Acyclic traversal of affine hyperplane arrangements

Brennan Dury

Abstract

Consider a locally finite affine hyperplane arrangement. Given a flag of faces with codimension r base face, small perturbations transverse to the flag of a point in the base face determine a set of 2^r neighboring chambers. A function from chambers to flags then induces a directed graph on the chambers, with edges corresponding to the neighboring chambers determined by the flags. Given an affine objective function, we first show that the supremum of the objective on each chamber is attained, and that the optimizer set is a closed face of the arrangement, generalizing the usual polyhedral linear-programming statement to a class of feasible sets which need not be polyhedra. We then construct a function from chambers to flags such that the base of the flag is the optimal face for the objective, and the resulting directed graph is acyclic. Consequently, we show that from any chamber with finite optimal value, all sufficiently long paths in the graph eventually reach a chamber whose optimal value is greater than the starting chamber. If the arrangement has finitely many normal spaces, we show that from any chamber with finite optimal value, all sufficiently long paths in the graph eventually reach a chamber whose optimal value is greater than any prescribed threshold.

1 Introduction

Let A be a finite-dimensional real affine space with translation vector space V_A. Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A, and let $\Omega(\mathcal{P})$ be the set of chambers of its complement.

We consider a crossing rule at a codimension r face F and chamber $\mathfrak{c}$, encoded by a flag of faces with base F and top $\mathfrak{c}$. Indexed by $\alpha \in \{\pm 1\}^r$, the choices of signs in the successive one-dimensional transverse directions of the flag determine 2^r chambers, each containing F in its closure, one of which is the source chamber $\mathfrak{c}$. Thus a function from chambers to flags induces a directed graph with edges from the source chamber to the target chambers determined by the flag.

The crossing rule is motivated by the jump formula for the iterated residue [5] derived in our companion paper [2], which relates a residue term with complex contour integrals over sets of the form $c(\alpha) + i\mathbb{R}^k$, where $c(\alpha) \in \mathbb{R}^k$ is an element of the chamber indexed by α. Iteratively crossing faces corresponds to iteratively applying the jump formula, which gives new formulas for the contour integrals. To make these formulas nontrivial, one has to organize the choices of flags to at least guarantee the graph is acyclic. Further, in some applications such as Mellin inversion, one would like bounds on the location of the terminal chambers in order to treat the terminal contour integrals as small error terms.

Given a chamber $\mathfrak{c}$ and a nonconstant affine objective ω, set $\lambda(\mathfrak{c}, \omega) = \sup_{x \in \overline{\mathfrak{c}}} \omega(x)$. The main results can be organized as follows.

First, we show that if an affine objective is bounded above on a closed face of a locally finite affine hyperplane arrangement, then the supremum is attained, and the optimizer set is again a closed face of the arrangement. This generalizes the usual polyhedral linear-programming statement to a

class of feasible sets which need not be polyhedra. We also give a counterexample showing that the weaker assumption that the feasible set looks locally like a polyhedron is not sufficient.

Second, we show that a codimension r flag canonically determines 2^r chambers neighboring its base. These chambers are described both by sign patterns and by a perturbations of a point of the base face.

Third, assuming the base of the chosen flag is the optimal face for ω, acyclicity of the induced graph implies uniform progress from the current value: for each finite-λ starting chamber $\mathfrak{c}_0$, all sufficiently long paths satisfy $\lambda(\mathfrak{c}_L, \omega) > \lambda(\mathfrak{c}_0, \omega)$. This makes strict progress from any current level, but it does not by itself rule out asymptoting to a finite value. If $\mathcal{P}$ has finitely many normal spaces, acyclicity gives the stronger escaping statement: for every threshold η, all sufficiently long paths satisfy $\lambda(\mathfrak{c}_L, \omega) > \eta$.

Finally, we construct a function from chambers to flags such that the base is the optimal face for ω and the induced graph is acyclic. Thus the conclusions of the previous result hold.

2 Locally finite arrangements

We establish notation and basic geometric facts for locally finite affine hyperplane arrangements. These facts have standard analogues in the study of polyhedra.

In this paper, all affine spaces are finite dimensional. For integers $a \le b$, write $[a, b] := \{a, a+1, \dots, b\}$.

Definition 1. *An affine hyperplane arrangement in a real affine space A is a collection $\mathcal{P}$ of affine hyperplanes in A. Its support is $|\mathcal{P}| := \bigcup_{H \in \mathcal{P}} H$. We say that $\mathcal{P}$ is locally finite if every compact subset of A meets only finitely many $H \in \mathcal{P}$.*

Definition 2. *Let $\mathrm{Aff}(A, \mathbb{R})$ be the set of real affine functions on A, and let $\mathrm{Con}(A, \mathbb{R})$ be the set of real constant functions. For $q \in \mathrm{Aff}(A, \mathbb{R})$, define $Z(q) := \{x \in A : q(x) = 0\}$.*

Definition 3. *Let $\mathcal{P}$ be an affine hyperplane arrangement in A. Define $\mathfrak{Q}(\mathcal{P}) := \{q \in \mathrm{Aff}(A, \mathbb{R}) \setminus \mathrm{Con}(A, \mathbb{R}) : Z(q) \in \mathcal{P}\}$.*

Definition 4. *Let $\mathcal{P}$ be an affine hyperplane arrangement in A. For $q \in \mathfrak{Q}(\mathcal{P})$, let $dq \in V_A^* \setminus \{0\}$ be its linear part. The normal spaces of $\mathcal{P}$ are the one-dimensional vector spaces $\{\mathbb{R}dq : q \in \mathfrak{Q}(\mathcal{P})\}$. We say that $\mathcal{P}$ has finitely many normal spaces if this set is finite.*

Definition 5. *Let $\mathcal{P}$ be an affine hyperplane arrangement in A. Let $\Omega(\mathcal{P})$ denote the set of connected components of $A \setminus |\mathcal{P}|$. We call these connected components chambers.*

Definition 6. *A flat of $\mathcal{P}$ is a nonempty intersection of hyperplanes in $\mathcal{P}$, with A allowed as the intersection over an empty index set. If L is a flat, let $\mathcal{P}^L := \{H \cap L : H \in \mathcal{P},\ L \not\subseteq H,\ H \cap L \neq \emptyset\}$ be the restricted arrangement in L. A face of $\mathcal{P}$ is a connected component of $L \setminus |\mathcal{P}^L|$ for some flat L. The faces obtained from the flat $L = A$ are exactly the chambers.*

If F is a face, then $\overline{F}$ is called a closed face. A face consisting of one point is called a vertex.

Definition 7. *Let F be a face. Its affine span is $\mathrm{aff}(F)$, and $\mathrm{lin}(F)$ is the translation space of $\mathrm{aff}(F)$. Define*

$$\dim(F) := \dim(\mathrm{aff}(F)), \qquad \mathrm{codim}(F) := \dim(A) - \dim(F), \qquad \mathrm{Ann}(F) := \{\ell \in V_A^* : \ell|_{\mathrm{lin}(F)} = 0\}.$$

Thus chambers are exactly the codimension-zero faces.

Definition 8. *Let $\mathcal{P}$ be an affine hyperplane arrangement in A, and let F be a face. Define*

$$\mathcal{H}(F, \mathcal{P}) := \{H \in \mathcal{P} : F \subseteq H\}, \qquad \mathfrak{Q}(F, \mathcal{P}) := \{q \in \mathfrak{Q}(\mathcal{P}) : F \subseteq Z(q)\}.$$

The map $q \mapsto Z(q)$ takes $\mathfrak{Q}(F, \mathcal{P})$ onto $\mathcal{H}(F, \mathcal{P})$, with two forms defining the same hyperplane exactly when they differ by a nonzero scalar. If $\mathcal{P}$ is locally finite, Lemma 3 below shows that $\mathcal{H}(F, \mathcal{P})$ is finite, and hence that $\mathfrak{Q}(F, \mathcal{P})$ is finite up to nonzero scalar rescaling.

Definition 9. *For the locally finite arrangement $\mathcal{P}$ and $x \in A$, define the sign function* $\operatorname{sign}_{\mathcal{P}}(x) : \mathfrak{Q}(\mathcal{P}) \to \{\pm 1, 0\}$ *by* $\operatorname{sign}_{\mathcal{P}}(x)(q) := \operatorname{sign}(q(x))$. *If F is a face, by Lemma 1, this sign function is constant on F; write* $\operatorname{sign}_F(q) := \operatorname{sign}_{\mathcal{P}}(x)(q)$ *for any $x \in F$.*

Lemma 1. *Faces are exactly the nonempty fibres of $x \to \operatorname{sign}_{\mathcal{P}}(x)$. In particular, the faces partition A.*

Proof. Let F be a face. By definition, F is a connected component of $L \setminus |\mathcal{P}^L|$ for some flat L. If $H = Z(q)$ contains L, then q vanishes on F. If H does not contain L, then F is contained in one connected component of $L \setminus (H \cap L)$, so q has constant nonzero sign on F. Thus every face is contained in a sign fibre.

Conversely, let S be a nonempty sign fibre and choose $x \in S$. Let

$$L := \bigcap_{H \in \mathcal{P},\ x \in H} H.$$

Local finiteness makes this a finite intersection, so L is a flat. If $y \in S$ and $H = Z(q)$ contains x, then $q(x) = 0$, so equality of the sign functions at x and y gives $q(y) = 0$. Hence $y \in H$, and therefore $S \subseteq L$.

For the reverse inclusion in the claimed description of S, suppose that $y \in L$ and that y satisfies

$$q(y) > 0 \text{ if } q(x) > 0, \qquad q(y) < 0 \text{ if } q(x) < 0$$

for every $q \in \mathfrak{Q}(\mathcal{P})$. If $q(x) = 0$, then $x \in Z(q)$, so the definition of L gives $L \subseteq Z(q)$ and hence $q(y) = 0$. Thus x and y have the same sign function. We have therefore proved that

$$S = \{y \in L : q(y) > 0 \text{ if } q(x) > 0,\ q(y) < 0 \text{ if } q(x) < 0\},$$

where again the conditions range over $q \in \mathfrak{Q}(\mathcal{P})$. As an intersection of L with open affine half-spaces, S is convex. Moreover, if $H \cap L \in \mathcal{P}^L$, then H does not contain L, so $x \notin H$; otherwise H would be one of the hyperplanes defining L. The strict sign condition associated to a defining form for H therefore ensures that $S \cap H = \emptyset$. Hence $S \subseteq L \setminus |\mathcal{P}^L|$.

Let C be the connected component of $L \setminus |\mathcal{P}^L|$ containing x. The first paragraph gives $C \subseteq S$, while connectedness of S gives $S \subseteq C$. Thus $S = C$, so S is a face. □

Lemma 2. *Let $K \subseteq A$ be a nonempty convex subset. Then* $\operatorname{relint}(K)$ *is nonempty and dense in K. If $f \in \operatorname{Aff}(A, \mathbb{R})$ satisfies $f \geq 0$ on K, then either f vanishes on K or $f > 0$ on* $\operatorname{relint}(K)$.

Proof. Choose affinely independent points of K spanning $\operatorname{aff}(K)$. Their simplex lies in K, so $\operatorname{relint}(K)$ is nonempty. The density of $\operatorname{relint}(K)$ in K follows from the fact that if $x \in \operatorname{relint}(K)$ and $y \in K$, then $(1-t)x + ty \in \operatorname{relint}(K)$ for $0 \leq t < 1$.

Suppose $x \in \operatorname{relint}(K)$ and $f(x) = 0$. For any $z \in K$, $x - \epsilon(z - x) \in K$ for all sufficiently small $\epsilon > 0$. Hence

$$0 \leq f(x - \epsilon(z - x)) = -\epsilon f(z) \leq 0,$$

so $f(z) = 0$. Thus a zero on $\operatorname{relint}(K)$ forces $f|_K = 0$, and otherwise f is strictly positive on $\operatorname{relint}(K)$. □

Lemma 3. *Let F be a face.*

1. *F and $\overline{F}$ are convex.*
2. *F is an open connected component in $\operatorname{aff}(F) \setminus |\mathcal{P}^{\operatorname{aff}(F)}|$.*
3. *The set $\mathcal{H}(F,\mathcal{P})$ is finite and $\operatorname{aff}(F) = \bigcap_{H\in\mathcal{H}(F,\mathcal{P})} H$.*
4. *Let G be a face with $G \subseteq \overline{F}$ and $G \neq F$. Then $\operatorname{aff}(G) \subsetneq \operatorname{aff}(F)$ and $\mathcal{H}(F,\mathcal{P}) \subsetneq \mathcal{H}(G,\mathcal{P})$.*

Proof. (1). By Lemma 1, F is a sign fibre. Thus F is the intersection of the hyperplanes $Z(q)$ for which $\operatorname{sign}_F(q) = 0$ and the open affine half-spaces

$$\{x \in A : \operatorname{sign}_F(q)q(x) > 0\}$$

for which $\operatorname{sign}_F(q) \neq 0$. Hence F is convex, and therefore so is $\overline{F}$.

(2). By definition of a face, there exists a flat L such that F is a connected component of $L \setminus |\mathcal{P}^L|$. This complement is locally connected, so F is open in L. Since F is nonempty and open in L, it contains a relative open ball in L. Such a ball affinely spans L, so $\operatorname{aff}(F) = L$.

(3). Choose $x \in F$. Every hyperplane in $\mathcal{H}(F,\mathcal{P})$ contains x, so local finiteness makes this set finite. Moreover, every such hyperplane contains $\operatorname{aff}(F) = L$, and hence

$$L \subseteq \bigcap_{H\in\mathcal{H}(F,\mathcal{P})} H.$$

Since L is a flat, choose $\mathcal{S} \subseteq \mathcal{P}$ such that $L = \bigcap_{H\in\mathcal{S}} H$. For every $H \in \mathcal{S}$, we have $F \subseteq L \subseteq H$, so $\mathcal{S} \subseteq \mathcal{H}(F,\mathcal{P})$. Therefore

$$\bigcap_{H\in\mathcal{H}(F,\mathcal{P})} H \subseteq \bigcap_{H\in\mathcal{S}} H = L.$$

Together with the first inclusion and $\operatorname{aff}(F) = L$, this proves the asserted formula for $\operatorname{aff}(F)$.

(4). Let G and F be faces with $G \subseteq \overline{F}$ and $G \neq F$. Suppose toward a contradiction that $\operatorname{aff}(G) = \operatorname{aff}(F) = L$, and put $X := L \setminus |\mathcal{P}^L|$. By (2), F and G are connected components of X. Connected components are closed in X, so

$$G \subseteq \overline{F} \cap X = F.$$

Thus the two components are equal, a contradiction. Thus $\operatorname{aff}(G) \neq \operatorname{aff}(F)$. Since $G \subseteq \overline{F}$, we have $\operatorname{aff}(G) \subsetneq \operatorname{aff}(F)$. Every hyperplane containing F contains $\operatorname{aff}(F)$, hence contains G, so $\mathcal{H}(F,\mathcal{P}) \subseteq \mathcal{H}(G,\mathcal{P})$. If these sets were equal, (3) would give $\operatorname{aff}(F) = \operatorname{aff}(G)$. The inclusion is therefore strict. □

Lemma 4. *Let*

$$\mathfrak{Q}(\mathcal{P}) = \mathfrak{Q}^0 \sqcup \mathfrak{Q}^+ \sqcup \mathfrak{Q}^-$$

be a partition, and set

$$S_{\mathfrak{Q}^\bullet} := \{x \in A : q(x) = 0 \text{ for all } q \in \mathfrak{Q}^0,\ q(x) \geq 0 \text{ for all } q \in \mathfrak{Q}^+,\ q(x) \leq 0 \text{ for all } q \in \mathfrak{Q}^-\}.$$

If $S_{\mathfrak{Q}^\bullet}$ is nonempty, then $\operatorname{relint}(S_{\mathfrak{Q}^\bullet})$ is nonempty. For $x \in \operatorname{relint}(S_{\mathfrak{Q}^\bullet})$, let $F_{\mathfrak{Q}^\bullet}$ be the face containing x. Then $\overline{F_{\mathfrak{Q}^\bullet}} = S_{\mathfrak{Q}^\bullet}$.

Consequently, a nonempty subset $K \subseteq A$ is a closed face of the arrangement if and only if it has the displayed weak-sign form $S_{\mathfrak{Q}^\bullet}$ for some partition.

Proof. Assume $S_{\mathfrak{Q}^\bullet}$ is nonempty. It is convex by definition. By Lemma 2, it has nonempty relative interior. Choose $x \in \operatorname{relint}(S_{\mathfrak{Q}^\bullet})$, let $F_{\mathfrak{Q}^\bullet}$ be the face containing x. We first show that $F_{\mathfrak{Q}^\bullet} \subseteq S_{\mathfrak{Q}^\bullet}$. If $y \in F_{\mathfrak{Q}^\bullet}$, then $\operatorname{sign}(q(y)) = \operatorname{sign}(q(x))$ for every $q \in \mathfrak{Q}(\mathcal{P})$. For $q \in \mathfrak{Q}^0$, the condition $x \in S_{\mathfrak{Q}^\bullet}$ gives $q(x) = 0$, hence $q(y) = 0$. For $q \in \mathfrak{Q}^+$, it gives $q(x) \geq 0$, hence $q(y) \geq 0$. For $q \in \mathfrak{Q}^-$, it gives $q(x) \leq 0$, hence $q(y) \leq 0$. Thus $y \in S_{\mathfrak{Q}^\bullet}$. Since $S_{\mathfrak{Q}^\bullet}$ is closed, $\overline{F_{\mathfrak{Q}^\bullet}} \subseteq S_{\mathfrak{Q}^\bullet}$.

Conversely, let $a \in S_{\mathfrak{Q}^\bullet}$ and $0 \leq t < 1$. We show that $z_t := (1-t)x + ta$ has the same sign pattern as x. If $q(x) > 0$, then $q \in \mathfrak{Q}^+$, because $q \in \mathfrak{Q}^0$ would give $q(x) = 0$ and $q \in \mathfrak{Q}^-$ would give $q(x) \leq 0$. Hence $q(a) \geq 0$, so $q(z_t) = (1-t)q(x) + tq(a) > 0$. The case $q(x) < 0$ is analogous: then $q \in \mathfrak{Q}^-$, $q(a) \leq 0$, and $q(z_t) < 0$.

It remains to consider $q(x) = 0$. If $q \in \mathfrak{Q}^0$, then $q(a) = 0$. Suppose $q \in \mathfrak{Q}^+$. Since $q \geq 0$ on $S_{\mathfrak{Q}^\bullet}$ and $q(x) = 0$, Lemma 2 gives $q|_{S_{\mathfrak{Q}^\bullet}} = 0$, so $q(a) = 0$. If $q \in \mathfrak{Q}^-$, the same argument applied to $-q$ gives $q(a) = 0$. Therefore $q(z_t) = 0$ whenever $q(x) = 0$. Hence $\operatorname{sign}(q(z_t)) = \operatorname{sign}(q(x))$ for every q, so $z_t \in F_{\mathfrak{Q}^\bullet}$. Since $z_t \to a$, we have $a \in \overline{F_{\mathfrak{Q}^\bullet}}$. Thus $S_{\mathfrak{Q}^\bullet} \subseteq \overline{F_{\mathfrak{Q}^\bullet}}$, so $S_{\mathfrak{Q}^\bullet} = \overline{F_{\mathfrak{Q}^\bullet}}$.

