = 8$)					
Utility K-means	1.325 ± 0.062	0.162 ± 0.016	0.970 ± 0.093	0.186 ± 0.026	0.033 ± 0.005
Anchored K-means	0.750 ± 0.107	$\mathbf{0.103 \pm 0.007}$	0.378 ± 0.105	$\mathbf{0.048 \pm 0.008}$	$\mathbf{0.032 \pm 0.006}$
Strategic Medoid	0.698 ± 0.075	0.200 ± 0.032	$\mathbf{0.236 \pm 0.041}$	0.103 ± 0.005	0.036 ± 0.011
Strategic Minimax	$\mathbf{0.648 \pm 0.094}$	0.198 ± 0.042	0.290 ± 0.190	0.104 ± 0.018	0.241 ± 0.043
4 bits ($K = 16$)					
Utility K-means	1.329 ± 0.059	0.164 ± 0.016	1.045 ± 0.190	0.184 ± 0.030	0.041 ± 0.003
Anchored K-means	0.633 ± 0.111	$\mathbf{0.085 \pm 0.008}$	0.273 ± 0.070	$\mathbf{0.027 \pm 0.005}$	0.047 ± 0.003
Strategic Medoid	0.581 ± 0.055	0.175 ± 0.012	0.237 ± 0.022	0.092 ± 0.012	$\mathbf{0.039 \pm 0.007}$
Strategic Minimax	$\mathbf{0.546 \pm 0.046}$	0.150 ± 0.037	$\mathbf{0.225 \pm 0.077}$	0.074 ± 0.012	0.522 ± 0.102
5 bits ($K = 32$)					
Utility K-means	1.334 ± 0.086	0.166 ± 0.016	1.097 ± 0.205	0.178 ± 0.026	$\mathbf{0.069 \pm 0.003}$
Anchored K-means	0.562 ± 0.067	$\mathbf{0.067 \pm 0.007}$	0.205 ± 0.046	$\mathbf{0.016 \pm 0.003}$	0.087 ± 0.011
Strategic Medoid	0.468 ± 0.030	0.112 ± 0.023	0.207 ± 0.071	0.054 ± 0.021	0.075 ± 0.004
Strategic Minimax	$\mathbf{0.462 \pm 0.020}$	0.101 ± 0.032	$\mathbf{0.191 \pm 0.043}$	0.048 ± 0.032	1.174 ± 0.198

Table 4: Public projects, concentrated preferences: every communication budget and method. Mean $\pm$ sample standard deviation over five seeds. Maximum distortion and deviation are empirical maxima; welfare loss is a profile mean.

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
2 bits ($K = 4$)					
Utility K-means	1.191 ± 0.190	0.550 ± 0.084	0.744 ± 0.184	0.538 ± 0.083	$\mathbf{0.021 \pm 0.005}$
Anchored K-means	0.918 ± 0.162	0.377 ± 0.058	0.604 ± 0.202	0.164 ± 0.037	0.024 ± 0.004
Strategic Medoid	0.758 ± 0.100	0.389 ± 0.061	$\mathbf{0.352 \pm 0.072}$	$\mathbf{0.155 \pm 0.022}$	0.029 ± 0.004
Strategic Minimax	$\mathbf{0.689 \pm 0.085}$	$\mathbf{0.363 \pm 0.050}$	0.454 ± 0.125	0.156 ± 0.028	0.155 ± 0.023
3 bits ($K = 8$)					
Utility K-means	1.185 ± 0.199	0.551 ± 0.084	0.976 ± 0.280	0.529 ± 0.097	$\mathbf{0.028 \pm 0.004}$
Anchored K-means	0.765 ± 0.136	0.309 ± 0.048	0.319 ± 0.060	0.108 ± 0.016	0.032 ± 0.006
Strategic Medoid	0.622 ± 0.087	0.317 ± 0.054	$\mathbf{0.315 \pm 0.050}$	0.104 ± 0.023	0.036 ± 0.007
Strategic Minimax	$\mathbf{0.574 \pm 0.061}$	$\mathbf{0.301 \pm 0.047}$	0.316 ± 0.102	$\mathbf{0.097 \pm 0.020}$	0.335 ± 0.071
4 bits ($K = 16$)					
Utility K-means	1.258 ± 0.226	0.551 ± 0.083	0.978 ± 0.226	0.510 ± 0.085	$\mathbf{0.045 \pm 0.009}$
Anchored K-means	0.604 ± 0.068	0.262 ± 0.041	0.268 ± 0.059	0.071 ± 0.010	0.047 ± 0.004
Strategic Medoid	$\mathbf{0.504 \pm 0.104}$	0.263 ± 0.044	0.268 ± 0.074	0.069 ± 0.021	0.049 ± 0.015
Strategic Minimax	0.517 ± 0.094	$\mathbf{0.249 \pm 0.040}$	$\mathbf{0.242 \pm 0.046}$	$\mathbf{0.066 \pm 0.018}$	0.578 ± 0.069
5 bits ($K = 32$)					
Utility K-means	1.219 ± 0.189	0.531 ± 0.068	0.949 ± 0.243	0.419 ± 0.044	$\mathbf{0.067 \pm 0.007}$
Anchored K-means	0.507 ± 0.070	0.219 ± 0.035	0.245 ± 0.061	0.050 ± 0.004	0.072 ± 0.006
Strategic Medoid	0.422 ± 0.072	0.221 ± 0.033	$\mathbf{0.200 \pm 0.033}$	0.051 ± 0.007	0.075 ± 0.017
Strategic Minimax	$\mathbf{0.408 \pm 0.075}$	$\mathbf{0.208 \pm 0.032}$	0.235 ± 0.036	$\mathbf{0.047 \pm 0.006}$	0.896 ± 0.074

Table 5: Bundle selection, uniform preferences: every communication budget and method. Mean $\pm$ sample standard deviation over five seeds. Maximum distortion and deviation are empirical maxima; welfare loss is a profile mean.

Method	Max. dist. ↓	Mean dist. ↓	Max. dev. ↓	Welfare loss ↓	Fit (s) ↓
2 bits ($K = 4$)					
Utility K-means	1.040 ± 0.205	0.126 ± 0.017	0.531 ± 0.122	0.143 ± 0.031	0.018 ± 0.002
Anchored K-means	0.889 ± 0.223	$\mathbf{0.104 \pm 0.019}$	0.391 ± 0.082	$\mathbf{0.078 \pm 0.022}$	$\mathbf{0.016 \pm 0.002}$
Strategic Medoid	0.742 ± 0.146	0.229 ± 0.143	$\mathbf{0.305 \pm 0.081}$	0.100 ± 0.023	0.024 ± 0.007
Strategic Minimax	$\mathbf{0.704 \pm 0.152}$	0.234 ± 0.086	0.324 ± 0.080	0.099 ± 0.016	0.122 ± 0.030
3 bits ($K = 8$)					
Utility K-means	1.035 ± 0.205	0.127 ± 0.018	0.616 ± 0.161	0.144 ± 0.031	0.029 ± 0.002
Anchored K-means	0.782 ± 0.157	$\mathbf{0.089 \pm 0.013}$	0.369 ± 0.084	$\mathbf{0.047 \pm 0.008}$	0.027 ± 0.002
Strategic Medoid	0.662 ± 0.176	0.145 ± 0.020	$\mathbf{0.212 \pm 0.018}$	0.086 ± 0.010	$\mathbf{0.027 \pm 0.006}$
Strategic Minimax	$\mathbf{0.602 \pm 0.174}$	0.145 ± 0.026	0.277 ± 0.035	0.077 ± 0.014	0.174 ± 0.027
4 bits ($K = 16$)					
Utility K-means	1.036 ± 0.203	0.129 ± 0.017	0.701 ± 0.043	0.146 ± 0.018	0.041 ± 0.005
Anchored K-means	0.681 ± 0.145	$\mathbf{0.075 \pm 0.012}$	0.320 ± 0.053	$\mathbf{0.032 \pm 0.006}$	0.045 ± 0.004
Strategic Medoid	0.565 ± 0.218	0.149 ± 0.019	$\mathbf{0.197 \pm 0.018}$	0.072 ± 0.011	$\mathbf{0.038 \pm 0.009}$
Strategic Minimax	$\mathbf{0.515 \pm 0.167}$	0.120 ± 0.034	0.247 ± 0.096	0.070 ± 0.023	0.429 ± 0.064
5 bits ($K = 32$)					
Utility K-means	1.046 ± 0.189	0.134 ± 0.017	0.743 ± 0.084	0.149 ± 0.025	0.069 ± 0.009
Anchored K-means	0.526 ± 0.102	$\mathbf{0.065 \pm 0.008}$	0.248 ± 0.086	$\mathbf{0.024 \pm 0.006}$	0.075 ± 0.005
Strategic Medoid	0.446 ± 0.116	0.111 ± 0.016	0.219 ± 0.055	0.068 ± 0.008	$\mathbf{0.059 \pm 0.015}$
Strategic Minimax	$\mathbf{0.414 \pm 0.144}$	0.088 ± 0.023	$\mathbf{0.185 \pm 0.102}$	0.035 ± 0.015	0.804 ± 0.159

Table 6: Bundle selection, concentrated preferences: every communication budget and method. Mean $\pm$ sample standard deviation over five seeds. Maximum distortion and deviation are empirical maxima; welfare loss is a profile mean.

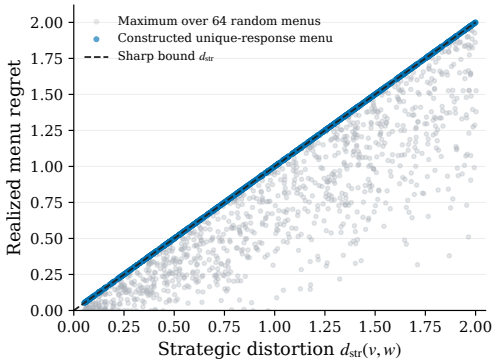

(a) Constructed prices approach the all-menu supremum; sampling a finite random menu bank can underestimate it. The construction uses a positive unique-response margin.

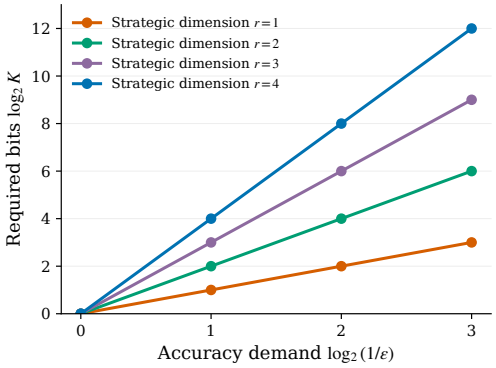

(b) Covering and packing counts agree. Every constructed cell vertex is checked, and the minimum packing separation is computed independently.

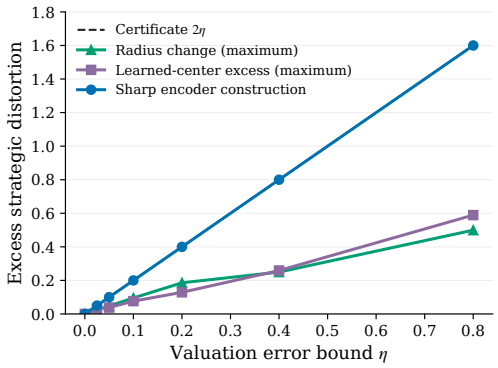

(c) Perturbation bounds are checked directly; a two-outcome construction nearly attains the factor two for estimated encoding.

Figure 2: Basic property checks for sharp menu regret, covering complexity, and perturbation stability.

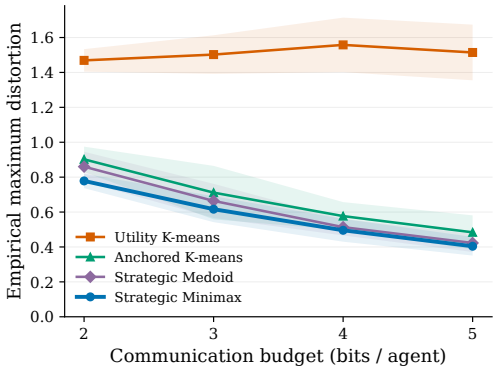

(a) Public, uniform preferences: empirical maximum distortion. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

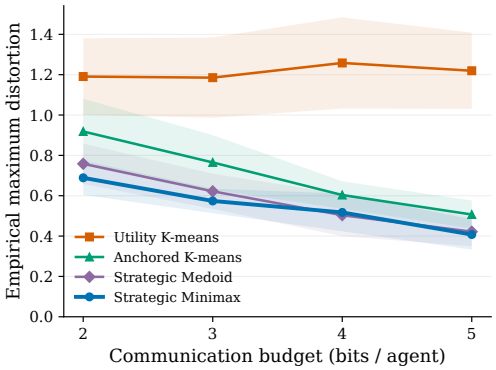

(b) Bundles, uniform preferences: empirical maximum distortion. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

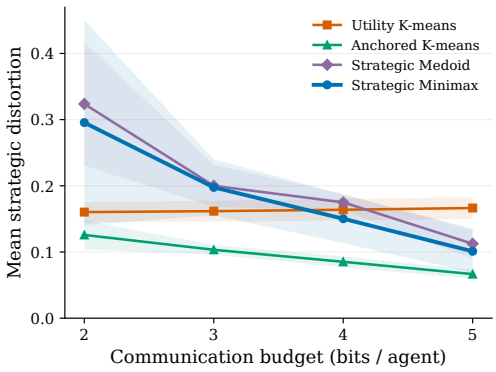

(c) Public, concentrated preferences: mean strategic distortion. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

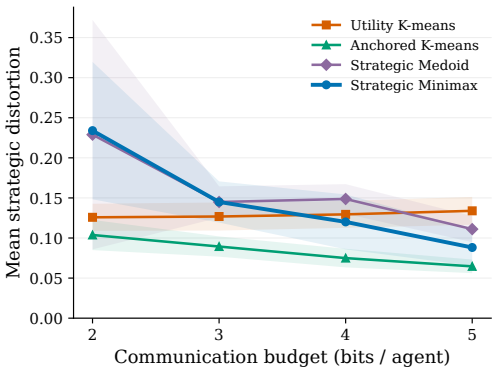

(d) Bundles, concentrated preferences: mean strategic distortion. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

Figure 3: Empirical strategic distortion across public-project and bundle domains.

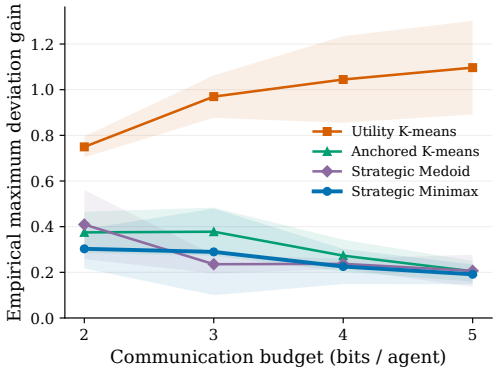

(a) Public, concentrated preferences: empirical maximum deviation gain. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

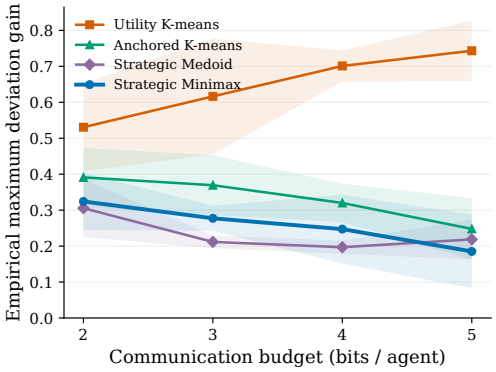

(b) Bundles, concentrated preferences: empirical maximum deviation gain. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

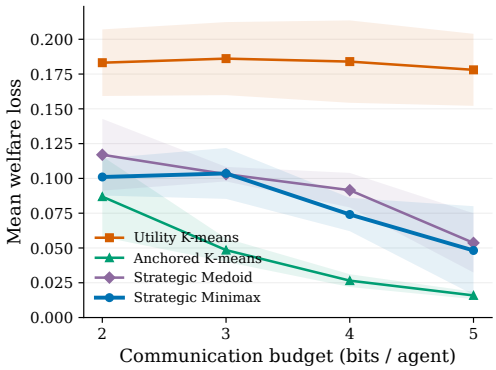

(c) Public, concentrated preferences: mean welfare loss. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

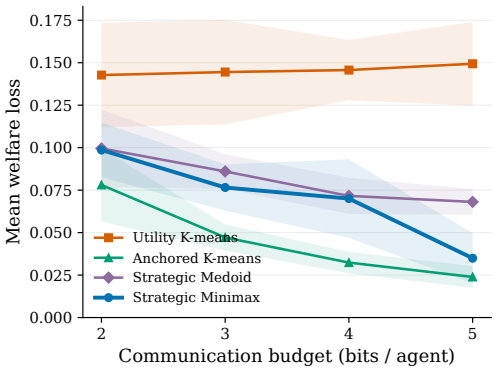

(d) Bundles, concentrated preferences: mean welfare loss. Lines are seed means; bands are one sample standard deviation. Test maxima concern the finite held-out set.

