Bundling Complements*

Weijie Zhong

Abstract

I develop a duality-based multi-dimensional screening framework with a geometric characterization of combinatorial preferences. For a mechanism to be optimal, the type distribution pins down *required* directions of binding feasibility constraints, while the complementarity among bundles determines the *covered* directions; optimality reduces to full coverage of required directions. I apply the framework to a one-parameter family in which every bundle containing a fixed *core* of items earns a complementarity premium. Two thresholds organize the optimum: above a lower threshold the grand bundle must be offered; above a higher threshold a *core-peripheral* menu — a bundled core with optional add-ons that are not sold standalone — is optimal. The tight distributional condition for finiteness of the higher threshold is *inclusivity*, that the menu exclude no near-top buyer.

Keywords: multi-dimensional screening, bundling, complementarity, optimal transport, mechanism design, core-peripheral mechanisms.

JEL Classification: D42, D82, D86.

*Stanford Graduate School of Business. Email: weijie.zhong@stanford.edu. I thank Roberto Corrao and Laura Doval for valuable comments and suggestions.

1 Introduction

Many products and services are sold with a core component paired with optional add-ons unavailable standalone: AI assistants bundle premium reasoning models with a base subscription, enterprise vendors license security and compliance modules only atop their core platform, and cable operators have long sold premium channels only as add-ons to a basic tier. In each, an add-on is worth more paired with the core than alone — a premium model inside a workflow that also runs faster base models for routine steps, a security module integrated with the platform whose data it monitors. Such cross-good interactions generate *combinatorial preferences*: a bundle's value is not the sum of its parts, and the interaction — complementarity, in the case these examples share — drives the bundling decision.

This paper traces the comparative statics of the optimal screening mechanism as the strength of complementarity among goods varies. The model considers a monopoly seller who sells multiple indivisible goods to a buyer with a multidimensional private type. A valuation map determines the buyer type's value for each bundle of goods, which can exhibit complementarity, substitutability or arbitrary combinatorial preferences.

The methodology extends the optimal-transport duality of Daskalakis et al. (2017) and Kleiner and Manelli (2019) and surfaces a geometric structure of combinatorial preferences hidden in the dual cost. In the dual problem, the type distribution and preference jointly determine a *differential virtual value* (positive where providing buyer surplus increases revenue, negative where it decreases). The differential virtual value pins down the *required directions* of the dual transport — the directions of binding feasibility constraints under an optimal mechanism. The allocation feasibility constraint attaches to each bundle a cone of *covered directions* along which allocating the bundle maximizes the surplus differential. Optimality of a candidate menu reduces to a geometric coverage condition: every required direction must lie inside the covered cone of the allocated bundle. Under linear valuation, the problem pulls back to the primitive type space and the two sides fully separate — required directions are determined solely by the type distribution, and the covered cones are determined solely by the valuation map.

This separation reframes the comparative-statics landscape. The additive-preferences literature traces how the optimum responds to the type distribution at fixed pref-

erences — working entirely through the distribution-to-required-directions side. The bundle-to-cone mapping opens the symmetric question, how the optimum responds to preferences, through the preference-to-cone side: a bundle's cone widens when its items exhibit complementarity and narrows when they exhibit substitutability. The application of the framework exploits this preference-side handle to trace the optimum as complementarity varies.

I apply the framework to a one-parameter slice of combinatorial preferences. Complementarity is parameterized by a single scalar $\alpha > 0$ multiplying the value of every bundle that contains a fixed set of *core items* $C \subseteq [N]$; two thresholds organize how the optimum responds as α varies. A lower threshold α_*, reflecting the spirit of free disposal of the grand bundle, marks the level above which the grand bundle is forced into the optimal menu (Theorem 3). An upper threshold α^*, derived from sufficient cone *coverage*, marks the level above which a *core-peripheral* menu is optimal: every menu option either excludes the buyer or sells a bundle containing the entire core (Theorem 2). I identify *inclusivity* — all top-value types being included by the core-peripheral mechanism — as the binding distributional condition for finiteness of α^* (Theorem 4). The gap between α_* and α^* — generically nonempty — is the intermediate regime in which the menu may carry both the grand bundle and Ã la carte options.

The paper's two contributions are the methodology, portable beyond the one-parameter slice studied here, and the comparative-statics characterization it enables. Applying the methodology identifies settings with mild, or explicit, primitive sufficient conditions for core-peripheral bundling. Pure bundling is the sharpest case: the certifying hypothesis of Theorem 2 — *alignment regularity*, a stochastic-dominance condition on the menu–density pair — reduces to the canonical Myersonian regularity of the bundle-value distribution, and pure bundling is optimal for all large α exactly when the distribution is inclusive (Theorem 5). An explicit two-tier family with one add-on is solved under log-concave core marginals (Section 5.2), and a tractable iid Beta family shows the conditions are met in higher dimensions, with closed-form solutions (Section 5.1). A 2-good iid uniform example illustrates the arc at all three regimes: the lower threshold $\alpha_* = 1/2$ is explicit, an upper-threshold bound $\alpha^* \lesssim 1.253$ is constructed by a dual-transport line family (Section S.5), and the intermediate regime — between separate sales and pure bundling — contains the additive optimum of Pavlov (2011) as one slice.

The rest of the paper proceeds as follows. Section 2 formulates the problem and develops the dual representation, including the bundle-to-cone mapping and the saddle-point criterion. Section 3 illustrates the geometry through four canonical two-good examples. Section 4 states the upper- and lower-threshold theorems and the necessity of the inclusivity condition. Section 5 applies them: pure bundling characterized, the iid Beta benchmark, a two-tier family, and the market segments the thresholds organize. Section 6 concludes. Proofs are in the appendix and the Online Supplement.

1.1 Literature review

Optimal mechanisms under additive multi-dimensional preferences. The structural foundations of multidimensional screening were laid by McAfee and McMillan (1988), Armstrong (1996), and Rochet and Chone (1998), who established the logic of exclusion, bunching, and convexity once the buyer has several dimensions of private information. A central feature of the additive multi-good optimum, emphasized by Manelli and Vincent (2007) and traced in two-good detail by Pavlov (2011), is that randomization (lotteries over bundles) is intrinsic rather than a technical artifact: deterministic posted prices are generically suboptimal. Manelli and Vincent (2006) characterizes when deterministic *bundle* pricing nevertheless solves the multi-good monopoly problem. Rochet (2024) surveys recent developments in this literature.[1]

Optimal-transport duality. A methodological strand within this literature characterizes the additive multi-good optimum through a transport dual. Daskalakis et al. (2017) and Kleiner and Manelli (2019) establish the duality and the saddle-point characterization of the optimum against a dual cost c, applying the framework to a range of canonical settings. In these foundational developments the cost c is taken as a primitive object, with the preference structure of the screening problem implicit in its definition. This paper builds on these foundations by decomposing the cost: under linearity, c admits an explicit decomposition over its gradient polytope, with each vertex (a bundle's gradient) carrying its own cone of covered directions. The bundle-to-cone mapping is the geometric representation of this

[1] The complexity of the additive optimum has also been studied through approximation lenses: Armstrong (1999) shows that simple tariffs can be approximately optimal in product spaces of large dimension; Hart and Nisan (2017, 2019) quantify how well finite-menu simple mechanisms perform and how large their menus must grow; and Briest et al. (2015) exhibit a regime in which lottery pricing dominates deterministic item pricing by an unbounded factor with three or more goods.

decomposition, surfacing the preference primitives as explicit geometric objects in the analysis.

Sufficient conditions for pure bundling. A separate strand identifies sufficient conditions on the type distribution under which the additive optimum collapses to pure bundling or to separate sales. McAfee et al. (1989) show that bundling dominates separate sales when two non-complementary goods have stochastically independent values. Giannakopoulos and Koutsoupias (2018) give exact two-good sufficient conditions for pure bundling under uniform distributions. Menicucci et al. (2015) establishes pure-bundling optimality for two additive, independently distributed goods under non-negative virtual valuations. A parallel literature on robust mechanism design, set in additive or nearly additive environments, characterizes the optimum under various ambiguity sets: Carroll (2017) shows separate sales are robust when only marginal distributions are known; Deb and Roesler (2023) shows pure bundling is informationally robust when the prior is exchangeable across items; Che and Zhong (2025) traces categorical bundling through a hierarchy of moment-based ambiguity sets. Relatedly, Frick et al. (2026) show that as the seller's information about the buyer grows precise, pure bundling becomes asymptotically optimal and strictly beats separate sales. Hart and Reny (2015) observe that revenue can fail to be monotone in the buyer's value even in additive multi-good settings. These papers locate the bundling pressure entirely in the type distribution or in the seller's information about it — preferences are additive, so the force comes from how the distribution tilts the marginal-value comparison toward the bundle. This paper introduces complementarity as a separate preference primitive and justifies the practically prevalent core-peripheral menu — which nests pure bundling as a special case — as the optimum under sufficient complementarity.

Combinatorial preferences and bundling extensions. The closer comparators are papers that allow non-additive bundle valuations. Yang (2025) develops a theory of nested bundling in a one-dimensional type-space setting amenable to the Myersonian approach. Haghpanah and Hartline (2021) works in multi-dimensional types with general non-additive bundle valuations and identifies a monotone-ratio condition — via a path decomposition reducing the problem to a family of one-

dimensional restrictions — under which pure bundling is optimal.[2] Thanassoulis (2004) studies substitutes, and Bikhchandani and Mishra (2024) shows that selling heterogeneous objects under symmetric additive values is equivalent to selling identical objects under decreasing marginal values. Relative to this literature and the additive-preferences literature above, the contribution here is to make complementarity a geometric primitive — the bundle cone's covered directions — that varies independently of the distribution, opening the preference-side comparative statics these papers, working at fixed preferences, do not pursue.

2 Model and Dual Representation

This section formalizes the buyer–seller environment, develops the optimal-transport dual of the seller's problem, and isolates two technical conditions — tangency and linearity — on which the geometric analysis of later sections rests.

2.1 Preferences, mechanisms, and feasibility

The model takes the buyer's bundle valuations as the primary preference primitive. Consider N goods and a single buyer. Let $[N] = \{1, \dots, N\}$ index goods and let $\mathcal{P} := 2^{[N]} \setminus \{\varnothing\}$ denote the full family of nonempty bundles. The buyer's primitive type is $x \in D \subset \mathbb{R}^d$, where D is compact with piecewise-smooth boundary and density $f \in C^2(D)$ that is strictly positive on D.[3] [4] The seller offers a subfamily $\mathcal{F} \subseteq \mathcal{P}$ of bundles; let $K := |\mathcal{F}|$. For each $B \in \mathcal{F}$, the buyer's valuation $v_B : D \to \mathbb{R}$ is smooth. Collect these into the valuation map

$$v(x) := \big(v_B(x)\big)_{B \in \mathcal{F}} \in \mathbb{R}^K,$$

whose image $Y := v(D)$ is the *valuation manifold*. Separating D and Y isolates two distinct ingredients: D carries the distributional data — the density f and the differential virtual value introduced in Section 2.2 — while Y carries the preference

[2] Neither setup satisfies the tangency condition of Definition 1, so both characterizations are parallel to, rather than specializations of, Theorem 2.

[3] These conditions are maintained for the necessity theorems (Theorems 3 and 4). The sufficiency results — the optimality certificates of Theorem 2 and its applications — extend to densities that are only C^1 and positive on $\operatorname{int} D$ with integrable boundary behavior, covering the singular Beta family of Section 5.1; the extension is given in Section S.1 of the Online Supplement.

[4] The type-space dimension d is decoupled from the number of goods N: the canonical case is $d = N$, but $d < N$ (for instance, a one-dimensional Myersonian type) and $D = Y$ (types stated directly as valuation vectors) are recovered as special cases of the same framework.

data, the bundle-value coordinates and the quantity set $\mathcal{Q}$ defined below. This separation is innocuous in the additive benchmark but essential once preferences are combinatorial.

Two conventions on (D, v) are maintained throughout, both structural requirements for the dual construction of Section 2.2 rather than mere normalizations. The domain has a *bottom type* $x_0 \in D$ that is coordinatewise-least, $x_0 \preceq x$ for every $x \in D$, and each valuation v_B is nondecreasing, so that $v(x_0)$ is the coordinatewise-least point of Y; the bottom type is where individual rationality pins the buyer's surplus and where the dual measure carries its atom. Every application takes D to be a box and $x_0 = 0$.

The seller posts a *menu* $M = \{(\sigma_\ell, p_\ell)\}_{\ell=0}^{L}$: option $\ell \geq 1$ offers a bundle lottery $\sigma_\ell \in \Delta(\mathcal{F})$ at price $p_\ell > 0$, and $\ell = 0$ is the outside option — the empty bundle $\varnothing$ ($v_\varnothing \equiv 0$) at price 0. A buyer of type x takes the option of greatest surplus $\sum_B \sigma_\ell(B)\, v_B(x) - p_\ell$, ties broken in the seller's favor.[5] Option ℓ's *marginal-inclusion vector* w_ℓ records the probability of allocating each item: $w_{\ell,i} := \sum_{B \ni i} \sigma_\ell(B)$.

Equivalently, M is a *direct mechanism*: an allocation $q(\cdot) \in \mathcal{Q} := \{q \in \mathbb{R}_+^K : \mathbf{1}^\top q \leq 1\}$ — the simplex of bundle lotteries, q_B the probability of bundle B — and a transfer $t(\cdot)$. By Rochet and Chone (1998), incentive compatibility and individual rationality hold exactly when the indirect utility $u : Y \to \mathbb{R}_+$ is the restriction to Y of a convex function on $\mathbb{R}^K$ whose subgradients all lie in $\mathcal{Q}$, and, for a measurable selection $q(y) \in \partial u(y) \cap \mathcal{Q}$,

$$u(y) = \max_{y' \in Y} \{y \cdot q(y') - t(y')\}, \qquad t(y) = q(y) \cdot y - u(y). \tag{1}$$

The design problem is then the choice of such a convex u.

2.2 Revenue as a linear functional

The expected-revenue integral mixes u with its gradient through the IC formula (1); turning it into a linear functional of u alone requires one technical condition on the valuation map.

Definition 1 (Tangency condition)**.** *The valuation map v satisfies the* tangency condition *if there exists a smooth vector field $a : D \to \mathbb{R}^d$ such that*

$$Jv(x)\, a(x) = v(x) \qquad \text{for every } x \in D,$$

[5] The highest-revenue option in the buyer's indifference set, with bundle offers preferred to exclusion.

where $Jv(x)$ is the Jacobian of v at x.

Tangency is satisfied by every linear valuation, and by every homogeneous-of-degree-one valuation wherever it is differentiable (Euler's identity);[6] the linear case leads the rest of the paper. Tangency does two things at once. First, it removes a notational ambiguity: because Y is at most d-dimensional inside $\mathbb{R}^K$, the subdifferential $\partial u(y)$ in (1) is generally not a singleton, and the quantity $q(y) \in \partial u(y)$ is a measurable selection from a fiber of admissible values rather than a unique gradient. Tangency makes the choice of selection revenue-irrelevant (Lemma 2), so the standard abuse $q(y) = \nabla u(y)$ is harmless. Second, tangency underlies the divergence-theorem rewriting of revenue developed below.

Under Definition 1, the transfer of any incentive-compatible mechanism implementing u satisfies $t(x) = a(x)\cdot\nabla U(x) - U(x)$ at λ_D-a.e. x, where $U := u\circ v$ (Lemma 2, proved in Section A by Rademacher's theorem and the convex chain rule). Revenue therefore depends on the mechanism only through u, and equals

$$\int_D \big[a(x)\cdot\nabla U(x) - U(x)\big] f(x)\,\mathrm{d}x.$$

The divergence theorem then produces a signed measure μ on D defined by

$$\mu(A) := \mathbf{1}_A(x_0) + \int_{\partial D} \mathbf{1}_A(x) f(x)\, a(x)\cdot n(x)\,\mathrm{d}\sigma(x)^{7} - \int_A \big[\nabla\cdot(f(x)a(x)) + f(x)\big]\,\mathrm{d}x,$$

where x_0 is the bottom type of Section 2.1 and $n(x)$ is the outward normal on ∂D. Since $v(x_0)$ is the least point of Y and u is nondecreasing, $U = u\circ v$ attains its minimum over D at x_0; individual rationality leaves no surplus to the lowest type, so it is without loss to normalize $U(x_0) = 0$. With the atom placed at x_0, μ

[6] Nonlinear HD1 examples — CES aggregators $\big(\sum_{i\in B} x_i^p\big)^{1/p}$, Leontief $\min_{i\in B} x_i$, Cobb–Douglas $\prod_{i\in B} x_i^{1/|B|}$ — are tangent with $a(x) = x$ where differentiable; each is non-smooth on some diagonal or lower face, and no nonlinear HD1 valuation is differentiable at the origin, so global smoothness on D is special to the linear case the theorems use. More generally, any shifted-HD1 valuation $v(x) = w(x - z_0)$ with w HD1 is tangent with $a(x) = x - z_0$; for instance, the affine form $v_B(x) = \sum_{i\in B}(x_i - c)$ with per-good baseline $c \neq 0$ is tangent but not HD1. The shift moves only the vector field a: the dual measure constructed below keeps its atom at the bottom type x_0 (the origin in every application), since that is where individual rationality pins U.

[7] Here σ denotes Lebesgue surface measure on ∂D. In the standing applications, where $D = \prod_j [0, \bar{v}_j]$ is a box and $a(x) = x$, the integrand $f\, a\cdot n$ vanishes on the bottom faces (there $a\cdot n = -x_j = 0$), so only the top faces $F_j := \{x_j = \bar{v}_j\}$ contribute; σ_j denotes the restriction of σ to F_j, and these per-face measures appear in the OS proofs.

depends only on (f, a, D) and not on the utility being optimized; the atom contributes $U(x_0) = 0$ to the revenue, and, as f integrates to one, $\mu(D) = 0$ regardless of where x_0 lies. Because $\mu(D) = 0$, adding a constant to U leaves $\int_D U \, \mathrm{d}\mu$ unchanged, so the normalization $U(x_0) = 0$ selects a representative rather than constraining the admissible class (cf. Section A). Let $\nu := v_{\#}\mu$ be the pushforward of μ onto $Y = v(D)$. The signed measure ν — the *differential virtual value* of the screening problem — encodes how marginal revenue flows from raising buyer surplus across the type space: its positive part identifies types where leaving the buyer high surplus increases revenue, its negative part identifies types where it decreases.[8] The seller's problem becomes the choice of an admissible utility from $\mathcal{U} := \{u : Y \to \mathbb{R}_+ : u = \bar{u}|_Y$ for some convex $\bar{u} : \mathbb{R}^K \to \mathbb{R}$ with $\partial \bar{u}(z) \subseteq \mathcal{Q}$ for all $z\}$:

$$\sup_{u \in \mathcal{U}} \int_Y u(y) \, \mathrm{d}\nu(y). \tag{2}$$

The simplex $\mathcal{Q}$ encodes the allocation primitives, and the seller's problem is to widen the surplus gap between revenue-rewarding and revenue-absorbing buyers subject to the feasibility cap on $\mathcal{Q}$. Section 2.3 reformulates this in Y as a transport problem whose cost is the support function of $\mathcal{Q}$. The Y-space formulation does not yet separate the two sides: ν depends on (f, a, D) and v jointly through the pushforward, and Section 2.4 delivers the separation by specializing to linear valuations and pulling back to D, where μ depends only on (f, a, D) while the bundle gradients live in a polytope determined entirely by V.

2.3 Transport representation and the saddle-point criterion

Reformulating (2) via the support function of $\mathcal{Q}$ converts the primal constraint into a geometric transport cost. Because $\mathcal{Q}$ is the simplex of bundle lotteries, its support function is the bundle-coordinate one-sided max:

$$c(y, y') := \left|(y - y')_+\right|_\infty = \max\left\{0, \max_{B \in \mathcal{F}} (y_B - y'_B)\right\} = \sup_{q \in \mathcal{Q}} q \cdot (y - y').$$

The dual transport cost is this one-sided max cost. A quantity q *covers* a displacement $y - y'$ when $q \cdot (y - y') = c(y, y')$ — equivalently, when q maximizes the surplus differential along $y - y'$ among all allocations in $\mathcal{Q}$. The set of displacements covered by q is its *cone of covered directions*, made explicit on D in Section 2.4; complementing it, the dual program below identifies the *required directions* along which

[8] In one dimension, $\nu([\theta, 1]) = \varphi(\theta) f(\theta)$, where φ is Myerson's virtual valuation, so ν is exactly the *differential* of Myerson's virtual revenue density, extended to multi-dimensional settings.

the seller must transport mass between revenue-rewarding and revenue-absorbing types in ν. As in Daskalakis et al. (2017); Kleiner and Manelli (2019), the dual variable is a transport plan. Assume from here on that the valuation manifold Y is convex — as in the standing applications, where D is a convex box and v is linear, so that $Y = VD$ is a linear image of a convex set.[9] The dual program is

$$\inf_{\gamma \in \Gamma_+(Y \times Y)} \int_{Y \times Y} c(y, y') \, \mathrm{d}\gamma(y, y') \quad \text{s.t.} \quad \gamma_1 - \gamma_2 \succeq_{\mathrm{mcvx}} \nu, \tag{3}$$

where $\Gamma_+(Y \times Y)$ denotes the finite positive Borel measures on $Y \times Y$, γ_1 and γ_2 are the marginals, and $\eta \succeq_{\mathrm{mcvx}} \nu$ (*monotone-convex stochastic dominance*) means $\int \phi \, \mathrm{d}\eta \geq \int \phi \, \mathrm{d}\nu$ for every coordinatewise nondecreasing convex ϕ on Y.

Theorem 1 (Strong duality)**.** *The primal problem* (2) *and the dual problem* (3) *have the same value, and both attain their optima.*

The proof, given in Section A, adapts the conic-linear-programming argument of Daskalakis et al. (2017) and Kleiner and Manelli (2019) to the valuation manifold and one-sided max cost; convexity of the domain is what its function-space argument uses.[10]

The following complementary-slackness condition pins down a primal-dual saddle point. It gives sufficient — and generally not necessary — conditions for a finite-menu mechanism and a transport plan to form an optimal primal-dual pair.

Proposition 1. *Suppose $u \in \mathcal{U}$ is piecewise linear on a finite polyhedral partition $\{C_\ell\}_{\ell=0}^{L}$ of Y, with*

$$u(y) = a_\ell + q^\ell \cdot y \qquad \text{for } y \in C_\ell,$$

where each quantity $q^\ell \in \mathcal{Q}$. Let $\gamma \in \Gamma_+(Y \times Y)$ be such that:

(i) *for γ-almost every (y, y'), both y and y' lie in the closure of one linearity cell C_ℓ;*[11]

[9] For nonconvex Y — possible for the nonlinear tangent valuations of Definition 1 — pose the primal and the dual on $\mathrm{conv}(Y)$, with the admissibility constraints extended there and the objective unchanged; the proof in Section A is written in that generality.

[10] The saddle-point criteria below (Propositions 1 and 2) consume only weak duality together with an explicitly constructed plan supported on $Y \times Y$, so no downstream result depends on the convexity assumption.

[11] This condition is not an independent restriction at optimality: when u is primal-optimal, any dual-optimal γ whose sink marginal is absolutely continuous satisfies it automatically — see Remark 2. I retain it in the statement because the criterion is applied to candidate pairs whose optimality is not known in advance.

(ii) *for γ-almost every (y, y') and any cell ℓ with $y, y' \in \overline{C_\ell}$ (as provided by (i)),*

$$u(y) - u(y') = q^\ell \cdot (y - y') = c(y, y');$$

(iii) $\gamma_1 - \gamma_2 = \nu$.

Then (u, γ) is a primal-dual saddle point. In particular, u solves the primal and γ solves the dual.

Geometrically, the complementary-slackness equality $q^\ell \cdot (y - y') = c(y, y')$ requires the displacement $y - y'$ to lie in the cone of directions that the active quantity q^ℓ covers.

Remark 1 (The no-ironing condition). The dual program (3) constrains the plan only by the dominance $\gamma_1 - \gamma_2 \succeq_{\text{mcvx}} \nu$; condition (iii) sharpens the dominance to equality, so the plan transports ν itself rather than a dominating rearrangement of it. This is the framework's *no-ironing condition* — the multi-dimensional analog of ruling out the ironing pathology of one-dimensional Myersonian screening.

2.4 Linear utility and the bundle-to-cone mapping

For the remainder of the paper, valuations are linear in the buyer's type. Linearity separates the two primitives of the problem: the type distribution and the buyer's preferences enter through distinct terms, so the effect of each can be isolated. That separation is what Section 4's construction exploits.[12]

Assumption 1 (Linearity)**.** *The valuation map v is linear in x: $v(x) = Vx$ for a constant matrix $V \in \mathbb{R}^{K \times d}$ whose rows V_B encode the bundle valuations $v_B(x) = V_B \cdot x$.*

Under Assumption 1, the Jacobian $Jv \equiv V$ is constant, so the gradients $V^\top q$ on primitive types live in the polytope

$$\mathcal{P}_V := V^\top \mathcal{Q} = \text{conv}\{0,\, V_B^\top : B \in \mathcal{F}\} \subset \mathbb{R}^d.$$

Its support function pulls back the OT cost from Y to D and singles out, for each bundle, the displacement cone on which that bundle's gradient is active.

[12] The construction extends to valuations tangent to a linear map at each type; the main theorems use linearity throughout.

Lemma 1 (Bundle-to-cone mapping)**.** *Under Assumption 1, the OT cost is translation-invariant on D,*

$$c(v(x), v(x')) \;=\; c_V(x - x'), \qquad \text{where } c_V(h) := \sup_{p \in \mathcal{P}_V} p \cdot h = \max\Big\{0,\, \max_{B \in \mathcal{F}} v_B(h)\Big\}.$$

The gradient $V_B^\top$ attains the supremum at h if and only if h lies in the bundle cone

$$K_B^{\mathcal{F}} := \big\{h \in \mathbb{R}^d : v_B(h) \geq v_{B'}(h) \text{ for every } B' \in \mathcal{F} \cup \{\varnothing\}\big\},$$

where $v_\varnothing \equiv 0$.

The proof, in Section A, follows from support-function calculus on the polytope $\mathcal{P}_V$. The assignment $B \mapsto K_B^{\mathcal{F}}$ is the *bundle-to-cone mapping*; the cone $K_B^{\mathcal{F}}$ is the set of directions B covers against the alternatives in $\mathcal{F}$. It translates the saddle-point criterion of Proposition 1 into primitive coordinates.

Proposition 2 (Bundle-cone saddle-point criterion)**.** *Under Assumption 1, let finite menu $M = \{(\sigma_\ell, p_\ell)\}_{\ell=0}^L$ be incentive compatible and individually rational, with indirect utility u_M. If $\gamma \in \Gamma_+(D \times D)$ satisfies*

(i) *both points of γ-almost every (x, x') lie in the closure of one cell C_ℓ;*

(ii) *for γ-almost every (x, x') and any cell ℓ with $x, x' \in \overline{C_\ell}$ (as provided by (i)), the displacement $x - x'$ lies in every bundle cone $K_B^{\mathcal{F}}$ option ℓ may allocate ($B \in \operatorname{supp} \sigma_\ell$);*

(iii) *$\gamma_1 - \gamma_2 = \mu$;*

then $(u_M, v_\# \gamma)$ is a primal-dual saddle point; in particular M is optimal.

The proof, given in Section A, reduces the criterion to Proposition 1 through the valuation map. The proposition is the framework's methodological pivot: optimality of a candidate mechanism factors into two halves that the conditions separate. The *required directions* — those along which the seller must widen the surplus differential between revenue-rewarding and revenue-absorbing types — are set entirely by the type distribution, and condition (iii)'s marginal identity $\gamma_1 - \gamma_2 = \mu$ pins them down. The *covered directions* — displacements along which the cell's allocation maximizes the surplus differential — are set entirely by the menu's bundle cones, and condition (ii) confines each displacement to them. The cell partition is where the two meet: condition (i) couples them, and the mechanism is optimal whenever, cell by cell, the active option's cone covers every required direction — the criterion is sufficient, not necessary.

Relation to Daskalakis et al. (2017); Kleiner and Manelli (2019). In the additive case — $v_B(x) = \sum_{i\in B} x_i$ over the full family $\mathcal{F} = \mathcal{P}$ on a box-shaped D, so that $\mathcal{P}_V = [0,1]^N$ and the cost is the separable $c_V(h) = \sum_i (h_i)_+$ — this proposition specializes to their duality. The combinatorial generalization acts on the cost: c extends from this one-sided ℓ^1 cost to the one-sided max over arbitrary bundle values, and Lemma 1 decomposes c_V over the gradient polytope, each vertex (a bundle's gradient) carrying its own cone of covered directions. That geometric content — absent when the cost is a primitive, as in Daskalakis et al. (2017); Kleiner and Manelli (2019) — is what exposes complementarity as a parameter scaling the gradients.

Standing core-peripheral form. The main characterization results specialize to the *core-peripheral form*

$$v_B(x) = \alpha^{\mathbf{1}\{C\subseteq B\}} \sum_{i\in B} x_i, \tag{4}$$

where $C \subseteq [N]$ is a fixed nonempty set of *core items*, $P := [N]\setminus C$ the complementary *peripheral items*, and $\alpha > 0$ parameterizes the premium generated by complementarity: $\alpha > 1$ is complementary, $\alpha = 1$ additive, and $0 < \alpha < 1$ substitute-like. When $C = [N]$, only the grand bundle earns the α-multiplier; when $|C| < N$, every bundle containing the core earns it. The bundle cone specializes to

$$K_B^{\mathcal{F}} = \left\{ h \in \mathbb{R}^d : \alpha^{\mathbf{1}\{C\subseteq B\}} \sum_{i\in B} h_i \geq \alpha^{\mathbf{1}\{C\subseteq B'\}} \sum_{i\in B'} h_i \text{ for every } B' \in \mathcal{F}\cup\{\varnothing\} \right\},$$

and complementarity widens the cones of core-containing bundles.

3 Canonical Two-Item Examples

This section traces the bundle-to-cone geometry through four canonical two-good examples, each isolating one piece of the comparative-statics arc — crossing the upper threshold α^*, crossing the lower threshold α_*, showing inclusivity's role in keeping α^* finite, and extending to partial bundling — and closes by assembling the three regimes in the iid uniform benchmark (Section 3.5).