If K is a closed arrangement face, write $K = \overline{F}$ and take $\mathfrak{Q}^0 = \{q : \operatorname{sign}_F(q) = 0\}$, $\mathfrak{Q}^+ = \{q : \operatorname{sign}_F(q) = 1\}$, and $\mathfrak{Q}^- = \{q : \operatorname{sign}_F(q) = -1\}$. The closure $\overline{F}$ is contained in the resulting set $S_{\mathfrak{Q}^\bullet}$. For the reverse inclusion, choose $x \in F$ and $a \in S_{\mathfrak{Q}^\bullet}$. The same sign computation as above shows that $(1-t)x + ta$ has the sign pattern of F for every $0 \leq t < 1$. Hence this open segment lies in F, and letting $t \to 1$ gives $a \in \overline{F}$. Thus $K = \overline{F} = S_{\mathfrak{Q}^\bullet}$. Conversely, if K has the displayed weak-sign form for a partition, the preceding paragraph writes it as the closure of a face. □

Lemma 5. *Let E and H be faces of $\mathcal{P}$. Then $E \subseteq \overline{H}$ if and only if, for every $q \in \mathfrak{Q}(\mathcal{P})$,*

$$\operatorname{sign}_E(q) \in \begin{cases} \{0\}, & \operatorname{sign}_H(q) = 0, \\ \{0, \operatorname{sign}_H(q)\}, & \operatorname{sign}_H(q) \neq 0. \end{cases}$$

Proof. By Lemma 4,

$$\overline{H} = \{y \in A : q(y) = 0 \text{ if } \operatorname{sign}_H(q) = 0,\ \operatorname{sign}_H(q)q(y) \geq 0 \text{ if } \operatorname{sign}_H(q) \neq 0\}.$$

Since signs are constant on E, the containment $E \subseteq \overline{H}$ is exactly the displayed condition on $\operatorname{sign}_E(q)$. □

Definition 10. *Let $C \subseteq A$ be convex. A convex subset $K \subseteq C$ is a convex face of C if every segment in C whose relative interior meets K has both endpoints in K.*

Definition 11. *Let $C \subseteq A$ be convex. A convex subset $K \subseteq C$ is an exposed face of C if $K = \operatorname{argmax}_{y \in C} \ell(y)$ for some affine functional $\ell \in \operatorname{Aff}(A, \mathbb{R})$.*

Lemma 6. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, let $\mathfrak{c} \in \Omega(\mathcal{P})$, and put $C = \overline{\mathfrak{c}}$. For a nonempty subset $K \subseteq C$, the following are equivalent:*

1. *$K = \overline{F}$ for some face $F \subseteq C$.*
2. *K is a closed convex face of C.*
3. *K is a closed exposed face of C.*
4. *There is a subset $\mathcal{Z} \subseteq \mathfrak{Q}(\mathcal{P})$ such that $K = \{x \in C : q(x) = 0 \text{ for all } q \in \mathcal{Z}\}$.*

Proof. Write $s_q := \operatorname{sign}_{\mathfrak{c}}(q)$. Since $\mathfrak{c}$ is a chamber, $s_q \in \{\pm 1\}$. Lemma 4 applied to the partition determined by $\mathfrak{c}$ gives $C = \{x \in A : s_q q(x) \geq 0 \text{ for all } q \in \mathfrak{Q}(\mathcal{P})\}$.

For (3) $\Rightarrow$ (2), suppose $K = \operatorname{argmax}_{y \in C} \ell(y)$ for some affine ℓ. Choose $p \in K$ and put $m := \ell(p)$. Then $\ell \leq m$ on C and $K = C \cap \ell^{-1}(m)$. Hence K is closed and convex. If $tx + (1-t)y \in K$ with

$x, y \in C$ and $t \in (0,1)$, then $m = \ell(tx + (1-t)y) = t\ell(x) + (1-t)\ell(y)$, while $\ell(x), \ell(y) \le m$. Thus $\ell(x) = \ell(y) = m$, so $x, y \in K$. Hence K is a closed convex face of C.

For (2) $\Rightarrow$ (1), assume K is a closed convex face of C. We prove that K is the closure of an arrangement face. By Lemma 2, choose $x \in \operatorname{relint}(K)$. For each $q \in \mathfrak{Q}(\mathcal{P})$, put $f_q := s_q q$. Then $f_q \ge 0$ on C, hence on K, so Lemma 2 says that f_q either vanishes on K or is strictly positive on $\operatorname{relint}(K)$.

It follows that every $q \in \mathfrak{Q}(\mathcal{P})$ has constant sign on $\operatorname{relint}(K)$, namely either 0 or s_q. Hence $\operatorname{relint}(K)$ lies in a single sign fibre; call this arrangement face G, using Lemma 1. Since K is closed, the density statement in Lemma 2 gives $K = \overline{\operatorname{relint}(K)}$, so $K \subseteq \overline{G}$.

We next show $G \subseteq K$. The signs on G are either 0 or s_q, so $G \subseteq C$. If $z \in G$, then $x, z \in G$, and G is relatively open in $\operatorname{aff}(G)$ by Lemma 3; hence, for sufficiently small $\epsilon > 0$, the point $w := x - \epsilon(z - x)$ lies in $G \subseteq C$. Since $x = (1+\epsilon)^{-1} w + \epsilon(1+\epsilon)^{-1} z$, the point $x \in K$ lies in the relative interior of the segment from w to z. Because K is a convex face of C, this forces $w, z \in K$. Thus $G \subseteq K$, and since K is closed, $\overline{G} \subseteq K$. Therefore $K = \overline{G}$, an arrangement closed face.

For (1) $\Rightarrow$ (4), assume $K = \overline{F}$ for an arrangement face $F \subseteq C$. For each $H \in \mathcal{H}(F, \mathcal{P})$, choose an affine form $q_H \in \mathfrak{Q}(\mathcal{P})$ with $Z(q_H) = H$, and put $\mathcal{Z} := \{q_H : H \in \mathcal{H}(F, \mathcal{P})\}$. Lemma 3 gives $\operatorname{aff}(F) = \bigcap_{H \in \mathcal{H}(F,\mathcal{P})} H$. Therefore, for $x \in C$, the conditions $q(x) = 0$ for all $q \in \mathcal{Z}$ are equivalent to $x \in \operatorname{aff}(F)$. We claim that $C \cap \operatorname{aff}(F) = \overline{F}$. The inclusion $\overline{F} \subseteq C \cap \operatorname{aff}(F)$ is immediate. Conversely, if $a \in C \cap \operatorname{aff}(F)$, then $q(a) = 0$ for every q vanishing on F, since such a q vanishes on $\operatorname{aff}(F)$. If q has nonzero sign on F, then that sign agrees with the chamber sign because $F \subseteq C$. Since $a \in C$, it follows that $q(a) \ge 0$ whenever $q > 0$ on F, and $q(a) \le 0$ whenever $q < 0$ on F. The closed-face description in Lemma 4 gives $a \in \overline{F}$. Therefore

$$K = \overline{F} = C \cap \operatorname{aff}(F) = \{x \in C : q(x) = 0 \text{ for all } q \in \mathcal{Z}\}.$$

Finally, we show (4) $\Rightarrow$ (3). Since A is finite dimensional and K is nonempty, there are finitely many $q_1, \dots, q_m \in \mathcal{Z}$ such that

$$K = \{x \in C : q_i(x) = 0 \text{ for } i = 1, \dots, m\},$$

with the empty list allowed. Indeed, choose $a \in K$. The differentials $\{dq : q \in \mathcal{Z}\}$ span a finite-dimensional subspace of V_A^*, so choose $q_1, \dots, q_m \in \mathcal{Z}$ whose differentials span it. If $q \in \mathcal{Z}$, then $dq = \sum_i c_i\, dq_i$, and since $q(a) = q_i(a) = 0$, the affine function $q - \sum_i c_i q_i$ is the zero constant function. Hence the full common zero set inside C is cut out by this finite list. Define the affine functional

$$\ell := -\sum_{i=1}^{m} s_{q_i} q_i.$$

Each $s_{q_i} q_i$ is nonnegative on C, so $\ell \le 0$ on C, and equality holds exactly when every q_i vanishes. Thus $K = \operatorname{argmax}_{y \in C} \ell(y)$, so K is a closed exposed face of C. □

Lemma 7. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. If F is any face of $\mathcal{P}$ and $H_1, \dots, H_m \in \mathcal{P}$, then $\overline{F} \cap H_1 \cap \cdots \cap H_m$ is either empty or $\overline{G}$ for a unique face G of the arrangement.*

Proof. The case $m = 0$ is clear, and by induction on m, it suffices to prove the case $m = 1$. Write $H = Z(h)$ with $h \in \mathfrak{Q}(\mathcal{P})$, and assume $K := \overline{F} \cap H$ is nonempty. By Lemma 4,

$$\overline{F} = \{y \in A : q(y) = 0 \text{ if } \operatorname{sign}_F(q) = 0,\ \operatorname{sign}_F(q) q(y) \ge 0 \text{ if } \operatorname{sign}_F(q) \ne 0\}.$$

Therefore

$$K = \{y \in A : h(y) = 0,\ q(y) = 0 \text{ if } \operatorname{sign}_F(q) = 0,\ \operatorname{sign}_F(q)q(y) \geq 0 \text{ if } \operatorname{sign}_F(q) \neq 0\}.$$

Equivalently, K has the weak-sign form from Lemma 4 for the partition

$$\begin{aligned}
\mathfrak{Q}_H^0 &:= \{q : \operatorname{sign}_F(q) = 0 \text{ or } Z(q) = H\},\\
\mathfrak{Q}_H^+ &:= \{q : \operatorname{sign}_F(q) = 1 \text{ and } Z(q) \neq H\},\\
\mathfrak{Q}_H^- &:= \{q : \operatorname{sign}_F(q) = -1 \text{ and } Z(q) \neq H\}.
\end{aligned}$$

Since K is nonempty, Lemma 4 writes K as $\overline{G}$ for a face G. This face is unique: by Lemma 3, every face is a relatively open convex subset of its affine span, and therefore $G = \operatorname{relint}(\overline{G})$. Thus any face whose closure is K must equal $\operatorname{relint}(K)$. □

Lemma 8. *Let F be a face and let $x \in F$. Let $H \in \mathcal{P}$. Then $x \in H$ if and only if $F \subseteq H$ if and only if* $\operatorname{aff}(F) \subseteq H$.
Moreover, there is a precompact convex open neighborhood U of x such that every arrangement hyperplane meeting U contains F. Consequently, for every $q \in \mathfrak{Q}(\mathcal{P}) \setminus \mathfrak{Q}(F, \mathcal{P})$,

$$Z(q) \cap U = \emptyset \qquad \text{and} \qquad \operatorname{sign}(q(y)) = \operatorname{sign}_F(q) \quad (y \in U).$$

If $\mathfrak{e}$ is a chamber with $F \subseteq \overline{\mathfrak{e}}$, then this common sign is also $\operatorname{sign}_{\mathfrak{e}}(q)$.

Proof. If $\operatorname{aff}(F) \subseteq H$, then $F \subseteq H$, and if $F \subseteq H$, then $x \in H$. Conversely, suppose $x \in H$ and choose $q \in \mathfrak{Q}(\mathcal{P})$ with $H = Z(q)$. Then $q(x) = 0$, so $\operatorname{sign}_F(q) = 0$. Hence $F \subseteq H$, and since H is an affine subspace, $\operatorname{aff}(F) \subseteq H$.
For the neighborhood statement, choose any ball B around x with compact closure. Because $\mathcal{P}$ is locally finite, only finitely many arrangement hyperplanes meet $\overline{B}$. Among these, the ones passing through x are exactly the hyperplanes containing F. x has a positive distance to any given closed set, so shrinking B, and then replacing it by a smaller convex open ball, gives the neighborhood U and avoids the remaining finitely many hyperplanes.
If $q \notin \mathfrak{Q}(F, \mathcal{P})$, then $Z(q)$ does not contain F, so $Z(q) \cap U = \emptyset$. Since U is connected, the sign of q is constant on U. It is $\operatorname{sign}_F(q)$, because $x \in F$. If $F \subseteq \overline{\mathfrak{e}}$, then the chamber $\mathfrak{e}$ meets every neighborhood of x, so this constant is also $\operatorname{sign}_{\mathfrak{e}}(q)$. □

Definition 12. *Let G and H be arrangement faces. We say that H covers G if $G \subsetneq \overline{H}$ and there is no arrangement face K such that*

$$G \subsetneq \overline{K} \subsetneq \overline{H}.$$

Lemma 9. *The face poset of a locally finite hyperplane arrangement, ordered by closure containment, is graded by codimension. Equivalently, if $G \subseteq \overline{M}$, then every maximal chain from G to M has length* $\operatorname{codim}(G) - \operatorname{codim}(M)$. *In particular, a face H with $G \subsetneq \overline{H} \subseteq \overline{M}$ covers G in this interval if and only if* $\operatorname{codim}(H) = \operatorname{codim}(G) - 1$.

Proof. Every strict closure containment strictly decreases codimension by Lemma 3, so a chain from G to M has length at most $\operatorname{codim}(G) - \operatorname{codim}(M)$.
It remains to show that no maximal chain can contain a jump of codimension greater than one. Suppose $E \subseteq \overline{B}$, and put $d := \dim B - \dim E$. If $d > 1$, we construct an intermediate face N with

$$E \subsetneq \overline{N} \subsetneq \overline{B}.$$

Choose $x \in E$, and choose a linear complement N_E to $\operatorname{lin}(E)$ in V_A. The finite set $\mathcal{H}(E, \mathcal{P})$ defines a finite central arrangement on N_E: for each $H \in \mathcal{H}(E, \mathcal{P})$, choose $q_H \in \mathfrak{Q}(E, \mathcal{P})$ with $Z(q_H) = H$, and use the restricted form $\widehat{q}_H = dq_H|_{N_E}$. The common kernel of these restricted forms is zero, because Lemma 3 gives $\operatorname{aff}(E) = \bigcap_{H \in \mathcal{H}(E,\mathcal{P})} H$.

Identify the chosen complement N_E with $V_A/\operatorname{lin}(E)$. For every face K satisfying $E \subseteq \overline{K}$, let C_K be the sign fibre of the central arrangement whose sign on $\widehat{q}_H$ is $\operatorname{sign}_K(q_H)$. This fibre is nonempty. Indeed, choose $y \in K$ arbitrarily close to x, write $y - x = e + v$ with $e \in \operatorname{lin}(E)$ and $v \in N_E$, and observe that $q_H(y) = \widehat{q}_H(v)$ for every $H \in \mathcal{H}(E, \mathcal{P})$.

The map $K \mapsto C_K$ is a bijection from the faces whose closures contain E to the faces of this central arrangement. The central signs determine all signs on hyperplanes containing E; the signs on the remaining hyperplanes are forced to be those of E by Lemma 5. Conversely, if v belongs to a central-arrangement face, then for sufficiently small $\epsilon > 0$, Lemma 8 makes the signs of $x + \epsilon v$ on hyperplanes not containing E equal to those of E, while its remaining signs are the prescribed central signs. This constructs the inverse map. The closure criterion shows that the bijection preserves closure containment, and the same local parametrization identifies $\operatorname{lin}(C_K)$ with $\operatorname{lin}(K)/\operatorname{lin}(E)$. Hence

$$\dim C_K = \dim K - \dim E.$$

In particular, the signs of B define a face C_B of dimension d, and the interval $[E, B]$ corresponds to the faces contained in $\overline{C_B}$. The closed cone $\overline{C_B}$ is pointed: if $v, -v \in \overline{C_B}$, then all restricted forms vanish on v, so $v = 0$. Since $d > 1$, the cone $\overline{C_B}$ is a nonzero pointed polyhedral cone, so it has an extreme ray R. Let N be the arrangement face corresponding to the relative interior of R. Then

$$E \subseteq \overline{N} \subseteq \overline{B}, \qquad \dim N = \dim E + 1.$$

Thus $E \subsetneq \overline{N}$, and $N \neq B$ because $d > 1$. Hence a maximal chain cannot jump directly from E to B unless $\dim B - \dim E = 1$.

Therefore every step in a maximal chain increases dimension by exactly one, so every maximal chain from G to M has length $\operatorname{codim}(G) - \operatorname{codim}(M)$. □

Lemma 10. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. If $B \subseteq A$ is compact, then only finitely many face closures $\overline{F}$ meet B. In particular, only finitely many chamber closures meet B.*

Proof. Choose an open ball U with $B \subset U$ and compact closure. Let $\mathcal{P}_U := \{H \in \mathcal{P} : H \cap U \neq \emptyset\}$. This set is finite by local finiteness, since every hyperplane meeting U also meets $\overline{U}$. The finite arrangement $\mathcal{P}_U$ has only finitely many faces: it has only finitely many flats, and on each flat its restricted arrangement is finite, so the complement is partitioned by finitely many strict sign choices by Lemma 1.

Let F be a face of $\mathcal{P}$ with $\overline{F} \cap B \neq \emptyset$, and choose $a \in \overline{F} \cap B$. Since $B \subset U$, the open set U is a neighborhood of a. The face F is dense in $\overline{F}$, so $F \cap U \neq \emptyset$. Choose $y \in F \cap U$, and let $E_U(F)$ be the face of the finite arrangement $\mathcal{P}_U$ containing y.

The map $F \mapsto E_U(F)$ is injective. Indeed, suppose F and F' meet U in points $y \in F$ and $y' \in F'$ that lie in the same face of $\mathcal{P}_U$. For hyperplanes in $\mathcal{P}_U$, the signs at y and y' agree. For hyperplanes outside $\mathcal{P}_U$, the hyperplane misses the convex set U, so its defining forms have constant nonzero sign on U; hence the signs at y and y' again agree. So by Lemma 1, $F = F'$.

Therefore the faces F whose closures meet B inject into the finite set of faces of $\mathcal{P}_U$. The statement for chambers is the codimension zero case. □

3 Nonpolyhedral linear programming

We prove in Theorem 1 that an affine objective bounded above on a closed face of a locally finite affine hyperplane arrangement attains its supremum, and that its optimizer set is again a closed arrangement face. We first give an example showing the role of local finiteness.

Definition 13. *Let A be a real affine space. A closed convex set $C \subseteq A$ is locally polyhedral if there exist an affine hyperplane arrangement $\mathcal{P}$ and a chamber $\mathfrak{c}$ of $\mathcal{P}$ such that $C = \overline{\mathfrak{c}}$, and every compact $K \subseteq C$ meets only finitely many hyperplanes of $\mathcal{P}$.*

Consider the example drawn in Figure 1. In $A = \mathbb{R}^2$, write $p_n = (n, -2^{-n})$ for $n \geq 0$, and let H_n be the line through p_n and p_{n+1}:

$$H_n = \{(x, y) : y = \ell_n(x)\}, \qquad \ell_n(x) := 2^{-(n+1)}(x - n - 2).$$

Set

$$C := \{(x, y) \in \mathbb{R}^2 : x \geq 0,\ y \leq \ell_n(x) \text{ for all } n \geq 0\}.$$

Let $V := \{(x, y) : x = 0\}$. The set C is locally polyhedral: it is the closure of a chamber for $\{H_n : n \geq 0\} \cup \{V\}$, and every compact subset of C meets only finitely many of these lines. Let $\omega(x, y) = y$. The optimal value is 0, but it is never achieved.

However, the hyperplane arrangement $\{H_n : n \geq 0\} \cup \{V\}$ is not locally finite because infinitely many hyperplanes meet the compact set $\{(x, y) : x = 0, y \in [-1, 0]\}$. Thus, if the optimal value is not achieved for a closed arrangement face, the proof must force infinitely many hyperplanes to meet some compact set, possibly outside C.

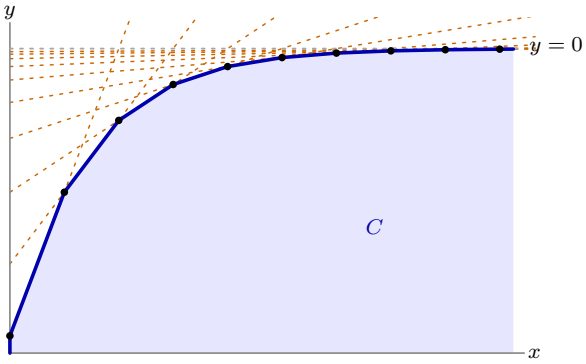


Figure 1: A locally polyhedral feasible set with unattained optimum. The boundary points $p_n = (n, -2^{-n})$ escape to infinity and approach the level $y = 0$, while the supporting lines H_n accumulate outside the feasible set C.

Lemma 11. *Let $\mathcal{L}$ be a locally finite affine line arrangement in an affine plane Π, and let $D \subset \Pi$ be the closure of a face of $\mathcal{L}$. Let $\omega \in \mathrm{Aff}(\Pi, \mathbb{R})$ be nonconstant and suppose $D \subset \{\omega < 0\}$. Let $R \subset \Pi$ be an affine line on which ω is nonconstant, choose $e \in \mathrm{lin}(R)$ with $d\omega(e) = 1$, and let $p_0 \in D \cap R$ maximize ω on $D \cap R$. If $p_1 \in D \setminus R$ and $\omega(p_1) > \omega(p_0)$, then there is an arrangement line $L \in \mathcal{L}$ supporting D at a point z with $\omega(z) \geq \omega(p_1)$, meaning $z \in D \cap L$ and D lies in one of the closed half-planes bounded by L, and L meets the compact segment*

$$S_R := \{p_0 + te : 0 \leq t \leq -\omega(p_0)\} \subset R.$$

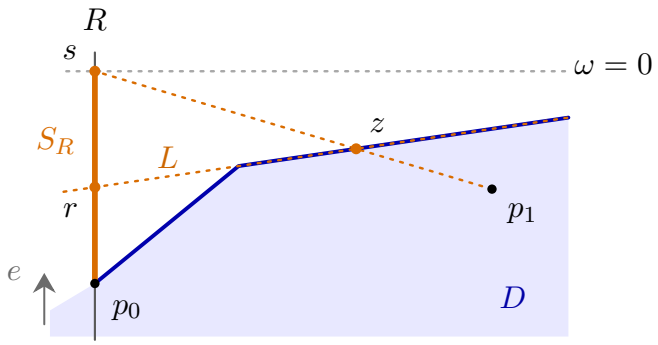


Figure 2: The full-dimensional case of the planar step. Let s be the top endpoint of S_R. The segment from p_1 to s leaves the closed face D at z. An active arrangement line L at z separates the outgoing direction toward s from D, and L meets S_R.