Figure 4: Mechanism-level consequences of strategic compression: deviation gains and welfare losses in the public-project and bundle domains.

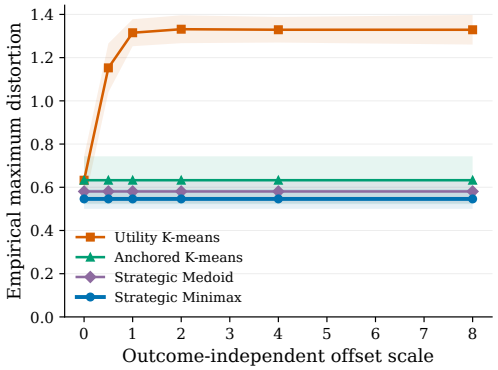

(a) Public-project concentrated distribution, $K = 16$. Utility K-means is refitted at each offset; the invariant methods use fixed anchored codebooks on identical relative valuations.

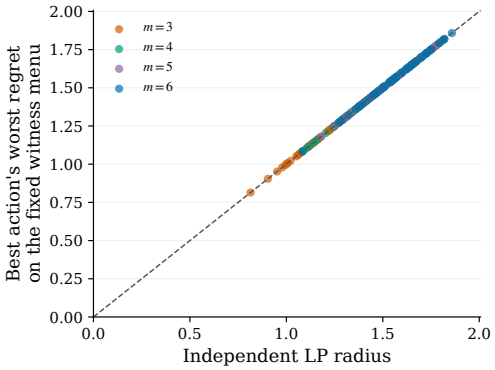

(b) The same cycle-derived menu is evaluated against every deterministic action. Its minimax decision regret agrees with the independent center LP.

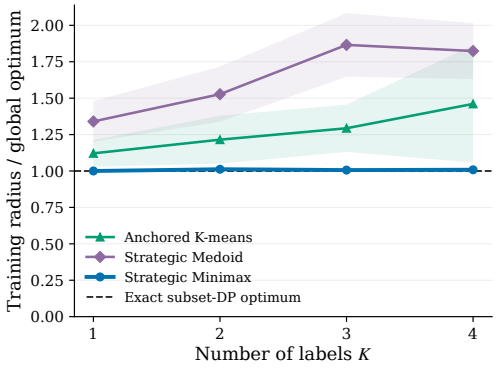

(c) Eight-point cells with five outcomes. Subset dynamic programming gives the global K-center optimum. Alternating optimization is not assumed globally optimal.

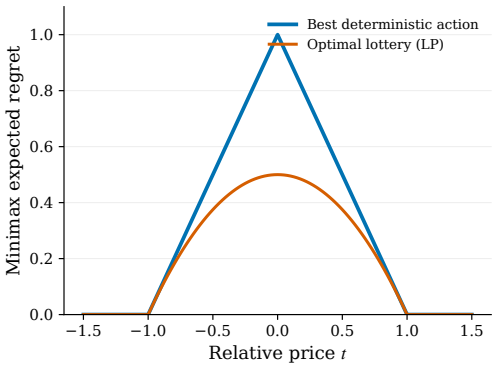

(d) The two-type example demonstrates why the deterministic-response restriction matters: worst-menu values are 1 and $1/2$.

Figure 5: Additional structural checks: gauge invariance, cell certificates, global optimization, and the scope of deterministic-response guarantees.

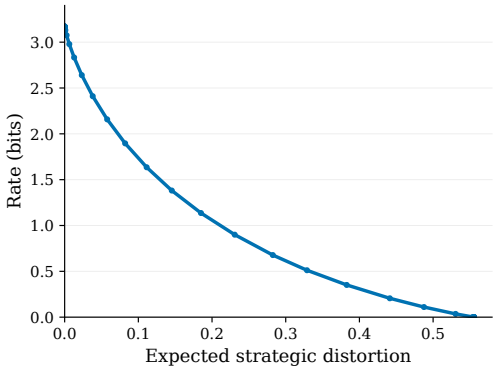

(a) Blahut–Arimoto curve for a fixed nine-point source and the same nine-point reproduction library; no continuous-alphabet optimality claim.

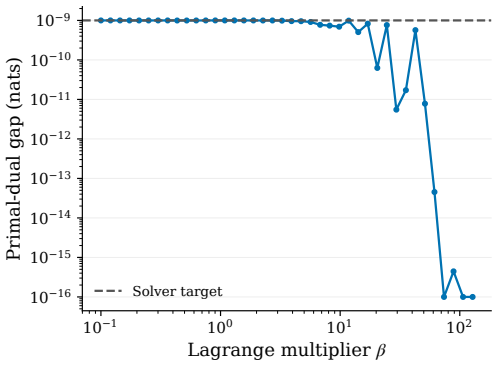

(b) Feasible primal–dual Lagrangian gaps. Values below 10^{-16} are drawn at that display floor, without modifying stored results.

Figure 6: Finite-library rate–distortion computation and numerical optimality checks.

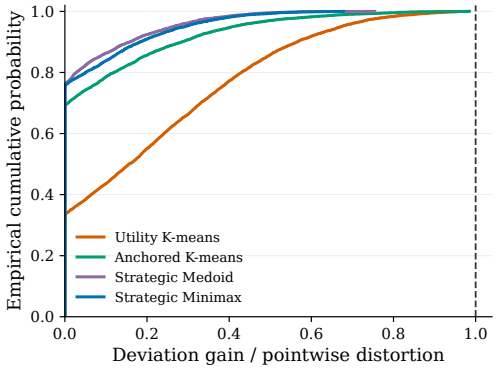

(a) Public-project concentrated preferences, $K = 16$, all seeds. Every unilateral report is enumerated. The vertical line is the theoretical upper bound.

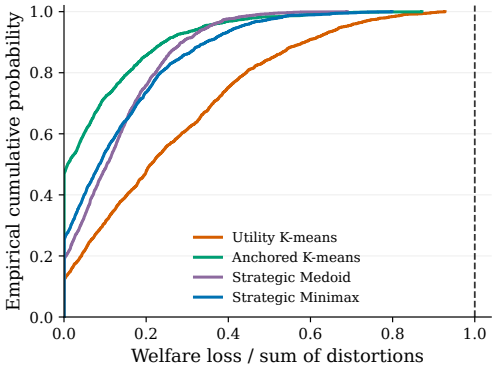

(b) Public-project concentrated preferences, $K = 16$, all seeds. Every unilateral report is enumerated. The vertical line is the theoretical upper bound.

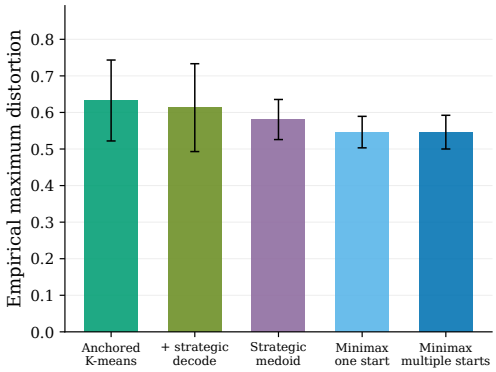

(c) Public, concentrated preferences, $K = 16$. Decoding and initialization controls; error bars are sample standard deviations across seeds.

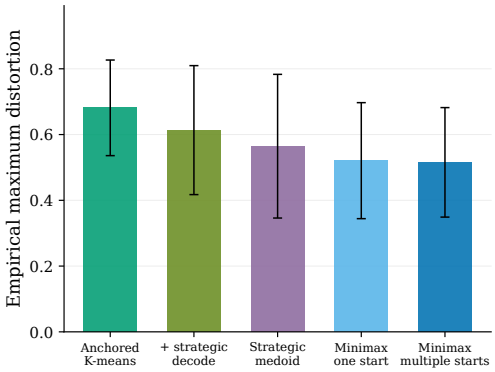

(d) Bundles, concentrated preferences, $K = 16$. Decoding and initialization controls; error bars are sample standard deviations across seeds.

Figure 7: Mechanism certificates and implementation ablations. The certificate experiments exhaustively enumerate unilateral reports, while the ablations isolate decoding and initialization choices.

Subtracting the latter inequality from the former gives

$$\begin{aligned} d_{\mathrm{str}}(u,w) &= \mathrm{osc}(u-w) \\ &= \mathrm{osc}\big((u-v)+(v-w)\big) \\ &= \mathrm{osc}(a+b) \\ &= \max_x(a_x+b_x) - \min_x(a_x+b_x) \\ &\le \big(\max_x a_x - \min_x a_x\big) + \big(\max_x b_x - \min_x b_x\big) \\ &= \mathrm{osc}(a) + \mathrm{osc}(b) \\ &= d_{\mathrm{str}}(u,v) + d_{\mathrm{str}}(v,w). \end{aligned}$$

Hence d_{str} is a pseudometric on $\mathbb{R}^m$.

We next characterize its zero-distance classes. By definition,

$$d_{\mathrm{str}}(v,w) = 0 \quad \Longleftrightarrow \quad \max_x(v_x - w_x) - \min_x(v_x - w_x) = 0.$$

Since the maximum is always at least the minimum, this is equivalent to

$$\max_x(v_x - w_x) = \min_x(v_x - w_x).$$

Because $\mathcal{X}$ is finite, the maximum and minimum are attained. Equality of the two therefore holds if and only if every coordinate of $v - w$ is equal to one common scalar c. Thus

$$d_{\mathrm{str}}(v,w) = 0 \quad \Longleftrightarrow \quad v - w = c\mathbf{1} \quad \text{for some } c \in \mathbb{R} \quad \Longleftrightarrow \quad v - w \in \mathrm{span}\{\mathbf{1}\}.$$

It remains to verify that d_{str} is well defined on the quotient $\mathcal{Q} = \mathbb{R}^m / \mathrm{span}\{\mathbf{1}\}$. Suppose that v' and w' are different representatives of the same respective quotient classes as v and w. Then there exist $c, c' \in \mathbb{R}$ such that

$$v' = v + c\mathbf{1}, \qquad w' = w + c'\mathbf{1}.$$

Consequently,

$$v' - w' = (v - w) + (c - c')\mathbf{1}.$$

Adding the same scalar to every coordinate shifts the maximum and the minimum by the same amount, so

$$\begin{aligned} \mathrm{osc}\big((v-w)+(c-c')\mathbf{1}\big) &= \Big(\max_x(v_x - w_x) + (c - c')\Big) \\ &\quad - \Big(\min_x(v_x - w_x) + (c - c')\Big) \\ &= \max_x(v_x - w_x) - \min_x(v_x - w_x) \\ &= \mathrm{osc}(v - w). \end{aligned}$$

Hence

$$d_{\mathrm{str}}(v', w') = d_{\mathrm{str}}(v, w),$$

and therefore

$$d_{\mathrm{str}}([v],[w]) := d_{\mathrm{str}}(v,w)$$

is well defined on $\mathcal{Q}$. The preceding zero-distance characterization shows that

$$d_{\mathrm{str}}([v],[w]) = 0 \quad \Longleftrightarrow \quad [v] = [w].$$

Symmetry and the triangle inequality are inherited from the pseudometric on $\mathbb{R}^m$. Thus the induced function is a metric on $\mathcal{Q}$.

For the priced-comparison identity, fix arbitrary $x, y \in \mathcal{X}$ and arbitrary prices $p_x, p_y \in \mathbb{R}$. The difference between the two priced pairwise comparisons is

$$\begin{aligned} &\big([v(y) - p_y] - [v(x) - p_x]\big) - \big([w(y) - p_y] - [w(x) - p_x]\big) \\ &\quad = v(y) - p_y - v(x) + p_x - w(y) + p_y + w(x) - p_x \\ &\quad = v(y) - v(x) - w(y) + w(x) \\ &\quad = [v(y) - w(y)] - [v(x) - w(x)] \\ &\quad = a_y - a_x. \end{aligned}$$

In particular, the price terms cancel exactly, so the resulting error does not depend on p_x or p_y. It follows that

$$\sup_{\substack{x,y\in\mathcal{X}\\ p_x,p_y\in\mathbb{R}}} \left|\big([v(y)-p_y]-[v(x)-p_x]\big)-\big([w(y)-p_y]-[w(x)-p_x]\big)\right| = \max_{x,y\in\mathcal{X}} |a_y-a_x|.$$

Let

$$a_{\max} := \max_{z\in\mathcal{X}} a_z, \qquad a_{\min} := \min_{z\in\mathcal{X}} a_z.$$

For every $x, y \in \mathcal{X}$,

$$a_y - a_x \le a_{\max} - a_{\min}$$

and

$$a_x - a_y \le a_{\max} - a_{\min},$$

so

$$|a_y - a_x| \le a_{\max} - a_{\min}.$$

On the other hand, because $\mathcal{X}$ is finite, there exist $x_{\min}, x_{\max} \in \mathcal{X}$ such that

$$a_{x_{\min}} = a_{\min}, \qquad a_{x_{\max}} = a_{\max}.$$

For this pair,

$$|a_{x_{\max}} - a_{x_{\min}}| = a_{\max} - a_{\min}.$$

Therefore

$$\max_{x,y\in\mathcal{X}} |a_y - a_x| = a_{\max} - a_{\min} = \operatorname{osc}(a) = d_{\mathrm{str}}(v,w).$$

Thus $d_{\mathrm{str}}(v,w)$ is a valid uniform additive bound for every priced pairwise comparison, and no smaller constant can be a uniform bound because a maximizing pair of outcomes attains exactly $d_{\mathrm{str}}(v,w)$. This proves the claimed exact comparison geometry. □

B.1.2 Additional geometric characterizations

Proposition B.1 (Auxiliary quotient identities). *For any $v, w \in \mathbb{R}^m$,*

$$d_{\mathrm{str}}(v,w) = 2 \inf_{c\in\mathbb{R}} \|v - w - c\mathbf{1}\|_\infty .$$

Moreover,

$$\mathrm{BR}_v(M) = \mathrm{BR}_w(M) \quad \textit{for every } M \in \mathcal{M} \quad \Longleftrightarrow \quad v - w \in \operatorname{span}\{\mathbf{1}\}.$$

Hence exact preservation of all priced best-response correspondences requires distinguishing distinct elements of $\mathcal{Q}$.