All four share $D = [0,1]^2$, the pulled-back cost $c_\alpha(h) = \max\{0,\, v_1(h),\, v_2(h),\, v_{\{1,2\}}(h)\}$, and the source/sink split of μ (top-face mass on $F_i = \{x_i = 1\}$ and the origin atom positive, the Lebesgue interior negative). Each pairs a primitive-space cell diagram — showing the required directions $h = x_{\text{src}} - x_{\text{snk}}$ the bundle cones must cover — with displacement-space cone diagrams at two values of α (color convention

in the figures). Throughout, cost and cones are computed against the full family $\mathcal{P} = \{\{1\}, \{2\}, \{1,2\}\}$ of nonempty bundles; a *candidate menu* M offers a subfamily of these as options, and the question is whether M is certified optimal by Proposition 2.

3.1 Upper threshold: pure bundling in the uniform benchmark

The first example crosses the upper threshold in the simplest setting. The symmetric two-good iid uniform environment with a pure-bundling menu identifies a sufficient threshold α^* above which pure bundling is optimal.

Example 1 (Symmetric two-good uniform benchmark)**.** Let $D = [0,1]^2$, $f \equiv 1$, and $v(x) = (x_1, x_2, \alpha(x_1 + x_2))$ with $\alpha > 0$, where the first two coordinates are singleton values and the third is the value of the grand bundle. The parameter α measures complementarity: $\alpha = 1$ is additive, $\alpha > 1$ complementary, $\alpha < 1$ substitutable. The induced pullback of the one-sided max cost is

$$c_\alpha(h) = \max\{0,\ h_1,\ h_2,\ \alpha(h_1 + h_2)\},$$

and the transformed measure on primitive type space coincides with the iid additive uniform benchmark, $\mu = \delta_{(0,0)} + \sigma_{\{x_1=1\}} + \sigma_{\{x_2=1\}} - 3\,\mathcal{L}^2|_{[0,1]^2}$, so any change in the optimal mechanism arises from geometry alone.

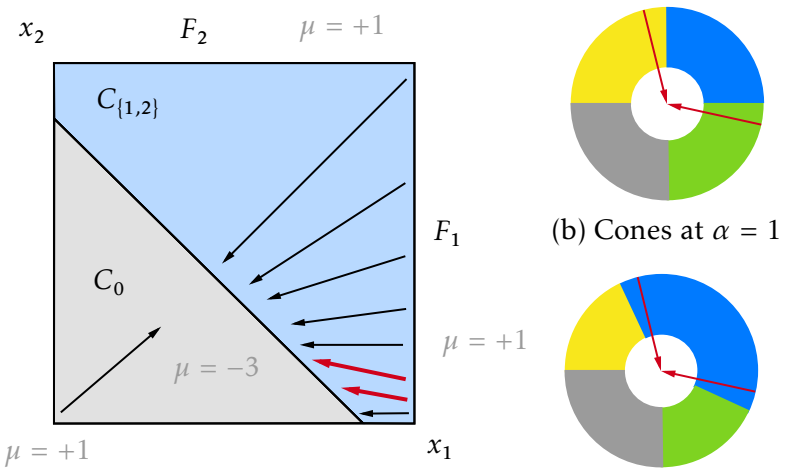
The pure-bundling candidate menu M offers the grand bundle $\{1,2\}$ alone, at price $\alpha\,\hat{s}$ with $\hat{s} = \sqrt{2/3}$. Its linearity cells are the two triangles

$$C_0 = \{x \in D : x_1 + x_2 \le \hat{s}\}, \qquad C_{\{1,2\}} = \{x \in D : x_1 + x_2 \ge \hat{s}\},$$

separated by the anti-diagonal $\{x_1 + x_2 = \hat{s}\}$.

Required transport. Inside C_0 the exclusion option is active, and the origin atom moves along nonnegative directions into the interior. Inside $C_{\{1,2\}}$ the grand bundle is the only active option, and the top-face mass moves along oblique directions into the interior: by symmetry, F_2-mass toward the lower-right (required directions in the second quadrant, $h_1 < 0,\ h_2 > 0$) and F_1-mass toward the upper-left (fourth quadrant, $h_1 > 0,\ h_2 < 0$). Panel (a) of Figure 1 marks in red the required directions the grand bundle's cone fails to cover at $\alpha = 1$.

The cones. At $\alpha = 1$ the bundle's cone is exactly the first quadrant, so the oblique required directions (second and fourth quadrants) fall into the singleton cones, outside the menu's coverage, and Proposition 2 fails. At $\alpha = 2$ the bundle cone has widened into both mixed-sign quadrants — boundary rays $h_2 = -h_1/2$ and $h_1 = -h_2/2$ — and now covers them, so pure bundling is certified (panels (b), (c) of Figure 1).



(a) Cells and transports in D

(b) Cones at $\alpha = 1$

(c) Cones at $\alpha = 2$

Figure 1: Pure bundling in the uniform benchmark. (a) Cells of the pure-bundling candidate, with transport arrows. (b) Direction cones at $\alpha = 1$: required direction outside the bundle's cone. (c) Direction cones at $\alpha = 2$: required direction inside the bundle's cone.

Complementarity widens the cone at rate $1 - 1/\alpha$ in each mixed-sign quadrant; the explicit dual-transport line construction of Section S.5 certifies pure bundling once $\alpha \geq \alpha^*_{\text{unif}} := (9 + 4\sqrt{6})/15 \approx 1.253$ (sufficient, not known to be sharp).

3.2 Lower threshold: necessity of the grand bundle

While the first example crossed the upper threshold, this one identifies the *lower* threshold above which the grand bundle must enter the menu *at all*: in the environment of Example 1, separate sales cease to be certifiable once $\alpha > 1/2$ (Theorem 3).

Consider the separate-sales menu M offering the two singletons $\{1\}, \{2\}$ in the uniform environment. By symmetry the optimal prices satisfy $p_1 = p_2 = p^\star$, and the indirect utility is $u(x) = \max\{0, x_1 - p^\star, x_2 - p^\star\}$. The three linearity cells are

$$C_0 = \{x_1 \leq p^\star,\ x_2 \leq p^\star\}, \quad C_{\{1\}} = \{x_1 \geq p^\star,\ x_1 \geq x_2\}, \quad C_{\{2\}} = \{x_2 \geq p^\star,\ x_2 \geq x_1\},$$

separated along the diagonal segment $E = \{x \in D : x_1 = x_2 \geq p^\star\}$ shown in panel (a) of Figure 2.

Required transport along the diagonal. The cells $C_{\{1\}}$ and $C_{\{2\}}$ each carry top-face mass (on F_1 and F_2) that must be moved into their interiors, and symmetry forces the flow pattern to be symmetric across E — in particular, the net flow across E is zero, so the obstruction is not a mass imbalance between the two selling cells. It is directional. Near the top corner $(1, 1)$ the two faces together carry source mass of order δ in a δ-neighborhood against interior absorption of order δ^2, so most of the corner's source mass must travel an order-one distance; the deep sinks of the selling cells pool around the bottom vertex $(p^\star, p^\star)$ of E. Sources high on the diagonal are therefore paired with sinks near $(p^\star, p^\star)$, so $h = x_{\text{src}} - x_{\text{snk}} \propto (1, 1)$ — a $+45^\circ$ required direction.

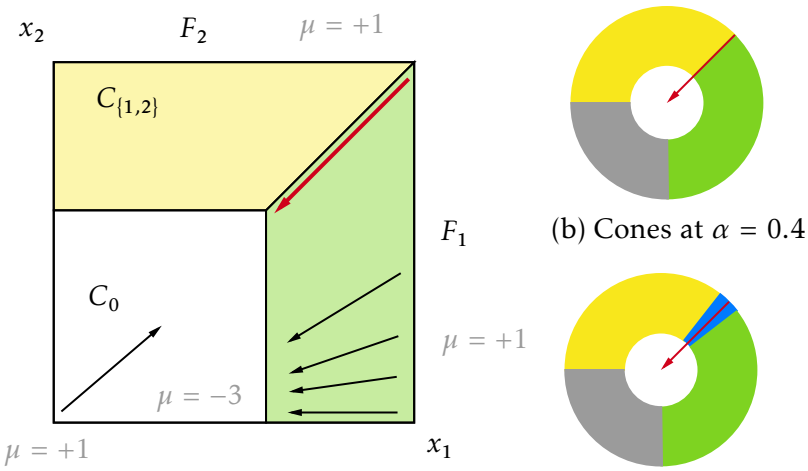
The cones: where does the $+45^\circ$ direction live? The one-sided cost at $h = (1, 1)$ is

$$c_\alpha(1, 1) = \max\{0,\ 1,\ 1,\ 2\alpha\} = \begin{cases} 1, & \alpha \leq 1/2, \\ 2\alpha, & \alpha > 1/2. \end{cases}$$

For $\alpha \leq 1/2$ the singleton gradients $(1, 0, 0)$ and $(0, 1, 0)$ both cover the $+45^\circ$ direction, which sits on the shared boundary of their cones (panel (b), $\alpha = 0.4$: the bundle cone degenerates to the origin). For $\alpha > 1/2$ only the grand-bundle gradient covers it — the $+45^\circ$ direction now lies in the *interior* of the bundle cone and in neither singleton cone (panel (c), $\alpha = 0.6$: the bundle cone is a wedge around $+45^\circ$). Since M offers only the singletons $\{1\}, \{2\}$, carrying only the singleton gradients, no offered bundle's cone covers the required direction and Proposition 2 fails: separate sales cannot be certified.

The threshold $\alpha = 1/2$ is exactly where the grand bundle first becomes strictly preferred by an interior buyer — weak free disposal, the lower-threshold condition of Section 4.2.

Certifying separate sales below $\alpha_* = 1/2$. Below the threshold the grand bundle is pointwise dominated, $\alpha(x_1 + x_2) \leq \max(x_1, x_2)$, so no buyer takes it and the cost degenerates to the unit-demand form $c_\alpha(h) = \max\{0, h_1, h_2\}$. Separate sales at the single-good revenue-maximizing price $p^\star = 1/\sqrt{3}$ is then certifiable. The dual plan is transparent: mass balance holds at $p^\star$ (its defining first-order condition, $\frac{\mathrm{d}}{\mathrm{d}p}\, p(1 - p^2) = 0$); the origin atom feeds the exclusion cell $C_0 = [0, p^\star)^2$ along downward



(a) Cells and transports in D

(b) Cones at $\alpha = 0.4$

(c) Cones at $\alpha = 0.6$

Figure 2: Once $\alpha > 1/2$, the $+45^\circ$ required direction lies strictly inside the bundle cone—neither singleton gradient can certify it. Omitting the grand bundle from the menu is no longer optimal.

directions in the no-purchase cone; and each top face F_i transports into its singleton cell $C_{\{i\}}$ along the cone $\{h : h_i \geq \max(0, h_j)\}$, the requisite Hall condition reducing to $p^\star \leq 2/3$, which holds. So Proposition 2 certifies separate sales. Consistently, the lottery deviation that could otherwise undercut it (a discounted split over the two goods) turns profitable only when the posted price exceeds $2/3$ — above the uniform optimum.

3.3 Without inclusivity, the upper threshold is infinite

The first two examples crossed the thresholds α^*_{unif} and $\alpha_* = 1/2$. This one shows that finiteness of the upper threshold rests on a distributional condition: if the density places enough mass near $(1, 1)$ that the crossing point $\hat{s}$ exceeds the top of each coordinate axis, no finite complementarity certifies pure bundling. The property that prevents this — the exclusion cell sitting strictly below every top face — is inclusivity (Definition 2).

Keep the valuation $v(x) = (x_1, x_2, \alpha(x_1 + x_2))$ and the pure-bundling menu M (the grand bundle alone), but take a density with crossing point $\hat{s} > 1$ — heavier near $(1, 1)$ than the uniform, for which $\hat{s} = \sqrt{2/3} < 1$. The picture below depends only on $\hat{s} > 1$, not on the specific f.

Cells under inclusivity failure. When $\hat{s} > 1$, the anti-diagonal $\{x_1 + x_2 = \hat{s}\}$ no longer connects opposite sides of D; it enters the square through the two top faces, at the points

$$(\hat{s} - 1,\ 1) \in F_2, \qquad (1,\ \hat{s} - 1) \in F_1.$$

The cell $C_{\{1,2\}}$ degenerates to the small upper-right triangle with vertices $(\hat{s} - 1, 1)$, $(1, \hat{s} - 1)$, and $(1, 1)$. The cell C_0 is everything else—crucially, C_0 now contains substantial segments of the top faces:

$$F_1 \cap C_0 = \{1\} \times [0, \hat{s} - 1], \qquad F_2 \cap C_0 = [0, \hat{s} - 1] \times \{1\}.$$

Panel (a) of Figure 3 shows this configuration.

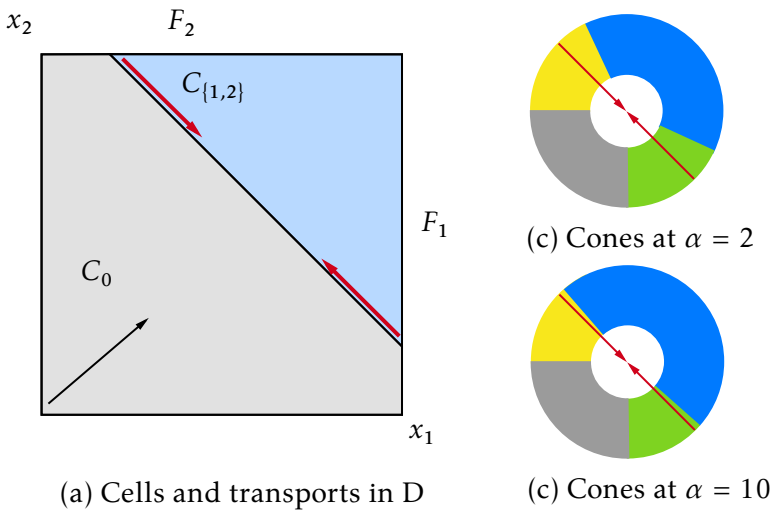
The required direction no cone covers. Focus on the positive top-face mass sitting in $F_i \cap C_0$. Inside C_0 the active option is exclusion, whose cone covers only coordinatewise-nonpositive displacements in v-space. So this mass cannot be consumed within C_0 — any transport into the C_0 interior has a strictly positive singleton component the exclusion option cannot absorb. The mass must instead be moved into $C_{\{1,2\}}$. However, a source at $(1, \hat{s} - 1) \in F_1 \cap \overline{C_0}$ paired with a sink $(u, v) \in C_{\{1,2\}}$ generates a displacement $h = (1-u,\ \hat{s}-1-v)$ with $h_1+h_2 = \hat{s}-(u+v) \leq 0$, with equality only on the anti-diagonal boundary. The required transport points along the negative anti-diagonal — a -45° direction.

Why no α helps. A direct computation gives bundle cone $\subseteq \{h : h_1 + h_2 \geq 0\}$ for every $\alpha > 0$ (equality as $\alpha \to \infty$): a displacement with $h_1 + h_2 < 0$ has $\alpha(h_1 + h_2) < 0$, so the grand-bundle gradient cannot attain the cost. As α grows the bundle cone widens into both mixed-sign quadrants but is always capped by the -45° line (panels (b), (c) of Figure 3, at $\alpha = 2$ and $\alpha = 10$), so the required -45° direction is never covered and α^* runs off to $+\infty$. Economically, inclusivity failure leaves an unexploited singleton margin at near-top types with sub-bundle sum — a positive-surplus pocket that no widening of the bundle cone can reach.

3.4 Upper threshold in partial bundling: the core-peripheral case

The first three examples were symmetric. The last generalizes the upper-threshold construction to a core-peripheral environment, where good 1 carries value on its own while good 2 is valuable mainly in combination. Attach the complementarity multiplier to good 1 and the grand bundle, leaving good 2 at baseline:

$$v(x) = \big(\alpha x_1,\ x_2,\ \alpha(x_1 + x_2)\big), \qquad \alpha > 1.$$



(a) Cells and transports in D

(c) Cones at $\alpha = 2$

(c) Cones at $\alpha = 10$

Figure 3: When inclusivity fails ($\hat{s} > 1$), the required transport directions lie outside the bundle cone for every α: the bundle cone's boundary at $\{h_1 + h_2 \geq 0\}$ is a hard structural limit, not a complementarity issue.

Good 1 is the *core* and good 2 is the *peripheral* good. The candidate menu now offers both the core good and the bundle,

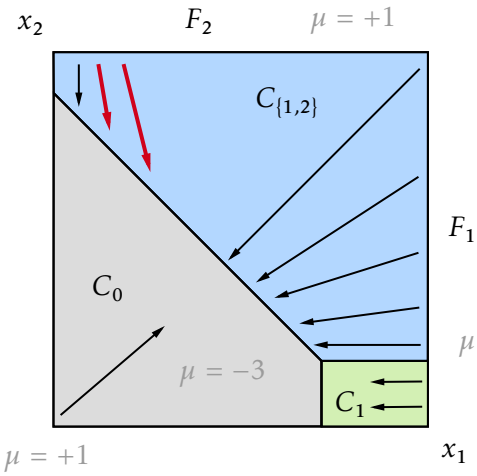
$$M = \{(\{1\},\ \alpha p_1),\ (\{1,2\},\ \alpha p_{\{1,2\}})\},$$

with cost and cones still taken against the full family $\mathcal{P}$. The upper threshold in this partial-bundling case is the smallest α at which M is certified.

Cell structure and required transport. The cell partition and source/sink decomposition parallel the pure-bundling case of Section 3.1, with one asymmetry: the menu now has an active singleton cell $C_{\{1\}} = \{x_1 \geq p_1,\ x_2 \leq p_{\{1,2\}} - p_1\}$ in addition to exclusion C_0 and bundle $C_{\{1,2\}}$, with axis-aligned boundaries (panel (a) of Figure 4). F_1-mass is split between $C_{\{1\}}$ (lower segment, fourth-quadrant displacements) and $C_{\{1,2\}}$ (upper segment, first-quadrant displacements — in the bundle cell the cone requires $h_2 \geq 0$, since the core good and the bundle carry the same multiplier and cancel); F_2-mass routes entirely to $C_{\{1,2\}}$ along second-quadrant displacements — the obstruction at $\alpha = 1$.

The cones. The asymmetry between core and peripheral shows up cleanly in the cones (panels (b) and (c) of Figure 4). The good-1 cone is the entire fourth quadrant $\{h_1 \geq 0,\ h_2 \leq 0\}$ *for every* α: the multiplier scales the core good and the bundle

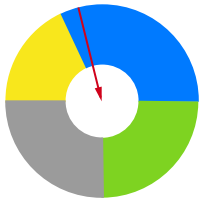
symmetrically and cancels in their comparison. Complementarity instead widens only the bundle cone, into the second quadrant at the expense of the good-2 cone, with boundary ray $h_1 = -\frac{\alpha-1}{\alpha} h_2$. At $\alpha = 1$ the bundle cone is the first quadrant and no menu-bundle's cone covers the required F_2-to-interior directions; once α is large enough that the widened bundle cone reaches them, the coverage obstruction disappears; Theorem 2 turns this picture into a finite-threshold guarantee under its hypotheses — mass balance, inclusivity, and alignment regularity (Definition 3). As in Section 3.1, the bundle cone widens at rate $1 - 1/\alpha$, but one-sidedly: the core good's cone absorbs all fourth-quadrant transport for free, so only the peripheral direction is sensitive to α.



(a) Cells and transports in D



(b) Cones at $\alpha = 1$



(c) Cones at $\alpha = 2$

Figure 4: Core-peripheral bundling with asymmetric valuations $v = (\alpha x_1, x_2, \alpha(x_1 + x_2))$ and menu M offering $\{1\}, \{1,2\}$. The bundle cone is sensitive to α only in the peripheral direction; the core-good cone is invariant.

3.5 Three regimes in the 2-good iid uniform benchmark

The first two examples bracket three regimes in the symmetric two-good iid uniform environment. Below $\alpha_* = 1/2$ separate sales is certified (Section 3.2): the grand bundle is pointwise dominated, the cost is unit-demand, and the singleton cones carry the transport at price $p^\star = 1/\sqrt{3}$. The lower endpoint $\alpha_* = 1/2$ is the level at which weak free disposal of the grand bundle first holds. Above the explicit dual-transport threshold $\alpha^*_{\text{unif}} \approx 1.253$, the line construction of Section S.5 certifies pure bundling (Section 3.1). In the intermediate range the optimum is

characterized here only at the additive point: at $\alpha = 1$ it is the mixed-bundling menu of Pavlov (2011) — the grand bundle priced against two singletons; away from $\alpha = 1$ no closed-form sufficient condition covers the range.

The 2-good uniform benchmark is illustrative; Theorems 2 and 3 carry the framework to general N, with a closed-form iid Beta benchmark in Section 5.1.

4 Core-Peripheral Bundling

This section turns the recurrent menu pattern described in the examples into a theorem. The argument is organized around *two complementarity thresholds* attached to the same core C and type distribution f. An upper threshold α^* is the level above which the optimal mechanism collapses onto core-containing bundles — *core-peripheral bundling* in the language of the introduction; Section 4.1 states the sufficiency theorem (Theorem 2) that produces it. A lower threshold α_* is the level above which the grand bundle is forced into the menu; Section 4.2 states the matching necessity theorem (Theorem 3). Two further subsections support the pair: Section 4.3 establishes that inclusivity is the binding distributional condition for finiteness of α^*, and Section 4.4 explains why the core-peripheral family is the natural carrier of the premium. The applications of the two theorems — pure bundling characterized, the iid Beta benchmark, and an explicit two-tier family — are developed in Section 5.

Throughout this section I take the *standing form*. The type space is the box $D = \prod_{i\in[N]}[0, \bar{v}_i]$; the density f is C^2 and strictly positive on D, with radial score $\Phi(x) := x \cdot \nabla \log f(x) + N + 1$ strictly positive; valuations take the core-peripheral form (4) for a fixed core C, $P := [N] \setminus C$; and the candidate is a *core-peripheral mechanism* — a menu of lotteries over the core-containing bundles $\mathcal{F}_C := \{B \in \mathcal{P} : C \subseteq B\}$, so each option σ_ℓ has marginal-inclusion vector $w_\ell := \mathbb{E}_{B\sim\sigma_\ell}[\mathbf{1}_B]$ and a strictly positive price. Ties are broken seller-favorably (Section 2.1). The differential virtual value of Section 2.2 splits into a boundary part on the top faces $F_j := \{x_j = \bar{v}_j\}$, where granting surplus rewards the seller, and an interior part where it is absorbed:

$$\mu = \mu^+ - \mu^-, \qquad \mu^+ = \delta_0 + \sum_j \bar{v}_j\, f|_{F_j}, \qquad \mu^- = \Phi f \cdot \lambda_D$$

4.1 Upper threshold: sufficiency of core-peripheral bundling

When is a core-peripheral mechanism optimal not merely among menus on $\mathcal{F}_C$ but against *every* mechanism — so that every active offer sells a bundle contain-

ing the entire core, or else excludes the buyer? Theorem 2 gives conditions on the density f and the candidate menu under which the answer is affirmative once complementarity exceeds a finite *upper threshold* α^*.

Beyond restricted optimality of the menu, the sufficiency theorem rests on two further conditions, *inclusivity* (on the menu's prices) and *alignment regularity* (on the menu–density pair), introduced in turn. Both are stated on the selling cells: option ℓ wins on a cell $C_\ell \subseteq D$, and prices scale with the premium, $p_\ell = \alpha \hat{s}_\ell$, so the *crossing points* $\hat{s}_\ell := p_\ell/\alpha$ and the cells $\{C_\ell\}_{\ell \geq 0}$ are α-independent (Section B).

Definition 2 (Inclusivity)**.** *A candidate menu M on $\mathcal{F}_C$, with exclusion cell C_0 (the no-purchase region), is* inclusive *if no top face meets the exclusion cell: $F_j \cap C_0 = \varnothing$ for every $j \in [N]$.*

Inclusivity is the key economic condition: the menu's crossing points are low enough to *include* every buyer who is near-maximal on a single coordinate. Were some F_j absorbed by C_0, a positive mass of buyers with x_j near $\bar{v}_j$ would buy nothing despite valuing item j near its maximum, and an item-j offer priced just below $\bar{v}_j$ would break in. Theorem 4 below makes the rationale formal: leaving a top face excluded is strictly suboptimal at every $\alpha > 0$. For pure bundling, inclusivity reduces to the single price inequality $\hat{s} \leq \min_i \bar{v}_i$.

The second condition generalizes Myersonian regularity — in one dimension, the statement that the revenue-rewarding part of the differential virtual value dominates the revenue-absorbing part in stochastic order. With several goods the same comparison is made, but only along the directions the cell's option can push surplus, and the option's allocation w_ℓ names them. It sells the full-weight peripherals $S_\ell^{(1)} := \{j \in P : w_{\ell,j} = 1\}$, so it covers *upward* moves in them; it omits the zero-weight items $Z_\ell := \{j \in P : w_{\ell,j} = 0\}$, so it covers *downward* moves; and it covers neither way on the fractional items $B_\ell^{\mathrm{frac}} := \{j \in P : 0 < w_{\ell,j} < 1\}$, which are frozen. The core, sold as a bundle, enters through its total, fanned by a margin $\kappa \in [0, 1)$ into the κ-tilted core sums: for each $T \subseteq C$,

$$\sigma_T^\kappa(x) := \sum_{i \in T} x_i + \kappa \sum_{i \in C \setminus T} x_i + \kappa \sum_{j \in S_\ell^{(1)}} x_j,$$

collected into vector $\sigma^\kappa(x) := \big(\sigma_T^\kappa(x)\big)_{\varnothing \neq T \subseteq C}$: each T credits a subset of the core at full value and the rest at the discounted rate κ, so the family asks regularity to

survive a core dip of up to a κ-fraction. Writing $\mu_\ell^\pm := \mu^\pm|_{C_\ell}$ for the rewarding and absorbing masses on the cell, *alignment regularity* asks for this Myersonian dominance in exactly these directions: *regular* as in Myerson, *aligned* to what the menu covers.

Definition 3 (Alignment regularity)**.** *For a level $\kappa \in [0,1)$ and a selling cell ℓ, write $x|_S$ for the restriction of a type $x \in \mathbb{R}^N$ to a coordinate block S. The* alignment score *is*

$$\Theta_\ell^\kappa(x) := \big(\, \sigma^\kappa(x),\ x|_{S_\ell^{(1)}},\ -x|_{Z_\ell},\ x|_{B_\ell^{\mathrm{frac}}},\ -x|_{B_\ell^{\mathrm{frac}}} \,\big),$$

and $\nu_{\ell,\kappa}^\pm$ is the law of $\mu_\ell^\pm$ under Θ_ℓ^κ.[13] *The pair (f, M) is* alignment-regular *at level $\kappa \in [0,1)$ if*

$$\nu_{\ell,\kappa}^- \preceq_{\mathrm{FOSD}} \nu_{\ell,\kappa}^+ \text{[14]}$$

at every selling cell ℓ, and alignment-regular *if this holds at some level κ.*

The condition is joint in (f, M) — a marginal-revenue ordering checked at the candidate menu, not a property of f alone — and is sufficient, not necessary, for the theorem below. At $C = [N]$ — pure bundling — it reduces to Myersonian regularity of the bundle value (Section 5.1). Alignment regularity is the framework's second layer of regularity: the no-ironing condition of Remark 1 rules out ironing at the level of the marginal ν, alignment regularity rules it out within each cell, and in one dimension the two collapse into regularity of Myerson's virtual valuation.

Theorem 2 (Optimality of core-peripheral bundling)**.** *Fix a core $C \subseteq [N]$ and let $M = \{(\sigma_\ell, p_\ell)\}_{\ell=1}^L$ be a finite menu that is* optimal in the restricted problem on $\mathcal{F}_C$. *If M is* inclusive *(Definition 2) and (f, M) is* alignment-regular *(Definition 3), then there exists a finite threshold $\alpha^*(C, M) < \infty$ such that for every $\alpha \geq \alpha^*(C, M)$, M is optimal for the unrestricted problem on $\mathcal{P}$.*[15]

Intuitively, the three hypotheses divide the labor. Restricted optimality fixes each selling cell's price at its revenue-maximizing level, so no rigid price change

[13] That is, $\nu_{\ell,\kappa}^\pm(A) := \mu_\ell^\pm\{x : \Theta_\ell^\kappa(x) \in A\}$. The score lists each fractional coordinate twice, once with each sign, so the dominance below binds in both directions on the B_ℓ^{frac} block, forcing the B_ℓ^{frac}-marginals of $\nu_{\ell,\kappa}^+$ and $\nu_{\ell,\kappa}^-$ to coincide; this is what freezes those coordinates ($h_j = 0$) in the coupling of Lemma 6.

[14] That is, $\int \varphi \, \mathrm{d}\nu_{\ell,\kappa}^- \leq \int \varphi \, \mathrm{d}\nu_{\ell,\kappa}^+$ for every bounded coordinatewise nondecreasing φ.

[15] Prices scale with α: M is the cell structure with α-independent crossing points $\hat{s}_\ell = p_\ell/\alpha$, posting prices $\alpha\hat{s}_\ell$ at complementarity α (Section B). The assertion is that this α-indexed family — fixed cells $\{C_\ell\}$ and crossing points $\hat{s}_\ell$ — is unrestricted-optimal once $\alpha \geq \alpha^*$.

gains revenue. Inclusivity keeps every near-top buyer served, so no singleton offer can profitably break in. Alignment regularity is the ordering the proof consumes: within every cell, the types the seller rewards with surplus — the top-face mass — sit above the types from which surplus is withheld — the interior mass. Complementarity then does the rest, as the proof sketch shows.

Proof sketch. The certificate is a dual transport plan (Proposition 2): in each cell it carries revenue-rewarding boundary mass onto revenue-absorbing interior mass along displacements the cell's option *covers*. Optimality is therefore a coverage question — do the displacements the distribution forces lie in the option's covered cone? — and complementarity is the lever, since it widens that cone. The three steps below fix the covered displacements, measure the shrinkage from the core-only cone to the full cone, and supply the distributional work that closes it with margin to spare; the construction is in Section B.

Step 1 — The covered cone fixes the displacement signs. On each selling cell, cone-forcing (Lemma 4) pins the sign pattern of every covered displacement from the cell's geometry alone, with no appeal to the distribution; inclusivity keeps the boundary mass off the exclusion cell.