Proof. If D has empty interior in Π, then the segment from p_0 to p_1 is contained in ∂D. Since D is the closure of a face and contains two distinct points, that face is one-dimensional; its affine span is an arrangement line. This line supports D, contains p_1, and meets S_R at p_0.

Assume now that D has nonempty interior. Put $s := p_0 - \omega(p_0)e$, the endpoint of S_R with $\omega(s) = 0$. Then $s \notin D$, while $p_1 \in D$. By convexity and closedness, the intersection $D \cap [p_1, s]$ is a closed segment $[p_1, z]$ for some $z \in [p_1, s)$. Since ω increases from p_1 to s, we have $\omega(z) \geq \omega(p_1)$.

The point z lies in ∂D. Let E be the face of $\mathcal{L}$ containing z. By Lemma 8, applied to $\mathcal{L}$, the lines not containing E keep their D-side signs near z. Since D is the closure of a chamber, for each arrangement line M through z choose an affine function ℓ_M with $M = Z(\ell_M)$ and $D \subseteq \{x : \ell_M(x) \geq 0\}$. The endpoint property $D \cap [p_1, s] = [p_1, z]$ forces at least one of these functions to satisfy $d\ell_M(s - z) < 0$. Indeed, suppose instead that $d\ell_M(s - z) \geq 0$ for every arrangement line through z. Choose a ball U around z whose closure meets only finitely many arrangement lines. For the lines through z, the chamber-side inequalities remain nonnegative at $z + \epsilon(s - z)$. For each of the finitely many other lines meeting U, orient a defining function to be nonnegative on D. Its value at z is strictly positive, so it remains positive at $z + \epsilon(s - z)$ when $\epsilon > 0$ is sufficiently small. All remaining arrangement lines miss U, and hence miss these nearby points. The weak-sign description of the chamber closure would therefore put $z + \epsilon(s - z)$ in D, contradicting the definition of z. Choose a line L with $d\ell_L(s - z) < 0$, and write $\ell = \ell_L$. Then $L = Z(\ell)$ and $D \subseteq \{x : \ell(x) \geq 0\}$. $\ell(z) = 0$ and $d\ell(s - z) < 0$, hence $\ell(s) < 0$. Also $\ell(p_0) \geq 0$, since $p_0 \in D$. By the intermediate value theorem on the segment $[p_0, s] = S_R$, there is $r \in S_R$ with $\ell(r) = 0$, so $r \in L$. Thus L supports D at z, with $\omega(z) \geq \omega(p_1)$, and L meets S_R. □

Theorem 1 (Attainment on closed arrangement faces)**.** *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, let F be a face, and set $C = \overline{F}$. If $\omega \in \mathrm{Aff}(A, \mathbb{R})$ is an affine functional bounded above on C, then there is $x \in C$ such that $\omega(x) = \sup_{y \in C} \omega(y)$. Moreover, the optimal set $\operatorname{argmax}_{y \in C} \omega(y)$ is the closure of a face of $\mathcal{P}$.*

Proof. We argue by induction on $\dim C$. The zero and one dimensional cases are immediate. Assume the result for closed arrangement faces of smaller dimension, and normalize so that $\sup_C \omega = 0$. If $d\omega$ vanishes on $\mathrm{lin}(C)$, then ω is constant on C; the optimal set is C, which is a closed face, so there is nothing to prove. Thus we can assume $d\omega$ does not vanish on $\mathrm{lin}(C)$, so we may choose $e \in \mathrm{lin}(C)$ with $d\omega(e) = 1$.

Assume, for contradiction, that $C \cap \{\omega = 0\} = \emptyset$. Choose $x_* \in C$. Along $x_* + \mathbb{R}e$, the functional ω is an affine coordinate with slope one. The nonempty set $C \cap (x_* + \mathbb{R}e)$ is closed, and its ω-values form a closed subset of $\mathbb{R}$ bounded above by zero; hence its maximum is attained. Replace x_* by such a maximizer and call it x_0. Put $\omega_0 := \omega(x_0) < 0$.

For each n, the equality $\sup_C \omega = 0$ gives a point of C with value greater than $-1/n$. Since F is dense in $C = \overline{F}$ and ω is continuous, we may choose $x_n \in F$ with $-1/n < \omega(x_n) < 0$.

For n large, $\omega(x_n) > \omega_0$. Let $\Pi_n \subset \mathrm{aff}(C)$ be the affine span of the line $x_0 + \mathbb{R}e$ and x_n. Since x_0 maximizes ω on $C \cap (x_0 + \mathbb{R}e)$, such an x_n cannot lie on the line $x_0 + \mathbb{R}e$. Hence Π_n is two-dimensional.

Let $\mathcal{L}_n := \{\Pi_n \cap H : H \in \mathcal{P},\ \Pi_n \cap H \text{ is an affine line}\}$. This is a locally finite line arrangement in Π_n. Put $D_n := C \cap \Pi_n$.

We claim that D_n is the closure of the chamber of $\mathcal{L}_n$ which contains x_n.

We have $x_n \in \Pi_n \cap F$. By Lemma 4, C is defined inside $\mathrm{aff}(C)$ by

$$\{z \in \mathrm{aff}(C) : q(x_n)q(z) \geq 0 \text{ for all } q \in \mathfrak{Q}(\mathcal{P})\}.$$

Indeed, if $q(x_n) = 0$, then q vanishes on F, hence on $\mathrm{aff}(C)$; otherwise $\mathrm{sign}_F(q)$ is the sign of $q(x_n)$.

The point x_n lies on no line of $\mathcal{L}_n$: if $x_n \in \Pi_n \cap H$ for some $H \in \mathcal{P}$, then Lemma 8 gives $F \subseteq H$, hence $\mathrm{aff}(C) \subseteq H$, so $\Pi_n \cap H = \Pi_n$, not a line. For each line $L \in \mathcal{L}_n$, choose an affine form $\ell_L \in \mathrm{Aff}(\Pi_n, \mathbb{R})$ with $Z(\ell_L) = L$ and $\ell_L(x_n) > 0$. If $q \in \mathfrak{Q}(\mathcal{P})$ and $\Pi_n \cap Z(q) = \emptyset$, then $q|_{\Pi_n}$ has constant nonzero sign, so $q(x_n)q(z) > 0$ for all $z \in \Pi_n$. If $\Pi_n \subseteq Z(q)$, the inequality is also automatic. The only remaining case is $\Pi_n \cap Z(q) = L \in \mathcal{L}_n$, where $q(x_n)q(z) \geq 0$ is equivalent to $\ell_L(z) \geq 0$. Therefore

$$D_n = \{z \in \Pi_n : \ell_L(z) \geq 0 \text{ for all } L \in \mathcal{L}_n\}.$$

This is the closure of the chamber of $\mathcal{L}_n$ containing x_n, by Lemma 4.

Applying Lemma 11 in Π_n, with reference line $x_0 + \mathbb{R}e$, gives a hyperplane $\widehat{H_n}$ of $\mathcal{L}_n$ and a point $y_n \in D_n \cap \widehat{H_n}$ such that $\omega(y_n) \geq \omega(x_n) > -1/n$ and $\widehat{H_n} \cap S \neq \emptyset$, where $S := \{x_0 + te : 0 \leq t \leq -\omega_0\}$.

Since $\widehat{H_n}$ is a hyperplane of $\mathcal{L}_n$, there is a hyperplane H_n of $\mathcal{P}$ such that $\Pi_n \cap H_n = \widehat{H_n}$. Then H_n does not contain Π_n, and thus does not contain $\mathrm{aff}(C)$. Also, $y_n \in D_n \cap \widehat{H_n} \subseteq C \cap H_n$.

Because $\mathcal{P}$ is locally finite, only finitely many distinct H_n meet the compact set S. Hence some hyperplane H occurs infinitely often among the H_n. Put $C_H := C \cap H$. This set is nonempty, because the corresponding subsequence y_n lies in it. Since H does not contain $\mathrm{aff}(C)$, C_H is a proper subset of C.

By Lemma 7, C_H is the closure of an arrangement face G. Since C_H is a proper subset of C, we have $G \neq F$. Thus Lemma 3 gives $\mathrm{aff}(G) \subsetneq \mathrm{aff}(F)$, hence $\dim C_H < \dim C$.

Since $C_H \subseteq C$ and $\sup_C \omega = 0$, we have $\sup_{C_H} \omega \leq 0$. The subsequence $y_n \in C_H$ with $\omega(y_n) > -1/n$ gives $\sup_{C_H} \omega \geq 0$, hence $\sup_{C_H} \omega = 0$. But $C \cap \{\omega = 0\} = \emptyset$, so this supremum is not attained on C_H, contradicting the induction hypothesis. Therefore $C \cap \{\omega = 0\} \neq \emptyset$, as required.

Let $K := \mathrm{argmax}_{y \in C}\, \omega(y)$. Then K is a nonempty closed exposed face of C. Apply Lemma 6 to the restricted arrangement $\mathcal{P}^{\mathrm{aff}(F)}$ in $\mathrm{aff}(F)$, where F is a chamber and $C = \overline{F}$. It follows that K is the closure of a face of the restricted arrangement. Here $\mathrm{aff}(F)$ is a flat by Lemma 3. Every flat of the restricted arrangement is a flat of $\mathcal{P}$ contained in $\mathrm{aff}(F)$, and the restricted sign fibres are exactly the sign fibres of $\mathcal{P}$ lying in $\mathrm{aff}(F)$. Thus its faces are precisely the faces of $\mathcal{P}$ contained in $\mathrm{aff}(F)$, and K is the closure of a face of $\mathcal{P}$. $\square$

4 From face flags to target chambers

We now work locally around an arbitrary face and convert a flag based there into 2^r target chambers, where r is the codimension of the base. Positive adapted coordinates allow us to realize these chambers by separated transverse perturbations. The main result, Theorem 2, also characterizes the targets by their sign patterns. This construction does not require an objective or a choice of optimal face.

4.1 Face flags and sign patterns

Definition 14. *For a face F of $\mathcal{P}$, define its chamber star by* $\mathrm{Star}_\Omega(F) := \{\mathfrak{c} \in \Omega(\mathcal{P}) : F \subseteq \bar{\mathfrak{c}}\}$.

Lemma 12. $\mathrm{Star}_\Omega(F)$ *is finite.*

Proof. If $x \in F$, then every chamber in $\mathrm{Star}_\Omega(F)$ has closure that meets $\{x\}$. Since $\{x\}$ is compact, only finitely many chamber closures meet $\{x\}$ by Lemma 10. □

Definition 15. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. A flag of faces is a sequence of faces $F_r, \dots, F_0$ such that $F_i \subsetneq \overline{F_{i-1}}$ for all $i \in [1, r]$, and* $\mathrm{codim}(F_i) = i$. *We say that F_r is the base and F_0 is the top.*

Definition 16. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. Consider any flag of faces with base F and top $\mathfrak{c}$, written $F = F_r, \dots, F_0 = \mathfrak{c}$. The hyperplane layers attached to the flag are $\mathcal{H}_i := \mathcal{H}(F_i, \mathcal{P}) \setminus \mathcal{H}(F_{i-1}, \mathcal{P})$.*

For every $i \in [1, r]$, Lemma 3 gives $\mathcal{H}(F_{i-1}, \mathcal{P}) \subsetneq \mathcal{H}(F_i, \mathcal{P})$: the containment follows from $F_i \subseteq \overline{F_{i-1}}$, and it is strict because the two faces have different affine spans. Also $\mathcal{H}(F_0, \mathcal{P}) = \emptyset$, since $F_0 = \mathfrak{c}$ is a chamber. Thus the $\mathcal{H}_i$ are nonempty finite sets and form a disjoint partition

$$\mathcal{H}(F, \mathcal{P}) = \mathcal{H}_r \sqcup \cdots \sqcup \mathcal{H}_1.$$

Equivalently, the sets $\mathfrak{Q}(F_i, \mathcal{P})$ form a strictly increasing chain from $\mathfrak{Q}(F_0, \mathcal{P}) = \emptyset$ to $\mathfrak{Q}(F_r, \mathcal{P}) = \mathfrak{Q}(F, \mathcal{P})$, and $\mathcal{H}_i$ records the zero hyperplanes of the forms added at the i-th step.

Definition 17. *Assume the setting of Definition 16. For any $q \in \mathfrak{Q}(F, \mathcal{P})$, define* $\mathrm{gi}(q)$ *as the unique index such that $Z(q) \in \mathcal{H}_{\mathrm{gi}(q)}$.*

Theorem 2 translates a flag with base F and a sign vector $\alpha \in \{\pm 1\}^r$ into another flag of faces with base F. The top of that flag is a chamber $\mathfrak{b}_{FL}(\alpha) \in \mathrm{Star}_\Omega(F)$. We defer the proof until after the coordinate construction, where Corollary 2 also identifies an explicit point in each target face.

Theorem 2 (Opposite flags)**.** *(→ proof) Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A, and let $\mathfrak{c} \in \Omega(\mathcal{P})$. Let $F = F_r, \dots, F_0 = \mathfrak{c}$ be a flag of faces of $\mathcal{P}$ with base F and top $\mathfrak{c}$, with* $\mathrm{codim}(F) = r > 0$, *and denote it by FL. For each $\alpha \in \{\pm 1\}^r$, there exists a chamber $\mathfrak{b}_{FL}(\alpha) \in \Omega(\mathcal{P})$ and another flag of faces $F = F_r(\alpha), \dots, F_0(\alpha) = \mathfrak{b}_{FL}(\alpha)$ with signs*

$$\mathrm{sign}_{F_i(\alpha)}(q) = \begin{cases} 0 & : q \in \mathfrak{Q}(F_i, \mathcal{P}) \\ \alpha_{\mathrm{gi}(q)} \, \mathrm{sign}_{\mathfrak{c}}(q) & : q \in \mathfrak{Q}(F, \mathcal{P}) \setminus \mathfrak{Q}(F_i, \mathcal{P}) \\ \mathrm{sign}_{\mathfrak{c}}(q) & : q \in \mathfrak{Q}(\mathcal{P}) \setminus \mathfrak{Q}(F, \mathcal{P}) \end{cases}.$$

We write $FL(\alpha)$ for this sign-indexed flag. In particular, $F_i(1^r) = F_i$ for all i, so $\mathfrak{b}_{FL}(1^r) = \mathfrak{c}$.

Definition 18. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. Let $\mathfrak{c} \in \Omega(\mathcal{P})$, and let $F = F_r, \dots, F_0 = \mathfrak{c}$ be a flag of faces with base F and top $\mathfrak{c}$, with $r > 0$. Denote the flag by FL.*

Apply Theorem 2 to FL. The chambers determined by the flag are $\mathfrak{b}_{FL}(\alpha)$, $\alpha \in \{\pm 1\}^r$. *They lie in* $\mathrm{Star}_\Omega(F)$, *and* $\mathfrak{b}_{FL}(1^r) = \mathfrak{c}$. *The target chambers of FL are* $\mathfrak{b}_{FL}(\alpha)$, $\alpha \in \{\pm 1\}^r \setminus \{1^r\}$.

The source and its $2^r - 1$ targets are mutually distinct. Indeed, if $\alpha \neq \beta$, choose i with $\alpha_i \neq \beta_i$, then choose $H \in \mathcal{H}_i$ and $q \in \mathfrak{Q}(F, \mathcal{P})$ with $Z(q) = H$. The sign formula in Theorem 2 gives opposite signs for q on $\mathfrak{b}_{FL}(\alpha)$ and $\mathfrak{b}_{FL}(\beta)$.

Figure 3 shows the target chambers in a two-dimensional example. The arrangement in $A = \mathbb{R}^2$ has hyperplanes $x_2 = 0$, $x_1 = 0$, $x_1 + \sqrt{3}x_2 = 0$, $x_2 + x_1 = 0$, and $\sqrt{3}x_1 + x_2 = 0$. For the lower-left chamber, suppose the optimal face is $F_2 = \{(0, 0)\}$. The chamber closure has two codimension-one faces whose closures contain it, so there are two choices for F_1. For $0 < \epsilon_1 \ll \epsilon_2$, the plotted points are the lexicographically perturbed points in each chamber that appear in the proof of Theorem 2, determined by the four sign vectors $\alpha \in \{\pm 1\}^2$ for each flag.

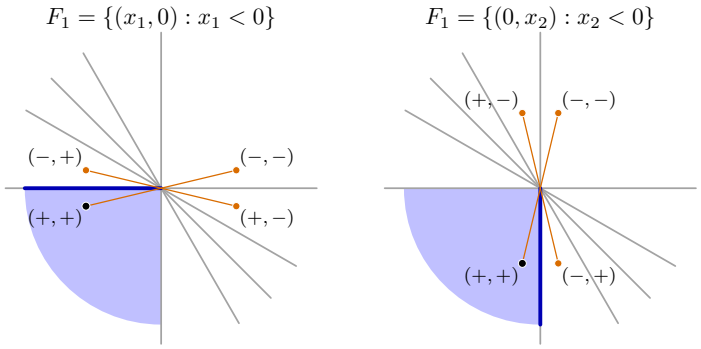


Figure 3: The two possible flags at the starting chamber in a two-dimensional example. In each panel $0 < \epsilon_1 \ll \epsilon_2$; the black point is $\alpha = (+, +)$, which recovers the source chamber, and the orange points are the three target signs.

We now prove Theorem 2. The proof first turns a flag into a positive flag-adapted coordinate basis. In those coordinates, separated perturbations of a point in F, with signs indexed by $\alpha \in \{\pm 1\}^r$, produce points in the faces $F_i(\alpha)$ and hence determine the sign-indexed chambers.

4.2 Positive adapted coordinates

Definition 19. *Let* $\mathcal{P}$ *be a locally finite affine hyperplane arrangement, and let* F *be a face of positive codimension* r. *Define* $\mathrm{Aff}_0(F) := \{\rho \in \mathrm{Aff}(A, \mathbb{R}) : F \subseteq Z(\rho)\}$.

Lemma 13. *(→ proof) Assume the setting of Definition 19. For* $\ell \in \mathrm{Ann}(F)$, *choose any* $x_F \in F$ *and set* $\rho_\ell(y) := \ell(y - x_F)$. *This is independent of the choice of* $x_F \in F$, *and the map* $\mathrm{Ann}(F) \to \mathrm{Aff}_0(F)$, $\ell \mapsto \rho_\ell$, *is an isomorphism of real vector spaces. In particular,* $\mathrm{Aff}_0(F)$ *has dimension* r, *an ordered list* $\rho_{\ell_1}, \ldots, \rho_{\ell_r}$ *is a basis of* $\mathrm{Aff}_0(F)$ *exactly when* $\ell_1, \ldots, \ell_r$ *is a basis of* $\mathrm{Ann}(F)$, *and every* $q \in \mathfrak{Q}(F, \mathcal{P})$ *has a unique expansion in any ordered basis of* $\mathrm{Aff}_0(F)$.

Proof of Lemma 13. If $x_F, x'_F \in F$, then $x'_F - x_F \in \mathrm{lin}(F)$, so $\ell(y - x_F) = \ell(y - x'_F)$. Thus ρ_ℓ is well-defined and vanishes on F. The map $\ell \mapsto \rho_\ell$ is linear.

If $\rho_\ell = 0$, then for all $y, y' \in A$ we have $\ell(y - y') = \rho_\ell(y) - \rho_\ell(y') = 0$. Hence $\ell = 0$, so the map is injective. For surjectivity, let $q \in \mathrm{Aff}_0(F)$. If $x, y \in F$, then $0 = q(x) - q(y) = dq(x - y)$. Since such differences span $\mathrm{lin}(F)$, dq vanishes on $\mathrm{lin}(F)$. For any $x_F \in F$, $q(y) - \rho_{dq}(y) = q(y) - dq(y - x_F) = q(x_F) = 0$, so $q = \rho_{dq}$.