Proof. Let

$$a := v - w, \qquad \alpha := \min_{x\in\mathcal{X}} a_x, \qquad \beta := \max_{x\in\mathcal{X}} a_x.$$

For any $c \in \mathbb{R}$,

$$\|a - c\mathbf{1}\|_\infty = \max_{x\in\mathcal{X}} |a_x - c|.$$

Since α and β are attained by coordinates of a, the maximum on the right must in particular dominate the deviations of these two extreme coordinates. Hence

$$\|a - c\mathbf{1}\|_\infty \ge \max\{|\alpha - c|, |\beta - c|\}.$$

The two endpoint deviations cannot both be smaller than half the length of the interval $[\alpha, \beta]$. Indeed,

$$\begin{aligned}\beta - \alpha &= |(\beta - c) - (\alpha - c)| \\ &\le |\beta - c| + |\alpha - c| \\ &\le 2\max\{|\alpha - c|, |\beta - c|\}.\end{aligned}$$

Therefore, for every $c \in \mathbb{R}$,

$$\|a - c\mathbf{1}\|_\infty \ge \frac{\beta - \alpha}{2}.$$

Taking the infimum over c gives

$$\inf_{c\in\mathbb{R}} \|a - c\mathbf{1}\|_\infty \geq \frac{\beta-\alpha}{2}.$$

For the reverse inequality, choose the midpoint

$$c^* := \frac{\alpha+\beta}{2}.$$

For every $x \in \mathcal{X}$, the definition of α and β gives

$$\alpha \leq a_x \leq \beta.$$

Subtracting c^* yields

$$\alpha - \frac{\alpha+\beta}{2} \leq a_x - c^* \leq \beta - \frac{\alpha+\beta}{2},$$

or equivalently,

$$-\frac{\beta-\alpha}{2} \leq a_x - c^* \leq \frac{\beta-\alpha}{2}.$$

Hence

$$|a_x - c^*| \leq \frac{\beta-\alpha}{2} \qquad \text{for every } x \in \mathcal{X},$$

and therefore

$$\|a - c^*\mathbf{1}\|_\infty \leq \frac{\beta-\alpha}{2}.$$

Combining this with the lower bound valid for every c gives

$$\inf_{c\in\mathbb{R}} \|a - c\mathbf{1}\|_\infty = \frac{\beta-\alpha}{2}.$$

Since

$$\beta - \alpha = \mathrm{osc}(a) = d_{\mathrm{str}}(v, w),$$

we conclude that

$$d_{\mathrm{str}}(v, w) = 2 \inf_{c\in\mathbb{R}} \|v - w - c\mathbf{1}\|_\infty .$$

We now prove the best-response characterization. Suppose first that

$$v - w = c\mathbf{1}$$

for some $c \in \mathbb{R}$. Then, for every menu $M = (S, p) \in \mathcal{M}$ and every $x \in S$,

$$v(x) - p_x = w(x) + c - p_x = [w(x) - p_x] + c.$$

Thus the entire collection of feasible utilities under v is obtained from the corresponding collection under w by adding the same constant c to every outcome. Adding a common constant cannot change any maximizer, and so

$$\mathrm{BR}_v(M) = \mathrm{BR}_w(M) \qquad \text{for every } M \in \mathcal{M}.$$

Conversely, suppose that

$$v - w \notin \mathrm{span}\{\mathbf{1}\}.$$

Then $a = v - w$ is not constant. Hence there exist distinct $x, y \in \mathcal{X}$ such that

$$a_y \neq a_x.$$

Define the corresponding unpriced pairwise differences by

$$\Delta_v := v(y) - v(x), \qquad \Delta_w := w(y) - w(x).$$

Their difference is

$$\begin{aligned}\Delta_v - \Delta_w &= [v(y) - v(x)] - [w(y) - w(x)] \\ &= [v(y) - w(y)] - [v(x) - w(x)] \\ &= a_y - a_x,\end{aligned}$$

which is nonzero. Thus

$$\Delta_v \neq \Delta_w.$$

Set

$$t := \frac{\Delta_v + \Delta_w}{2}$$

and consider the two-outcome menu

$$M = (\{x, y\}, p), \qquad p_x := 0, \qquad p_y := t.$$

Under valuation v, the utility advantage of y over x is

$$\begin{aligned}[v(y) - p_y] - [v(x) - p_x] &= v(y) - v(x) - t \\ &= \Delta_v - \frac{\Delta_v + \Delta_w}{2} \\ &= \frac{\Delta_v - \Delta_w}{2}.\end{aligned}$$

Under valuation w, the corresponding utility advantage is

$$\begin{aligned}[w(y) - p_y] - [w(x) - p_x] &= w(y) - w(x) - t \\ &= \Delta_w - \frac{\Delta_v + \Delta_w}{2} \\ &= -\frac{\Delta_v - \Delta_w}{2}.\end{aligned}$$

Because $\Delta_v - \Delta_w \neq 0$, these two quantities are both nonzero and have opposite signs. If $\Delta_v > \Delta_w$, then

$$[v(y) - p_y] - [v(x) - p_x] > 0$$

while

$$[w(y) - p_y] - [w(x) - p_x] < 0,$$

so

$$\mathrm{BR}_v(M) = \{y\}, \qquad \mathrm{BR}_w(M) = \{x\}.$$

If $\Delta_v < \Delta_w$, the inequalities and therefore the two singleton best-response sets are reversed. In either case,

$$\mathrm{BR}_v(M) \neq \mathrm{BR}_w(M).$$

Thus equality of the best-response correspondence for every priced menu cannot hold when $v - w$ is nonconstant. We have therefore proved

$$\mathrm{BR}_v(M) = \mathrm{BR}_w(M) \quad \text{for every } M \in \mathcal{M} \quad \Longleftrightarrow \quad v - w \in \mathrm{span}\{\mathbf{1}\}.$$

Since two valuations represent the same element of $\mathcal{Q}$ exactly when they differ by an element of $\mathrm{span}\{\mathbf{1}\}$, exact preservation of all priced best-response correspondences requires distinguishing distinct elements of $\mathcal{Q}$. □

B.1.3 Proof of Theorem 2.3

Proof. Fix $v, w \in \mathbb{R}^m$, $\delta \geq 0$, a menu $M = (S, p) \in \mathcal{M}$, and an action $x \in S$ such that

$$\mathrm{Reg}_w(x; M) \leq \delta.$$

Write

$$d := d_{\mathrm{str}}(v, w).$$

For any $y \in S$, the utility difference under v can be decomposed by adding and subtracting the corresponding valuation difference under w:

$$\begin{aligned}[v(y) - p_y] - [v(x) - p_x] &= v(y) - v(x) - (p_y - p_x) \\ &= w(y) - w(x) - (p_y - p_x) \\ &\quad + [v(y) - v(x)] - [w(y) - w(x)] \\ &= [w(y) - p_y] - [w(x) - p_x] \\ &\quad + \big([v(y) - v(x)] - [w(y) - w(x)]\big).\end{aligned}$$

By the definition of regret,

$$\mathrm{Reg}_w(x; M) = \max_{z \in S} \big([w(z) - p_z] - [w(x) - p_x] \big).$$

Since this maximum is at most δ, every term being maximized is also at most δ. Hence, for every $y \in S$,

$$[w(y) - p_y] - [w(x) - p_x] \le \delta.$$

Moreover, by the definition of strategic distortion,

$$d = \max_{r,s \in \mathcal{X}} \left| [v(s) - v(r)] - [w(s) - w(r)] \right|.$$

Therefore the particular pair (x, y) satisfies

$$\left| [v(y) - v(x)] - [w(y) - w(x)] \right| \le d,$$

and hence, in particular,

$$[v(y) - v(x)] - [w(y) - w(x)] \le d.$$

Substituting the two preceding inequalities into the decomposition of the v-utility difference yields

$$[v(y) - p_y] - [v(x) - p_x] \le \delta + d \qquad \text{for every } y \in S.$$

Taking the maximum over $y \in S$ gives

$$\begin{aligned} \mathrm{Reg}_v(x; M) &= \max_{y \in S} \big([v(y) - p_y] - [v(x) - p_x] \big) \\ &\le \delta + d \\ &= \delta + d_{\mathrm{str}}(v, w). \end{aligned}$$

This proves the stated transfer inequality.

We now prove that this upper bound is sharp. Let

$$a := v - w.$$

Because $\mathcal{X}$ is finite, there exist outcomes $x, y \in \mathcal{X}$ satisfying

$$a_x = \min_{z \in \mathcal{X}} a_z, \qquad a_y = \max_{z \in \mathcal{X}} a_z.$$

For these outcomes,

$$\begin{aligned} a_y - a_x &= \max_{z \in \mathcal{X}} a_z - \min_{z \in \mathcal{X}} a_z \\ &= \mathrm{osc}(a) \\ &= d_{\mathrm{str}}(v, w) \\ &= d. \end{aligned}$$

If $d > 0$, the maximizing and minimizing outcomes are necessarily distinct. If $d = 0$, all coordinates of a coincide; because $m \ge 2$, we may choose any two distinct outcomes x and y, and still have $a_y - a_x = d = 0$.

Consider the two-outcome menu

$$M^* = (\{x, y\}, p)$$

with prices

$$p_x := 0, \qquad p_y := w(y) - w(x) - \delta.$$

Under the surrogate valuation w, the utility of y relative to the utility of x is

$$\begin{aligned} [w(y) - p_y] - [w(x) - p_x] &= w(y) - \big(w(y) - w(x) - \delta \big) - w(x) \\ &= w(y) - w(y) + w(x) + \delta - w(x) \\ &= \delta. \end{aligned}$$

The menu contains only x and y. Relative to x, the utility difference for choosing x itself is zero, while the utility difference for choosing y is δ. Since $\delta \ge 0$,

$$\begin{aligned} \mathrm{Reg}_w(x; M^*) &= \max \{0,\, [w(y) - p_y] - [w(x) - p_x]\} \\ &= \max\{0, \delta\} \\ &= \delta. \end{aligned}$$

Thus x satisfies the feasibility condition

$$\mathrm{Reg}_w(x; M^*) \le \delta$$

with equality.

Under the true valuation v, the utility advantage of y over x is

$$\begin{aligned}[v(y) - p_y] - [v(x) - p_x] &= v(y) - v(x) - \big(w(y) - w(x) - \delta\big) \\ &= \delta + \big(v(y) - v(x) - w(y) + w(x)\big) \\ &= \delta + \big([v(y) - w(y)] - [v(x) - w(x)]\big) \\ &= \delta + (a_y - a_x) \\ &= \delta + d.\end{aligned}$$

Since $\delta + d \ge 0$, it follows that

$$\begin{aligned}\mathrm{Reg}_v(x; M^*) &= \max\{0,\ [v(y) - p_y] - [v(x) - p_x]\} \\ &= \max\{0, \delta + d\} \\ &= \delta + d.\end{aligned}$$

Therefore this particular menu and action give

$$\sup_{M=(S,p)\in\mathcal{M}} \sup_{\substack{x\in S \\ \mathrm{Reg}_w(x;M)\le\delta}} \mathrm{Reg}_v(x; M) \ge \delta + d.$$

On the other hand, the transfer inequality already proved applies to every menu and every action satisfying $\mathrm{Reg}_w(x; M) \le \delta$. Hence

$$\sup_{M=(S,p)\in\mathcal{M}} \sup_{\substack{x\in S \\ \mathrm{Reg}_w(x;M)\le\delta}} \mathrm{Reg}_v(x; M) \le \delta + d.$$

Combining the lower and upper bounds yields

$$\sup_{M=(S,p)\in\mathcal{M}} \sup_{\substack{x\in S \\ \mathrm{Reg}_w(x;M)\le\delta}} \mathrm{Reg}_v(x; M) = \delta + d = \delta + d_{\mathrm{str}}(v, w),$$

which establishes the exact general-δ supremum.

For $\delta = 0$, the same construction uses

$$p_y = w(y) - w(x).$$

Thus

$$[w(y) - p_y] - [w(x) - p_x] = 0,$$

so x and y tie under w, and in particular

$$x \in \mathrm{BR}_w(M^*).$$

The calculation above becomes

$$\mathrm{Reg}_v(x; M^*) = d.$$

Therefore

$$\sup_{M\in\mathcal{M}} \sup_{x\in\mathrm{BR}_w(M)} \mathrm{Reg}_v(x; M) \ge d.$$

The transfer inequality with $\delta = 0$ gives, for every $x \in \mathrm{BR}_w(M)$,

$$\mathrm{Reg}_v(x; M) \le d.$$

Hence

$$\sup_{M\in\mathcal{M}} \sup_{x\in\mathrm{BR}_w(M)} \mathrm{Reg}_v(x; M) = d = d_{\mathrm{str}}(v, w).$$

For completeness, the same supremum remains valid if one further requires x to be the unique best response under w. When $d > 0$, fix any $\varepsilon \in (0, d)$ and use the same two outcomes x, y, but set

$$p_x := 0, \qquad p_y := w(y) - w(x) + \varepsilon.$$

Then

$$\begin{aligned}[w(y) - p_y] - [w(x) - p_x] &= w(y) - \big(w(y) - w(x) + \varepsilon\big) - w(x) \\ &= -\varepsilon \\ &< 0.\end{aligned}$$

Thus x is the unique best response under w. Under v,

$$\begin{aligned}[v(y) - p_y] - [v(x) - p_x] &= v(y) - v(x) - \big(w(y) - w(x) + \varepsilon\big) \\ &= (a_y - a_x) - \varepsilon \\ &= d - \varepsilon \\ &> 0.\end{aligned}$$

Consequently,

$$\mathrm{Reg}_v(x; M) = d - \varepsilon.$$

Because $\varepsilon > 0$ can be arbitrarily small,

$$\sup_{M \in \mathcal{M},\, x \text{ unique in } \mathrm{BR}_w(M)} \mathrm{Reg}_v(x; M) \geq d.$$

The transfer inequality at $\delta = 0$ gives the reverse bound, so the restricted supremum is also d. If $d = 0$, the same upper bound already forces every such regret to be zero. Thus requiring a unique surrogate best response does not change the supremum, although when $d > 0$ the value d is approached through a vanishing positive response margin rather than through the exact tie used above. □

B.2 Proofs for Section 3

B.2.1 Proof of Theorem 3.1

Proof. For brevity, write $A_{xy} := A_{xy}(C)$. Since C is bounded and $\mathcal{X}$ is finite, every A_{xy} is finite. Fix $w \in \mathbb{R}^m$. For any $v \in C$, the set of ordered pairs (x, y) contains the reverse of every pair, and hence

$$\begin{aligned}d_{\mathrm{str}}(v, w) &= \max_{x,y \in \mathcal{X}} |[v(y) - v(x)] - [w(y) - w(x)]| \\ &= \max_{x,y \in \mathcal{X}} \left([v(y) - v(x)] - [w(y) - w(x)]\right).\end{aligned}$$

Indeed, if the expression associated with (x, y) is negative, then the expression associated with the reversed pair (y, x) is its positive opposite. Because $\mathcal{X}^2$ is finite, the supremum over v commutes with this finite maximum. Therefore

$$\begin{aligned}\sup_{v \in C} d_{\mathrm{str}}(v, w) &= \max_{x,y \in \mathcal{X}} \sup_{v \in C} \left([v(y) - v(x)] - [w(y) - w(x)]\right) \\ &= \max_{x,y \in \mathcal{X}} \left(A_{xy} - (w_y - w_x)\right).\end{aligned}$$

It follows that, for any real number r,

$$\sup_{v \in C} d_{\mathrm{str}}(v, w) \leq r$$

is equivalent to

$$A_{xy} - (w_y - w_x) \leq r \qquad \text{for every } x, y \in \mathcal{X},$$

or, equivalently,

$$w_y - w_x \geq A_{xy} - r \qquad \text{for every } x, y \in \mathcal{X}.$$

Thus the optimization defining $\rho(C)$ is exactly the difference-constraint program in equation 2, apart from the harmless anchoring condition $w_0 = 0$. The anchoring condition is harmless because adding the same constant to every coordinate of w changes neither $d_{\mathrm{str}}(v, w)$ nor any difference $w_y - w_x$.

We next characterize the values of r for which these difference constraints are feasible. Define

$$B_{xy}^{(r)} := A_{xy} - r.$$

Suppose first that a vector w satisfies

$$w_y - w_x \geq B_{xy}^{(r)} \qquad \forall x, y \in \mathcal{X}.$$

Let

$$\gamma = (x_1, x_2, \dots, x_k, x_1)$$

be any directed cycle, where $x_{k+1} := x_1$. Summing the k inequalities corresponding to the edges of the cycle gives

$$\begin{aligned} 0 &= \sum_{j=1}^{k} (w_{x_{j+1}} - w_{x_j}) \\ &\geq \sum_{j=1}^{k} B^{(r)}_{x_j x_{j+1}} \\ &= \sum_{j=1}^{k} A_{x_j x_{j+1}} - kr. \end{aligned}$$

Hence every feasible r must satisfy

$$r \geq \frac{1}{k} \sum_{j=1}^{k} A_{x_j x_{j+1}}$$

for every directed cycle. In particular, it must satisfy this inequality for every simple directed cycle in Γ.

Conversely, suppose that $\sum_{(x,y)\in\gamma} B^{(r)}_{xy} \leq 0$ for every simple directed cycle γ. Every closed directed walk can be decomposed into simple directed cycles. Consequently, every closed directed walk also has nonpositive total $B^{(r)}$-weight. For each $y \in \mathcal{X}$, consider all directed walks from 0 to y and define w_y as the largest total $B^{(r)}$-weight of such a walk. This quantity is finite and the maximum is attained. Indeed, whenever a walk repeats a vertex, the portion between two successive occurrences of that vertex is a closed walk of nonpositive weight. Deleting that portion does not decrease the total weight. Repeating this deletion reduces any walk to a simple path with weight at least as large, and there are only finitely many simple paths in the finite graph on $\mathcal{X}$.

For $y = 0$, the empty walk has weight zero, while every nonempty closed walk has weight at most zero. Hence $w_0 = 0$. Now fix $x, y \in \mathcal{X}$. Take a maximum-weight path from 0 to x and append the edge $x \to y$. The resulting walk has weight $w_x + B^{(r)}_{xy}$. If this walk repeats a vertex, deleting its nonpositive closed subwalks yields a simple path from 0 to y with weight at least $w_x + B^{(r)}_{xy}$. By the definition of w_y as the maximum path weight, $w_y \geq w_x + B^{(r)}_{xy}$. Therefore

$$w_y - w_x \geq B^{(r)}_{xy} = A_{xy} - r.$$

Thus the complete system of difference constraints is feasible whenever every simple directed cycle has nonpositive total $B^{(r)}$-weight.