Step 2 — Lifting to the full menu shrinks each cell's cone. The criterion requires displacements in the *full cone* $K_\ell^{\mathcal{P}} \subseteq K_\ell^{\mathcal{F}_C}$. Because the certificate is supported on within-cell pairs by construction (Lemma 5(ii)), only this within-cell shrinkage must be absorbed.

Step 3 — Alignment regularity supplies the coupling; complementarity absorbs the shrinkage. The scores are calibrated to the covered cone: dominance of $\nu^+_{\ell,\kappa}$ over $\nu^-_{\ell,\kappa}$ transfers, through a monotone-coupling form of Strassen's theorem, into an α-independent within-cell coupling of μ^+_ℓ to μ^-_ℓ whose displacements respect the cell's sign pattern and dip into the core by at most a κ-fraction of their aligned gains (Lemma 6). Such displacements lie in the covered cone with margin, and the level controls the worst non-core deviation: once $\alpha \geq 1/(1-\kappa)$, every coupled displacement lies in the full cone. Assembling the cell couplings — with the exclusion cell handled by inclusivity — yields the certificate, with $\alpha^*(C, M) \leq 1/(1-\kappa)$.

4.2 Lower threshold: necessity of the grand bundle

Theorem 2 identified the upper threshold α^*. The complementary question is when complementarity is high enough to *force* the grand bundle into the menu.

Throughout, $M_{\mathcal{F}} = \{(\sigma_\ell, p^*_\ell)\}_{\ell=1}^{L^*}$ is a candidate mechanism on a family $\mathcal{F} \not\ni [N]$, with highest posted price $\bar{p} := \max_\ell p^*_\ell$. Extend the bundle-to-cone construction (Lemma 1) to the grand bundle by its *strict-preference set*

$$U_{[N]}^{M_{\mathcal{F}}} := \left\{h \in \mathbb{R}^N : v_{[N]}(h) > \bar{v}_\ell(h) \text{ for every cell } \ell \geq 1, \text{ and } v_{[N]}(h) > 0\right\}, \qquad v_{[N]}(h) = \alpha \textstyle\sum_i h_i,$$

the open cone of directions along which the grand bundle strictly beats every menu option and the outside option ($\bar{v}_\ell(h) = \mathbb{E}_{B\sim\sigma_\ell}[v_B(h)]$ is option ℓ's expected valuation, linear in h by (4)); it widens as α scales the core-containing gradients. The necessity condition asks that a positive mass of buyers in this cone sit at the seller's most profitable offer.

Definition 4 (Top-cell wedge)**.** *Let ℓ^* be a top-priced cell, $p^*_{\ell^*} = \bar{p}$. The mechanism $M_{\mathcal{F}}$ satisfies the* top-cell wedge *condition if*

$$\lambda_D(V^*) > 0, \qquad \textit{where } V^* := C_{\ell^*} \cap \operatorname{int} D \cap U_{[N]}^{M_{\mathcal{F}}} \cap \mathbb{R}^N_{>0}.$$

Read this in two steps. That $U_{[N]}^{M_{\mathcal{F}}} \cap \mathbb{R}^N_{>0}$ is nonempty is a *weak free disposal* condition — some interior buyer strictly prefers the grand bundle to everything on offer; it is weaker than pointwise free disposal (which in the core-peripheral form needs $\alpha \geq 1$) and already holds at $\alpha > 1/2$ in the two-good case of Section 3.2. The wedge strengthens this by placing such buyers inside the top-priced cell C_{ℓ^*} — the segment from which an upward nudge extracts the most revenue and, the anchor being top-priced, draws no cannibalization.

Theorem 3 (Necessity of the grand bundle)**.** *Suppose the restricted family $\mathcal{F}$ excludes $[N]$ and $M_{\mathcal{F}}$ is* any *candidate mechanism on $\mathcal{F}$ with a finite menu satisfying the top-cell wedge condition (Definition 4). Then $M_{\mathcal{F}}$ is not optimal in the unrestricted problem.*

I write $\alpha_*(C, \mathcal{F})$ for the infimum of the set of $\alpha > 0$ at which the top-cell wedge holds for every restricted-optimal mechanism on $\mathcal{F}$; Theorem 3 forces the grand bundle at every α in that set (whether the set is an interval is not needed). In the two-good iid uniform case with separate sales, the wedge at the restricted optimum opens exactly as weak free disposal does — types near the top corner strictly prefer the grand bundle once $\alpha > 1/2$ (Section 3.2) — so $\alpha_* = 1/2$.

Proof sketch. The necessity argument constructs a profitable deviation. The strict-preference set $U^{M_{\mathcal{F}}}_{[N]}$ gives directions on which an offered $[N]$ strictly dominates every alternative in $\mathcal{F}$ — surplus the existing menu leaves unextracted. The top-cell wedge (Definition 4) localizes this gap to a positive-measure subset $V^* \subseteq C_{\ell^*}$ of the top-priced cell, and a Lebesgue-density argument extracts a small ball $W \subseteq V^*$ with uniform preference and margin. Augmenting $M_{\mathcal{F}}$ with the grand bundle $([N], p^*_{\ell^*} + \delta)$: for δ below the margin every buyer in W switches to it, contributing bulk revenue $\Theta(\delta)$; cannibalization from other cells is non-negative because the top-priced anchor means no other cell carries a higher price. The bulk gain is strictly positive and the cannibalization non-negative, so the deviation is profitable. The full proof is in Section S.3. *Q.E.D.*

4.3 Finiteness of the upper threshold: necessity of inclusivity

Theorem 2 took inclusivity as a hypothesis; this subsection shows the same condition — no top face left in the exclusion cell — is also necessary, and is the binding distributional condition for finiteness of α^*. If a candidate menu excludes some top face, a singleton offer for that item profitably breaks in at every $\alpha > 0$. The pure-bundling necessity result is an immediate corollary.

Definition 5 (Excluded top face)**.** *Let $M_{\mathcal{F}}$ be a candidate mechanism on a restricted family $\mathcal{F} \subseteq \mathcal{P}$, with exclusion cell C_0 (the no-purchase region). The top face $F_i = \{x \in D : x_i = \bar{v}_i\}$ is* excluded *by $M_{\mathcal{F}}$ if a positive $(N-1)$-dimensional measure of it lies in the exclusion cell,*

$$\lambda_{N-1}(F_i \cap C_0) > 0.$$

An excluded top face is a positive-measure failure of inclusivity (Definition 2): a near-top-i population is pushed to no purchase despite valuing item i near its maximum. Under the seller-favorable tie convention the exclusion cell is relatively open, so excluded top faces and inclusivity failures coincide. For pure bundling, $C_0 = \{\sum_j x_j < \hat{s}_1\}$, so F_i is excluded exactly when $\bar{v}_i < \hat{s}_1$.

Theorem 4 (Necessity of inclusivity)**.** *Let $N \geq 2$, let $M_{\mathcal{F}}$ be any candidate mechanism on a restricted family $\mathcal{F} \subseteq \mathcal{P}$ with a finite menu, and let $f \in C^2(D)$ be strictly positive on D. If some top face F_i is excluded (Definition 5), then $M_{\mathcal{F}}$ is not optimal in the unrestricted problem, at any $\alpha > 0$.*

Proof sketch. Fix $\alpha > 0$ and price item i just below its standalone value $V_i := \alpha^{\mathbf{1}\{C\subseteq\{i\}\}}\bar{v}_i = \max_x v_{\{i\}}(x)$: add the singleton $(\{i\}, r_\varepsilon)$ with $r_\varepsilon := V_i - \varepsilon$. A buyer switches iff $v_{\{i\}}(x) - r_\varepsilon > u_{M_\mathcal{F}}(x)$, i.e. $\alpha^{\mathbf{1}\{C\subseteq\{i\}\}}x_i > V_i - \varepsilon$, so every switcher has x_i within $\varepsilon' := \varepsilon/\alpha^{\mathbf{1}\{C\subseteq\{i\}\}}$ of the top — a slab of width ε'. Over the positive-measure base $F_i \cap C_0$ the near-top types are excluded ($u_{M_\mathcal{F}} = 0$), so each switches and pays $r_\varepsilon \approx V_i$ where it paid nothing before: the gain is $\Theta(V_i\varepsilon') = \Theta(\varepsilon)$. The only offsetting loss is from former buyers of some option ℓ who switch; such a type lies within ε' of the top in x_i *and* within ε of indifference for option ℓ, i.e. $\bar{v}_\ell(x) \in [p_\ell, p_\ell + \varepsilon]$ — two transverse constraints cutting an $O(\varepsilon^2)$ slab, unless ℓ is the deterministic singleton $\{i\}$, which an excluded F_i forces to price V_i (so it sells on a null set, no loss). So the gain $\Theta(\varepsilon)$ dominates the loss $O(\varepsilon^2)$ and the deviation is profitable. Section S.4 gives the bookkeeping. *Q.E.D.*

4.4 From geometry to algebra: why core-peripheral?

Theorem 2 certifies the core-peripheral menu — supported on $\mathcal{F}_C = \{B : C \subseteq B\}$ — as optimal above α^*. Why this specific structure? The framework reveals an asymmetry across bundles, organized by size: *large bundles are hard to exclude from a certifiable menu even when their cones are narrow; small bundles are easy to exclude even when their cones are wide.* The grand bundle $[N]$ is the extreme case: by Theorem 3, once complementarity makes it a needed cover — the top-cell wedge — it cannot be dropped, its protection structural rather than a matter of cone width. (Inclusivity, Theorem 4, is a separate force: it requires every item's top face to be served but does not by itself pull in the grand bundle, since separate sales can leave every top face served.) A singleton $\{i\}$, in contrast, has a wide cone but is replaced for free whenever some larger premium bundle covers item i: cone-wideness is geometric slack the menu does not need.

Algebraic mirror: closure under supersets. The asymmetry has a direct algebraic counterpart. The proof of Theorem 2 scales the gradients of the core-containing — *premium* — bundles by α; within-premium cones are α-invariant (premium gradients scale together) while between-premium cones widen. This invariance requires the premium set to be *closed under supersets* — if B is premium, every $B' \supseteq B$ must be too, or a non-premium superset breaks the cancellation. The natural such family is $\mathcal{F}_C$, and every non-empty superset-closed family must contain $[N]$: the set-theoretic universality of the grand bundle is exactly its algebraic universality.

Open questions. Two natural extensions preserve the asymmetric structure but enlarge the premium family. *Coexisting complementarity systems:* multiple disjoint cores, each with its own premium (Microsoft Office and Azure as parallel ecosystems). *Competing vertical hierarchies:* multiple incomparable maximal chains (cable, satellite, streaming as parallel ladders), each admitting a core-peripheral form. Both are conjectured core-peripheral within their own systems and reduce to Theorem 2 in the single-system case, but require multi-parameter analogs of alignment regularity.

5 Applications

This section applies Theorem 2 to settings where its hypotheses can be discharged from primitives, and closes with the markets the theory organizes. The standing core-peripheral form of Section 4 — (4), a box domain, and an admissible density with nonnegative radial score — remains in force throughout. Pure bundling needs no distributional assumption beyond Myersonian regularity and inclusivity, with the iid Beta family making its conditions explicit in every dimension (Section 5.1); an explicit two-tier family with one peripheral good is solved under log-concave core marginals (Section 5.2); and the two thresholds map observed menus across market segments (Section 5.3).

5.1 Pure bundling

At $C = [N]$ the menu is a single grand-bundle price $p^* = \alpha\hat{s}$ — the restricted optimum on $\mathcal{F}_{[N]} = \{[N]\}$, with α-independent crossing point $\hat{s}$. Call the bundle-value distribution of $\sum_i x_i$ *strictly Myersonian regular* if its virtual value is strictly positive above the crossing point, $H(r) := r\,\gamma(r)/\bar{F}(r) > 1$ on the active range ($\gamma, \bar{F}$ the bundle value's density and survivor); classical regularity together with the price first-order condition implies it.

Theorem 5 (Pure bundling)**.** *Suppose the bundle-value distribution is strictly Myersonian regular. Then pure bundling at p^* is optimal in the unrestricted problem for all sufficiently large α if f is inclusive ($\hat{s} < \min_i \bar{v}_i$), and is strictly suboptimal at every $\alpha \geq 1$ if f is non-inclusive ($\min_i \bar{v}_i < \hat{s}$).*

The two directions give inclusivity as a necessary and sufficient condition for pure bundling, with the knife edge $\hat{s} = \min_i \bar{v}_i$ unclassified. *Sufficiency* is the $C = [N]$ case of Theorem 2: single-cell mass balance is the price first-order condition

(Lemma 3), and strict Myersonian regularity supplies alignment regularity through the single-cell certificate (Proposition 4); conversely that alignment regularity forces the weak form $H \geq 1$ (Remark 3), so the hypothesis is pinched between the weak and strict classical conditions. *Necessity* is the pure-bundling case of Theorem 4, strengthening the additive observation of Manelli and Vincent (2007); Pavlov (2011): the force is geometric, not pricing. Complementarity widens the bundle cone but never moves the crossing point, so a buyer with $\bar{v}_i < \hat{s}$ pays nothing at every α and an item-i singleton captures the leakage.

The iid Beta family. The conditions of Theorem 5 are satisfiable in every dimension, with closed forms. For the iid Beta$(\beta, 1)$ family $f_\beta(x) = \beta^N \prod_i x_i^{\beta-1}$ on $[0,1]^N$, the bundle value $\sum_i x_i$ has the explicit law $G(t) = t^{N\beta}/c_N(\beta)$ on $[0,1]$, with $c_N(\beta) = \Gamma(1+N\beta)/\Gamma(1+\beta)^N$. Writing $A_N(\beta) := (1+N\beta)/c_N(\beta)$, the interior first-order condition for the pure-bundling price yields the candidate crossing point

$$\hat{s}_{N,\beta} = A_N(\beta)^{-1/(N\beta)},$$

which lies strictly below 1 exactly when $A_N(\beta) > 1$; the strict Myersonian regularity below places the global revenue maximizer on $[0,1]$, so this root is the crossing point and the menu is then inclusive. A small exponent secures inclusivity, and more: it pushes mass toward zero on each coordinate (the *bottom-heavy* regime that inclusivity asks for), and for β small enough the bundle-value distribution is strictly Myersonian regular in *every* dimension, so pure bundling is optimal for all large α. The dimension-adapted choice $\beta = 1/N$ keeps the per-coordinate tail mass scale-invariant and gives a benchmark whose crossing point $\hat{s}_{N,1/N} = 1/\big(2\,\Gamma(1+1/N)^N\big)$ rises monotonically toward $e^{\gamma_E}/2 \approx 0.89$. The formal analysis is deferred to a supplementary note.[16]

5.2 Two-tier menus with one peripheral good

The second application is the simplest menu in which the core-peripheral structure is genuinely visible: a *two-tier menu* offering the bare core C and the grand bundle $C \cup \{p\}$ adding a single peripheral good. Unlike pure bundling, the certificate is built in two cells at once and the peripheral good forces transport *across* peripheral levels. Normalize to $D = [0,1]^N$; the core is C with $|C| = N - 1 \geq 1$, $P = \{p\}$, types

[16] The two-good uniform benchmark is the $N = 2$, $\beta = 1$ member; its explicit threshold $\alpha^*_{\text{unif}} \approx 1.253$ is constructed in Section S.5.

$x = (\omega, z)$ with $z := x_p$ and core sum $t(\omega) := \sum_{i \in C} \omega_i$, and $H_C(q) := q\, g_C(q)/\bar{G}_C(q)$ is the core sum's virtual-value statistic under f_C (Section S.1).

The *two-tier setting* fixes three things. (S1) The density is a product $f = f_C \otimes f_p$ with $f_C = \bigotimes_{i \in C} f_i$, each f_i log-concave, C^1, and positive on $(0, 1]$ — so f_C is an admissible core density (Definition 7) — and the core radial score $\Phi_C(\omega) := |C| + \omega \cdot \nabla \log f_C(\omega)$ is nonnegative. (S2) The peripheral marginal is $f_p(z) = \beta z^{\beta-1}$ for some $\beta > 0$. (S3) The menu $M = \{(C, \alpha \hat{s}_C),\ (C \cup \{p\}, \alpha \hat{s}_G)\}$ has crossing points $0 < \hat{s}_C < \hat{s}_G < 1$ and mass-balanced selling cells; write $d := \hat{s}_G - \hat{s}_C$ and $\tau(z) := (\hat{s}_G - z)^+$.[17]

Normalizing utilities by α, the bare-core tier wins on $C_1 = \{t \geq \hat{s}_C,\ z \leq d\}$ and the grand tier on $C_2 = \{z \geq d,\ t \geq \tau(z)\}$: the upgrade threshold depends on z alone (buy the add-on when $z \geq d$), while the participation boundary $t = \tau(z)$ falls one-for-one in z to zero at $z = \hat{s}_G$. *Inclusivity is free* from the (S3) crossing points: every core top face has $t \geq 1 > \hat{s}_C \geq \tau(z)$ and the peripheral face $\{z = 1\}$ sits in C_2 since $\hat{s}_G < 1$ — both read directly off the geometric (face-meets-cell) form of Definition 2. And mass balance on the bare-core cell is equivalent to

$$H_C(\hat{s}_C) \;=\; 1 + \beta$$

(Lemma 17), the two-tier analogue of the pure-bundling first-order condition $H(\hat{s}) = 1$ of Proposition 4; the constant rises from 1 to $1 + \beta$ because the cell's interior mass now carries the peripheral block's share of the radial score, $\Phi = \Phi_C + 1 + \beta$.

Theorem 6 (Optimality of the two-tier menu)**.** *Under the two-tier setting* (S1)–(S3)*, the pair (f, M) is alignment-regular at some level $\kappa^* < 1$; hence the two-tier menu is optimal in the unrestricted problem for all sufficiently large α.*

Intuitively, the certificate splits each cell's transport in two. *Stage W* works at frozen peripheral coordinate: in each slice the core faces are carried onto the slice's interior share by the single-cell couplings (Theorem 8) at entry level $q = \hat{s}_C$ (bare-core cell) or $q = \tau(z)$ (grand cell) — pure bundling, slice by slice — at a uniform level $\kappa_W < 1$ (Proposition 3). *Stage G* serves the peripheral face by sliding down the peripheral coordinate, the core sum dipping by at most a κ-fraction per unit of descent — the rate the grand cell's budget relation permits; the binding slice $z = d$ fixes the closed form $\hat{\kappa} = \hat{s}_C/(1 - d)$, which rises to 1 as the grand slack $1 - \hat{s}_G$ vanishes. The proof and supporting remarks are in Section S.2.2.

[17] For $\beta < 1$, f falls outside the closed-box standing form but is an admissible core density to which the sufficiency chain extends (Lemma 14).

5.3 Three regimes in observed markets

The two thresholds traced in Sections 4.1 and 4.2 produce a comparative-statics arc — separate sales below α_*, a mixed-bundling intermediate regime, core-peripheral bundling above α^* — observable where the same goods are sold to segments differing in their integration premium, the complementarity premium α of (4).

AI providers maintain two parallel price lists for the same models. The consumer segment — individuals using a chat interface for everyday tasks at a flat monthly fee — receives a tiered, bundled mechanism: the ChatGPT Plus / Pro ladder. The developer segment — engineers building applications, agents, and back-end tools that call the models programmatically — receives an unbundled per-token price list, with each model priced separately by usage.[18] The two channels target observably distinct segments: a consumer's value of having both models inside one subscription comes from the unified conversational interface, the shared memory, and the multi-model workflow the provider performs automatically, so the integration premium is large; a developer's routing logic, context management, and model substitution live in the developer's own code, so the integration premium is small.

Enterprise software shows the same segmentation at industry scale. Large traditional firms — banks, retailers, manufacturers, with limited internal engineering capacity — typically license integrated SAP, Oracle, or Salesforce suites at the bundled list price. Tech firms — Meta, Stripe, Netflix, with deep internal engineering capacity — typically forgo the integrated suite and assemble a custom stack out of best-of-breed components. Between the two extremes, mid-sized enterprises license a core platform together with selected optional modules: an integrated suite for the workflows tightest with the platform's data, separately priced add-ons for the rest.

The three segments map naturally onto the three regimes of the comparative-statics arc. Per-token API access and best-of-breed assembly sit below α_*, where the integration premium is too small to force the grand bundle into the menu — the bundle's cone covers no required direction that some Ã la carte option cannot already cover. Mid-sized enterprise licensing sits in the intermediate (α_*, α^*) band, where the grand bundle covers some required directions but core-only cones miss

[18] Compare https://openai.com/chatgpt/pricing with https://openai.com/api/pricing/; Anthropic's consumer and API channels show the same split.

others, so both the bundle and Ã la carte options remain in the menu. Consumer AI tiers and integrated enterprise suites sit above α^*, where the core-containing cones jointly cover every required direction and Ã la carte alternatives are squeezed out. The same product, faced with three different integration premia, supports three qualitatively different mechanisms.

6 Conclusion

This paper develops a bundle-to-cone framework for optimal screening under combinatorial preferences: each bundle carries a cone of covered directions, the differential virtual value identifies the required directions, and optimality is certified by matching required to covered, cell by cell, with complementarity widening the cones of core-containing bundles. Applied to a scalar complementarity premium α multiplying every bundle that contains a fixed core, the framework identifies two thresholds, each a cover-side event in α: a lower threshold α_* at which the grand bundle's cone first becomes a needed cover (Theorem 3), and an upper threshold α^* at which the cones of core-containing bundles jointly cover every required direction (Theorem 2), certified by alignment regularity — the stochastic-dominance condition that orders revenue-rewarding above revenue-absorbing types in every alignment score. Inclusivity is the binding distributional condition for finiteness of α^* (Theorem 4). The applications discharge these hypotheses from primitives. For pure bundling the picture is sharpest: alignment regularity is supplied by standard Myersonian regularity, and inclusivity characterizes optimality outright — up to a knife edge, the distribution is inclusive if and only if pure bundling is optimal for all large α (Theorem 5). The two-tier family (Section 5.2) and the iid Beta benchmark (Section 5.1) extend the picture to one add-on and to higher dimensions.

The methodology is portable to richer slices of the combinatorial preference space; Section 4.4 sketches the natural extensions — coexisting complementarity systems and competing vertical hierarchies — as open questions on the algebraic structure of premium bundles.

References

Armstrong, M. (1996): "Multiproduct Nonlinear Pricing," *Econometrica*, 64, 51–75.

——— (1999): "Price Discrimination by a Many-Product Firm," *The Review of Economic Studies*, 66, 151–168.

Bagnoli, M. and T. Bergstrom (2005): "Log-concave probability and its applications," *Economic Theory*, 26, 445–469.

Bikhchandani, S. and D. Mishra (2024): "Rank-preserving multidimensional mechanisms: An equivalence between identical-object and heterogeneous-object models," *Journal of Economic Theory*, 222, Paper 105912.

Briest, P., S. Chawla, R. Kleinberg, and S. M. Weinberg (2015): "Pricing Lotteries," *Journal of Economic Theory*, 156, 144–174.

Carroll, G. (2017): "Robustness and Separation in Multidimensional Screening," *Econometrica*, 85, 453–488.

Che, Y.-K. and W. Zhong (2025): "Robustly Optimal Mechanisms for Selling Multiple Goods," *The Review of Economic Studies*, 92, 2923–2951.

Daskalakis, C., A. Deckelbaum, and C. Tzamos (2017): "Strong Duality for a Multiple-Good Monopolist," *Econometrica*, 85, 735–767.

Deb, R. and A.-K. Roesler (2023): "Multi-Dimensional Screening: Buyer-Optimal Learning and Informational Robustness," *Review of Economic Studies*.

Efron, B. (1965): "Increasing properties of Pólya frequency functions," *The Annals of Mathematical Statistics*, 36, 272–279.

Frick, M., R. Iijima, and Y. Ishii (2026): "Multidimensional Screening with Precise Seller Information," *Econometrica*, 94, 35–70.

Giannakopoulos, Y. and E. Koutsoupias (2018): "Selling Two Goods Optimally," *Information and Computation*, 261, 432–445.

Gretsky, N. E., J. M. Ostroy, and W. R. Zame (2002): "Subdifferentiability and the Duality Gap," *Positivity*, 6, 261–274.

Haghpanah, N. and J. Hartline (2021): "When is pure bundling optimal?" *The Review of Economic Studies*, 88, 1127–1156.

Hart, S. and N. Nisan (2017): "Approximate Revenue Maximization with Multiple Items," *Journal of Economic Theory*, 172, 313–347.

——— (2019): "Selling Multiple Correlated Goods: Revenue Maximization and Menu-Size Complexity," *Journal of Economic Theory*, 183, 991–1029.

Hart, S. and P. J. Reny (2015): "Maximal revenue with multiple goods: Nonmonotonicity and other observations," *Theoretical Economics*, 10, 893–922.

Kellerer, H. G. (1984): "Duality Theorems for Marginal Problems," *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 67, 399–432.

Kleiner, A. and A. Manelli (2019): "Strong Duality in Monopoly Pricing," *Econometrica*, 87, 1391–1396.

Kuratowski, K. and C. Ryll-Nardzewski (1965): "A general theorem on selectors," *Bulletin de l'Académie Polonaise des Sciences, Série des Sciences Mathématiques, Astronomiques et Physiques*, 13, 397–403.

Manelli, A. M. and D. R. Vincent (2006): "Bundling as an optimal selling mechanism for a multiple-good monopolist," *Journal of Economic Theory*, 127, 1–35.

——— (2007): "Multidimensional Mechanism Design: Revenue Maximization and the Multiple-Good Monopoly," *Journal of Economic Theory*, 137, 153–185.

McAfee, R. P. and J. McMillan (1988): "Multidimensional Incentive Compatibility and Mechanism Design," *Journal of Economic Theory*, 46, 335–354.

McAfee, R. P., J. McMillan, and M. D. Whinston (1989): "Multiproduct Monopoly, Commodity Bundling, and Correlation of Values," *The Quarterly Journal of Economics*, 104, 371–383.

Menicucci, D., S. Hurkens, and D.-S. Jeon (2015): "On the Optimality of Pure Bundling for a Monopolist," *Journal of Mathematical Economics*, 60, 33–42.

Pavlov, G. (2011): "Optimal Mechanism for Selling Two Goods," *The B.E. Journal of Theoretical Economics*, 11.

Prékopa, A. (1973): "On logarithmic concave measures and functions," *Acta Scientiarum Mathematicarum*, 34, 335–343.

Rochet, J.-C. (2024): "Multidimensional Screening After 37 Years," *Journal of Mathematical Economics*, 113, 103010.

Rochet, J.-C. and P. Chone (1998): "Ironing, Sweeping, and Multidimensional Screening," *Econometrica*, 66, 783–826.

Strassen, V. (1965): "The Existence of Probability Measures with Given Marginals," *Annals of Mathematical Statistics*, 36, 423–439.

Thanassoulis, J. (2004): "Haggling Over Substitutes," *Journal of Economic Theory*, 117, 217–245.

Yang, F. (2025): "Nested Bundling," *American Economic Review*, 115, 2970–3013.

A Proofs of strong duality

Lemma 2 (Revenue formula)**.** *Under Definition 1, for any incentive-compatible mechanism implementing u, the transfer satisfies $t(x) = a(x)\cdot\nabla U(x) - U(x)$ at λ_D-a.e. $x \in D$, where $U(x) := u(v(x))$; expected revenue therefore depends on the mechanism only through u.*

Proof of Lemma 2. The constraint $\partial u(y) \subseteq \mathcal{Q}$ makes u Lipschitz on Y, so $U = u \circ v$ is Lipschitz on the compact set D. By Rademacher's theorem, U is differentiable at λ_D-a.e. $x \in D$. At any such point, $\partial U(x) = \{\nabla U(x)\}$ is a singleton, and the convex chain rule $\partial U(x) = Jv(x)^\top \partial u(v(x))$ forces

$$Jv(x)^\top q \;=\; \nabla U(x) \qquad \text{for every } q \in \partial u(v(x)).$$

Tangency $Jv(x)\,a(x) = v(x)$ then gives

$$q(x)\cdot v(x) \;=\; q(x)\cdot Jv(x)\,a(x) \;=\; \big(Jv(x)^\top q(x)\big)\cdot a(x) \;=\; a(x)\cdot\nabla U(x).$$

Incentive compatibility provides $t(x) = q(x)\cdot v(x) - U(x)$, which together with the previous display yields the claimed formula. *Q.E.D.*

Proof of Lemma 1. $c(v(x), v(x')) = \sup_{q\in\mathcal{Q}} q\cdot V(x-x') = \sup_{p\in\mathcal{P}_V} p\cdot(x-x') = c_V(x-x')$, with the closed form coming from the polytope vertices $\{0\} \cup \{V_B^\top : B \in \mathcal{F}\}$. The vertex $V_B^\top$ attains the supremum at h if and only if $V_B \cdot h \geq V_{B'} \cdot h$ for every $B' \in \mathcal{F}$ and $V_B \cdot h \geq 0$, i.e., $h \in K_B^{\mathcal{F}}$. *Q.E.D.*

Proof of Theorem 1. *Roadmap.* The proof has three steps. Step 1 rewrites the primal as a conic linear program using the operator A that encodes the one-sided constraint $u(y) - u(y') \leq c(y, y')$. Step 2 derives a sup-norm Lipschitz bound to obtain compactness by ArzelÃ –Ascoli, giving primal attainment. Step 3 applies the Gretsky–Ostroy–Zame criterion, together with an infimal-convolution argument for the one-sided cost, to obtain no duality gap and dual attainment.