Finally, $\dim \mathrm{Ann}(F) = \mathrm{codim}(F) = r$, and the remaining statements follow from the isomorphism and the inclusion $\mathfrak{Q}(F, \mathcal{P}) \subseteq \mathrm{Aff}_0(F)$. □

Definition 20 (Positive flag-adapted coordinates)**.** *Assume the setting of Definition 16, and denote the flag by $FL = (F_r, \dots, F_0)$. An ordered basis $\rho = (\rho_r, \dots, \rho_1)$ of $\mathrm{Aff}_0(F)$ is FL-adapted if, for each $i \in [1, r]$, $\rho_i, \rho_{i-1}, \dots, \rho_1$ is a basis of $\mathrm{Aff}_0(F_i)$. It is positive if, in addition, $\mathrm{sign}_{F_{i-1}}(\rho_i) = +1$ for $i = 1, \dots, r$.*

Lemma 14. *Any two positive FL-adapted coordinate bases differ by a triangular change-of-basis matrix with positive diagonal entries with respect to the order $r, \dots, 1$. Consequently the orientation of $d\rho_r \wedge \cdots \wedge d\rho_1$ is independent of the chosen positive adapted basis.*

Proof. Let ρ and ρ' be two positive FL-adapted bases. Since $\rho'_i, \rho'_{i-1}, \dots, \rho'_1$ and $\rho_i, \rho_{i-1}, \dots, \rho_1$ are both bases of $\mathrm{Aff}_0(F_i)$, each ρ'_i is a linear combination of $\rho_i, \rho_{i-1}, \dots, \rho_1$, and the diagonal coefficient is nonzero. On F_{i-1}, all lower coordinates $\rho_1, \dots, \rho_{i-1}$ vanish, while both ρ_i and ρ'_i are positive. Hence the diagonal coefficient is positive. The determinant is therefore positive. □

Lemma 15. *(→ proof) Every flag $FL = (F_r, \dots, F_0)$ admits a positive FL-adapted coordinate basis.*

Proof of Lemma 15. By Lemma 13, $\dim \mathrm{Aff}_0(F_i) = i$, and the inclusions $\mathrm{Aff}_0(F_{i-1}) \subsetneq \mathrm{Aff}_0(F_i)$ follow from $F_i \subseteq \overline{F_{i-1}}$ and the codimension increase. Choose $\widetilde{\rho}_i \in \mathrm{Aff}_0(F_i) \setminus \mathrm{Aff}_0(F_{i-1})$. Then $\widetilde{\rho}_i$ is nonzero on F_{i-1}, with constant sign there. To see this, choose an arrangement hyperplane containing F_i but not F_{i-1}, which exists by Lemma 3. Its intersection with $\mathrm{aff}(F_{i-1})$ and the zero set of $\widetilde{\rho}_i|_{\mathrm{aff}(F_{i-1})}$ are both hyperplanes containing $\mathrm{aff}(F_i)$, hence are equal. The face F_{i-1} lies on one open side of this hyperplane. Replace $\widetilde{\rho}_i$ by its negative if necessary so that this sign is positive. The resulting $\rho_i, \rho_{i-1}, \dots, \rho_1$ form a basis of $\mathrm{Aff}_0(F_i)$ for each i, so $\rho_r, \dots, \rho_1$ is positive and FL-adapted. □

Positive adapted coordinates recover the preferred transverse side for each adjacent pair of flag faces. If $FL = (F_r, \dots, F_0)$, ρ is positive and FL-adapted, $i \in [1, r]$, and $v \in \mathrm{lin}(F_{i-1})$, then $d\rho_i(v) > 0$ characterizes the side pointing from F_i toward F_{i-1}. Indeed, ρ_i vanishes on F_i and is positive on F_{i-1}, while the preceding triangular-change lemma shows that the criterion is independent of the chosen positive adapted coordinates.

Corollary 1. *(→ proof) Assume the setting of Definition 16, denote the flag by FL, and let* gi *be the greatest-index function from Definition 17. Let $\rho^{FL} = (\rho^{FL}_r, \dots, \rho^{FL}_1)$ be any positive FL-adapted coordinate basis. If $q = \sum_{i=1}^{r} a^{FL}_i(q)\rho^{FL}_i$ is the expansion of $q \in \mathfrak{Q}(F, \mathcal{P})$ in this basis, then $a^{FL}_i(q) = 0$ for all $i > \mathrm{gi}(q)$, $a^{FL}_{\mathrm{gi}(q)}(q) \neq 0$, and $\mathrm{sign}(a^{FL}_{\mathrm{gi}(q)}(q)) = \mathrm{sign}_{\mathfrak{c}}(q)$. Thus $\mathrm{gi}(q)$ is the greatest index with nonzero coefficient in any positive flag-adapted coordinate basis.*

Proof of Corollary 1. Since $Z(q) \in \mathcal{H}_{\mathrm{gi}(q)}$, we have $q \in \mathfrak{Q}(F_{\mathrm{gi}(q)}, \mathcal{P})$, so the linear part of q lies in $\mathrm{Ann}(F_{\mathrm{gi}(q)})$. Under Lemma 13, this means $q \in \mathrm{Aff}_0(F_{\mathrm{gi}(q)}) = \mathrm{span}(\rho^{FL}_{\mathrm{gi}(q)}, \dots, \rho^{FL}_1)$. Therefore $a^{FL}_i(q) = 0$ for all $i > \mathrm{gi}(q)$. If $a^{FL}_{\mathrm{gi}(q)}(q) = 0$, then the expansion uses only ρ^{FL}_i with $i < \mathrm{gi}(q)$. These terms vanish on $F_{\mathrm{gi}(q)-1}$, so $Z(q) \in \mathcal{H}(F_{\mathrm{gi}(q)-1}, \mathcal{P})$, contradicting $Z(q) \in \mathcal{H}_{\mathrm{gi}(q)}$.

On $F_{\mathrm{gi}(q)-1}$, the terms ρ^{FL}_i with $i < \mathrm{gi}(q)$ vanish, while $\rho^{FL}_{\mathrm{gi}(q)}$ is positive by Definition 20. Hence the sign of q on $F_{\mathrm{gi}(q)-1}$ is the sign of $a^{FL}_{\mathrm{gi}(q)}(q)$. Since $F_{\mathrm{gi}(q)-1} \subseteq \bar{\mathfrak{c}}$ and $q \notin \mathfrak{Q}(F_{\mathrm{gi}(q)-1}, \mathcal{P})$, this sign is $\mathrm{sign}_{\mathfrak{c}}(q)$. □

4.3 Separated perturbations and target chambers

The decomposition from Corollary 1 shows how each $q \in \mathfrak{Q}(\mathcal{P})$ is affected as we perturb a point of F by prescribing positive adapted coordinate values on widely separated scales. The separation condition is designed so that for each $q \in \mathfrak{Q}(F, \mathcal{P})$, the sign of q is controlled by its greatest nonzero coordinate. This is what allows a sign vector $\alpha \in \{\pm 1\}^r$ to determine the chamber $\mathfrak{b}_{FL}(\alpha)$. Before specializing to the flag basis, we record certain notation for an arbitrary ordered basis of $\mathrm{Aff}_0(F)$.

Definition 21. *Assume the setting of Definition 19. Fix an ordered basis $\rho = (\rho_r, \dots, \rho_1)$ of $\mathrm{Aff}_0(F)$. Then every $q \in \mathfrak{Q}(F, \mathcal{P})$ has a unique expansion $q = \sum_{i=1}^{r} a_i^\rho(q)\rho_i$. Define $\mathrm{gi}_\rho(q)$ to be the greatest index i such that $a_i^\rho(q) \neq 0$, and define $\mathrm{gc}_\rho(q) := a^\rho_{\mathrm{gi}_\rho(q)}(q)$ and $\mathrm{tail}_\rho(q) := q - \mathrm{gc}_\rho(q)\rho_{\mathrm{gi}_\rho(q)} = \sum_{i=1}^{\mathrm{gi}_\rho(q)-1} a_i^\rho(q)\rho_i$.*

$\mathrm{gi}_\rho(q)$ is the greatest index of q in the ordered basis ρ, and $\mathrm{gc}_\rho(q)$ is the coefficient of that greatest-index term.

Definition 22. *Assume the setting of Definition 21. Choose a finite set $\mathcal{Q}_F \subseteq \mathfrak{Q}(F, \mathcal{P})$ containing one representative for each hyperplane in $\{Z(q) : q \in \mathfrak{Q}(F, \mathcal{P})\}$. Such a set exists by Definition 8. Choose $\tau_{0,\rho} > 0$ so large that for all $q \in \mathcal{Q}_F$,*

$$\frac{\sum_{i=1}^{\mathrm{gi}_\rho(q)-1} |a_i^\rho(q)|\tau^{i-\mathrm{gi}_\rho(q)}}{|\,\mathrm{gc}_\rho(q)|} \leq \frac{1}{2}$$

holds whenever $\tau \geq \tau_{0,\rho}$. Such a choice is possible because every exponent $i - \mathrm{gi}_\rho(q)$ in the numerator is negative, so the displayed ratio tends to zero as $\tau \to \infty$, and $\mathcal{Q}_F$ is finite. Since this inequality is invariant under multiplying q by a nonzero scalar, it then holds for all $q \in \mathfrak{Q}(F, \mathcal{P})$.

Any uniform bound strictly below 1 would suffice. The choice $\frac{1}{2}$ ensures that the leading term and the full affine form have the same sign in the perturbation argument below.

For the arguments below, choose a positive FL-adapted coordinate basis ρ^{FL}. Corollary 1 shows that $\mathrm{gi}_{\rho^{FL}}(q) = \mathrm{gi}(q)$, where $\mathrm{gi}(q)$ is the combinatorial index from Definition 17. After this specialization, we suppress ρ^{FL} from the notation and write ρ_i, $a_i(q)$, $\mathrm{gc}(q)$, $\mathrm{tail}(q)$, and τ_0.

Definition 23. *We say that $\epsilon \in \mathbb{R}^r_{>0}$ is τ-separated for $\tau > 0$ if, for all $i \in [1, r-1]$, $\tau\epsilon_i < \epsilon_{i+1}$.*

We now construct the sign-indexed chambers by assigning the positive adapted coordinates on τ_0-separated scales. The sign vector $\alpha \in \{\pm 1\}^r$ controls the sign of the greatest-index term for each $q \in \mathfrak{Q}(F, \mathcal{P})$, and therefore controls the sign of q itself. It also determines an associated codimension i face $F_i(\alpha)$ for each $i \in [0, r]$, not only the chamber $F_0(\alpha)$.

Proof of Theorem 2. Fix $\alpha \in \{\pm 1\}^r$. Choose positive FL-adapted coordinates $\rho = (\rho_r, \dots, \rho_1)$. For each $i \in [1, r]$, the affine forms $\rho_i, \dots, \rho_1$ form a basis of $\mathrm{Aff}_0(F_i)$, so $\mathrm{aff}(F_i) = \bigcap_{m=1}^{i} Z(\rho_m)$. Let $D_\rho : V_A \to \mathbb{R}^r$ be the surjective linear map $D_\rho(v) = (d\rho_r(v), \dots, d\rho_1(v))$, and choose a linear section $\sigma : \mathbb{R}^r \to V_A$. Let $e_r, \dots, e_1$ be the standard basis of $\mathbb{R}^r$, indexed in this coordinate order.

Let $x \in F$. By Lemma 8, choose a precompact convex open neighborhood U of x on which every $q \in \mathfrak{Q}(\mathcal{P}) \setminus \mathfrak{Q}(F, \mathcal{P})$ has constant sign $\mathrm{sign}_{\mathfrak{c}}(q)$.

Choose $\epsilon \in \mathbb{R}^r_{>0}$ that is τ_0-separated and small enough for the points below to lie in U. For $i \in [0, r]$, set $y_i(\alpha) := x + \sigma(\sum_{j=i+1}^{r} \epsilon_j \alpha_j e_j)$. By our choice of ϵ, $y_i(\alpha) \in U$ for all i and all α.

We now compute the signs of all arrangement forms at $y_i(\alpha)$. First, $y_i(\alpha) \in \mathrm{aff}(F_i)$: for $i = 0$ this is $\mathrm{aff}(F_0) = A$, while for $i \geq 1$ and $m \leq i$ we have $\rho_m(y_i(\alpha)) = 0$, so $y_i(\alpha) \in \bigcap_{m=1}^{i} Z(\rho_m) = \mathrm{aff}(F_i)$. Hence, if $q \in \mathfrak{Q}(F_i, \mathcal{P})$, then $q(y_i(\alpha)) = 0$.

Next let $q \in \mathfrak{Q}(F,\mathcal{P}) \setminus \mathfrak{Q}(F_i,\mathcal{P})$. Then $i < \mathrm{gi}(q)$, $\rho_{\mathrm{gi}(q)}(y_i(\alpha)) = \alpha_{\mathrm{gi}(q)}\epsilon_{\mathrm{gi}(q)}$, and $q = \mathrm{gc}(q)\rho_{\mathrm{gi}(q)} + \mathrm{tail}(q)$, so

$$q(y_i(\alpha)) = \mathrm{gc}(q)\rho_{\mathrm{gi}(q)}(y_i(\alpha))\Big(1 + \frac{\mathrm{tail}(q)(y_i(\alpha))}{\mathrm{gc}(q)\rho_{\mathrm{gi}(q)}(y_i(\alpha))}\Big).$$

The separation condition gives

$$\Big|\frac{\mathrm{tail}(q)(y_i(\alpha))}{\mathrm{gc}(q)\rho_{\mathrm{gi}(q)}(y_i(\alpha))}\Big| \le \frac{\sum_{m=1}^{\mathrm{gi}(q)-1}|a_m(q)|\epsilon_m}{|\,\mathrm{gc}(q)|\epsilon_{\mathrm{gi}(q)}} < \frac{\sum_{m=1}^{\mathrm{gi}(q)-1}|a_m(q)|\tau_0^{m-\mathrm{gi}(q)}}{|\,\mathrm{gc}(q)|} \le \frac{1}{2}.$$

Thus the last factor is positive, and $\mathrm{sign}(q(y_i(\alpha))) = \mathrm{sign}(\mathrm{gc}(q))\alpha_{\mathrm{gi}(q)}$. By Corollary 1, $\mathrm{sign}(\mathrm{gc}(q)) = \mathrm{sign}_{\mathfrak{c}}(q)$. Therefore $\mathrm{sign}(q(y_i(\alpha))) = \alpha_{\mathrm{gi}(q)}\,\mathrm{sign}_{\mathfrak{c}}(q)$.

Finally, if $q \in \mathfrak{Q}(\mathcal{P}) \setminus \mathfrak{Q}(F,\mathcal{P})$, then $y_i(\alpha) \in U$, so $\mathrm{sign}(q(y_i(\alpha))) = \mathrm{sign}_{\mathfrak{c}}(q)$.

Let $F_i(\alpha)$ be the face containing $y_i(\alpha)$. For $i = 0$, no arrangement form vanishes at $y_0(\alpha)$, so $y_0(\alpha)$ lies in a unique chamber; denote this chamber by $\mathfrak{b}_{FL}(\alpha)$, so $F_0(\alpha) = \mathfrak{b}_{FL}(\alpha)$. The three sign computations give exactly the displayed sign formula. In particular, $F_i(\alpha)$ has the same vanishing forms as F_i, so $\mathrm{codim}(F_i(\alpha)) = \mathrm{codim}(F_i) = i$. If $\alpha = 1^r$, then the displayed sign formula is exactly the sign pattern of F_i, hence $F_i(1^r) = F_i$.

It remains only to check that the faces form a flag. For $i \in [1,r]$, the sign pattern of $F_i(\alpha)$ is obtained from the sign pattern of $F_{i-1}(\alpha)$ by changing exactly the forms q with $Z(q) \in \mathcal{H}_i$ from the nonzero sign $\alpha_i\,\mathrm{sign}_{\mathfrak{c}}(q)$ to 0. Thus $F_i(\alpha) \subseteq \overline{F_{i-1}(\alpha)}$. The containment is proper because $\mathcal{H}_i \neq \emptyset$. □

Theorem 2 characterizes $F_i(\alpha)$ by the sign of each $q \in \mathfrak{Q}(\mathcal{P})$. We also have a characterization by identifying a specific point in the face.

Corollary 2. *(→ proof) Assume the setting of Theorem 2. Choose a positive FL-adapted coordinate basis $\rho = (\rho_r,\dots,\rho_1)$, and choose a linear section $\sigma : \mathbb{R}^r \to V_A$ of $D_\rho = (d\rho_r,\dots,d\rho_1)$, meaning $D_\rho \circ \sigma = \mathrm{Id}_{\mathbb{R}^r}$. Let $e_r,\dots,e_1$ be the standard basis of $\mathbb{R}^r$, indexed in this coordinate order, and let $x \in F$. For every $\alpha \in \{\pm 1\}^r$, every $\epsilon \in \mathbb{R}^r_{>0}$ sufficiently small and τ_0-separated, and every $i \in [0,r]$, $F_i(\alpha)$ contains $x + \sigma(\sum_{j=i+1}^r \epsilon_j\alpha_j e_j)$.*

Proof of Corollary 2. This is shown in the proof of Theorem 2. □

5 Desiderata for flag selection rules

The target construction turns a choice of flag in each chamber into a directed graph. We require the chosen flag to be based at the optimal face whenever the objective is bounded above. This requirement forces weak increase of the objective value along edges. We then distinguish acyclicity, progress beyond the starting value, and escape beyond every prescribed threshold, including pathwise and uniform versions of the latter two properties. We prove the relationships among these conditions and show that finitely many normal spaces suffice to deduce uniform escape from acyclicity.

5.1 Optimal-face selection and the induced graph

Definition 24. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. Let $\omega \in \mathrm{Aff}(A,\mathbb{R}) \setminus \mathrm{Con}(A,\mathbb{R})$. Let $\mathfrak{c} \in \Omega(\mathcal{P})$. Define the linear program $\lambda(\mathfrak{c},\omega) := \sup_{x\in\overline{\mathfrak{c}}}\omega(x)$.*

Definition 25. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A. Let $\omega \in \mathrm{Aff}(A,\mathbb{R}) \setminus \mathrm{Con}(A,\mathbb{R})$. Let $\mathfrak{c} \in \Omega(\mathcal{P})$ be such that $\lambda(\mathfrak{c},\omega) < \infty$. The linear program achieves its optimal value and the optimal set is a face of $\overline{\mathfrak{c}}$ by Theorem 1. Define the optimal face $F(\mathfrak{c},\omega)$ to be the face whose closure is $\overline{F(\mathfrak{c},\omega)} = \{x \in \overline{\mathfrak{c}} : \omega(x) = \lambda(\mathfrak{c},\omega)\}$. This optimizer set does not meet $\mathfrak{c}$. Indeed, if $x \in \mathfrak{c}$, choose $v \in V_A$ with $d\omega(v) > 0$. For all sufficiently small $t > 0$, one has $x + tv \in \mathfrak{c}$ and $\omega(x + tv) > \omega(x)$, so x is not optimal. Therefore*

$$\overline{F(\mathfrak{c},\omega)} \subseteq \partial\mathfrak{c} \subseteq \overline{\mathfrak{c}} \cap |\mathcal{P}|.$$

Definition 26 (Flag selection strategy)**.** *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement in A, and let $\omega \in \mathrm{Aff}(A,\mathbb{R}) \setminus \mathrm{Con}(A,\mathbb{R})$. A flag selection strategy for $(\mathcal{P},\omega)$ assigns to every chamber $\mathfrak{c} \in \Omega(\mathcal{P})$ a flag with top $\mathfrak{c}$, $\mathcal{F}(\mathfrak{c}) = (F_r(\mathfrak{c}), \ldots, F_0(\mathfrak{c}) = \mathfrak{c})$ subject to the following rules. If $\lambda(\mathfrak{c},\omega) < \infty$, then the flag must have the optimal face as its base: $F_r(\mathfrak{c}) = F(\mathfrak{c},\omega)$. If $\lambda(\mathfrak{c},\omega) = \infty$, then $\mathcal{F}(\mathfrak{c})$ is the codimension-zero flag ($F_0(\mathfrak{c}) = \mathfrak{c}$).*

Definition 27. *Let $\mathcal{F}$ be a flag selection strategy for $(\mathcal{P},\omega)$. Define a directed graph $G_{\mathcal{F}}$ whose vertices are $\Omega(\mathcal{P})$ as follows. For each chamber $\mathfrak{c}$, let $FL = \mathcal{F}(\mathfrak{c})$, and let r be the codimension of its base face. If $r = 0$, add no edges. If $r > 0$, add an edge from $\mathfrak{c}$ to $\mathfrak{b}_{FL}(\alpha)$ for each $\alpha \in \{\pm 1\}^r \setminus \{1^r\}$.*

Every chamber has finite out-degree: if $\mathcal{F}(\mathfrak{c})$ has codimension r, then there are $2^r - 1$ outgoing edges. By Definition 18, these outgoing edges are distinct and are never self-loops.

Lemma 16. *Let $\mathcal{F}$ be a flag selection strategy for $(\mathcal{P},\omega)$, and let $\mathfrak{c} \to \mathfrak{d}$ be an edge in $G_{\mathcal{F}}$. Then $\lambda(\mathfrak{d},\omega) \geq \lambda(\mathfrak{c},\omega)$. Moreover, if equality holds, then the optimal face weakly grows: $\overline{F(\mathfrak{c},\omega)} \subseteq \overline{F(\mathfrak{d},\omega)}$.*

Proof. Because $\mathfrak{c}$ has an outgoing edge, $\lambda(\mathfrak{c},\omega) < \infty$. By the definition of a flag selection strategy, $\mathcal{F}(\mathfrak{c})$ has base $F(\mathfrak{c},\omega)$. Every target chamber contains that base face in its closure by Definition 18, so $F(\mathfrak{c},\omega) \subseteq \overline{\mathfrak{d}}$. Thus, for any $x \in F(\mathfrak{c},\omega)$, $\lambda(\mathfrak{d},\omega) \geq \omega(x) = \lambda(\mathfrak{c},\omega)$.
If equality holds, then every point of $F(\mathfrak{c},\omega)$ lies in $\overline{\mathfrak{d}}$ and has ω-value $\lambda(\mathfrak{d},\omega)$. Hence $\overline{F(\mathfrak{c},\omega)} \subseteq \overline{F(\mathfrak{d},\omega)}$. □

Lemma 17. *Let $\mathcal{F}$ be a flag selection strategy for $(\mathcal{P},\omega)$. Let $(\mathfrak{c}_0, \ldots, \mathfrak{c}_m)$ be a directed path in $G_{\mathcal{F}}$ with $\lambda(\mathfrak{c}_0,\omega) < \infty$. If $\lambda(\mathfrak{c}_m,\omega) = \lambda(\mathfrak{c}_0,\omega)$, then all chambers on the path have this same λ-value and*

$$\overline{F(\mathfrak{c}_0,\omega)} \subseteq \overline{F(\mathfrak{c}_i,\omega)} \qquad (0 \leq i \leq m).$$

In particular, each $\mathfrak{c}_i$ lies in the finite chamber star $\mathrm{Star}_\Omega(F(\mathfrak{c}_0,\omega))$. If the path is a directed cycle, then all optimal face closures along the cycle are equal.