Define

$$r^\star := \max_{\gamma\in\Gamma} \frac{1}{|\gamma|} \sum_{(x,y)\in\gamma} A_{xy}.$$

The maximum exists because Γ is finite. Moreover,

$$A_{xx} = \sup_{v\in C} \{v(x) - v(x)\} = 0,$$

so every self-loop has mean zero and therefore $r^\star \geq 0$. If $r < r^\star$, let $\gamma^\star$ be a cycle attaining the maximum. Then

$$\begin{aligned} \sum_{(x,y)\in\gamma^\star} \big(A_{xy} - r\big) &= \sum_{(x,y)\in\gamma^\star} A_{xy} - |\gamma^\star| r \\ &= |\gamma^\star|(r^\star - r) \\ &> 0. \end{aligned}$$

By the necessary cycle condition proved above, the difference constraints are therefore infeasible.

On the other hand, if $r = r^\star$, then for every simple cycle γ,

$$\begin{aligned}\sum_{(x,y)\in\gamma}\big(A_{xy} - r^\star\big) &= \sum_{(x,y)\in\gamma} A_{xy} - |\gamma| r^\star \\ &= |\gamma|\left(\frac{1}{|\gamma|}\sum_{(x,y)\in\gamma} A_{xy} - r^\star\right) \\ &\le 0.\end{aligned}$$

The preceding construction therefore produces an anchored feasible potential $w^\star$ satisfying

$$w^\star_y - w^\star_x \ge A_{xy} - r^\star \qquad \forall x, y \in \mathcal{X}.$$

By the equivalence between the difference constraints and strategic distortion,

$$\sup_{v\in C} d_{\mathrm{str}}(v, w^\star) \le r^\star.$$

No $r < r^\star$ is feasible, so no surrogate can have worst-case strategic distortion strictly below $r^\star$. Consequently,

$$\rho(C) = r^\star = \max_{\gamma\in\Gamma} \frac{1}{|\gamma|}\sum_{(x,y)\in\gamma} A_{xy}(C).$$

In particular, the infimum defining $\rho(C)$ is attained by $w^\star$. The same equivalence also shows that equation 2 has optimal value $r^\star$, and that the w-component of every optimizer is an optimal anchored surrogate.

It remains to identify this geometric radius with the minimax value over all deterministic menu-response rules. Fix an optimal surrogate $w^\star$ and a deterministic tie-breaking rule that selects one element of $\mathrm{BR}_{w^\star}(M)$ for every menu M. Denote the resulting deterministic response rule by $a^\star$. For every $v \in C$ and every menu M, Theorem 2.3 with $\delta = 0$ gives

$$\mathrm{Reg}_v(a^\star(M); M) \le d_{\mathrm{str}}(v, w^\star).$$

Hence

$$\begin{aligned}\sup_{v\in C}\sup_{M\in\mathcal{M}} \mathrm{Reg}_v(a^\star(M); M) &\le \sup_{v\in C} d_{\mathrm{str}}(v, w^\star) \\ &= \rho(C).\end{aligned}$$

Since $\mathfrak{R}(C)$ takes the infimum over all deterministic response rules,

$$\mathfrak{R}(C) \le \rho(C).$$

For the reverse inequality, let

$$\gamma^\star = (x_1, x_2, \ldots, x_k, x_1)$$

be any cycle attaining the maximum cycle mean, and again write $x_{k+1} := x_1$. Since

$$\frac{1}{k}\sum_{j=1}^{k} A_{x_j x_{j+1}} = r^\star,$$

we have

$$\sum_{j=1}^{k}\big(A_{x_j x_{j+1}} - r^\star\big) = 0.$$

Hence there exist prices on the vertices of this cycle satisfying

$$p_{x_{j+1}} - p_{x_j} = A_{x_j x_{j+1}} - r^\star, \qquad j = 1, \ldots, k,$$

with indices interpreted cyclically. Explicitly, one may set

$$p_{x_1} = 0$$

and then recursively define

$$p_{x_{j+1}} = p_{x_j} + A_{x_j x_{j+1}} - r^\star, \qquad j = 1, \ldots, k-1.$$

For the remaining edge $x_k \to x_1$,

$$\begin{aligned} p_{x_1} - p_{x_k} &= -\sum_{j=1}^{k-1} \big(A_{x_j x_{j+1}} - r^\star\big) \\ &= A_{x_k x_1} - r^\star, \end{aligned}$$

where the last equality follows from the zero-sum identity around $\gamma^\star$. Thus these prices are consistent around the entire cycle.

Let

$$M_{\gamma^\star} = \big(\{x_1, \dots, x_k\}, p\big).$$

Consider an arbitrary deterministic response rule $a \in \mathcal{A}$. Because

$$a(M_{\gamma^\star}) \in \{x_1, \dots, x_k\},$$

there exists an index j such that $a(M_{\gamma^\star}) = x_j$. For every $v \in C$, the regret from selecting x_j is at least the utility advantage of the particular feasible alternative x_{j+1} over x_j:

$$\begin{aligned} \mathrm{Reg}_v(x_j; M_{\gamma^\star}) &= \max_{y \in \{x_1, \dots, x_k\}} \Big([v(y) - p_y] - [v(x_j) - p_{x_j}] \Big) \\ &\geq [v(x_{j+1}) - p_{x_{j+1}}] - [v(x_j) - p_{x_j}] \\ &= [v(x_{j+1}) - v(x_j)] - [p_{x_{j+1}} - p_{x_j}] \\ &= [v(x_{j+1}) - v(x_j)] - A_{x_j x_{j+1}} + r^\star. \end{aligned}$$

Taking the supremum over $v \in C$ gives

$$\begin{aligned} \sup_{v \in C} \mathrm{Reg}_v(x_j; M_{\gamma^\star}) &\geq \sup_{v \in C} \big([v(x_{j+1}) - v(x_j)] - A_{x_j x_{j+1}} + r^\star \big) \\ &= A_{x_j x_{j+1}} - A_{x_j x_{j+1}} + r^\star \\ &= r^\star. \end{aligned}$$

This calculation is important because it uses only the defining supremum

$$A_{x_j x_{j+1}} = \sup_{v \in C} \{v(x_{j+1}) - v(x_j)\}.$$

It does not require the existence of a valuation in C that actually attains this supremum. In particular, the conclusion remains valid when C is bounded but not closed.

Since

$$a(M_{\gamma^\star}) = x_j,$$

we have proved

$$\sup_{v \in C} \mathrm{Reg}_v(a(M_{\gamma^\star}); M_{\gamma^\star}) \geq r^\star = \rho(C) \qquad \forall a \in \mathcal{A}.$$

The menu $M_{\gamma^\star}$ is fixed before the response rule's selected outcome is known; only the valuation witnessing the supremum may depend on that selected outcome. It follows that every $a \in \mathcal{A}$ satisfies

$$\sup_{v \in C} \sup_{M \in \mathcal{M}} \mathrm{Reg}_v(a(M); M) \geq \rho(C).$$

Taking the infimum over a yields

$$\mathfrak{R}(C) \geq \rho(C).$$

Combining this inequality with the previously established upper bound gives

$$\mathfrak{R}(C) = \rho(C) = \max_{\gamma \in \Gamma} \frac{1}{|\gamma|} \sum_{(x,y) \in \gamma} A_{xy}(C).$$

The response rule obtained by deterministic tie-breaking from $\mathrm{BR}_{w^\star}(M)$ therefore attains the minimax value, while the menu $M_{\gamma^\star}$ provides the claimed single-menu converse for every deterministic response rule. □

B.2.2 PROOF OF COROLLARY 3.2

Proof. We first express the strategic diameter in terms of the quantities $A_{xy}(C)$. For any $u, v \in C$, the presence of both ordered pairs (x, y) and (y, x) implies

$$d_{\mathrm{str}}(u, v) = \max_{x,y \in \mathcal{X}} \left([u(y) - u(x)] - [v(y) - v(x)] \right).$$

Since $\mathcal{X}^2$ is finite,

$$\begin{aligned}
\Delta(C) &= \sup_{u,v \in C} d_{\mathrm{str}}(u, v) \\
&= \sup_{u,v \in C} \max_{x,y \in \mathcal{X}} \left([u(y) - u(x)] - [v(y) - v(x)] \right) \\
&= \max_{x,y \in \mathcal{X}} \sup_{u,v \in C} \left([u(y) - u(x)] + [v(x) - v(y)] \right) \\
&= \max_{x,y \in \mathcal{X}} \left(\sup_{u \in C} [u(y) - u(x)] + \sup_{v \in C} [v(x) - v(y)] \right) \\
&= \max_{x,y \in \mathcal{X}} \left(A_{xy}(C) + A_{yx}(C) \right).
\end{aligned}$$

For every pair of distinct outcomes x, y, the directed two-cycle

$$x \to y \to x$$

has mean

$$\frac{1}{2} \left(A_{xy}(C) + A_{yx}(C) \right).$$

Therefore the largest two-cycle mean is exactly

$$\frac{1}{2} \Delta(C).$$

Since Theorem 3.1 identifies $\rho(C)$ with the maximum mean over all simple directed cycles,

$$\rho(C) \geq \frac{1}{2} \Delta(C).$$

We now prove the upper bound. Let

$$\gamma = (x_1, x_2, \ldots, x_k, x_1)$$

be a simple directed cycle of length $k \geq 2$, and put $x_{k+1} := x_1$. Define

$$d_i(v) := v(x_{i+1}) - v(x_i), \qquad A_i := A_{x_i x_{i+1}}(C).$$

For every $v \in C$, the valuation differences telescope around the cycle:

$$\begin{aligned}
\sum_{i=1}^{k} d_i(v) &= \sum_{i=1}^{k} \big(v(x_{i+1}) - v(x_i) \big) \\
&= v(x_2) - v(x_1) + v(x_3) - v(x_2) + \cdots \\
&\quad + v(x_1) - v(x_k) \\
&= 0.
\end{aligned}$$

Fix one edge index

$$j \in \{1, \ldots, k\}.$$

By the definition of A_j as a supremum, there exists a sequence $v^{(n)} \in C$ such that

$$d_j(v^{(n)}) \longrightarrow A_j.$$

No compactness or attainment assumption on C is needed for this sequence.

For every edge i and every n,

$$\begin{aligned}
A_i - d_i(v^{(n)}) &= \sup_{u \in C} d_i(u) - d_i(v^{(n)}) \\
&= \sup_{u \in C} \big(d_i(u) - d_i(v^{(n)}) \big).
\end{aligned}$$

For a fixed $u \in C$,

$$d_i(u) - d_i(v^{(n)}) = [u(x_{i+1}) - u(x_i)] - [v^{(n)}(x_{i+1}) - v^{(n)}(x_i)].$$

The strategic distortion $d_{\mathrm{str}}(u, v^{(n)})$ is the maximum absolute error over all pairwise valuation differences. Hence

$$d_i(u) - d_i(v^{(n)}) \le d_{\mathrm{str}}(u, v^{(n)}).$$

Taking the supremum over $u \in C$ yields

$$\begin{aligned} A_i - d_i(v^{(n)}) &\le \sup_{u \in C} d_{\mathrm{str}}(u, v^{(n)}) \\ &\le \Delta(C). \end{aligned}$$

Using

$$\sum_{i=1}^{k} d_i(v^{(n)}) = 0,$$

we obtain

$$\begin{aligned} \sum_{i=1}^{k} A_i &= \sum_{i=1}^{k} \Big(A_i - d_i(v^{(n)}) \Big) + \sum_{i=1}^{k} d_i(v^{(n)}) \\ &= \sum_{i=1}^{k} \Big(A_i - d_i(v^{(n)}) \Big) \\ &= \Big(A_j - d_j(v^{(n)}) \Big) + \sum_{i \neq j} \Big(A_i - d_i(v^{(n)}) \Big) \\ &\le \Big(A_j - d_j(v^{(n)}) \Big) + (k-1)\Delta(C). \end{aligned}$$

Letting $n \to \infty$ gives

$$A_j - d_j(v^{(n)}) \longrightarrow 0,$$

and therefore

$$\sum_{i=1}^{k} A_i \le (k-1)\Delta(C).$$

Consequently the mean of γ satisfies

$$\frac{1}{k} \sum_{i=1}^{k} A_i \le \frac{k-1}{k} \Delta(C).$$

A self-loop has mean

$$A_{xx}(C) = 0,$$

so the same conclusion is trivial when $k = 1$.

Every simple directed cycle on $\mathcal{X}$ has length at most m. Since

$$\frac{k-1}{k} = 1 - \frac{1}{k}$$

is increasing in k,

$$\frac{k-1}{k} \le \frac{m-1}{m}.$$

Thus every simple directed cycle satisfies

$$\frac{1}{|\gamma|} \sum_{(x,y) \in \gamma} A_{xy}(C) \le \frac{m-1}{m} \Delta(C).$$

Taking the maximum over $\gamma \in \Gamma$ and applying Theorem 3.1 gives

$$\rho(C) \le \frac{m-1}{m} \Delta(C).$$

It remains to show that neither constant can be improved. For the lower constant, let

$$z = (0, 1, 0, \dots, 0)$$

and consider the two-point anchored cell

$$C_{\mathrm{mid}} := \{0, z\}.$$

Its diameter is

$$\begin{aligned}\Delta(C_{\mathrm{mid}}) &= d_{\mathrm{str}}(0, z) \\ &= \mathrm{osc}(z) \\ &= 1.\end{aligned}$$

Consider the ambient midpoint

$$w := \frac{1}{2}z.$$

By homogeneity of the oscillation,

$$\begin{aligned}d_{\mathrm{str}}(0, w) &= \mathrm{osc}\left(\frac{1}{2}z\right) = \frac{1}{2}, \\ d_{\mathrm{str}}(z, w) &= \mathrm{osc}\left(z - \frac{1}{2}z\right) = \mathrm{osc}\left(\frac{1}{2}z\right) = \frac{1}{2}.\end{aligned}$$

Therefore

$$\rho(C_{\mathrm{mid}}) \le \frac{1}{2}.$$

The lower bound already established gives

$$\rho(C_{\mathrm{mid}}) \ge \frac{1}{2}\Delta(C_{\mathrm{mid}}) = \frac{1}{2}.$$

Hence

$$\rho(C_{\mathrm{mid}}) = \frac{1}{2} = \frac{1}{2}\Delta(C_{\mathrm{mid}}),$$

which proves that the lower constant $1/2$ is sharp.

For the upper constant, define

$$b := \frac{1}{m}(0, m-1, m-2, \dots, 1) \in \mathbb{R}^m.$$

Let

$$b^{(0)}, b^{(1)}, \dots, b^{(m-1)}$$

be the m cyclic rotations of b, and anchor each rotation by setting

$$\widetilde{b}^{(s)} := b^{(s)} - b_0^{(s)}\mathbf{1}.$$

Consider the cell

$$C_{\mathrm{cyc}} := \{\widetilde{b}^{(s)} : s = 0, \dots, m-1\}.$$

Anchoring does not change any pairwise valuation difference, and hence does not change d_{str}. It is therefore enough to analyze the unanchored rotations $b^{(s)}$.

