Deletion-Star Ceilings and a 1.28249 . . . Lower Bound for Cost-Preserving Single-Source Unsplittable Flows

Sergey Nikolenko[1,2]

[1]St. Petersburg Department of the Steklov Institute of Mathematics, St. Petersburg, Russia
[2]St. Petersburg State University, St. Petersburg, Russia

July 31, 2026

Abstract

For a single-source unsplittable flow we ask for a rounding of a feasible fractional flow $\mathbf{x}$ to one path per terminal that satisfies both the cost constraint $\mathbf{c}^\mathsf{T}\mathbf{y} \leq \mathbf{c}^\mathsf{T}\mathbf{x}$ and the overload constraint $y_a \leq x_a + Cd_{\max}$ on every arc. Goemans conjectured $C = 1$; this was disproved in July 2026 by a seven-vertex instance with critical constant 16/15, and our previous work [14] raised the bound to $(299 - 41\sqrt{41})/32 = 1.1397\ldots$ with a four-terminal planar family. We continue that program on three fronts.

Records: we present a seventeen-terminal common-point interval instance that certifies

$$C \;\geq\; \frac{1282494797984843521}{10^{18}} = 1.282494797984843521\ldots,$$

verified independently by enumerating all 2^{17} routings from the raw 67-arc list, together with an 18-atom exact convex-hull certificate.

Upper bounds: we prove the first *unconditional* upper bound below 2 for a nontrivial family of the extremal cells. Representing the class as a weighted two-permutation prefix system, we introduce the deletion-star residual B, show that a bound $B \leq \beta$ yields congestion at most $1 + \beta$, and prove that on the codimension two cost face defined by the equality $\sum_i f_i = k - 2$ one always has $B \leq \frac{1}{2}$, for arbitrary orders, arbitrary demands (zero demands allowed) and arbitrary shares. This gives congestion at most $\frac{3}{2}$ on every common-point cell whose first cost-good routing rank is $k - 1$. The constant $\frac{1}{2}$ is asymptotically sharp at every fixed codimension q, with the exact value $B_{mq,q} = \frac{1}{2} - \frac{1}{2m}$, and the sharpness persists through a non-divisible two-scale demand perturbation. Throughout the paper, theorems are named by the mechanism that proves them (proportional credits, a single deletion star, the multi-star threshold), and subsequent negative results (of which this program has produced several) make the scope of these methods more narrow rather than retract theorems.

Structure: we show that the naive continuum limit of the class has value *zero*, so the class supremum is a genuinely microscopic quantity; we give a cost-preserving transfer through Knuth's two-way rounding network with exact radius $B(\mathbf{d}) = \min_{\lambda \geq 0}(\lambda + \sum_i |d_i - \lambda|)$; we prove the lower-box statement (LB) for all 2^{k-1} deque orders and for a maximal switched-interval class of signed tree networks; and we record exact minimal obstructions showing that every scalar two-order induction fails, leaving a convex Pareto transfer as the surviving mechanism. Finally, literal three-order overlays are dominated at 6/5.

Every numerical claim is certified in exact rational or symbolic arithmetic by programs that rebuild each instance from its raw arc list; the artifacts are published at `https://github.com/snikolenko/unsplittable-flows`.

1 Introduction

An instance of *single-source unsplittable flow* (SSUF) consists of an acyclic digraph $D = (V, A)$, a source s, terminals $t_1, \dots, t_k$ with demands $d_i > 0$, and nonnegative arc costs $\mathbf{c} \in \mathbb{R}^{A}_{\geq 0}$. Let $\mathbf{x}$ be a feasible fractional flow sending d_i units from s to t_i, and let $d_{\max} = \max_i d_i$. An *unsplittable* flow picks one $s \to t_i$ path per terminal and sends all d_i units along it; we write $\mathbf{y}$ for its vector of arc loads.

Dinitz, Garg and Goemans [3] proved that some unsplittable $\mathbf{y}$ always satisfies $y_a \leq x_a + d_{\max}$ on every arc in the absence of costs. Goemans conjectured the cost-preserving strengthening, that the same additive bound holds together with $\mathbf{c}^{\mathsf{T}}\mathbf{y} \leq \mathbf{c}^{\mathsf{T}}\mathbf{x}$, and this conjecture stood for a quarter of a century. It is known to hold when demands form a divisibility chain [17, 12] and on series-parallel digraphs [9]; for planar (and bounded-genus) instances $2d_{\max}$ suffices [21, 22]; the work [19] converts cost-free error bounds into cost-preserving ones.

In July 2026, the exact conjecture was disproved by a seven-vertex instance of critical constant 16/15 [15], and our previous work [14] raised the lower bound to

$$C^{\star}_{\mathrm{I}} = \frac{299 - 41\sqrt{41}}{32} = 1.139747070789\dots$$

with a four-terminal planar family, identified *baseline loads* as the mechanism behind all known counterexamples, and conjectured that the exact supremum of forced constants is 2.

This paper is a continuation of this work. It is written to be self-contained and readable on its own: Section 3 restates, with proofs, everything from [14] that is used here, and only the numerical record ladder is quoted rather than re-derived.

Our results organize into three layers with three conjectures. *Lower bounds*: every known extremal instance is k intervals with one common arc, i.e., two interleaved rankings of the terminals, and the certified record ladder (Fig. 1) climbs toward a conjectured class supremum of $\frac{4}{3}$. *Upper bounds*: the deletion-star calculus bounds what any cell of complement mass q can force; we prove the exact frontier $\frac{1}{2}$ at $q = 2$ (hence congestion $\frac{3}{2}$), show that it is sharp at every codimension, and also show that it fails beyond that: the ceilings march toward 2. As for universal upper bounds, the lower-box route would give the constant 2 for all instances, and it holds on every record numerically with deficit below 10^{-9}, but proving it remains an open problem.

All lower bound families in [14] and here are *common-point interval systems.* This means that each terminal i has exactly two simple $s \to t_i$ paths, which we call an *early* and a *late* exit from the directed spine, and switching it from early to late adds d_i on a consecutive interval $I_i = [\ell_i, r_i]$ of spine arcs; all intervals share one common arc. Lemma 3.4 below shows that the shared rows of such a system are exactly the prefixes of *two* orders of the terminals. The abstract problem is therefore a weighted two-permutation prefix rounding problem with one preserved linear functional, and this is the object that we study in this work.

We make the following contributions, which split into two halves: items 1–2 push the *lower* bound up and delimit how far that pushing can go, items 3–5 push *upper* bounds down on ever larger pieces of the class.

Lower bounds, and the limits of the lower-bound method

1. A seventeen-terminal record. Section 4 certifies

$$C \;\geq\; \frac{1282494797984843521}{10^{18}} = 1.282494797984843521\dots \qquad \text{(Theorem 4.1)},$$

by a $k = 17$, rank-seven common-point instance, found by joint optimization in the coordinates $(d_i,\, p_i = d_i f_i)$ in which every fixed-witness overload inequality is linear. The complete instance is in Appendix A, and Figure 5 shows it as a system of intervals. We provide a certificate that enumerates all 131072 routings from the raw 67-arc list, exhibits the 15 minimizers, checks that every strictly better routing has rank at most six, and closes the argument with an 18-atom exact convex-hull mixture. Together with the results of [14] this gives the ladder of Figure 1.

2. The naive continuum is degenerate. We prove (Theorem 5.1) that on an atomless measure space, the two-order minimax with one preserved functional is identically zero. This means that the class supremum is not determined by any macroscopic limit object, it is a first-order atomic quantity. This is a negative result, and it means that an empirical fit of the form $4/3 - \alpha/k$ that would be very tempting to read off the ladder of counterexamples has no continuum justification.

Upper bounds for common-point interval systems

3. Deletion-star ceilings. Section 6 is the technical core. We define the *deletion-star residual* $B(\pi, \sigma, \mathbf{d}, \mathbf{f})$ (Definition 6.1) and prove a support reduction lemma: $B \le \beta$ forces congestion at most $1 + \beta$ (Lemma 6.2). The natural target $\beta = 1/3$, which would give a $4/3$ ceiling, is refuted exactly since already at $k = 8$ one has $B = 3/8$ (Example 6.5). The main result is that

$$\text{if} \quad \sum_i f_i = k - 2 \quad \text{then} \quad B \le \tfrac{1}{2} \qquad \text{(Theorem 6.6)},$$

with *no* assumption on the orders, the demands (zero demands are allowed) or the shares, and hence

$$\text{congestion} \;\le\; \tfrac{3}{2} \;\; \text{on every cell of first cost-good rank } k-1 \qquad \text{(Corollary 6.8)}.$$

To our knowledge this is the first unconditional upper bound below 2 for a nontrivial family of the extremal cells. Theorem 6.9 shows that the constant $\frac{1}{2}$ is not an artifact of codimension two: for identity and reverse orders, unit demands and uniform shares, $B_{k,q} = \frac{1}{2} - \frac{q}{2k}$ exactly when an explicit capacity condition holds, so $B_{mq,q} = \frac{1}{2} - \frac{1}{2m} \to \frac{1}{2}$ for every fixed q; finally, Theorem 6.10 shows that the same frontier carries over through a two-scale demand perturbation that is not a divisibility chain.

4. Two-order upper-bound mechanisms. Section 7 considers a different route to upper bounds. These statements are about *rounding*, not flows, and Table 1 summarizes how each of our results bounds C.

There are two conversions. First, by Observation 3.2 the cost constraint is a single signed linear functional of $\big(d_i(z_i - f_i)\big)_i$, so a rounding theorem producing a $\mathbf{z}$ that is prefix-accurate *and* respects that functional *is* a cost-preserving routing, with congestion equal to its prefix accuracy. A min-cost version of Knuth's two-way rounding network does exactly this (Theorem 7.2), with accuracy $B(\mathbf{d}) = \min_{\lambda \ge 0}(\lambda + \sum_i |d_i - \lambda|)$; hence $C \le B(\mathbf{d})/d_{\max}$ for every instance of the class with that demand vector. Proposition 7.3 evaluates $B(\mathbf{d})$ exactly and shows that $B(\mathbf{d}) = d_{\max}$, i.e., this argument alone gives the conjecturally optimal $C \le 1$ precisely when all demands except one maximum coincide; note that this scope includes incommensurable two-scale systems and therefore does not follow from the divisibility-chain theorem.

The second conversion goes through the lower box (Definition 3.8). Theorem 7.5 (all 2^{k-1} deque orders, *arbitrary* demands) and Theorem 7.7 (a maximal switched-interval class of signed tree networks) produce a routing that is two-sided $d_{\max}$-good; such a routing verifies (LB) on those

Instances covered	Upper bound	Source	Condition
(a) restricted classes: these bound C_{cls}, not C			
common-point cells of complement mass q	$\max(1, q-\eta)$	Prop. 3.5	
common-point cells of complement mass 2	$\frac{3}{2}$	Cor. 6.8	
two-order systems, any demands	$B(\mathbf{d})/d_{\max}$	Thm. 7.2	
which is 1 iff $d_{(1)} = \cdots = d_{(k-1)}$	1	Prop. 7.3	
deque orders, any demands	2	Thm. 7.5	
switched-interval signed trees	2	Thm. 7.7	
(b) all instances: these bound C itself			
all planar instances	2	[21]	Known
all instances	2	Prop. 3.10	the statement (LB)

Table 1: Ceilings on the cost-preserving constant proved or used here. The split between the two blocks matters. A *universal* constant C must serve every instance, so only block (b) bounds C, and there only conditionally except on planar supports. The rows of block (a) bound $C_{\mathrm{cls}} := \sup C^{\star}(\mathbf{x})$ over the stated class of instances only, i.e., they show how far our program of constructing counterexamples can go. The best lower bound of Section 4 is $C \geq C_{\mathrm{cls}} \geq 1.2825\ldots$, so the gap remaining for C is [1.2825, 2].

classes, and Proposition 3.10 converts (LB) into $C \leq 2$. The conclusion is weaker than the first route but does not make any assumptions about the demands.

Section 7.4 then marks the boundary: three minimal exact instances on which every *scalar* strengthening of the one-order induction fails. Section 7.5 isolates the one statement that survives, a convex Pareto transfer, and identifies the exact missing operation.

The two conjectures. The results above organize the remaining uncertainty into two displayed conjectures, which the rest of the paper treats as its coordinate system.

Conjecture 1.1 (class supremum)**.** $C_{\mathrm{cls}} = \frac{4}{3}$*: the supremum of critical constants over common-point cells equals* 4/3*, approached but not attained along the record ladder of Figure 1.*

Conjecture 1.2 (universal supremum [14])**.** *The supremum of forced cost-preserving constants over* all *instances equals* 2*, and every instance admits a routing with* $\mathbf{y} \leq \mathbf{x} + 2d_{\max}\mathbf{1}$ *and* $\mathbf{c}^{\top}\mathbf{y} \leq \mathbf{c}^{\top}\mathbf{x}$.

We emphasize that the deletion-star ceilings of Section 6 cannot separate Conjecture 1.1 from Conjecture 1.2: the ceiling of a cell of complement mass q tends to 2 as q grows, and the record cells have growing complement mass (the $k = 17$ record has $q = 11$), so any proof of a class ceiling below 2 must use a mechanism beyond a single deletion star per support; one natural candidate could be the multi-star threshold of Section 6.

5. Three-order overlays are dominated. In Section 8, we show that on the smallest literal three-replica equality overlay, the fractional load is an exact convex combination of routings of congestion at most 6/5, so no signed arc-toll vector can separate anything above 6/5. Any escape from the two-order world must place positive fractional baseline mass on hybrid paths.

Verification. Every numerical claim below has been recomputed in exact rational or symbolic arithmetic by at least two, and for the headline results four, independently written programs that rebuild each instance from its raw arc list, rediscover all simple source–terminal paths by depth-first search, and enumerate all integral routings; one of these auditors is a separate model instance working only from the prose statements and certificates, never from the original code.

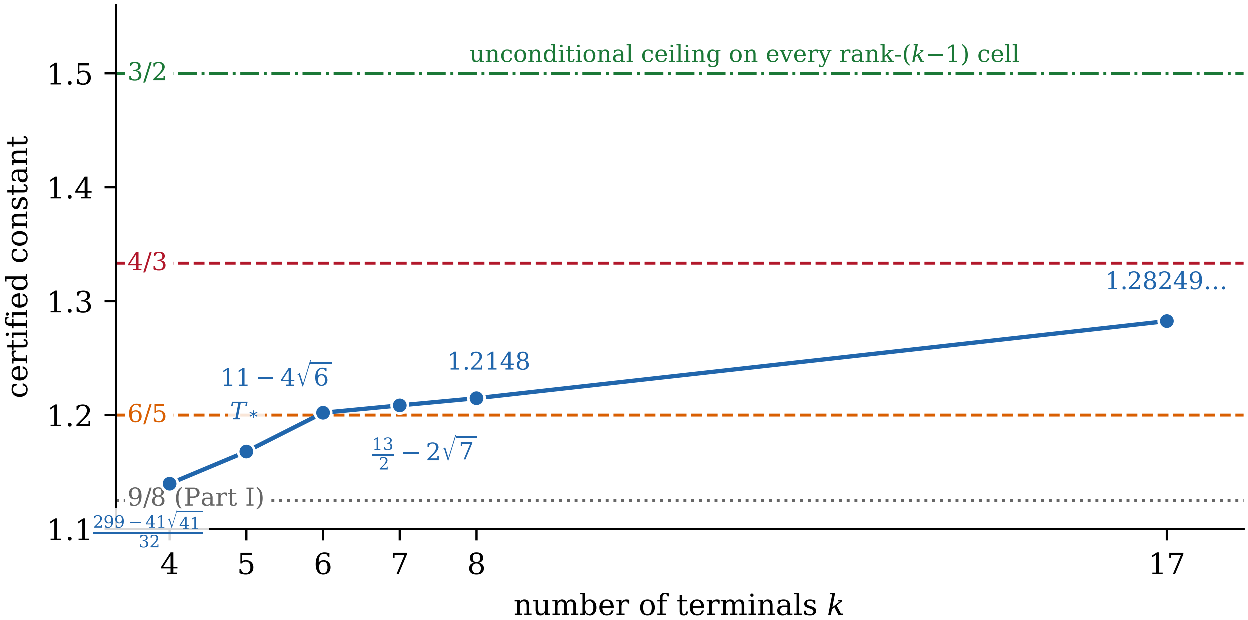

Figure 1: The certified record ladder. Each point is the exact supremum of an explicit common-point interval family with k terminals; $k = 4$ is from [14]. The dash-dotted line is the unconditional ceiling of Corollary 6.8, which applies to cells of complement mass two (first cost-good rank $k - 1$); the ladder entries have complement mass at least three, so the ceiling does not bind on them.

Every certificate comes with an executable convention block fixing its orders, normalizations, and sign conventions in code. Section 9 describes the pipeline, and all artifacts are available at `https://github.com/snikolenko/unsplittable-flows`.

AI use. Similar to the previous paper [14], this work was produced in close collaboration with frontier LLMs; the construction search, exact certificate pipeline, and much of the theory development were carried out with GPT-5.6 Sol, Claude Fable 5, and Claude Opus 5. As I said in [14] and am happy to repeat here, this work is much more theirs than my own.

The paper is organized as follows. Section 2 surveys related work, and Section 3 is a self-contained preliminaries section. Section 4 gives the new record lower bound, Section 5 presents the continuum theorem (negative result). Section 6 discusses the deletion-star ceilings, Section 7 presents the upper bound mechanisms, and Section 8 the three-order frontier. Finally, Section 10 lists open problems and concludes the work.

2 Related work

Additive rounding for unsplittable flows. Dinitz, Garg and Goemans [3] introduced the additive formulation and proved the $+d_{\max}$ congestion theorem; Goemans' cost-preserving strengthening is stated there. Kolliopoulos and Stein [5] and Skutella [17] gave the classical bicriteria guarantees $y_a \le 2x_a + d_{\max}$ with $\mathbf{c}^\top \mathbf{y} \le 2\mathbf{c}^\top \mathbf{x}$ and $\mathbf{c}^\top \mathbf{y} \le \mathbf{c}^\top \mathbf{x}$ respectively; see [16] for the exact historical statements. Note that a guarantee relative to $2\mathbf{x}$ yields *no* finite additive constant relative to $\mathbf{x}$. Martens, Salazar and Skutella [10] gave convex decompositions for rounded demands. Morell and Skutella [12] proved the one-sided lower bound $y_a \ge x_a - d_{\max}$, the two-sided $x_a/2 - d_{\max} \le y_a \le 2x_a + d_{\max}$, and the two-sided $\pm d_{\max}$ statement for divisibility chains, conjecturing the two-sided $\pm d_{\max}$ statement in general. Traub, Vargas Koch and Zenklusen [21, 22] proved that conjecture for planar and bounded-genus instances and deduced the cost-preserving planar bound $|y_a - x_a| \le 2d_{\max}$. Almoghrabi, Skutella and Warode [9] proved that in a series-parallel digraph every multiflow is a convex combination of unsplittable multiflows with $|y_a - x_a| < d_{\max}$; averaging gives Goemans' conjecture there. Swamy, Traub, Vargas Koch and Zenklusen [19] showed that a face-preserving rounding algorithm with symmetric error $\alpha d_{\max}$ yields, for every $\lambda \in (0, 1]$, a routing with $\mathbf{c}^\mathsf{T}\mathbf{y} \le \lambda^{-1}\mathbf{c}^\mathsf{T}\mathbf{x}$ and $|y_a - x_a| \le (1 + \lambda)\alpha d_{\max}$; at $\lambda = 1$ this turns the Morell–Skutella conjecture into a cost-preserving $2d_{\max}$ guarantee. Rybin [15] disproved the exact conjecture, and our previous work [14] gave the first lower bound above 9/8 (which is the natural limit of Rybin's family of counterexamples).

Two-way rounding and discrepancy. The abstract projection of the class studied here is a one-sided weighted two-permutation prefix discrepancy problem with one preserved linear functional. For unit weights, the two-order problem is exactly Knuth's *two-way rounding* [4], whose network formulation we use in Section 7 in a min-cost variant so that a linear functional is preserved. Tijdeman's chairman assignment theorem [20] bounds the prefix discrepancy of a *common*-order assignment to m agents by $1 - \frac{1}{2m-2}$; Liu and Reis [7] extended this to weighted items with the same constant, which in particular confirms the Morell–Skutella conjecture in the common-order ("carpooling") case. Their theorem is the driving force behind our Theorems 7.5 and 7.7. For general discrepancy background, see [2, 18, 11]; the three-permutation lower bound of Newman and Nikolov [13] refutes Beck's conjecture. Asano, Katoh, Obokata and Tokuyama [1] prove that rounding for a union of two laminar families is exactly solvable in the unit-weight case; Section 7.5 explains precisely how column weights make that argument invalid.

Purification. Theorem 5.1 is an application of Lyapunov's convexity theorem [8] (see [6] for a short proof) in the following form: the range of a finite-dimensional atomless vector measure is convex and coincides with the set of values of its fractional relaxations.

3 Preliminaries

3.1 Routings, critical constants, baseline cancellation

For a finite index set E and $F \subseteq E$ we write $\boldsymbol{\chi}(F) \in \{0, 1\}^E$ for the incidence vector of F; we use this for sets of arcs, for a path P we write $\boldsymbol{\chi}(P)$ for the incidence vector of its arc set, and in Section 7 we also use $\boldsymbol{\chi}(F)$ notation for sets of terminal positions.

We write $\mathbf{1}$ and $\mathbf{0}$ for the all-ones and all-zero vectors, of whatever dimension the context requires, and $\mathbf{e}_i$ for the i-th unit vector; when it carries a subscripted condition, as in $\mathbf{1}_{\{i \in O\}}$, the same symbol denotes the indicator of that condition.

A *routing* assigns to each terminal i one simple path P_i from s to t_i and sends all d_i units of demand along it; its vector of arc loads is therefore $\mathbf{y} = \sum_i d_i \boldsymbol{\chi}(P_i)$. Given a feasible fractional flow $\mathbf{x}$ and $C \ge 0$, a routing is *C-good* if $y_a \le x_a + Cd_{\max}$ for every arc a, and *cost-good* if $\mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}$. When both sides matter we call a routing *upper-good* if $\mathbf{y} \le \mathbf{x} + d_{\max}\mathbf{1}$, *lower-good* if $\mathbf{y} \ge \mathbf{x} - d_{\max}\mathbf{1}$, and *two-sided $d_{\max}$-good* if it is both. Throughout the paper, we normalize $d_{\max} = 1$ whenever convenient.