Proof. Weak increase of λ along edges, Lemma 16, forces every chamber on the path to have the common endpoint value. Applying the equality case of the same lemma along each edge and composing the inclusions gives the displayed containment. This containment puts $F(\mathfrak{c}_0,\omega)$ in the closure of every chamber on the path, so the finite-star assertion follows from Lemma 12. If the path is a directed cycle, the displayed inclusions also hold after starting at any vertex of the cycle, so the optimal face closures are all mutually contained and hence equal. □

5.2 Progress and escape conditions

With the graph defined, we distinguish four ways in which paths may eventually increase the value of λ.

Definition 28 (Progress and escape conditions)**.**

1. *$\mathcal{F}$ is pathwise progressive if, for every finite-λ chamber $\mathfrak{d}_0$, every infinite path $(\mathfrak{d}_0, \mathfrak{d}_1, \ldots)$ in $G_{\mathcal{F}}$ eventually reaches a chamber of larger value: there is some $i \geq 0$ such that $\lambda(\mathfrak{d}_i, \omega) > \lambda(\mathfrak{d}_0, \omega)$.*
2. *$\mathcal{F}$ is uniform progressive if, for every finite-λ chamber $\mathfrak{d}$, there exists $L_0 \in \mathbb{Z}_{\geq 1}$ such that every path $(\mathfrak{d} = \mathfrak{c}_0, \ldots, \mathfrak{c}_L)$ in $G_{\mathcal{F}}$ of length $L \geq L_0$ satisfies $\lambda(\mathfrak{c}_L, \omega) > \lambda(\mathfrak{d}, \omega)$.*
3. *$\mathcal{F}$ is pathwise escaping if, for every finite-λ chamber $\mathfrak{d}_0$, every infinite path $(\mathfrak{d}_0, \mathfrak{d}_1, \ldots)$ in $G_{\mathcal{F}}$, and every $\eta \in \mathbb{R}$, there is some $i \geq 0$ such that $\lambda(\mathfrak{d}_i, \omega) > \eta$.*
4. *$\mathcal{F}$ is uniform escaping if, for every finite-λ chamber $\mathfrak{d}$ and every $\eta \in \mathbb{R}$, there exists $L_0 \in \mathbb{Z}_{\geq 1}$ such that every path $(\mathfrak{d} = \mathfrak{c}_0, \ldots, \mathfrak{c}_L)$ in $G_{\mathcal{F}}$ of length $L \geq L_0$ satisfies $\lambda(\mathfrak{c}_L, \omega) > \eta$.*

Lemma 18. *Let $\mathcal{F}$ be a flag selection strategy for $(\mathcal{P}, \omega)$. Then $\mathcal{F}$ is uniform progressive if and only if it is pathwise progressive. Also, $\mathcal{F}$ is uniform escaping if and only if it is pathwise escaping. If $\mathcal{F}$ is pathwise escaping or uniform escaping, then $\mathcal{F}$ is pathwise progressive and uniform progressive.*

Proof. Uniform statements immediately imply their pathwise versions. We prove the converse in a common form for both the progressive and escaping properties. Fix a finite-λ chamber $\mathfrak{d}_0$ and a threshold $\eta \in \mathbb{R}$, and assume every infinite path from $\mathfrak{d}_0$ eventually has value $> \eta$. If $\lambda(\mathfrak{d}_0, \omega) > \eta$, Lemma 16 gives the uniform bound $L_0 = 1$.

Assume $\lambda(\mathfrak{d}_0, \omega) \leq \eta$. If no such L_0 existed, then for arbitrarily large L there would be a path $(\mathfrak{d}_0, \ldots, \mathfrak{d}_L)$ with $\lambda(\mathfrak{d}_L, \omega) \leq \eta$. By Lemma 16, every vertex of such a path has value at most η. The tree of finite paths from $\mathfrak{d}_0$ whose vertices all have value at most η is finitely branching and has arbitrarily long branches, so König's infinity lemma [1, Ch. 8] gives an infinite path which never exceeds η, a contradiction.

Applying this with $\eta = \lambda(\mathfrak{d}_0, \omega)$ proves that pathwise progressive implies uniform progressive. Applying it for each $\eta \in \mathbb{R}$ proves that pathwise escaping implies uniform escaping. Once the uniform and pathwise versions are equivalent, the last statement is immediate. □

Corollary 3. *Assume that $G_{\mathcal{F}}$ is acyclic. Then $\mathcal{F}$ is pathwise progressive and uniform progressive.*

Proof. We show that $\mathcal{F}$ is pathwise progressive. Let $(\mathfrak{d}_0, \mathfrak{d}_1, \ldots)$ be an infinite path in $G_{\mathcal{F}}$, with $\lambda(\mathfrak{d}_0, \omega) = \tau < \infty$. Suppose, for contradiction, that no chamber in the path has λ-value $> \tau$. Lemma 17 applied to every finite prefix puts all chambers of the path in the finite star of $F(\mathfrak{d}_0, \omega)$. Hence some chamber repeats, producing a directed cycle, contrary to acyclicity. Thus $\mathcal{F}$ is pathwise progressive, and Lemma 18 gives uniform progress. □

5.3 Escaping with finitely many normal spaces

We showed that if $G_{\mathcal{F}}$ is acyclic, then $\mathcal{F}$ is pathwise progressive and uniform progressive. We leave open whether acyclicity of $G_{\mathcal{F}}$ alone implies that $\mathcal{F}$ is pathwise escaping and uniform escaping. However, in this subsection we provide a sufficient condition under which this holds: that $\mathcal{P}$ has finitely many normal spaces. In that case, the closure of every chamber is a polyhedron. The simplex algorithm supplies a path within each polyhedral chamber, and the hypothesis forces this path to use only finitely many possible directions. This gives a uniform lower bound on how much ω increases per unit distance traveled. We make this argument precise in this section.

Lemma 19. *(→ proof) If $\mathcal{P}$ is a locally finite affine hyperplane arrangement in A with finitely many normal spaces, then the closure of every chamber of $\mathcal{P}$ is a polyhedron.*

Proof of Lemma 19. Fix a chamber $\mathfrak{c}$, and put $C = \overline{\mathfrak{c}}$. Fix one normal space, choose $\nu \in V_A^* \setminus \{0\}$ spanning it, and choose $a_\nu \in \mathrm{Aff}(A, \mathbb{R})$ with linear part ν. Write the hyperplanes with this normal space as $H_a = \{x \in A : a_\nu(x) = a\}$ for $a \in S_\nu$.

Suppose C meets three hyperplanes H_a, H_b, H_c with $a < b < c$. Let $q \in \mathfrak{Q}(\mathcal{P})$ be such that $H_b = Z(q)$. Then q attains both positive and negative values on C, since C meets both H_a and H_c, contradicting the weak-sign description of chamber closures in Lemma 4.

Thus C meets at most two hyperplanes with normal space ν. Since there are only finitely many normal spaces, C meets only finitely many hyperplanes; call this finite set $\mathcal{H}_C$.

For each $H \in \mathcal{H}_C$, choose $q_H \in \mathfrak{Q}(\mathcal{P})$ with $Z(q_H) = H$, multiplied by -1 if necessary so that $q_H > 0$ on $\mathfrak{c}$. Put

$$C' := \{y \in A : q_H(y) \geq 0 \text{ for all } H \in \mathcal{H}_C\}.$$

The closed-face description of $\overline{\mathfrak{c}}$ gives $C \subseteq C'$. We prove the reverse inclusion. Suppose $y \notin C$, and choose $x \in \mathfrak{c}$. The segment $[x, y]$ meets only finitely many arrangement hyperplanes, by local finiteness. Among the hyperplanes whose chamber-side inequality is violated at y, choose one whose first intersection time with the segment is minimal. At that first intersection point, no chamber-side inequality has yet become negative, so the point lies in C. Hence the chosen hyperplane lies in $\mathcal{H}_C$, while y violates its defining inequality for C'. Therefore $y \notin C'$. Thus $C = C'$, a finite intersection of closed halfspaces. □

Lemma 20. *(→ proof) Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, and let $\mathfrak{c}, \mathfrak{b} \in \Omega(\mathcal{P})$. If $x \in \overline{\mathfrak{c}} \cap \overline{\mathfrak{b}}$ is a vertex of $\overline{\mathfrak{c}}$, then x is a vertex of $\overline{\mathfrak{b}}$.*

Proof of Lemma 20. Fix any norm on V_A. Let $\mathcal{H}_x$ be the set of hyperplanes of the arrangement $\mathcal{P}$ which contain x. The set $\mathcal{H}_x$ is finite, since every member meets the compact set $\{x\}$. Again by local finiteness, there is $\epsilon > 0$ such that every hyperplane of the arrangement which meets $B_\epsilon(x)$ lies in $\mathcal{H}_x$.

We first show that $\bigcap_{H \in \mathcal{H}_x} H = \{x\}$. Otherwise, this intersection contains a line $x + \mathbb{R}u$ with $u \neq 0$. For sufficiently small $t > 0$, the points $x + tu$ and $x - tu$ lie in $B_\epsilon(x)$. They lie on every hyperplane in $\mathcal{H}_x$ and on the same side of every remaining hyperplane as x. Since $x \in \overline{\mathfrak{c}}$, this implies $x + tu, x - tu \in \overline{\mathfrak{c}}$, contradicting that x is a vertex of $\overline{\mathfrak{c}}$.

By Definition 10, it remains to show that every segment in $\overline{\mathfrak{b}}$ containing x in its relative interior is trivial. Suppose $x = ty + (1-t)z$ for some $y, z \in \overline{\mathfrak{b}}$ and $t \in (0, 1)$. Replacing y, z by $x + \delta(y - x), x + \delta(z - x)$ for sufficiently small $0 < \delta \leq 1$, we may assume $y, z \in B_\epsilon(x)$ while preserving the displayed convex combination. For each $H \in \mathcal{H}_x$, $\overline{\mathfrak{b}}$ is contained in one of the two closed halfspaces bounded by H, and $x \in H$. The displayed convex combination can lie on H only if both y and z lie on H. Thus $y, z \in \bigcap_{H \in \mathcal{H}_x} H = \{x\}$. Therefore $y = z = x$, so x is a vertex of $\overline{\mathfrak{b}}$. □

The quotient below removes the directions invisible to all arrangement hyperplanes. After this reduction the normal spaces span the dual, so finite-normal chamber closures are pointed polyhedra and the vertex-neighbor argument can be applied.

Lemma 21 (Normal quotient reduction)**.** *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, and let $\omega \in \mathrm{Aff}(A, \mathbb{R}) \setminus \mathrm{Con}(A, \mathbb{R})$. Set $W_\mathcal{P} := \{v \in V_A : dq(v) = 0 \text{ for all } q \in \mathfrak{Q}(\mathcal{P})\}$, and let $\rho : A \to A^\rho := A/W_\mathcal{P}$ be the quotient map. For each $q \in \mathfrak{Q}(\mathcal{P})$, there is a unique $q^\rho \in \mathrm{Aff}(A^\rho, \mathbb{R}) \setminus \mathrm{Con}(A^\rho, \mathbb{R})$ such that $q = q^\rho \circ \rho$. The collection $\mathcal{P}^\rho := \{\rho(H) : H \in \mathcal{P}\}$ is a locally finite affine hyperplane arrangement in A^ρ whose normal spaces span $V_{A^\rho}^*$. The braces mean that $\mathcal{P}^\rho$ is*

treated as a set of image hyperplanes; if the original collection is presented with repeated hyperplanes, repetitions do not add sign data. If $\mathcal{P}$ has finitely many normal spaces, then so does $\mathcal{P}^\rho$.

If some chamber has finite $\lambda(-,\omega)$, then $d\omega$ vanishes on $W_{\mathcal{P}}$. Let $\omega^\rho \in \mathrm{Aff}(A^\rho,\mathbb{R})$ be the unique affine functional with $\omega = \omega^\rho \circ \rho$. The map $\mathfrak{c} \mapsto \mathfrak{c}^\rho := \rho(\mathfrak{c})$ is a bijection from chambers of $\mathcal{P}$ to chambers of $\mathcal{P}^\rho$, and $\mathfrak{c} = \rho^{-1}(\mathfrak{c}^\rho)$, $\overline{\mathfrak{c}} = \rho^{-1}(\overline{\mathfrak{c}^\rho})$, and $\rho(\overline{\mathfrak{c}}) = \overline{\mathfrak{c}^\rho}$. Moreover $\lambda(\mathfrak{c},\omega) = \lambda(\mathfrak{c}^\rho,\omega^\rho)$. When this common value is finite, optimal faces correspond by $\overline{F(\mathfrak{c},\omega)} = \rho^{-1}\big(\overline{F(\mathfrak{c}^\rho,\omega^\rho)}\big)$ and $\rho\big(\overline{F(\mathfrak{c},\omega)}\big) = \overline{F(\mathfrak{c}^\rho,\omega^\rho)}$.

Proof. We first construct the affine forms on the quotient and identify the image hyperplanes. For $q \in \mathfrak{Q}(\mathcal{P})$, the differential dq vanishes on $W_{\mathcal{P}}$, so q is constant on the affine fibers of ρ. Thus the universal mapping property gives a unique $q^\rho \in \mathrm{Aff}(A^\rho,\mathbb{R})$ with $q = q^\rho \circ \rho$. This q^ρ is nonconstant, since otherwise q would be constant.

If $H = Z(q)$, then $\rho(H) = Z(q^\rho)$: if $x^\rho \in \rho(H)$, then $x^\rho = \rho(x)$ for some $x \in H$, so $q^\rho(x^\rho) = q^\rho(\rho(x)) = q(x) = 0$, so $x^\rho \in Z(q^\rho)$. Conversely, if $x^\rho \in Z(q^\rho)$, choose $x \in A$ with $\rho(x) = x^\rho$. Then $q(x) = q^\rho(\rho(x)) = q^\rho(x^\rho) = 0$, so $x \in H$ and $x^\rho = \rho(x) \in \rho(H)$. Hence every element of $\mathcal{P}^\rho$ is an affine hyperplane in A^ρ.

Next we prove local finiteness. Let $K^\rho \subseteq A^\rho$ be compact, and choose an affine section $\sigma : A^\rho \to A$ of ρ. Then $\sigma(K^\rho)$ is compact in A. Suppose $\rho(H) \in \mathcal{P}^\rho$ meets K^ρ. Choose $q \in \mathfrak{Q}(\mathcal{P})$ with $H = Z(q)$, and choose $x^\rho \in \rho(H) \cap K^\rho$. Since $\rho(H) = Z(q^\rho)$, we have $q^\rho(x^\rho) = 0$. Put $x := \sigma(x^\rho)$. Since $\rho \circ \sigma = \mathrm{id}_{A^\rho}$, we have $q(x) = q^\rho(\rho(\sigma(x^\rho))) = q^\rho(x^\rho) = 0$. Thus $x \in H \cap \sigma(K^\rho)$, so H meets $\sigma(K^\rho)$. Local finiteness of $\mathcal{P}$ leaves only finitely many possibilities for H, hence only finitely many possibilities for $\rho(H)$. Therefore $\mathcal{P}^\rho$ is locally finite.

We now record the normal-space consequences. Since $dq = dq^\rho \circ d\rho$ with $d\rho$ surjective, if the normal lines $\mathbb{R}dq$ of $\mathcal{P}$ form a finite set, then the normal lines $\mathbb{R}dq^\rho$ of $\mathcal{P}^\rho$ also form a finite set. To prove the spanning assertion, suppose $v^\rho \in V_{A^\rho}$ is killed by every dq^ρ. Any lift $v \in V_A$ is then killed by every dq, so $v \in W_{\mathcal{P}}$ and $v^\rho = 0$. Thus the common kernel of the normal spaces of $\mathcal{P}^\rho$ is zero, so these normal spaces span $V_{A^\rho}^*$.

We next show that a finite-λ chamber forces the objective to factor through the quotient. Every chamber inequality is constant along $W_{\mathcal{P}}$. Hence $\overline{\mathfrak{c}} + W_{\mathcal{P}} = \overline{\mathfrak{c}}$ for every chamber $\mathfrak{c}$. If $\lambda(\mathfrak{c},\omega) < \infty$ for some chamber and $d\omega$ did not vanish on $W_{\mathcal{P}}$, translating inside $\overline{\mathfrak{c}}$ along a vector of $W_{\mathcal{P}}$ would make ω unbounded above. Thus $d\omega|_{W_{\mathcal{P}}} = 0$, so there is a unique affine functional $\omega^\rho \in \mathrm{Aff}(A^\rho,\mathbb{R})$ with $\omega = \omega^\rho \circ \rho$. Since $d\omega \neq 0$, the quotient objective ω^ρ is nonconstant.

We then compare chambers and their closures. The signs of all defining forms factor through ρ. Conversely, if a nonempty sign fibre for the image forms q^ρ occurs in A^ρ, its inverse image is a nonempty sign fibre for the original forms in A; duplicate image hyperplanes impose the same sign condition. Thus a sign pattern occurs in A exactly when the corresponding sign pattern for the forms q^ρ occurs in A^ρ. By the sign-fibre description of chambers, $\mathfrak{c} \mapsto \rho(\mathfrak{c})$ is a bijection from chambers of $\mathcal{P}$ to chambers of $\mathcal{P}^\rho$, with inverse $\mathfrak{d}^\rho \mapsto \rho^{-1}(\mathfrak{d}^\rho)$. The weak-sign closure formula in Lemma 4 gives the displayed closure identities.

Finally, we compare objective values and optimal faces. Since $\omega = \omega^\rho \circ \rho$, the closure identity $\overline{\mathfrak{c}} = \rho^{-1}(\overline{\mathfrak{c}^\rho})$ preserves feasible objective values in both directions, and the two linear programs have the same supremum. If the common supremum is finite, the same identity identifies the two maximizer sets. Definition 25 identifies these maximizer sets with the closures of the optimal faces, giving the final display. □

Definition 29. *Let $C \subseteq A$ be a polyhedron. Distinct vertices v, v' of C are neighboring vertices if the segment $[v, v']$ is a one-dimensional face of C.*

Lemma 22. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement with finitely many normal spaces, and assume that these normal spaces span V_A^*. Fix any norm on V_A. Then there are unit vectors $u_1, \dots, u_J$, depending only on the arrangement and the norm, such that whenever v and v' are neighboring vertices of the closure of any chamber, then $v' - v = tu_j$ for some $t \in \mathbb{R}_{\neq 0}$ and some $j \in \{1, \dots, J\}$.*

Proof. Let $\mathcal{N} \subseteq V_A^* \setminus \{0\}$ be a finite set representing the normal spaces of $\mathcal{P}$. Let $C = \bar{\mathfrak{c}}$ be a chamber closure, and let $E = [v, v']$ be a neighboring-vertex edge of C. By Lemma 6, $E = \overline{F}$ for an arrangement face $F \subseteq C$. Lemma 3 gives $\operatorname{aff}(F) = \bigcap_{H \in \mathcal{H}(F,\mathcal{P})} H$, so $\operatorname{lin}(E) = \operatorname{lin}(F)$ is the common kernel of the normal covectors of the hyperplanes in $\mathcal{H}(F, \mathcal{P})$. Those covectors lie in the finite set of normal lines represented by $\mathcal{N}$, so $\operatorname{lin}(E)$ is a one-dimensional member of $\{\bigcap_{n \in I} \ker n : I \subseteq \mathcal{N}\}$. Only finitely many such lines occur. Choose one unit vector on each of them. □

Lemma 23. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement with finitely many normal spaces, and let $\omega \in \operatorname{Aff}(A, \mathbb{R}) \setminus \operatorname{Con}(A, \mathbb{R})$. Suppose some chamber has finite $\lambda(-, \omega)$. Let $\rho : A \to A^\rho$, $\mathcal{P}^\rho$, and ω^ρ be the quotient data from Lemma 21, and fix any norm on V_{A^ρ}. Then, after identifying chambers of $\mathcal{P}$ with chambers of $\mathcal{P}^\rho$ through ρ, the following hold:*

1. *the closure of every chamber of $\mathcal{P}^\rho$ is a polyhedron;*
2. *every chamber $\mathfrak{c}^\rho$ with $\lambda(\mathfrak{c}^\rho, \omega^\rho) < \infty$ has an optimal vertex;*
3. *there are unit vectors $u_1, \dots, u_J \in V_{A^\rho}$, depending only on $\mathcal{P}^\rho$ and the chosen norm, such that whenever v and v' are neighboring vertices of the closure of any chamber of $\mathcal{P}^\rho$, then $v' - v = tu_j$ for some $t \in \mathbb{R}_{\neq 0}$ and some $j \in \{1, \dots, J\}$.*

Consequently every vertex path in the quotient obtained by repeatedly moving to a neighboring vertex with larger ω^ρ-value uses steps parallel to this same finite list of directions.