Consider two distinct rotations whose cyclic offset is

$$q \in \{1, \dots, m-1\}.$$

The unscaled coordinates of b are

$$0, m-1, m-2, \dots, 1,$$

which are exactly the residue classes modulo m, written in decreasing cyclic order. Moving by an offset q adds q modulo m. Consequently, at each coordinate the unscaled difference between the two rotations is either

$$q$$

when no modular wrap occurs, or

$$q - m$$

when a modular wrap occurs. Since q is strictly between 0 and m, both cases occur among the m coordinates. Dividing by m, the coordinatewise difference between the two rotations therefore takes exactly the two values

$$\frac{q}{m} \qquad \text{and} \qquad \frac{q}{m} - 1.$$

Its oscillation is

$$\frac{q}{m} - \left(\frac{q}{m} - 1\right) = 1.$$

Thus every two distinct elements of C_{cyc} have strategic distance 1, and hence

$$\Delta(C_{\mathrm{cyc}}) = 1.$$

Now consider the directed cycle

$$0 \to 1 \to \cdots \to m - 1 \to 0.$$

In the base vector b, the adjacent difference along the edge $0 \to 1$ is

$$b_1 - b_0 = \frac{m-1}{m},$$

whereas every other adjacent difference around the same cyclic ordering is

$$-\frac{1}{m}.$$

Cyclically rotating b moves the unique positive jump $(m-1)/m$ to an arbitrary edge of the directed cycle. Therefore, for every x with the successor $x + 1$ interpreted modulo m,

$$A_{x,x+1}(C_{\mathrm{cyc}}) = \frac{m-1}{m}.$$

The mean of the cycle

$$0 \to 1 \to \cdots \to m - 1 \to 0$$

is consequently

$$\begin{aligned}\frac{1}{m}\sum_{x=0}^{m-1} A_{x,x+1}(C_{\mathrm{cyc}}) &= \frac{1}{m}\sum_{x=0}^{m-1}\frac{m-1}{m} \\ &= \frac{m-1}{m}.\end{aligned}$$

Theorem 3.1 therefore implies

$$\rho(C_{\mathrm{cyc}}) \geq \frac{m-1}{m}.$$

On the other hand, the upper bound already proved and $\Delta(C_{\mathrm{cyc}}) = 1$ imply

$$\rho(C_{\mathrm{cyc}}) \leq \frac{m-1}{m}.$$

Hence

$$\rho(C_{\mathrm{cyc}}) = \frac{m-1}{m} = \frac{m-1}{m}\Delta(C_{\mathrm{cyc}}).$$

Thus the upper constant $(m-1)/m$ is also sharp. □

B.2.3 A THREE-OUTCOME CYCLE OBSTRUCTION

Consider

$$C = \{(0, 2, 1),\ (0, -1, 1),\ (0, -1, -2)\}\,.$$

The differences between the first and second valuations, the second and third valuations, and the first and third valuations are, respectively,

$$(0, 3, 0), \qquad (0, 0, 3), \qquad (0, 3, 3),$$

up to an overall sign. Each has oscillation 3. Therefore

$$\Delta(C) = 3.$$

Moreover,

$$\begin{aligned} A_{01}(C) &= \sup_{v \in C}\{v(1) - v(0)\} \\ &= \max\{2, -1, -1\} \\ &= 2, \\ A_{12}(C) &= \sup_{v \in C}\{v(2) - v(1)\} \\ &= \max\{-1, 2, -1\} \\ &= 2, \end{aligned}$$

and

$$\begin{aligned} A_{20}(C) &= \sup_{v \in C}\{v(0) - v(2)\} \\ &= \max\{-1, -1, 2\} \\ &= 2. \end{aligned}$$

Thus the directed cycle

$$0 \to 1 \to 2 \to 0$$

has mean

$$\frac{1}{3}\left(A_{01}(C) + A_{12}(C) + A_{20}(C)\right) = 2.$$

Theorem 3.1 therefore gives

$$\rho(C) \geq 2.$$

At the same time, Corollary 3.2, with $m = 3$, gives

$$\rho(C) \leq \frac{2}{3}\Delta(C) = 2.$$

Consequently,

$$\rho(C) = 2 > \frac{\Delta(C)}{2} = \frac{3}{2}.$$

The same obstruction can be seen directly through a single priced menu. Take the full menu with zero prices,

$$S = \{0, 1, 2\}, \qquad p_0 = p_1 = p_2 = 0.$$

If a deterministic rule selects outcome 0, the valuation $(0, 2, 1)$ gives

$$[v(1) - p_1] - [v(0) - p_0] = 2,$$

so its regret is at least 2. If the rule selects outcome 1, the valuation $(0, -1, 1)$ gives

$$[v(2) - p_2] - [v(1) - p_1] = 2.$$

If the rule selects outcome 2, the valuation $(0, -1, -2)$ gives

$$[v(0) - p_0] - [v(2) - p_2] = 2.$$

Thus whichever outcome is selected, some valuation in the cell prefers the next outcome on the cycle by 2. This shows concretely why the two-cycle quantity $\Delta(C)/2$ can fail to capture a genuinely multi-outcome obstruction.

B.2.4 WHY DETERMINISTIC RESPONSE RULES MATTER

Consider the two-outcome cell

$$C = \{v^-, v^+\}, \qquad v^- = (0, -1), \qquad v^+ = (0, 1).$$

Singleton menus have zero regret, so it suffices to consider the full menu $S = \{0, 1\}$. Since adding the same constant to both prices has no effect on choice or regret, write

$$t := p_1 - p_0$$

for the relative price. Under v^-, the utility advantage of outcome 1 over outcome 0 is

$$\begin{aligned}[v^-(1) - p_1] - [v^-(0) - p_0] &= -1 - (p_1 - p_0) \\ &= -1 - t,\end{aligned}$$

whereas under v^+ it is

$$\begin{aligned}[v^+(1) - p_1] - [v^+(0) - p_0] &= 1 - (p_1 - p_0) \\ &= 1 - t.\end{aligned}$$

We first verify the deterministic value directly. At $t = 0$, any deterministic rule must select either outcome 0 or outcome 1. If it selects outcome 0, then v^+ prefers outcome 1 by 1, so the regret is 1. If it selects outcome 1, then v^- prefers outcome 0 by 1, so again the regret is 1. Therefore every deterministic response rule has worst-case regret at least 1.

Conversely, consider the deterministic rule that selects outcome 1 when $t < 0$ and outcome 0 when $t \geq 0$. If $t \leq -1$, then $-1 - t \geq 0 \qquad$ and $\qquad 1 - t > 0$, so both valuations weakly prefer outcome 1, and the selected outcome has zero regret. If $-1 < t < 0$, then v^+ prefers outcome 1, while v^- prefers outcome 0. The only positive regret is therefore the regret of v^- from selecting outcome 1, which equals

$$\begin{aligned}[v^-(0) - p_0] - [v^-(1) - p_1] &= 1 + t \\ &< 1.\end{aligned}$$

If $0 \leq t < 1$, the rule selects outcome 0. Valuation v^- prefers outcome 0, while v^+ prefers outcome 1. The only positive regret is

$$\begin{aligned}[v^+(1) - p_1] - [v^+(0) - p_0] &= 1 - t \\ &\leq 1.\end{aligned}$$

Finally, if $t \geq 1$, both valuations weakly prefer outcome 0, so the regret is zero. Consequently this deterministic rule has worst-case regret at most 1. Together with the lower bound at $t = 0$, this proves that the deterministic minimax value is exactly 1.

Now allow the menu-response rule to randomize, and measure regret in expectation over the rule's randomization. For $t \in [-1, 1]$, valuation v^- weakly prefers outcome 0, while v^+ weakly prefers outcome 1. Suppose the rule selects outcome 1 with probability $q \in [0, 1]$ and outcome 0 with probability $1 - q$.

For v^-, choosing outcome 0 creates no regret, while choosing outcome 1 creates regret

$$[v^-(0) - p_0] - [v^-(1) - p_1] = 1 + t.$$

Its expected regret is therefore

$$R_-(q, t) = q(1 + t).$$

For v^+, choosing outcome 1 creates no regret, while choosing outcome 0 creates regret

$$[v^+(1) - p_1] - [v^+(0) - p_0] = 1 - t.$$

Its expected regret is therefore

$$R_+(q, t) = (1 - q)(1 - t).$$

For a fixed $t \in [-1, 1]$, choose

$$q = q(t) := \frac{1 - t}{2}.$$

This is a valid probability because

$$0 \le \frac{1-t}{2} \le 1 \qquad \text{when } t \in [-1,1].$$

The two expected regrets are then equal:

$$\begin{aligned} R_-(q(t),t) &= \frac{1-t}{2}(1+t) \\ &= \frac{1-t^2}{2}, \end{aligned}$$

while

$$\begin{aligned} R_+(q(t),t) &= \left(1-\frac{1-t}{2}\right)(1-t) \\ &= \frac{1+t}{2}(1-t) \\ &= \frac{1-t^2}{2}. \end{aligned}$$

Thus, for every $t \in [-1,1]$, the worst-case expected regret is

$$\max\{R_-(q(t),t), R_+(q(t),t)\} = \frac{1-t^2}{2} \le \frac{1}{2}.$$

If $t < -1$, then $-1-t > 0$ and $1-t > 0$, so both valuations strictly prefer outcome 1. Selecting outcome 1 with probability one therefore gives zero expected regret. If $t > 1$, then $-1-t < 0$ and $1-t < 0$, so both valuations strictly prefer outcome 0, and selecting outcome 0 with probability one again gives zero expected regret. At the boundary values $t = -1$ and $t = 1$, one valuation is indifferent and the same rule also has zero regret. Hence the randomized response rule just constructed has worst-case expected regret at most $\frac{1}{2}$ over all priced menus.

It remains to show that no randomized rule can obtain a smaller value. Take again the menu with $t = 0$. Suppose an arbitrary randomized response rule selects outcome 1 with probability q. For valuation v^-, outcome 0 is optimal and the regret from outcome 1 is exactly 1, so its expected regret is q. For valuation v^+, outcome 1 is optimal and the regret from outcome 0 is exactly 1, so its expected regret is $1-q$. Therefore the worst-case expected regret at this single menu is

$$\max\{q, 1-q\}.$$

For every $q \in [0,1]$,

$$\max\{q, 1-q\} \ge \frac{q+(1-q)}{2} = \frac{1}{2}.$$

Thus every randomized menu-response rule has worst-case expected regret at least $1/2$. Combining this lower bound with the preceding construction shows that the randomized minimax value is exactly $\frac{1}{2}$. Hence randomization strictly reduces the minimax value from 1 to $1/2$ in this example, which is why Theorem 3.1 is stated specifically for deterministic response rules.

B.3 Proofs for Section 4

B.3.1 Proof of Theorem 4.1

Proof. Fix an agent i, a true type $\theta_i \in \Theta_i$, an arbitrary profile $z_{-i} \in Z_{-i}$ of the other agents' labels, and an arbitrary alternative label $z_i' \in Z_i$. Let

$$z_i := E_i(\theta_i), \qquad x^0 := x(z_i, z_{-i}), \qquad x^1 := x(z_i', z_{-i}),$$

and abbreviate

$$p^0 := p_i(z_i, z_{-i}), \qquad p^1 := p_i(z_i', z_{-i}), \qquad v_i^\theta(x) := v_i(\theta_i, x), \qquad w_i^z(x) := w_{i,z_i}(x).$$

With this notation, the true utility gain from deviating from z_i to z_i' is

$$\begin{aligned} &u_i(\theta_i; x^1, p^1) - u_i(\theta_i; x^0, p^0) \\ &\quad = \left[v_i^\theta(x^1) - p^1\right] - \left[v_i^\theta(x^0) - p^0\right] \\ &\quad = \left[w_i^z(x^1) - p^1\right] - \left[w_i^z(x^0) - p^0\right] \\ &\qquad + \Big(\left[v_i^\theta(x^1) - v_i^\theta(x^0)\right] - \left[w_i^z(x^1) - w_i^z(x^0)\right]\Big). \end{aligned}$$

The first two terms on the right-hand side are precisely the surrogate utility gain obtained by changing the report from z_i to z_i' while holding z_{-i} fixed. Since the latent mechanism is δ_i-DSIC on the surrogate label space, equation 3, applied to the pair (z_i, z_i'), gives

$$\left[w_i^z(x^1) - p^1\right] - \left[w_i^z(x^0) - p^0\right] \le \delta_i.$$

By the definition of the pointwise certificate and of strategic distortion,

$$\begin{aligned} D_i(\theta_i) &= d_{\mathrm{str}}\big(v_i(\theta_i, \cdot), w_{i,E_i(\theta_i)}\big) \\ &= \max_{a,b \in \mathcal{X}} \left|\left[v_i(\theta_i, b) - v_i(\theta_i, a)\right] - \left[w_{i,z_i}(b) - w_{i,z_i}(a)\right]\right|. \end{aligned}$$

Therefore the particular ordered pair of outcomes (x^0, x^1) satisfies

$$\left|\left[v_i^\theta(x^1) - v_i^\theta(x^0)\right] - \left[w_i^z(x^1) - w_i^z(x^0)\right]\right| \le D_i(\theta_i),$$

and hence, in the direction needed for an upper bound,

$$\left[v_i^\theta(x^1) - v_i^\theta(x^0)\right] - \left[w_i^z(x^1) - w_i^z(x^0)\right] \le D_i(\theta_i).$$

Substituting the preceding two inequalities into the exact decomposition of the true deviation gain yields

$$\begin{aligned} &u_i(\theta_i; x(z_i', z_{-i}), p_i(z_i', z_{-i})) \\ &\qquad - u_i(\theta_i; x(E_i(\theta_i), z_{-i}), p_i(E_i(\theta_i), z_{-i})) \\ &\qquad \le \delta_i + D_i(\theta_i), \end{aligned}$$

which is exactly equation 4. Only one distortion term is needed: the deviation compares the same two outcomes x^0 and x^1 under the true and surrogate valuations, while the payment difference is identical in the two utility comparisons and therefore cancels before the strategic-distortion bound is applied.

If $\sup_{\theta_i \in \Theta_i} D_i(\theta_i) \le \varepsilon_i$, then $D_i(\theta_i) \le \varepsilon_i$ for every $\theta_i \in \Theta_i$. Since z_{-i} and z_i' were arbitrary, the preceding inequality implies

$$\begin{aligned} &u_i(\theta_i; x(z_i', z_{-i}), p_i(z_i', z_{-i})) \\ &\qquad - u_i(\theta_i; x(E_i(\theta_i), z_{-i}), p_i(E_i(\theta_i), z_{-i})) \le \delta_i + \varepsilon_i \end{aligned}$$

for every true type θ_i, every profile of the other agents' reports, and every unilateral label deviation. Thus reporting $E_i(\theta_i)$ is $(\delta_i + \varepsilon_i)$-approximately dominant.

Now assume additionally that the allocation rule maximizes total surrogate welfare at every label profile. Fix an arbitrary true type profile $\theta = (\theta_1, \ldots, \theta_n)$ and write

$$z := E(\theta) = (E_1(\theta_1), \ldots, E_n(\theta_n)), \qquad \widehat{x} := x(z), \qquad x^\star := x^\star(\theta).$$

For each agent i, the definition of $D_i(\theta_i)$ applied to the ordered pair $(\widehat{x}, x^\star)$ gives

$$\left|\left[v_i(\theta_i, x^\star) - v_i(\theta_i, \widehat{x})\right] - \left[w_{i,z_i}(x^\star) - w_{i,z_i}(\widehat{x})\right]\right| \le D_i(\theta_i).$$

In particular,

$$v_i(\theta_i, x^\star) - v_i(\theta_i, \widehat{x}) \le w_{i,z_i}(x^\star) - w_{i,z_i}(\widehat{x}) + D_i(\theta_i).$$

Summing this inequality over all agents yields

$$\begin{aligned} &\sum_{i=1}^n v_i(\theta_i, x^\star) - \sum_{i=1}^n v_i(\theta_i, \widehat{x}) \\ &\qquad \le \sum_{i=1}^n w_{i,z_i}(x^\star) - \sum_{i=1}^n w_{i,z_i}(\widehat{x}) + \sum_{i=1}^n D_i(\theta_i). \end{aligned}$$

By the assumed definition of the allocation rule,

$$\widehat{x} = x(z) \in \arg\max_{x \in \mathcal{X}} \sum_{i=1}^{n} w_{i,z_i}(x),$$

and therefore $\sum_{i=1}^{n} w_{i,z_i}(\widehat{x}) \geq \sum_{i=1}^{n} w_{i,z_i}(x)$ for every $x \in \mathcal{X}$. In particular, taking $x = x^\star$ gives

$$\sum_{i=1}^{n} w_{i,z_i}(x^\star) - \sum_{i=1}^{n} w_{i,z_i}(\widehat{x}) \leq 0.$$

Substitution into the preceding welfare comparison gives

$$\sum_{i=1}^{n} v_i(\theta_i, x^\star) - \sum_{i=1}^{n} v_i(\theta_i, \widehat{x}) \leq \sum_{i=1}^{n} D_i(\theta_i).$$

Recalling that $x^\star = x^\star(\theta)$ and $\widehat{x} = x(E(\theta))$, we obtain

$$\sum_{i=1}^{n} v_i(\theta_i, x^\star(\theta)) - \sum_{i=1}^{n} v_i(\theta_i, x(E(\theta))) \leq \sum_{i=1}^{n} D_i(\theta_i),$$

which is equation 5. Finally, if $D_i(\theta_i) \leq \varepsilon_i$ for every i and every θ_i, then

$$\sum_{i=1}^{n} D_i(\theta_i) \leq \sum_{i=1}^{n} \varepsilon_i,$$

so the true welfare loss is at most $\sum_i \varepsilon_i$, as claimed. □

B.3.2 Surrogate Groves specialization

Corollary B.2 (Surrogate Groves mechanism)**.** *For every label profile $z \in Z$, let*

$$x(z) \in \arg\max_{x \in \mathcal{X}} \sum_{i=1}^{n} w_{i,z_i}(x),$$

and define

$$p_i(z) = h_i(z_{-i}) - \sum_{j \neq i} w_{j,z_j}(x(z)),$$

where h_i does not depend on z_i. Then the latent mechanism is 0-DSIC on the surrogate label space. Consequently, the true ex-post deviation gain is at most $D_i(\theta_i)$, and the allocation satisfies equation 5.