Definition 3.1 (per-instance constant)**.** The *critical constant* of a pair (instance, $\mathbf{x}$) is

$$C^\star(\mathbf{x}) \;:=\; \min\{\max_a (y_a - x_a)/d_{\max} \;:\; \mathbf{y} \text{ is a routing with } \mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}\}.$$

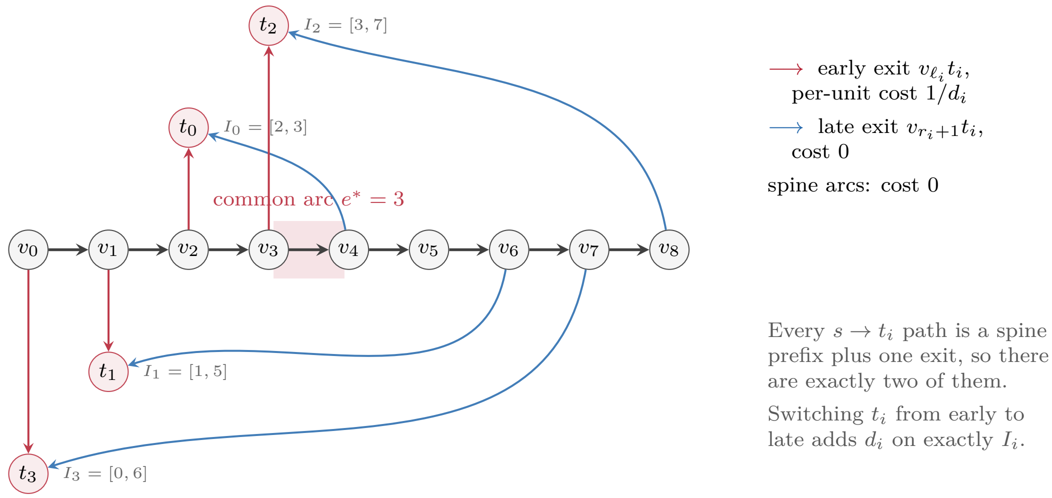

Figure 2: A common-point interval system (Definition 3.3). The digraph is a directed spine plus two private exit arcs per terminal, so every $s \to t_i$ path is a spine prefix followed by one exit and there are exactly two of them. Terminal i is drawn at the left endpoint of its interval; the four intervals shown all contain the common arc e^*.

Any universal constant C for which the cost-preserving rounding statement holds satisfies $C \geq C^\star(\mathbf{x})$ for every instance and every feasible $\mathbf{x}$.

The following identity is used throughout; it is Observation 2.2 of [14].

Observation 3.2 (baseline cancellation)*.* Suppose terminal i has exactly two simple $s \to t_i$ paths P_i^0, P_i^1, the fractional flow sends $d_i f_i$ on P_i^1 and $d_i(1-f_i)$ on P_i^0, and a routing selects $P_i^{z_i}$. Put $\boldsymbol{\Delta}_i = \boldsymbol{\chi}(P_i^1) - \boldsymbol{\chi}(P_i^0) \in \{-1,0,1\}^A$. Then

$$\mathbf{y} - \mathbf{x} = \sum_i d_i(z_i - f_i)\boldsymbol{\Delta}_i.$$

In particular, every arc used by *both* routes of i contributes nothing to $\boldsymbol{\Delta}_i$, although it does contribute to the absolute load x_a. Moreover, if $\kappa_i^r = \mathbf{c}^\mathsf{T}\boldsymbol{\chi}(P_i^r)$, then $\mathbf{c}^\mathsf{T}(\mathbf{y}-\mathbf{x}) = \sum_i d_i(\kappa_i^1 - \kappa_i^0)(z_i - f_i)$, so the cost constraint is a single signed linear functional of $\big(d_i(z_i - f_i)\big)_i$.

3.2 Common-point interval systems and two chains

Definition 3.3 (common-point interval system)**.** Let $v_0 = s \to v_1 \to \cdots \to v_N$ be a directed spine, whose arcs are indexed $0, \ldots, N-1$. For each terminal i fix an interval $I_i = [\ell_i, r_i]$ of spine arcs and add two private exit arcs: $v_{\ell_i} t_i$ (*early*) and $v_{r_i+1} t_i$ (*late*); there are no other arcs. The system is a *common-point* system if $\bigcap_i I_i \neq \varnothing$.

Figure 2 shows the shape of such an instance. Exactly as in [14], every path $s \to t_i$ is a spine prefix followed by one exit arc, so each terminal has exactly *two* simple paths, and the analysis cannot be invalidated by unintended path switching. Switching terminal i from early to late adds d_i on precisely the arcs of I_i, so by Observation 3.2 the overload of spine arc e under a routing $\mathbf{z} \in \{0,1\}^k$ is

$$y_e - x_e = \sum_{i \in R_e} d_i(z_i - f_i), \qquad R_e = \{i : \ e \in I_i\}, \tag{1}$$

while the overload of an exit arc of terminal i is $d_i(1-f_i)$ if $z_i = 1$ and $d_i f_i$ if $z_i = 0$ — in both cases at most $d_i \leq d_{\max}$.

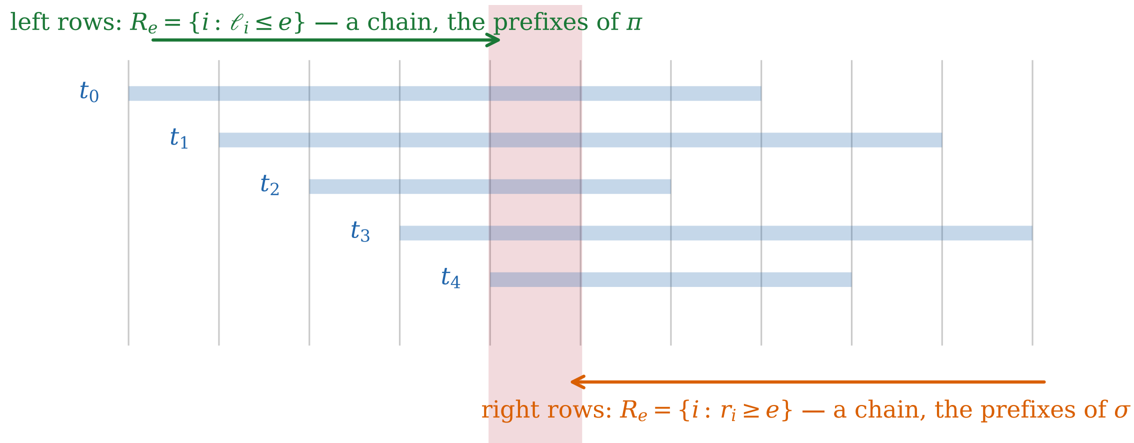

Figure 3: Lemma 3.4. Every interval contains the common arc, so the row supports left of it grow monotonically and those right of it shrink monotonically: the shared rows of a common-point system are the prefixes of two orders of the terminals.

Lemma 3.4 (two chains)**.** *Let a common-point interval system have common arc e^*. Then $\{R_e : e \le e^*\}$ is a chain, increasing in e, and $\{R_e : e \ge e^*\}$ is a chain, decreasing in e. Consequently there are orders π, σ of $[k]$ such that every shared row is a prefix of π or a prefix of σ:*

$$\{R_e\}_e \ \subseteq\ \mathrm{Pref}(\pi) \cup \mathrm{Pref}(\sigma).$$

Proof. Fix $e \le e^*$. Since $e^* \in I_i$ for every i, we have $r_i \ge e^* \ge e$, so $i \in R_e$ if and only if $\ell_i \le e$; hence $R_e = \{i : \ell_i \le e\}$, which is nondecreasing in e. Let π order the terminals by nondecreasing ℓ_i (ties broken arbitrarily); then every R_e with $e \le e^*$ is a prefix of π. Symmetrically, for $e \ge e^*$ we have $\ell_i \le e^* \le e$, so $R_e = \{i : r_i \ge e\}$, nonincreasing in e; ordering by nonincreasing r_i gives σ. $\square$

Figure 3 illustrates the lemma. Thus a common-point system is exactly a *weighted two-permutation prefix system*: the shared constraints are the prefixes of two orders, weighted by the demands, and the cost constraint is one preserved linear functional (Observation 3.2).

3.3 Cost faces and rank

Give the early exit of terminal i per-unit cost $1/d_i$ and every other arc cost zero, so that every full early path costs exactly 1 and every late path costs 0. Then

$$\mathbf{c}^\mathsf{T}\mathbf{x} = \sum_i (1 - f_i) = k - \textstyle\sum_i f_i, \qquad \mathbf{c}^\mathsf{T}\mathbf{y} = k - \textstyle\sum_i z_i,$$

so a routing is cost-good precisely when $\sum_i z_i \ge \sum_i f_i$. Writing

$$\sum_i f_i = m - 1 + \eta, \qquad 0 < \eta \le 1, \tag{2}$$

the cost-good routings are exactly those that have $\sum_i z_i \ge m$. We call the value $\sum_i z_i$ the *rank* of a routing, and we call m the *first cost-good rank* of the cell and $q = k - m + 1$ its *complement mass*; note that $\sum_i (1 - f_i) = q - \eta$. Throughout, a *cell* means one instance of the class together with its fractional point, i.e. a choice of the two orders π, σ, the demands $\mathbf{d}$ and the shares $\mathbf{f}$; this determines the cost face and with it m, η and q. Because all shared rows in (1) have nonnegative coefficients, adding a late choice can only increase every shared overload, so it suffices to analyse routings of rank exactly m.

The complement mass already gives a rough upper bound.

Proposition 3.5 (trivial upper bound)**.** *Every common-point cell satisfies $C^\star(\mathbf{x}) \le \max(1,\, q-\eta)$. In particular a cell of complement mass $q=1$ has $C^\star \le 1$, and a cell can force a constant above γ only if $q > \gamma$.*

Proof. The all-late routing $\mathbf{z} = \mathbf{1}$ has rank $k \ge m$, so it is cost-good. Its overload on a shared row is $\sum_{i\in R_e} d_i(1-f_i) \le \sum_{i=1}^{k} d_i(1-f_i) \le \sum_i (1-f_i) = q-\eta$, using $d_i \le d_{\max} = 1$, and its overload on an exit arc is at most $d_i \le 1$. □

Example 3.6 (the counterexample from [14])**.** The four-interval family that we introduced in [14] is the $k=4$ member of this class. Its intervals are

$$I_0 = [0,2], \quad I_1 = [0,4], \quad I_2 = [1,4], \quad I_3 = [2,3],$$

which all contain arc 2, so it is a common-point system with $e^* = 2$; Figure 4 draws the digraph. In the normalization of [14],

$$\mathbf{d} = (1, (1-t)^2, 1-t, 1), \qquad \mathbf{f} = (2t-t^2,\, 1-4t+t^2,\, t,\, t+\varepsilon), \qquad \sum_i f_i = 1+\varepsilon.$$

Hence by (2) we get $m-1+\eta = 1+\varepsilon$, which means that $m=2$, $\eta=\varepsilon$, and $q = k-m+1 = 3$: the first cost-good rank is two (a routing is cost-good exactly when at least two terminals take their late path) and the complement mass is three. Proposition 3.5 gives the ceiling $3-\varepsilon$, and the true value is $C(t)-\varepsilon$, at most $(299-41\sqrt{41})/32 = 1.1397\ldots$.

The instance has $k=4$ terminals but a spine of five arcs, so there are four intervals and five shared rows: the endpoints are $(\ell_i) = (0,0,1,2)$ and $(r_i) = (2,4,4,3)$, while $R_e = \{i : e \in I_i\}$ runs over $e = 0,\ldots,4$. Explicitly,

$$R_0 = \{0,1\}, \quad R_1 = \{0,1,2\}, \quad R_2 = \{0,1,2,3\}, \quad R_3 = \{1,2,3\}, \quad R_4 = \{1,2\},$$

which is a direct illustration of Lemma 3.4. The left chain is the set of rows with $e \le e^* = 2$: sorting the terminals by nondecreasing ℓ_i gives $\pi = (0,1,2,3)$, and $R_0 \subset R_1 \subset R_2$ are its prefixes of sizes $2,3,4$. The right chain is the set of rows with $e \ge e^*$: sorting by nonincreasing r_i gives $\sigma = (1,2,3,0)$, and $R_4 \subset R_3 \subset R_2$ are its prefixes of sizes $2,3,4$. The two chains have three rows each and share the common arc's row R_2, which is how five shared rows arise from two orders of four terminals, so the whole of Section 6 applies to this instance.

3.4 The hull characterization and the lower box

Two general facts about acyclic instances are used throughout. Both proofs are short, so we give them here to keep this paper self-contained.

The first removes the cost vector from the definition of the critical constant.

Proposition 3.7 (hull characterization)**.** *Let G be acyclic, $\mathbf{x}$ feasible, and for $C \ge 0$ let $U_C(\mathbf{x}) = \{$unsplittable $\mathbf{y} :\ y_a \le x_a + C d_{\max}\ \forall a\}$. Then the following are equivalent:*

(i) *for every $\mathbf{c} \ge 0$ some $\mathbf{y} \in U_C(\mathbf{x})$ has $\mathbf{c}^\top \mathbf{y} \le \mathbf{c}^\top \mathbf{x}$;*

(ii) *$\mathbf{x} \in \operatorname{conv} U_C(\mathbf{x})$;*

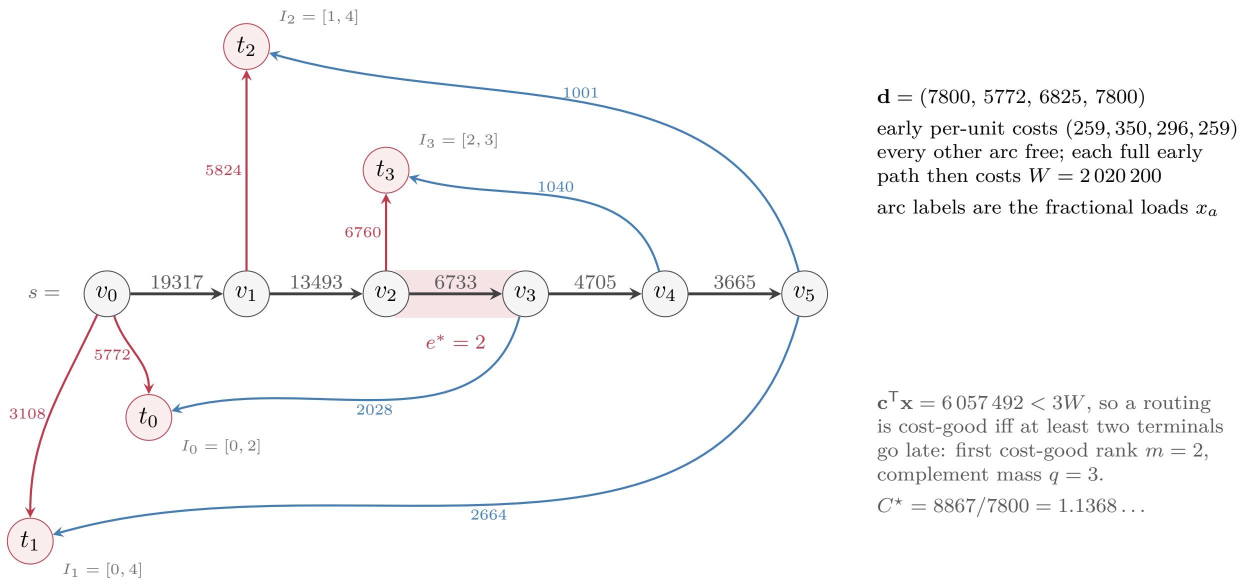

Figure 4: Example 3.6: the four-interval instance of [14] drawn as a digraph, at the exact rational point with $C^\star = 8867/7800$. Spine labels are the fractional arc loads x_a; exit labels are the amounts $d_i(1 - f_i)$ (early, in red) and $d_i f_i$ (late, in blue). All four intervals contain arc $e^* = 2$, so this is a common-point system; it is the $k = 4$ member of the class studied here.

(iii) there are convex coefficients $\lambda_\mathbf{y}$ with $\sum_\mathbf{y} \lambda_\mathbf{y}\mathbf{y} \le \mathbf{x}$.

Proof. (ii)⇒(i) by averaging and (ii)⇒(iii) trivially. For (iii)⇒(ii), let $\mathbf{q} = \sum_\mathbf{y} \lambda_\mathbf{y}\mathbf{y} \le \mathbf{x}$; every unsplittable load vector and $\mathbf{x}$ have the same divergence, so $\mathbf{h} = \mathbf{x} - \mathbf{q} \ge \mathbf{0}$ is a nonnegative circulation, and an acyclic digraph has none except $\mathbf{0}$. If (iii) fails, strong separation of $\operatorname{conv} U_C(\mathbf{x}) + \mathbb{R}^A_{\ge 0}$ from $\mathbf{x}$ produces $\mathbf{c} \ge 0$ with $\mathbf{c}^\mathsf{T}\mathbf{y} > \mathbf{c}^\mathsf{T}\mathbf{x}$ for all $\mathbf{y} \in U_C(\mathbf{x})$, so (i) fails. □

The second is the structural statement whose validity would settle the cost-preserving question at the value 2.

Fix an instance, a feasible $\mathbf{x}$, and costs $\mathbf{c} \ge 0$, and let F denote the fractional-routing polytope. Consider the capacitated minimum-cost routing

$$\mathbf{q} \in \arg\min\{\mathbf{c}^\mathsf{T}\mathbf{p} :\ \mathbf{p} \in F,\ \mathbf{p} \le \mathbf{x} + d_{\max}\mathbf{1}\},$$

which exists because $\mathbf{x}$ itself is feasible, and let $\boldsymbol{\lambda} \ge 0$ be optimal dual multipliers for the arc upper bounds. Complementary slackness gives $\boldsymbol{\lambda}^\mathsf{T}\mathbf{q} = \boldsymbol{\lambda}^\mathsf{T}\mathbf{x} + d_{\max}\|\boldsymbol{\lambda}\|_1$, and since $\mathbf{q}$ minimises $(\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{p}$ over F while $\mathbf{x} \in F$,

$$\mathbf{q} \in \arg\min_{\mathbf{p} \in F}(\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{p}, \qquad \mathbf{c}^\mathsf{T}\mathbf{q} \ \le\ \mathbf{c}^\mathsf{T}\mathbf{x} - d_{\max}\|\boldsymbol{\lambda}\|_1. \tag{3}$$

Since D is acyclic, $\mathbf{q}$ has no positive-flow cycles, and there is a node potential ϕ whose reduced costs $(\mathbf{c} + \boldsymbol{\lambda})_{uv} + \phi_u - \phi_v$ are nonnegative on all arcs and zero on the support of $\mathbf{q}$.

Write $G[\mathbf{q} > 0]$ for the subgraph of arcs carrying positive $\mathbf{q}$-flow. A *support path* for terminal i is a simple path $s \to t_i$ all of whose arcs lie in $G[\mathbf{q} > 0]$, and a *support routing* is a routing that assigns to every terminal a support path; as in Section 3, its load vector is $\mathbf{y} = \sum_i d_i \boldsymbol{\chi}(P_i)$. Support paths exist for every terminal, since $\mathbf{q}$ decomposes into $s \to t_i$ paths carrying d_i in total, and each of those lies in $G[\mathbf{q} > 0]$.

Because the reduced costs vanish on $G[\mathbf{q} > 0]$, every support path for terminal i is $(\mathbf{c}+\boldsymbol{\lambda})$-shortest, of length $\phi_{t_i} - \phi_s$; hence every support routing has the same $(\mathbf{c}+\boldsymbol{\lambda})$-cost $\sum_i d_i(\phi_{t_i} - \phi_s) = (\mathbf{c}+\boldsymbol{\lambda})^\mathsf{T}\mathbf{q}$ as $\mathbf{q}$ itself. We denote

$$U(\mathbf{q}) = \{\mathbf{y} : \ \mathbf{y} \text{ is a support routing, } \ \mathbf{y} \le \mathbf{q} + d_{\max}\mathbf{1}\}$$

and call its members the *upper-good* support routings. This set is nonempty: applying the theorem of Dinitz, Garg and Goemans [3] to $\mathbf{q}$ on $G[\mathbf{q} > 0]$ yields an unsplittable $\mathbf{y} \le \mathbf{q} + d_{\max}\mathbf{1}$ whose paths lie in the support.

Definition 3.8 ((**LB**), the lower box)**.** Statement (LB): for every instance, every feasible $\mathbf{x}$ and every $\mathbf{c} \ge 0$, with $\mathbf{q}$ as above,

$$\operatorname{conv} U(\mathbf{q}) \ \cap \ \{\mathbf{m} : \ \mathbf{m} \ge \mathbf{q} - d_{\max}\mathbf{1}\} \ \neq \ \varnothing.$$

Lemma 3.9 (separation form)**.** (LB) *holds for a given* $\mathbf{q}$ *if and only if*

$$\max_{\mathbf{y} \in U(\mathbf{q})} \mathbf{w}^\mathsf{T}\mathbf{y} \ \ge \ \mathbf{w}^\mathsf{T}\mathbf{q} - d_{\max}\|\mathbf{w}\|_1 \qquad \textit{for every } \mathbf{w} \ge 0.$$

Proof. The second set in Definition 3.8 is a translate of the nonnegative orthant, so it is unbounded in every positive coordinate direction and any separating normal $\mathbf{w}$ must be nonnegative. Its minimum over that set is $\mathbf{w}^\mathsf{T}(\mathbf{q} - d_{\max}\mathbf{1})$, while its maximum over $\operatorname{conv} U(\mathbf{q})$ is attained at a member of $U(\mathbf{q})$. The two convex sets are disjoint precisely when some such $\mathbf{w}$ separates them strictly, which is the negation of the display. □

Proposition 3.10. (LB) *implies the universal cost-preserving additive constant* 2*: for every instance and every feasible* $\mathbf{x}$ *there is an unsplittable* $\mathbf{y}$ *with* $\mathbf{y} \le \mathbf{x} + 2d_{\max}\mathbf{1}$ *and* $\mathbf{c}^\mathsf{T}\mathbf{y} \le \mathbf{c}^\mathsf{T}\mathbf{x}$.

Proof. Apply Lemma 3.9 with $\mathbf{w} = \boldsymbol{\lambda}$: some $\mathbf{y} \in U(\mathbf{q})$ satisfies $\boldsymbol{\lambda}^\mathsf{T}(\mathbf{q} - \mathbf{y}) \le d_{\max}\|\boldsymbol{\lambda}\|_1$. Support optimality gives $(\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{y} = (\mathbf{c} + \boldsymbol{\lambda})^\mathsf{T}\mathbf{q}$, so

$$\mathbf{c}^\mathsf{T}\mathbf{y} = \mathbf{c}^\mathsf{T}\mathbf{q} + \boldsymbol{\lambda}^\mathsf{T}(\mathbf{q} - \mathbf{y}) \ \le \ \mathbf{c}^\mathsf{T}\mathbf{q} + d_{\max}\|\boldsymbol{\lambda}\|_1 \ \le \ \mathbf{c}^\mathsf{T}\mathbf{x}$$

by (3), while $\mathbf{y} \le \mathbf{q} + d_{\max}\mathbf{1} \le \mathbf{x} + 2d_{\max}\mathbf{1}$. □

Remark 3.11. Proposition 3.10 is close in spirit to Theorem 1.6 of [19], which derives a cost-preserving $2d_{\max}$ guarantee from the two-sided cost-free conjecture of Morell and Skutella [12]. What (LB) isolates is the apparently minimal structural core: only the *lower* box is needed, only relative to the capacitated optimum $\mathbf{q}$, and only over support routings.