Proof. Lemma 21 says that $\mathcal{P}^\rho$ is locally finite, has finitely many normal spaces, and has normal spaces spanning $V_{A^\rho}^*$. Hence Lemma 19 gives (1). These chamber closures are pointed: any line in one of them would have direction killed by every normal covector of $\mathcal{P}^\rho$, hence by the spanning assertion would have zero direction.

For (2), let $\mathfrak{c}^\rho$ have finite $\lambda(\mathfrak{c}^\rho, \omega^\rho)$. Then $\overline{\mathfrak{c}^\rho}$ is a nonempty pointed polyhedron on which ω^ρ is bounded above, so linear programming gives an optimal vertex [3, Corollary 2.11]. Finally, (3) is Lemma 22 applied to $\mathcal{P}^\rho$. □

The proof of Theorem 3 essentially runs the simplex algorithm in each chamber. The theorem is independent of the frame-based construction below: it applies to any acyclic strategy whose edges cross source optimal faces. We will use the following standard result [3, Chapter 4, pp. 239–240].

Lemma 24. *Let $C \subseteq A$ be a nonempty polyhedron, and let $\omega \in \operatorname{Aff}(A, \mathbb{R})$ be bounded above on C. If v is a vertex of C which is not optimal for maximizing ω, then there is a neighboring vertex v' of C such that $\omega(v') > \omega(v)$.*

Theorem 3 (Acyclic optimal-face graphs are uniformly escaping with finitely many normals)**.** *Suppose $\mathcal{P}$ has finitely many normal spaces, and let $\mathcal{F}$ be a flag selection strategy. If $G_\mathcal{F}$ is acyclic, then $\mathcal{F}$ is uniform escaping, and hence satisfies all four progress and escape conditions.*

Proof. It is enough to prove uniform escaping. Then Lemma 18 gives pathwise escaping, and uniform escaping gives uniform progressive by taking η to be the starting value; applying Lemma 18 again gives pathwise progressive.

Fix a finite-λ chamber $\mathfrak{c}_0$ and a threshold $\eta \in \mathbb{R}$. Let $\rho : A \to A^\rho$, $\mathcal{P}^\rho$, and ω^ρ be the quotient data from Lemma 21. Since $\lambda(\mathfrak{c}_0, \omega) < \infty$, the objective factors as $\omega = \omega^\rho \circ \rho$, chambers are identified with quotient chambers, and $\lambda(\mathfrak{c}, \omega) = \lambda(\mathfrak{c}^\rho, \omega^\rho)$, where $\mathfrak{c}^\rho := \rho(\mathfrak{c})$. We keep paths in the original graph $G_{\mathcal{F}}$, while we use the quotient to compute $\lambda(\mathfrak{c}, \omega)$.

Fix a norm on V_{A^ρ}. By Lemma 23, every quotient chamber closure is a polyhedron, every finite-λ quotient chamber has an optimal vertex, and there are unit vectors $u_1, \dots, u_J \in V_{A^\rho}$ such that every neighboring-vertex edge in any quotient chamber closure is parallel to one of the u_j. Let $S := \{|d\omega^\rho(u_j)| : j \in [1, J]\} \cap \mathbb{R}_{>0}$.

If $\lambda(\mathfrak{c}_0, \omega) > \eta$, Lemma 16 gives $\lambda(\mathfrak{c}_L, \omega) > \eta$ for every path from $\mathfrak{c}_0$, so take $L_0 = 1$. Thus assume $\lambda(\mathfrak{c}_0, \omega) \le \eta$.

Choose an optimal vertex x_0^ρ of $\overline{\mathfrak{c}_0^\rho}$. If $S \neq \emptyset$, put $s := \min S$. Define

$$B^\rho := \begin{cases} \{x_0^\rho\}, & S = \emptyset, \\ \{y^\rho \in A^\rho : |y^\rho - x_0^\rho| \le (\eta - \lambda(\mathfrak{c}_0, \omega))/s\}, & S \neq \emptyset, \end{cases}$$

The set B^ρ is compact. By Lemma 10, applied to $\mathcal{P}^\rho$, only finitely many quotient chamber closures meet B^ρ. Let N be their number and set $L_0 := N + 1$.

Let $(\mathfrak{c}_0, \dots, \mathfrak{c}_L)$ be a path in $G_{\mathcal{F}}$ with $L \ge L_0$. Suppose, for a contradiction, that $\lambda(\mathfrak{c}_L, \omega) \le \eta$. By Lemma 16, every chamber on the path has finite λ-value and $\lambda(\mathfrak{c}_i, \omega) \le \eta$.

Suppose x_i^ρ has been chosen as an optimal vertex of $\overline{\mathfrak{c}_i^\rho}$. Since the edge $\mathfrak{c}_i \to \mathfrak{c}_{i+1}$ crosses the source optimal face, $\overline{F(\mathfrak{c}_i, \omega)} \subseteq \overline{\mathfrak{c}_{i+1}}$. Applying Lemma 21 to optimal faces gives $\overline{F(\mathfrak{c}_i^\rho, \omega^\rho)} \subseteq \overline{\mathfrak{c}_{i+1}^\rho}$. Thus $x_i^\rho \in \overline{\mathfrak{c}_{i+1}^\rho}$. Lemma 20, applied in the quotient arrangement, shows that x_i^ρ is a vertex of $\overline{\mathfrak{c}_{i+1}^\rho}$.

If x_i^ρ is already optimal on $\overline{\mathfrak{c}_{i+1}^\rho}$, set $x_{i+1}^\rho := x_i^\rho$, $L_i := 0$, and $v_{i0}^\rho := x_i^\rho$. Otherwise, iterating Lemma 24 inside the quotient polyhedron $\overline{\mathfrak{c}_{i+1}^\rho}$ gives a finite vertex path, because this polyhedron has finitely many vertices and the objective strictly increases at each step: $x_i^\rho = v_{i0}^\rho, \dots, v_{iL_i}^\rho = x_{i+1}^\rho$. Here $v_{i(l-1)}^\rho$ and v_{il}^ρ are neighboring vertices of $\overline{\mathfrak{c}_{i+1}^\rho}$, $\omega^\rho(v_{il}^\rho) > \omega^\rho(v_{i(l-1)}^\rho)$ for all $l \in [1, L_i]$, and x_{i+1}^ρ is optimal on $\overline{\mathfrak{c}_{i+1}^\rho}$.

If $S = \emptyset$, no neighboring-vertex direction in the quotient has positive ω^ρ-slope, so the nontrivial alternative above never occurs. Hence $x_j^\rho = x_0^\rho \in B^\rho$ for every j.

If $S \neq \emptyset$, each nonzero step satisfies $v_{il}^\rho - v_{i(l-1)}^\rho = t u_a$ for some $t \neq 0$ and some $a \in \{1, \dots, J\}$, and the positivity of the objective increase implies $\operatorname{sign}(t) d\omega^\rho(u_a) \in S$. Therefore

$$\omega^\rho(x_{i+1}^\rho) - \omega^\rho(x_i^\rho) = \sum_{l=1}^{L_i} \big(\omega^\rho(v_{il}^\rho) - \omega^\rho(v_{i(l-1)}^\rho)\big) \ge s \sum_{l=1}^{L_i} |v_{il}^\rho - v_{i(l-1)}^\rho| \ge s |x_{i+1}^\rho - x_i^\rho|.$$

For every $j \in [0, L]$, summing these inequalities and using Lemma 16 gives

$$s \sum_{i=0}^{j-1} |x_{i+1}^\rho - x_i^\rho| \le \omega^\rho(x_j^\rho) - \omega^\rho(x_0^\rho) = \lambda(\mathfrak{c}_j, \omega) - \lambda(\mathfrak{c}_0, \omega) \le \eta - \lambda(\mathfrak{c}_0, \omega).$$

Thus $|x_j^\rho - x_0^\rho| \le (\eta - \lambda(\mathfrak{c}_0, \omega))/s$, so $x_j^\rho \in B^\rho$ for every j.

In either case, $B^\rho \cap \overline{\mathfrak{c}_j^\rho} \neq \emptyset$ for every $j \in [0, L]$. Hence the quotient chambers $\mathfrak{c}_j^\rho$, and therefore the original chambers $\mathfrak{c}_j$, lie in a fixed finite set of size N. Since $L \ge N + 1$, some chamber repeats along the path, producing a directed cycle in $G_{\mathcal{F}}$, contradicting acyclicity. Therefore $\lambda(\mathfrak{c}_L, \omega) > \eta$. This proves uniform escaping. $\square$

Summarizing, the pathwise and uniform versions of progress are equivalent, as are the pathwise and uniform versions of escape. Escape implies progress. Acyclicity implies progress and, when the arrangement has finitely many normal spaces, also implies escape.

6 An acyclic flag-selection rule

We construct a greedy rule that selects an optimal-face flag in each chamber with finite objective value. Its acyclicity is proved in Section 7.

6.1 Example

The difficulty is that several chambers in $\mathrm{Star}_\Omega(F(\mathfrak{c},\omega))$ may have the same optimal value, and the same optimal face, as $\mathfrak{c}$. Consequently, an arbitrary graph might contain cycles and remain stuck at the same optimal value. For instance, if $\mathfrak{c} \to \mathfrak{d}$ is an edge with sign vector α, put $FL := \mathcal{F}(\mathfrak{c})$. The target chamber is $\mathfrak{d} = \mathfrak{b}_{FL}(\alpha)$, and $FL(\alpha)$ is a flag with top $\mathfrak{d}$. If $F(\mathfrak{c},\omega) = F(\mathfrak{d},\omega)$, a strategy could therefore choose $\mathcal{F}(\mathfrak{d}) = FL(\alpha)$. Applying the same sign vector a second time restores every transverse sign, since $\alpha_i^2 = 1$, so Theorem 2 gives $\mathcal{F}(\mathfrak{d})(\alpha) = FL$. Thus the same α produces the reverse edge $\mathfrak{d} \to \mathfrak{c}$, giving a 2-cycle.

Return to the two-dimensional arrangement introduced in Figure 3. The path diagram in Figure 4 adds the objective and a traversal whose terminal chambers all have greater objective value. The dotted line represents $\omega(x_1, x_2) = x_1 + x_2$, whose value can be read by orthogonally projecting onto the dotted line. Chambers with optimal value $\lambda(\mathfrak{c},\omega) = 0$ are shaded light blue, while chambers with optimal value $\lambda(\mathfrak{c},\omega) > 0$ are shaded light orange. The darker blue shading marks the starting chamber $\mathfrak{c}_0$.

As seen in Figure 3, in two dimensions there are exactly two possible flags at a vertex. In this example, the first step chooses $(\{0\}, \{(x_1, 0) : x_1 < 0\}, \mathfrak{c}_0)$ as the flag. This determines $2^2 - 1 = 3$ target chambers in $\mathrm{Star}_\Omega(\{(0,0)\})$. Two of these target chambers already satisfy $\lambda(\mathfrak{c},\omega) > 0$. One target chamber still has optimal value 0 and the same optimal face at the origin. Writing that target chamber as $\mathfrak{c}_1$, the second step chooses $(\{0\}, \{(x_1, x_2) : x_1 + \sqrt{3}x_2 = 0, x_1 < 0\}, \mathfrak{c}_1)$ as the flag. Again, one chamber with optimal value 0 remains, but now its optimal face $\{(x_1, x_2) : x_1 + x_2 = 0,\ x_1 < 0\}$ has codimension one. Crossing that face yields a single new chamber, since $2^1 - 1 = 1$, and this chamber has $\lambda(\mathfrak{c},\omega) > 0$. At that point, this part of the graph has left the initial level set.

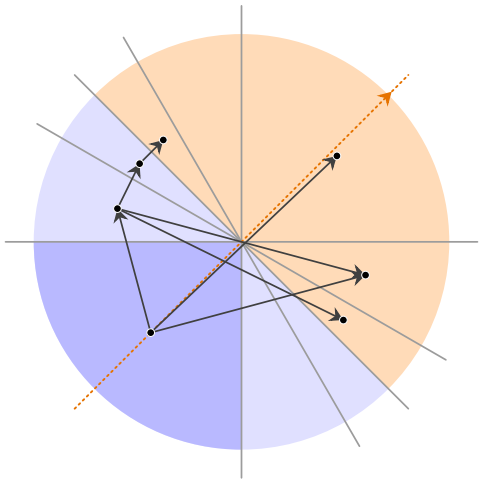


Figure 4: A two-dimensional arrangement showing the paths from the starting chamber to chambers with larger optimal value for an example strategy.

In dimension two, this picture suggests a general flag selection strategy. If the optimal face has codimension one, there is only one flag to choose from. If the optimal face is a vertex, orient a small circle around it. The chambers for which that vertex remains optimal meet the circle inside the open semicircle cut out by $-d\omega > 0$, so the orientation induces a linear order on these same-optimal-face chambers. At each step there are two flags to choose from; choose the flag so that the adjacent target that can remain in this class comes later in the induced order. Same-level edges then strictly increase this order and cannot form a cycle. This argument does not generalize directly to higher dimensions because an orientation of a sphere does not naturally order the chambers.

However, the following perspective recovers the traversal in the figure and generalizes to higher dimensions. Choose a generic secondary objective ω'; in this example, $\omega'(x_1, x_2) = x_2$ suffices. Since $-\omega$ is positive on $\overline{\mathfrak{c}_0} \setminus \{0\}$, consider the segment $\{x \in \overline{\mathfrak{c}_0} : -\omega(x) = 1\}$. For generic ω', the maximum of ω' on this segment is attained at a unique endpoint, corresponding to one of the codimension-one faces in $\overline{\mathfrak{c}_0}$. In Figure 5, this point is labeled p. Choose the flag with $F_1 = \mathbb{R}_{>0}p$, and apply the same strategy at each step. Let v be the arrow from p to the other vertex of the polytope, so $p + v$ lies on the other boundary face of $\mathfrak{c}_0$. Then $\omega'(v) < 0$, and $-v$ points from p into the other chamber sharing the face $\mathbb{R}_{>0}p$. Stretching the slice $-\omega = 1$ in the direction $-v$ reaches the next face at $p - tv$ for some $t > 0$, and $\omega'(p - tv) > \omega'(p)$. Thus the secondary value strictly increases at each same-level step, ruling out cycles. The rest of this section generalizes this strategy to higher dimensions and makes the proof precise.

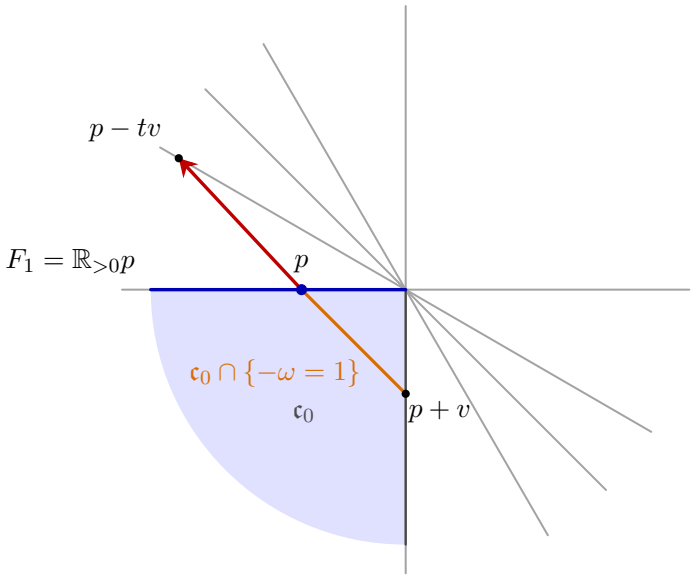


Figure 5: The shaded lower-left chamber is $\mathfrak{c}_0$. The orange segment is $\mathfrak{c}_0 \cap \{-\omega = 1\}$, with endpoints p and $p + v$. The chosen boundary ray is labeled $F_1 = \mathbb{R}_{>0}p$, and the red arrow continues from p to the point $p - tv$ in the neighboring chamber.

6.2 Transverse geometry at a face

At a face G, directions parallel to G do not affect the next face of a flag. The local choice therefore lives in the normal quotient

$$N_G := V_A / \operatorname{lin}(G).$$

The hyperplanes containing G induce a finite essential central arrangement in N_G. The starting chamber corresponds to a pointed cone in the quotient, and the covering faces correspond to the extreme rays. A positive covector cuts this cone in a compact polytope, generalizing the segment in Figure 5.

Definition 30. *Let G be an arrangement face, and let*

$$\pi_G : V_A \longrightarrow N_G := V_A / \operatorname{lin}(G)$$

be the quotient map. If $K \in \mathcal{H}(G, \mathcal{P})$, choose $x \in \operatorname{aff}(G)$ and define

$$K_G^{\perp} := (K - x) / \operatorname{lin}(G) \subseteq N_G.$$

This hyperplane is independent of x. The transverse arrangement at G is

$$\mathcal{P}_G^{\perp} := \{K_G^{\perp} : K \in \mathcal{H}(G, \mathcal{P})\}.$$

For $q \in \mathfrak{Q}(G, \mathcal{P})$, the differential dq descends to a covector $dq_G^{\perp} \in N_G^*$, canonically identifying N_G^* with $\operatorname{Ann}(G)$.

If H is a face with $G \subseteq \overline{H}$, define its transverse face at G by

$$R_G(H) := \left\{ \xi \in N_G : \begin{array}{ll} dq_G^{\perp}(\xi) = 0 & q \in \mathfrak{Q}(H, \mathcal{P}), \\ \operatorname{sign}_H(q) dq_G^{\perp}(\xi) > 0 & q \in \mathfrak{Q}(G, \mathcal{P}) \setminus \mathfrak{Q}(H, \mathcal{P}) \end{array} \right\}.$$

The conditions are unchanged if q is rescaled. Define the ambient lift sets

$$D_G(H) := \pi_G^{-1}(R_G(H)), \qquad D_G(\overline{H}) := \pi_G^{-1}(\overline{R_G(H)}).$$

For $\theta \in N_G^*$ and $v \in V_A$, write $\theta(v) := \theta(\pi_G(v))$.

Lemma 25. *The arrangement $\mathcal{P}_G^{\perp}$ is finite, central, and essential. The map*

$$H \longmapsto R_G(H)$$

is a closure-order-preserving bijection from the faces H of $\mathcal{P}$ satisfying $G \subseteq \overline{H}$ to the faces of $\mathcal{P}_G^{\perp}$. Moreover,

$$\dim R_G(H) = \dim H - \dim G, \qquad \lin(\overline{R_G(H)}) = \lin(H)/\lin(G).$$

Proof. Choose $x \in G$. Every hyperplane of $\mathcal{H}(G, \mathcal{P})$ contains x, so local finiteness makes this set finite. The induced hyperplanes are central. Their common intersection is zero: if $\xi = \pi_G(v)$ lies in every $K_G^{\perp}$, then $dq(v) = 0$ for every $q \in \mathfrak{Q}(G, \mathcal{P})$. Since $\operatorname{aff}(G) = \bigcap_{K \in \mathcal{H}(G,\mathcal{P})} K$ by Lemma 3, the common kernel of these differentials is $\lin(G)$. Thus $v \in \lin(G)$ and $\xi = 0$, so the transverse arrangement is essential.

The displayed equalities and inequalities in Definition 30 are the sign description of a face of $\mathcal{P}_G^{\perp}$, so $R_G(H)$ is such a face. If $G \subseteq \overline{H}$, every hyperplane not containing G has the same nonzero sign on H as on G by Lemma 5. Consequently the transverse signs determine the full sign pattern of H, proving injectivity.

Conversely, let R be a face of $\mathcal{P}_G^{\perp}$, choose $\xi \in R$, and lift it to $v \in V_A$. For all sufficiently small $\epsilon > 0$, the forms whose hyperplanes do not contain G have at $x + \epsilon v$ the same signs as on G, by Lemma 8; the remaining signs are those of R. The face H containing $x + \epsilon v$ therefore satisfies $G \subseteq \overline{H}$ and $R_G(H) = R$. This proves surjectivity. The same sign description and Lemma 5 show that closure order is preserved.