Step 0: normalization. The proof is written for a possibly nonconvex Y: pose the primal and the dual on $\widehat{Y} := \operatorname{conv}(Y)$, compact — hence closed — because Y is compact; under the convexity maintained in the body, $\widehat{Y} = Y$ and the statement is Theorem 1 verbatim. The transformed measure on D is

$$\mu = \delta_{x_0} + (f\,a\cdot n)\,\sigma_{\partial D} - \big[\nabla\cdot(f\,a) + f\big]\,\mathcal{L}^N|_D,$$

so the divergence theorem gives

$$\mu(D) = 1 + \int_{\partial D} f(x)\,a(x)\cdot n(x)\,\mathrm{d}\sigma(x) - \int_D \big[\nabla\cdot(f(x)\,a(x)) + f(x)\big]\,\mathrm{d}x = 1 - \int_D f(x)\,\mathrm{d}x = 0.$$

Hence the pushforward measure also has zero total mass:

$$\nu(Y) = \mu(D) = 0.$$

Adding a constant to a feasible utility therefore does not change the objective. Fix an arbitrary anchor point $y_0 \in Y$ and define

$$\mathcal{U}_0 := \big\{u \in C(\widehat{Y}) : u \text{ is convex and coordinatewise nondecreasing, and } u(y_0) = 0\big\}.$$

Because both $\gamma_1 - \gamma_2$ and ν have total mass zero, the constant normalization is immaterial. Since

$$\sup_{q\in\mathcal{Q}} q\cdot z = \max\Big\{0, \max_{B\in\mathcal{F}} z_B\Big\} = |z_+|_\infty,$$

the primal constraint $\partial u(y) \subseteq \mathcal{Q}$ implies the global bound

$$u(y) - u(y') \leq c(y, y'), \qquad y, y' \in \widehat{Y}.$$

Conversely, any u satisfying the bound coincides on $\widehat{Y}$ with its one-sided envelope $y \mapsto \inf_{z\in\widehat{Y}}\{u(z) + c(y,z)\}$, which extends u to a convex function on all of $\mathbb{R}^K$ with subdifferential in $\mathcal{Q}$ everywhere (the computation in Step 3 below, applied with $g = b$); this is exactly the admissibility defining $\mathcal{U}$. The two constraint sets therefore carry the same value of the normalized primal, and the argument works with the global bound: the normalized primal is equivalent to maximizing $\int_Y u \,\mathrm{d}\nu$ over $u \in \mathcal{U}_0$ subject to the one-sided bound $u(y) - u(y') \le c(y,y')$, the constraint Step 1 encodes. Exchanging y and y' in the bound yields the two-sided Lipschitz estimate

$$\left|u(y) - u(y')\right| \le \left|y - y'\right|_\infty, \qquad y, y' \in \widehat{Y},$$

which supplies the equicontinuity for Step 2.

Step 1: conic-LP form. Work in the Banach spaces $C(\widehat{Y})$ and $C(\widehat{Y}\times\widehat{Y})$ equipped with the sup norm. Let

$$\mathcal{V} := \{g \in C(\widehat{Y}\times\widehat{Y}) : g \ge 0\},$$

and define the linear map $A : C(\widehat{Y}) \to C(\widehat{Y}\times\widehat{Y})$ by

$$(Au)(y,y') := u(y) - u(y'),$$

together with the cost function $b(y,y') := c(y,y')$. The primal problem is

$$\sup_{u\in\mathcal{U}_0} \inf_{\gamma\in\Gamma_+(\widehat{Y}\times\widehat{Y})} \Big\{\langle \nu, u\rangle + \langle \gamma, b - Au\rangle\Big\}.$$

Interchanging the order of optimization yields the formal dual

$$\inf_{\gamma\in\Gamma_+(\widehat{Y}\times\widehat{Y})} \sup_{u\in\mathcal{U}_0} \Big\{\langle \gamma, b\rangle - \langle A^*\gamma - \nu, u\rangle\Big\}.$$

A direct calculation shows that

$$A^*\gamma = \gamma_1 - \gamma_2,$$

where γ_1 and γ_2 are the marginals of γ. Hence the dual reads

$$\inf_{\gamma\ge 0} \int_{\widehat{Y}\times\widehat{Y}} c(y,y')\,\mathrm{d}\gamma(y,y') \quad \text{s.t.} \quad \langle u, \gamma_1 - \gamma_2 - \nu\rangle \ge 0 \quad \forall\, u \in \mathcal{U}_0.$$

Since zero-mass measures are unaffected by constant shifts, this constraint is exactly

$$\gamma_1 - \gamma_2 \succeq_{\mathrm{mcvx}} \nu,$$

which is the dual problem (3).

Step 2: primal attainment. Let

$$F(b) := \{u \in \mathcal{U}_0 : Au \le b\}.$$

If $u \in F(b)$, Step 0 already gives $\left|u(y) - u(y')\right| \le \left|y - y'\right|_\infty$. Thus every $u \in F(b)$ is 1-Lipschitz with respect to the bundle-coordinate sup norm, so $F(b)$ is equicontinuous. Because each $u \in F(b)$ satisfies $u(y_0) = 0$ and $\widehat{Y}$ is compact, the family is also uniformly bounded. Convexity and coordinatewise monotonicity are preserved under uniform limits, so by ArzelÃ –Ascoli, $F(b)$ is compact in $C(\widehat{Y})$. The objective map $u \mapsto \int_Y u(y)\,\mathrm{d}\nu(y)$ is continuous on $C(\widehat{Y})$, hence attains a maximum on $F(b)$. The primal therefore attains its optimum.

Step 3: no duality gap and dual attainment. For $g \in C(\widehat{Y} \times \widehat{Y})$ define

$$F(g) := \{u \in \mathcal{U}_0 : Au \le g\}, \qquad V(g) := \sup\left\{\int_Y u(y)\,\mathrm{d}\nu(y) : u \in F(g)\right\}.$$

Fix $\bar{u} \in F(g)$ and define the one-sided envelope

$$u(y) := \inf_{z \in \widehat{Y}}\left\{\bar{u}(z) + c(y,z)\right\}, \qquad y \in \widehat{Y}.$$

Because $c(y,z) = \sup_{q \in \mathcal{Q}} q \cdot (y - z)$ is the support function of the compact convex set $\mathcal{Q}$, support-function calculus yields

$$u(y) = \sup_{q \in \mathcal{Q}}\left\{q \cdot y - \bar{u}^*(q)\right\}, \qquad \bar{u}^*(q) := \sup_{z \in \widehat{Y}}\left\{q \cdot z - \bar{u}(z)\right\}.$$

Hence u is convex and coordinatewise nondecreasing. The infimum defining $u(y')$ is attained because $\widehat{Y}$ is compact and $z \mapsto \bar{u}(z) + c(y',z)$ is continuous. For any $y, y' \in \widehat{Y}$ and any minimizer z for $u(y')$,

$$u(y) - u(y') \le c(y,z) - c(y',z) \le c(y,y'),$$

so $Au \le b$. Choosing $z = y$ gives $u(y) \le \bar{u}(y)$, and because $\bar{u} \in F(g)$,

$$\bar{u}(y) - \bar{u}(z) \le g(y,z) \qquad \text{for every } y, z \in \widehat{Y}.$$

Therefore

$$u(y) \ge \inf_{z \in \widehat{Y}}\left\{\bar{u}(y) - g(y,z) + c(y,z)\right\} \ge \bar{u}(y) - \left\|g - b\right\|_\infty,$$

so

$$0 \le \bar{u}(y) - u(y) \le \left\|g - b\right\|_\infty.$$

Set $\tilde{u} := u - u(y_0)$. Then $A\tilde{u} = Au \leq b$ and $\tilde{u} \in \mathcal{U}_0$. Since $\bar{u}(y_0) = 0$ and $\left|u(y_0)\right| \leq \left\|g - b\right\|_\infty$,

$$\left\|\bar{u} - \tilde{u}\right\|_\infty \leq 2 \left\|g - b\right\|_\infty,$$

hence

$$V(g) - V(b) \leq 2 \left\|\nu\right\|_{TV} \left\|g - b\right\|_\infty.$$

The value function V is concave: $\mathcal{U}_0$ is convex, the constraint $Au \leq g$ is jointly linear in (u, g) — so $\lambda F(g_1) + (1-\lambda)F(g_2) \subseteq F(\lambda g_1 + (1-\lambda)g_2)$ for $\lambda \in [0,1]$ — and the objective is linear in u, so $-V$ is convex. The bound $V(g) - V(b) \leq 2 \left\|\nu\right\|_{TV} \left\|g - b\right\|_\infty$, which holds for every $g \in C(\widehat{Y} \times \widehat{Y})$, shows that the upper difference quotient of V at b is bounded above — Condition 3 (p. 266) of Gretsky et al. (2002) for the subdifferential of $-V$ at b to be non-empty.[19] By the Gretsky–Ostroy–Zame theorem, non-emptiness of that subdifferential is equivalent to the absence of a duality gap together with attainment of the dual optimum. The primal and dual values therefore coincide, and the dual problem attains an optimizer. *Q.E.D.*

Proof of Proposition 1. Weak duality gives

$$\int_Y u \, \mathrm{d}\nu = \int_{Y \times Y} \big(u(y) - u(y')\big) \, \mathrm{d}\gamma \leq \int_{Y \times Y} c(y, y') \, \mathrm{d}\gamma,$$

where the identity uses $\gamma_1 - \gamma_2 = \nu$ from condition (iii). Under conditions (i) and (ii), the inequality holds with equality γ-almost surely: inside each transported linearity cell, the affine piece's quantity q^ℓ exactly supports the transported direction and attains the one-sided cost. The display in (ii) is exact complementary slackness; in the case $q^\ell = 0$ it forces $c(y, y') = 0$, so the transport is weakly downward in bundle values. Hence the candidate's primal value equals the dual cost. Optimality of both sides now follows from weak duality alone. Any primal-feasible u' extends to $\widehat{Y}$ by its one-sided envelope $\inf_{z \in Y}\{u'(z) + c(\cdot, z)\}$ — convex with quantities in $\mathcal{Q}$ by the support-function calculus of the strong-duality proof, and equal to u' on Y, hence objective-preserving; for the extension, monotone convexity and dual feasibility of any γ' give

$$\int_Y u' \, \mathrm{d}\nu \;\leq\; \int \big(u'(y) - u'(y')\big) \, \mathrm{d}\gamma' \;\leq\; \int c(y, y') \, \mathrm{d}\gamma'.$$

[19] V is *not* continuous at b: whenever $g(y, y) < 0$ for some y the feasible set $F(g)$ is empty, so $V(g) = -\infty$ arbitrarily close to b. The one-sided bound, not two-sided Lipschitz continuity, is what the criterion requires. This is the route of Kleiner and Manelli (2019), whose argument the present proof adapts: the interiority (Slater) condition fails here because A carries every admissible u to the boundary of the positive cone, $Au(y, y) = 0$.

Thus every primal value is at most every dual cost; the pair (u, γ) is feasible on both sides with matching values, hence u solves the primal and γ solves the dual — no appeal to Theorem 1, and hence no convexity of Y, is needed. So (u, γ) is a primal-dual saddle point. *Q.E.D.*

Proof of Proposition 2. The criterion reduces to Proposition 1 through the valuation map. *Admissibility.* $u_M(y) = \max_\ell \{q^\ell \cdot y - p_\ell\}$ is a maximum of finitely many affine functions with quantities $q^\ell \in \mathcal{Q}$, so it is convex with $\partial u_M(y) = \operatorname{conv}\{q^\ell : \ell \text{ active at } y\} \subseteq \mathcal{Q}$, and the outside option gives $u_M \geq q^0 \cdot y - p_0 = 0$; hence $u_M \in \mathcal{U}$, piecewise linear on the polyhedral cells of the menu. *Cell correspondence.* Under Assumption 1, $(V^\top q^\ell)\cdot x - p_\ell = q^\ell \cdot v(x) - p_\ell$, so $U_M = u_M \circ v$ and the maximizing option at x is the maximizing option at $v(x)$: the cells correspond, $C_\ell = v^{-1}(\widetilde{C}_\ell)$ with $\widetilde{C}_\ell$ the linearity cells of u_M on Y. *The pushed-forward plan.* Let $\tilde{\gamma} := (v \times v)_\# \gamma \in \Gamma_+(Y \times Y)$. For condition (i): if x, x' lie in the closure of one cell C_ℓ, then $v(x), v(x')$ lie in the closure of $\widetilde{C}_\ell$ by continuity of v. For condition (ii): the displacement $h := x - x'$ lies in $K_B^{\mathcal{F}}$ for every $B \in \operatorname{supp} \sigma_\ell$, so $V_B \cdot h = c_V(h)$ at each such B by Lemma 1; averaging against σ_ℓ (each option allocates with total probability one) gives the cost identity

$$c\big(v(x), v(x')\big) = c_V(h) = (V^\top q^\ell) \cdot h = q^\ell \cdot \big(v(x) - v(x')\big),$$

and within the cell the common price cancels, $u_M(v(x)) - u_M(v(x')) = q^\ell \cdot (v(x) - v(x'))$ — exactly the complementary-slackness display of Proposition 1(ii). The outside option is the case $\ell = 0$: $\operatorname{supp} \sigma_0 = \{\varnothing\}$ with $V_\varnothing = 0$, so $K_\varnothing^{\mathcal{F}} = \{h : c_V(h) = 0\}$ and the same display reads $c_V(h) = 0$. For condition (iii): the marginals push forward, $\tilde{\gamma}_1 - \tilde{\gamma}_2 = v_\#(\gamma_1 - \gamma_2) = v_\# \mu = \nu$. Proposition 1 therefore applies to $(u_M, \tilde{\gamma})$: u_M solves the primal and, by Lemma 2, the menu M implementing u_M is optimal. *Q.E.D.*

Remark 2 (Within-cell support). Condition (i) is not an independent hypothesis *at optimal pairs*: if u is primal-optimal of the stated piecewise-linear form, then any dual-optimal γ with absolutely continuous sink marginal satisfies it. Indeed, when the primal and dual values coincide — Theorem 1 — the chain $\int_Y u \, \mathrm{d}\nu \leq \int u \, \mathrm{d}(\gamma_1 - \gamma_2) = \int (u(y) - u(y')) \, \mathrm{d}\gamma \leq \int c \, \mathrm{d}\gamma$ collapses to equalities, so the complementary slackness $u(y) - u(y') = c(y, y')$ holds γ-a.e. Applied to a pair $y \in C_k$, $y' \in C_\ell$ with $u = \max_m (a_m + q^m \cdot \cdot)$ and $c(y, y') \geq q^k \cdot (y - y')$, it forces $a_k + q^k \cdot y' \geq a_\ell + q^\ell \cdot y'$; since $y' \in C_\ell$ gives the reverse inequality, y' lands on the indifference set $\{(q^k - q^\ell) \cdot z = a_\ell - a_k\}$, which for distinct affine pieces is a proper hyperplane (when $q^k \neq q^\ell$) or

empty (when $q^k = q^\ell$, $a_k \neq a_\ell$), hence Lebesgue-null. Hence $\gamma(C_k \times C_\ell) = 0$ for all $k \neq \ell$: γ is supported on within-cell pairs. Primal optimality of u is essential — for a suboptimal candidate the chain stays strict and no support restriction follows. The applied version is the one for Proposition 2 in the coordinates of D, where the sink marginal μ^- is absolutely continuous with respect to λ_D and the argument runs verbatim, with the exclusion cell as the case $\ell = 0$, $(q^0, a_0) = (0, 0)$.

B Proof of the core-bundle theorem

This appendix proves Theorem 2; the standing form is in force throughout. The candidate menu $M = \{(\sigma_\ell, p_\ell)\}_{\ell=1}^L$ has selling cells $\{C_\ell\}_{\ell \geq 1}$ and exclusion cell C_0, with the tie convention assigning each type to one cell. For a cell ℓ write w_ℓ for the marginal-inclusion vector and $s_\ell(x) := w_\ell \cdot x$ for the effective cell sum, and classify its peripherals by weight: full-weight $S_\ell^{(1)} := \{j \in P : w_{\ell,j} = 1\}$, fractional $B_\ell^{\text{frac}} := \{j \in P : 0 < w_{\ell,j} < 1\}$, zero-weight Z_ℓ the rest. The *active set* $A_\ell := \{i : w_{\ell,i} > 0\} = C \cup S_\ell^{(1)} \cup B_\ell^{\text{frac}}$ collects the coordinates of positive weight; set $B_\ell^{(1)} := C \cup S_\ell^{(1)}$ and $t_C(x) := \sum_{i \in C} x_i$, with corner $\bar{t}_C := \sum_{i \in C} \bar{v}_i$. The boundary and interior parts $\mu^\pm$ of the differential virtual value are as in Section 4.1, with cell restrictions $\mu_\ell^\pm := \mu^\pm|_{C_\ell}$. The pulled-back one-sided cost splits as $c_\alpha = \max(\alpha\, c^{\mathcal{F}_C}, c^P)$ into a *core-only cost* $c^{\mathcal{F}_C}(h) := \max\{0, \max_{B \in \mathcal{F}_C} \sum_{i \in B} h_i\}$ (the α-premium shown explicitly) and a *non-core cost* $c^P(h) := \max\{0, \max_{B \notin \mathcal{F}_C} \sum_{i \in B} h_i\}$, which carries no premium. Cell ℓ's *core-only cone* is $K_\ell^{\mathcal{F}_C} := \{h \in \mathbb{R}^N : s_\ell(h) \geq 0 \text{ and } s_\ell(h) \geq \sum_{i \in B} h_i \text{ for every } B \in \mathcal{F}_C\}$ — the displacements on which the option's marginal-inclusion vector w_ℓ attains the core-only cost, so that $c^{\mathcal{F}_C}(h) = s_\ell(h)$ there.

The construction is a transport plan certifying Proposition 2(iii); it consumes the candidate's optimality through one first-order condition, cell mass balance:

$$\text{(MB)} \qquad \mu(C_\ell) = 0, \quad \text{i.e.} \quad \mu^+(C_\ell) = \mu^-(C_\ell), \qquad \ell \geq 1. \tag{5}$$

We may take the candidate menu to be *face-generic*: no indifference hyperplane between two options, or between an option and exclusion, coincides with a top face $F_j := \{x_j = \bar{v}_j\}$. This is without loss. Were there such a coincidence, the two options would differ only in item j's marginal inclusion ($w_\ell - w_m \parallel e_j$), and, the tie hyperplane lying on $\{x_j = \bar{v}_j\}$, their surplus difference would be $\propto (x_j - \bar{v}_j)$ — single-signed on D; one option would then weakly dominate the other throughout D (equality only on the λ_D-null face F_j), and deleting the dominated, redundant

option leaves the indirect utility, hence revenue, unchanged. Iterating over the finite menu yields a face-generic representative of equal revenue.

Lemma 3. *If M is optimal in the restricted problem on $\mathcal{F}_C$, then each selling cell satisfies* (5).

Proof. For an $\mathcal{F}_C$-mechanism with option utilities $\bar{v}_m(x) - p_m$, the indirect utility is $U_p(x) = \max\{0,\ \max_m(\bar{v}_m(x) - p_m)\}$ and seller revenue is the primal objective $R = \int_D U_p \, \mathrm{d}\mu$ (Lemma 2 and (2), with μ the differential virtual value of Section 2.2). Fix a selling cell ℓ and perturb its price alone, $p_\ell \mapsto p_\ell + \epsilon$, holding the options $\{\sigma_m\}$ and the other prices fixed; the perturbed mechanism is again supported on $\mathcal{F}_C$. The map $\epsilon \mapsto U_{p+\epsilon e_\ell}(x)$ is 1-Lipschitz (a single price moves a finite maximum by at most $|\epsilon|$), and for λ_D-a.e. x the maximizing option is unique, so the derivative exists with $\partial U_p / \partial p_\ell = -\mathbf{1}_{\operatorname{int} C_\ell}(x)$. The exceptional set — where two options, or an option and exclusion, tie — is a finite union of affine hyperplanes: λ_D-null, hence μ^--null, and (by the tie convention and face-genericity) meeting each top face in dimension $\leq N - 2$ and avoiding the origin atom, hence μ^+-null. Dominated convergence against the finite measures $\mu^\pm$ (difference quotient bounded by 1) gives

$$\frac{\mathrm{d}}{\mathrm{d}p_\ell} R \;=\; \int_D \frac{\partial U_p}{\partial p_\ell} \, \mathrm{d}\mu \;=\; -\mu(\operatorname{int} C_\ell) \;=\; -\mu(C_\ell),$$

the last equality because ∂C_ℓ is μ-null. A selling cell has $\lambda_D(C_\ell) > 0$ (face-genericity excludes the degenerate null cells), so option ℓ wins on a set of positive measure under every small perturbation; p_ℓ is thus an interior maximizer of R along this $\mathcal{F}_C$-family, and restricted optimality forces $\mathrm{d}R/\mathrm{d}p_\ell = 0$, i.e. $\mu(C_\ell) = 0$. *Q.E.D.*

Global balance $\mu(D) = 0$ (divergence theorem on the measure μ of Section 2.2) extends (5) to the exclusion cell, where inclusivity makes $\mu^+|_{C_0}$ the origin atom of mass 1; so each pair $\mu_\ell^\pm$ has equal total mass, the prerequisite for the couplings below. The construction consumes optimality only through (5), hence applies to any finite, inclusive, alignment-regular menu whose selling cells are mass-balanced — the form in which the applications of Section 5 are verified.

Normalization in α. Since (4) scales each core-containing bundle value by α, the candidate is specified scale-invariantly (the theorem's footnote): prices scale linearly, $p_\ell = \alpha \, \hat{s}_\ell$, with α-invariant crossing points $\hat{s}_\ell := p_\ell / \alpha$; the cells $\{C_\ell\}$, the

vectors w_ℓ, and the sums s_ℓ are therefore α-independent. The saddle pair is positively homogeneous, so it suffices to verify Proposition 2 in *normalized units*: write $\tilde{u}_M := u_M/\alpha$, so that $\tilde{u}_M(x) = s_\ell(x) - \hat{s}_\ell$ on a selling cell and $\tilde{u}_M = 0$ on C_0. Dividing the cost by α as well — writing $\tilde{c}_\alpha := c_\alpha/\alpha = \max\big(c^{\mathcal{F}_C}, c^P/\alpha\big)$ for the normalized full cost — the within-cell complementary-slackness identity reads $\tilde{u}_M(y^1) - \tilde{u}_M(y^2) = s_\ell(h) = c^{\mathcal{F}_C}(h)$ on $K_\ell^{\mathcal{F}_C}$, and α *enters the rest of the proof through a single comparison*: for a within-cell displacement $h \in K_\ell^{\mathcal{F}_C}$, full-cost slackness holds iff $c^P(h) \le \alpha\, s_\ell(h)$, equivalently iff h lies in the full cone

$$K_\ell^{\mathcal{P}} = \big\{h \in K_\ell^{\mathcal{F}_C} : c^P(h) \le \alpha\, s_\ell(h)\big\}.$$

Every cone membership, coupling, and slackness identity below is stated in these α-free units; raising α widens $K_\ell^{\mathcal{P}}$ toward $K_\ell^{\mathcal{F}_C}$, and $\alpha^*(C,M)$ is the level at which the comparison holds for γ'-a.e. pair.

Roadmap. Lemma 4 pins the sign pattern of within-cell displacements on the cell's core-only cone and identifies the normalized core-only cost as $s_\ell(h)$. By (5) each pair $\mu_\ell^\pm$ has equal mass, so Lemma 5 assembles any family of within-cell couplings $\{\gamma'_\ell\}$ into a single transport realizing condition (iii) of Proposition 2, supported on within-cell pairs by construction. Lemma 6 then builds the selling-cell couplings under alignment regularity at level κ: every coupled displacement lies in the budget relation $K_{\ell,\kappa}$ — a subcone of the core-only cone carrying a uniform budget slack — on which $c^P(h) \le \alpha\, s_\ell(h)$ for every $\alpha \ge 1/(1-\kappa)$. Lemma 10 converts this into full-cost complementary slackness, and the concluding verification of Theorem 2 closes condition (ii) of Proposition 2 once $\alpha \ge \alpha^*(C,M) := 1/(1-\kappa)$.

B.1 Cone-forcing and cost anatomy

Lemma 4 (Cone-forcing and cost anatomy)**.** *For every selling cell C_ℓ and every $h \in K_\ell^{\mathcal{F}_C}$:*

(1) $h_j = 0$ *for every* $j \in B_\ell^{\mathrm{frac}}$;

(2) $h_j \le 0$ *for every* $j \notin A_\ell$;

(3) $h_l \ge 0$ *for every* $l \in S_\ell^{(1)}$, *and hence* $\sum_T h_l \ge 0$ *for every* $T \subseteq S_\ell^{(1)}$;

(4) *the core-only cost satisfies* $c^{\mathcal{F}_C}(h) = s_\ell(h)$ *with* $s_\ell(h) = \sum_{i \in B_\ell^{(1)}} h_i$;

(5) *the non-core cost satisfies* $c^P(h) \le s_\ell(h) + \sum_{i\in C}(-h_i)^+$, *equivalently* $c^P(h) \le \sum_{i\in B_\ell^{(1)}} h_i^+$.

Proof. Every $B \in \mathcal{F}_C$ contains C and $w_{\ell,i} = 1$ on C, so the cone-defining inequality $\alpha\, s_\ell(h) \ge \alpha \sum_{i\in B} h_i$ cancels the common core terms; writing $B = C \cup S$ for $S \subseteq P$, it reduces to

$$\sum_{j\in P} w_{\ell,j}\, h_j \;\ge\; \sum_{j\in S} h_j \qquad \text{for every } S \subseteq P. \tag{6}$$

Sign analysis. Specialize (6) to $S^+ := \{j \in P : h_j > 0\}$. Splitting P into S^+ and $P \setminus S^+$ and rearranging,

$$\sum_{j\in S^+} (w_{\ell,j} - 1)\, h_j \;+\; \sum_{j\in P\setminus S^+} w_{\ell,j}\, h_j \;\ge\; 0.$$

For $j \in S^+$ each summand is ≤ 0 (since $w_{\ell,j} \le 1$ and $h_j > 0$); for $j \in P \setminus S^+$ each summand is ≤ 0 (since $w_{\ell,j} \ge 0$ and $h_j \le 0$). The non-negative sum forces every summand to vanish:

$$h_j > 0 \Rightarrow w_{\ell,j} = 1, \qquad h_j < 0 \Rightarrow w_{\ell,j} = 0, \qquad 0 < w_{\ell,j} < 1 \Rightarrow h_j = 0.$$

Parts (1)–(3) follow: a fractional item has $0 < w_{\ell,j} < 1$, so $h_j = 0$; for $j \notin A_\ell$ ($w_{\ell,j} = 0$, hence $j \in P$ as $C \subseteq A_\ell$), contrapositive gives $h_j \le 0$; for $l \in S_\ell^{(1)}$ ($w_{\ell,l} = 1$), contrapositive gives $h_l \ge 0$, and summing non-negatives yields the subset-sum statement.

Cost identities. On $K_\ell^{\mathcal{F}_C}$ the cell sum collapses: fractional terms vanish by (1), zero-weight terms vanish because $w_{\ell,j} = 0$, so

$$s_\ell(h) = w_\ell \cdot h = \sum_{i\in C} h_i + \sum_{j\in S_\ell^{(1)}} h_j = \sum_{i\in B_\ell^{(1)}} h_i.$$

Since $B_\ell^{(1)} \in \mathcal{F}_C$ (it contains C), $\max_{B\in\mathcal{F}_C} \sum_B h_i \ge s_\ell(h)$, while (6) gives the reverse inequality; together with $s_\ell(h) \ge 0$ on $K_\ell^{\mathcal{F}_C}$, this proves (4).

For (5), take any non-core $B \ne \varnothing$; then $C \not\subseteq B$. Decompose $\sum_B h_i = \sum_{B\cap C} h_i + \sum_{B\cap S_\ell^{(1)}} h_j + \sum_{B\cap B_\ell^{\mathrm{frac}}} h_j + \sum_{B\cap(P\setminus A_\ell)} h_j$. By (1) the fractional sum is 0; by (2) the zero-weight sum is ≤ 0; by (3) the full-weight peripheral sum is $\le \sum_{S_\ell^{(1)}} h_j^+$ (with $h_j = h_j^+$ on $S_\ell^{(1)}$ since $h_j \ge 0$ there). The core piece is $\sum_{B\cap C} h_i \le \sum_{i\in C} h_i^+$ termwise. Thus

$$\sum_B h_i \;\le\; \sum_C h_i^+ + \sum_{S_\ell^{(1)}} h_j^+ \;=\; \sum_{B_\ell^{(1)}} h_i^+,$$

giving $c^P(h) \le \sum_{B_\ell^{(1)}} h_i^+$. Using $h_i^+ = h_i + (-h_i)^+$ and $\sum_{S_\ell^{(1)}}(-h_i)^+ = 0$ (by (3)), this rewrites as $s_\ell(h) + \sum_C (-h_i)^+$, proving (5). The lemma pins the geometry of permissible within-cell displacements — zero motion on fractional peripherals, downward on zero-weight items, upward on full-weight peripherals, core free — and identifies the normalized core-only cost as $s_\ell(h)$, both identities feeding the downstream within-cell ratio bound. *Q.E.D.*

B.2 Within-cell coupling freedom

Lemma 5 (Within-cell coupling freedom)**.** *Write $\mu_\ell^\pm := \mu^\pm|_{C_\ell}$; by (5) on the selling cells and global balance $\mu(D) = 0$ on the exclusion cell, the two have equal total mass on every cell. Suppose that for each cell $\ell \ge 0$ a coupling γ'_ℓ of μ_ℓ^+ and μ_ℓ^- supported on $C_\ell \times C_\ell$ is given, with $h := y^1 - y^2 \in K_\ell^{\mathcal{F}_C}$ for γ'_ℓ-a.e. pair on each selling cell ($\ell \ge 1$). Then $\gamma' := \sum_{\ell \ge 0} \gamma'_\ell$ (i) transports μ^+ to μ^-; (ii) is supported on within-cell pairs; and (iii) satisfies the core-only complementary-slackness identity $c^{\mathcal{F}_C}(h) = \tilde{u}_M(y^1) - \tilde{u}_M(y^2)$ for γ'-a.e. pair on each selling cell.*

Proof. Equal cell masses, from (5), make each γ'_ℓ a genuine coupling; the summands have pairwise disjoint supports $C_\ell \times C_\ell$, so $\gamma' \ge 0$ and (ii) holds.

(i) The tie-breaking convention assigns each type to exactly one cell, so $\{C_\ell\}_{\ell \ge 0}$ partitions D as Borel sets and $\mu^\pm = \sum_\ell \mu_\ell^\pm$ exactly — for μ^- the cell boundaries are in addition Lebesgue-null, while any μ^+ face mass on a cell boundary is assigned by the convention. Marginalizing, $(\mathrm{pr}_1)_*\gamma' = \sum_\ell (\mathrm{pr}_1)_*\gamma'_\ell = \sum_\ell \mu_\ell^+ = \mu^+$, and likewise $(\mathrm{pr}_2)_*\gamma' = \mu^-$.