B.3.3 Proof of Corollary B.2

Proof. Fix an agent i, a label $z_i \in Z_i$, an arbitrary profile $z_{-i} \in Z_{-i}$ of the other agents' labels, and an arbitrary alternative report $z_i' \in Z_i$. Write

$$x^0 := x(z_i, z_{-i}), \qquad x^1 := x(z_i', z_{-i}).$$

Under the surrogate valuation w_{i,z_i}, agent i's utility from reporting z_i is

$$\begin{aligned} w_{i,z_i}(x^0) - p_i(z_i, z_{-i}) &= w_{i,z_i}(x^0) - \left(h_i(z_{-i}) - \sum_{j \neq i} w_{j,z_j}(x^0) \right) \\ &= w_{i,z_i}(x^0) + \sum_{j \neq i} w_{j,z_j}(x^0) - h_i(z_{-i}) \\ &= \sum_{j=1}^{n} w_{j,z_j}(x^0) - h_i(z_{-i}). \end{aligned}$$

If the same surrogate type z_i instead reports z_i', its surrogate utility is

$$\begin{aligned} w_{i,z_i}(x^1) - p_i(z_i', z_{-i}) &= w_{i,z_i}(x^1) - \left(h_i(z_{-i}) - \sum_{j \neq i} w_{j,z_j}(x^1) \right) \\ &= w_{i,z_i}(x^1) + \sum_{j \neq i} w_{j,z_j}(x^1) - h_i(z_{-i}). \end{aligned}$$

The crucial point is that the first term in this expression remains $w_{i,z_i}(x^1)$ rather than $w_{i,z_i'}(x^1)$: incentive compatibility evaluates the deviating report using agent i's original surrogate type z_i.

By construction,

$$x^0 = x(z_i, z_{-i}) \in \arg\max_{x \in \mathcal{X}} \left\{ w_{i,z_i}(x) + \sum_{j \neq i} w_{j,z_j}(x) \right\}.$$

Hence, for every $x \in \mathcal{X}$,

$$w_{i,z_i}(x^0) + \sum_{j \neq i} w_{j,z_j}(x^0) \geq w_{i,z_i}(x) + \sum_{j \neq i} w_{j,z_j}(x).$$

The deviating outcome $x^1 = x(z_i', z_{-i})$ is an element of $\mathcal{X}$, so we may take $x = x^1$ in this inequality and obtain

$$w_{i,z_i}(x^0) + \sum_{j \neq i} w_{j,z_j}(x^0) \geq w_{i,z_i}(x^1) + \sum_{j \neq i} w_{j,z_j}(x^1).$$

Subtracting the common quantity $h_i(z_{-i})$ from both sides gives

$$\begin{aligned} &w_{i,z_i}(x^0) + \sum_{j \neq i} w_{j,z_j}(x^0) - h_i(z_{-i}) \\ &\quad \geq w_{i,z_i}(x^1) + \sum_{j \neq i} w_{j,z_j}(x^1) - h_i(z_{-i}). \end{aligned}$$

Using the two surrogate-utility identities derived above, this becomes

$$w_{i,z_i}(x^0) - p_i(z_i, z_{-i}) \geq w_{i,z_i}(x^1) - p_i(z_i', z_{-i}),$$

or equivalently,

$$w_{i,z_i}(x(z_i', z_{-i})) - p_i(z_i', z_{-i}) \leq w_{i,z_i}(x(z_i, z_{-i})) - p_i(z_i, z_{-i}).$$

This is precisely equation 3 with $\delta_i = 0$. Since the choices of i, z_i, z_i', and z_{-i} were arbitrary, the latent mechanism is 0-DSIC on the entire surrogate label space.

The allocation rule in the corollary also satisfies

$$x(z) \in \arg\max_{x \in \mathcal{X}} \sum_{j=1}^{n} w_{j,z_j}(x),$$

which is exactly the additional allocation hypothesis in Theorem 4.1. Applying that theorem with $\delta_i = 0$ therefore gives, for every true type θ_i, every z_{-i}, and every alternative label z_i',

$$\begin{aligned} &u_i(\theta_i; x(z_i', z_{-i}), p_i(z_i', z_{-i})) \\ &\quad - u_i(\theta_i; x(E_i(\theta_i), z_{-i}), p_i(E_i(\theta_i), z_{-i})) \\ &\quad \leq D_i(\theta_i). \end{aligned}$$

Thus the true ex-post deviation gain is at most $D_i(\theta_i)$. The same application of Theorem 4.1 gives, for every true type profile θ,

$$\sum_{i=1}^{n} v_i(\theta_i, x^\star(\theta)) - \sum_{i=1}^{n} v_i(\theta_i, x(E(\theta))) \leq \sum_{i=1}^{n} D_i(\theta_i),$$

which is exactly equation 5. This proves the corollary. □

B.4 Proofs for Section 5

B.4.1 Proof of Theorem 5.1

Proof. We prove the two inequalities in equation 8 by converting, in one direction, an ambient strategic cover into a deterministic representation and, in the other direction, every deterministic representation into an ambient strategic cover.

Let

$$K = N_{\mathrm{amb}}(\mathcal{V}, \varepsilon),$$

and choose centers $c_1, \dots, c_K \in \mathcal{Q}$ satisfying

$$\mathcal{V} \subseteq \bigcup_{z=1}^{K} \overline{B}_{d_{\mathrm{str}}}(c_z, \varepsilon).$$

For each c_z, choose its anchored representative $w_z \in \mathbb{R}^m$. For every $q = [v] \in \mathcal{V}$, the covering property guarantees that the set

$$I(q) := \{z \in \{1, \dots, K\} : d_{\mathrm{str}}(q, c_z) \le \varepsilon\}$$

is nonempty. Define a deterministic encoder by a fixed tie-breaking rule, for example

$$E(q) := \min I(q).$$

For each label z and menu $M = (S, p)$, define $a_z(M)$ to be a fixed deterministic selection from $\mathrm{BR}_{w_z}(M)$. Thus

$$a_z(M) \in \mathrm{BR}_{w_z}(M) \qquad \text{and hence} \qquad \mathrm{Reg}_{w_z}(a_z(M); M) = 0.$$

Now fix an arbitrary $q = [v] \in \mathcal{V}$ and an arbitrary menu $M \in \mathcal{M}$. By the definition of E,

$$d_{\mathrm{str}}(v, w_{E(q)}) = d_{\mathrm{str}}(q, c_{E(q)}) \le \varepsilon,$$

where the first equality uses the fact that d_{str} is the quotient metric and therefore does not depend on the chosen representatives. Applying Theorem 2.3 with $\delta = 0$ gives

$$\mathrm{Reg}_v\big(a_{E(q)}(M); M\big) \le d_{\mathrm{str}}(v, w_{E(q)}) \le \varepsilon.$$

Since q and M were arbitrary,

$$\sup_{[v] \in \mathcal{V}} \sup_{M \in \mathcal{M}} \mathrm{Reg}_v\big(a_{E([v])}(M); M\big) \le \varepsilon.$$

Therefore a cover by $N_{\mathrm{amb}}(\mathcal{V}, \varepsilon)$ ambient balls induces an admissible representation with the same number of labels, and consequently

$$K_\varepsilon \le N_{\mathrm{amb}}(\mathcal{V}, \varepsilon).$$

For the reverse inequality, consider any deterministic representation with K labels satisfying equation 6. For every label z, let

$$C_z := \{\widetilde{v} : [v] \in \mathcal{V},\ E([v]) = z\},$$

where $\widetilde{v}$ denotes the anchored representative of $[v]$. Some of these cells may be empty. Discard every empty label and let $J \le K$ denote the number of remaining nonempty cells. Because $\mathcal{V}$ is bounded in d_{str}, each nonempty C_z is a nonempty bounded cell of anchored valuations, so Theorem 3.1 applies to it. The same response rule a_z appearing in the representation satisfies

$$\sup_{v \in C_z} \sup_{M \in \mathcal{M}} \mathrm{Reg}_v(a_z(M); M) \le \varepsilon.$$

By the definition of the deterministic cell minimax value,

$$\mathfrak{R}(C_z) = \inf_{a \in \mathcal{A}} \sup_{v \in C_z} \sup_{M \in \mathcal{M}} \mathrm{Reg}_v(a(M); M) \le \sup_{v \in C_z} \sup_{M \in \mathcal{M}} \mathrm{Reg}_v(a_z(M); M) \le \varepsilon.$$

Theorem 3.1 gives the exact identity

$$\rho(C_z) = \mathfrak{R}(C_z),$$

and also states that the infimum in $\rho(C_z)$ is attained. Hence there is an anchored surrogate $w_z^\star \in \mathbb{R}^m$ such that

$$\sup_{v \in C_z} d_{\mathrm{str}}(v, w_z^\star) = \rho(C_z) = \mathfrak{R}(C_z) \le \varepsilon.$$

Equivalently, if $c_z^\star = [w_z^\star] \in \mathcal{Q}$, then

$$\{[v] : v \in C_z\} \subseteq \overline{B}_{d_{\mathrm{str}}}(c_z^\star, \varepsilon).$$

The nonempty encoder cells partition $\mathcal{V}$. Therefore the J ambient balls centered at the classes $c_z^\star$ cover all of $\mathcal{V}$, and thus

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) \le J \le K.$$

This conclusion holds for every K for which an admissible deterministic representation exists. Taking the smallest such K yields

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) \le K_\varepsilon.$$

Together with the first inequality, this proves

$$K_\varepsilon = N_{\mathrm{amb}}(\mathcal{V}, \varepsilon).$$

It remains only to translate the alphabet size into a fixed-length binary budget. A binary word of length B can encode at most 2^B distinct labels, so every fixed-length description of K_ε labels must satisfy

$$2^B \ge K_\varepsilon, \qquad \text{hence} \qquad B \ge \lceil \log_2 K_\varepsilon \rceil .$$

Conversely, $\lceil \log_2 K_\varepsilon \rceil$ bits provide at least K_ε binary codewords, so assigning one distinct word to each label achieves this bound. Therefore

$$B_\varepsilon = \lceil \log_2 K_\varepsilon \rceil .$$

□

B.4.2 Proof of Proposition 5.2

Proof. For $h \in L$, define

$$\|h\|_{\mathrm{str},L} := \max_{x,y \in \mathcal{X}} |\langle h, \psi(x) - \psi(y) \rangle| .$$

We first verify that this is a genuine norm on L. Absolute homogeneity and the triangle inequality follow immediately from the corresponding properties of the absolute value and the maximum. If $\|h\|_{\mathrm{str},L} = 0$, then

$$\langle h, \psi(x) - \psi(y) \rangle = 0 \qquad \forall x, y \in \mathcal{X}.$$

Since L is the span of the vectors $\psi(x) - \psi(y)$, the preceding equalities imply $h \perp L$. But $h \in L$, and hence $h = 0$. Thus $\|\cdot\|_{\mathrm{str},L}$ is positive definite on L and is a norm.

For arbitrary $\theta, \theta' \in \mathbb{R}^d$, every difference $\psi(x) - \psi(y)$ lies in L, so its inner product with the component in $L^\perp$ vanishes. Therefore

$$\begin{aligned} d_{\mathrm{str}}(v_\theta, v_{\theta'}) &= \max_{x,y \in \mathcal{X}} |\langle \theta - \theta', \psi(x) - \psi(y) \rangle| \\ &= \max_{x,y \in \mathcal{X}} |\langle P_L\theta - P_L\theta', \psi(x) - \psi(y) \rangle| \\ &= \|P_L\theta - P_L\theta'\|_{\mathrm{str},L} . \end{aligned}$$

In particular, $[v_\theta] = [v_{P_L\theta}]$. If $h, h' \in L$ and $[v_h] = [v_{h'}]$, then the left-hand side above is zero, so $\|h - h'\|_{\mathrm{str},L} = 0$ and therefore $h = h'$. Consequently the map

$$T : L \to \mathcal{Q}, \qquad T(h) := [v_h],$$

is injective and isometric when L is equipped with $\|\cdot\|_{\mathrm{str},L}$. Moreover,

$$\mathcal{V}_\Theta = T(P_L\Theta).$$

Thus all strategically visible variation of the parametric family is exactly the variation of $P_L\Theta$ inside the r-dimensional normed space $(L, \|\cdot\|_{\mathrm{str},L})$.

Because L is finite dimensional and $r > 0$, all norms on L are equivalent. Hence there exist constants $0 < \alpha \le \beta < \infty$, depending only on L and the feature vectors, such that

$$\alpha \|h\|_2 \le \|h\|_{\mathrm{str},L} \le \beta \|h\|_2 \qquad \forall h \in L.$$

Write $A := P_L\Theta$. By assumption, A is bounded. After enlarging a constant if necessary, there exists $R < \infty$ such that

$$A \subseteq B_L(0, R),$$

where $B_L(0, R)$ is the Euclidean ball of radius R in the r-dimensional space L. Fix $\varepsilon > 0$ and put

$$\delta := \frac{\varepsilon}{\beta}.$$

The compact set $\overline{A}$ admits a finite maximal δ-separated set $\{h_1, \dots, h_N\}$ in Euclidean distance. Maximality means that every $h \in \overline{A}$ lies within Euclidean distance δ of at least one h_j; otherwise h could be added while preserving δ-separation. Hence, for every $h \in A$, there is a j for which

$$\|h - h_j\|_{\mathrm{str},L} \le \beta \|h - h_j\|_2 \le \beta\delta = \varepsilon.$$

Applying the isometry T, the ambient centers $T(h_j) \in \mathcal{Q}$ therefore form a strategic ε-cover of $\mathcal{V}_\Theta$. To bound N, note that the Euclidean balls $B_L(h_j, \delta/2)$ are pairwise disjoint, while each is contained in $B_L(0, R + \delta/2)$. If λ_r denotes r-dimensional Lebesgue measure on L, then

$$N\, \lambda_r(B_L(0, \delta/2)) \le \lambda_r(B_L(0, R + \delta/2))\,.$$

Since Euclidean ball volume scales as the rth power of the radius,

$$N \le \left(\frac{R + \delta/2}{\delta/2}\right)^r = \left(1 + \frac{2R}{\delta}\right)^r = \left(1 + \frac{2R\beta}{\varepsilon}\right)^r.$$

For all sufficiently small ε, this gives a constant $C_2 < \infty$, independent of ε, such that

$$N_{\mathrm{amb}}(\mathcal{V}_\Theta, \varepsilon) \le C_2 \varepsilon^{-r}.$$

We next prove a matching lower bound that remains valid even though the covering centers are allowed to lie anywhere in the ambient quotient space. Because A contains a nonempty relative open ball in L, there exist $h_0 \in L$ and $\rho > 0$ such that

$$B_L(h_0, \rho) \subseteq A.$$

In particular, the closed Euclidean ball $\overline{B}_L(h_0, \rho/2)$ is contained in A. For sufficiently small ε, set

$$\eta := \frac{4\varepsilon}{\alpha} < \frac{\rho}{2},$$

and choose a finite maximal η-separated set $\{g_1, \dots, g_M\}$ in $\overline{B}_L(h_0, \rho/2)$. By maximality, the Euclidean η-balls centered at these points cover $\overline{B}_L(h_0, \rho/2)$. Hence

$$\lambda_r(B_L(0, \rho/2)) \le M\, \lambda_r(B_L(0, \eta))\,,$$

and therefore

$$M \ge \left(\frac{\rho}{2\eta}\right)^r = \left(\frac{\alpha\rho}{8\varepsilon}\right)^r.$$

For distinct i, j, the separation and the lower norm-equivalence bound give

$$\begin{aligned} d_{\mathrm{str}}(v_{g_i}, v_{g_j}) &= \|g_i - g_j\|_{\mathrm{str},L} \\ &\ge \alpha \|g_i - g_j\|_2 \\ &\ge \alpha\eta = 4\varepsilon > 2\varepsilon. \end{aligned}$$

Now consider any ambient strategic ball $\overline{B}_{d_{\mathrm{str}}}(c, \varepsilon)$ in $\mathcal{Q}$. If it contained both $T(g_i)$ and $T(g_j)$, the triangle inequality in the quotient metric would imply

$$d_{\mathrm{str}}(T(g_i), T(g_j)) \le d_{\mathrm{str}}(T(g_i), c) + d_{\mathrm{str}}(c, T(g_j)) \le 2\varepsilon,$$

contradicting the preceding strict inequality. Thus one ambient ε-ball can contain at most one of the M points $T(g_j)$. Since all of these points belong to $\mathcal{V}_\Theta$, every ambient ε-cover must contain at least M balls. Consequently there is a constant $C_1 > 0$, independent of ε, such that, for all sufficiently small ε,

$$N_{\mathrm{amb}}(\mathcal{V}_\Theta, \varepsilon) \geq C_1 \varepsilon^{-r}.$$

Combining the upper and lower estimates with Theorem 5.1 yields

$$C_1 \varepsilon^{-r} \leq K_\varepsilon \leq C_2 \varepsilon^{-r}$$

for all sufficiently small ε. Taking base-two logarithms gives

$$r \log_2(1/\varepsilon) + \log_2 C_1 \leq \log_2 K_\varepsilon \leq r \log_2(1/\varepsilon) + \log_2 C_2,$$

which is exactly

$$\log_2 K_\varepsilon = r \log_2(1/\varepsilon) + O(1), \qquad \varepsilon \downarrow 0.$$

The proof also shows why the interior condition is needed: the normed space L identifies the largest strategically visible linear space, but if $P_L\Theta$ occupies only a lower-dimensional subset of L, the packing lower bound above need not scale as ε^{-r}. □

B.4.3 An explicit family with exact covering complexity

Let $m = 2r + 1$ and

$$v_a = (0, a_1, -a_1, \ldots, a_r, -a_r), \qquad a \in [0,1]^r.$$

Then

$$d_{\mathrm{str}}(v_a, v_b) = 2 \left\| a - b \right\|_\infty,$$

and, for every $0 < \varepsilon \leq 1$,

$$K_\varepsilon = \left\lceil \frac{1}{\varepsilon} \right\rceil^r.$$

This remains exact even when covering centers may lie anywhere in the ambient quotient space.