4 The seventeen-terminal record

In this section, we present our currently best counterexample to the Goemans conjecture.

Theorem 4.1. *Every universal cost-preserving additive constant satisfies*

$$C \ \ge \ \frac{1282494797984843521}{10^{18}} = 1.282494797984843521\ldots \ > \ \frac{6}{5}.$$

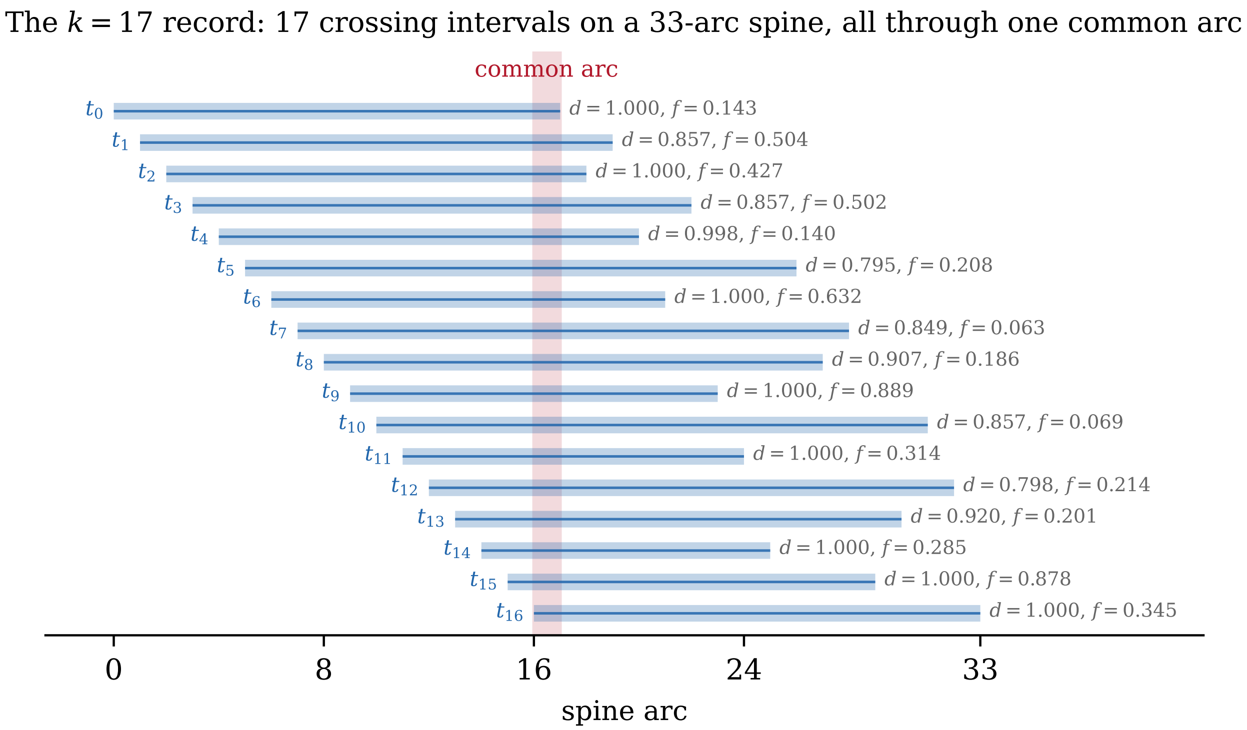

Figure 5: The $k = 17$ record instance, as an interval system. Seventeen mutually crossing intervals on a 33-arc spine, all through one common arc; the first cost-good rank is seven. Complete data in Appendix A.

The witness is the $k = 17$ common-point interval instance of Appendix A, drawn as an interval system in Figure 5 and as a raw digraph in Figure 6. Its intervals are $I_i = [i, r_i]$ with

$$(r_0, \dots, r_{16}) = (16, 18, 17, 21, 19, 25, 20, 27, 26, 22, 30, 23, 31, 29, 24, 28, 32),$$

so the spine has 33 arcs, all intervals contain arc 16, and the raw digraph has $33 + 34 = 67$ arcs. The demands and shares are rationals with denominator 10^9; they satisfy $\max_i d_i = 1$ and

$$\sum_{i=0}^{16} f_i = 6 + 10^{-9},$$

so by (2) the first cost-good rank is $m = 7$ and the complement mass is $q = 11$. In the notation of Lemma 3.4, π is the identity and σ is the order by nonincreasing r_i, namely

$$\sigma = (16, 12, 10, 13, 15, 7, 8, 5, 14, 11, 9, 3, 6, 4, 1, 2, 0).$$

Proof of Theorem 4.1. The instance is finite and rational, so the claim is a finite computation. From the raw arc list, depth-first search finds exactly two simple paths $s \to t_i$ for every terminal, and the intervals recovered from those paths are the ones displayed. We have enumerated all $2^{17} = 131072$ routings in exact integer arithmetic: multiplying by 10^{18} turns every overload (1) into an integer. Of these, 109294 are cost-good, i.e. of rank at least 7; their minimum all-arc overload is exactly $1282494797984843521/10^{18}$, attained by 15 routings. Since $\max_i d_i = 1 < 1.28$, no private exit arc can be the maximizer, so the binding constraints are spine rows everywhere. Finally, 8885 routings have strictly smaller overload and all of them have rank at most 6; hence the value is not an artifact of the rank threshold. By Definition 3.1, $C^\star(\mathbf{x})$ equals the displayed constant. $\square$

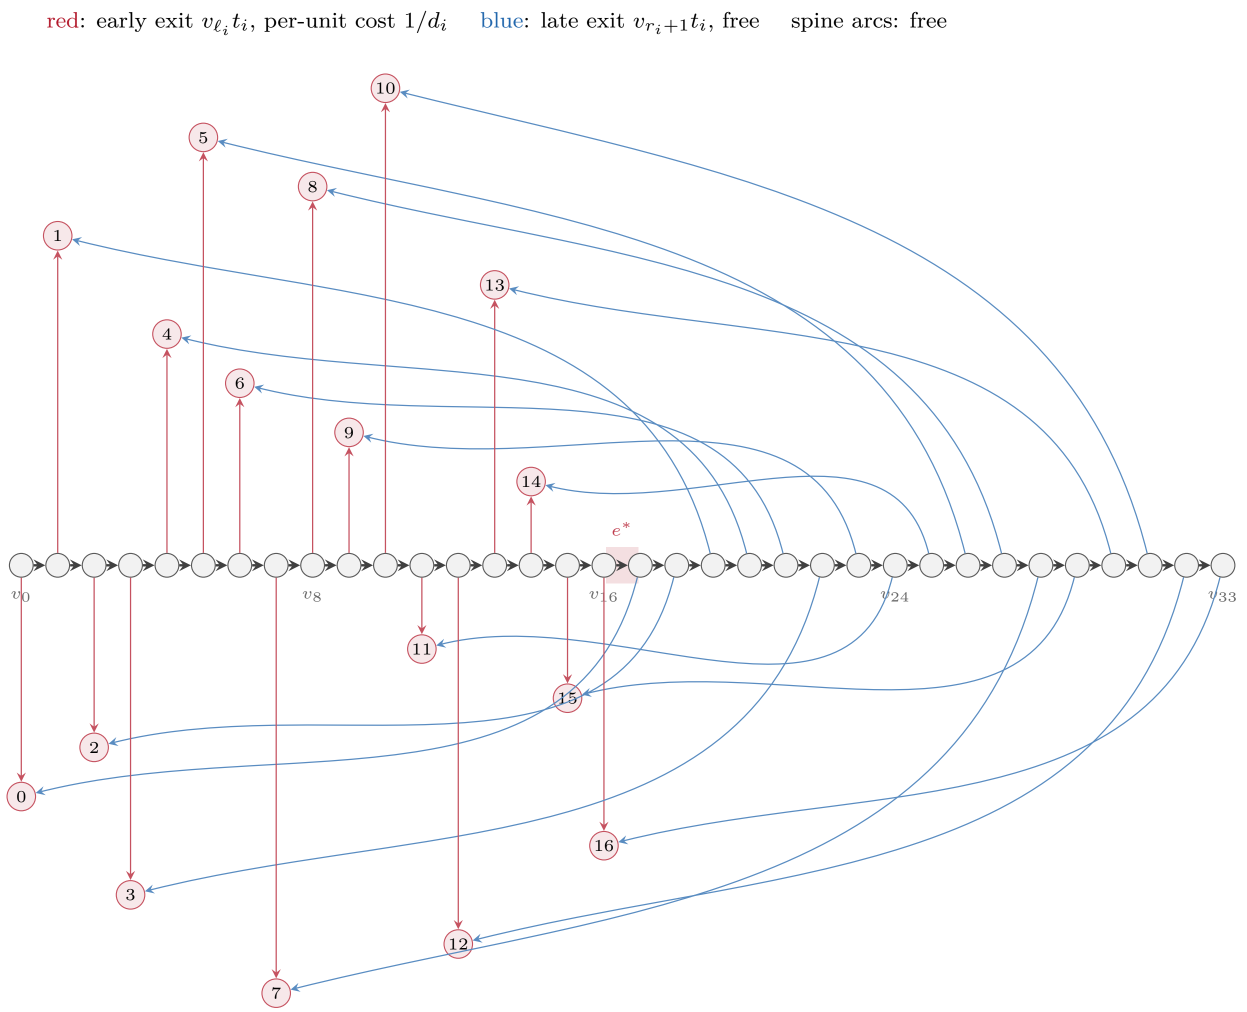

Figure 6: The instance from Fig. 5 shown as a digraph: 34 spine vertices, 17 terminals, 67 arcs. Terminal i is drawn at its left endpoint $\ell_i = i$; its red arc is the early exit $v_i t_i$ and its blue arc the late exit $v_{r_i+1} t_i$; the shaded spine arc is the common arc $e^* = 16$; demands and shares are in Appendix A.

Two further certificates strengthen the statement.

Proposition 4.2 (the value is the all-cost critical constant)**.** *There are* 18 *routings* $\mathbf{z}^{(1)}, \dots, \mathbf{z}^{(18)}$, *each of overload at most* $C^\star$, *and positive rationals* $\lambda_1, \dots, \lambda_{18}$ *summing to* 1 *with* $\sum_j \lambda_j \mathbf{z}^{(j)} = \mathbf{f}$.

By the hull characterization (Proposition 3.7), Proposition 4.2 means that $C^\star$ is not merely the constant forced by *this* cost vector: it is exactly the first overload level at which the retained routings contain the fractional point in their convex hull, hence the largest constant forced by *any* nonnegative cost vector on this fractional flow.

Proposition 4.3 (the lower box holds on the record)**.** *For the* $k = 17$ *instance, with* $U(\mathbf{q})$ *the set of upper-good support routings,*

$$\delta(\mathbf{q}) := \min_{\mathbf{m} \in \operatorname{conv} U(\mathbf{q})} \max_a (q_a - m_a) = \max_{\mathbf{w} \geq 0,\ \|\mathbf{w}\|_1 = 1} \Big(\mathbf{w}^\top \mathbf{q} - \max_{\mathbf{y} \in U(\mathbf{q})} \mathbf{w}^\top \mathbf{y} \Big)$$
$$= 6.3627205289 \cdot 10^{-10} \; < \; d_{\max}.$$

There are 4799 upper-good routings; a 17-atom primal mixture and a 14-arc dual weight vector both give the displayed value, so (LB) — whose validity would imply the universal cost-preserving constant 2 (Proposition 3.10) — holds here with a very large margin. The same is true of the earlier $k = 11, 13, 15$ records, whose exact deficits are $7.18 \cdot 10^{-10}$, $7.18 \cdot 10^{-10}$ and $7.09 \cdot 10^{-10}$. The strict-rank and lower-box obstructions are thus nearly orthogonal on every sharp instance.

Remark 4.4 (how the record was found). The instance was found by joint optimization in the coordinates $(d_i,\, p_i = d_i f_i)$, in which every fixed-witness overload inequality $\sum_{i \in P}(d_i z_i - p_i) \geq C$ is linear and the only nonlinearity is the cost equation $\sum_i p_i/d_i = m - 1 + \eta$. The record is *not* claimed to be a stationary cell: the full multiplier LP is infeasible at tolerances through 10^{-5}. One-column continuations ($k = 18$, all 18^2 insertion positions at ranks seven and eight, quotiented by boundary-share deletion) and a bounded two-column screen (1200 sampled $k = 19$ extensions) do not improve it; this is a finite search and it does not prove optimality.

Combining with [14] and intermediate families that we found on the way to this record, we get the following ladder of certified counterexamples:

k	4	5	6	7	8	17
exact supremum	$\dfrac{299 - 41\sqrt{41}}{32}$	T_*	$11 - 4\sqrt{6}$	$\dfrac{13}{2} - 2\sqrt{7}$	≈ 1.2148	$\dfrac{1282494797984843521}{10^{18}}$
decimal	1.13975	1.16798	1.20204	1.20850	1.21480	1.28249

where $T_* = 1.1679802157706246\ldots$ is a root of the polynomial

$$6T^6 - 94T^5 - 1229T^4 - 3056T^3 + 5380T^2 + 4704T - 5488 = 0.$$

The $k = 6$ and $k = 7$ entries are limits of Pell families: with $(5 + 2\sqrt{6})^n = p_n + q_n\sqrt{6}$ the six-interval family has exact constants $(p_n - 4q_n)^2/(2q_n^2) \uparrow 11 - 4\sqrt{6}$, and with $(8 + 3\sqrt{7})^n = p_n + Q_n\sqrt{7}$ the seven-interval family has $(p_n - 8Q_n)(p_n - 4Q_n)/(6Q_n^2) \uparrow \frac{13}{2} - 2\sqrt{7}$. Every entry is realized by a common-point interval system with two simple paths per terminal and demands that are not a divisibility chain. Appendix C draws the $k = 5, 6, 7, 8$ instances (Figures 11–14), each both as a raw digraph and as a family of intervals; note how they have a single common arc in every case, and intervals that cross rather than nest.

We close this section with a remark on two further finite experiments, both negative, whose details are deferred to Appendix D.

The first asks whether the record can be *squeezed* rather than merely attained. For each certified record one can relax the problem by allowing, in addition to every routing of overload at most $d_{\max}$, all of its one-bit upward perturbations; the first level at which the fractional point enters the convex hull of that larger vertex set is then an upper relaxation of the critical constant. On all seven records in the chain $k = 7, 8, 11, 13, 15, 16, 17$ this relaxation is *exactly equal* to the record itself, so the two-sided bracket has width zero. That is a strong internal consistency check, but it does not bound the class supremum: the upper endpoint is computed at the very same fractional point as the lower one, so it says nothing about neighbouring cells.

The second asks whether the $k = 17$ geometry can simply be scaled up. The most literal microscopic ansatz is to clone coordinates of the record and interleave the clones in the second order. Over a scheduled family of 69 endpoints with k up to 40, every single endpoint turns out to admit an explicit cost-good routing whose exact overload is *below* the $k = 17$ record. Naive cloning therefore does not produce a limit above 1.2825, and in particular gives no support for a limiting value of 4/3. This exhausts that specific schedule, not all cloning schemes.

5 The atomless continuum has value zero

It is tempting to look for the class supremum as the limit of a macroscopic object, in this case a permuton (a limit of permutation matrices, here carrying a density of terminals) with two order coordinates. The following theorem says that this cannot work because the limit problem is trivial.

Theorem 5.1. *Let $(\Omega, \mathcal{A}, \nu)$ be a finite atomless measure space, let $\ell, r : \Omega \to \mathbb{R}$ be measurable (the two order coordinates), let $d, f : \Omega \to [0,1]$ be measurable, and let $w \in L^1(\nu)$. Then for every $\varepsilon > 0$ there is a measurable set $A \subseteq \Omega$ with*

$$\int_A w\, d\nu = \int_\Omega wf\, d\nu,$$

$$\sup_{t\in\mathbb{R}} \Big| \int_{\{\ell \le t\}} (\mathbf{1}_A - f)\, d\, d\nu \Big| < \varepsilon, \qquad \sup_{t\in\mathbb{R}} \Big| \int_{\{r \le t\}} (\mathbf{1}_A - f)\, d\, d\nu \Big| < \varepsilon.$$

Consequently the atomless two-order minimax with one preserved linear functional is identically zero.

Proof. Let ρ be the finite measure $d\, d\nu$ on Ω, and let $\mu_\ell = \ell_* \rho$ and $\mu_r = r_* \rho$ be its pushforwards to $\mathbb{R}$.

Step 1: a good grid. We claim there are finite grids $s_1 < \cdots < s_N$ and $t_1 < \cdots < t_M$ such that every set of the form $\{s_j < \ell \le s_{j+1}\}$, $\{\ell \le s_1\}$ or $\{\ell > s_N\}$ has ρ-measure $< \varepsilon$, *after* adjoining to the grid the finitely many levels that carry a large atom. To be precise, μ_ℓ has at most $\lceil \mu_\ell(\mathbb{R})/\varepsilon \rceil$ atoms of mass $\ge \varepsilon$; for each such atom at a level c we adjoin *two* constraints, using the sets $\{\ell < c\}$ and $\{\ell \le c\}$ instead of a single grid level. The remaining mass of μ_ℓ is carried by an interval decomposition in which no single point has mass $\ge \varepsilon$, so a greedy left-to-right accumulation cuts it into finitely many pieces of mass $< \varepsilon$ each. Then we can do the same for μ_r. If ℓ and r have no large level sets, for instance if ρ is nonatomic along both coordinates, this is just the usual fine grid.

Step 2: purification. Consider the finite-dimensional vector measure

$$m = \big(m_0,\ m_1,\ m_2^{(1)}, \dots, m_2^{(N')},\ m_3^{(1)}, \dots, m_3^{(M')}\big)$$

on $(\Omega, \mathcal{A})$ with

$$m_0(E) = \nu(E), \quad m_1(E) = \int_E w\, d\nu, \quad m_2^{(j)}(E) = \rho(E \cap L_j), \quad m_3^{(j)}(E) = \rho(E \cap R_j),$$

where L_j and R_j run over the finitely many prefix sets selected in Step 1. Each coordinate is absolutely continuous with respect to ν, which is atomless, so m is an atomless finite-dimensional vector measure. By Lyapunov's convexity theorem [8, 6] its range $\{m(E) : E \in \mathcal{A}\}$ is convex and compact, and it coincides with $\{\int_\Omega g\, dm :\ 0 \le g \le 1\}$. Since $0 \le f \le 1$, the point $\int_\Omega f\, dm$ lies in that set, so there is a measurable A with

$$m(A) = \int_\Omega f\, dm.$$

Coordinate m_1 gives $\int_A w\, d\nu = \int wf\, d\nu$; coordinates m_2 and m_3 give $\int_{L_j} (\mathbf{1}_A - f)\, d\rho = 0$ and $\int_{R_j} (\mathbf{1}_A - f)\, d\rho = 0$ for every selected prefix.

Step 3: arbitrary levels. Fix $t \in \mathbb{R}$ and let L_j be the selected prefix immediately below $\{\ell \le t\}$, so that $\{\ell \le t\} \setminus L_j$ is contained in one of the slabs of Step 1 and therefore has ρ-measure $< \varepsilon$. Then

$$\Big| \int_{\{\ell \le t\}} (\mathbf{1}_A - f)\, d\rho \Big| = \Big| \underbrace{\int_{L_j} (\mathbf{1}_A - f)\, d\rho}_{=0} + \int_{\{\ell \le t\} \setminus L_j} (\mathbf{1}_A - f)\, d\rho \Big| \le \rho(\{\ell \le t\} \setminus L_j) < \varepsilon,$$

using the fact that $|\mathbf{1}_A - f| \le 1$. The same holds for r. $\square$

Remark 5.2. Step 1 is not a formality. If ℓ is constant on a set of positive ρ-measure, no grid can make every slab light, and the naive argument fails; the repair is to constrain the two prefixes $\{\ell < c\}$ and $\{\ell \le c\}$ around each heavy level exactly, which is possible because Lyapunov's theorem accommodates any finite number of linear constraints.

Theorem 5.1 says that the supremum C_{cls} of the common-point class is indeed an *atomic* quantity: at k atoms the normalized discrepancy is C_k/k and it tends to zero, so the informative quantity is C_k itself, a first-order term. A macroscopic limit object without a microscopic recovery rule cannot determine it. In particular, an empirical fit of the shape $4/3 - \alpha/k$, which one can easily notice in the ladder of counterexamples, has no continuum justification; the rigorous statement proven in Section 4 is $C_{\text{cls}} \ge 1.282494797984843521\ldots$.

6 Deletion stars and the first unconditional upper bounds

Everything so far has been about *lower* bounds: explicit instances that force a universal constant to be large. We now turn the question around and ask how large the constant can possibly be *on this class*. Recall what the class is. By Lemma 3.4, a common-point interval system is completely described by two orders π, σ of the terminals, demands $d_i \in (0, d_{\max}]$ and shares $\mathbf{f} \in [0,1]^k$: the shared arcs are the prefixes of π and of σ, weighted by the demands, and the cost constraint says that a routing must take at least m late choices, where $\sum_i f_i = m - 1 + \eta$ by (2). Normalize $d_{\max} = 1$ from now on. Proving an upper bound for the class therefore means: exhibit *one* rank-m vector $\mathbf{z} \in \{0,1\}^k$ whose weighted error $\sum_{i \in P} d_i(z_i - f_i)$ is small on every prefix P of both orders.

The obstacle is that this error can be large for a trivial reason. Take $\mathbf{z} = \mathbf{1}_S$ for an m-set S; on a prefix P the error is $\sum_{i \in P \cap S} d_i - \sum_{i \in P} d_i f_i$, and the first sum has no reason to be matched by the second. The device that fixes this is to *give one unit of demand back*, spread over the support: instead of measuring $\mathbf{z}$ itself we measure $\mathbf{z} - \mathbf{c}$ for a probability vector $\mathbf{c}$ carried by S. One unit is exactly the right amount, because a routing of rank m overshoots the fractional rank $m - 1 + \eta$ by just under one.