A transverse hyperplane $K_G^{\perp}$ contains $R_G(H)$ if and only if K contains H. Therefore

$$\lin(\overline{R_G(H)}) = \bigcap_{K \in \mathcal{H}(H,\mathcal{P})} K_G^{\perp} = \lin(H)/\lin(G),$$

where the second equality follows from Lemma 3. Taking dimensions gives

$$\dim R_G(H) = \dim H - \dim G.$$

□

Definition 31. *Let $G \subseteq \overline{\mathfrak{e}}$, where G is a face and $\mathfrak{e}$ is a chamber. Define*

$$C_G(\mathfrak{e}) := \overline{R_G(\mathfrak{e})} \subseteq N_G$$

and

$$\operatorname{Cov}_G(\mathfrak{e}) := \{H : G \subseteq \overline{H} \subseteq \overline{\mathfrak{e}},\ \operatorname{codim}(H) = \operatorname{codim}(G) - 1\}.$$

By Lemma 9, these are exactly the faces covering G in $\overline{\mathfrak{e}}$.

Lemma 26. *Let $G \subseteq \overline{\mathfrak{e}}$. Then $C_G(\mathfrak{e})$ is a full-dimensional pointed polyhedral cone in N_G. Its faces have relative interiors*

$$R_G(H), \qquad G \subseteq \overline{H} \subseteq \overline{\mathfrak{e}}.$$

In particular, the extreme rays of $C_G(\mathfrak{e})$ are

$$\rho_G(H) := \overline{R_G(H)}, \qquad H \in \mathrm{Cov}_G(\mathfrak{e}).$$

These extreme rays generate $C_G(\mathfrak{e})$. Distinct extreme rays contain linearly independent nonzero vectors.

Proof. The cone $C_G(\mathfrak{e})$ is the closure of a chamber of the finite central arrangement $\mathcal{P}_G^\perp$, so it is full-dimensional and polyhedral. It is pointed because the arrangement is essential: if $\xi, -\xi \in C_G(\mathfrak{e})$, every defining form vanishes on ξ, hence $\xi = 0$. The transverse face correspondence identifies the relative interiors of the cone faces with the displayed $R_G(H)$. Its dimension formula identifies the one-dimensional faces with the covering faces. Rockafellar's extreme-ray representation for closed pointed convex cones shows that the cone is generated by these rays [4, Corollary 18.5.2]. Two distinct extreme rays of a pointed cone cannot lie in one line, so their nonzero vectors are linearly independent. □

After choosing an extreme ray $\rho_G(H)$, the recursion passes from N_G to N_H. The next three lemmas record the compatibility needed for this passage: first for the transverse arrangements and then for ambient representatives of consecutive directions.

Lemma 27. *If $G \subseteq \overline{H}$, put*

$$L_G(H) := \mathrm{lin}(H)/\mathrm{lin}(G) = \mathrm{lin}\big(\overline{R_G(H)}\big) \subseteq N_G.$$

There is a canonical identification

$$N_H \cong N_G/L_G(H),$$

under which $\mathcal{P}_H^\perp$ is the arrangement induced by those hyperplanes of $\mathcal{P}_G^\perp$ which contain $L_G(H)$. Dually,

$$N_H^* \cong L_G(H)^\perp \subseteq N_G^*.$$

Moreover, if $H \subseteq \overline{M}$, the quotient map $N_G \to N_H$ sends $R_G(M)$ onto $R_H(M)$.

Proof. The vector-space isomorphism

$$\big(V_A/\mathrm{lin}(G)\big)/\big(\mathrm{lin}(H)/\mathrm{lin}(G)\big) \cong V_A/\mathrm{lin}(H)$$

is canonical. A hyperplane containing G induces a hyperplane in the second quotient exactly when it also contains H, which proves the arrangement statement. The dual identification follows.

For the final assertion, the signs defining $R_G(M)$ on the hyperplanes containing H are exactly the signs defining $R_H(M)$, so the quotient maps $R_G(M)$ into $R_H(M)$. Conversely, take $\zeta \in R_H(M)$, lift it to $\widetilde{\zeta} \in N_G$, and choose $\xi \in R_G(H)$. For all sufficiently small $t > 0$, the signs of $\xi + t\widetilde{\zeta}$ are those of M: the signs on hyperplanes not containing H are fixed by ξ, while the remaining signs are determined by $\widetilde{\zeta}$. Thus $\xi + t\widetilde{\zeta} \in R_G(M)$. Its image is $t\zeta$. Since $R_G(M)$ is a central-arrangement face and hence is closed under multiplication by positive scalars, $t^{-1}\xi + \widetilde{\zeta} \in R_G(M)$ maps to ζ. This proves surjectivity. □

Lemma 28. *Let G, H, M be faces with $G \subseteq \overline{H} \subseteq \overline{M}$. If $u \in D_G(H)$ and $w \in D_H(\overline{M})$, then*

$$u + tw \in D_G(\overline{M})$$

for all sufficiently small $t > 0$.

Proof. For $q \in \mathfrak{Q}(G,\mathcal{P}) \setminus \mathfrak{Q}(H,\mathcal{P})$, one has $\operatorname{sign}_H(q) = \operatorname{sign}_M(q) \neq 0$ and $\operatorname{sign}_M(q)dq(u) > 0$. There are only finitely many relevant hyperplanes, so these strict inequalities persist for $u + tw$ when $t > 0$ is sufficiently small.

If $q \in \mathfrak{Q}(H,\mathcal{P})$, then $dq(u) = 0$. Since $w \in D_H(\overline{M})$, one has $dq(w) = 0$ for $q \in \mathfrak{Q}(M,\mathcal{P})$ and $\operatorname{sign}_M(q)dq(w) \geq 0$ otherwise. Thus $u + tw$ satisfies precisely the weak transverse sign conditions defining $\overline{R_G(M)}$, and the claim follows from the definition of the ambient lift set. □

Lemma 29. *Suppose*

$$G \subseteq \overline{H} \subseteq \overline{M} \subseteq \bar{\mathfrak{e}}, \qquad H \in \operatorname{Cov}_G(\mathfrak{e}), \qquad M \in \operatorname{Cov}_H(\mathfrak{e}).$$

Then there is $L \in \operatorname{Cov}_G(\mathfrak{e})$, $L \neq H$, with $L \subseteq \overline{M}$ and

$$D_G(L) \cap D_H(M) \neq \emptyset.$$

Proof. The cone $\overline{R_G(M)}$ is a two-dimensional face of $C_G(\mathfrak{e})$, and $\rho_G(H)$ is one of its extreme rays. A two-dimensional pointed polyhedral cone has a second extreme ray, which is $\rho_G(L)$ for some $L \in \operatorname{Cov}_G(\mathfrak{e})$ with $L \subseteq \overline{M}$ and $L \neq H$.

By Lemma 27, the canonical quotient $N_G \to N_H = N_G/L_G(H)$ maps $R_G(M)$ onto $R_H(M)$. Its kernel ray in $\overline{R_G(M)}$ is $\rho_G(H)$. Quotienting the two-dimensional cone by the span of this boundary ray leaves a one-dimensional cone. The image of the other boundary ray is nonzero and lies on the same open half-line as the image of $R_G(M)$; that half-line is $R_H(M)$. Hence the relative interior of the second ray maps to $R_H(M)$. Choose $\xi \in R_G(L)$ and an ambient lift $v \in V_A$. Then $\pi_G(v) = \xi$ and $\pi_H(v) \in R_H(M)$, so $v \in D_G(L) \cap D_H(M)$. □

Let $G \subseteq \bar{\mathfrak{e}}$ have positive codimension, and let $\eta \in N_G^*$ be positive on $C_G(\mathfrak{e}) \setminus \{0\}$. To compare the extreme rays, normalize them by setting

$$P_\eta(\mathfrak{e}, G) := \{\xi \in C_G(\mathfrak{e}) : \eta(\xi) = 1\}$$

Lemma 30. *The set $P_\eta(\mathfrak{e}, G)$ is a compact polytope.*

Proof. The cone $C_G(\mathfrak{e})$ is pointed and polyhedral by Lemma 26, so the section is a closed polyhedron. Suppose it were unbounded. Fix a norm on N_G, and choose $\xi_n \in P_\eta(\mathfrak{e}, G)$ with $\|\xi_n\| \to \infty$. Put

$$u_n := \frac{\xi_n}{\|\xi_n\|}.$$

After passing to a subsequence, compactness of the unit sphere gives $u_n \to u$ with $\|u\| = 1$. Since $C_G(\mathfrak{e})$ is a closed cone, $u \in C_G(\mathfrak{e})$. On the other hand,

$$\eta(u_n) = \frac{1}{\|\xi_n\|} \longrightarrow 0,$$

so $\eta(u) = 0$. This contradicts the strict positivity of η on $C_G(\mathfrak{e}) \setminus \{0\}$. Hence the section is bounded and therefore a compact polytope. □

Each extreme ray $\rho_G(H)$, $H \in \operatorname{Cov}_G(\mathfrak{e})$, meets this section in the single point

$$p_H := \frac{\xi_H}{\eta(\xi_H)}, \qquad \xi_H \in R_G(H),$$

and these points are precisely its vertices.

6.3 The ray-regular greedy step

At a current face G, the first covector normalizes the transverse chamber cone and the second covector selects a vertex of the normalized polytope. After that vertex is chosen, the covectors are eliminated along its ray and descend to the next normal quotient. The following convenient sufficient condition makes the maximizer unique and is preserved by this elimination.

Definition 32 (Ray-regular frame)**.** *Let $\mathfrak{e}$ be a chamber and let $G \subseteq \bar{\mathfrak{e}}$ be a face of codimension d. If $d = 0$, the empty frame is ray-regular at $(\mathfrak{e}, G)$. Suppose $d > 0$. An ordered basis $\Theta = (\theta_1, \dots, \theta_d)$ of N_G^* is ray-regular at $(\mathfrak{e}, G)$ if $\theta_1(\xi) > 0$ for every $\xi \in C_G(\mathfrak{e}) \setminus \{0\}$, and, for every $1 \le k \le d$ and every collection*

$$H_1, \dots, H_k \in \mathrm{Cov}_G(\mathfrak{e})$$

with linearly independent choices $\xi_b \in R_G(H_b)$, one has

$$\det(\theta_a(\xi_b))_{a,b=1}^{k} \neq 0.$$

The determinant condition is independent of the ray representatives up to multiplication by a positive scalar.

Positivity of the first covector gives the compact polytope of Lemma 30; the two-row determinants make its maximizing vertex unique; and the higher determinants ensure that the frame left after eliminating that ray remains ray-regular.

Lemma 31 (Ray-regular greedy step)**.** *Let $\mathfrak{e}$ be a chamber, let $G \subseteq \bar{\mathfrak{e}}$ be a face, and put $d := \mathrm{codim}(G) > 1$. Suppose*

$$\Theta = (\eta, \varphi, \psi_1, \dots, \psi_{d-2})$$

is ray-regular at $(\mathfrak{e}, G)$. For each $H' \in \mathrm{Cov}_G(\mathfrak{e})$, choose $\xi_{H'} \in R_G(H')$ and set

$$S_{\eta,\varphi}(H') := \frac{\varphi(\xi_{H'})}{\eta(\xi_{H'})}.$$

These values are independent of the representatives, pairwise distinct, and have a unique maximizer

$$H = \mathrm{argmax}_{H' \in \mathrm{Cov}_G(\mathfrak{e})} S_{\eta,\varphi}(H').$$

Put $s := S_{\eta,\varphi}(H)$. Then

$$\frac{\varphi(\xi)}{\eta(\xi)} \le s \qquad (\xi \in C_G(\mathfrak{e}) \setminus \{0\}),$$

with equality only if $\xi \in \rho_G(H) \setminus \{0\}$.

Choose any $\xi_* \in R_G(H)$, and define

$$\eta^+ := s\eta - \varphi, \qquad \psi_a^+ := \psi_a - \frac{\psi_a(\xi_*)}{\eta(\xi_*)}\eta \quad (1 \le a \le d-2).$$

The covectors ψ_a^+ are independent of the choice of ξ_*. All the updated covectors vanish on $L_G(H)$, so they descend through the canonical identification $N_H^* \cong L_G(H)^\perp \subseteq N_G^*$. As covectors on N_H, η^+ is positive on $C_H(\mathfrak{e}) \setminus \{0\}$, and

$$(\eta^+, \psi_1^+, \dots, \psi_{d-2}^+)$$

is ray-regular at $(\mathfrak{e}, H)$.

Proof. Ray-regularity gives $\eta > 0$ on $C_G(\mathfrak{e}) \setminus \{0\}$, so all denominators are positive. The value attached to a covering face is independent of the representative because $R_G(H')$ is an open ray. If $H_1 \neq H_2$ and $\xi_b \in R_G(H_b)$, then ξ_1, ξ_2 are linearly independent by Lemma 26. The two-row ray-regularity condition therefore gives

$$\det \begin{pmatrix} \eta(\xi_1) & \eta(\xi_2) \\ \varphi(\xi_1) & \varphi(\xi_2) \end{pmatrix} \neq 0,$$

so the two ratios are distinct. By Lemma 30, $P_\eta(\mathfrak{e}, G)$ is compact; its vertices are the normalized covering rays $p_{H'}$. Their φ-values are the displayed ratios, so φ has a unique maximizing vertex p_H. Its maximizing face is therefore $\{p_H\}$. Scaling any nonzero $\xi \in C_G(\mathfrak{e})$ into $P_\eta(\mathfrak{e}, G)$ gives the asserted inequality, with equality exactly on $\rho_G(H) \setminus \{0\}$.

Every updated covector vanishes on ξ_*, hence on $L_G(H) = \operatorname{span}(\xi_*)$, and therefore defines a covector on N_H. Let $\zeta \in C_H(\mathfrak{e}) \setminus \{0\}$, choose $u_* \in D_G(H)$ with $\pi_G(u_*) = \xi_*$, and choose $w \in D_H(\bar{\mathfrak{e}})$ with $\pi_H(w) = \zeta$. By Lemma 28, $u_* + tw \in D_G(\bar{\mathfrak{e}})$ for all sufficiently small $t > 0$. Since $\eta^+(u_*) = 0$, the preceding inequality gives

$$0 \leq \eta^+(u_* + tw) = t\eta^+(w).$$

If equality held, then $\pi_G(u_* + tw)$ would lie in $\rho_G(H) = L_G(H) \cap C_G(\mathfrak{e})$, so its image in $N_H = N_G/L_G(H)$ would be zero. That image is $t\zeta$, a contradiction. Thus $\eta^+(\zeta) = \eta^+(w) > 0$. Independence of ξ_* follows because all its choices are positive multiples of one another.

It remains to verify the determinant condition. Fix $M_1, \ldots, M_k \in \operatorname{Cov}_H(\mathfrak{e})$, with $1 \leq k \leq d-1$, and choose linearly independent $\zeta_b \in R_H(M_b)$. Put

$$(\chi_1, \ldots, \chi_k) := (\eta^+, \psi_1^+, \ldots, \psi_{k-1}^+).$$

By Lemma 29, for each b there is a covering face L_b of G and a vector $u_b \in D_G(L_b) \cap D_H(M_b)$. Both $\pi_H(u_b)$ and ζ_b lie in the open ray $R_H(M_b)$, so a unique positive rescaling of u_b makes $\pi_H(u_b) = \zeta_b$. This rescaling preserves membership in both ambient lift sets. Put $\widehat{\xi}_b := \pi_G(u_b) \in R_G(L_b)$.

The vectors $\xi_*, \widehat{\xi}_1, \ldots, \widehat{\xi}_k$ in N_G are linearly independent: a relation first descends to N_H, forcing the coefficients of the $\widehat{\xi}_b$ to vanish, and then forces the coefficient of ξ_* to vanish. Ray-regularity of Θ gives

$$\det \begin{pmatrix} \eta(\xi_*) & \eta(\widehat{\xi}_1) & \cdots & \eta(\widehat{\xi}_k) \\ \varphi(\xi_*) & \varphi(\widehat{\xi}_1) & \cdots & \varphi(\widehat{\xi}_k) \\ \psi_1(\xi_*) & \psi_1(\widehat{\xi}_1) & \cdots & \psi_1(\widehat{\xi}_k) \\ \vdots & \vdots & & \vdots \\ \psi_{k-1}(\xi_*) & \psi_{k-1}(\widehat{\xi}_1) & \cdots & \psi_{k-1}(\widehat{\xi}_k) \end{pmatrix} \neq 0.$$

Replacing φ by $\eta^+ = s\eta - \varphi$, and each ψ_a by $\psi_a^+ = \psi_a - \frac{\psi_a(\xi_*)}{\eta(\xi_*)}\eta$, is an invertible row operation. Since every updated covector vanishes on ξ_*, expansion along the first column shows that the resulting determinant is

$$\eta(\xi_*) \det\big(\chi_a(\zeta_b)\big)_{a,b=1}^{k}.$$

Here the updated covectors have the same values on $\widehat{\xi}_b$ and its image ζ_b. The determinant is nonzero, and $\eta(\xi_*) > 0$. Hence the updated frame is ray-regular at $(\mathfrak{e}, H)$. □

6.4 Recursive selection and generic auxiliary data

The greedy step can now be iterated from the optimal face to the chamber. We first define that recursion from fixed auxiliary data. We then show that data for which every step is valid are generically available.

At an optimizer face, the objective supplies the positive covector that initializes the recursion.

Lemma 32. *Let $\mathfrak{c} \in \Omega(\mathcal{P})$, and suppose $F = F(\mathfrak{c}, \omega)$. Then $d\omega|_{\mathrm{lin}(F)} = 0$, so*

$$h_F := -d\omega \in N_F^* \cong \mathrm{Ann}(F).$$

Moreover,

$$h_F(\xi) > 0 \qquad (\xi \in C_F(\mathfrak{c}) \setminus \{0\}).$$

Proof. The objective is constant on the optimizer face F. Since F affinely spans $\mathrm{aff}(F)$, its differential vanishes on $\mathrm{lin}(F)$, so h_F descends to N_F.

Choose $x \in F$, let $\xi \in C_F(\mathfrak{c})$, and lift ξ to $v \in V_A$. For the hyperplanes containing F, membership in $C_F(\mathfrak{c})$ gives the weak chamber signs of $x + \epsilon v$. The forms whose hyperplanes do not contain F retain their chamber signs near x by Lemma 8. Hence $x + \epsilon v \in \bar{\mathfrak{c}}$ for all sufficiently small $\epsilon > 0$. Optimality gives $d\omega(v) \le 0$.

If equality held for $\xi \neq 0$, then $x + \epsilon v$ would also maximize ω on $\bar{\mathfrak{c}}$, hence would lie in the optimizer set $\overline{F}$. Therefore $\epsilon v \in \mathrm{lin}(F)$, contradicting $\xi = \pi_F(v) \neq 0$. Thus $h_F(\xi) > 0$. □

Definition 33 (Admissible auxiliary data)**.** *Fix $(\mathcal{P}, \omega)$. For every face F of codimension r which occurs as $F(\mathfrak{c}, \omega)$ for some chamber $\mathfrak{c}$, choose covectors $\theta_{F,r-1}, \dots, \theta_{F,1} \in N_F^*$ and put*

$$\Theta_F := (h_F, \theta_{F,r-1}, \dots, \theta_{F,1}).$$

The resulting family $\Theta = (\Theta_F)_F$ is called admissible *if Θ_F is ray-regular at $(\mathfrak{c}, F)$ for every chamber $\mathfrak{c}$ with $F(\mathfrak{c}, \omega) = F$.*

Definition 34. *Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, let $\omega \in \mathrm{Aff}(A, \mathbb{R}) \setminus \mathrm{Con}(A, \mathbb{R})$, fix admissible auxiliary data Θ, and let $\mathfrak{c} \in \Omega(\mathcal{P})$. If $\lambda(\mathfrak{c}, \omega) = \infty$, set*

$$FL_\Theta(\mathfrak{c}) := (\mathfrak{c}).$$

If $\lambda(\mathfrak{c}, \omega) < \infty$, let $F = F(\mathfrak{c}, \omega)$ and $r = \mathrm{codim}(F)$. Since ω is nonconstant and is constant on its optimizer face, $r > 0$. Write $\Theta_F = (\theta_{F,r}, \dots, \theta_{F,1})$, where $\theta_{F,r} = h_F$, and set

$$G_r := F, \qquad \eta_r := \theta_{F,r}, \qquad \theta_{F,\ell}^{[r]} := \theta_{F,\ell} \quad (1 \le \ell \le r-1).$$

Suppose $G_r, \dots, G_i$ and the ray-regular frame

$$(\eta_i, \theta_{F,i-1}^{[i]}, \dots, \theta_{F,1}^{[i]})$$

at $(\mathfrak{c}, G_i)$ have been constructed, where $i \ge 2$. Apply Lemma 31 with $\varphi = \theta_{F,i-1}^{[i]}$. Let G_{i-1} be the unique face in $\mathrm{Cov}_{G_i}(\mathfrak{c})$ maximizing

$$S_i(H) := \frac{\theta_{F,i-1}^{[i]}(\xi_H)}{\eta_i(\xi_H)}, \qquad \xi_H \in R_{G_i}(H),$$

and put $s_i(\mathfrak{c}) := S_i(G_{i-1})$. Choose any $\xi_i \in R_{G_i}(G_{i-1})$, and update

$$\eta_{i-1} := s_i(\mathfrak{c})\eta_i - \theta_{F,i-1}^{[i]},$$

and, when $i > 2$,

$$\theta_{F,\ell}^{[i-1]} := \theta_{F,\ell}^{[i]} - \frac{\theta_{F,\ell}^{[i]}(\xi_i)}{\eta_i(\xi_i)}\eta_i \qquad (1 \le \ell \le i-2).$$

The same lemma makes these updates independent of ξ_i and makes the updated frame ray-regular at $(\mathfrak{c}, G_{i-1})$, so the construction continues. At codimension one, set $G_0 := \mathfrak{c}$, the unique chamber face of $\bar{\mathfrak{c}}$.