(iii) On a selling cell fix $y^1, y^2 \in C_\ell$ with $h \in K_\ell^{\mathcal{F}_C}$. There $\tilde{u}_M(x) = s_\ell(x) - \hat{s}_\ell$, so $\tilde{u}_M(y^1) - \tilde{u}_M(y^2) = s_\ell(h)$ (crossing points cancel), and Lemma 4(4) gives $c^{\mathcal{F}_C}(h) = s_\ell(h)$. So CS holds. *Q.E.D.*

B.3 Within-cell coupling under alignment regularity

Alignment regularity delivers the within-cell coupling for every selling cell, through the monotone-coupling form of Strassen's theorem applied to the cell's alignment-score laws.

Theorem 7 (Relational Strassen (Strassen, 1965; Kellerer, 1984))**.** *Let X, Y be compact metric spaces, $\mu \in \mathcal{P}(X)$, $\nu \in \mathcal{P}(Y)$, and $G \subseteq X \times Y$ closed. A coupling $\pi \in \Pi(\mu, \nu)$ with $\pi(G) = 1$ exists if and only if*

$$\nu(B) \le \mu\big(G^{-1}(B)\big) \qquad \text{for every closed } B \subseteq Y, \quad G^{-1}(B) := \mathrm{proj}_X\big(G \cap (X \times B)\big).$$

In particular, for finite Borel measures ν_1, ν_2 of equal mass on a compact subset of $\mathbb{R}^m$: $\nu_1 \preceq_{\mathrm{FOSD}} \nu_2$ if and only if there is a coupling of (ν_2, ν_1) concentrated on $\{(u, v) : u \geq v \text{ coordinatewise}\}$.[20]

Lemma 6 (Within-cell coupling)**.** *Suppose (f, M) is alignment-regular at level $\kappa \in [0, 1)$ (Definition 3) and the selling cells are mass-balanced (5). For each selling cell C_ℓ there is an α-independent coupling γ'_ℓ of μ^+_ℓ, μ^-_ℓ such that, for γ'_ℓ-a.e. pair (writing $h = y^1 - y^2$):*

(a) *$h \in K_\ell^{\mathcal{F}_C}$ and $s_\ell(h) \geq (1-\kappa)\big(\sum_{i\in C} h_i^+ + \sum_{j\in S_\ell^{(1)}} h_j\big) \geq 0$;*

(b) *$c^P(h) \leq \alpha\, s_\ell(h)$ for every $\alpha \geq 1/(1-\kappa)$.*

The proof of Lemma 6 goes through the *budget relation*: a polyhedral cone of displacements whose membership is encoded, coordinate by coordinate, by the alignment scores of Definition 3. Fix a selling cell ℓ and a level $\kappa \in [0, 1)$; the standing form is in force, and each win region is a closed polytopal subset of the box (an intersection of the affine expected-utility half-spaces with D), so couplings live on its closure with marginals assigned by the tie convention.

Definition 6 (Budget relation)**.** *For $\kappa \in [0, 1)$ and a selling cell ℓ, let*

$$K_{\ell,\kappa} := \Big\{h \in \mathbb{R}^N :\ h_j \geq 0\ \big(j \in S_\ell^{(1)}\big),\ \ h_j \leq 0\ (j \in Z_\ell),\ \ h_j = 0\ \big(j \in B_\ell^{\mathrm{frac}}\big),\ \ \sum_{i\in C} h_i^- \leq \kappa\Big(\sum_{j\in S_\ell^{(1)}} h_j + \sum_{i\in C} h_i^+\Big)\Big\}.$$

The sign clauses replicate the pattern cone-forcing imposes on covered displacements (Lemma 4(1)–(3)); the budget clause caps core dips at a κ-fraction of the aligned gains.

Lemma 7 (Generators)**.** *$h \in K_{\ell,\kappa}$ if and only if every coordinate functional of Θ^κ_ℓ is nonnegative at h. Consequently $K_{\ell,\kappa}$ is a closed convex polyhedral cone, and for $x, y \in \mathbb{R}^N$: $y - x \in K_{\ell,\kappa}$ if and only if $\Theta^\kappa_\ell(y) \geq \Theta^\kappa_\ell(x)$ coordinatewise.*

Proof. ($\Rightarrow$) Let $h \in K_{\ell,\kappa}$. The peripheral coordinates of Θ^κ_ℓ are nonnegative at h by the sign clauses (the fractional pairs $\pm h_j$ both vanish). For $\varnothing \neq T \subseteq C$,

$$\sigma^\kappa_T(h) = \sum_{i\in C} c_i\, h_i + \kappa \sum_{j\in S_\ell^{(1)}} h_j, \qquad c_i = 1\ (i \in T), \qquad c_i = \kappa\ (i \in C \setminus T),$$

[20] After common normalization, apply the first part with G the coordinatewise order: indicators of closed up-sets are bounded monotone Borel, so $\nu_1 \preceq_{\mathrm{FOSD}} \nu_2$ gives $\nu_1(U) \leq \nu_2(U)$ for every closed up-set U; $G^{-1}(B)$ is the up-closure of B, closed by compactness, so the Hall criterion reads $\nu_1(B) \leq \nu_1(G^{-1}(B)) \leq \nu_2(G^{-1}(B))$. The converse direction is immediate from the coupling.

and since $c_i \in [\kappa, 1]$, termwise $c_i h_i \geq \kappa h_i^+ - h_i^-$ (for $h_i \geq 0$ use $c_i \geq \kappa$; for $h_i < 0$ use $c_i \leq 1$); summing,

$$\sigma_T^\kappa(h) \;\geq\; \kappa \sum_{i \in C} h_i^+ \;-\; \sum_{i \in C} h_i^- \;+\; \kappa \sum_{j \in S_\ell^{(1)}} h_j \;\geq\; 0$$

by the budget clause. ($\Leftarrow$) The peripheral coordinates recover the sign clauses, the fractional pairs forcing $h_j = 0$. If the dipping set $T^* := \{i \in C : h_i < 0\}$ is empty the budget clause holds trivially; otherwise $\sum_{i \in T^*} h_i = -\sum_C h_i^-$ and $\sum_{i \in C \setminus T^*} h_i = \sum_C h_i^+$, so $\sigma_{T^*}^\kappa(h) \geq 0$ reads $-\sum_C h_i^- + \kappa \sum_C h_i^+ + \kappa \sum_{S_\ell^{(1)}} h_j \geq 0$, which is the budget clause. Thus $K_{\ell,\kappa}$ is the set where finitely many linear functionals are nonnegative — a closed convex polyhedral cone — and the pair form follows by linearity of each coordinate functional. *Q.E.D.*

Lemma 8 (Admissibility)**.** *Every $h \in K_{\ell,\kappa}$ satisfies:* (i) *the covered-cone inequalities $w_\ell \cdot h \geq 0$ and $w_\ell \cdot h \geq \sum_{i \in B} h_i$ for every $B \in \mathcal{F}_C$, so $h \in K_\ell^{\mathcal{F}_C}$;* (ii) *$s_\ell(h) \geq (1 - \kappa)\big(\sum_{i \in C} h_i^+ + \sum_{j \in S_\ell^{(1)}} h_j\big) \geq 0$;* (iii) *$c^P(h) \leq \alpha\, s_\ell(h)$ for every $\alpha \geq 1/(1-\kappa)$.*

Proof. Write $G := \sum_C h_i^+ + \sum_{S_\ell^{(1)}} h_j \geq 0$ (nonnegative by the sign clause on $S_\ell^{(1)}$), $L := \sum_C h_i^-$, and $m := \min_{i \in C} h_i^+ \geq 0$; the budget clause reads $L \leq \kappa G$.

(ii) On $K_{\ell,\kappa}$ the fractional terms of $s_\ell(h) = w_\ell \cdot h$ vanish and the zero-weight items carry weight 0, so $s_\ell(h) = \sum_C h_i + \sum_{S_\ell^{(1)}} h_j = G - L \geq (1 - \kappa)\, G \geq 0$.

(i) The first inequality is (ii). For $B = C \cup S \in \mathcal{F}_C$ with $S \subseteq P$, the common core terms cancel and

$$w_\ell \cdot h - \sum_{i \in B} h_i \;=\; \sum_{j \in S_\ell^{(1)} \setminus S} h_j \;-\; \sum_{j \in S \cap Z_\ell} h_j \;\geq\; 0,$$

each block nonnegative by the sign clauses (fractional terms vanish on both sides). Membership in $K_\ell^{\mathcal{F}_C} = \bigcap_{B : \sigma_\ell(B) > 0} K_B^{\mathcal{F}_C}$ follows by a mean-equals-max argument: $w_\ell \cdot h = \mathbb{E}_{B \sim \sigma_\ell}\big[\sum_{i \in B} h_i\big]$ with every $B \in \operatorname{supp} \sigma_\ell$ in $\mathcal{F}_C$, so the average weakly dominating the maximum forces $\sum_{i \in B} h_i = w_\ell \cdot h \geq \max\big(0, \max_{B' \in \mathcal{F}_C} \sum_{i \in B'} h_i\big)$ for every B in the support, i.e. $h \in K_B^{\mathcal{F}_C}$ for each of them.

(iii) A non-core bundle B omits some core item i_0, so $\sum_{i \in B} h_i \leq \sum_{k \in [N]} h_k^+ - h_{i_0}^+ \leq G - m$, using $\sum_k h_k^+ = G$ under the sign clauses; hence $c^P(h) \leq (G - m)^+$. For $\alpha \geq 1/(1-\kappa)$,

$$\alpha\, s_\ell(h) - (G - m) \;=\; \alpha(G - L) - G + m \;\geq\; \big[\alpha(1 - \kappa) - 1\big] G + m \;\geq\; 0$$

by the budget clause $L \le \kappa G$, and $\alpha\, s_\ell(h) \ge 0$ by (ii); together $c^P(h) \le \max\{0,\, G - m\} \le \alpha\, s_\ell(h)$. *Q.E.D.*

Lemma 9 (Transfer and certification)**.** *Let the cell be mass-balanced* (5)*. The score dominance* $\nu^-_{\ell,\kappa} \preceq_{\mathrm{FOSD}} \nu^+_{\ell,\kappa}$ *holds* if and only if *there is a coupling* γ'_ℓ *of* (μ^+_ℓ, μ^-_ℓ) *with* $y^1 - y^2 \in K_{\ell,\kappa}$ *for* γ'_ℓ*-a.e. pair. Since* $K_{\ell,\kappa}$ *grows with* κ*, the levels at which dominance holds form an up-interval, so the existential level in Definition 3 is well posed.*

Proof. ($\Rightarrow$) By (5), $\nu^+_{\ell,\kappa}$ and $\nu^-_{\ell,\kappa}$ have equal total mass on the compact set $\Theta^\kappa_\ell(\overline{C_\ell}) \subset \mathbb{R}^{m_\ell}$. By the monotone-coupling form of Theorem 7 there is a coupling π of $(\nu^+_{\ell,\kappa}, \nu^-_{\ell,\kappa})$ with $u \ge v$ coordinatewise for π-a.e. (u, v). Disintegrate over score fibers (regular conditional kernels exist on Polish spaces): $\mu^+_\ell(dy) = \int K^+(u, dy)\, \nu^+_{\ell,\kappa}(du)$ with $K^+(u,\cdot)$ concentrated on the fiber $(\Theta^\kappa_\ell)^{-1}(u)$ for $\nu^+_{\ell,\kappa}$-a.e. u, and similarly K^- for μ^-_ℓ. Glue:

$$\gamma'_\ell(dy^1, dy^2) \;:=\; \int K^+(u, dy^1)\, K^-(v, dy^2)\, \pi(du, dv).$$

The marginals are μ^+_ℓ and μ^-_ℓ by Fubini and the disintegration identities. For γ'_ℓ-a.e. pair, $\Theta^\kappa_\ell(y^1) = u \ge v = \Theta^\kappa_\ell(y^2)$ coordinatewise, hence $y^1 - y^2 \in K_{\ell,\kappa}$ by Lemma 7. ($\Leftarrow$) Given such a γ'_ℓ, for every bounded coordinatewise nondecreasing φ,

$$\int \varphi\, d\nu^-_{\ell,\kappa} = \int \varphi\big(\Theta^\kappa_\ell(y^2)\big)\, d\gamma'_\ell \le \int \varphi\big(\Theta^\kappa_\ell(y^1)\big)\, d\gamma'_\ell = \int \varphi\, d\nu^+_{\ell,\kappa},$$

using $\Theta^\kappa_\ell(y^1) \ge \Theta^\kappa_\ell(y^2)$ on the support (Lemma 7); this is the dominance. The monotonicity in κ is immediate from $K_{\ell,\kappa} \subseteq K_{\ell,\kappa'}$ for $\kappa \le \kappa'$. *Q.E.D.*

The converse is how the applications of Section 5 verify alignment regularity: they construct a $K_{\ell,\kappa}$-coupling and read off the dominance.

Proof of Lemma 6. Lemma 9 produces, in each selling cell, a coupling γ'_ℓ concentrated on $K_{\ell,\kappa}$ displacements; it is α-independent because the cells, the weights, the scores, and $\mu^\pm_\ell$ are all α-free. Lemma 8(i)–(ii) deliver part (a), and Lemma 8(iii) delivers part (b). *Q.E.D.*

B.4 Full-cost dominance

Lemma 10 (Full-cost dominance)**.** *In the setting of Lemma 6 — (f, M) alignment-regular at level κ, with γ'_ℓ the couplings it provides — for every $\alpha \ge 1/(1-\kappa)$ and for γ'_ℓ-a.e. within-cell pair, full-cost slackness holds: $c^P(h) \le \alpha\, s_\ell(h)$, so $\tilde{c}_\alpha(h) = s_\ell(h) = \tilde{u}_M(y^1) - \tilde{u}_M(y^2)$.*

Proof. Lemma 6 supplies a coupling γ'_ℓ satisfying parts (a)/(b). Part (a) gives $h \in K_\ell^{\mathcal{F}_C}$ and $s_\ell(h) \geq 0$ for γ'_ℓ-a.e. pair; part (b) gives $c^P(h) \leq \alpha\, s_\ell(h)$ for every $\alpha \geq 1/(1-\kappa)$, i.e., $h \in K_\ell^{\mathcal{P}}$. By Lemma 4(4) the normalized full cost $\tilde{c}_\alpha(h) = \max\big(s_\ell(h),\, c^P(h)/\alpha\big) = s_\ell(h) = c^{\mathcal{F}_C}(h)$. Lemma 5(iii) gives the core-only identity $c^{\mathcal{F}_C}(h) = \tilde{u}_M(y^1) - \tilde{u}_M(y^2)$; combining, the full-cost CS holds. This is exactly condition (ii) of Proposition 2: since every offered bundle contains the core, the option gradient is $V^\top q^\ell = \alpha\, w_\ell$ and $c_V = \max(\alpha\, c^{\mathcal{F}_C},\, c^P)$, so on $K_\ell^{\mathcal{F}_C}$ (where $c^{\mathcal{F}_C} = s_\ell$) membership $h \in K_\ell^{\mathcal{P}}$ reads $(V^\top q^\ell)\cdot h = \alpha\, s_\ell(h) = c_V(h)$, i.e. the displacement lies in the cone option ℓ covers. *Q.E.D.*

B.5 Verification of Theorem 2

Proof of Theorem 2. By Lemma 3, restricted optimality makes every selling cell mass-balanced (5), so each pair $\mu_\ell^\pm$ has equal total mass. Let $\kappa \in [0,1)$ be a level at which (f, M) is alignment-regular (Definition 3) and set $\alpha^*(C,M) := 1/(1-\kappa) < \infty$: Lemma 6 supplies, for every selling cell $\ell \geq 1$, an α-independent coupling γ'_ℓ of μ_ℓ^+, μ_ℓ^- satisfying its parts (a)–(b). On the exclusion cell — which contains the origin, since at $x = 0$ every option yields $-p_\ell < 0$ (standing positive prices) — inclusivity makes μ_0^+ the origin atom, and global balance gives it equal mass with μ_0^-, so a coupling γ'_0 exists (e.g. $\delta_0 \otimes \mu_0^-/\mu_0^-(C_0)$). Fix $\alpha \geq \alpha^*(C,M)$. Lemma 5 assembles $\gamma' := \sum_{\ell\geq 0} \gamma'_\ell$, which transports μ^+ to μ^- (i), is supported on within-cell pairs (ii), and satisfies the core-only CS identity on selling cells (iii).

Case split for γ'-typical pairs. By (ii) every γ'-typical pair lies in a single cell. Two families carry mass.

(W) Within a selling cell: $y^1, y^2 \in C_\ell$ ($\ell \geq 1$). Lemma 10 (using $\alpha \geq 1/(1-\kappa)$ and the within-cell coupling from Lemma 6) gives $\tilde{c}_\alpha(h) = \tilde{u}_M(y^1) - \tilde{u}_M(y^2)$ a.e.

(E) Within the exclusion cell: $y^1, y^2 \in C_0$. Inclusivity keeps every top face out of C_0, so the only μ^+-mass there is the origin atom; hence $y^1 = 0$ and $h = -y^2$. Both parts of the cost vanish, $\tilde{c}_\alpha(h) = 0$, since $v_B(-y^2) = -v_B(y^2) \leq 0$ for every bundle B; this matches $\tilde{u}_M(y^1) - \tilde{u}_M(y^2) = 0 - 0 = 0$, so complementary slackness holds trivially.

Conclusion. Every γ'-typical pair lies within a single cell, and full-cost complementary slackness holds there — by Lemma 10 on selling cells (W) and trivially on the exclusion cell (E); cross-cell pairs are γ'-null by construction (ii). The marginal identity $\gamma'_1 - \gamma'_2 = \mu$ holds by construction (i). All three conditions of Proposition 2 therefore hold — within-cell support (i), the full-cost CS identity (ii), and

the marginals (iii) — so $(u_M, v_\# \gamma')$ is a primal–dual saddle point and M is optimal in the unrestricted problem. *Q.E.D.*

Online Supplement

S.1 The single-cell certificate

This section constructs the *single-cell couplings* that the applications of Section 5 convert into alignment regularity through the certification direction of Lemma 9. The setting is menu-free: a single core block carrying a density, and a single level q playing the role of a crossing point. Throughout, C is a nonempty index set, $D_C := \prod_{i\in C}[0,\bar v_i]$, $t(\omega) := \sum_{i\in C}\omega_i$ is the core sum with corner value $\bar t_C := \sum_{i\in C}\bar v_i$ and lowest top $\underline{t}_C := \min_{i\in C}\bar v_i$, and $F_c^C := \{\omega \in D_C : \omega_c = \bar v_c\}$ are the core top faces.

Definition 7 (Admissible core density)**.** *A density ρ on D_C is* admissible *if:* (i) $\rho \in C^1(\operatorname{int} D_C)$, $\rho > 0$ *on* $\operatorname{int} D_C$*, and ρ together with each $\omega_c\,\partial_c\rho$ ($c \in C$) is integrable on D_C — hence each $\partial_c(\omega_c\rho)$ and $\Phi_\rho\,\rho := \sum_c \partial_c(\omega_c\rho)$ are integrable, where $\Phi_\rho(\omega) := |C| + \omega\cdot\nabla\log\rho(\omega)$;* (ii) *$\rho$ extends continuously to D_C off the lower faces $\{\omega_c = 0\}$, with integrable top-face traces, and $\omega_c\,\rho(\omega) \to 0$ as $\omega_c \downarrow 0$ for each $c \in C$;* (iii) *the law of t under $\rho\,d\omega$ has a continuous, strictly positive density γ on $(0,\bar t_C)$ with $q\,\gamma(q) \to 0$ as $q \downarrow 0$, and for $|C| \geq 2$ each top-face trace law of t is atomless;* (iv) *ρ is continuous and strictly positive on a neighborhood of the top corner $(\bar v_c)_{c\in C}$, and $\Phi_\rho\,\rho$ is bounded on that neighborhood.*

The standing baseline at $C = [N]$ ($\rho = f$, with $f \in C^2(D)$, $f > 0$; the standing form) is admissible — every clause holds with room to spare on a compact box — and so is the iid Beta$(\beta,1)$ family for every $\beta > 0$: there $\Phi_\rho \equiv |C|\beta$ is constant, $\omega_c\partial_c\rho = (\beta-1)\rho \in L^1$, the singularities sit only on the lower faces, the core-sum density is continuous and positive on $(0,\bar t_C)$ with $q\,\gamma(q) \asymp q^{|C|\beta} \to 0$ (continuity and positivity follow by induction on $|C|$: convolving with $\beta x^{\beta-1}$ preserves both, by dominated convergence across the integrable endpoint singularity), and the face trace laws are absolutely continuous. Independent core blocks with admissible structure are instantiated where used (Section 5.2). Write $\bar F(r) := \int_{\{t\geq r\}}\rho\,d\omega$ for the (unnormalized) survivor of the core sum, $h(r) := \gamma(r)/\bar F(r)$ for its hazard on $(0,\bar t_C)$, and

$$H(q) := q\,h(q) = \frac{q\,\gamma(q)}{\bar F(q)}, \qquad H(0) := 0,$$

continuous on $[0,\bar t_C)$ by Definition 7(iii). The objects of this section, for a level

$q \in [0, \underline{t}_C)$, are the *certificate measures*

$$\nu_q^+ := \sum_{c \in C} \bar{v}_c \, \rho\big|_{F_c^C} \quad \text{on } F_c^C \cap \{t \geq q\}, \qquad \nu_q^- := \big(\Phi_\rho + H(q)\big)\rho \, d\omega \quad \text{on } \Omega_q := \{\omega \in D_C : t(\omega) \geq q\},$$

and the *core wedge* at level κ, $W_\kappa := \{h \in \mathbb{R}^C : \sum_i h_i^- \leq \kappa \sum_i h_i^+\}$.

Lemma 11 (Divergence identity)**.** *For admissible ρ and every $r \in (0, \bar{t}_C)$,*

$$\int_{\{t \geq r\}} \Phi_\rho \, \rho \, d\omega \;=\; B(r) - r\,\gamma(r), \qquad B(r) := \sum_{c \in C} \int_{F_c^C \cap \{t \geq r\}} \bar{v}_c \, \rho \, d\sigma_c,$$

with σ_c the surface (Lebesgue) measure on F_c^C; at $r = 0$ the identity holds with $r\gamma(r)$ replaced by 0. In particular, taking $r = q$ and adding $H(q)\bar{F}(q) = q\,\gamma(q)$: the certificate measures have equal total mass, $\nu_q^-(\Omega_q) = B(q) = \nu_q^+(\Omega_q)$.

Proof. Coordinatewise integration by parts. Fix $c \in C$ and ω_{-c} with $t_{-c}(\omega_{-c}) := \sum_{i \neq c} \omega_i$; the fiber $\{\omega_c : (\omega_c, \omega_{-c}) \in \Omega_r\}$ is $[a_c, \bar{v}_c]$ with $a_c := \max\{0,\, r - t_{-c}(\omega_{-c})\}$ (empty if $a_c > \bar{v}_c$). On a nonempty fiber, $\omega_c \mapsto \omega_c \rho$ is C^1 on the interior with endpoint limits given by Definition 7(ii), so

$$\int_{a_c}^{\bar{v}_c} \partial_c\big(\omega_c \, \rho\big)\, d\omega_c \;=\; \bar{v}_c \, \rho(\bar{v}_c, \omega_{-c}) \;-\; a_c \, \rho(a_c, \omega_{-c}),$$

the lower term vanishing when $a_c = 0$ by (ii). Integrating over ω_{-c} and summing over c: the left side is $\int_{\Omega_r} \sum_c \partial_c(\omega_c \rho) = \int_{\Omega_r} \Phi_\rho \, \rho$; the first right-hand terms sum to $B(r)$ (the interchange of sum and integral is licensed by Definition 7(i), each $\partial_c(\omega_c \rho) \in L^1$); and the subtracted terms sum, by Fubini on the graph parametrizations of the cut $\{t = r\}$, to $\int\big(\sum_c \omega_c\big)\rho$ over the cut, i.e. $r\,\gamma(r)$ (each c-term integrates $\omega_c \rho$ over the same cut, and $\gamma(r)$ is the cut integral of ρ in the graph variables). The Fubini identity holds for a.e. r a priori; both sides are continuous in r on $(0, \bar{t}_C)$ (γ by Definition 7(iii), B by the atomless face trace laws — trivially constant for $|C| = 1$ — and the left side because $\Phi_\rho \rho \, d\omega$ is absolutely continuous), so it holds at every r. The mass identity follows since $\nu_q^-(\Omega_q) = \int_{\Omega_q} \Phi_\rho \rho + H(q)\bar{F}(q) = B(q) - q\gamma(q) + q\gamma(q)$. *Q.E.D.*

Lemma 12 (Gap identity and regularity)**.** *For admissible ρ, $q \in [0, \underline{t}_C)$, and $r \in [q, \bar{t}_C)$, define the tails $I_q(r) := \nu_q^-(\{t \geq r\})$ and $B(r) := \nu_q^+(\{t \geq r\})$ (consistent with Lemma 11 since all face mass has $t \geq \underline{t}_C > q$). Then*

$$\Delta_q(r) \;:=\; B(r) - I_q(r) \;=\; r\,\gamma(r) - H(q)\,\bar{F}(r) \;=\; \bar{F}(r)\big[H(r) - H(q)\big] \quad (q \leq r < \bar{t}_C).$$

If ρ is regular above q — $H(r) > H(q)$ *for every* $r \in (q, \bar{t}_C)$ — *then* $\Delta_q \geq 0$ *on* $[q, \bar{t}_C)$, *strictly positive on* $(q, \bar{t}_C)$, *with a strictly positive minimum on every compact subinterval of* $(q, \bar{t}_C)$.

Proof. By Lemma 11, $I_q(r) = \big[B(r) - r\gamma(r)\big] + H(q)\bar{F}(r)$, giving the identity. Under regularity above q, $\Delta_q(r) = \bar{F}(r)[H(r) - H(q)] > 0$ for $r \in (q, \bar{t}_C)$ since $\bar{F}(r) > 0$ there; at $r = q$ it vanishes. Δ_q is continuous on $[q, \bar{t}_C)$ (γ and $\bar{F}$ continuous by Definition 7(iii)), so on a compact subinterval of $(q, \bar{t}_C)$ the strictly positive continuous function attains a strictly positive minimum. *Q.E.D.*

Lemma 13 (Shifted-tail inequality)**.** *Let* $|C| \geq 2$, *let* ρ *be admissible and regular above* $q \in [0, \underline{t}_C)$ *with* $\Phi_\rho + H(q) \geq 0$ *on* Ω_q. *Then there is* $\lambda \in (0, 1/2]$ *such that*

$$I_q(r) \;\leq\; B\big(T_\lambda(r)\big) \qquad \textit{for every } r \in [q, \bar{t}_C], \qquad T_\lambda(r) := (1-\lambda)\, r + \lambda\, \bar{t}_C.$$

Proof. Three bands, split at $q_0 := (q + \underline{t}_C)/2 \in (q, \underline{t}_C)$ and at a corner cutoff $\bar{t}_C - r_0$.

Corner band. By Definition 7(iv) pick a corner radius $u_0 \in (0, \min_i \bar{v}_i]$ and constants $0 < \underline{\rho} \leq \bar{\rho}_{\mathrm{loc}} < \infty$ with $\underline{\rho} \leq \rho \leq \bar{\rho}_{\mathrm{loc}}$ and $\Phi_\rho \rho \leq \bar{c}$ on $D_C \cap \prod_i [\bar{v}_i - u_0, \bar{v}_i]$. For $s \in (0, u_0]$, so that $s \leq \min_i \bar{v}_i$, the region $\{t \geq \bar{t}_C - s\}$ is the full corner simplex $\{\sum_i (\bar{v}_i - \omega_i) \leq s\}$, of volume $s^{|C|}/|C|!$, so

$$I_q(\bar{t}_C - s) \;\leq\; \big(\bar{c} + H(q)\,\bar{\rho}_{\mathrm{loc}}\big)\, \frac{s^{|C|}}{|C|!};$$

on each face F_c^C the set $\{t \geq \bar{t}_C - u\}$ is the face corner simplex of volume $u^{|C|-1}/(|C|-1)!$, so for any $\lambda \leq 1/2$, $T_\lambda(\bar{t}_C - s) = \bar{t}_C - (1-\lambda)s \leq \bar{t}_C - s/2$ and

$$B\big(T_\lambda(\bar{t}_C - s)\big) \;\geq\; B(\bar{t}_C - s/2) \;\geq\; \Big(\sum_c \bar{v}_c\Big)\, \underline{\rho}\, \frac{(s/2)^{|C|-1}}{(|C|-1)!}.$$

The ratio of the bounds is $O(s)$; choose $r_0 \in (0, u_0]$ so that $I_q(\bar{t}_C - s) \leq B(T_\lambda(\bar{t}_C - s))$ for all $s \leq r_0$ and all $\lambda \leq 1/2$.

Low band $[q, q_0]$. Set $\lambda_{\mathrm{low}} := (\underline{t}_C - q_0)/(\bar{t}_C - q_0) > 0$. For $\lambda \leq \lambda_{\mathrm{low}}$ and $r \leq q_0$, $T_\lambda(r) \leq T_\lambda(q_0) \leq \underline{t}_C$, where B is constant at its total value (every face point has $t \geq \underline{t}_C$): $B(T_\lambda(r)) = B(q) = I_q(q) \geq I_q(r)$, by the mass identity of Lemma 11 and monotonicity of I_q (sinks nonnegative).

Middle band $[q_0, \bar{t}_C - r_0]$ (empty if $q_0 > \bar{t}_C - r_0$). By Lemma 12 the gap has a strictly positive minimum m on the band. For $|C| \geq 2$ the face trace laws of t are

atomless (Definition 7(iii)), so B is continuous, hence uniformly continuous on the compact $[0, \bar{t}_C]$; pick $\delta > 0$ with $B(r) - B(r') \le m/2$ whenever $0 \le r' - r \le \delta$, and set $\lambda_{\text{mid}} := \delta/(\bar{t}_C - q_0)$. For $\lambda \le \lambda_{\text{mid}}$ and r in the band, $T_\lambda(r) - r = \lambda(\bar{t}_C - r) \le \delta$, so

$$B\big(T_\lambda(r)\big) \;\ge\; B(r) - \tfrac{m}{2} \;=\; I_q(r) + \Delta_q(r) - \tfrac{m}{2} \;\ge\; I_q(r) + \tfrac{m}{2} \;>\; I_q(r).$$

Take $\lambda := \min\{\lambda_{\text{low}}, \lambda_{\text{mid}}, 1/2\}$ (omitting λ_{mid} when the middle band is empty); the bands cover $[q, \bar{t}_C]$. *Q.E.D.*

Theorem 8 (Single-cell certificate)**.** *Let ρ be admissible and let $q \in [0, \underline{t}_C)$ satisfy: ρ is regular above q, and $\Phi_\rho + H(q) \ge 0$ on Ω_q. Then there exists $\kappa_{\text{pb}}(\rho, q) \in [0, 1)$ such that for every $\kappa \ge \kappa_{\text{pb}}(\rho, q)$ there is a coupling of (ν_q^+, ν_q^-) concentrated on $\{(y, x) : y - x \in W_\kappa\}$. Quantitatively, $\kappa_{\text{pb}}(\rho, q) \le 1 - \lambda$ with λ from Lemma 13; and $\kappa_{\text{pb}}(\rho, q) = 0$ when $|C| = 1$.*

Proof. *Case $|C| = 1$.* All source mass sits at the single top point $\bar{v}$, all sinks lie below it, and the masses agree by Lemma 11; the product coupling $\delta_{\bar{v}} \otimes \nu_q^-/\nu_q^-(\Omega_q)$ has displacements $h = \bar{v} - x \ge 0$, which lie in $W_0 \subseteq W_\kappa$ for every $\kappa \ge 0$.