Definition 6.1 (deletion star)**.** Let π, σ be orders of $[k]$, let $0 \le d_i \le 1$, and let $\mathbf{f} \in [0,1]^k$ with $\sum_i f_i = m - 1$ (the limiting cost face). For an m-set $S \subseteq [k]$ and *deletion credits* $c_i \ge 0$ supported on S with $\sum_i c_i = 1$, put $\mathbf{z} = \mathbf{1}_S$ and $\mathbf{u} = \mathbf{z} - \mathbf{c}$. The *deletion-star residual* is

$$B(\pi, \sigma, \mathbf{d}, \mathbf{f}) \;=\; \min_{S, \mathbf{c}} \; \max_{P \in \mathrm{Pref}(\pi) \cup \mathrm{Pref}(\sigma)} \; \sum_{i \in P} d_i(u_i - f_i).$$

Note $\sum_i u_i = m - 1 = \sum_i f_i$, so the residual $\mathbf{u} - \mathbf{f}$ has total weight zero in the unweighted sense; the credits are exactly what makes this possible.

Lemma 6.2 (support reduction)**.** *If $B(\pi, \sigma, \mathbf{d}, \mathbf{f}) \le \beta$, then the rank-$m$ routing $\mathbf{z} = \mathbf{1}_S$ attaining the minimum has overload at most $1 + \beta$ on every shared row and at most 1 on every private exit arc.*

Proof. Since $\mathbf{z} = \mathbf{u} + \mathbf{c}$, for every prefix P,

$$\sum_{i \in P} d_i(z_i - f_i) = \underbrace{\sum_{i \in P} d_i(u_i - f_i)}_{\le \beta} + \sum_{i \in P} d_i c_i \;\le\; \beta + \max_i d_i \sum_i c_i = \beta + 1.$$

Every shared row is such a prefix by Lemma 3.4. The exit-arc bound was noted after (1). $\square$

The two definitions are easiest to see on a small instance, which we will reuse as a running example.

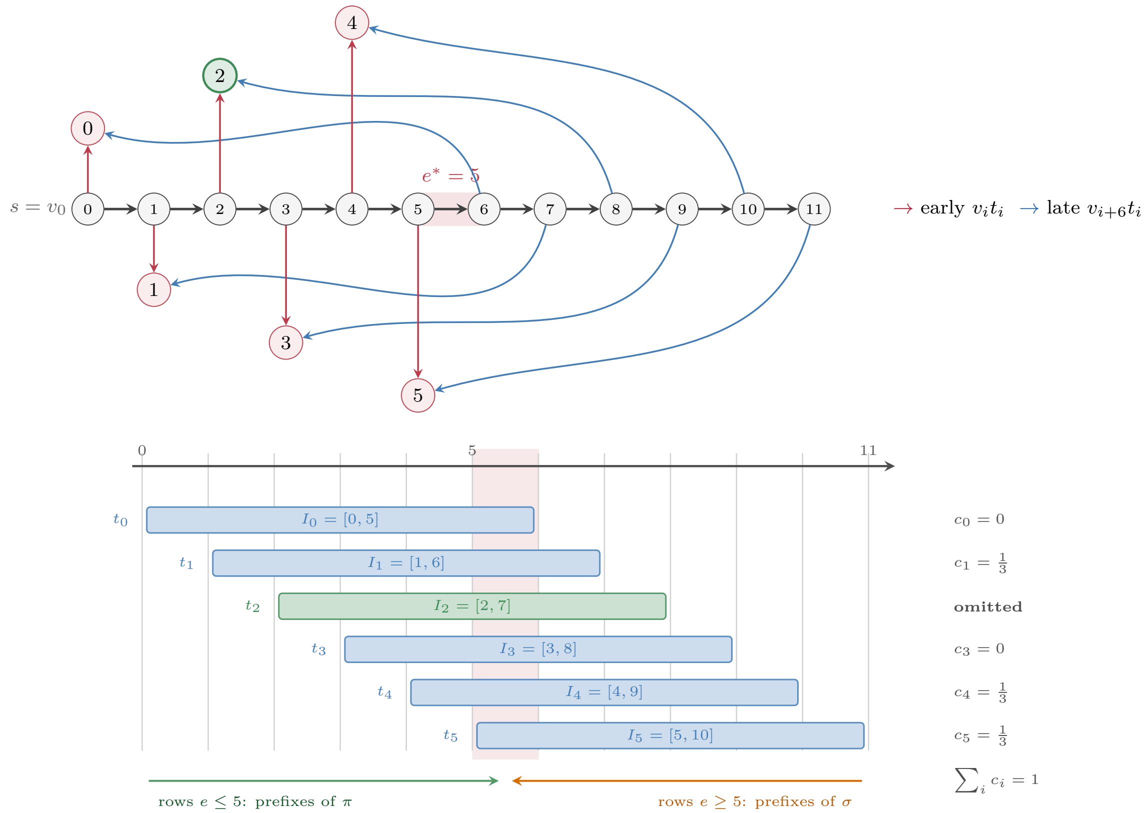

Figure 7: The instance of Example 6.3, as a digraph (top) and as an interval system (bottom). Six terminals of unit demand on a spine of eleven arcs, with $I_i = [i, i+5]$; every interval contains the common arc $e^* = 5$, so by Lemma 3.4 the eleven shared rows split into the six prefixes of π (rows $e \le 5$, left arrow) and the six prefixes of σ (rows $e \ge 5$, right arrow), overlapping in R_5. Green: the coordinate t_2 omitted by the deletion star. The right-hand column lists the deletion credits, which sum to one and are spread over the five selected coordinates.

Example 6.3 (a deletion star, explicitly)**.** Take $k = 6$ terminals with unit demands, π the identity order, σ its reverse, and $f_i = \frac{2}{3}$ for every i, so that $\sum_i f_i = 4 = k - 2$. This is the limiting codimension-two face of Definition 6.1, on which $m = \sum_i f_i + 1 = 5$ and the star uses a 5-set; it is the boundary of the slab $k - 2 \le \sum_i f_i \le k - 1$ of Corollary 6.8, whose cells have complement mass $q = 2$ and need at least five of the six terminals late to be cost-good.

Concretely the two orders are realized by the staircase interval system $I_i = [i, i+5]$ on a spine of eleven arcs, drawn in Figure 7: the six rows left of the common arc $e^* = 5$ are the prefixes $\{0\}, \{0, 1\}, \ldots, \{0, \ldots, 5\}$ of π, the six rows to its right are the prefixes $\{5\}, \{4, 5\}, \ldots$ of σ, and the two chains share the row $R_5 = \{0, \ldots, 5\}$ — eleven shared rows from two orders of six terminals.

Choose the 5-set $S = \{0, 1, 3, 4, 5\}$, i.e. omit $j = 2$. The routing $\mathbf{z} = \mathbf{1}_S$ by itself is badly unbalanced: its errors are $z_i - f_i = \frac{1}{3}$ on S and $-\frac{2}{3}$ at j, and the identity prefix $P = \{0, 1, 2, 3, 4, 5\}$ accumulates $5 \cdot \frac{1}{3} - \frac{2}{3} = 1$. That is the left panel of Figure 8.

Now hand one unit of demand back, taking the credits

$$\mathbf{c} = (0,\ \tfrac{1}{3},\ 0,\ 0,\ \tfrac{1}{3},\ \tfrac{1}{3}), \qquad \textstyle\sum_i c_i = 1,$$

so that $\mathbf{u} = \mathbf{z} - \mathbf{c}$ has residual $\mathbf{u} - \mathbf{f} = \left(\frac{1}{3}, 0, -\frac{2}{3}, \frac{1}{3}, 0, 0\right)$. Its running sums are

$$\tfrac{1}{3},\ \tfrac{1}{3},\ -\tfrac{1}{3},\ 0,\ 0,\ 0$$

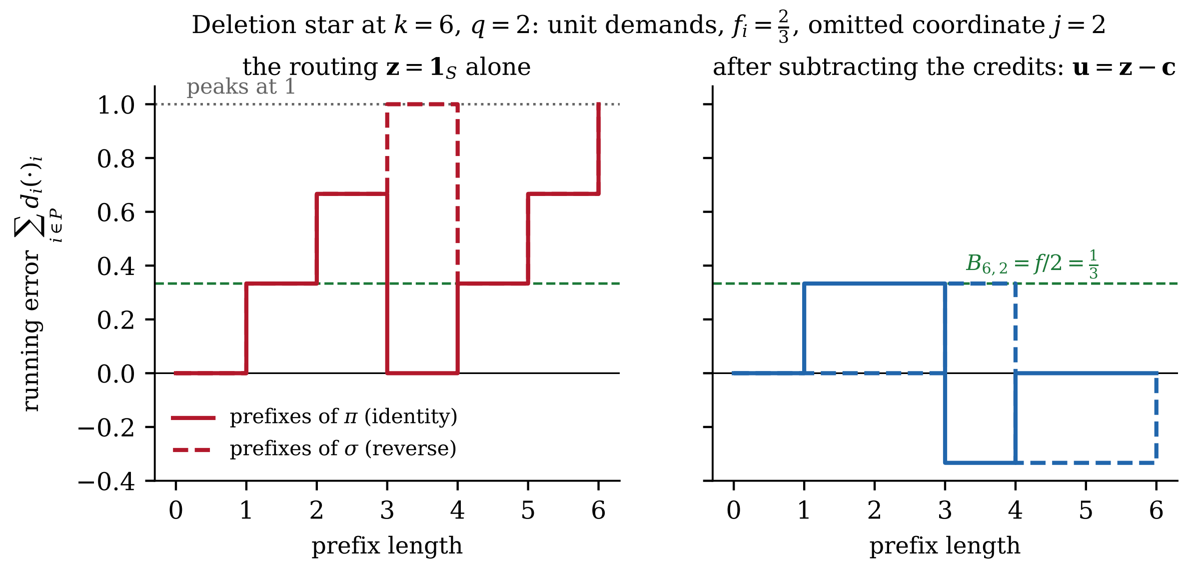

Figure 8: Example 6.3. Left: the running error of the rank-5 routing $\mathbf{z} = \mathbf{1}_S$ along the two orders; it climbs to 1. Right: after subtracting the deletion credits, the same data becomes a sawtooth that never leaves $[-\frac{1}{3}, \frac{1}{3}]$. Lemma 6.2 then bounds the routing's own overload by $1 + \frac{1}{3}$. The sawtooth shape is exactly the construction used in the proof of Theorem 6.9.

along π, and the reverse order sees the negatives of these; so every prefix of either order has residual at most $\frac{1}{3}$. Hence $B(\pi, \sigma, \mathbf{d}, \mathbf{f}) \le \frac{1}{3}$ — and in fact $B = \frac{1}{3}$ exactly, by Theorem 6.9 with $k = 6, q = 2$.

Lemma 6.2 now converts this into a statement about the flow: the rank-5 routing $\mathbf{1}_S$ is cost-good and has overload at most $1 + \frac{1}{3} = \frac{4}{3}$ on every arc of the instance. The credits are pure bookkeeping — they are not part of the routing — but they are what lets us bound the routing.

Remark 6.4 (a strengthening that fails). Definition 6.1 normalizes the credits by $\mathbf{1}^\mathsf{T}\mathbf{c} = 1$, which forces the *unweighted* residual $\mathbf{1}^\mathsf{T}(\mathbf{u} - \mathbf{f})$ to vanish but not the weighted one $\mathbf{d}^\mathsf{T}(\mathbf{u} - \mathbf{f})$. It is natural to ask for the stronger program in which both vanish. That program is genuinely different, and the half bound is *false* for it: already at $k = 4$ and complement mass three there is an exact instance whose strengthened optimum is $\frac{1}{2} + \frac{139}{458752}$, while the rank-only optimum of Definition 6.1 on the same instance is $1606605/67108864 < \frac{1}{2}$. Appendix E gives that instance, an exact cell whose strengthened supremum is $4 - 2\sqrt{3} = 0.5358\ldots$, and the associated exact certificates. Everything in the rest of this paper uses Definition 6.1, so none of it is affected.

So a residual bound β buys a congestion ceiling $1+\beta$. A bound $\beta = 1/3$ would give the appealing ceiling $4/3$. It is false.

Example 6.5 ($B = 3/8$; the 4/3 route is refuted)**.** Take $k = 8$, π the identity and σ its reverse, $d_i = 1$, and $f_i = 3/4$, so $\sum_i f_i = 6 = k - 2$ and $m = 7$. Every support omits exactly one index j, and $u_j - f_j = -3/4$. Because $\sum_i (u_i - f_i) = 0$, the identity prefix P strictly before j and the reverse prefix Q strictly after j satisfy

$$\Big(\sum_{i \in P}(u_i - f_i)\Big) + (u_j - f_j) + \Big(\sum_{i \in Q}(u_i - f_i)\Big) = 0, \quad \text{so} \quad \sum_{i \in P}(u_i - f_i) + \sum_{i \in Q}(u_i - f_i) = \tfrac{3}{4}.$$

Hence some prefix has residual at least 3/8, for every support and every credit allocation: $B \ge 3/8$. Equality is attained by omitting $j = 3$ and taking residuals

$$(u_i - f_i)_{i=0}^{7} = \left(\tfrac{1}{8}, \tfrac{1}{8}, \tfrac{1}{8}, -\tfrac{3}{4}, \tfrac{3}{32}, \tfrac{3}{32}, \tfrac{3}{32}, \tfrac{3}{32}\right),$$

i.e. $c_i = 1/8$ for $i \le 2$ and $c_i = 5/32$ for $i \ge 4$, whose partial sums in both orders never exceed $3/8$. So $B = 3/8 > 1/3$, and no argument that bounds a single deletion star by $1/3$ can yield the ceiling $4/3$. The instance itself is harmless as a lower bound — its demands are equal, hence a divisibility chain — but it is a genuine obstruction to the proof route.

6.1 The sharp codimension-two theorem

Example 6.3 produced a deletion star by hand, on a symmetric instance where the right choice of omitted coordinate was obvious. The theorem of this subsection says that a good choice always exists, on the whole boundary slice $\sum_i f_i = k-2$, with no symmetry and no assumption whatever on the two orders, the demands or the shares — zero demands are allowed. Its conclusion is that the residual never exceeds $\frac{1}{2}$, which by Lemma 6.2 means that these instances never force a constant above $\frac{3}{2}$.

Why is this the natural slice to attack? By Proposition 3.5 an instance of complement mass q cannot force more than $q-\eta$, so $q=1$ is already settled at 1 and $q=2$ is the first case that is not automatic. It is also the first case in which the argument can be carried out in closed form: on $\sum_i f_i = k-2$ the complement $\mathbf{x} = \mathbf{1} - \mathbf{f}$ has total mass exactly 2, so once one full complement coordinate is removed, what is left is exactly one probability vector, which is precisely the object the credits are allowed to be. That coincidence is what makes the proof below work, and it is also why the same proof does not generalize to $q \ge 3$.

The whole difficulty is in choosing which coordinate to omit. The proof does it by a counting argument: it identifies a set of candidates that are safe for the prefixes *before* the omitted coordinate (those whose weighted predecessor mass is at most one, in both orders at once), shows that this set is heavy, shows that the candidates which are unsafe for the full prefix are light, and concludes that a candidate safe for both must exist.

Theorem 6.6. *Let π, σ be arbitrary orders of $[k]$, let $0 \le d_i \le 1$ with $\max_i d_i = 1$, and let $\mathbf{f} \in [0,1]^k$ satisfy $\sum_i f_i = k-2$. Then*

$$B(\pi, \sigma, \mathbf{d}, \mathbf{f}) \;\le\; \frac{1}{2}.$$

No lower bound on the positive demands is required; zero demands are allowed.

Proof. Complement form. Put

$$x_i = 1 - f_i \in [0,1], \qquad \sum_i x_i = 2, \qquad \mu_i = d_i x_i, \qquad X = \sum_i \mu_i \;\le\; 2,$$

the last inequality because $d_i \le 1$. Choosing $S = [k] \setminus \{j\}$ and credits $\mathbf{c}$ is the same as choosing $\mathbf{w} = \mathbf{1} - \mathbf{u} = e_j + \mathbf{c}$, and if

$$g_i = d_i(w_i - x_i) = -d_i(u_i - f_i),$$

then the assertion $B \le \frac{1}{2}$ is exactly

$$g(P) := \sum_{i \in P} g_i \;\ge\; -\tfrac{1}{2} \qquad \text{for every prefix } P \text{ of } \pi \text{ and of } \sigma. \tag{4}$$

So we must choose one full complement coordinate j and one probability vector $\mathbf{c}$ on the rest, in such a way that the weighted running total never lags the target $\mathbf{x}$ by more than $1/2$.

The median sets. For an order τ put

$$A_\tau = \Big\{ j : \sum_{i \prec_\tau j} \mu_i \le 1 \Big\},$$

a prefix of τ. If $X > 1$ then $\mu(A_\tau) > 1$: indeed, either $A_\tau = [k]$, in which case $\mu(A_\tau) = X > 1$, or the τ-successor j' of the last element of A_τ exists and satisfies $\mu(A_\tau) = \sum_{i \prec_\tau j'} \mu_i > 1$. Consequently, by inclusion–exclusion,

$$\mu(A_\pi \cap A_\sigma) \;\geq\; \mu(A_\pi) + \mu(A_\sigma) - X \;>\; 2 - X. \tag{5}$$

Bad coordinates. For a candidate j set

$$\alpha_j = \frac{1 - x_j}{2 - x_j} \in [0, \tfrac{1}{2}]$$

(decreasing in x_j, equal to $\frac{1}{2}$ at $x_j = 0$ and to 0 at $x_j = 1$), and call j *bad* if

$$\alpha_j\,(X - 2d_j) > \tfrac{1}{2}. \tag{6}$$

Bad-mass lemma: if $X > 1$ then $(A_\pi \cap A_\sigma)$ contains a coordinate that is not bad. Let Q be the bad set and $s = \sum_{j \in Q} x_j$. A bad coordinate has $\alpha_j > 0$, since $\alpha_j = 0$ makes the left-hand side of (6) vanish; so we may divide by α_j. As $\alpha_j \leq \frac{1}{2}$, (6) forces $X - 2d_j > 1/(2\alpha_j) \geq 1$, hence

$$d_j < \frac{X - 1}{2} \qquad (j \in Q), \tag{7}$$

and in particular Q is empty unless $X > 1$. Multiplying (7) by $x_j \geq 0$ and summing,

$$\mu(Q) \;\leq\; \frac{X - 1}{2}\, s, \tag{8}$$

with strict inequality when $s > 0$. On the complement, $d_i \leq 1$ gives

$$\mu([k] \setminus Q) \;\leq\; \sum_{i \notin Q} x_i = 2 - s. \tag{9}$$

We distinguish two cases according to whether $X < 2$ or $X = 2$.

Case $1 < X < 2$. We show $\mu(Q) < 2 - X$. If $s = 0$ then every bad coordinate has $x_j = 0$, hence $\mu_j = 0$ and $\mu(Q) = 0 < 2 - X$. If $s > 0$, add the strict form of (8) to (9):

$$X < \frac{X - 1}{2} s + 2 - s = 2 - \frac{3 - X}{2}\, s \quad\Longrightarrow\quad s < \frac{2(2 - X)}{3 - X},$$

and therefore, using (8) again,

$$\mu(Q) < \frac{X - 1}{2} \cdot \frac{2(2 - X)}{3 - X} = \frac{(X - 1)(2 - X)}{3 - X} \;\leq\; 2 - X,$$

the last step because $X \leq 2$ gives $X - 1 \leq 3 - X$. Combining with (5), $\mu(Q) < 2 - X < \mu(A_\pi \cap A_\sigma)$, so $A_\pi \cap A_\sigma \not\subseteq Q$.

Case $X = 2$. Here $\sum_i (1 - d_i) x_i = \sum_i x_i - X = 0$ with all terms nonnegative, so $d_i = 1$ whenever $x_i > 0$; such a coordinate has $\alpha_i(X - 2d_i) = \alpha_i \cdot 0 = 0$ and is therefore not bad. By (5) the set $A_\pi \cap A_\sigma$ has positive μ-mass, so it contains a coordinate with $\mu_i > 0$, hence with $x_i > 0$, hence non-bad.

In both cases the lemma follows.

The construction. If $X \leq 1$ pick any j; if $X > 1$ pick a non-bad $j \in A_\pi \cap A_\sigma$. In either case set

$$S = [k] \setminus \{j\}, \qquad c_i = \frac{x_i}{2 - x_j}\ (i \neq j), \qquad c_j = 0.$$

These are valid credits: $\sum_{i \neq j} c_i = (2 - x_j)/(2 - x_j) = 1$ and $0 \leq c_i \leq 1$ because $x_i \leq 1 \leq 2 - x_j$. Write $\alpha = \alpha_j$. For $i \neq j$,

$$c_i - x_i = x_i\Big(\frac{1}{2 - x_j} - 1\Big) = -\alpha x_i, \qquad \text{so} \qquad g_i = -\alpha\mu_i \leq 0,$$

while $g_j = d_j(1 - x_j) \geq 0$.

Prefixes not containing j. Such a prefix P (in either order) is contained in the set of τ-predecessors of j, so

$$g(P) = -\alpha\,\mu(P) \;\geq\; -\alpha \;\geq\; -\tfrac{1}{2},$$

using $\mu(P) \leq 1$ — which holds because $j \in A_\pi \cap A_\sigma$ when $X > 1$, and because $\mu(P) \leq X \leq 1$ when $X \leq 1$.

Prefixes containing j. Once j has been absorbed, every further increment is $g_i = -\alpha\mu_i \leq 0$, so g is nonincreasing and its minimum over such prefixes is the full sum. Substituting,

$$g([k]) = d_j(1 - x_j) - \alpha(X - \mu_j) = \alpha\big(d_j(2 - x_j) - X + d_j x_j\big) = \alpha\,(2d_j - X),$$

where the middle step used $d_j(1 - x_j) = \alpha\,d_j(2 - x_j)$. If $X > 1$ then j was chosen non-bad, which is precisely $\alpha(2d_j - X) \geq -\frac{1}{2}$. If $X \leq 1$ then $g([k]) \geq -\alpha X \geq -\frac{1}{2}$ since $\alpha \leq \frac{1}{2}$.

In all cases (4) holds, which is the theorem. □

Remark 6.7. The mass estimate is the step that removes any assumption of the form $d_i \geq \frac{1}{2}$, which earlier arguments needed. A low-demand candidate may well have a bad full-prefix error; the point is that all such candidates together are too light to cover the two-order median intersection. The special feature of codimension two is that $\sum_i x_i = 2$: after fixing one full complement coordinate, what remains is exactly one probability vector, and proportional credits give every pre-j increment the same sign.