Proof. For this subsection, write

$$\mathcal{V} := \{[v_a] : a \in [0,1]^r\} \subset \mathcal{Q}.$$

Fix $a, b \in [0,1]^r$ and write $d_i := a_i - b_i$. Then

$$v_a - v_b = (0, d_1, -d_1, \ldots, d_r, -d_r).$$

If $D := \max_{1 \leq i \leq r} |d_i|$, the coordinates of this vector have maximum D and minimum $-D$. By the oscillation characterization of strategic distortion,

$$\begin{aligned} d_{\mathrm{str}}(v_a, v_b) &= \mathrm{osc}(v_a - v_b) \\ &= \max_x (v_a(x) - v_b(x)) - \min_x (v_a(x) - v_b(x)) \\ &= D - (-D) \\ &= 2D = 2 \left\| a - b \right\|_\infty. \end{aligned}$$

Let

$$q := \left\lceil \frac{1}{\varepsilon} \right\rceil.$$

For the upper bound, partition each coordinate interval $[0,1]$ into q intervals of length $1/q$ and use their midpoints

$$t_j := \frac{2j-1}{2q}, \qquad j = 1, \ldots, q.$$

For every multi-index $\mathbf{j} = (j_1, \ldots, j_r) \in \{1, \ldots, q\}^r$, let

$$c_{\mathbf{j}} := (t_{j_1}, \ldots, t_{j_r})$$

and use the valuation class $[v_{c_\mathbf{j}}]$ as a covering center. Given any $a \in [0,1]^r$, choose in each coordinate a midpoint t_{j_i} from the interval containing a_i. Then

$$\|a - c_\mathbf{j}\|_\infty \le \frac{1}{2q},$$

so the metric identity just proved yields

$$d_{\mathrm{str}}(v_a, v_{c_\mathbf{j}}) = 2\,\|a - c_\mathbf{j}\|_\infty \le \frac{1}{q} \le \varepsilon.$$

Thus q^r strategic balls of radius ε cover the family, and

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) \le q^r.$$

It remains to prove that allowing centers outside the family, and even outside its r-dimensional affine image in $\mathcal{Q}$, cannot reduce this number. If $q = 1$, then the desired lower bound is the trivial inequality $N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) \ge 1$. Suppose therefore that $q \ge 2$. Because

$$q - 1 < \frac{1}{\varepsilon},$$

we have

$$\frac{1}{q-1} > \varepsilon.$$

Consider the parameter grid

$$G := \left\{0, \frac{1}{q-1}, \frac{2}{q-1}, \dots, 1\right\}^r.$$

It contains exactly q^r points. If $a, b \in G$ are distinct, then they differ in at least one coordinate by at least $1/(q-1)$, and hence

$$\|a - b\|_\infty \ge \frac{1}{q-1} > \varepsilon.$$

Therefore

$$d_{\mathrm{str}}(v_a, v_b) = 2\,\|a - b\|_\infty > 2\varepsilon.$$

Now let $c \in \mathcal{Q}$ be an arbitrary ambient center. If a single ball $\overline{B}_{d_{\mathrm{str}}}(c, \varepsilon)$ contained both $[v_a]$ and $[v_b]$ for distinct $a, b \in G$, the triangle inequality would give

$$d_{\mathrm{str}}(v_a, v_b) \le d_{\mathrm{str}}([v_a], c) + d_{\mathrm{str}}(c, [v_b]) \le 2\varepsilon,$$

which is impossible. Hence every ambient ball of radius ε covers at most one of the q^r grid points. Any ambient cover of the whole family must therefore use at least q^r balls:

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) \ge q^r.$$

Combining the upper and lower bounds gives

$$N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) = q^r = \left\lceil \frac{1}{\varepsilon} \right\rceil^r.$$

Finally, Theorem 5.1 converts this exact ambient covering number into the claimed exact message complexity,

$$K_\varepsilon = N_{\mathrm{amb}}(\mathcal{V}, \varepsilon) = \left\lceil \frac{1}{\varepsilon} \right\rceil^r.$$

□

B.4.4 Operational finite-alphabet coding statement

Let $V_1, \dots, V_N$ be i.i.d. copies of V. For any block code

$$f_N : \mathrm{supp}(V)^N \to \{1, \dots, J\}, \qquad g_N : \{1, \dots, J\} \to \mathcal{W}^N,$$

write $\widehat{W}^N = g_N(f_N(V^N))$ and

$$D_N := \frac{1}{N} \sum_{t=1}^{N} \mathbb{E}[d_{\mathrm{str}}(V_t, \widehat{W}_t)].$$

Then every such code satisfies

$$\log_2 J \geq N R_{\mathcal{W}}(D_N).$$

Conversely, for every test channel $q(w \mid v)$ and every $\xi > 0$, sufficiently large block lengths admit codes satisfying

$$\frac{1}{N} \log_2 J \leq I(V; W) + \xi, \qquad D_N \leq \mathbb{E}[d_{\mathrm{str}}(V, W)] + \xi.$$

B.4.5 Proof of Proposition 5.3

Proof. Let

$$Q := [V]$$

denote the quotient class of the source valuation. The key observation is that, for every $v \in \mathrm{supp}(V)$ and $w \in \mathcal{W}$, the distortion $d_{\mathrm{str}}(v, w)$ depends on v only through $[v]$. We first show that this pointwise invariance is also sufficient at the level of the rate–distortion optimization.

Take an arbitrary test channel $q(w \mid v)$ and let it induce a joint law for (V, Q, W). For every quotient class q_0 with $\Pr(Q = q_0) > 0$, define

$$\bar{q}(w \mid q_0) := \Pr(W = w \mid Q = q_0) = \sum_{v : [v] = q_0} \Pr(V = v \mid Q = q_0) q(w \mid v).$$

Now define a new channel on the original source alphabet by

$$q^\circ(w \mid v) := \bar{q}(w \mid [v]).$$

Under q°, the conditional distribution of W given Q is exactly the same as under the original channel. Hence the joint distribution of (Q, W) is unchanged. Since $d_{\mathrm{str}}(v, w)$ is a function only of $([v], w)$,

$$\mathbb{E}_{q^\circ}[d_{\mathrm{str}}(V, W)] = \mathbb{E}_q[d_{\mathrm{str}}(V, W)].$$

Under the original channel, Q is a deterministic function of V, so the data-processing inequality gives

$$I_q(Q; W) \leq I_q(V; W).$$

Under q°, the variables form the Markov chain

$$V \longrightarrow Q \longrightarrow W,$$

and Q is still a deterministic function of V. Therefore

$$\begin{aligned} I_{q^\circ}(V; W) &= I_{q^\circ}(V, Q; W) \\ &= I_{q^\circ}(Q; W) + I_{q^\circ}(V; W \mid Q) \\ &= I_{q^\circ}(Q; W) \\ &= I_q(Q; W) \\ &\leq I_q(V; W). \end{aligned}$$

Thus every feasible channel can be replaced, without changing its expected distortion and without increasing its mutual information, by a channel that depends on v only through $[v]$. Conversely, every channel $\bar{q}(w \mid Q)$ can be lifted to the source alphabet by setting $q(w \mid v) = \bar{q}(w \mid [v])$, in which case $I(V; W) = I(Q; W)$. Consequently

$$R_{\mathcal{W}}(D) = \min_{\bar{q}(w|q) :\, \mathbb{E}[d_{\mathrm{str}}(Q, W)] \leq D} I(Q; W),$$

where $d_{\mathrm{str}}(Q, W)$ denotes the well-defined quotient distortion $d_{\mathrm{str}}(v, w)$ for any representative v of Q. This expression depends on the source only through the distribution of $Q = [V]$.

We next establish monotonicity and convexity. If $D_1 \le D_2$, every channel feasible for distortion level D_1 is also feasible for D_2, and therefore

$$R_{\mathcal{W}}(D_2) \le R_{\mathcal{W}}(D_1).$$

Hence $R_{\mathcal{W}}$ is nonincreasing. Because both source and reproduction alphabets are finite, the set of stochastic matrices is compact, expected distortion is continuous and affine in the channel, and mutual information is continuous. Thus whenever the distortion constraint is feasible, the minimum in equation 9 is attained. Let D_1, D_2 belong to the effective domain, let q_1 and q_2 attain $R_{\mathcal{W}}(D_1)$ and $R_{\mathcal{W}}(D_2)$, respectively, and fix $\lambda \in [0, 1]$. Define the mixture channel

$$q_\lambda(w \mid v) := \lambda q_1(w \mid v) + (1 - \lambda) q_2(w \mid v).$$

Its distortion satisfies

$$\begin{aligned}\mathbb{E}_{q_\lambda}[d_{\mathrm{str}}(V, W)] &= \lambda \mathbb{E}_{q_1}[d_{\mathrm{str}}(V, W)] + (1 - \lambda)\mathbb{E}_{q_2}[d_{\mathrm{str}}(V, W)] \\ &\le \lambda D_1 + (1 - \lambda) D_2.\end{aligned}$$

For a fixed source distribution, mutual information is convex as a function of the conditional distribution $q(w \mid v)$. Hence

$$I_{q_\lambda}(V; W) \le \lambda I_{q_1}(V; W) + (1 - \lambda) I_{q_2}(V; W) = \lambda R_{\mathcal{W}}(D_1) + (1 - \lambda) R_{\mathcal{W}}(D_2).$$

Since q_λ is feasible at distortion $\lambda D_1 + (1 - \lambda) D_2$, we obtain

$$R_{\mathcal{W}}(\lambda D_1 + (1 - \lambda) D_2) \le \lambda R_{\mathcal{W}}(D_1) + (1 - \lambda) R_{\mathcal{W}}(D_2),$$

which proves convexity on the effective domain.

Assume now that $\mathcal{W}$ contains at least one representative of every quotient class in the support of V. We prove the zero-distortion identity. If a test channel satisfies

$$\mathbb{E}[d_{\mathrm{str}}(V, W)] = 0,$$

then nonnegativity of d_{str} implies

$$d_{\mathrm{str}}(V, W) = 0 \qquad \text{almost surely.}$$

By Proposition 2.2, zero strategic distortion is equivalent to equality of quotient classes, so

$$[W] = [V] = Q \qquad \text{almost surely.}$$

Because $[W]$ is a deterministic function of W, this shows that Q is a deterministic function of W and hence

$$H(Q \mid W) = 0.$$

Since Q is also a deterministic function of V, the data-processing inequality and the preceding identity give

$$I(V; W) \ge I(Q; W) = H(Q) - H(Q \mid W) = H(Q).$$

Thus every zero-distortion channel has rate at least $H(Q)$. For the reverse inequality, for each quotient class q in the support of Q, choose one $w(q) \in \mathcal{W}$ satisfying $[w(q)] = q$, and define the deterministic test channel

$$W = w(Q).$$

Then $d_{\mathrm{str}}(V, W) = 0$ almost surely. Distinct quotient classes must be assigned to distinct reproduction symbols, because one vector w has only one quotient class, so Q and W determine one another on their supports. Consequently

$$I(V; W) = I(Q; W) = H(Q).$$

The lower and upper bounds coincide, proving

$$R_{\mathcal{W}}(0) = H(Q) = H([V]).$$

We now prove the finite-block converse stated above. Let

$$\widehat{W}^N = g_N(f_N(V^N))$$

be the reproduction sequence generated by an arbitrary deterministic block code with J messages. Since $\widehat{W}^N$ can take at most J distinct values,

$$H(\widehat{W}^N) \le \log_2 J.$$

Therefore

$$\log_2 J \ge H(\widehat{W}^N) \ge I(V^N; \widehat{W}^N).$$

Because $V_1, \dots, V_N$ are independent copies of V,

$$H(V^N) = \sum_{t=1}^{N} H(V_t).$$

Using the chain rule for conditional entropy and the fact that conditioning cannot increase entropy,

$$\begin{aligned} H(V^N \mid \widehat{W}^N) &= \sum_{t=1}^{N} H(V_t \mid V_1, \dots, V_{t-1}, \widehat{W}^N) \\ &\le \sum_{t=1}^{N} H(V_t \mid \widehat{W}^N) \\ &\le \sum_{t=1}^{N} H(V_t \mid \widehat{W}_t). \end{aligned}$$

Subtracting this upper bound from $H(V^N)$ gives

$$\begin{aligned} I(V^N; \widehat{W}^N) &= H(V^N) - H(V^N \mid \widehat{W}^N) \\ &\ge \sum_{t=1}^{N} \Big(H(V_t) - H(V_t \mid \widehat{W}_t) \Big) \\ &= \sum_{t=1}^{N} I(V_t; \widehat{W}_t). \end{aligned}$$

Define the single-letter distortions

$$D_t := \mathbb{E}[d_{\mathrm{str}}(V_t, \widehat{W}_t)].$$

For each t, the joint law of $(V_t, \widehat{W}_t)$ induces a legitimate single-letter test channel from the source alphabet to $\mathcal{W}$. Since V_t has the same marginal law as V, the definition of $R_{\mathcal{W}}$ implies

$$I(V_t; \widehat{W}_t) \ge R_{\mathcal{W}}(D_t).$$

Combining the preceding inequalities and then applying convexity of $R_{\mathcal{W}}$ gives

$$\begin{aligned} \log_2 J &\ge \sum_{t=1}^{N} I(V_t; \widehat{W}_t) \\ &\ge \sum_{t=1}^{N} R_{\mathcal{W}}(D_t) \\ &\ge N R_{\mathcal{W}}\left(\frac{1}{N} \sum_{t=1}^{N} D_t \right) \\ &= N R_{\mathcal{W}}(D_N). \end{aligned}$$

This proves the claimed finite-block converse.

Finally, we give the standard finite-alphabet achievability argument in the present notation. Fix a test channel $q(w \mid v)$, let

$$P_{VW}(v, w) = P_V(v) q(w \mid v), \qquad \bar{D} := \mathbb{E}[d_{\mathrm{str}}(V, W)],$$

and fix $\xi > 0$. Because the alphabets are finite,

$$d_{\max} := \max_{v \in \mathrm{supp}(V),\, w \in \mathcal{W}} d_{\mathrm{str}}(v, w) < \infty.$$

Choose a sufficiently small typicality tolerance and, for each block length N, generate

$$J_N := \left\lceil 2^{N(I(V;W)+\xi/2)} \right\rceil$$

reproduction sequences independently according to the product marginal P_W^N. Given a source sequence v^N, the encoder searches the codebook for a reproduction sequence w^N that is jointly typical with v^N under P_{VW} and sends the index of the first such sequence; if none exists, it sends a fixed default index. The decoder outputs the indexed reproduction sequence.