Case $|C| \ge 2$. Let λ be as in Lemma 13 and define the closed relation $G := \{(y, x) \in \overline{\Omega}_q \times \overline{\Omega}_q : t(y) \ge T_\lambda(t(x))\}$. For closed $A \subseteq \overline{\Omega}_q$ with $r_A := \min_A t \ge q$,

$$\nu_q^-(A) \;\le\; \nu_q^-\big(\{t \ge r_A\}\big) \;=\; I_q(r_A) \;\le\; B\big(T_\lambda(r_A)\big) \;=\; \nu_q^+\big(\{t \ge T_\lambda(r_A)\}\big) \;=\; \nu_q^+\big(G^{-1}(A)\big),$$

using Lemma 13 and $G^{-1}(A) = \{y : t(y) \ge T_\lambda(r_A)\}$ (T_λ is increasing). The measures have equal total mass (Lemma 11), so Theorem 7 (after normalization) yields a coupling γ of (ν_q^+, ν_q^-) concentrated on G.

Wedge membership. For $(y, x) \in G$ write $h := y - x$ and $\varepsilon := \bar{t}_C - t(x) \ge 0$. Then $t(h) = t(y) - t(x) \ge \lambda\varepsilon$ and

$$\sum_i h_i^- \;\le\; \sum_i \big(\bar{v}_i - y_i\big) \;=\; \bar{t}_C - t(y) \;\le\; (1 - \lambda)\,\varepsilon,$$

while $\sum_i h_i^+ = t(h) + \sum_i h_i^- \ge \lambda\varepsilon$. Since $x \mapsto x/(c + x)$ is increasing for $c > 0$,

$$\frac{\sum_i h_i^-}{\sum_i h_i^+} \;\le\; \frac{(1 - \lambda)\varepsilon}{\lambda\varepsilon + (1 - \lambda)\varepsilon} \;=\; 1 - \lambda \qquad (\varepsilon > 0),$$

and $\varepsilon = 0$ forces x at the corner, $y = x$, $h = 0 \in W_\kappa$. Hence γ is concentrated on $W_{1-\lambda}$-displacements, and on W_κ for every $\kappa \ge 1 - \lambda$ (W_κ is nondecreasing in κ).

Define $\kappa_{\mathrm{pb}}(\rho, q)$ as the infimum of levels admitting a single-cell coupling. The set of such levels is an up-set (a coupling at κ works verbatim at $\kappa' \geq \kappa$, W_κ being nondecreasing), and it contains its infimum: take levels $\kappa_n \downarrow \kappa_{\mathrm{pb}}$ with couplings π_n on the fixed compact square $\overline{\Omega}_q \times \overline{\Omega}_q$; any weak-$*$ limit π has the same marginals, and since π_n gives full mass to the closed set $\{(y, x) : \sum_i (y_i - x_i)^- \leq \kappa_m \sum_i (y_i - x_i)^+\}$ for every $n \geq m$, Portmanteau gives π full mass on it for every m, hence on their intersection, which is the $W_{\kappa_{\mathrm{pb}}}$ displacement set. So couplings exist at every $\kappa \geq \kappa_{\mathrm{pb}}$, and $\kappa_{\mathrm{pb}} \leq 1 - \lambda < 1$. *Q.E.D.*

Proposition 3 (Uniformity over compact level ranges)**.** *Let ρ be admissible with $\Phi_\rho \geq 0$ on D_C and H strictly increasing on $(0, \bar{t}_C)$. Then for every $\bar{q} < \underline{t}_C$,*

$$\sup_{q \in [0, \bar{q}]} \kappa_{\mathrm{pb}}(\rho, q) \;<\; 1.$$

Proof. For $|C| = 1$, $\kappa_{\mathrm{pb}} \equiv 0$. Let $|C| \geq 2$. The hypotheses of Theorem 8 hold at every $q \in [0, \bar{q}]$: $\Phi_\rho + H(q) \geq \Phi_\rho \geq 0$, and regularity above q follows from strict monotonicity of H. It suffices to produce a single $\lambda > 0$ valid in Lemma 13 for all $q \in [0, \bar{q}]$ simultaneously; then $\sup_q \kappa_{\mathrm{pb}} \leq 1 - \lambda$.

Corner band: the corner constants $u_0, \underline{\rho}, \bar{\rho}_{\mathrm{loc}}, \bar{c}$ are q-free, and $H(q) \leq H(\bar{q})$ for $q \leq \bar{q}$, so a single r_0 works for all $q \in [0, \bar{q}]$.

Low band: $q_0(q) = (q + \underline{t}_C)/2$ gives $\lambda_{\mathrm{low}}(q) = (\underline{t}_C - q_0(q))/(\bar{t}_C - q_0(q))$, which is decreasing in q; its value at $\bar{q}$ is a positive uniform choice.

Middle band: the parameter set $\Lambda := \{(q, r) : q \in [0, \bar{q}],\ q_0(q) \leq r \leq \bar{t}_C - r_0\}$ is compact, and on it $(q, r) \mapsto \Delta_q(r) = \bar{F}(r)\big[H(r) - H(q)\big]$ is continuous (H continuous on $[0, \bar{t}_C)$ by Definition 7(iii)) and strictly positive ($r \geq q_0(q) > q$ and H strictly increasing; $\bar{F}(r) \geq \bar{F}(\bar{t}_C - r_0) > 0$). Hence $m := \min_\Lambda \Delta_q(r) > 0$ (if $\Lambda = \varnothing$ the middle band is empty for every q and this step is vacuous), and the uniform-continuity choice of δ in Lemma 13, together with $\lambda_{\mathrm{mid}} := \delta/(\bar{t}_C - q_0(0))$ — calibrated to the widest band, attained at $q = 0$ — is valid for every $q \in [0, \bar{q}]$ simultaneously, since $T_\lambda(r) - r \leq \lambda\,(\bar{t}_C - q_0(q)) \leq \lambda\,(\bar{t}_C - q_0(0)) = \delta$ on each band.

Take $\lambda := \min\{\lambda_{\mathrm{low}}(\bar{q}), \lambda_{\mathrm{mid}}, 1/2\}$. *Q.E.D.*

Lemma 14 (Extension to admissible singular densities)**.** *Let $C = [N]$ and let f be an admissible core density on D (Definition 7) with $\Phi = \Phi_\rho + 1 > 0$ on $\operatorname{int} D$. Then the boundary–interior decomposition $\mu = \delta_0 + \sum_j \bar{v}_j\, f|_{F_j} - \Phi f \cdot \lambda_D$ of Section 4.1 and the resulting revenue representation hold for f, and every result of the sufficiency chain*

— Propositions 1 and 2, Section B, the present section, Theorem 2, and the sufficiency direction of Theorem 5 — holds verbatim with the C^2/strict-positivity of the standing form replaced by admissibility. The necessity theorems — including the suboptimality (non-inclusive) direction of Theorem 5, which is the pure-bundling case of Theorem 4 — retain their own stated hypotheses.

Proof. Both parts of μ are finite measures: the top-face traces are integrable by Definition 7(ii), and $\Phi f = \sum_j \partial_j(x_j f) + f \in L^1$ by (i). For the decomposition, integrate $\int (x\cdot\nabla U)\, f$ by parts on the shrunken box $D_\varepsilon := \prod_j [\varepsilon, \bar{v}_j]$, where $f \in C^1$, against any Lipschitz U (the indirect utilities of Lemma 2 qualify):

$$\int_{D_\varepsilon} (x\cdot\nabla U)\, f \;=\; \sum_j \bar{v}_j \int_{F_j\cap D_\varepsilon} U f \;-\; \varepsilon \sum_j \int_{\{x_j=\varepsilon\}\cap D_\varepsilon} U f \;-\; \int_{D_\varepsilon} U\,(x\cdot\nabla f + N f).$$

As $\varepsilon \downarrow 0$: the top-face terms converge to the full face integrals (integrable traces); the interior term converges by dominated convergence (U bounded, $x\cdot\nabla f + Nf = \Phi_\rho f \in L^1$); and each lower-face term vanishes because, for a.e. x_{-j}, the map $x_j \mapsto x_j f$ is absolutely continuous on $(0, \bar{v}_j]$ with limit 0 at 0 (Definition 7(i)–(ii)), so $\varepsilon \int_{\{x_j=\varepsilon\}} f \le \int_{\{x_j\le\varepsilon\}} |\partial_j(x_j f)| \to 0$ by absolute continuity of the integral. The limit reads $\int (x\cdot\nabla U)\, f = \sum_j \bar{v}_j \int_{F_j} Uf - \int U\,\Phi_\rho f$. Lemma 2 writes revenue as $\int (a\cdot\nabla U - U)\, f$ with the tangency field $a(x) = x$; the extra $-\int Uf$ turns Φ_ρ into $\Phi = \Phi_\rho + 1$, so revenue $= U(0)\cdot 1 + \sum_j \bar{v}_j \int_{F_j} Uf - \int U\,\Phi f$, i.e. $\int U\,\mathrm{d}\mu$ with the stated μ (the atom from the $U(0)$ normalization term). This is the same μ and the same revenue representation as in the baseline case. Every downstream argument of the sufficiency chain consumes only the finite measures $\mu^{\pm}$, the cell geometry, and the admissibility clauses — the saddle-point criterion and the couplings of Lemma 6 and the present section never evaluate f pointwise on the closed box — so each proof applies verbatim. *Q.E.D.*

S.2 Proofs for the applications

S.2.1 Pure bundling

This subsection proves Theorem 5 and records the converse pinch. Here $C = [N]$, $P = \varnothing$, the menu is the single option $([N], p^*)$ with crossing point $\hat{s} = p^*/\alpha$, the selling cell is $C_1 = \{x : t(x) \ge \hat{s}\}$ with $t(x) = \sum_i x_i$ (up to the tie convention), and $\rho := f$ is an admissible core density at $C = [N]$ (Definition 7; the standing form qualifies,

as does the Beta family of Section 5.1, via Lemma 14). Write $\Phi_\rho = N + x \cdot \nabla \log f$, so the standing radial score is $\Phi = \Phi_\rho + 1$, and let $\gamma, \bar{F}, H$ be the core-sum objects of Section S.1 — here the *bundle-value* density, survivor, and virtual-value statistic. *Strict Myersonian regularity* of the bundle-value distribution means $H(r) > 1$ for every $r \in (\hat{s}, \bar{t}_C)$.

Proposition 4 (Pure bundling is alignment-regular)**.** *In the setting above, suppose the single selling cell is mass-balanced, the price is strictly inclusive ($\hat{s} < \min_i \bar{v}_i$), the radial score is positive ($\Phi = \Phi_\rho + 1 > 0$ on C_1), and the bundle-value distribution is strictly Myersonian regular. Then:* (i) *mass balance is equivalent to $H(\hat{s}) = 1$;* (ii) *(f, M) is alignment-regular at level $\kappa_{\mathrm{pb}}(f, \hat{s}) < 1$ from Theorem 8;* (iii) *Theorem 2 applies, and pure bundling at $p^* = \alpha\hat{s}$ is unrestricted-optimal for every $\alpha \geq \alpha^*_{\mathrm{pure}} := 1/\big(1-\kappa_{\mathrm{pb}}(f, \hat{s})\big)$.*

Proof. *(i)* By inclusivity every top face lies inside the selling cell ($t \geq \bar{v}_j > \hat{s}$ on F_j) and the origin atom lies in C_0 (standing positive prices), so $\mu_1^+ = \sum_j \bar{v}_j f|_{F_j} = \nu_{\hat{s}}^+$ with total mass $B(\hat{s})$. The interior part is $\mu_1^- = \Phi f = (\Phi_\rho + 1) f$ on $\{t \geq \hat{s}\}$, so by the divergence identity (Lemma 11), $\mu_1^-(C_1) = B(\hat{s}) - \hat{s}\,\gamma(\hat{s}) + \bar{F}(\hat{s})$. Mass balance $\mu_1^+(C_1) = \mu_1^-(C_1)$ is therefore $\hat{s}\,\gamma(\hat{s}) = \bar{F}(\hat{s})$, i.e. $H(\hat{s}) = 1$.

(ii) With $H(\hat{s}) = 1$, the certificate measures of Section S.1 at level $q = \hat{s}$ are the cell parts themselves: $\nu_{\hat{s}}^+ = \mu_1^+$ and $\nu_{\hat{s}}^- = (\Phi_\rho + H(\hat{s}))\, f = \Phi f = \mu_1^-$ on C_1. Strict Myersonian regularity says $H(r) > 1 = H(\hat{s})$ on $(\hat{s}, \bar{t}_C)$ — exactly regularity above $\hat{s}$ — the sink density is $\Phi_\rho + H(\hat{s}) = \Phi > 0$, and inclusivity places $\hat{s} \in [0, \underline{t}_C)$. Theorem 8 yields, for every $\kappa \geq \kappa_{\mathrm{pb}}(f, \hat{s})$, a coupling of (μ_1^+, μ_1^-) concentrated on wedge displacements. At $P = \varnothing$ the budget relation *is* the wedge, $K_{1,\kappa} = W_\kappa$ (Definition 6: no peripheral clauses remain, and the budget clause is the wedge inequality), so the certification direction of Lemma 9 converts the coupling into $\nu_{1,\kappa}^- \preceq_{\mathrm{FOSD}} \nu_{1,\kappa}^+$: (f, M) is alignment-regular at level $\kappa_{\mathrm{pb}}(f, \hat{s}) < 1$.

(iii) The menu is finite, mass-balanced (its single price is the restricted optimum on $\mathcal{F}_{[N]}$, with mass balance its first-order condition by Lemma 3), inclusive, and alignment-regular, so Theorem 2 applies; tracking the threshold through Lemma 10 and the assembly gives optimality at every $\alpha \geq 1/(1-\kappa)$ for any valid level κ, in particular at every $\alpha \geq \alpha^*_{\mathrm{pure}}$. *Q.E.D.*

Remark 3 (The converse pinch: weak regularity is necessary). If the mass-balanced, inclusive pure-bundling menu is alignment-regular at *any* level $\kappa \in [0, 1)$, then the bundle-value distribution is weakly Myersonian regular on the active range: $H(r) \geq$

1 for $r \in [\hat{s}, \bar{t}_C)$. Indeed, at $C = [N]$ the score σ_C is the bundle value t itself, so for $r \geq \hat{s}$ the indicator $\mathbf{1}\{u_{\sigma_C} \geq r\}$ is a bounded coordinatewise nondecreasing function of the scores, and dominance gives $\mu_1^-(\{t \geq r\}) \leq \mu_1^+(\{t \geq r\})$ — by Lemma 12 and part (i), $\bar{F}(r)\,[H(r) - 1] = \Delta_{\hat{s}}(r) \geq 0$. Alignment regularity of the pure-bundling menu is thus pinched between the weak and strict forms of the classical condition.

S.2.2 Two-tier menus

This subsection proves Theorem 6. Throughout, the two-tier setting (S1)–(S3) is in force on the unit box $D = [0,1]^N$, with core C, $|C| = N - 1 \geq 1$, single peripheral good p, and types $x = (\omega, z)$. Option 1 is the bare-core tier $(C, \alpha\hat{s}_C)$ with $w_1 = \mathbf{1}_C$, option 2 the grand tier $(C \cup \{p\}, \alpha\hat{s}_G)$ with $w_2 = \mathbf{1}_{C\cup\{p\}}$; both options are deterministic and contain the core, so the menu lives on $\mathcal{F}_C$ and the weight classes of Section 4.1 are

$$S_1^{(1)} = \varnothing, \quad Z_1 = \{p\}, \qquad S_2^{(1)} = \{p\}, \quad Z_2 = \varnothing, \qquad B_1^{\mathrm{frac}} = B_2^{\mathrm{frac}} = \varnothing,$$

so the budget relations of Definition 6 read

$$K_{1,\kappa} = \Big\{h : h_p \leq 0,\ \sum_{i\in C} h_i^- \leq \kappa \sum_{i\in C} h_i^+\Big\}, \qquad K_{2,\kappa} = \Big\{h : h_p \geq 0,\ \sum_{i\in C} h_i^- \leq \kappa\Big(h_p + \sum_{i\in C} h_i^+\Big)\Big\}.$$

The core-sum objects are those of Section S.1 at $\rho = f_C$ (admissibility is Lemma 15 below): g_C, G_C, $\bar{G}_C$ are the density, cdf, and survivor of $t(\omega) = \sum_{i\in C} \omega_i$ under f_C, with $\bar{t}_C = |C|$, lowest top $\underline{t}_C = 1$, statistic $H_C(q) = q\, g_C(q)/\bar{G}_C(q)$, face-mass total $\Gamma_C := \sum_{i\in C} f_i(1)$, and certificate measures $\nu_q^{\pm}$ on $\Omega_q = \{t \geq q\}$; since every core face has $t \geq 1$, the faces are uncut at every level $q < 1$, so $\nu_q^+ = \nu_0^+ =: \nu^+$ is the full core-face measure with total $B(q) = \Gamma_C$ for every $q \in [0,1)$. Write c_t for the conditional law of ω given $t(\omega) = t$ under f_C (the canonical fiber disintegration, $f_C\,\mathrm{d}\omega = \int c_t(\cdot)\, g_C(t)\,\mathrm{d}t$).

Virtual density. The radial score splits along the product:

$$\Phi(x) \;=\; N + 1 + x \cdot \nabla \log f(x) \;=\; \underbrace{\big(|C| + \omega \cdot \nabla \log f_C(\omega)\big)}_{\Phi_C(\omega)} \;+\; \underbrace{\big(2 + z\,(\log f_p)'(z)\big)}_{\Phi_p \equiv 1+\beta},$$

using $z(\log f_p)'(z) = \beta - 1$ from (S2). So $\mu^- = (\Phi_C + 1 + \beta)\, f$, a genuine finite measure by Lemma 15(i), bounded below by $(1+\beta)f$; the boundary part is $\mu^+ = \delta_0 + \sum_{i\in C} f|_{F_i} + f|_{F_p}$ with $F_i = \{\omega_i = 1\}$ and $F_p = \{z = 1\}$. The peripheral face carries total mass $\mu^+(F_p) = f_p(1) \int f_C\,\mathrm{d}\omega = \beta$, and write $\mu_p^+ := \mu^+|_{F_p} = \beta\, f_C(\omega)\,\mathrm{d}\omega \otimes \delta_{z=1}$.

Geometry and inclusivity. In normalized units the option utilities are $\tilde{u}_1 = t - \hat{s}_C$ and $\tilde{u}_2 = t + z - \hat{s}_G$, so $\tilde{u}_2 \geq \tilde{u}_1 \iff z \geq d$ and the cells are

$$C_1 = \{\, t \geq \hat{s}_C,\ z \leq d \,\}, \qquad C_2 = \{\, z \geq d,\ t \geq \tau(z) \,\}, \qquad \tau(z) = (\hat{s}_G - z)^+,$$

up to the tie convention; the tie sets are Lebesgue-null, hence μ^--null, and carry no top-face mass (each face measure is $(N-1)$-dimensional, and the tie sets cut it in codimension one). Every core face satisfies $t \geq 1 > \hat{s}_C$, so $F_i \cap \{z \leq d\} \subseteq C_1$ and $F_i \cap \{z \geq d\} \subseteq C_2$ (there $t \geq 1 > \hat{s}_C = \tau(d) \geq \tau(z)$); the peripheral face F_p lies strictly inside C_2's win region, since at $z = 1$ both $\tilde{u}_2 = t + 1 - \hat{s}_G > 0$ and $\tilde{u}_2 - \tilde{u}_1 = 1 - d > 0$; and the origin lies in the exclusion cell. Hence no top face meets C_0: *the margins of* (S3) *imply inclusivity* (Definition 2). Write $\mu_\ell^\pm := \mu^\pm|_{C_\ell}$; so μ_1^+ is the core-face mass over $z \leq d$, with total $\Gamma_C F_p(d) = \Gamma_C\, d^\beta$, and $\mu_2^+ = \mu_{2,C}^+ + \mu_p^+$, where $\mu_{2,C}^+$ is the core-face mass over $z \geq d$.

The proof runs through seven lemmas. Lemmas 15 and 16 are distributional groundwork; Lemma 17 converts mass balance into two price identities; Lemma 18 splits the grand cell's sinks into a stage-W part, served slice by slice at frozen z by the single-cell certificate (Lemma 19), and a stage-G part, served from the peripheral face by an Efron lift of a sheared two-dimensional coupling (Lemmas 20 and 21); the assembly closes the argument through the certification direction of Lemma 9 and Theorem 2.

Lemma 15 (Admissibility and the core statistic)**.** *Under* (S1)*:* (i) f_C *is an admissible core density (Definition 7) with* $\Phi_\rho = \Phi_C$*; in particular* $\Phi_C f_C \in L^1(D_C)$*;* (ii) g_C *is log-concave, continuous, bounded, and strictly positive on* $(0, \bar{t}_C)$*;* (iii) H_C *is continuous on* $[0, \bar{t}_C)$ *with* $H_C(0) = 0$*, and strictly increasing on* $(0, \bar{t}_C)$*.*

Proof. Two facts about each marginal are used repeatedly. First, f_i is *bounded*: writing $f_i = e^{-V_i}$ with V_i convex on $(0,1)$, V_i lies above one of its tangent lines, and an affine function on a bounded interval is bounded, so V_i is bounded below and $f_i \leq e^{-\inf V_i} < \infty$. Second, f_i has *bounded variation*: a log-concave density is unimodal ($-V_i$ quasiconcave), so on $(0,1)$ it increases and then decreases, giving $\mathrm{TV}(f_i) \leq 2 \sup f_i < \infty$; since $f_i \in C^1((0,1])$, this reads $\int_0^1 |f_i'|\, \mathrm{d}\omega_i < \infty$.

(i), clause by clause. (i): $f_C \in C^1$ and > 0 on $\operatorname{int} D_C$ as a product of C^1 positive factors; and $\omega_c\, \partial_c f_C = \big(\omega_c f_c'(\omega_c)\big) \prod_{i \neq c} f_i(\omega_i)$ is integrable by Fubini, since $\int_0^1 |\omega_c f_c'| \leq \int_0^1 |f_c'| < \infty$ and the other factors integrate to one. (ii): each f_i is continuous on $(0,1]$,

so f_C extends continuously to D_C off the lower faces; the trace on F_c^C is $f_c(1)\prod_{i\neq c} f_i$, integrable with total $f_c(1)$; and $\omega_c f_C \to 0$ as $\omega_c \downarrow 0$ pointwise in ω_{-c} because f_c is bounded. (iii): g_C is log-concave by Prékopa's theorem (Prékopa, 1973) (marginals and convolutions of log-concave functions are log-concave), hence continuous on the interior of its support; the support is $[0, \bar{t}_C]$ and $g_C > 0$ on $(0, \bar{t}_C)$ by induction on $|C|$ (for $|C| = 1$, $g_C = f_1 > 0$ on $(0, 1)$; for the inductive step, $g_C(q) = \int f_1(s)\, g_{C\setminus 1}(q - s)\, \mathrm{d}s$ has strictly positive integrand on a set of positive length whenever $q \in (0, \bar{t}_C)$). Moreover $g_C \leq \sup f_1 < \infty$ (Young's bound for a convolution against a probability density; for $|C| = 1$, directly), so $q\, g_C(q) \to 0$ as $q \downarrow 0$. For $|C| \geq 2$ the trace law of t on F_c^C is the law of $1 + \sum_{i\neq c} \omega_i$ under an absolutely continuous law, hence atomless. (iv): each f_i is continuous on $(0, 1]$ with $f_i(1) > 0$, so on a corner neighborhood $\prod_i [1-\varepsilon, 1]$ with ε small, f_C is continuous and bounded away from zero; and $\Phi_C f_C = \sum_c \big(f_C + \omega_c f_c' \prod_{i\neq c} f_i\big)$ is bounded there because each f_c' is continuous on $[1 - \varepsilon, 1]$ (the C^1-up-to-the-endpoint clause), hence bounded. This verifies every clause of Definition 7; $\Phi_\rho = |C| + \omega \cdot \nabla \log f_C = \Phi_C$ by definition, and integrability of $\Phi_C f_C = \sum_c \partial_c(\omega_c f_C)$ follows from clause (i).

(ii) was established along the way.

(iii) Continuity on $[0, \bar{t}_C)$ with $H_C(0) = 0$ is Definition 7(iii) as recorded after that definition ($q g_C(q) \to 0$ and $\bar{G}_C(q) \to 1$ as $q \downarrow 0$). For strict monotonicity: g_C log-concave implies $\bar{G}_C$ log-concave, hence the failure rate $h_C := g_C/\bar{G}_C$ is nondecreasing on $(0, \bar{t}_C)$ (Bagnoli and Bergstrom, 2005); it is strictly positive there since $g_C > 0$. Therefore $H_C(q) = q\, h_C(q)$ is a product of a strictly increasing positive factor and a nondecreasing positive factor, hence strictly increasing on $(0, \bar{t}_C)$. *Q.E.D.*

Lemma 16 (IFR superadditivity)**.** *For all $a, b \geq 0$: $\bar{G}_C(a + b) \leq \bar{G}_C(a)\, \bar{G}_C(b)$.*

Proof. If $\bar{G}_C(a + b) = 0$ the claim is trivial. Otherwise all three values are positive, and $\psi := -\log \bar{G}_C$ is convex on $[0, a+b]$ ($\bar{G}_C$ is log-concave by Bagnoli and Bergstrom (2005), the survivor of the log-concave g_C of Lemma 15(ii)) with $\psi(0) = 0$ ($t \geq 0$ almost surely). Convexity through the origin gives $\psi(a) \leq \frac{a}{a+b}\, \psi(a + b)$ and $\psi(b) \leq \frac{b}{a+b}\, \psi(a+b)$; adding, $\psi(a)+\psi(b) \leq \psi(a+b)$, which exponentiates to the claim. *Q.E.D.*

Lemma 17 (Price identities)**.** *Under the two-tier setting (S1)–(S3):* (i) *mass balance on the bare-core cell is equivalent to*

$$H_C(\hat{s}_C) \;=\; 1 + \beta. \tag{7}$$

(ii) *Given* (7), *define* $\varphi(z) := 1+\beta-H_C(\tau(z))$ *on* $[d,1]$. *Then mass balance on the grand cell is equivalent to*

$$\int_d^1 f_p(z)\,\varphi(z)\,\bar{G}_C(\tau(z))\,\mathrm{d}z \;=\; \beta. \tag{8}$$

Proof. Both identities follow from the divergence identity (Lemma 11) at $\rho = f_C$, where $B(q) = \Gamma_C$ for every $q \in [0,1)$ (faces uncut) and $q\,g_C(q) = H_C(q)\,\bar{G}_C(q)$:

$$\int_{\{t\ge q\}} \Phi_C\, f_C\,\mathrm{d}\omega \;=\; \Gamma_C - q\,g_C(q), \qquad q\in[0,\hat{s}_C]. \tag{9}$$

(i) The cell's boundary mass is the core-face mass over $z \le d$ (the peripheral face has $z = 1 > d$, and the origin atom is excluded): $\mu^+(C_1) = \Gamma_C\,F_p(d)$ with $F_p(d) = d^\beta$. The interior mass factorizes over the product:

$$\mu^-(C_1) \;=\; \int_0^d f_p(z)\,\mathrm{d}z \int_{\{t\ge\hat{s}_C\}} \big(\Phi_C+1+\beta\big) f_C\,\mathrm{d}\omega \;=\; F_p(d)\Big[\Gamma_C - \hat{s}_C\,g_C(\hat{s}_C) + (1+\beta)\,\bar{G}_C(\hat{s}_C)\Big]$$

by (9) at $q = \hat{s}_C$. Since $F_p(d) = d^\beta > 0$, mass balance $\mu^+(C_1) = \mu^-(C_1)$ is equivalent to $\hat{s}_C\,g_C(\hat{s}_C) = (1+\beta)\,\bar{G}_C(\hat{s}_C)$, which is (7).

(ii) The grand cell's boundary mass is $\mu^+(C_2) = \Gamma_C\big(1-F_p(d)\big)+\beta$ (core faces over $z \ge d$, plus the peripheral face). Its interior mass, slicing at fixed z and applying (9) at $q = \tau(z) \in [0,\hat{s}_C]$,

$$\mu^-(C_2) \;=\; \int_d^1 f_p(z)\Big[\Gamma_C - \tau(z)\,g_C(\tau(z)) + (1+\beta)\,\bar{G}_C(\tau(z))\Big]\mathrm{d}z.$$

Since $\int_d^1 f_p = 1 - F_p(d)$, the Γ_C terms cancel from mass balance, which becomes $\beta = \int_d^1 f_p\big[(1+\beta)\bar{G}_C(\tau) - \tau g_C(\tau)\big]\mathrm{d}z$. Substituting $(1+\beta)\,\bar{G}_C(\tau) = \varphi\,\bar{G}_C(\tau) + H_C(\tau)\,\bar{G}_C(\tau) = \varphi\,\bar{G}_C(\tau) + \tau\,g_C(\tau)$ turns this into (8). *Q.E.D.*

Lemma 18 (Sink partition and books)**.** *Assume* (7) *and define, on* C_2,

$$\mu_W^- := \big(\Phi_C(\omega) + H_C(\tau(z))\big)\,f, \qquad \mu_G^- := \varphi(z)\,f, \qquad \varphi(z) = 1+\beta-H_C(\tau(z)).$$

Then: (i) $\mu_2^- = \mu_W^- + \mu_G^-$ *with both parts nonnegative, and* $\varphi(d) = 0$; (ii) *the stage-W books close slice by slice: for every* $z \in [d,1]$, *with* $q = \tau(z) \in [0,\hat{s}_C]$,

$$\int_{\{t\ge q\}} \big(\Phi_C + H_C(q)\big)\,f_C\,\mathrm{d}\omega \;=\; \Gamma_C,$$

the slice's core-face mass; (iii) *the bare-core cell needs no partition:* $\mu_1^- = (\Phi_C + H_C(\hat{s}_C))\,f$ *on* $\{t \ge \hat{s}_C\}\times[0,d]$, *with the slice books of* (ii) *at* $q = \hat{s}_C$; (iv) *the stage-G books close globally:* $\mu_G^-(C_2) = \beta = \mu^+(F_p)$ *whenever* (8) *holds.*

Proof. *(i)* The sum telescopes by construction since $\Phi = \Phi_C + 1 + \beta$. Nonnegativity of μ_W^-: $\Phi_C \geq 0$ by (S1) and $H_C \geq 0$. Nonnegativity of μ_G^-: for $z \in [d, 1]$, $\tau(z) \leq \tau(d) = \hat{s}_C$, so $H_C(\tau(z)) \leq H_C(\hat{s}_C) = 1 + \beta$ by monotonicity of H_C (Lemma 15(iii)) and (7); hence $\varphi \geq 0$, with equality at $z = d$.