Corollary 6.8. *Every common-point interval cell with*

$$k - 2 \;\leq\; \sum_i f_i \;\leq\; k - 1$$

admits a cost-good routing of overload at most $\frac{3}{2}$, and hence cannot witness a constant larger than $\frac{3}{2}$. In particular this covers every cell of complement mass $q = 2$, i.e. every cell whose first cost-good rank is $m = k - 1$: by (2) those are exactly the cells with $k - 2 < \sum_i f_i \leq k - 1$.

Proof. Choose $\mathbf{f}' \leq \mathbf{f}$ coordinatewise with $\sum_i f'_i = k - 2$; this is possible because $\sum_i f_i \geq k - 2$ and $\mathbf{f} \geq \mathbf{0}$. Apply Theorem 6.6 to $\mathbf{f}'$, obtaining S of size $k - 1$ and credits $\mathbf{c}$ with $\sum_{i \in P} d_i(u_i - f'_i) \leq \frac{1}{2}$ on all prefixes. Since $d_i \geq 0$ and $f_i \geq f'_i$,

$$\sum_{i \in P} d_i(u_i - f_i) \;\leq\; \sum_{i \in P} d_i(u_i - f'_i) \;\leq\; \tfrac{1}{2},$$

and Lemma 6.2 adds one credit unit, giving overload at most $\frac{3}{2}$ for the rank-$(k-1)$ routing $\mathbf{1}_S$. Finally $\mathbf{1}_S$ is cost-good: $\sum_i z_i = k - 1 \geq \sum_i f_i$. □

This is the first unconditional ceiling below 2 for a nontrivial family of extremal cells; its scope is the complement-mass-two slab. Note that this slab is not where the current records are (e.g., the $k = 17$ record has complement mass eleven), and since $3/2$ exceeds every certified value, the corollary does not constrain them. But combined with the trivial ceiling $C^\star \leq \max(1, q - \eta)$ obtained from the always-cost-good all-late routing, it shows that any family exceeding $3/2$, in particular any family approaching the conjectured supremum 2, must have complement mass at least three.

6.2 Sharpness at every codimension

Theorem 6.6 gives $B \le \frac{1}{2}$ on the codimension-two slice, and Example 6.5 shows the constant cannot be improved to $\frac{1}{3}$ there. But one might still hope that $\frac{1}{2}$ is an accident of $q = 2$ — that the same method, pushed to higher complement mass, would give something smaller and so a ceiling below $\frac{3}{2}$. The next theorem rules that out. On the uniform identity/reverse instances it computes B *exactly*, for every k and every q, and the answer is $\frac{1}{2} - \frac{q}{2k}$ whenever an explicit capacity condition holds. Fixing q and letting $k = qm$ grow therefore drives B to $\frac{1}{2}$ from below, separately for each codimension: the half is a property of the deletion-star device itself, not of the slice we proved it on.

The proof also explains *why* $\frac{1}{2}$: each omitted coordinate makes the running residual jump down by f, and the two orders between them force the values on the two sides of that jump to be symmetric about zero, so no allocation can do better than $f/2$. The matching construction is the sawtooth already seen in Example 6.3 and Figure 8; the capacity condition is exactly the statement that there are enough selected coordinates to climb back up between consecutive jumps.

Theorem 6.9. *Fix integers $2 \le q \le k$, take π the identity and σ its reverse, $d_i = 1$, and uniform shares $f_i = f = (k-q)/k$, so that $\sum_i f_i = k - q$ and the deletion star omits $q-1$ coordinates. Put $n = k - q$ and*

$$M_{k,q} = 2\left\lceil \frac{n}{2q} \right\rceil + (q-2)\left\lceil \frac{n}{q} \right\rceil.$$

Then

$$B_{k,q} \;\ge\; \frac{f}{2} = \frac{1}{2} - \frac{q}{2k}, \qquad \textit{and} \qquad B_{k,q} = \frac{f}{2} \iff n = 0 \;\; \textit{or} \;\; M_{k,q} \le n+1.$$

In particular $k = mq$ always satisfies the criterion, so

$$B_{mq,q} = \frac{1}{2} - \frac{1}{2m} \xrightarrow[m\to\infty]{} \frac{1}{2}$$

for every fixed q (Figure 10).

The proof has three parts — the jump lower bound, the necessity of the capacity condition, and the sawtooth construction that matches it — and is given in Appendix F. The construction is the one already seen in Example 6.3; Figure 9 shows it at $k = 12$, $q = 3$.

Equal demands are a divisibility chain, so Theorem 6.9 is a sharpness statement about the *method*, not an SSUF lower bound. The frontier nevertheless survives the demand firewall.

Theorem 6.10. *Fix $q \ge 2$ and $k = qm$, take identity and reverse orders and uniform shares $f_i = f = (k-q)/k$ as above, but set $d_0 = \frac{2}{3}$ and $d_i = 1$ for $i > 0$. (After scaling, the demand set is $\{2, 3\}$, which is not a divisibility chain.) Then*

$$\frac{1}{2} - \frac{1}{2m} - \frac{1}{6m} \;\le\; B^{\{2/3,1\}}_{qm,q} \;\le\; \frac{1}{2} - \frac{1}{2m}.$$

Both bounds tend to $\frac{1}{2}$ as $m \to \infty$.

The proof, in Appendix F, is the same sawtooth with the low-demand item parked on a zero-residual coordinate, plus a two-case argument for the lower bound.

Together with Corollary 6.8, this delimits the deletion-star method exactly: it yields the ceiling $\frac{3}{2}$ on the codimension-two slab, it cannot yield anything below $\frac{3}{2}$ by improving the constant $\frac{1}{2}$, and any ceiling below $\frac{3}{2}$ for the whole class must use a mechanism beyond one deletion star per support.

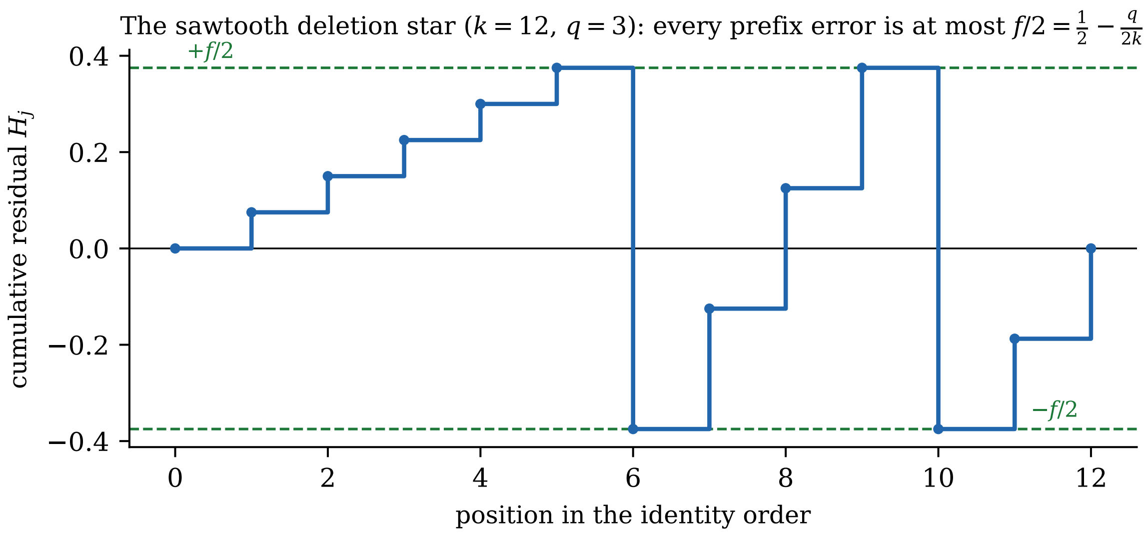

Figure 9: The sawtooth deletion star of Theorem 6.9 ($k = 12$, $q = 3$). The cumulative residual rises to $+f/2$, drops by exactly f at each omitted coordinate, and returns to zero; every prefix of the identity order and of the reverse order is therefore bounded by $f/2 = \frac{1}{2} - \frac{q}{2k}$.

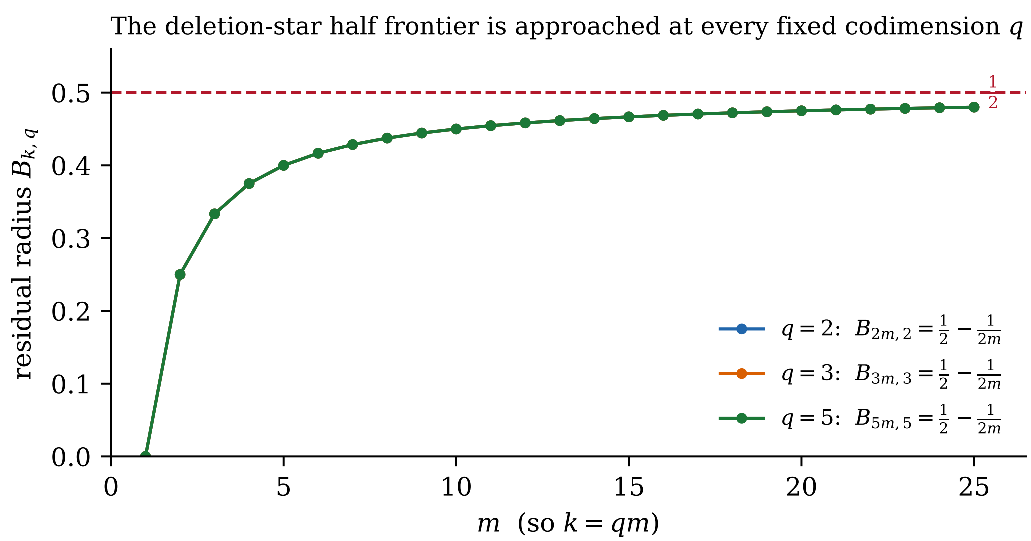

Figure 10: Theorem 6.9: at every fixed codimension q the exact residual radius is $\frac{1}{2} - \frac{1}{2m}$ along $k = qm$, so the half frontier is approached for each q separately.

7 Two-order upper-bound mechanisms

Section 6 bounded the constant on one slice of the class by constructing a good routing directly. This section collects the other route to upper bounds, and it is worth saying at the outset how it connects to C, since the individual statements can look unrelated to flows.

There are two ways a result about two-order rounding turns into a bound on the cost-preserving constant.

Directly, through the cost functional. By Observation 3.2 the cost constraint is a single signed linear functional $\mathbf{g}$ of the vector $\big(d_i(z_i - f_i)\big)_i$. So a rounding theorem that produces a $\mathbf{z}$ which is prefix-accurate *and* satisfies $\mathbf{g}^\top(\mathbf{z} - \mathbf{f}) \geq 0$ is exactly a cost-preserving routing, and its prefix accuracy is exactly the congestion. This is what Theorem 7.2 does: min-cost two-way rounding gives a routing that preserves an arbitrary functional and is accurate to within $B(\mathbf{d}) = \min_{\lambda \geq 0}(\lambda + \sum_i |d_i - \lambda|)$ on every prefix of both orders, so $C \leq B(\mathbf{d})/d_{\max}$ on any instance of the class with that demand vector. Proposition 7.3 then evaluates $B(\mathbf{d})$ exactly and identifies when it equals $d_{\max}$, i.e. when this argument alone already gives the conjectured-optimal $C \leq 1$: precisely when all demands but one maximum coincide. That locus contains incommensurable two-scale systems, so it is not covered by the divisibility-chain theorem.

Indirectly, through the lower box. Theorem 7.5 and Theorem 7.7 do not preserve a cost functional at all. Instead they produce a routing that is two-sided $d_{\max}$-good, which puts it in $U(\mathbf{q})$ and simultaneously above $\mathbf{q} - d_{\max}\mathbf{1}$, and so verifies (LB) of Definition 3.8 on their respective classes. By Proposition 3.10, (LB) yields the cost-preserving bound $C \leq 2$ there. This is a weaker conclusion than the first route, but it applies for *arbitrary* demands, which the first route does not.

The rest of the section then maps the boundary of what these mechanisms can reach: Section 7.4 exhibits minimal instances on which every *scalar* strengthening of the one-order induction fails, and Section 7.5 isolates the one statement that does survive.

Throughout, π, σ are orders of $[k]$, $\mathbf{d} \in (0, d_{\max}]^k$, $\mathbf{f} \in [0,1]^k$, and $E_P(\mathbf{z}) = \sum_{i \in P} d_i(z_i - f_i)$.

7.1 A cost-preserving Knuth transfer

Knuth's two-way rounding theorem [4] states that for unit demands there is $\mathbf{z} \in \{0,1\}^k$ with $|\sum_{i \in P}(z_i - f_i)| < 1$ on every prefix of both orders. Knuth proves it by a network flow argument: build a network with a unit-capacity middle arc $u_i v_i$ per item, so that the natural fractional maximum flow uses that arc to extent f_i and every integral maximum flow uses it to extent $z_i \in \{0,1\}$. We need a cost-preserving version.

Lemma 7.1 (min-cost two-way rounding)**.** *For unit demands, two orders, and* any *signed functional* $\mathbf{g} \in \mathbb{R}^k$*, there is* $\mathbf{z} \in \{0,1\}^k$ *with*

$$\Big|\sum_{i \in P}(z_i - f_i)\Big| < 1 \quad \text{on every prefix of both orders,} \qquad \mathbf{g}^{\mathsf{T}}\mathbf{z} \;\geq\; \mathbf{g}^{\mathsf{T}}\mathbf{f}.$$

Proof. Give the middle arc $u_i v_i$ cost $-g_i$ and leave all other costs zero. The natural fractional maximum flow has cost $-\mathbf{g}^{\mathsf{T}}\mathbf{f}$. The minimum-cost maximum-flow polytope of an integral network is integral, so it has an integral optimum, whose cost is therefore at most $-\mathbf{g}^{\mathsf{T}}\mathbf{f}$; that is, the corresponding $\mathbf{z}$ satisfies $\mathbf{g}^{\mathsf{T}}\mathbf{z} \geq \mathbf{g}^{\mathsf{T}}\mathbf{f}$. Being an integral maximum flow, it also satisfies Knuth's prefix bound. If $\sum_i f_i$ is not an integer, append one zero-cost dummy coordinate at the end of both orders to make the total integral; deleting it afterwards preserves both conclusions. □

Theorem 7.2 (cost-preserving Knuth transfer)**.** *For two orders, demands* $d_i \leq d_{\max}$*, shares* $\mathbf{f}$ *and any signed functional* $\mathbf{g}$*, there is* $\mathbf{z} \in \{0,1\}^k$ *with* $\mathbf{g}^{\mathsf{T}}(\mathbf{z} - \mathbf{f}) \geq 0$ *and*

$$|E_P(\mathbf{z})| \;<\; B(\mathbf{d}) := \min_{\lambda \geq 0}\Big(\lambda + \sum_i |d_i - \lambda|\Big)$$

on every prefix of both orders.

Proof. Take $\mathbf{z}$ from Lemma 7.1 and put $e_i = z_i - f_i$, so $|e_i| \leq 1$ and $|\sum_{i \in P} e_i| < 1$. For any $\lambda \geq 0$,

$$\Big|\sum_{i \in P} d_i e_i\Big| = \Big|\lambda \sum_{i \in P} e_i + \sum_{i \in P}(d_i - \lambda)e_i\Big| \leq \lambda\Big|\sum_{i \in P} e_i\Big| + \sum_{i \in P}|d_i - \lambda|\,|e_i| < \lambda + \sum_i |d_i - \lambda|.$$

Minimizing over λ gives the claim. □

Proposition 7.3 (the exact radius)**.** *Write $d_{(1)} \le \cdots \le d_{(k)} = d_{\max}$. Then the minimum defining $B(\mathbf{d})$ is attained at a lower median of $\mathbf{d}$, and*

$$d_{\max} \;\le\; B(\mathbf{d}) \;\le\; \left\lceil \frac{k}{2} \right\rceil d_{\max}, \qquad B(\mathbf{d}) = d_{\max} \iff d_{(1)} = \cdots = d_{(k-1)}.$$

Moreover, for two scales $d_i \in \{a d_{\max}, d_{\max}\}$ with ℓ low and h high demands,

$$\frac{B(\mathbf{d})}{d_{\max}} = \min\{\, a + h(1-a),\; 1 + \ell(1-a) \,\}.$$

The proof is a short computation with the piecewise-linear convex function $\varphi(\lambda) = \lambda + \sum_i |d_i - \lambda|$; we give it in Appendix F.

So whenever all demands except possibly one maximum coincide — including *incommensurable* two-scale systems, which are not divisibility chains — a single routing is simultaneously arbitrary-cost-preserving and two-sided $d_{\max}$-good in both orders.

Remark 7.4. The obstruction to a direct weighted version of Knuth's network is exact: replacing the unit middle arc by one of capacity d_i permits partial integral amounts $1, \ldots, d_i - 1$, whereas the rounding problem requires an all-or-nothing choice.

7.2 Deque orders

Recall the lower box (LB) and the upper-good set $U(\mathbf{q})$ from Section 3.4; by Proposition 3.10 its validity would give the universal cost-preserving constant 2. The following class settles it outright.

Theorem 7.5 (deque orders)**.** *Suppose every prefix of σ is an interval of π; equivalently, after its first element σ repeatedly appends the immediate unused π-predecessor or π-successor. There are exactly 2^{k-1} such orders. Then for arbitrary demands there is a single routing that is two-sided $d_{\max}$-good on all prefixes of both orders, and hence the weighted upper-good inequality and (LB) hold on the whole deque class.*

Proof. Count first. Read the construction backwards: reversing σ and peeling off its elements one at a time, each peeled element is the leftmost or the rightmost element of the π-interval that remains, and the final element admits only one choice. This is a bijection between deque orders and binary strings of length $k-1$, so there are exactly 2^{k-1} of them. Forwards, a prefix of σ is exactly the complement of the interval not yet peeled, hence a π-interval; conversely if every prefix of σ is a π-interval then consecutive prefixes differ by one element at one of the two ends, which is the deque rule.

Now apply Liu–Reis [7] with $m = 2$ assignment rows in the single order π: their constant is $1 - \frac{1}{2m-2} = \frac{1}{2}$, so the resulting integral assignment has two-sided error at most $d_{\max}/2$ on every π-prefix. Each σ-prefix is a π-interval $[a, b]$, hence the difference of the two π-prefixes ending at b and at $a-1$; its error is therefore at most $d_{\max}/2 + d_{\max}/2 = d_{\max}$ in absolute value. The property is hereditary under deleting coordinates with $f_i \in \{0, 1\}$, so the routing stays inside the positive support. Being two-sided $d_{\max}$-good, it lies in $U(\mathbf{q})$ and satisfies $\mathbf{y} \ge \mathbf{q} - d_{\max}\mathbf{1}$, so the single-point mixture $\mathbf{m} = \mathbf{y}$ witnesses (LB). ☐

The deque class strictly exceeds the radius-$d_{\max}$ locus of Proposition 7.3: there is an explicit seven-scale deque instance with $B(\mathbf{d}) = 71631337/43732568 > d_{\max}$, for which Theorem 7.5 nonetheless applies.

7.3 A maximal switched-interval lift

Common-point systems are not the only exact-two-path instances. In general, if terminal i has exactly two simple paths P_i^0, P_i^1, then by Observation 3.2 everything is governed by the route-difference matrix whose i-th column is $\boldsymbol{\Delta}_i = \boldsymbol{\chi}(P_i^1) - \boldsymbol{\chi}(P_i^0)$, restricted to the shared arcs.

Definition 7.6 (signed tree network)**.** A matrix $A \in \{-1, 0, 1\}^{E \times k}$ is a *signed tree network* if there is a tree T whose edge set is E and, for each i, an ordered pair (q_i, p_i) of nodes of T, such that the i-th column of A is the signed incidence vector of the unique undirected q_i–p_i path in T: the entry A_{ei} is $+1$ if e is traversed forwards on that path, -1 if backwards, and 0 if e is not on it.

For a common-point interval system T is a path and every column is a $+1$-interval, so Definition 7.6 strictly generalizes the class of Section 3.

Theorem 7.7 (switched-interval lift)**.** *Let $A \in \{-1, 0, 1\}^{E \times k}$ be the route-difference matrix of an exact-two-path instance whose shared part is a signed tree network in the sense of Definition efdef:signedtree. Suppose there are one terminal order τ, column switches $\varepsilon_i \in \{\pm 1\}$ and row signs $\rho_e \in \{\pm 1\}$ such that every switched row $(\rho_e A_{ei} \varepsilon_i)_{i=1}^k$ is the $0/1$ indicator of one τ-interval. Then one routing is simultaneously upper- and lower-good, so (LB) holds; with Proposition 3.10 this gives the cost-preserving bound $C \leq 2$ on this class.*

Proof. Complement the coordinates with $\varepsilon_i = -1$ (replace f_i by $1 - f_i$ and z_i by $1 - z_i$); this turns every switched row into a nonnegative interval indicator without changing any $|E_P|$. Apply Liu–Reis in the single order τ with two rows, obtaining two-sided error at most $d_{\max}/2$ on every τ-prefix. An interval is a difference of two prefixes, so every switched row has absolute error at most $d_{\max}$; private exits are singleton rows with error at most $d_i \leq d_{\max}$. Undoing the complementation and the row signs changes only signs, so the routing is two-sided $d_{\max}$-good on every row. $\square$

The hypothesis is maximal for any proof that uses one common order and only the Liu–Reis prefix bound, in the following exact sense.

Proposition 7.8. *For a row $\mathbf{a} \in \mathbb{R}^k$ define the prefix-representation norm*

$$V_\tau(\mathbf{a}) = |a_k| + \sum_{j=1}^{k-1} |a_j - a_{j+1}|.$$

Then $\mathbf{a} = \sum_j c_j \boldsymbol{\chi}(\{1, \ldots, j\})$ with $c_j = a_j - a_{j+1}$ ($j < k$) and $c_k = a_k$, so $\sum_j |c_j| = V_\tau(\mathbf{a})$ and the error of row $\mathbf{a}$ is at most $(d_{\max}/2) V_\tau(\mathbf{a})$. For a nonzero $\mathbf{a} \in \{-1, 0, 1\}^k$ one has $V_\tau(\mathbf{a}) \leq 2$ if and only if $\mathbf{a}$ is, after one row sign, the indicator of a single τ-interval; every row supported on two separated intervals has $V_\tau(\mathbf{a}) \geq 3$.