The resulting flag

$$FL_\Theta(\mathfrak{c}) := (G_r, G_{r-1}, \dots, G_0 = \mathfrak{c})$$

is the selected chamber flag.

Lemma 33 (Generic existence of admissible auxiliary data)**.** *(→ proof) Let F be a positive-codimension face for which $F = F(\mathfrak{c}, \omega)$ for at least one chamber $\mathfrak{c}$, and put $r = \operatorname{codim}(F)$. There is a proper algebraic set $E_F \subseteq (N_F^*)^{r-1}$ such that every tuple*

$$(\theta_{F,r-1}, \dots, \theta_{F,1}) \in (N_F^*)^{r-1} \setminus E_F$$

makes $\Theta_F = (h_F, \theta_{F,r-1}, \dots, \theta_{F,1})$ ray-regular at every pair $(\mathfrak{c}, F)$ with $F(\mathfrak{c}, \omega) = F$. When $r = 1$, take $E_F = \emptyset \subseteq (N_F^*)^0$. Consequently, admissible auxiliary data exist, and at each relevant face their possible values form the complement of a proper algebraic subset.

Proof of Lemma 33. There are only finitely many chambers whose closures contain F, by Lemma 12. Lemma 32 shows that h_F is positive on $C_F(\mathfrak{c}) \setminus \{0\}$ for every chamber $\mathfrak{c}$ with optimal face F.

If $r = 1$, this positivity is the only condition in Definition 32, so the claim holds with $E_F = \emptyset$. Assume $r > 1$. For every chamber $\mathfrak{c}$ with $F(\mathfrak{c}, \omega) = F$, every $1 \le k \le r$, and every collection $H_1, \dots, H_k \in \operatorname{Cov}_F(\mathfrak{c})$ admitting linearly independent choices, choose linearly independent vectors $\xi_b \in R_F(H_b)$. The determinant

$$\det\big(\chi_a(\xi_b)\big)_{a,b=1}^{k}, \qquad \chi_1 := h_F, \quad \chi_a := \theta_{F,r-a+1} \ (2 \le a \le k),$$

is a polynomial in the auxiliary covectors, and it is not identically zero. Indeed, put

$$W := \operatorname{span}(\xi_1, \dots, \xi_k) \subseteq N_F.$$

Since $h_F|_W$ is a nonzero covector, complete it to a basis $h_F|_W, \beta_2, \dots, \beta_k$ of W^*. The restriction map $N_F^* \to W^*$ is surjective, so choose extensions $\theta_{F,r-a+1} \in N_F^*$ with $\theta_{F,r-a+1}|_W = \beta_a$ for $2 \le a \le k$. With any choices of the remaining covectors, the displayed determinant is nonzero.

Let E_F be the union of the zero loci of these finitely many determinant polynomials. This finite union is the zero locus of their product. The product is a nonzero polynomial, so E_F is a proper algebraic set. Outside E_F, the determinant condition in Definition 32 holds for every relevant pair $(\mathfrak{c}, F)$. The same determinant conditions also give the basis condition in Definition 32. Indeed, for any chamber $\mathfrak{c}$ with $F(\mathfrak{c}, \omega) = F$, the cone $C_F(\mathfrak{c})$ is full-dimensional in N_F and is generated by its extreme rays. Hence we may choose r covering faces $H_1, \dots, H_r \in \operatorname{Cov}_F(\mathfrak{c})$ with linearly independent ray directions. The nonzero $r \times r$ determinant for this collection shows that $(h_F, \theta_{F,r-1}, \dots, \theta_{F,1})$ restricts to a basis of $N_F^* \cong \operatorname{Ann}(F)$. For each relevant face, choose a tuple outside the corresponding E_F. The resulting family is admissible by Definition 33. □

Definition 35. *Fix $(\mathcal{P}, \omega)$ and admissible auxiliary data Θ. Define the flag selection strategy $\mathcal{F}_\Theta$ by $\mathcal{F}_\Theta(\mathfrak{c}) := FL_\Theta(\mathfrak{c})$. Thus $G_{\mathcal{F}_\Theta}$ is the directed graph associated to this strategy by Definition 27.*

7 Acyclicity of the greedy strategy

Fix admissible auxiliary data Θ, and let $G_{\mathcal{F}_\Theta}$ be the graph of the greedy strategy. Every possible cycle lies in a subgraph $\mathcal{G}_F$ induced by chambers with one fixed optimal face F. Within $\mathcal{G}_F$, we compare the source and target greedy executions at their first differing choice and show that the selected execution score strictly increases. This rules out cycles. The progress and escape conclusions then follow from the abstract results of Section 5.

7.1 Reduction to a fixed optimal face

Lemma 16 gives weak increase of $\lambda(-,\omega)$ along every edge and weak growth of optimal-face closures along equal-λ edges. Thus, by Lemma 17, every cycle is contained in the subgraph induced by the chambers with one fixed optimal face.

Definition 36. *Fix a positive-codimension face F. Let $\mathcal{G}_F$ be the directed graph whose vertices are the chambers $\mathfrak{c}$ such that $\lambda(\mathfrak{c},\omega) < \infty$ and $F(\mathfrak{c},\omega) = F$, and whose edges are directed edges of $G_{\mathcal{F}_\Theta}$ whose source and target both lie in this vertex set.*

Lemma 34. *($\to$ proof) If $\mathcal{G}_F$ is acyclic for every positive-codimension face F, then $G_{\mathcal{F}_\Theta}$ is acyclic.*

Proof of Lemma 34. Suppose there is a directed cycle $\mathfrak{c}'_0 \to \mathfrak{c}'_1 \to \cdots \to \mathfrak{c}'_m = \mathfrak{c}'_0$. Infinite-$\lambda$ chambers have no outgoing edges by Definitions 35 and 27, so all chambers in the cycle have finite λ-value. By Lemma 17, applied to this cycle, all optimal face closures along the cycle are equal. For arrangement faces, equality of closures determines the face: each face is the relative interior of its closure in its affine span, equivalently the face with a given sign vector is unique. Hence all optimal faces along the cycle are equal. Thus the directed cycle lies in $\mathcal{G}_F$ for $F = F(\mathfrak{c}'_0,\omega)$, contradiction. □

7.2 The first changed face of a target flag

For a sign-indexed target, the largest flipped coordinate determines the first face at which the target flag can leave the source flag. The two flags agree down to that face and then leave it in opposite transverse directions.

Lemma 35. *($\to$ proof) Assume the setting of Theorem 2. Suppose $\alpha \neq 1^r$, and let i^* be the largest index such that $\alpha_{i^*} = -1$. Then $F_i = F_i(\alpha)$ for $i^* \le i \le r$. In particular, $F_i \subseteq \overline{\mathfrak{b}_{FL}(\alpha)} \cap \bar{\mathfrak{c}}$ for $i^* \le i \le r$.*

Moreover, $\operatorname{lin}(F_{i^*-1}(\alpha)) = \operatorname{lin}(F_{i^*-1})$, and there is $w \in \operatorname{lin}(F_{i^*-1}) \setminus \operatorname{lin}(F_{i^*})$ such that $w \in D_{F_{i^*}}(F_{i^*-1})$ and $-w \in D_{F_{i^*}}(F_{i^*-1}(\alpha))$.

Proof of Lemma 35. Fix $i \in [i^*, r]$. We compare the sign descriptions in Theorem 2. If $q \in \mathfrak{Q}(F_i,\mathcal{P})$, then both faces have sign 0. If $q \in \mathfrak{Q}(F,\mathcal{P}) \setminus \mathfrak{Q}(F_i,\mathcal{P})$, then $Z(q) \in \mathcal{H}_k$ for some $k > i$, so $\operatorname{gi}(q) = k > i \ge i^*$. Hence $\alpha_{\operatorname{gi}(q)} = 1$ by maximality of i^*. Thus the two nonzero signs are both $\operatorname{sign}_{\mathfrak{c}}(q)$. For $q \notin \mathfrak{Q}(F,\mathcal{P})$, both signs are $\operatorname{sign}_{\mathfrak{c}}(q)$. Hence F_i and $F_i(\alpha)$ have the same sign pattern, so $F_i = F_i(\alpha)$.

The inclusion follows because $F_i(\alpha) \subseteq \overline{F_0(\alpha)} = \overline{\mathfrak{b}_{FL}(\alpha)}$ and $F_i \subseteq \overline{F_0} = \bar{\mathfrak{c}}$.

Put $i := i^*$. The sign formula of Theorem 2 shows that

$$\mathfrak{Q}(F_{i-1}(\alpha),\mathcal{P}) = \mathfrak{Q}(F_{i-1},\mathcal{P}).$$

Lemma 3 therefore gives $\operatorname{aff}(F_{i-1}(\alpha)) = \operatorname{aff}(F_{i-1})$, hence equality of their linear spans.

For $q \in \mathfrak{Q}(F_i, \mathcal{P}) \setminus \mathfrak{Q}(F_{i-1}, \mathcal{P})$, one has $\mathrm{gi}(q) = i$, so the signs of F_{i-1} and $F_{i-1}(\alpha)$ on q are opposite. Their zero forms are the same. Consequently

$$R_{F_i}(F_{i-1}(\alpha)) = -R_{F_i}(F_{i-1}).$$

Choose $\xi \in R_{F_i}(F_{i-1})$. Since this ray spans $\mathrm{lin}(F_{i-1})/\mathrm{lin}(F_i)$, choose $w \in \mathrm{lin}(F_{i-1})$ with $\pi_{F_i}(w) = \xi$. Then $w \notin \mathrm{lin}(F_i)$, while the displayed equality of rays gives $w \in D_{F_i}(F_{i-1})$ and $-w \in D_{F_i}(F_{i-1}(\alpha))$. □

Lemma 36. *(→ proof) Let $FL_\Theta(\mathfrak{c}) = (F_r, \ldots, F_0)$, and let $\mathfrak{d} = \mathfrak{b}_{FL_\Theta(\mathfrak{c})}(\alpha)$ be a target of a finite-λ source chamber $\mathfrak{c}$. Suppose $F(\mathfrak{d}, \omega) = F(\mathfrak{c}, \omega) = F$. Then $\alpha_r = 1$.*

Proof of Lemma 36. Since F is the optimal face of $\mathfrak{c}$, Lemma 32 says that $h_F = -d\omega$ is positive on $C_F(\mathfrak{c}) \setminus \{0\}$. If $\alpha_r = -1$, then Lemma 35 gives a vector $w \notin \mathrm{lin}(F)$ with $w \in D_F(F_{r-1}) \subseteq D_F(\overline{\mathfrak{c}})$ and $-w \in D_F(\overline{\mathfrak{d}})$. Thus $h_F(\pi_F(w)) > 0$, so $h_F(\pi_F(-w)) < 0$. If F were also the optimal face of the target chamber, then Lemma 32 would make h_F positive on $\pi_F(-w)$, a contradiction. Hence $\alpha_r = 1$. □

7.3 Selected-score progress at the first changed face

We compare the source and target greedy executions directly. Above the first changed face, every source choice is also a candidate for the target. If the target first chooses a different covering face there, uniqueness of the maximizer gives a strict score increase. Otherwise the two executions reach the first changed face with the same updated covectors. The opposite transverse direction below that face then produces a target direction whose score is strictly larger than the source score.

Definition 37 (Selected execution score)**.** *Fix a positive-codimension optimal face F, put $r := \mathrm{codim}(F)$, and let $\mathfrak{c}$ be a chamber with $F(\mathfrak{c}, \omega) = F$. The selected execution score is the tuple of maximizing values produced by the recursion in Definition 34:*

$$\widehat{\sigma}_F(\mathfrak{c}) := \big(s_r(\mathfrak{c}), s_{r-1}(\mathfrak{c}), \ldots, s_2(\mathfrak{c})\big).$$

It is the empty tuple when $r \le 1$. We compare these tuples lexicographically only among chambers whose optimal face is F.

Proposition 1 (Score increase within $\mathcal{G}_F$)**.** *(→ proof) Fix a positive-codimension face F, and suppose Θ_F is ray-regular at $(\mathfrak{e}, F)$ for every vertex $\mathfrak{e}$ of $\mathcal{G}_F$. If $\mathfrak{c} \to \mathfrak{d}$ is an edge in $\mathcal{G}_F$, then $\widehat{\sigma}_F(\mathfrak{d}) >_{\mathrm{lex}} \widehat{\sigma}_F(\mathfrak{c})$.*

Proof of Proposition 1. Put $r := \mathrm{codim}(F)$, and write $FL = FL_\Theta(\mathfrak{c}) = (F_r, \ldots, F_0 = \mathfrak{c})$. By Definition 27, there is a sign vector $\alpha \in \{\pm 1\}^r \setminus \{1^r\}$ such that $\mathfrak{d} = \mathfrak{b}_{FL}(\alpha)$. Since $\mathfrak{c}$ and $\mathfrak{d}$ are both vertices of $\mathcal{G}_F$, their optimal faces agree: $F(\mathfrak{d}, \omega) = F(\mathfrak{c}, \omega) = F$. Lemma 36 gives $\alpha_r = 1$. Since $\alpha \neq 1^r$, this forces $r > 1$. Let $i = i^* := \max\{a : \alpha_a = -1\}$, so $i < r$. Let $FL(\alpha) = (F_r(\alpha), \ldots, F_0(\alpha) = \mathfrak{d})$ be the sign-indexed flag from Theorem 2. Lemma 35 gives

$$F_a = F_a(\alpha) \subseteq \overline{\mathfrak{c}} \cap \overline{\mathfrak{d}} \qquad (i \le a \le r).$$

Let $FL_\mathfrak{d} := FL_\Theta(\mathfrak{d}) = (G_r, \ldots, G_0 = \mathfrak{d})$. First suppose that the two executions make different choices before descending to level $i+1$. Let j be the level of their first differing choice. Thus, as the recursions descend from level r, they arrive at the same face $F_j = G_j$, with the same updated covectors, but select different covering faces F_{j-1} and G_{j-1}. Necessarily $i + 2 \le j \le r$, and all higher-indexed faces agree:

$$G_a = F_a \quad (j \le a \le r), \qquad G_{j-1} \neq F_{j-1}.$$

The common-face relation above puts both F_j and F_{j-1} in $\overline{\mathfrak{d}}$, and their codimensions differ by one. Hence $F_{j-1} \in \mathrm{Cov}_{F_j}(\mathfrak{d})$. Thus the source choice F_{j-1} is a candidate in the target maximization, with value $s_j(\mathfrak{c})$. The target instead chooses the distinct face G_{j-1}, its unique maximizer, and hence

$$s_j(\mathfrak{d}) > s_j(\mathfrak{c}).$$

All selected values with index greater than j agree, so $\widehat{\sigma}_F(\mathfrak{d}) >_{\mathrm{lex}} \widehat{\sigma}_F(\mathfrak{c})$, and we are done.
It remains to consider the case in which no such earlier difference occurs, namely

$$G_a = F_a \qquad (i+1 \le a \le r).$$

The preceding choices, and hence the updated covectors at F_{i+1}, are then the same for both executions. Write

$$\eta := \eta_{i+1}, \qquad \varphi := \theta_{F,i}^{[i+1]}, \qquad s := s_{i+1}(\mathfrak{c}), \qquad \delta := s\eta - \varphi = \eta_i.$$

The common updated frame is ray-regular at both $(\mathfrak{c}, F_{i+1})$ and $(\mathfrak{d}, F_{i+1})$. The source step chooses F_i, so Lemma 31 shows that δ vanishes on $L_{F_{i+1}}(F_i)$ and, as a covector on N_{F_i}, is positive on $C_{F_i}(\mathfrak{c}) \setminus \{0\}$.
Choose $u \in D_{F_{i+1}}(F_i)$. Since $F_i \in \mathrm{Cov}_{F_{i+1}}(\mathfrak{d})$, we have $\eta(u) > 0$. Lemma 35 supplies a vector

$$w \in D_{F_i}(F_{i-1}) \setminus \mathrm{lin}(F_i) \quad \text{with} \quad -w \in D_{F_i}(F_{i-1}(\alpha)).$$

Thus $\delta(w) > 0$. Also $-w \in D_{F_i}(\overline{\mathfrak{d}})$, so Lemma 28 gives

$$z_t := u - tw \in D_{F_{i+1}}(\overline{\mathfrak{d}})$$

for all sufficiently small $t > 0$. For such t, $\eta(z_t) > 0$. Since δ vanishes on $D_{F_{i+1}}(F_i)$, we have $\delta(u) = 0$, and therefore

$$\frac{\varphi(z_t)}{\eta(z_t)} - s = -\frac{\delta(z_t)}{\eta(z_t)} = \frac{t\delta(w)}{\eta(z_t)} > 0.$$

The target form of Lemma 31, applied to the nonzero direction $\pi_{F_{i+1}}(z_t) \in C_{F_{i+1}}(\mathfrak{d})$, now gives

$$s_{i+1}(\mathfrak{d}) \ge \frac{\varphi(z_t)}{\eta(z_t)} > s.$$

The selected values with index $> i+1$ agree, so $\widehat{\sigma}_F(\mathfrak{d}) >_{\mathrm{lex}} \widehat{\sigma}_F(\mathfrak{c})$. □

Theorem 4 (Acyclicity of the greedy strategy)**.** *(→ proof) Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, and let $\omega \in \mathrm{Aff}(A, \mathbb{R}) \setminus \mathrm{Con}(A, \mathbb{R})$. If the auxiliary data Θ are admissible, then $G_{\mathcal{F}_\Theta}$ is acyclic.*

Proof of Theorem 4. Admissibility of Θ says exactly that Θ_F is ray-regular at $(\mathfrak{e}, F)$ for every chamber $\mathfrak{e}$ with $F(\mathfrak{e}, \omega) = F$. Thus Proposition 1 applies to every edge $\mathfrak{c} \to \mathfrak{d}$ in every $\mathcal{G}_F$, giving $\widehat{\sigma}_F(\mathfrak{d}) >_{\mathrm{lex}} \widehat{\sigma}_F(\mathfrak{c})$. A directed cycle in $\mathcal{G}_F$ would force a strict lexicographic increase after returning to its starting chamber, which is impossible. Hence every $\mathcal{G}_F$ is acyclic. Lemma 34 gives acyclicity of $G_{\mathcal{F}_\Theta}$. □

Theorem 5 (Progress and escape for the greedy strategy)**.** *(→ proof) Let $\mathcal{P}$ be a locally finite affine hyperplane arrangement, and let $\omega \in \mathrm{Aff}(A, \mathbb{R}) \setminus \mathrm{Con}(A, \mathbb{R})$. At every relevant face, the auxiliary covectors defining the greedy strategy may be chosen outside a proper algebraic subset. Every resulting admissible choice Θ defines a strategy $\mathcal{F}_\Theta$ which is pathwise progressive and uniform progressive. If $\mathcal{P}$ has finitely many normal spaces, then $\mathcal{F}_\Theta$ is additionally pathwise escaping and uniform escaping.*

Proof of Theorem 5. Lemma 33 gives the asserted facewise-generic choices. Fix any resulting admissible auxiliary data Θ. By Definition 35, $G_{\mathcal{F}_\Theta}$ is the graph of a global flag selection strategy. Theorem 4 shows that this graph is acyclic, so Corollary 3 gives both pathwise and uniform progress. If $\mathcal{P}$ has finitely many normal spaces, Theorem 3 gives uniform escaping, and Lemma 18 then gives pathwise escaping. □

References

[1] Reinhard Diestel. *Graph Theory*, volume 173 of *Graduate Texts in Mathematics*. Springer, Berlin, 5th edition, 2017. doi:10.1007/978-3-662-53622-3.

[2] Brennan Dury. Affine residues and chamber jumps for hyperplane arrangements. Manuscript, 2026.

[3] Arkadi Nemirovski. *Introduction to Linear Optimization*. World Scientific, Singapore, 2024. Postprint dated March 17, 2024. URL: https://www2.isye.gatech.edu/~nemirovs/WSPrinted.pdf, doi:10.1142/13457.

[4] R. Tyrrell Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1970.

[5] András Szenes. Iterated residues and multiple Bernoulli polynomials. *International Mathematics Research Notices*, 1998(18):937–956, 1998. arXiv:hep-th/9707114.