(ii) This is the mass identity of Lemma 11 at $\rho = f_C$, level $q = \tau(z)$: $\nu_q^-(\Omega_q) = B(q) = \Gamma_C$, using $B(q) = \Gamma_C$ for $q \leq \hat{s}_C < 1$.

(iii) Immediate from (7): $\Phi_C + 1 + \beta = \Phi_C + H_C(\hat{s}_C)$, and C_1's slices are $\{t \geq \hat{s}_C\}$ at every $z \in [0, d]$.

(iv) $\mu_G^-(C_2) = \int_d^1 f_p(z)\,\varphi(z) \int_{\{t \geq \tau(z)\}} f_C \,\mathrm{d}\omega\,\mathrm{d}z = \int_d^1 f_p\,\varphi\,\bar{G}_C(\tau)\,\mathrm{d}z = \beta$ by (8), and $\mu^+(F_p) = \beta$ was computed above. *Q.E.D.*

Lemma 19 (Stage W: slice couplings)**.** *Let* $\kappa_W := \sup_{q \in [0, \hat{s}_C]} \kappa_{\mathrm{pb}}(f_C, q)$, *with* κ_{pb} *the single-cell threshold of Theorem 8. Then* $\kappa_W < 1$, *and for every* $\kappa \geq \kappa_W$ *there exist:* (i) *a coupling* γ_1 *of* (μ_1^+, μ_1^-) *with* $y - x \in K_{1,\kappa}$ *for* γ_1*-a.e. pair* (y, x); (ii) *a coupling* γ_W *of* $(\mu_{2,C}^+, \mu_W^-)$ *with* $y - x \in K_{2,\kappa}$ *for* γ_W*-a.e. pair.*

Proof. *Hypotheses of the single-cell machinery.* By Lemma 15 f_C is admissible with $\Phi_\rho = \Phi_C$ and H_C strictly increasing on $(0, \bar{t}_C)$, and $\Phi_C \geq 0$ on D_C by (S1). Hence for every $q \in [0, \hat{s}_C]$: f_C is regular above q ($H_C(r) > H_C(q)$ for $r > q$), $\Phi_C + H_C(q) \geq 0$ on Ω_q, and $q < 1 = \underline{t}_C$, so Theorem 8 applies at q and yields, for every $\kappa \geq \kappa_{\mathrm{pb}}(f_C, q)$, a coupling of (ν_q^+, ν_q^-) concentrated on $\{(y_\omega, x_\omega) : y_\omega - x_\omega \in W_\kappa\}$; and Proposition 3 with $\bar{q} = \hat{s}_C < \underline{t}_C$ gives $\kappa_W < 1$. Fix $\kappa \geq \kappa_W$.

Slice structure. Both cells' sources and stage-W sinks disintegrate over z into certificate problems: for $z \in [0, d]$ (cell 1) the slice problem is $(\nu_{\hat{s}_C}^+, \nu_{\hat{s}_C}^-)$ by Lemma 18(iii), and for $z \in [d, 1]$ (cell 2) it is $(\nu_{\tau(z)}^+, \nu_{\tau(z)}^-)$, balanced slice by slice with common total Γ_C (Lemma 18(ii)); the faces are uncut at every such level, so ν_q^+ is the full core-face measure throughout. Explicitly, with $\iota_z(y_\omega, x_\omega) := \big((y_\omega, z), (x_\omega, z)\big)$,

$$\mu_1^+ = \int_0^d (\nu^+ \otimes \delta_z) f_p(z)\,\mathrm{d}z, \qquad \mu_1^- = \int_0^d (\nu_{\hat{s}_C}^- \otimes \delta_z) f_p(z)\,\mathrm{d}z,$$

and similarly on $[d, 1]$ with $\nu_{\tau(z)}^-$ and $\mu_{2,C}^+, \mu_W^-$.

Cell 1. The slice problem does not depend on z: pick one coupling $\gamma^{\hat{s}_C}$ of $(\nu_{\hat{s}_C}^+, \nu_{\hat{s}_C}^-)$ on core-wedge displacements from Theorem 8 and set $\gamma_1 := \int_0^d (\iota_z)_\# \gamma^{\hat{s}_C} f_p(z)\,\mathrm{d}z$. The marginals are μ_1^+ and μ_1^- by the displayed disintegrations, and every displacement is $h = (h_\omega, 0)$ with $h_\omega \in W_\kappa$.

Cell 2, measurable selection. Here the slice problem moves with z through $q = \tau(z)$. Embed all slice couplings in the space $\mathcal{M}_{\Gamma_C}(D_C \times D_C)$ of nonnegative Borel measures of total mass Γ_C — compact and metrizable in the weak topology — and consider the correspondence

$$A(z) \;:=\; \big\{\gamma \in \mathcal{M}_{\Gamma_C}(D_C \times D_C) \;:\; \gamma \text{ couples } (\nu^+, \nu^-_{\tau(z)}) \text{ and } \gamma(R_\kappa) = \Gamma_C\big\}, \qquad R_\kappa := \{(y_\omega, x_\omega) : y_\omega - x_\omega \in W_\kappa\}$$

for $z \in [d, 1]$. Its values are nonempty (Theorem 8 at $q = \tau(z) \in [0, \hat{s}_C]$, since $\kappa \geq \kappa_W \geq \kappa_{\mathrm{pb}}(f_C, \tau(z))$), convex (the marginal constraints are affine and R_κ-concentration is preserved under mixtures), and weakly compact (closed subsets of a compact set: marginal constraints pass to weak limits, and R_κ is closed, so concentration survives by the portmanteau theorem). The graph is closed: if $z_n \to z$ and $\gamma_n \in A(z_n)$ with $\gamma_n \rightharpoonup \gamma$, then the second marginals $\nu^-_{\tau(z_n)} \rightharpoonup \nu^-_{\tau(z)}$ — the densities $(\Phi_C + H_C(\tau(z_n)))f_C\,\mathbf{1}_{\{t \geq \tau(z_n)\}}$ converge in L^1 by dominated convergence, since H_C is continuous on $[0, \hat{s}_C]$, $\{t = \tau(z)\}$ is Lebesgue-null, and $(\Phi_C + H_C(\hat{s}_C))f_C \in L^1$ dominates — while the first marginal is fixed at ν^+; both marginal identities and the R_κ-concentration pass to the limit as before, so $\gamma \in A(z)$. A correspondence into a compact metric space with closed graph is upper hemicontinuous with compact values, hence measurable, and the Kuratowski–Ryll-Nardzewski selection theorem (Kuratowski and Ryll-Nardzewski, 1965) provides a measurable selection $z \mapsto \gamma^W_z \in A(z)$. Set $\gamma_W := \int_d^1 (\iota_z)_\# \gamma^W_z \, f_p(z)\,\mathrm{d}z$; its marginals are $\mu^+_{2,C}$ and μ^-_W, and every displacement is $h = (h_\omega, 0)$ with $h_\omega \in W_\kappa$.

Budget membership in both cells. A stage-W displacement has $h_p = 0$ and $h_\omega \in W_\kappa$. In cell 1 ($S^{(1)}_1 = \varnothing$, $Z_1 = \{p\}$): the sign clause asks $h_p \leq 0$, satisfied with equality, and the budget clause is $\sum_C h_i^- \leq \kappa \sum_C h_i^+$ — the wedge inequality itself. In cell 2 ($S^{(1)}_2 = \{p\}$): the sign clause asks $h_p \geq 0$, again satisfied with equality, and the budget clause $\sum_C h_i^- \leq \kappa(h_p + \sum_C h_i^+)$ reduces to the wedge inequality at $h_p = 0$. So γ_1 is concentrated on $K_{1,\kappa}$ and γ_W on $K_{2,\kappa}$. *Q.E.D.*

It remains to carry the peripheral face μ^+_p onto the residual sinks μ^-_G. Both measures have the *same* ω-conditional given their (t, z) coordinates — the fiber law c_t: on the source side by independence, on the sink side because $\varphi(z)\, f_p(z)$ does not depend on ω. This is what the partition was designed to achieve: the Φ_C-tilt sits entirely in stage W. The stage-G problem therefore aggregates to the two coordinates (t, z). Define

$$m^+ := \beta\, g_C(t)\,\mathrm{d}t \otimes \delta_{z=1}, \qquad m^- := \varphi(z)\, f_p(z)\, g_C(t)\,\mathrm{d}t\,\mathrm{d}z \quad \text{on } \tilde{\Omega} := \{(t, z) : d \leq z \leq 1,\ \tau(z) \leq t \leq \bar{t}_C\},$$

the pushforwards of μ_p^+ and μ_G^- under $x \mapsto (t(\omega), z)$, with equal totals β by Lemma 18(iv). For $\kappa \in [0,1)$ define the *sheared relation* on $\tilde{\Omega}$:

$$(t', z') \preceq_\kappa (t, z) \quad :\Longleftrightarrow \quad z' \le z \ \text{ and } \ t' \le t + \kappa\,(z - z'),$$

a closed preorder, nondecreasing in κ. Descent in the peripheral coordinate buys core dips at rate κ.

Lemma 20 (The Efron lift)**.** *Let $\kappa \in [0,1)$ and let $\tilde{\pi}$ be a coupling of (m^+, m^-) with $(t', z') \preceq_\kappa (t, 1)$ for $\tilde{\pi}$-a.e. pair $\big((t,1),(t',z')\big)$. Then there is a coupling γ_G of (μ_p^+, μ_G^-) with $y - x \in K_{2,\kappa}$ for γ_G-a.e. pair (y, x).*

Proof. *Monotone fiber kernels.* The coordinates of ω are independent with log-concave densities, so by Efron's theorem (Efron, 1965) the map $t \mapsto \int \phi \, \mathrm{d}c_t$ is nondecreasing for every bounded coordinatewise nondecreasing ϕ: the fiber laws satisfy $c_{t'} \preceq_{\mathrm{FOSD}} c_t$ whenever $t' \le t$. By the monotone-coupling form of Theorem 7 (the "in particular" clause), for each pair $t' \le t$ there is a coupling of $(c_t, c_{t'})$ concentrated on the closed set $O := \{(\omega, \omega') : \omega \ge \omega'\}$, and for $t' > t$ one of $(c_t, c_{t'})$ concentrated on $O' := \{(\omega, \omega') : \omega \le \omega'\}$.

Measurable selection. On $\mathcal{P}(D_C) \times \mathcal{P}(D_C)$, the correspondence $(\mu, \nu) \mapsto \{Q \in \mathcal{P}(D_C^2) : Q \text{ couples } (\mu, \nu),\ Q(O) = 1\}$ has convex, weakly compact values and closed graph on the set where it is nonempty (marginal constraints and concentration on the closed set O pass to weak limits by the portmanteau theorem), hence is a measurable closed-valued correspondence there; the same holds with O'. The map $(t, t') \mapsto (c_t, c_{t'})$ is Borel (disintegration kernel), defined for Lebesgue-almost every coordinate — which suffices, since both marginals of $\tilde{\pi}$ are absolutely continuous in their t-coordinates (m^+ with density βg_C, m^- with density $\varphi f_p g_C$) — so composing and splitting the parameter domain into the Borel pieces $\{t' \le t\}$ and $\{t' > t\}$, the Kuratowski–Ryll-Nardzewski theorem (Kuratowski and Ryll-Nardzewski, 1965) yields a measurable family $(t, t') \mapsto Q_{t,t'}$ of couplings of $(c_t, c_{t'})$, concentrated on O when $t' \le t$ and on O' when $t' > t$.

Gluing. Let $\jmath_{z'}(\omega, \omega') := \big((\omega, 1), (\omega', z')\big)$ and set

$$\gamma_G \ := \ \int (\jmath_{z'})_{\#}\, Q_{t,t'} \ \ \tilde{\pi}\big(\mathrm{d}(t,1),\, \mathrm{d}(t',z')\big).$$

The first marginal is $\int c_t \, m^+(\mathrm{d}(t,1)) \otimes \delta_{z=1} = \beta f_C \, \mathrm{d}\omega \otimes \delta_{z=1} = \mu_p^+$, and the second is $\int c_{t'} \, m^-(\mathrm{d}(t',z')) = \varphi(z') f_p(z') f_C(\omega') \, \mathrm{d}\omega' \, \mathrm{d}z' = \mu_G^-$, both by the disintegration $f_C \, \mathrm{d}\omega = \int c_t \, g_C(t) \, \mathrm{d}t$ and Fubini.

Displacements. A γ_G-typical pair has $h = (\omega - \omega',\ 1 - z')$ with $h_p = 1 - z' \geq 0$, satisfying $K_{2,\kappa}$'s sign clause. If $t' \leq t$: $\omega \geq \omega'$, so there are no core dips, $\sum_C h_i^- = 0$, and the budget clause holds with room to spare. If $t' > t$: $\omega' \geq \omega$, so every core coordinate weakly falls, $\sum_C h_i^+ = 0$ and

$$\sum_{i \in C} h_i^- \;=\; \sum_{i \in C} (\omega_i' - \omega_i) \;=\; t' - t \;\leq\; \kappa\,(1 - z') \;=\; \kappa\Big(h_p + \sum_{i \in C} h_i^+\Big),$$

using the sheared inequality at $z = 1$. Either way $h \in K_{2,\kappa}$. *Q.E.D.*

Lemma 21 (The sheared Hall condition and the grand threshold)**.** *For $\kappa \in [0, 1)$ define the score $\theta := t + \kappa z$ on $\tilde{\Omega}$ and the excess*

$$\mathcal{E}(c; \kappa) \;:=\; m^-\big(\{\theta \leq c\}\big) \;-\; m^+\big(\{\theta \leq c\}\big), \qquad c \in \mathbb{R},$$

and let $V := \{\kappa \in [0, 1) : \mathcal{E}(\cdot\,; \kappa) \geq 0\}$. Then: (i) *a coupling $\tilde{\pi}$ as in Lemma 20 exists if and only if $\kappa \in V$;* (ii) *V is a closed up-interval of $[0, 1)$ containing $[\hat{\kappa}, 1)$, where $\hat{\kappa} = \hat{s}_C/(1 - d)$; in particular $\kappa_G := \min V$ exists and $\kappa_G \leq \hat{\kappa}$, with $\hat{\kappa} < 1$ if and only if $\hat{s}_G < 1$.*

Proof. *(i), sufficiency (completeness).* Suppose $\mathcal{E}(\cdot\,; \kappa) \geq 0$. Apply Theorem 7 after normalizing both measures by their common total β (Lemma 18(iv)), with $X := \tilde{\Omega}$ carrying m^-, $Y := \tilde{\Omega} \cap \{z = 1\}$ carrying m^+, and the closed relation $G := \{((t', z'), (t, 1)) : (t', z') \preceq_\kappa (t, 1)\}$. Let $B \subseteq Y$ be closed and nonempty (the empty case is trivial), and let $t_B := \max\{t : (t, 1) \in B\}$, attained by compactness. All sources lie on the single layer $z = 1$ — this is what makes the score family complete — and within the layer the relation is monotone in the source's t, so

$$G^{-1}(B) \;=\; \big\{(t', z') \in \tilde{\Omega} :\ t' \leq t_B + \kappa\,(1 - z')\big\} \;=\; \{\theta \leq c\} \cap \tilde{\Omega}, \qquad c := t_B + \kappa,$$

since membership for some $(t, 1) \in B$ is implied by membership for $(t_B, 1)$. On the layer $\theta = t + \kappa$, so $m^+(B) \leq m^+(\{t \leq t_B\} \times \{1\}) = m^+(\{\theta \leq c\}) \leq m^-(\{\theta \leq c\}) = m^-(G^{-1}(B))$, the middle inequality being $\mathcal{E}(c; \kappa) \geq 0$. The Hall criterion of Theorem 7 holds, a coupling of (m^-, m^+) concentrated on G exists, and flipping it to the order (source, sink) gives $\tilde{\pi}$.

(i), necessity. The score is $\preceq_\kappa$-monotone: if $(t', z') \preceq_\kappa (t, z)$ then $\theta' = t' + \kappa z' \leq t + \kappa(z - z') + \kappa z' = \theta$. So if $\tilde{\pi}$ exists, every source with $\theta \leq c$ has its sink partner in $\{\theta \leq c\}$, whence $m^+(\{\theta \leq c\}) = \tilde{\pi}\big(\{\theta \leq c\} \times \tilde{\Omega}\big) \leq \tilde{\pi}\big(\tilde{\Omega} \times \{\theta \leq c\}\big) = m^-(\{\theta \leq c\})$ for every c: $\kappa \in V$.

(ii), up-interval and closedness. For $\kappa \le \kappa'$ the relation grows, $\preceq_\kappa \subseteq \preceq_{\kappa'}$ (the dip budget $\kappa(z-z')$ only relaxes), so a coupling at κ is one at κ' and necessity in (i) gives $\kappa' \in V$: V is an up-set. It is closed in $[0,1)$ because $\kappa \mapsto \mathcal{E}(c;\kappa)$ is continuous for each fixed c: the m^- term by dominated convergence (the boundary $\{t+\kappa z = c\}$ is m^--null, m^- being absolutely continuous on $\tilde{\Omega}$), and the m^+ term equals $\beta\, G_C\big((c-\kappa)^+ \wedge \bar{t}_C\big)$ with G_C continuous. A closed nonempty up-set of $[0,1)$ is an up-interval containing its infimum, so $\kappa_G := \min V$ exists once $V \neq \varnothing$, which the next step shows; and then $\kappa_G \le \hat{\kappa}$.

(ii), score positivity at $\kappa \ge \hat{\kappa}$. Fix $\kappa \in [\hat{\kappa}, 1)$ and $c \in \mathbb{R}$; write $a_z := c - \kappa z$ for the slice cut, so that the layer cut is $a_1 = c - \kappa$. Slicing,

$$m^-(\{\theta \le c\}) \;=\; \int_d^1 f_p(z)\,\varphi(z)\,\big[G_C(a_z \wedge \bar{t}_C) - G_C(\tau(z))\big]^+ \mathrm{d}z, \qquad m^+(\{\theta \le c\}) \;=\; \beta\, G_C\big(a_1^+ \wedge \bar{t}_C\big).$$

Allocate the source mass across slices with weights $\bar{G}_C(\tau(z))$ via the stage-G books (8), $\beta = \int_d^1 f_p\,\varphi\,\bar{G}_C(\tau)\,\mathrm{d}z$:

$$\mathcal{E}(c;\kappa) \;=\; \int_d^1 f_p(z)\,\varphi(z)\; b(z,c)\,\mathrm{d}z, \qquad b(z,c) := \big[G_C(a_z \wedge \bar{t}_C) - G_C(\tau(z))\big]^+ - \bar{G}_C(\tau(z))\, G_C\big(a_1^+ \wedge \bar{t}_C\big).$$

Since $\varphi \ge 0$ (Lemma 18(i)), it suffices that $b \ge 0$ pointwise on $[d,1] \times \mathbb{R}$.

If $c \le \kappa$ then $a_1^+ = 0$ and $G_C(0) = 0$ (f_C is absolutely continuous, so $t > 0$ almost surely), giving $b \ge 0$. Let $c > \kappa$, so $a_1 > 0$.

Activity: $a_z \ge \tau(z)$ for every $z \in [d,1]$. For $z \ge \hat{s}_G$: $\tau(z) = 0$ and $a_z = c - \kappa z > \kappa(1-z) \ge 0$. For $z < \hat{s}_G$: $a_z \ge \tau(z)$ rearranges to $c \ge \hat{s}_G - (1-\kappa)z$, and

$$c \;>\; \kappa \;\ge\; \hat{s}_G - (1-\kappa)\,d \;\ge\; \hat{s}_G - (1-\kappa)\,z,$$

where the middle inequality is $\kappa(1-d) \ge \hat{s}_G - d = \hat{s}_C$, i.e. $\kappa \ge \hat{\kappa}$. So the positive part in b's first bracket can be dropped.

If $a_z \ge \bar{t}_C$: $b = \bar{G}_C(\tau) - \bar{G}_C(\tau)\, G_C(a_1 \wedge \bar{t}_C) = \bar{G}_C(\tau)\, \bar{G}_C(a_1 \wedge \bar{t}_C) \ge 0$.

If $a_z < \bar{t}_C$: then also $a_1 < \bar{t}_C$, since $a_1 \le a_z$ ($z \le 1$), and writing $G_C = 1 - \bar{G}_C$ throughout,

$$b \;=\; G_C(a_z) - G_C(\tau) - \bar{G}_C(\tau)\, G_C(a_1) \;=\; \bar{G}_C(\tau)\, \bar{G}_C(a_1) - \bar{G}_C(a_z),$$

so $b \ge 0$ if and only if $\bar{G}_C(a_z) \le \bar{G}_C(\tau)\, \bar{G}_C(a_1)$. For $z \ge \hat{s}_G$: $\tau = 0$, $\bar{G}_C(0) = 1$, and $a_z \ge a_1$ gives the inequality by monotonicity of $\bar{G}_C$. For $z \in [d, \hat{s}_G)$:

$$a_z - \tau(z) - a_1 \;=\; (c - \kappa z) - (\hat{s}_G - z) - (c - \kappa) \;=\; \kappa\,(1-z) - (\hat{s}_G - z),$$

and $\kappa(1-z) \geq \hat{s}_G - z$ if and only if $\kappa \geq \frac{\hat{s}_G - z}{1-z}$. The map $z \mapsto \frac{\hat{s}_G - z}{1-z}$ has derivative $\frac{\hat{s}_G - 1}{(1-z)^2} < 0$ — this is where the grand slack $1 - \hat{s}_G > 0$ enters — so its maximum over $[d, \hat{s}_G)$ is at $z = d$, with value $\frac{\hat{s}_G - d}{1-d} = \hat{\kappa} \leq \kappa$. Hence $a_z \geq \tau(z) + a_1$ with $\tau(z), a_1 \geq 0$, and the IFR superadditivity of Lemma 16 gives

$$\bar{G}_C(a_z) \;\leq\; \bar{G}_C\big(\tau(z) + a_1\big) \;\leq\; \bar{G}_C(\tau(z))\, \bar{G}_C(a_1).$$

Note that the activity requirement and the IFR requirement are the *same* inequality $\kappa \geq \frac{\hat{s}_G - z}{1-z}$, binding at the entry slice $z = d$. This proves $\mathcal{E}(\cdot\,; \kappa) \geq 0$ for every $\kappa \geq \hat{\kappa}$, i.e. $[\hat{\kappa}, 1) \subseteq V$. Finally $\hat{\kappa} = \hat{s}_C/(1-d) < 1 \iff \hat{s}_C < 1 - d \iff \hat{s}_G < 1$. *Q.E.D.*

Proof of Theorem 6. Fix $\kappa \geq \kappa^* = \max(\kappa_W, \kappa_G)$, with $\kappa_W < 1$ from Lemma 19 and $\kappa_G \leq \hat{\kappa} < 1$ from Lemma 21, so that $\kappa^* < 1$ and the level range is nonempty.

Cell couplings. On the bare-core cell, Lemma 19(i) supplies a coupling γ_1 of (μ_1^+, μ_1^-) concentrated on $K_{1,\kappa}$ displacements. On the grand cell: since $\kappa \geq \kappa_G$ and V is an up-interval, $\kappa \in V$, so Lemma 21(i) produces the sheared coupling $\tilde{\pi}$ of (m^+, m^-), which Lemma 20 lifts to a coupling γ_G of (μ_p^+, μ_G^-) concentrated on $K_{2,\kappa}$ displacements; and Lemma 19(ii) supplies γ_W of $(\mu_{2,C}^+, \mu_W^-)$, likewise concentrated on $K_{2,\kappa}$. Set $\gamma_2 := \gamma_W + \gamma_G$: its first marginal is $\mu_{2,C}^+ + \mu_p^+ = \mu_2^+$ (the origin atom lies in the exclusion cell, so the grand cell's boundary mass is exactly the core faces over $z \geq d$ plus the peripheral face), its second is $\mu_W^- + \mu_G^- = \mu_2^-$ (Lemma 18(i)), and it is concentrated on $K_{2,\kappa}$ displacements.

Certification. By the certification direction of Lemma 9, exhibiting a coupling of (μ_ℓ^+, μ_ℓ^-) concentrated on $K_{\ell,\kappa}$ displacements certifies the cell's score dominance $\nu_{\ell,\kappa}^- \preceq_{\mathrm{FOSD}} \nu_{\ell,\kappa}^+$. Both selling cells are so equipped, so (f, M) is alignment-regular at level κ (Definition 3) — in particular at every level $\kappa \in [\kappa^*, 1)$, since the budget relation $K_{\ell,\kappa}$ widens with κ, as the theorem asserts.

Standing form. The two-tier setting (S1)–(S3) relaxes the C^2/closed-box-positivity of the standing form, so the sufficiency chain is invoked through its admissible-density extension (Lemma 14). The hypothesis is met: $f = f_C \otimes f_p$ is an admissible core density at $C = [N]$ — $\Phi = \Phi_C + 1 + \beta \geq 1 + \beta > 0$ on $\operatorname{int} D$ (virtual-density paragraph); the full-sum density of $t + z$ is continuous, positive on $(0, N)$, and bounded (convolution of g_C, continuous and bounded by Lemma 15(ii), with $f_p \in L^1$, and the analogue of Lemma 15's induction extends positivity across the z-block), with r times it vanishing at 0; $z f_p(z) = \beta z^\beta \to 0$ and the core lower faces

are handled by Lemma 15(i)–(ii); and the corner clause holds since $f_p \in C^1$ with $f_p(1) = \beta > 0$ and the core corner is Lemma 15(iv). *Conclusion.* The menu is finite on $\mathcal{F}_C$, its selling cells are mass-balanced by (S3), it is inclusive by the margins (geometry paragraph above), and (f, M) is alignment-regular at level κ^*. Since the proof of Theorem 2 consumes restricted optimality only through cell mass balance (Lemma 3), it applies to the mass-balanced two-tier menu — extended to the admissible density f by Lemma 14, and tracking the threshold through its proof (Section B.5, where $\alpha^*(C, M) := 1/(1-\kappa)$ at any valid level κ) gives optimality of the two-tier menu in the unrestricted problem for every $\alpha \geq 1/(1-\kappa^*)$. The equivalence $\hat{\kappa} < 1 \iff \hat{s}_G < 1$ is Lemma 21(ii). *Q.E.D.*

S.3 Proof of Theorem 3

The proof is a six-lemma chain. Lemmas 22, 23 and 24 isolate a closed bulk ball W on which the new option uniquely dominates and the per-buyer transfer rises by exactly δ. Lemma 25 disposes of exclusion-cell switchers. Lemma 26 shows the remaining cells contribute non-negatively under the top-priced anchor. Lemma 27 and the closing argument assemble the tradeoff. Throughout, $M_{\mathcal{F}} = \{(\sigma_\ell, p^*_\ell)\}_{\ell=1}^{L^*}$ is a candidate mechanism on $\mathcal{F} \not\ni [N]$ with a finite menu and the top-cell wedge, with marginal-inclusion vectors w_ℓ; ℓ^* is a top-priced cell with $p^*_{\ell^*} = \bar{p}$; V^* is the witness from Definition 4; and $M^\delta_{\mathcal{F}} := M_{\mathcal{F}} \cup \{([N], p^*_{\ell^*} + \delta)\}$. Without loss of generality no two options share a surplus function $\bar{v}_\ell - p^*_\ell$: among duplicates the seller-favorable tie convention selects the highest-priced, and deleting the others changes neither $u_{M_{\mathcal{F}}}$ nor $t_{M_{\mathcal{F}}}$.

Notational convention. For an option ℓ with lottery σ_ℓ over bundles in $\mathcal{F}$, write $\bar{v}_\ell(x) := \mathbb{E}_{B\sim\sigma_\ell}[v_B(x)]$ for the buyer's expected valuation; by (4) this is linear in x, with gradient the σ_ℓ-average of the bundle gradients. A deterministic option $\sigma_\ell = \delta_{B_\ell}$ gives $\bar{v}_\ell = v_{B_\ell}$, whose gradient is $\alpha\,\mathbf{1}_{B_\ell}$ for a core-containing bundle ($C \subseteq B_\ell$, α-scaled) and $\mathbf{1}_{B_\ell}$ for a non-core bundle (not α-scaled). The proof uses only that each $\bar{v}_\ell$ is linear and that the cone and wedge comparisons of (4) hold; in particular it does not assume the options are α-scaled. The augmenting offer is the deterministic grand bundle $v_{[N]}(x) = \alpha \sum_i x_i$, which is α-scaled since $C \subseteq [N]$.

Lemma 22 (Bulk witness ball)**.** *There exist* $\hat{x} \in V^* \cap \operatorname{int} C_{\ell^*}$ *and* $r_0 > 0$ *with* $\overline{B(\hat{x}, r_0)} \subseteq \operatorname{int} D \cap U^{M_{\mathcal{F}}}_{[N]} \cap \mathbb{R}^N_{>0} \cap \operatorname{int} C_{\ell^*}$.