Proof. The representation and the norm identity are immediate by telescoping. If $\mathbf{a} = \boldsymbol{\chi}([p, q])$ then $V_\tau = 2$ when $1 < p \leq q < k$, and $V_\tau = 1$ when $p = 1$ or $q = k$; sign changes do not affect V_τ. Conversely suppose $\mathbf{a}$ has two maximal nonzero blocks $[p_1, q_1]$ and $[p_2, q_2]$ with $q_1 + 1 < p_2$. Then $|a_{q_1} - a_{q_1+1}| = 1$, $|a_{p_2-1} - a_{p_2}| = 1$, and either $|a_{q_2} - a_{q_2+1}| = 1$ (if $q_2 < k$) or $|a_k| = 1$; these three contributions are distinct, so $V_\tau(\mathbf{a}) \geq 3$. A row that changes sign inside a single block likewise has $V_\tau \geq 3$. $\square$

Thus rows with two separated intervals are exactly the ones this proof cannot reach, and the bound 3 is attained (e.g. $\mathbf{a} = (1, 0, \ldots, 0, 1)$).

7.4 Exact walls for scalar induction

Three minimal exact instances close off every attempt to carry a *scalar* state through a two-order induction. We state them as facts about explicit small instances; each is verified by complete enumeration.

Proposition 7.9. *1.* Path coherence. *Three fully supported intervals* $[0,3],[1,5],[2,4]$ *on a six-vertex path, with* $\mathbf{d}=(1,\frac{51}{52},\frac{51}{52})$ *and* $\mathbf{f}=(\frac{1}{2},\frac{44}{51},\frac{29}{51})$*, admit radius-one routings (the minimum radius is* $29/52$*), but the minimum* oscillation $\max_e r_e-\min_e r_e$ *among radius-one routings is* $55/52>1$.

2. Signed global coherence. *Two signed columns on a rooted* K_5 *support, with* $\mathbf{d}=(1,\frac{15}{16})$ *and* $\mathbf{f}=(\frac{7}{16},\frac{1}{3})$*, force minimum row-error diameter* $17/16$ *among all radius-one routings. A complete CEGIS exhaustion over all rooted tree shapes of* K_4,K_5,K_6 *gives* 43 *satisfiable and* 16 *unsatisfiable cases, split as* $K_4:0/8$*,* $K_5:7/8$*,* $K_6:36/0$.

3. Diagonal bracket. *The diagonal first-crossing bracket holds through* $k=5$ *and first fails at* $k=6$*: for* $\sigma=(3,5,0,4,2,1)$ *with the demands and shares of Appendix B, exactly seven routings are good on all* proper *prefixes of both orders, and the closest opposite-sign values of the preserved functional on them are separated by* $45/44>d_{\max}$*. Of all* 720 *permutations,* 704 *satisfy the bracket and* 16 *fail.*

None of these refutes (LB) or the cost-preserving statement. All sixteen $k=6$ bracket failures have exact lower-box deficit zero, and the largest all-cost hull entry among them is $129/134<1$; the instance of (3) above has hull entry $41/44$. What they do show is that no single scalar carried along one order (e.g., a coherent path state, a global oscillation bound, or a diagonal crossing pair) suffices.

7.5 The convex Pareto transfer

What still holds is a statement about convex mixtures rather than about single routings. For routing errors $\mathbf{r}(\mathbf{z})=\mathbf{y}(\mathbf{z})-\mathbf{q}$ and $U(\mathbf{q})=\{\mathbf{z}:\mathbf{r}(\mathbf{z})\le d_{\max}\mathbf{1}\}$, put

$$\Delta(\mathbf{q})=\min_{\boldsymbol{\lambda}\in\Delta(U(\mathbf{q}))}\max_a\Big\{-\sum_{\mathbf{z}}\lambda_{\mathbf{z}}r_a(\mathbf{z})\Big\}.$$

By LP duality (the inner maximum is over a simplex and the outer over a polytope),

$$\Delta(\mathbf{q})=\max_{\mathbf{w}\ge 0,\ \|\mathbf{w}\|_1=1}\Big\{-\max_{\mathbf{z}\in U(\mathbf{q})}\mathbf{w}^{\mathsf{T}}\mathbf{r}(\mathbf{z})\Big\},\tag{10}$$

and (LB) is exactly $\Delta(\mathbf{q})\le d_{\max}$. Equivalently, by (10), the minimal surviving two-order theorem is

$$\max_{\mathbf{z}\in U(\mathbf{q})}\mathbf{w}^{\mathsf{T}}\mathbf{r}(\mathbf{z})\ \ge\ -d_{\max}\|\mathbf{w}\|_1\qquad(\mathbf{w}\ge 0),$$

which asks only for the downward convex Pareto image of the upper-good routings, strictly less than reproducing all terminal marginals.

There are two structural reductions in play here. Call the raw digraph of a common-point system a *crossed two-arm broom*: a spine carrying the two prefix chains of Lemma 3.4, with the two private exits of each terminal hanging off it as the two arms, crossed because the intervals need not be nested. On such a broom every nonnegative raw separator collapses to the terminal coefficients $d_i(A_i-B_i+\gamma_i^L-\gamma_i^R)$, where A_i,B_i are order-monotone suffix potentials of the two prefix systems and γ^L, γ^R are the private-exit weights; this removes all extraneous objective directions. Second, at a fractional vertex of the two-chain LP, differencing consecutive tight prefixes turns the active rows into disjoint blocks in each order, and after column scaling the tight matrix is a bipartite incidence matrix. Its *active-block graph* has one node per active block of each order, plus one *dummy node* per order absorbing the coordinates that no active block of that order covers, and one edge per fractional variable. Nonsingularity then forces this graph to be a forest with no component containing both dummy nodes; consequently every component has a non-dummy leaf carrying exactly one fractional variable.

The missing operation is therefore weighted all-or-nothing coalescence on these leaf edges while carrying the two correlated endpoint errors — the two-order analogue of matrix rounding for unions of two laminar families. In the unit-weight case that problem is totally unimodular and solved by [1]; column weights destroy the unimodularity, which is exactly the gap. The finite evidence is aligned with the statement: 30000 crossed-broom separation LPs at $k = 6, \ldots, 11$ all have numerical deficit zero, and 2000 exact strict-rank QF-LRA cells at $k = 6, \ldots, 9$ with demands in $\{1, \frac{15}{16}\}$ are all UNSAT for excluding every rank-$\lceil k/2 \rceil$ upper-good routing.

8 Three-order overlays are dominated at $6/5$

Since two-order systems are exactly the common-point interval class, escaping to a larger constant means leaving it. The obvious attempt is a three-order overlay: three replicas of the terminal set, with rows indexed by the nonconstant binary words. The smallest literal such equality overlay has six shared rows. It cannot help.

Proposition 8.1. *For the six-row three-replica equality overlay with* $\mathbf{d} = (1, \frac{9}{10}, \frac{4}{5})$ *and* $\mathbf{f} = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, *and for both the lexicographic row order and a strongest two-pattern order, the fractional load is an exact convex combination of* 16 *routings, each of congestion at most* $6/5$. *Consequently no arc-toll vector (not even a signed one) can force congestion above* $6/5$ *on this overlay.*

Proof. Let $\mathbf{x} = \sum_j \lambda_j \mathbf{y}^{(j)}$ with $\lambda_j > 0$, $\sum_j \lambda_j = 1$, and $\max_a(y_a^{(j)} - x_a) \le \frac{6}{5} d_{\max}$ for each j; such a mixture is exhibited explicitly. For an arbitrary $\mathbf{c} \in \mathbb{R}^A$ we have $\mathbf{c}^\mathsf{T}\mathbf{x} = \sum_j \lambda_j \, \mathbf{c}^\mathsf{T}\mathbf{y}^{(j)}$, so some j has $\mathbf{c}^\mathsf{T}\mathbf{y}^{(j)} \le \mathbf{c}^\mathsf{T}\mathbf{x}$. That routing is cost-good and has congestion at most $6/5$. ☐

An exhaustive sweep over all 720 global row orders shows that 576 orders protect exactly one mixed word above 3037/2500, 144 protect exactly two (always complementary), and none protects three or more; all 144 strongest orders also carry a componentwise-dominating mixture at the calibrated unequal point. This exhausts the literal global-row-order overlay: changing row order or tolls cannot work, and the minimal construction not covered is a three-path cell with positive fractional mass on one hybrid/baseline path per replica. This again, like in [14], shows that constants come from baselines, not from incidence.

9 Verification

All claims above were re-derived independently, in exact rational or symbolic arithmetic, by programs written from the statements rather than from the original search code. Specifically: the $k = 17$ record was rebuilt from its intervals alone (raw arc list, DFS path discovery, all 131072 routings in exact integers, the 15 minimizers, the 8885 strict-sublevel routings and their maximum rank, and the 18-atom hull certificate); the lower-box certificate of Proposition 4.3 was recomputed from scratch, both primal and dual, over the complete set of 4799 upper-good routings; the constructive proof of Theorem 6.6 was executed verbatim on 10 500 random exact instances with $3 \leq k \leq 9$ and on an exhaustive grid of 164 160 instances at $k = 4$, and never failed, with the bad set nonempty in a substantial fraction of cases so that the hard branch is genuinely exercised; Theorem 6.9 was checked against exact LP optima for all $3 \leq k \leq 8$ and $2 \leq q \leq k$, including the cases where the capacity criterion predicts strict inequality; Theorem 6.10 was checked for $q \in \{2, 3, 4\}$ and $m \in \{1, 2, 3\}$; Proposition 7.3 was checked on 4000 random exact demand vectors; the deque count and the interval characterization were checked exhaustively through $k = 7$; Proposition 7.8 was checked by exhausting $\{-1, 0, 1\}^k$ through $k = 6$; and the walls of Proposition 7.9 and the mixture of Proposition 8.1 were recomputed from their primitive data. The $q = 3$ zero-sum convention was audited independently from its embedded executable block: all 91 proportional pairs and all 91 unrestricted supports were reconstructed, all 63 admissible supports returned literal full residual zero, and the remaining 28 were rejected by exact moment infeasibility; the independent support table has the hash printed in Appendix E. The $k = 6$ Pell family was rebuilt from its defining formulas and its exact constants recomputed by full enumeration.

10 Discussion and open problems

Let us conclude by assembling the whole picture. On the lower side, every certified record, from the four-interval instance of [14] to the seventeen-terminal instance of Section 4, uses the same mechanism: mutually crossing intervals through one common arc, i.e., a weighted two-permutation prefix system on a strict-rank cost face, with first cost-good rank $\lceil k/2 \rceil$.

On the upper side, the same mechanism now has its first unconditional upper bounds: the deletion-star residual is exactly $\frac{1}{2}$ on the codimension-two face (congestion $\frac{3}{2}$ on every cell with rank $k-1$), the constant $\frac{1}{2}$ is approached at every fixed codimension, and it fails beyond codimension two, so the ladder of upper bounds rises toward the universal constant 2 conjectured in [14], while the record ladder rises toward the conjectured class supremum 4/3.

The two conjectures are not in tension, but single deletion stars cannot decide between them, and the observed coincidence of multi-star thresholds with the exact record constants identifies the invariant that can. In addition, the continuum theorem shows the limit is irreducibly microscopic, and the lower-box deficits of order 10^{-10} on every record keep the route to the universal constant 2 fully open. The open problems below are ordered by this geometry: close the residual frontier at every codimension, turn the multi-star identity into a theorem, evaluate C_{cls}, and settle the lower box on the exact-two-path class. The conjecture of [14] that the exact supremum of lower bounds is 2, remains open, and the gap [1.2825, 2] is still the main open question.

Open Problem 10.1 (the value of the class)**.** Evaluate C_{cls}, the supremum of critical constants over common-point interval systems. By Theorem 5.1 this requires a microscopic recovery rule; the finite-cell stationarity equations in the (d_i, p_i) coordinates are explicit and are the natural starting point.

Open Problem 10.2 (the multi-star identity)**.** Prove or refute that the multi-star threshold equals the exact critical constant of every common-point cell. The two quantities coincide on every cell tested in this paper, including all six ladder records; a proof (e.g., by exhibiting the hull mixtures of Proposition 4.2 as optimal dual certificates of the multi-star relaxation) would make C_{cls} the limit of a computable threshold sequence and turn Conjecture 1.1 into a limit evaluation.

Open Problem 10.3 (the residual frontier)**.** Determine the exact deletion-star residual frontier B_q at every complement mass q. Theorem 6.6 gives $B_2 = \frac{1}{2}$, and the zero-sum instance of Appendix E shows $B_3 \geq \frac{1}{2} + \frac{139}{458752}$; we conjecture $B_q = \frac{q-1}{q}$, which would give the cell ceiling $2 - \frac{1}{q}$ at every codimension, rising toward the universal constant of Conjecture 1.2 (but incapable of deciding Conjecture 1.1).

Exhaustive exact searches (10^5 random unequal-demand inputs, leader 133/344; the complete $k = 5$ half-grid over all 120 second orders, 1 139 400 instances, maximum 1/3) found nothing above $\frac{1}{2}$, and a separate rank-only campaign produced a $k = 14$ identity/reverse instance of exact residual radius 7/16. These statements do not impose the weighted moment (12). In the stronger zero-sum program, Appendix E gives both a repaired scalar wall above 1/2 and an unrestricted two-moment repair below 1/2; the remaining problem is therefore the full two-equality exchange, not scalar bracketing.

Open Problem 10.4 (the Pareto transfer)**.** Prove $\max_{\mathbf{z}\in U(\mathbf{q})} \mathbf{w}^\mathsf{T}\mathbf{r}(\mathbf{z}) \geq -d_{\max}\|\mathbf{w}\|_1$ for all $\mathbf{w} \geq 0$ on crossed two-arm brooms, i.e. supply the weighted all-or-nothing coalescence described in Section 7.5. This would give (LB) for the exact-two-path class, and with Proposition 3.10 the cost-preserving bound $C \leq 2$ there.

Open Problem 10.5 (three orders with baselines)**.** Construct a three-order family with positive fractional baseline mass on hybrid paths whose constant exceeds 6/5, or prove a two-order ceiling for all splice-closed supports.

Open Problem 10.6 (unconditional constants)**.** Prove *any* unconditional universal constant with exact cost preservation for general graphs; none is currently known. Corollary 6.8 is unconditional but applies only to the rank-$(k-1)$ slab of one class.

References

[1] Tetsuo Asano, Naoki Katoh, Koji Obokata, and Takeshi Tokuyama. Matrix rounding under the L_p-discrepancy measure and its application to digital halftoning. *SIAM Journal on Computing*, 32(6):1423–1435, 2003.

[2] József Beck and Tibor Fiala. “integer-making” theorems. *Discrete Applied Mathematics*, 3(1):1–8, 1981.

[3] Yefim Dinitz, Naveen Garg, and Michel X. Goemans. On the single-source unsplittable flow problem. *Combinatorica*, 19(1):17–41, 1999.

[4] Donald E. Knuth. Two-way rounding. *SIAM Journal on Discrete Mathematics*, 8(2):281–290, 1995. arXiv:math/9504228.

[5] Stavros G. Kolliopoulos and Clifford Stein. Approximation algorithms for single-source unsplittable flow. *SIAM Journal on Computing*, 31(3):919–946, 2002.

[6] Joram Lindenstrauss. A short proof of Liapounoff’s convexity theorem. *Journal of Mathematics and Mechanics*, 15:971–972, 1966.

[7] Siyue Liu and Victor Reis. Weighted chairman assignment and flow-time scheduling. In *17th Innovations in Theoretical Computer Science Conference (ITCS 2026)*, volume 362 of *LIPIcs*, pages 98:1–98:18, 2026. arXiv:2511.18546.

[8] Aleksandr A. Lyapunov. *Sur les fonctions-vecteurs compl‘etement additives.* Bull. Acad. Sci. URSS Sér. Math. 4 (1940) 465–478, 1940.

[9] Mohammed Majthoub Almoghrabi, Martin Skutella, and Philipp Warode. Integer and unsplittable multiflows in series-parallel digraphs. In *Integer Programming and Combinatorial Optimization (IPCO 2025)*, volume 15620 of *Lecture Notes in Computer Science*, pages 427–441, 2025. arXiv:2412.05182.

[10] Maren Martens, Fernanda Salazar, and Martin Skutella. Convex combinations of single source unsplittable flows. In *Algorithms – ESA 2007*, volume 4698 of *Lecture Notes in Computer Science*, pages 395–406, 2007.

[11] Jiří Matoušek. *Geometric Discrepancy: An Illustrated Guide*, volume 18 of *Algorithms and Combinatorics*. Springer, 1999.

[12] Sarah Morell and Martin Skutella. Single source unsplittable flows with arc-wise lower and upper bounds. *Mathematical Programming*, 192(1–2):477–496, 2022. Preliminary version in IPCO 2020, LNCS 12125, 339–350.

[13] Alantha Newman and Aleksandar Nikolov. A counterexample to Beck's conjecture on the discrepancy of three permutations, 2011. arXiv:1104.2922.

[14] Sergey Nikolenko. A planar lower bound $1.1397\ldots > 9/8$ for cost-preserving single-source unsplittable flows. Zenodo preprint, `https://doi.org/10.5281/zenodo.21701162`, 2026.

[15] Dmitry Rybin. A counterexample to Goemans' cost-preserving unsplittable flow conjecture. Announcement, 22 July 2026; record and independent restatement at `https://mathlab.drummerduck.com/p/goemans-unsplittable-flow` and `https://vibemathed.com/problem/dinitz-garg-goemans-unsplittable-flow`, 2026. Seven-vertex instance with $c^{\mathsf{T}}x = 58$ and every congestion-good routing of cost at least 60; not peer-reviewed.

[16] Fernanda Salazar and Martin Skutella. Single-source k-splittable min-cost flows. *Operations Research Letters*, 37(2):71–74, 2009.

[17] Martin Skutella. Approximating the single source unsplittable min-cost flow problem. *Mathematical Programming*, 91(3):493–514, 2002.

[18] Joel Spencer. Six standard deviations suffice. *Transactions of the American Mathematical Society*, 289(2):679–706, 1985.

[19] Chaitanya Swamy, Vera Traub, Laura Vargas Koch, and Rico Zenklusen. Unsplittable cost flows from unweighted error-bounded variants. In *Proceedings of the 2026 Symposium on Simplicity in Algorithms (SOSA)*, pages 512–523. SIAM, 2026. arXiv:2510.21287.

[20] Robert Tijdeman. The chairman assignment problem. *Discrete Mathematics*, 32(3):323–330, 1980.

[21] Vera Traub, Laura Vargas Koch, and Rico Zenklusen. Single-source unsplittable flows in planar graphs. In *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 639–668, 2024. arXiv:2308.02651.

[22] Vera Traub, Laura Vargas Koch, and Rico Zenklusen. Single-source unsplittable flows in planar and bounded-genus graphs. *Mathematical Programming*, 2026.

A The $k = 17$ record instance

The spine has 33 arcs, indexed $0, \ldots, 32$. Terminal i has interval $I_i = [\ell_i, r_i]$, early exit $v_{\ell_i} t_i$ and late exit $v_{r_i+1} t_i$; demands and shares are given as integers over 10^9. All intervals contain arc 16. The demands take nine distinct values and contain no divisibility relation; $\max_i d_i = 1$; and $\sum_i f_i = 6 + 10^{-9}$, so the first cost-good rank is $m = 7$.

i	I_i	$10^9 d_i$	$10^9 f_i$	i	I_i	$10^9 d_i$	$10^9 f_i$
0	$[0, 16]$	1000000000	142527184	9	$[9, 22]$	1000000000	888577908
1	$[1, 18]$	857472816	504331828	10	$[10, 30]$	856589319	68589186
2	$[2, 17]$	1000000000	427438147	11	$[11, 23]$	1000000000	313660163
3	$[3, 21]$	857472816	501514055	12	$[12, 31]$	797836555	213892339
4	$[4, 19]$	997583837	140450372	13	$[13, 29]$	919668646	201063088
5	$[5, 25]$	795420392	208466871	14	$[14, 24]$	1000000000	284850690
6	$[6, 20]$	1000000000	632017755	15	$[15, 28]$	1000000000	878167909
7	$[7, 27]$	849306191	63446846	16	$[16, 32]$	1000000000	344690630
8	$[8, 26]$	906842484	186315030				

Put per-unit cost $1/d_i$ on each early exit arc and zero elsewhere. Then $\mathbf{c}^\mathsf{T}\mathbf{x} = 17 - \sum_i f_i = 11 - 10^{-9}$, so a routing is cost-good exactly when it uses at most ten early paths, i.e. has rank at least seven.

B The $k = 6$ bracket instance

For Proposition 7.9(iii): π is the identity, $\sigma = (3, 5, 0, 4, 2, 1)$,

$$\mathbf{d} = \left(\tfrac{10}{11}, 1, \tfrac{43}{44}, \tfrac{21}{22}, \tfrac{43}{44}, 1\right), \qquad \mathbf{f} = \left(\tfrac{1}{4}, \tfrac{29}{44}, \tfrac{12}{43}, \tfrac{19}{21}, \tfrac{38}{43}, \tfrac{3}{44}\right).$$

All shares are strictly interior; the raw digraph has 24 arcs and exactly two simple paths per terminal.

C The ladder instances, drawn

The families behind the ladder of Figure 1 are all common-point interval systems, and the point of this appendix is to make that visible. For each of $k = 5, 6, 7, 8$ we draw the raw digraph — a spine with two private exit arcs per terminal, so that every terminal has exactly two simple $s \to t_i$ paths — and the same instance as a family of intervals, with the common arc shaded. The $k = 4$ member is the family of [14], already drawn in Figure 4, and the $k = 17$ record is Figures 5 and 6.

Two things are worth reading off these pictures. First, every family has a single common arc, so by Lemma 3.4 each one is a two-permutation prefix system, and the whole apparatus of Section 6 applies to it. Second, the intervals are actually crossing (not nested), which is what puts these systems outside the laminar case and is the mechanism identified in [14].

A caveat on the numbers. Except for $k = 8$, the instance drawn is a concrete rational *member* of its family, so its exact constant is slightly *below* the family's supremum quoted in the ladder; the suprema are approached but not attained. For $k = 8$ the optimum is an algebraic stationary cell rather than a single rational point, so only the combinatorial structure is drawn.