The finite-alphabet covering lemma states that, because the codebook rate is strictly larger than $I(V; W)$, the probability that a typical source sequence has no jointly typical codeword tends to zero as $N \to \infty$. The probability that V^N itself is not typical also tends to zero. Moreover, for every jointly typical pair (v^N, w^N), boundedness of the single-letter distortion and the typical average lemma give

$$\left| \frac{1}{N} \sum_{t=1}^{N} d_{\mathrm{str}}(v_t, w_t) - \bar{D} \right| \longrightarrow 0$$

uniformly as the typicality tolerance tends to zero. On the exceptional event where typical encoding fails, the per-letter distortion is at most $d_{\max}$. It follows that the expectation, over the random codebook and the i.i.d. source, of the achieved average distortion converges to at most $\bar{D}$. Hence, for all sufficiently large N, there exists at least one deterministic realization of the codebook for which

$$D_N \le \bar{D} + \xi = \mathbb{E}[d_{\mathrm{str}}(V, W)] + \xi.$$

For the same code,

$$\frac{1}{N} \log_2 J_N \le I(V; W) + \xi$$

for all sufficiently large N, because the ceiling contributes only a vanishing $O(1/N)$ term. We have therefore obtained deterministic finite-block codes satisfying

$$\frac{1}{N} \log_2 J \le I(V; W) + \xi, \qquad D_N \le \mathbb{E}[d_{\mathrm{str}}(V, W)] + \xi.$$

This is precisely the classical finite-alphabet rate–distortion achievability theorem specialized to the strategic distortion d_{str} and the fixed reproduction alphabet $\mathcal{W}$, completing the proof. □

B.5 Proofs and auxiliary certificates for Section 6

B.5.1 Computational complexity of the strategic minimax quantizer

Using

$$d_{\mathrm{str}}(v, w) = \max_x (v_x - w_x) - \min_x (v_x - w_x),$$

one evaluation of $d_{\mathrm{str}}(v, w)$ requires a single pass over the m coordinates of $v - w$, and therefore costs $O(m)$ time. For each of the N training valuations, the assignment step compares the valuation with all K current centers. Hence the complete assignment step costs

$$O(NKm).$$

Now fix the cells induced by these assignments. For a nonempty cell C, the exact center update of equation 2 depends on the quantities

$$A_{xy}(C) = \max_{v \in C} \{v(y) - v(x)\}, \qquad x, y \in \mathcal{X}.$$

For each $v \in C$, all m^2 ordered differences $v(y) - v(x)$ can be inspected in $O(m^2)$ time. Updating the corresponding running maxima therefore constructs the complete matrix $A(C)$ in

$$O(|C|m^2)$$

time. Since the cells form a partition of the N assigned valuations, their cardinalities satisfy

$$\sum_{z=1}^{K} |C_z| = N,$$

where empty cells contribute zero. Consequently, constructing all cell matrices costs

$$\sum_{z=1}^{K} O(|C_z| m^2) = O(Nm^2).$$

For a fixed nonempty cell, Theorem 3.1 reduces the center update to a maximum-cycle-mean problem on the complete directed graph on m outcomes, which has m vertices and m^2 directed edges. A standard maximum-cycle-mean algorithm runs in

$$O(|V||E|) = O(m^3)$$

time. Once the optimal value $r = \rho(C)$ is known, the associated difference constraints

$$w_y - w_x \geq A_{xy}(C) - r \qquad \forall x, y \in \mathcal{X}$$

form a feasible system with m^2 constraints. A feasible anchored potential can likewise be recovered by a Bellman–Ford-type longest-path computation after reversing signs, again in $O(m^3)$ time. Thus an exact center update costs $O(m^3)$ per nonempty cell and at most

$$O(Km^3)$$

over all cells. Empty-cell centers are left unchanged and add only lower-order bookkeeping.

Combining the assignment, cell-construction, and center-update costs gives

$$O(NKm + Nm^2 + Km^3)$$

per alternating iteration, up to lower-order bookkeeping terms.

B.5.2 Proof of Proposition 6.1

Proof. Because Proposition 2.2 shows that d_{str} induces a metric on $\mathcal{Q}$, we may use its triangle inequality throughout. For a candidate center $w \in \mathcal{Q}$, define

$$R_C(w) := \sup_{v \in C} d_{\mathrm{str}}(v, w), \qquad R_{\widetilde{C}}(w) := \sup_{\widetilde{v} \in \widetilde{C}} d_{\mathrm{str}}(\widetilde{v}, w).$$

Since the two cells are finite and indexed by the same set I, these quantities may equivalently be written as

$$R_C(w) = \max_{j \in I} d_{\mathrm{str}}(v^j, w), \qquad R_{\widetilde{C}}(w) = \max_{j \in I} d_{\mathrm{str}}(\widetilde{v}^j, w).$$

For every $j \in I$, the triangle inequality and the assumed matched perturbation bound give

$$\begin{aligned} d_{\mathrm{str}}(v^j, w) &\leq d_{\mathrm{str}}(v^j, \widetilde{v}^j) + d_{\mathrm{str}}(\widetilde{v}^j, w) \\ &\leq \eta + d_{\mathrm{str}}(\widetilde{v}^j, w). \end{aligned}$$

Taking the maximum over j yields

$$R_C(w) \leq \eta + R_{\widetilde{C}}(w).$$

The same argument with v^j and $\widetilde{v}^j$ interchanged gives

$$\begin{aligned} d_{\mathrm{str}}(\widetilde{v}^j, w) &\leq d_{\mathrm{str}}(\widetilde{v}^j, v^j) + d_{\mathrm{str}}(v^j, w) \\ &\leq \eta + d_{\mathrm{str}}(v^j, w), \end{aligned}$$

and therefore

$$R_{\widetilde{C}}(w) \leq \eta + R_C(w).$$

Thus, for every candidate center w,

$$\left| R_C(w) - R_{\widetilde{C}}(w) \right| \leq \eta.$$

By definition,

$$\rho(C) = \inf_{w \in \mathcal{Q}} R_C(w), \qquad \rho(\widetilde{C}) = \inf_{w \in \mathcal{Q}} R_{\widetilde{C}}(w).$$

Using the first of the two preceding inequalities before taking the infimum gives

$$\begin{aligned} \rho(C) &= \inf_{w \in \mathcal{Q}} R_C(w) \\ &\le \inf_{w \in \mathcal{Q}} \big(R_{\widetilde{C}}(w) + \eta \big) \\ &= \rho(\widetilde{C}) + \eta. \end{aligned}$$

Using the reverse inequality in exactly the same way gives

$$\rho(\widetilde{C}) \le \rho(C) + \eta.$$

The two inequalities together imply

$$\Big| \rho(C) - \rho(\widetilde{C}) \Big| \le \eta.$$

We next evaluate the true cell at an optimal center of the perturbed cell. By Theorem 3.1, an optimal center $\widetilde{w}^\star$ of $\widetilde{C}$ exists and satisfies

$$R_{\widetilde{C}}(\widetilde{w}^\star) = \rho(\widetilde{C}).$$

Applying the pointwise comparison of the two radius functions at this particular center gives

$$\begin{aligned} \sup_{v \in C} d_{\mathrm{str}}(v, \widetilde{w}^\star) &= R_C(\widetilde{w}^\star) \\ &\le R_{\widetilde{C}}(\widetilde{w}^\star) + \eta \\ &= \rho(\widetilde{C}) + \eta. \end{aligned}$$

The first part of the proposition already established that

$$\rho(\widetilde{C}) \le \rho(C) + \eta.$$

Substituting this bound into the preceding display yields

$$\sup_{v \in C} d_{\mathrm{str}}(v, \widetilde{w}^\star) \le \rho(C) + 2\eta,$$

which proves the second assertion. Notice that the two occurrences of η have distinct sources: one transfers the fixed center $\widetilde{w}^\star$ from the perturbed cell back to the true cell, while the other compares the optimal radii of the two cells.

It remains to prove the nearest-center statement. Fix a true valuation v, an estimated valuation $\widetilde{v}$ satisfying

$$d_{\mathrm{str}}(v, \widetilde{v}) \le \eta,$$

and an arbitrary fixed codebook $\{w_1, \dots, w_K\}$. Let

$$\widehat{z} := \widehat{z}(\widetilde{v}) \in \operatorname*{arg\,min}_{1 \le z \le K} d_{\mathrm{str}}(\widetilde{v}, w_z),$$

and independently choose

$$z^\star \in \operatorname*{arg\,min}_{1 \le z \le K} d_{\mathrm{str}}(v, w_z).$$

Since the codebook is finite, both minimizers exist. By the triangle inequality,

$$d_{\mathrm{str}}(v, w_{\widehat{z}}) \le d_{\mathrm{str}}(v, \widetilde{v}) + d_{\mathrm{str}}(\widetilde{v}, w_{\widehat{z}}).$$

The definition of $\widehat{z}$ implies

$$d_{\mathrm{str}}(\widetilde{v}, w_{\widehat{z}}) \le d_{\mathrm{str}}(\widetilde{v}, w_{z^\star}).$$

Applying the triangle inequality once more gives

$$d_{\mathrm{str}}(\widetilde{v}, w_{z^\star}) \le d_{\mathrm{str}}(\widetilde{v}, v) + d_{\mathrm{str}}(v, w_{z^\star}).$$

Combining these three inequalities and using symmetry of d_{str} yields

$$\begin{aligned} d_{\mathrm{str}}(v, w_{\widehat{z}}) &\le d_{\mathrm{str}}(v, \widetilde{v}) + d_{\mathrm{str}}(\widetilde{v}, w_{\widehat{z}}) \\ &\le d_{\mathrm{str}}(v, \widetilde{v}) + d_{\mathrm{str}}(\widetilde{v}, w_{z^\star}) \\ &\le d_{\mathrm{str}}(v, \widetilde{v}) + d_{\mathrm{str}}(\widetilde{v}, v) + d_{\mathrm{str}}(v, w_{z^\star}) \\ &\le 2\eta + d_{\mathrm{str}}(v, w_{z^\star}) \\ &= 2\eta + \min_{1 \le z \le K} d_{\mathrm{str}}(v, w_z). \end{aligned}$$

Therefore

$$d_{\mathrm{str}}(v, w_{\widehat{z}(\widetilde{v})}) \le \min_{1 \le z \le K} d_{\mathrm{str}}(v, w_z) + 2\eta,$$

as claimed. □

B.5.3 Proof of Proposition 6.2

Proof. Let

$$n := N_{\mathrm{val}}, \qquad D_j := D(V_j^{\mathrm{val}}), \qquad \widehat{D}_{\mathrm{val}} = \frac{1}{n} \sum_{j=1}^{n} D_j.$$

The encoder and codebook are fixed independently of the validation sample. Equivalently, if they were obtained from random training data or from randomness internal to the training algorithm, let $\mathcal{F}_{\mathrm{train}}$ denote the sigma-field containing all such training information and condition on $\mathcal{F}_{\mathrm{train}}$. Conditional on this sigma-field, the encoder E, the centers $w_1, \dots, w_K$, and hence the function

$$D(v) = d_{\mathrm{str}}(v, w_{E(v)})$$

are fixed.

Because the validation observations are i.i.d. and independent of the training information, conditional on $\mathcal{F}_{\mathrm{train}}$ the random variables

$$D_1, \dots, D_n$$

are i.i.d. copies of $D(V)$. By assumption,

$$0 \le D_j \le L_0 \qquad \text{almost surely.}$$

If $L_0 = 0$, then $D(V) = 0$ almost surely and the claimed inequality is immediate. Hence suppose $L_0 > 0$.

For the realized fixed encoder and decoder, write

$$\mu := \mathbb{E}[D(V) \mid \mathcal{F}_{\mathrm{train}}].$$

The one-sided Hoeffding inequality for independent random variables taking values in an interval of length L_0 gives, for every $t > 0$,

$$\Pr\Big(\mu - \widehat{D}_{\mathrm{val}} \ge t \;\Big|\; \mathcal{F}_{\mathrm{train}}\Big) \le \exp\left(-\frac{2nt^2}{L_0^2}\right).$$

Choose

$$t = L_0 \sqrt{\frac{\log(1/\alpha)}{2n}}.$$

Then

$$\begin{aligned} \exp\left(-\frac{2nt^2}{L_0^2}\right) &= \exp\left(-\frac{2n}{L_0^2} L_0^2 \frac{\log(1/\alpha)}{2n}\right) \\ &= \exp(-\log(1/\alpha)) \\ &= \alpha. \end{aligned}$$

Therefore, conditional on the fixed trained encoder and codebook,

$$\Pr\left(\mu \le \widehat{D}_{\mathrm{val}} + L_0 \sqrt{\frac{\log(1/\alpha)}{2N_{\mathrm{val}}}} \;\middle|\; \mathcal{F}_{\mathrm{train}}\right) \ge 1 - \alpha.$$

The bound holds for every realization of the training information for which the assumed boundedness condition holds. Taking expectation of the conditional probability with respect to $\mathcal{F}_{\text{train}}$ therefore preserves the same probability lower bound:

$$\Pr\left(\mu \le \widehat{D}_{\text{val}} + L_0\sqrt{\frac{\log(1/\alpha)}{2N_{\text{val}}}}\right) \ge 1-\alpha.$$

Once the trained encoder and codebook are fixed, μ is precisely the population mean distortion denoted in the proposition by $\mathbb{E}[D(V)]$. Hence, with probability at least $1-\alpha$,

$$\mathbb{E}[D(V)] \le \widehat{D}_{\text{val}} + L_0\sqrt{\frac{\log(1/\alpha)}{2N_{\text{val}}}},$$

which is the desired certificate. □

B.5.4 A FINITE-NET EXTENSION TO UNIFORM CERTIFICATION

Proposition B.3 (Finite-net extension)**.** *Let $\mathcal{V} \subset \mathcal{Q}$, let $\mathcal{S} \subset \mathcal{V}$ be an h-net in d_{str}, and let $\{w_1, \dots, w_K\} \subset \mathcal{Q}$ be a fixed codebook. If*

$$\max_{u\in\mathcal{S}} \min_{1\le z\le K} d_{\text{str}}(u, w_z) \le \varepsilon, \tag{11}$$

then

$$\sup_{v\in\mathcal{V}} \min_{1\le z\le K} d_{\text{str}}(v, w_z) \le \varepsilon + h. \tag{12}$$

Consequently, nearest-center surrogate decoding gives all-menu regret at most $\varepsilon + h$ uniformly over $\mathcal{V}$.

Proof. Fix an arbitrary valuation $v \in \mathcal{V}$. Since $\mathcal{S}$ is an h-net of $\mathcal{V}$ in the metric d_{str}, there exists some $u \in \mathcal{S}$ such that

$$d_{\text{str}}(v, u) \le h.$$

By the assumed covering property equation 11,

$$\min_{1\le z\le K} d_{\text{str}}(u, w_z) \le \varepsilon.$$

Because the codebook is finite, the minimum is attained. Choose $z_u \in \{1, \dots, K\}$ satisfying

$$d_{\text{str}}(u, w_{z_u}) = \min_{1\le z\le K} d_{\text{str}}(u, w_z) \le \varepsilon.$$

The triangle inequality for d_{str} then gives

$$\begin{aligned}\min_{1\le z\le K} d_{\text{str}}(v, w_z) &\le d_{\text{str}}(v, w_{z_u}) \\ &\le d_{\text{str}}(v, u) + d_{\text{str}}(u, w_{z_u}) \\ &\le h + \varepsilon.\end{aligned}$$

The valuation $v \in \mathcal{V}$ was arbitrary. Taking the supremum over all $v \in \mathcal{V}$ therefore yields

$$\sup_{v\in\mathcal{V}} \min_{1\le z\le K} d_{\text{str}}(v, w_z) \le \varepsilon + h,$$

which proves equation 12.

For completeness, consider the operational implication. Define a nearest-center encoder by choosing, for each $v \in \mathcal{V}$,

$$E(v) \in \arg\min_{1\le z\le K} d_{\text{str}}(v, w_z).$$

The bound just proved implies

$$d_{\text{str}}(v, w_{E(v)}) \le \varepsilon + h \qquad \text{for every } v \in \mathcal{V}.$$

Upon receiving the label $E(v)$ and observing a menu $M \in \mathcal{M}$, let the decoder choose any deterministic tie-broken best response

$$a_{E(v)}(M) \in \mathrm{BR}_{w_{E(v)}}(M).$$

Since this action is an exact best response to the surrogate valuation,

$$\mathrm{Reg}_{w_{E(v)}}(a_{E(v)}(M); M) = 0.$$

Applying Theorem 2.3 with $\delta = 0$, true valuation v, and surrogate $w_{E(v)}$ gives

$$\begin{aligned} \mathrm{Reg}_v(a_{E(v)}(M); M) &\le d_{\mathrm{str}}(v, w_{E(v)}) \\ &\le \varepsilon + h. \end{aligned}$$

This inequality holds simultaneously for every $v \in \mathcal{V}$ and every $M \in \mathcal{M}$. Hence nearest-center surrogate decoding has uniform all-menu regret at most $\varepsilon + h$ over $\mathcal{V}$. □