Proof. *Step 1 (the ambient open set).* For each menu option $\ell \geq 1$, $\{x \in \mathbb{R}^N : v_{[N]}(x) > \bar{v}_\ell(x)\}$ is the preimage of an open half-line under a continuous linear function (linearity by (4)), hence open, and likewise $\{x : v_{[N]}(x) > 0\}$; their finite intersection is the strict-preference set $U_{[N]}^{M_\mathcal{F}}$ of Section 4.2, open in $\mathbb{R}^N$. Since $\operatorname{int} D$ and $\mathbb{R}^N_{>0}$ are open,

$$\Omega := \operatorname{int} D \cap U_{[N]}^{M_\mathcal{F}} \cap \mathbb{R}^N_{>0}$$

is open in $\mathbb{R}^N$.

Step 2 (full-measure interior of C_{ℓ^*}). By piecewise linearity of $u_{M_\mathcal{F}}$, the cell C_{ℓ^*} is, up to a λ_D-null tie set, a closed convex polytope; its topological boundary ∂C_{ℓ^*} is a finite union of $(N-1)$-dimensional polyhedral faces, hence $\lambda_D(C_{\ell^*} \setminus \operatorname{int} C_{\ell^*}) = 0$. Setting $V^{**} := V^* \cap \operatorname{int} C_{\ell^*}$, monotonicity gives $\lambda_D(V^* \setminus V^{**}) \leq \lambda_D(C_{\ell^*} \setminus \operatorname{int} C_{\ell^*}) = 0$, so $\lambda_D(V^{**}) = \lambda_D(V^*) > 0$ by Definition 4.

Step 3 (Lebesgue density point). By the Lebesgue density theorem, λ_D-almost every $x \in V^{**}$ is a density-1 point of V^{**}, i.e., satisfies $\lim_{r \downarrow 0} \lambda_D(V^{**} \cap B(x,r))/\lambda_D(B(x,r)) = 1$. Since $\lambda_D(V^{**}) > 0$, the set of density-1 points is non-empty; pick any such $\hat{x} \in V^{**}$, so that $\hat{x} \in V^* \cap \operatorname{int} C_{\ell^*}$.

Step 4 (closed-ball nesting). Since $\hat{x} \in V^* \subseteq \Omega$ and Ω is open (Step 1), there exists $r_2 > 0$ with $\overline{B(\hat{x}, r_2)} \subseteq \Omega$. Since $\hat{x} \in \operatorname{int} C_{\ell^*}$ and $\operatorname{int} C_{\ell^*}$ is open, there exists $r_1 > 0$ with $\overline{B(\hat{x}, r_1)} \subseteq \operatorname{int} C_{\ell^*}$. Setting $r_0 := \min(r_1, r_2) > 0$ delivers $\overline{B(\hat{x}, r_0)} \subseteq \Omega \cap \operatorname{int} C_{\ell^*}$, which is the claimed nesting. *Q.E.D.*

Lemma 23 (Uniform margins)**.** *Set $W := B(\hat{x}, r_0/2)$. Then $\overline{W}$ lies in the same intersection as in Lemma 22, and there exist $\eta, m, F(W) > 0$ with*

$$\eta := \inf_{\overline{W}} (v_{[N]} - v_{B_{\ell^*}}), \quad m := \inf_{\overline{W}} \left[(v_{B_{\ell^*}} - p^*_{\ell^*}) - \max\{0, \max_{\ell' \neq \ell^*} (v_{B_{\ell'}} - p^*_{\ell'})\} \right], \quad F(W) := \int_W f \, \mathrm{d}\lambda_D.$$

Proof. By Lemma 22, $\overline{B(\hat{x}, r_0)} \subseteq \operatorname{int} D \cap U_{[N]}^{M_\mathcal{F}} \cap \mathbb{R}^N_{>0} \cap \operatorname{int} C_{\ell^*}$. Since $W = B(\hat{x}, r_0/2) \subseteq B(\hat{x}, r_0)$, monotonicity of closed balls gives $\overline{W} \subseteq \overline{B(\hat{x}, r_0)}$, which lies in the same open intersection. By the Heine–Borel theorem, $\overline{W}$ is compact.

Grand-bundle gap. Since $\overline{W} \subseteq U_{[N]}^{M_\mathcal{F}}$, every $x \in \overline{W}$ satisfies $v_{[N]}(x) > v_{B_{\ell^*}}(x)$. The map $\psi(x) := v_{[N]}(x) - v_{B_{\ell^*}}(x)$ is linear (continuous) by (4). By the extreme value theorem applied to ψ on the compact $\overline{W}$, the infimum $\eta := \inf_{\overline{W}} \psi$ is attained at some $x_* \in \overline{W}$, so $\eta = \psi(x_*) > 0$ since ψ is strictly positive on $\overline{W}$.

Surplus margin. Define

$$\Phi(x) := (v_{B_{\ell^*}}(x) - p^*_{\ell^*}) - \max\{0, \max_{\ell' \neq \ell^*} (v_{B_{\ell'}}(x) - p^*_{\ell'})\},$$

the surplus advantage of option ℓ^* over both the outside option and every other priced option. As a finite max of differences of continuous (linear) functions, Φ is continuous on $\overline{W}$. By Lemma 22, $\overline{W} \subseteq \operatorname{int} C_{\ell^*}$. On $\operatorname{int} C_{\ell^*}$, all cell-membership inequalities defining C_{ℓ^*} are strict, so $\Phi(x) > 0$ for all $x \in \overline{W}$. By the extreme value theorem, $m := \inf_{\overline{W}} \Phi$ is attained, so $m > 0$.

Bulk mass. Since $f \in C^2(D)$ and $f > 0$ on $D \supseteq \overline{W}$, the extreme value theorem gives $f_{\min} := \min_{\overline{W}} f > 0$. The open ball W has $\lambda_D(W) > 0$, hence $F(W) \geq f_{\min} \cdot \lambda_D(W) > 0$. *Q.E.D.*

Lemma 24 (Bulk gain)**.** *For every $\delta \in (0, \min\{\eta, m\}/2)$ and $x \in W$, the unique surplus-maximizer in $M_{\mathcal{F}}^{\delta}$ is $([N], p^*_{\ell^*}+\delta)$, and the per-buyer transfer change is δ. Hence $\int_W (t_{M_{\mathcal{F}}^{\delta}} - t_{M_{\mathcal{F}}}) f \, \mathrm{d}\lambda_D = \delta F(W)$.*

Proof. Fix $x \in W$ and $\delta \in (0, \min\{\eta, m\}/2)$, so that $\delta < \eta$, $\delta < m$, and $\delta < \eta/2$.

Step 1 (old transfer pinned). Since $W \subseteq \operatorname{int} C_{\ell^*}$ (Lemma 23), the cell-membership inequalities are strict at x, so the buyer's surplus-maximizer under $M_{\mathcal{F}}$ is uniquely ℓ^*. Hence $u_{M_{\mathcal{F}}}(x) = v_{B_{\ell^*}}(x) - p^*_{\ell^*}$ and $t_{M_{\mathcal{F}}}(x) = p^*_{\ell^*}$.

Step 2 (new-option surplus). The augmented menu $M_{\mathcal{F}}^{\delta}$ posts the new offer $([N], p^*_{\ell^*}+\delta)$. Its surplus at x is

$$\sigma_{\text{new}}(x) := v_{[N]}(x) - (p^*_{\ell^*} + \delta) = u_{M_{\mathcal{F}}}(x) + \big(v_{[N]}(x) - v_{B_{\ell^*}}(x)\big) - \delta.$$

By the definition of η in Lemma 23, $v_{[N]}(x) - v_{B_{\ell^*}}(x) \geq \eta$, so $\sigma_{\text{new}}(x) - u_{M_{\mathcal{F}}}(x) \geq \eta - \delta > \eta/2 > 0$: the new option strictly beats option ℓ^*.

Step 3 (new beats every other priced option). For $\ell' \neq \ell^*$, the definition of m in Lemma 23 gives $u_{M_{\mathcal{F}}}(x) - (v_{B_{\ell'}}(x) - p^*_{\ell'}) \geq m$, so combining with Step 2,

$$\sigma_{\text{new}}(x) - (v_{B_{\ell'}}(x) - p^*_{\ell'}) \geq (\eta - \delta) + m \geq m + \eta/2 > 0.$$

Hence the new option strictly beats every $\ell' \neq \ell^*$. Likewise $\sigma_{\text{new}}(x) \geq u_{M_{\mathcal{F}}}(x) + (\eta - \delta) \geq m + (\eta - \delta) > 0$, so the new option strictly beats the outside option.

Step 4 (uniqueness and transfer change). Steps 2–3 show the maximum in $u_{M_{\mathcal{F}}^{\delta}}(x)$ is uniquely attained by the new option, so the tie convention is moot. Therefore $t_{M_{\mathcal{F}}^{\delta}}(x) = p^*_{\ell^*} + \delta$, giving $t_{M_{\mathcal{F}}^{\delta}}(x) - t_{M_{\mathcal{F}}}(x) = \delta$ pointwise on W.

Step 5 (bulk integral). Both $t_{M_{\mathcal{F}}^{\delta}}$ and $t_{M_{\mathcal{F}}}$ are Borel-measurable as finite linear combinations of indicators of polyhedral cells. Since the integrand is the constant δ on W, linearity of the Lebesgue integral gives $\int_W (t_{M_{\mathcal{F}}^{\delta}} - t_{M_{\mathcal{F}}}) f \, \mathrm{d}\lambda_D = \delta F(W)$. *Q.E.D.*

Lemma 25 (Exclusion-cell switchers)**.** *For every* $\delta > 0$, $\int_{C_0} (t_{M^\delta_{\mathcal{F}}} - t_{M_{\mathcal{F}}})\, f \, \mathrm{d}\lambda_D \geq 0$.

Proof. On C_0, $u_{M_{\mathcal{F}}} = 0$ and $t_{M_{\mathcal{F}}} = 0$. A buyer at $x \in C_0$ switches to the new option iff $v_{[N]}(x) - (p^*_{\ell^*} + \delta) \geq 0$. Two sub-cases: (i) on the strict-switch set $\{x \in C_0 : v_{[N]}(x) > p^*_{\ell^*} + \delta\}$, the new option is uniquely surplus-maximizing (it has surplus > 0 while every old option has surplus ≤ 0 on C_0), so $t_{M^\delta_{\mathcal{F}}}(x) = p^*_{\ell^*} + \delta > 0$ and the integrand $t_{M^\delta_{\mathcal{F}}}(x) - t_{M_{\mathcal{F}}}(x) = p^*_{\ell^*} + \delta > 0$. (ii) On the tie set $\{v_{[N]} = p^*_{\ell^*} + \delta\}$, the tie convention may yield $t_{M^\delta_{\mathcal{F}}}(x) \in \{0, p^*_{\ell^*} + \delta\}$, both non-negative; this set is λ_D-null because $v_{[N]}$ is a non-zero linear function (otherwise $v_{[N]} \equiv 0$, contradicting the witness $V^* \subseteq U^{M_{\mathcal{F}}}_{[N]}$ where $v_{[N]} > 0$), so its level sets are codimension-1 affine hyperplanes. Elsewhere on C_0 the buyer's choice is unchanged and the integrand is 0. The integrand is non-negative λ_D-a.e. on C_0, so its integral against $f \geq 0$ is non-negative. *Q.E.D.*

Lemma 26 (Non-negative cannibalization)**.** *For every selling cell* $\ell \neq \ell^*$ *and every* $\delta > 0$, $\int_{C_\ell} (t_{M^\delta_{\mathcal{F}}} - t_{M_{\mathcal{F}}})\, f \, \mathrm{d}\lambda_D \geq 0$.

Proof. For a lottery option ℓ, read v_{B_ℓ} as the expected valuation $\bar{v}_\ell(x)$ (linear in x, with the per-option gradient above), which reduces to the deterministic bundle value when $\sigma_\ell = \delta_{B_\ell}$. Set $\varphi_\ell := v_{[N]} - v_{B_\ell}$ (a linear function on $\mathbb{R}^N$ by (4)) and $\tau_\ell := p^*_\ell - p^*_{\ell^*}$; the top-priced choice of ℓ^* (Definition 4) gives $\tau_\ell \leq 0$ for every ℓ.

On $\operatorname{int} C_\ell$, $u_{M_{\mathcal{F}}} = v_{B_\ell} - p^*_\ell$ and $t_{M_{\mathcal{F}}} = p^*_\ell$. The buyer switches to $([N], p^*_{\ell^*} + \delta)$ iff $v_{[N]}(x) - (p^*_{\ell^*} + \delta) \geq v_{B_\ell}(x) - p^*_\ell$, equivalently $\varphi_\ell(x) \geq \delta - \tau_\ell$. On the strict switch region $\{\varphi_\ell > \delta - \tau_\ell\}$ the new option is the unique maximizer and the per-buyer transfer change is $\delta - \tau_\ell \geq \delta > 0$; on the tie locus $\{\varphi_\ell = \delta - \tau_\ell\}$ the change lies in $\{0, \delta - \tau_\ell\}$ depending on the tie convention, both values non-negative. The integrand is therefore non-negative pointwise on $\operatorname{int} C_\ell$ — no nullity assumption on the tie locus is needed — and $\lambda_D(C_\ell \setminus \operatorname{int} C_\ell) = 0$ as in Lemma 22 Step 2, so the integral against $f \geq 0$ is non-negative. *Q.E.D.*

Lemma 27 (Bulk-vs-cannibalization tradeoff)**.** *For every* $\delta \in (0, \min\{\eta, m\}/2)$, $R(M^\delta_{\mathcal{F}}) - R(M_{\mathcal{F}}) \geq \delta\, F(W) > 0$.

Proof. Decompose the revenue change cell-by-cell:

$$R(M^\delta_{\mathcal{F}}) - R(M_{\mathcal{F}}) = \sum_{\ell=0}^{L^*} \int_{C_\ell} (t_{M^\delta_{\mathcal{F}}} - t_{M_{\mathcal{F}}})\, f \, \mathrm{d}\lambda_D,$$

the $\{C_\ell\}_{\ell=0}^{L^*}$ partitioning D up to a λ_D-null set.

Cell C_{ℓ^*}. On $W \subseteq \operatorname{int} C_{\ell^*}$ the integrand equals δ pointwise (Lemma 24, Step 4). On $C_{\ell^*} \setminus W$ the buyer is still at option ℓ^* under $M_{\mathcal{F}}$ (λ_D-a.e.) and either remains at ℓ^* (integrand zero) or strictly switches to the new offer (integrand $\delta - \tau_{\ell^*} = \delta > 0$ since $\tau_{\ell^*} = 0$); the tie locus $\{\varphi_{\ell^*} = \delta\}$ is λ_D-null because $\varphi_{\ell^*} = v_{[N]} - v_{B_{\ell^*}}$ is non-zero linear on $\mathbb{R}^N$ (it is strictly positive on $\overline{W}$ by the grand-bundle gap of Lemma 23). Hence $\int_{C_{\ell^*}} (t_{M^\delta_{\mathcal{F}}} - t_{M_{\mathcal{F}}})\, f \, \mathrm{d}\lambda_D \geq \int_W \delta\, f \, \mathrm{d}\lambda_D = \delta\, F(W)$ by monotonicity.

Cell C_0. Lemma 25 gives a non-negative contribution.

Cells $\ell \neq \ell^*$. Lemma 26 gives a non-negative contribution for every selling cell $\ell \neq \ell^*$.

Summing, $R(M^\delta_{\mathcal{F}}) - R(M_{\mathcal{F}}) \geq \delta\, F(W) > 0$. *Q.E.D.*

Proof of Theorem 3. *Step 1 (instantiation).* Apply Lemma 22 to obtain $\hat{x} \in V^* \cap \operatorname{int} C_{\ell^*}$ and $r_0 > 0$ with $\overline{B(\hat{x}, r_0)} \subseteq \operatorname{int} D \cap U^{M_{\mathcal{F}}}_{[N]} \cap \mathbb{R}^N_{>0} \cap \operatorname{int} C_{\ell^*}$. Apply Lemma 23 to obtain $W = B(\hat{x}, r_0/2)$ together with $\eta, m, F(W) > 0$.

Step 2 (choice of δ_0). Set $\delta_0 := \min\{\eta, m\}/2 > 0$. Then every $\delta \in (0, \delta_0)$ satisfies $\delta \in (0, \min\{\eta, m\}/2)$, so Lemmas 24, 25, 26 and 27 all apply.

Step 3 (final inequality chain). Lemma 27 gives

$$R(M^\delta_{\mathcal{F}}) - R(M_{\mathcal{F}}) \;\geq\; \delta\, F(W) > 0,$$

the strict inequality from $\delta > 0$ (Step 2) and $F(W) > 0$ (Lemma 23). Hence $R(M^\delta_{\mathcal{F}}) > R(M_{\mathcal{F}})$ for every $\delta \in (0, \delta_0)$: the wedge-satisfying candidate $M_{\mathcal{F}}$ is strictly improved by the grand-bundle deviation, so it is not optimal in the unrestricted screening problem. In particular, if the restricted-optimal mechanism on $\mathcal{F}$ satisfies the wedge, $[N] \in \mathcal{F}$ is necessary for unrestricted optimality. *Q.E.D.*

S.4 Proof of Theorem 4

Proof of Theorem 4. Fix the excluded item i and a complementarity level $\alpha > 0$, and write $\kappa := \mathbf{1}\{C \subseteq \{i\}\}$, so the singleton's valuation is $v_{\{i\}}(x) = \alpha^\kappa x_i$ with standalone value $V_i = \alpha^\kappa \bar{v}_i$. By Definition 5 the base

$$G := \left\{ x_{-i} \in \textstyle\prod_{j \neq i} [0, \bar{v}_j] : (\bar{v}_i, x_{-i}) \in C_0 \right\}$$

has $\lambda_{N-1}(G) =: c_G > 0$. Each option's expected valuation $\bar{v}_\ell(x) = \mathbb{E}_{B \sim \sigma_\ell}[v_B(x)]$ has nonnegative gradient, so $u_{M_{\mathcal{F}}} = \max_\ell\{\bar{v}_\ell - p_\ell, 0\}$ is nondecreasing in x_i; hence for $x_{-i} \in G$ the entire column $\{(t, x_{-i}) : t \leq \bar{v}_i\}$ lies in C_0, where $u_{M_{\mathcal{F}}} = 0$. Set

$$r_\varepsilon := V_i - \varepsilon, \qquad \varepsilon' := \varepsilon/\alpha^\kappa, \qquad M^\varepsilon_{\mathcal{F}} := M_{\mathcal{F}} \cup \{(\{i\}, r_\varepsilon)\},$$

and $\bar{p} := \max_\ell p_\ell$, $m_0 := \inf_D f > 0$, and the local bound $M_0 := \sup_{T_i} f$ over the top slab $T_i := \{x \in D : x_i \geq \bar{v}_i/2\}$, which contains every switcher on the range $\varepsilon \leq V_i/2$ to which the argument confines itself — only the density's size on the slab enters the loss bound.

Switchers lie in the top slab. A type x strictly prefers the new offer iff $v_{\{i\}}(x) - r_\varepsilon > u_{M_\mathcal{F}}(x) \geq 0$, i.e. $\alpha^\kappa x_i > V_i - \varepsilon$, which forces $x_i > \bar{v}_i - \varepsilon'$. So every switcher has $x_i \in (\bar{v}_i - \varepsilon', \bar{v}_i]$.

Gain. For $x_{-i} \in G$ and $x_i \in (\bar{v}_i - \varepsilon', \bar{v}_i]$ the type lies in C_0, so $u_{M_\mathcal{F}}(x) = 0 < v_{\{i\}}(x) - r_\varepsilon$: it switches and pays r_ε, where before it paid 0 (λ_D-a.e. on A_ε: a type in C_0 transacts only on the λ_D-null union of option-indifference hyperplanes $\{\bar{v}_\ell = p_\ell\}$, where the seller-favorable tie-break is immaterial). The gain region $A_\varepsilon := G \times (\bar{v}_i - \varepsilon', \bar{v}_i]$ has $\lambda_D(A_\varepsilon) = c_G\, \varepsilon'$, so for $\varepsilon \leq V_i/2$ the gross gain is at least

$$r_\varepsilon \cdot m_0 \cdot \lambda_D(A_\varepsilon) \;\geq\; \tfrac{1}{2} V_i\, m_0\, c_G\, \varepsilon' \;=\; \tfrac{1}{2} \bar{v}_i\, m_0\, c_G\, \varepsilon \;=\; \Theta(\varepsilon),$$

using $V_i\, \varepsilon' = \bar{v}_i\, \varepsilon$.

Loss. The only revenue at risk is from types who bought some option under $M_\mathcal{F}$ and switch. Let such a type buy option ℓ ($\ell \geq 1$), so on its cell $u_{M_\mathcal{F}}(x) = \bar{v}_\ell(x) - p_\ell$. Switching gives $v_{\{i\}}(x) - r_\varepsilon \geq \bar{v}_\ell(x) - p_\ell \geq 0$, and since $v_{\{i\}}(x) \leq V_i = r_\varepsilon + \varepsilon$ this forces

$$x_i \in (\bar{v}_i - \varepsilon', \bar{v}_i] \quad \text{and} \quad \bar{v}_\ell(x) \in [p_\ell, p_\ell + \varepsilon].$$

The gradient $\nabla \bar{v}_\ell = \mathbb{E}_{\sigma_\ell}[\alpha^{\mathbf{1}\{C \subseteq B\}} \mathbf{1}_B]$ is parallel to e_i only if every bundle in $\operatorname{supp} \sigma_\ell$ is contained in $\{i\}$, i.e. $\sigma_\ell = \delta_{\{i\}}$ (handled last). Otherwise $\nabla \bar{v}_\ell$ has a nonzero entry in some $j_0 \neq i$; applying Fubini in $(x_i, x_{j_0}, x_{\text{rest}})$, x_i ranges over an interval of length ε' and, for each fixed (x_i, x_{rest}), x_{j_0} over an interval of length $\leq \varepsilon/(\nabla \bar{v}_\ell)_{j_0}$. Hence

$$\lambda_D(\{\text{switchers}\} \cap C_\ell) \;\leq\; \varepsilon' \cdot \frac{\varepsilon}{(\nabla \bar{v}_\ell)_{j_0}} \cdot \operatorname{diam}(D)^{N-2} \;=\; O(\varepsilon^2),$$

and the cell's lost revenue is at most $\bar{p} \cdot M_0 \cdot O(\varepsilon^2) = O(\varepsilon^2)$.

The parallel case $\sigma_\ell = \delta_{\{i\}}$ (item i already offered) contributes no loss. Here $\bar{v}_\ell = v_{\{i\}}$, and an excluded F_i forces $p_\ell \geq V_i$: a near-top type $(\bar{v}_i, x_{-i})$ with $x_{-i} \in G$ is excluded, so its surplus from ℓ satisfies $V_i - p_\ell \leq 0$. Since $\max_x v_{\{i\}}(x) = V_i \leq p_\ell$, the cell $\{x : v_{\{i\}}(x) \geq p_\ell\} \subseteq \{x_i \geq \bar{v}_i\}$ on which ℓ is bought is λ_D-null, so switching from ℓ changes revenue only on a null set.

Net. Summing the $O(\varepsilon^2)$ losses over the finitely many cells (at most $L + 1$),

$$R(M_\mathcal{F}^\varepsilon) - R(M_\mathcal{F}) \;\geq\; \tfrac{1}{2} \bar{v}_i\, m_0\, c_G\, \varepsilon - C\, \varepsilon^2$$

for a constant $C = C(M_{\mathcal{F}}, f, D, \alpha)$. With $\varepsilon_0 := \min\{V_i/2,\ \bar{v}_i\, m_0\, c_G/(2C)\}$ the right-hand side is positive on $(0, \varepsilon_0)$. What is uniform in α is the gross gain $\frac{1}{2}\bar{v}_i\, m_0\, c_G\, \varepsilon$; the constant C, and with it the profitable window $(0, \varepsilon_0)$, may shrink as α grows. At every fixed $\alpha > 0$ the window is nonetheless nonempty and the deviation profitable; hence no mechanism leaving a top face excluded is optimal at any $\alpha > 0$. *Q.E.D.*

S.5 An explicit upper-threshold bound for the two-good uniform: dual-transport line construction

This section derives the explicit constant $\alpha^*_{\text{unif}} = (9 + 4\sqrt{6})/15$ as a sufficient upper-threshold bound for pure bundling in the 2-good iid uniform benchmark.

Setup. For $N = 2$, $f \equiv 1$ on $D = [0,1]^2$, and $v(x) = (x_1, x_2, \alpha(x_1 + x_2))$, the selling cell is $C_1 = \{x_1 + x_2 \geq \hat{s}\}$ ($\hat{s} = \sqrt{2/3}$, area $2/3$) and the exclusion cell is $C_0 = \{x_1 + x_2 < \hat{s}\}$. The dual measure carries the origin atom $\delta_0 \in C_0$, unit-density top-face traces on $F_1, F_2 \subset C_1$ (so $\mu^+(F_j) = 1$), and interior density $\mu^- = 3\,\lambda_D$. On C_0 the atom ($\mu^+(C_0) = 1$) couples to the interior sinks ($\mu^-(C_0) = 3 \cdot \frac{1}{3} = 1$) by displacements $-x' \leq 0$, which lie in the no-purchase cone $\{h : c_\alpha(h) = 0\}$; it remains to couple F_1, F_2 to the interior of C_1 ($\mu^-(C_1) = 2$).

Diagonal split and cone constraint. The main diagonal $\{x_2 = x_1\}$ splits C_1 into congruent halves of area $\frac{1}{3}$; couple F_1 to the lower half $T := C_1 \cap \{x_2 \leq x_1\}$ and F_2 to the upper half by reflection, so each half's sink mass $3 \cdot \frac{1}{3} = 1$ matches its top face. For a source $(1, t) \in F_1$, the grand-bundle cone $\alpha(h_1 + h_2) \geq \max\{0, h_1, h_2\}$ admits a sink x^2 exactly when x^2 lies below the *cone-extremal line* $\ell_\alpha(t) : x_2 = t + m_\alpha(x_1 - 1)$, $m_\alpha := -(\alpha - 1)/\alpha$. These lines share the slope m_α, so they are parallel and the reachable regions $R(t) := T \cap \{\text{below } \ell_\alpha(t)\}$ are nested increasing in t. A cone-respecting coupling of F_1 (uniform in t) onto T therefore exists iff the Hall condition $3\,A(t) \geq t$ holds for every $t \in (0,1)$, where $A(t) := \lambda_D(R(t))$ — a nested-reachability instance of Strassen's theorem (Strassen, 1965; Kellerer, 1984), with equality forced at $t = 1$ by total mass. (Nesting collapses the criterion to the down-set cuts $[0, t]$: the sinks reachable only from sources $\geq \theta$ form the up-set of ν-mass $1 - 3A(\theta)$, and feasibility there is the same inequality $3A(\theta) \geq \theta$.)

Threshold. Where the extremal line meets the anti-diagonal below the main diagonal — $a(t) \geq \hat{s}/2$, with $a(t) := (\hat{s} - t + m_\alpha)/(1 + m_\alpha)$ the x_1-coordinate of $\ell_\alpha(t) \cap \{x_1 +$

$x_2 = \hat{s}\}$ — the diagonal does not bind and direct integration (splitting at $x_1 = \hat{s}$) gives

$$A(t) = t(1-a) - \tfrac{m_\alpha}{2}(1-a)^2 - \tfrac{1}{2}(\hat{s}-a)^2. \tag{10}$$

The gap $g(t) := 3A(t) - t$ has $g'(t) = 3(1-a) - 1$, vanishing at $a(t^\star) = 2/3$, with $g''(t) = 3/(1+m_\alpha) > 0$; so its global minimum is $t^\star = \hat{s} - \frac{2}{3} + \frac{m_\alpha}{3} \approx 0.083$, which lies in this regime ($a(t^\star) = 2/3 > \hat{s}/2$). Setting $g(t^\star) = 0$ gives $m_{\alpha^*} = -9(\hat{s}-2/3)^2 = 4\sqrt{6} - 10$, hence

$$\alpha^*_{\text{unif}} = \frac{1}{1+m_{\alpha^*}} = \frac{1}{4\sqrt{6}-9} = \frac{9+4\sqrt{6}}{15} \approx 1.2532. \tag{11}$$

The Hall condition holds for every t. At $\alpha = \alpha^*_{\text{unif}}$ the minimum $t^\star$ sits in the below-diagonal regime, where $g(t) \geq g(t^\star) = 0$. For larger t — once $a(t) < \hat{s}/2$, i.e. $t > t_\times$, where at $t_\times$ the line passes through the corner $(\hat{s}/2, \hat{s}/2)$ — the diagonal trims $R(t)$: the extremal line then bounds $R(t)$ only over $x_1 \in [y_D(t), 1]$, where $y_D(t) := (t - m_\alpha)/(1-m_\alpha)$ is its diagonal crossing (the moving-endpoint boundary terms cancel because $\ell_\alpha(t)$ meets the diagonal there), so $A'(t) = 1 - y_D(t)$ and $g'(t) = 2 - 3\,y_D(t)$. Since y_D is increasing, g' has a single sign change on $(t_\times, 1)$, so g is unimodal there; its endpoint values are $g(t_\times) > 0$ (inherited from $g' = 2 - 3a > 0$ on $(t^\star, t_\times)$) and $g(1) = 3 \cdot \frac{1}{3} - 1 = 0$, so $g \geq 0$ on $(t_\times, 1)$. Combined with $g \geq g(t^\star) = 0$ on the below-diagonal regime, $g \geq 0$ for every t, with equality only at $t^\star$ and $t = 1$. The $F_1 \to T$ coupling therefore exists; reflection supplies the F_2 coupling, and with the exclusion-cell leg the result is a within-cell, cone-feasible plan γ with $\gamma_1 - \gamma_2 = \mu$, certifying pure bundling by Proposition 2. Raising α lifts each line $\ell_\alpha(t)$ pointwise while fixing $A(1) = \frac{1}{3}$, so g increases in α at every $t < 1$; the certificate hence persists, and pure bundling is optimal for all $\alpha \geq \alpha^*_{\text{unif}}$.

Necessity. For $\alpha < \alpha^*_{\text{unif}}$ the minimizer $t^\star_\alpha$ still lies in the below-diagonal regime, where A equals the unconfined below-line area; $g(t^\star_\alpha) < 0$ then forces $3A(t^\star_\alpha) < t^\star_\alpha$, so even the largest cone-respecting reachable set fails the Hall condition and no such transport exists. Hence α^*_{unif} is the exact threshold of the construction.