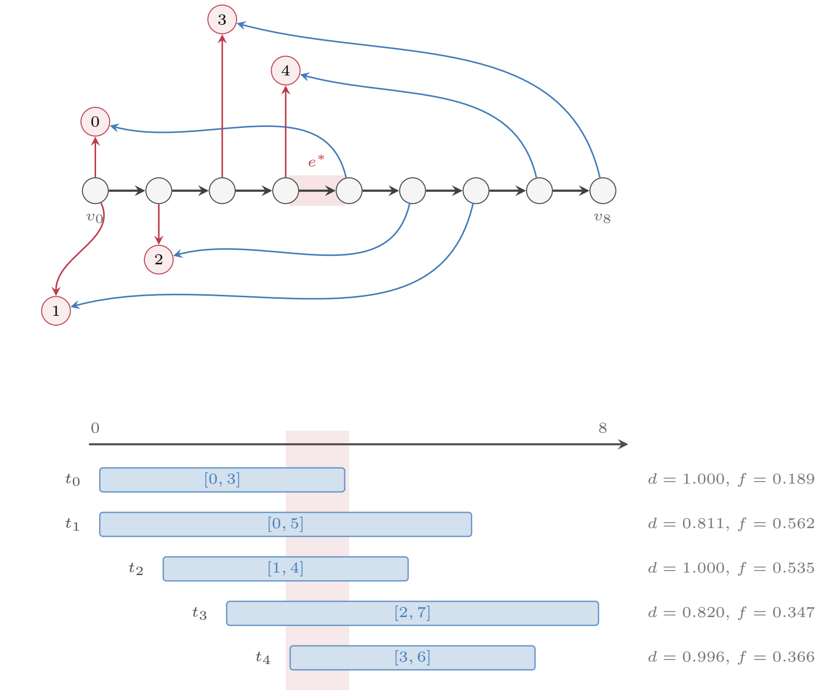

Figure 11: $k = 5$: intervals $[0, 3], [0, 5], [1, 4], [2, 7], [3, 6]$ on a spine of eight arcs; 18 arcs in all, common arc $e^* = 3$. The member drawn has $\sum_i f_i = 2 + 10^{-6}$, first cost-good rank $m = 3$, and exact critical constant $1.16591\ldots$; the family's supremum is $T_* = 1.16798\ldots$, the algebraic number of Section 4.

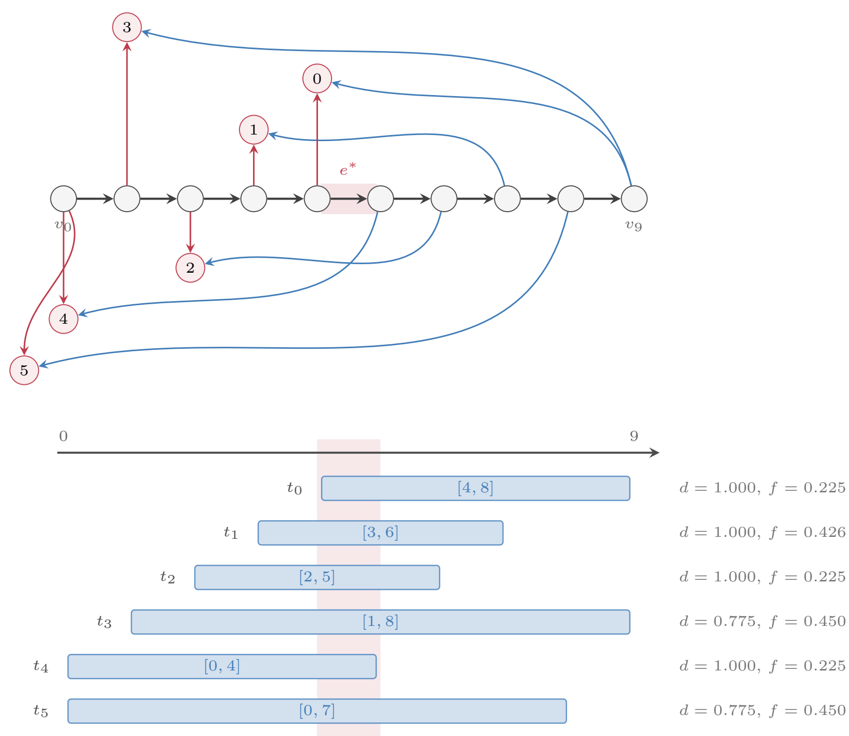

Figure 12: $k = 6$: intervals $[4, 8], [3, 6], [2, 5], [1, 8], [0, 4], [0, 7]$ on a spine of nine arcs; 21 arcs in all, common arc $e^* = 4$. The member drawn is the first Pell point $(p, q) = (49, 20)$, with $\eta = 1/800$, $m = 3$ and exact constant $961/800 = 1.20125$; the family increases to $11 - 4\sqrt{6} = 1.202041\ldots$.

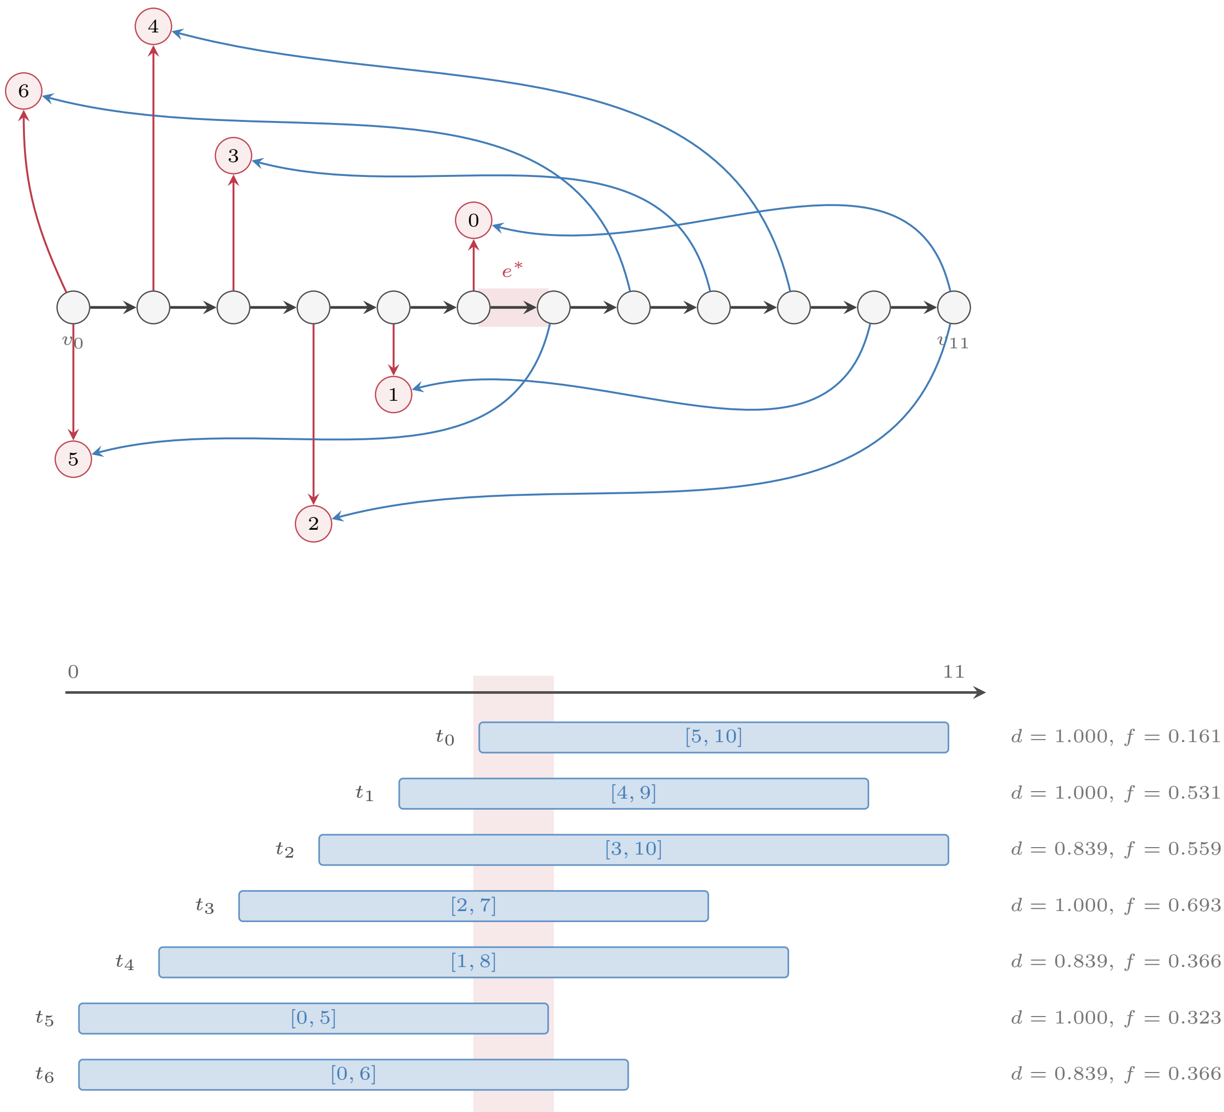

Figure 13: $k = 7$: intervals $[5, 10], [4, 9], [3, 10], [2, 7], [1, 8], [0, 5], [0, 6]$ on a spine of eleven arcs; 25 arcs in all, common arc $e^* = 5$. The member drawn is the first Pell point $(p, Q) = (127, 48)$, with $\eta = 1/6912$, $m = 4$ and exact constant $16705/13824 = 1.208405\ldots$; the family increases to $\frac{13}{2} - 2\sqrt{7} = 1.208497\ldots$.

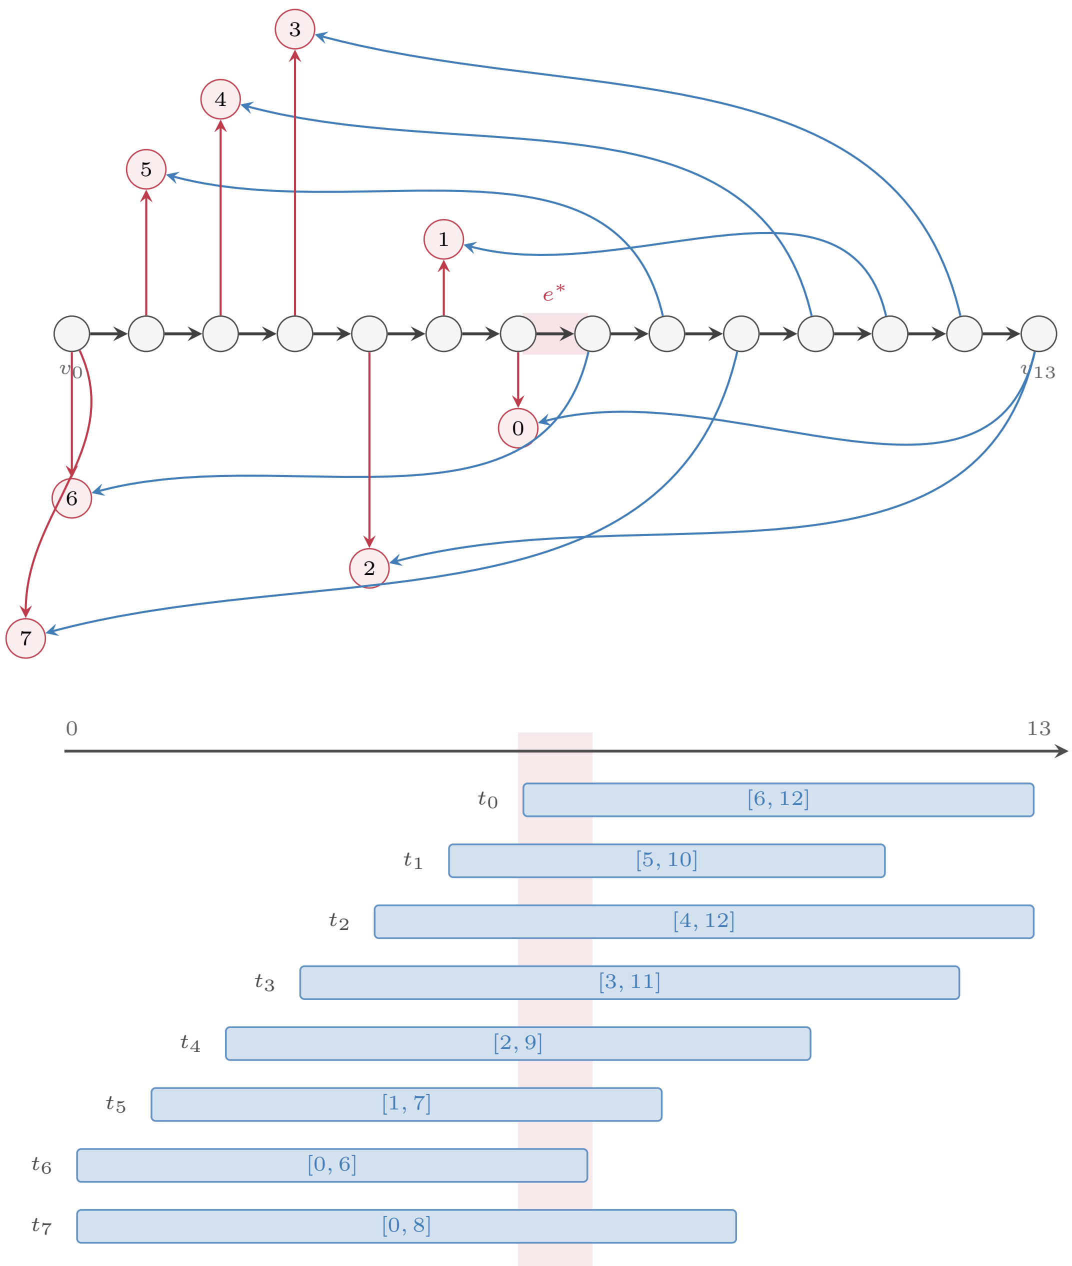

Figure 14: $k = 8$: intervals $[6, 12], [5, 10], [4, 12], [3, 11], [2, 9], [1, 7], [0, 6], [0, 8]$ on a spine of thirteen arcs; 29 arcs in all, common arc $e^* = 6$. Here the optimum is an interior algebraic stationary cell of value ≈ 1.2148 rather than a single rational member, so only the structure is drawn; the demands and shares solve the rank-four equalizer of that cell.

D The multi-star sandwich and the clone experiments

There is a useful exact upper relaxation on a *fixed* fractional point. Let $\rho(\mathbf{z})$ be the maximum overload of a binary routing, put

$$\mathcal{B} = \{\mathbf{z} \in \{0,1\}^k : \rho(\mathbf{z}) \le 1\}, \qquad \mathcal{V}^{\uparrow} = \mathcal{B} \cup \{\mathbf{z} + \mathbf{e}_i : \mathbf{z} \in \mathcal{B},\ z_i = 0\},$$

and let $C^{\uparrow}(\mathbf{f})$ be the first overload level at which $\mathbf{f}$ enters the convex hull of the corresponding vertices of $\mathcal{V}^{\uparrow}$. This is only an instancewise relaxation: it is not a class-wide upper bound.

Proposition D.1 (exact record-chain sandwich)**.** *For each certified record in the following table, the full hull-entry lower bound $C(\mathbf{f})$ and the upward multi-star threshold coincide exactly. Moreover, if m is the first cost-good rank and $\eta = \sum_i f_i - (m-1)$, the displayed exact hull mixture contains a rank-m atom $\mathbf{z}^+$ such that*

$$\mathbf{h} = \frac{\mathbf{f} - \eta \mathbf{z}^+}{1-\eta} \in \{\mathbf{u} \in [0,1]^k : \textstyle\sum_i u_i = m-1\}.$$

For both prefix systems, indeed for every shared row R,

$$\left| \sum_{i \in R} d_i (f_i - h_i) \right| \le \eta \sum_i d_i, \qquad \|\mathbf{f} - \mathbf{h}\|_1 \le \eta k. \tag{11}$$

k	$C(\mathbf{f}) = C^{\uparrow}(\mathbf{f})$	η	$\lvert\mathcal{V}^{\uparrow}\rvert$	hull atoms
7	1.208717093134040213	10^{-9}	96	8
8	1.214766411958000000	10^{-6}	146	9
11	1.255745974752202140	10^{-9}	857	12
13	1.261856423342816434	10^{-9}	2919	14
15	1.271981601327696511	10^{-9}	10016	16
16	1.274864467838565150	10^{-9}	13841	17
17	1.282494797984843521	10^{-9}	19566	18

Proof. This is an exact finite computation. For each row we enumerate all 2^k routings, retain every DGG-good base $\rho(\mathbf{z}) \le 1$, generate every one-bit upward star vertex, and sort their exact rational overloads. Exact linear feasibility first succeeds at the displayed record level; the returned mixture has the displayed barycenter. Below that level all retained vertices have rank at most $m-1$, whereas $\sum_i f_i = m - 1 + \eta$, giving the matching strict separator. The stated rank-m atom is read from the same exact mixture. Substitution gives $\sum_i h_i = m-1$ and exact coordinate checks give $0 \le h_i \le 1$; then $f_i - h_i = \eta(z_i^+ - h_i)$ proves both inequalities in (11). A separately written verifier rebuilds the rows and all routings and checks the barycenters and inequalities with rational arithmetic. □

Figure 15 plots both experiments. The zero width in Proposition D.1 is informative but does not yet squeeze the *class* supremum: its upper endpoint is tied to the same input point as its lower endpoint. We therefore tested the most literal microscopic limiting ansatz, cloning coordinates of the $k = 17$ point and interleaving the clones in the second order.

Proposition D.2 (exhausted finite clone/interleaving ansatz)**.** *Consider the 69 scheduled endpoints consisting of the* 48*-endpoint pilot ($k = 20, 24, 28, 32, 36, 40$, two clone-multiplicity allocations and four interleavings) and one continuation endpoint for every $k = 20, \ldots, 40$. After exact projection*

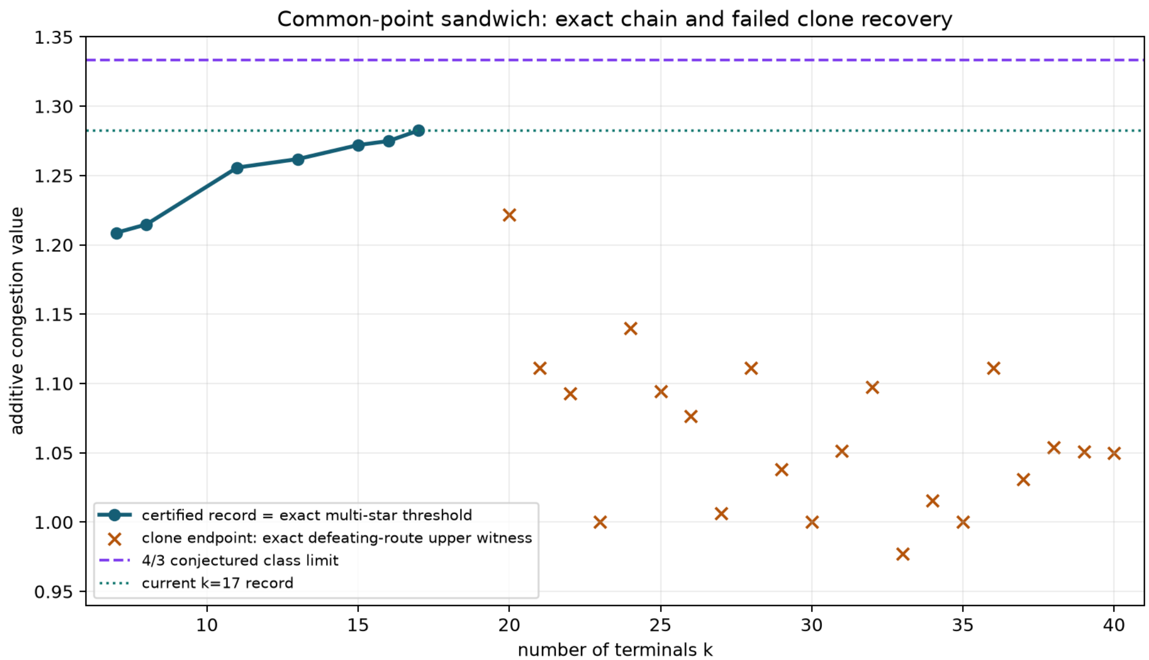

Figure 15: The exact instancewise record-chain sandwich and the finite clone/interleaving experiment. Crosses are defeating-route overloads for tested endpoints, not records or class upper bounds.

to the strict rank face, every endpoint has an explicit cost-good routing whose exact overload is below the $k = 17$ *record. The largest such defeating-route overload is*

$$\frac{1221623116867287239}{10^{18}} = 1.221623116867287239,$$

leaving the exact gap $30435840558778141/(5 \cdot 10^{17})$ *below the record.*

k	defeating overload	k	defeating overload	k	defeating overload
20	1.221623117	27	1.006526579	34	1.015719352
21	1.111388472	28	1.111114892	35	0.999993888
22	1.092611899	29	1.038033935	36	1.111233200
23	0.999998350	30	0.999996205	37	1.030929584
24	1.139959429	31	1.051268589	38	1.054067306
25	1.094425364	32	1.097631574	39	1.050798112
26	1.076606108	33	0.977342446	40	1.049881016

The entries in this last table are *only* exact upper witnesses defeating the tested endpoints. They are not multi-star thresholds, class upper bounds, or new lower-bound records. Thus Proposition D.2 exhausts precisely this finite cloning/interleaving schedule, not all multiplicities, all interleavings, or the continuous topology cells. Its conclusion is nevertheless sharp enough to reject naive cloning of the $k = 17$ geometry as the mechanism behind a limit at 4/3.

E A stronger zero-sum convention, and why it fails

Definition 6.1 is the *rank-normalized* deletion-star program: $\mathbf{1}^\top \mathbf{c} = 1$ forces $\mathbf{1}^\top(\mathbf{u} - \mathbf{f}) = 0$, but for unequal demands it does not force $\mathbf{d}^\top(\mathbf{u} - \mathbf{f}) = 0$. We retain this notion—in particular, Lemma 6.2

and all rank-only results below use exactly Definition 6.1. One may also ask for the following stronger, different program.

Specialize to complement mass three. Put $x_i = 1 - f_i$, choose an omitted pair O, put $S = [k] \setminus O$, and set $\mathbf{u} = \mathbf{1}_S - \mathbf{c}$. Besides $\mathbf{c} \geq 0$, $\operatorname{supp}(\mathbf{c}) \subseteq S$ and $\mathbf{1}^\mathsf{T}\mathbf{c} = 1$, impose the weighted moment

$$\sum_{i \in S} d_i c_i = A_O := \sum_i d_i x_i - \sum_{i \in O} d_i. \tag{12}$$

Then the authoritative residual

$$r_i = d_i(u_i - f_i) = d_i(x_i - \mathbf{1}_{\{i \in O\}} - c_i) \tag{13}$$

satisfies the hard full-row identity $\sum_i r_i = 0$. We denote the resulting minimum maximum-prefix value by B_{zm}. Thus B_{zm} does not replace B: it minimizes over the smaller set satisfying both the rank and weighted moments.

For fixed O, let

$$a_P(O) = \sum_{i \in P} d_i(x_i - \mathbf{1}_{\{i \in O\}}).$$

The exact two-equality primal and dual used in the certificates are

$$\begin{aligned} B_{\mathrm{zm}}(O) = \min_{t,\mathbf{c}} \quad & t \\ \text{s.t.} \quad & a_P(O) - \sum_{i \in P \cap S} d_i c_i \leq t \qquad (P \in \operatorname{Pref}(\pi) \cup \operatorname{Pref}(\sigma)), \\ & \sum_{i \in S} c_i = 1, \qquad \sum_{i \in S} d_i c_i = A_O, \qquad c_i \geq 0. \end{aligned} \tag{14}$$

$$\begin{aligned} B_{\mathrm{zm}}(O) = \max_{\lambda,\eta,\theta} \quad & \sum_P \lambda_P a_P(O) - \eta - \theta A_O \\ \text{s.t.} \quad & \lambda_P \geq 0, \qquad \sum_P \lambda_P = 1, \\ & \eta + \theta d_i \geq d_i \sum_{P \ni i} \lambda_P \qquad (i \in S), \qquad \eta, \theta \in \mathbb{R}. \end{aligned} \tag{15}$$

In particular, O is moment-admissible iff $A_O \in \operatorname{conv}\{d_i : i \in S\}$.

The distinction invalidates one former interpretation of the $k = 14$ calibration. Its legacy scalar credits $c_i = x_i/(3 - x(O))$ enforced $\sum c_i = 1$, but at the claimed pair $O = \{2, 4\}$ their exact full residual was

$$-\frac{10201126375319607}{154016500000000000} \neq 0.$$

The old certificates remain valid records for the rank-only program B, but are quarantined as zero-sum certificates.

There is nevertheless an exact zero-sum scalar wall nearby. Keep the legacy complement vector and right order and replace only

$$d_4 = \frac{5470154188105031}{6017794294600000}.$$

Among all 91 omitted pairs, the scalar proportional vector $c_i = x_i/(3 - x(O))$ satisfies both moments only for $O = \{2, 4\}$. At this unique admissible pair its exact value is

$$\frac{178129713169798083}{308033000000000000} = \frac{1}{2} + \frac{24113213169798083}{308033000000000000}. \tag{16}$$

This refutes the scalar proportional route under the zero-sum convention; it does not refute unrestricted HALF. Solving (14)–(15) for all 91 supports gives 63 admissible supports and 28 exact moment-infeasibility certificates, with

$$B_{\mathrm{zm}} = \frac{431852198179}{2000000000000} < \frac{1}{2} \qquad \text{at } O = \{1, 5\}. \tag{17}$$

Every primary and independent certificate embeds the exact orders, both normalizations, executable code for (13), all 28 ordered prefixes, the 27 deduplicated support-LP rows, and the full-row assertion. The independent recomputation reads only this block and reconstructs all pair and support optima; its support-table hash is `a6eafbdb2b3c8aafd5342636742f1146d21a0de6e2b84afc45af00b465af47fc`. The executable schema, quarantine, repaired proportional certificate, unrestricted audit and independent recomputation are, respectively, `certificates/q3_weighted_half_executable_convention_v1.schema.json`, `certificates/q3_k14_legacy_convention_inconsistency_exact.json`, `certificates/q3_zero_sum_proportional_k14_exact.json`, `certificates/q3_zero_sum_k14_unrestricted_91_support_exact.json`, and `certificates/q3_k14_executable_convention_recompute_exact.json`.

The associated two-moment base-exchange formulation is developed separately in the companion technical note `notes/weighted_half_polymatroid_base_exchange.md`.

Theorem E.1 (four-item zero-sum wall)**.** *The strengthened two-moment statement $B_{\mathrm{zm}} \le \frac{1}{2}$ is false already for $q = 3, k = 4$. For*

$$\begin{aligned} \mathbf{f} &= (13/64,\ 781/3136,\ 481/1024,\ 3919/50176), \\ \mathbf{d} &= (15/16,\ 975/1024,\ 61/64,\ 1), \\ \pi &= (0,1,2,3), \qquad \sigma = (2,3,0,1), \end{aligned}$$

the exact optimum is

$$B_{\mathrm{zm}} = \frac{229515}{458752} = \frac{1}{2} + \frac{139}{458752}.$$

In contrast, after dropping the weighted moment (12)*, the original rank-normalized deletion-star optimum on the same instance is*

$$B = \frac{1606605}{67108864} < \frac{1}{2}.$$

Hence this theorem refutes only the stronger zero-sum program, not the rank-only HALF conjecture of Definition 6.1.

Proof. Here $\mathbf{d}^{\mathsf{T}}\mathbf{f} = 3905/4096$. Of the six selected pairs, exactly $\{0,3\}, \{1,3\}, \{2,3\}$ bracket this demand moment; the remaining three are excluded by exact interval gaps. Both affine equations uniquely determine the state on each feasible pair. Enumerating all eight ordered prefix instances gives radii

$$\frac{2085}{4096}, \qquad \frac{229515}{458752}, \qquad \frac{98393}{196608}.$$

For the unique minimizing pair $\{1, 3\}$,

$$\mathbf{u} = (0, 191/196, 0, 5/196)$$

and

$$\mathbf{r} = (-195/1024,\ 316875/458752,\ -29341/65536,\ -377/7168).$$

The full sum is literally zero and the active row is $L{:}2$. Putting unit dual mass on that row and taking $\eta = 975/49, \theta = -975/49$ matches the primal value. Independent exact code reads the embedded executable convention, verifies its hash, executes both moments and every ordered row, and obtains the same table. Separate exact one-dimensional support enumeration gives the displayed rank-only value. $\square$

Theorem E.2 (exact active three-support cell)**.** *Assume*

$$d_0 < d_1 < d_2 < \mu < d_3 = 1, \qquad \pi = (0,1,2,3), \quad \sigma = (2,3,0,1),$$

and put

$$\delta = 1 - \mu, \qquad t_i = \frac{\delta}{1-d_i} \quad (i < 3).$$

On the cell where the three feasible supports $\{i,3\}$ have active rows $L{:}1, L{:}2, R{:}1$, respectively,

$$\sup B_{\mathrm{zm}} = 4 - 2\sqrt{3}.$$

This is an active-cell theorem; no global $q = 3$ ceiling is asserted.

Proof. The demand inequalities give $0 < t_0 < t_1 < t_2 < 1$. Moment preservation forces $u_i = t_i$ on support $\{i,3\}$ and

$$\frac{f_0}{t_0} + \frac{f_1}{t_1} + \frac{f_2}{t_2} = 1. \tag{Z1}$$

If the three active radii are at least C, their exact formulas imply

$$f_0 \le t_0 - C, \qquad f_0 + f_1 \le t_1 - C, \qquad f_2 \le t_2 - C. \tag{Z2}$$

Maximizing the left side of (Z1) subject to (Z2) first fills f_0, because $1/t_0 > 1/t_1$, and gives

$$1 \le 3 - \frac{C}{t_0} - \frac{t_0}{t_1} - \frac{C}{t_2}.$$

Consequently

$$C \le \frac{2 - t_0/t_1}{1/t_0 + 1/t_2}.$$

Writing $a = t_0/t_2$, then sending $t_1 \uparrow t_2$ and $t_2 \uparrow 1$, reduces the envelope to

$$g(a) = \frac{a(2-a)}{1+a}.$$

Since $g'(a) = (2 - 2a - a^2)/(1+a)^2$, the unique stationary point in $(0,1)$ is $a = \sqrt{3} - 1$, with value $4 - 2\sqrt{3}$.

For recovery, choose rational $(t_0, t_1, t_2) \to (\sqrt{3} - 1, 1, 1)$, set

$$C = \frac{2 - t_0/t_1}{1/t_0 + 1/t_2}, \qquad \mathbf{f} = (t_0 - C,\ t_1 - t_0,\ t_2 - C,\ 1 - t_1 - t_2 + 2C),$$

and take $d_i = 1 - \delta/t_i$, $d_3 = 1$, with rational $\delta \downarrow 0$. The limiting shares

$$(3\sqrt{3} - 5,\ 2 - \sqrt{3},\ 2\sqrt{3} - 3,\ 7 - 4\sqrt{3})$$

are positive, and the strict active branches persist. Thus the three support radii tend to $4 - 2\sqrt{3}$. $\square$

One exact rational member used in the certificates is

$$(t_0, t_1, t_2) = (183/250,\ 999/1000,\ 1999/2000), \qquad \delta = 10^{-5},$$
$$B_{\mathrm{zm}} = \frac{30874533011}{57658950000} = \frac{1}{2} + \frac{2045058011}{57658950000}.$$

The primary and independent artifacts are `certificates/q3_zero_sum_k4_counterexample_exact.json`, `certificates/q3_zero_sum_k4_counterexample_independent_audit.json`, `certificates/q3_zero_sum_k4_three_support_envelope_exact.json`, and `certificates/q3_zero_sum_k4_three_support_envelope_independent_audit.json`. Each embeds schema `ssuf.q3-weighted-half.executable-convention.v1`; the independent cell audit separately isolates the algebraic root in $(1/2, 3/5)$ by a Sturm sequence.

Proposition E.3 (a five-item two-moment record)**.** *The active-cell ceiling of Theorem E.2 is not global. The instance*

$$\begin{aligned}\mathbf{x} &= (1172, 1391, 2461, 4000, 2976)/4000,\\ \mathbf{d} &= (3655, 4000, 4000, 3445, 3605)/4000,\\ \pi &= (0, 1, 2, 3, 4), \qquad \sigma = (2, 1, 3, 4, 0)\end{aligned}$$

has

$$B_{\mathrm{zm}} = \frac{4918713}{8000000} = \frac{1}{2} + \frac{918713}{8000000} = 0.614839125.$$

Exactly six of its ten omitted pairs are moment-admissible, and the unique minimizer is $O = \{2, 3\}$.

Exact certificate. For each omitted pair, the primary verifier enumerates all rational primal vertices of (14); every moment-admissible support receives a matching exact dual of (15), while the other four receive exact demand-interval infeasibility gaps. All ten ordered prefix instances are rebuilt and the two full rows are located by their flags and checked to have literal residual zero.

The independent verifier uses a different reduction. A fixed support has three credit variables and the two affine moments, hence is a rational line segment. It parameterizes that segment from the embedded convention and minimizes the maximum of its ten prefix lines by enumerating the segment endpoints and every pairwise crossing. All ten supports reproduce the primary table and the displayed optimum. The artifacts are `certificates/q3_zero_sum_k5_full_credit_grid_leader_exact.json` and `certificates/q3_zero_sum_k5_full_credit_grid_leader_independent_audit.json`. ☐

The translation distinction is again extreme. If the demand moment is dropped on this five-item instance, separate exact vertex enumeration gives the rank-only value

$$B = -\frac{83277713}{1161600000}$$

at $O = \{1, 2\}$. Proposition E.3 is therefore a record for the stronger projected program only; it is neither an SSUF lower bound nor a counterexample to rank-only W1.

Theorem E.4 (a strict family approaching 2/3)**.** *Let*

$$d = (p, 1, 1, g, e), \qquad \frac{1}{2} < p < 1, \quad 0 < g < e < p,$$

and put

$$C = p\left(1 - \frac{2p - 1}{2p + e(1 - p)/(1 - e)}\right).$$

For orders

$$\pi = (0, 1, 2, 3, 4), \qquad \sigma = (2, 1, 3, 4, 0),$$

define the boundary complement vector

$$\mathbf{x} = (1 - C/p,\ 2C - p,\ C,\ 1,\ 1 - (2C - 1)/e).$$

Restrict to the boundary-equalizer subcell where all five displayed coordinates lie in $[0, 1]$*; the sequence in the proof is eventually in its strict interior. Move a mass* $\tau > 0$ *from* x_4 *to* x_1*. Along parameters tending to* $p, e \to 1$, $1 - p \sim 1 - e$*, followed by* $\tau \downarrow 0$*, the resulting strict instances satisfy*

$$B_{\mathrm{zm}} \longrightarrow \frac{2}{3}.$$

In particular the two-moment program has supremum at least $2/3$*.*

Proof. At $\tau = 0$,

$$\mathbf{1}^{\mathsf{T}}\mathbf{x} = 3, \qquad \mathbf{d}^{\mathsf{T}}(\mathbf{1} - \mathbf{x}) = p + 1.$$

The support omitting $\{3, 4\}$ is moment-feasible at an endpoint and has a low radius. After the perturbation the moment is

$$p + 1 - (1 - e)\tau,$$

so its credit-moment target becomes $1 + (1 - e)\tau > 1$, and that support is strictly infeasible. Six supports persist. At the boundary their fixed prefix lower bounds are

O	03	04	13	14	23	24
fixed value	C	$3C - 1$	C	C	C	$C.$

For parameters near the limit $C \geq 1/2$, hence every persistent support has limiting radius at least C. There is also a uniform upper witness. On support 03, put

$$q = \frac{1 - p}{1 - e} + \tau, \qquad (c_1, c_2, c_4) = (q, 0, 1 - q).$$

Both moments hold. The exposed values are $R{:}1 = R{:}4 = C$, while

$$r_1 = 2C - p - \frac{1 - p}{1 - e} = -\frac{(1 - p)e(1 - pe)}{(1 - e)(2p + e - 3pe)} < 0.$$

All remaining prefixes are at most C. Thus $B_{\mathrm{zm}} \leq C$ for every sufficiently small strict perturbation. Continuity of the other four fixed-support values now gives $B_{\mathrm{zm}} \to C$ as $\tau \downarrow 0$.

Set $a = 1 - p$, $b = 1 - e$, and $r = a/b$. Then $0 < a < b$ and $a < r < 1$, while direct simplification gives

$$C = \frac{(1 - a)(1 - a + r)}{2 - 3a + r}$$

and

$$\frac{2}{3} - C = \frac{(1 - r) + 3a(r - a)}{3(2 - 3a + r)} > 0.$$

Taking $a \downarrow 0$, $r \uparrow 1$ proves the claim. $\square$

For

$$p = 99/100, \quad e = 9899/10000, \quad g = 49/50, \quad \tau = 10^{-6},$$

the exact strict member has

$$B_{\mathrm{zm}} = \frac{282800721371}{427100000000} = 0.662141703046\ldots, \qquad \frac{2}{3} - B_{\mathrm{zm}} = \frac{5797835887}{1281300000000}.$$

The primary all-support primal/dual certificate is `certificates/q3_zero_sum_k5_two_thirds_member_exact.json`; the independent line-envelope and symbolic-family audit is `certificates/q3_zero_sum_k5_two_thirds_member_independent_audit.json`. The latter also obtains the rank-only value

$$-\frac{84705743294271}{8456152900000000}.$$

Thus Theorem E.4 remains completely separated from the rank-only HALF and SSUF questions.

Proposition E.5 (exact global 2/3 target and an order-free wall)**.** *Let*

$$H(\mathbf{x}, \mathbf{d}) = \{\mathbf{w} \in [0,1]^k : \mathbf{1}^\mathsf{T}\mathbf{w} = 3,\ \mathbf{d}^\mathsf{T}\mathbf{w} = \mathbf{d}^\mathsf{T}\mathbf{x}\}$$

and let Q_B be its intersection with all two-order prefix inequalities $\mathbf{d}^\mathsf{T}(\mathbf{x} - \mathbf{w})_P \leq B$. Then $B_{\mathrm{zm}} \leq B$ if and only if Q_B meets a face $v_a = v_b = 1$, equivalently if Q_B has a vertex with at least two upper-box coordinates. Requiring a vertex of the uncut slice $H(\mathbf{x}, \mathbf{d})$ to lie in Q_B is sufficient but stronger.

Moreover, the target $B = 2/3$ cannot be proved by bounding the order-free positive weighted variation. On the mass-two slice

$$\mathbf{x} = (4/5, 3/10, 3/10, 3/10, 3/10), \qquad \mathbf{d} = (1/3, 1, 1, 1, 1),$$

every valid upper-face state has positive weighted variation at least $7/10 > 2/3$.

Proof. A valid omitted-pair state is precisely $\mathbf{w} = \mathbf{1}_{\{a,b\}} + \mathbf{c}$, where $\mathbf{c} \geq 0$, $\mathbf{1}^\mathsf{T}\mathbf{c} = 1$, and $\mathbf{d}^\mathsf{T}\mathbf{c} = \mathbf{d}^\mathsf{T}\mathbf{x} - d_a - d_b$. Hence the valid set is the union of the faces $H(\mathbf{x}, \mathbf{d}) \cap \{v_a = v_b = 1\}$, proving the first equivalence. Whenever one such face meets Q_B, a vertex of their intersection is a vertex of Q_B and retains both active upper-box rows. Conversely such a vertex belongs to a valid face. A vertex of H itself has at most two fractional coordinates and therefore at least two unit coordinates, but an optimum of the convex prefix maximum on a valid face need not be a vertex of H; this is why the uncut-slice statement is strictly stronger.

For the displayed mass-two slice, the moment equations are

$$\mathbf{1}^\mathsf{T}\mathbf{w} = 2, \qquad \mathbf{d}^\mathsf{T}\mathbf{w} = 22/15.$$

Subtracting them gives $v_0 = 4/5$ and $\sum_{i=1}^{4} v_i = 6/5$. Every valid state fixes one of the four unit-demand coordinates, say $v_j = 1$; the other three have total value $1/5$. Their positive deficit is therefore at least

$$\sum_{\substack{i=1 \\ i \neq j}}^{4} (3/10 - v_i) \geq 9/10 - 1/5 = 7/10.$$

Equality occurs at $(4/5, 1, 1/5, 0, 0)$, up to permutation. □

F Deferred proofs

The proofs below are collected here to keep Sections 6 and 7 readable. None of them is used elsewhere in the paper.

Proposition 7.3 (the exact radius $B(\mathbf{d})$)

Proof. The objective $\varphi(\lambda) = \lambda + \sum_i |d_i - \lambda|$ is piecewise linear and convex with slope $1 + \#\{i : d_i < \lambda\} - \#\{i : d_i > \lambda\}$, which changes sign exactly at a lower median, so the minimum is attained there. For the radius-$d_{\max}$ statements, split off the largest demand:

$$\varphi(\lambda) = \bigl(\lambda + |d_{\max} - \lambda|\bigr) + \sum_{i \le k-1} |d_{(i)} - \lambda|.$$

If $\lambda \le d_{\max}$ the bracket equals $d_{\max}$; if $\lambda > d_{\max}$ it equals $2\lambda - d_{\max} > d_{\max}$. Hence $\varphi(\lambda) \ge d_{\max}$ for every λ, so $B(\mathbf{d}) \ge d_{\max}$; and $\varphi(\lambda) = d_{\max}$ forces both $\lambda \le d_{\max}$ and $d_{(1)} = \cdots = d_{(k-1)} = \lambda$. Thus $B(\mathbf{d}) = d_{\max}$ if and only if the $k-1$ smallest demands are equal, in which case $\lambda = d_{(1)}$ attains it.

For the upper bound write $k = 2r$ or $k = 2r+1$ and take $\lambda = d_{(r)}$, a lower median. Expanding the sum,

$$\varphi(d_{(r)}) = d_{(r)}(2r - k) + \sum_{i>r} d_{(i)} - \sum_{i<r} d_{(i)} \ \le\ \sum_{i>r} d_{(i)} \ \le\ (k-r)\, d_{\max},$$

since $2r - k \le 0$ and all $d_{(i)} \ge 0$; and $k - r = \lceil k/2 \rceil$ in both parities. The two-scale formula is φ evaluated at $\lambda = a d_{\max}$ and $\lambda = d_{\max}$, which are the only two candidate medians. $\square$

Theorem 6.9 (sharpness at every codimension)

Proof. Write $\mathbf{h} = \mathbf{u} - \mathbf{f}$ for the residual; $\sum_i h_i = 0$ and $h_i = -f$ at every omitted coordinate, while $h_i = q/k - c_i \le q/k$ at every selected coordinate. Let $H_j = \sum_{i \le j} h_i$ be the cumulative sums along the identity order, $H_0 = 0$, $H_k = 0$. An identity prefix has value H_j and the reverse prefix consisting of the coordinates after position j has value $-H_j$; hence the two prefix systems together control $|H_j|$, and

$$|H_j| \le B_{k,q} \qquad (0 \le j \le k). \tag{18}$$

Lower bound. Across an omitted coordinate at position j the cumulative sum drops by exactly f: $H_j = H_{j-1} - f$. By (18), $f = H_{j-1} - H_j \le 2B_{k,q}$, which is the claim. (If $j = 1$ or $j = k$ the single nonempty row already has value f, which is stronger.)

Necessity of the capacity condition. Assume $f > 0$ and equality $B_{k,q} = f/2$. Then across every omitted coordinate the drop $f = 2B_{k,q}$ must be realized between the extreme admissible values, forcing

$$H_{j-1} = \tfrac{f}{2}, \qquad H_j = -\tfrac{f}{2}.$$

The $q-1$ omitted coordinates split the $n+1 = k-q+1$ selected coordinates into q segments (some possibly empty). The forced boundary values say that the totals of $\mathbf{h}$ over those segments are

$$\tfrac{f}{2},\ \underbrace{f,\ \ldots,\ f}_{q-2 \text{ internal}},\ \tfrac{f}{2}.$$

Since each selected coordinate contributes at most q/k to a segment total, an edge segment needs at least $\lceil (f/2)/(q/k) \rceil = \lceil n/(2q) \rceil$ coordinates and an internal segment at least $\lceil f/(q/k) \rceil = \lceil n/q \rceil$. Summing gives $M_{k,q} \le n+1$. Note this argument allows negative h_i: they can only slow a segment down.

Sufficiency. Suppose $M_{k,q} \le n+1$. Allocate $\lceil n/(2q) \rceil$ selected coordinates to each of the two edge segments, $\lceil n/q \rceil$ to each internal segment, and distribute the remaining selected coordinates arbitrarily; place one omitted coordinate between consecutive segments. Inside a segment of N coordinates with target total $T \in \{f/2, f\}$, choose $0 \le h_i \le q/k$ with $\sum h_i = T$ — possible exactly by the capacity inequality (use increments q/k greedily, then one remainder, then zeros) — and set $c_i = q/k - h_i \geq 0$ on selected coordinates, $c_i = 0$ on omitted ones. The segment totals sum to $(q-1)f$, so

$$\sum_{\text{selected}} c_i = (k-q+1)\frac{q}{k} - (q-1)f = 1,$$

so these are valid credits. The cumulative sum now starts at 0, rises monotonically to $f/2$, drops to $-f/2$ at each omitted coordinate, rises monotonically through each internal segment back to $f/2$, and finally returns from $-f/2$ to 0 (Figure 9). Every partial sum lies in $[-f/2, f/2]$, so every identity and reverse prefix error is at most $f/2$.

Arithmetic form. Write $n = aq + \rho$ with $0 \leq \rho < q$. If $\rho = 0$ then $M_{k,q} = 2\lceil a/2 \rceil + (q-2)a \leq aq + 1 = n + 1$, so the criterion always holds; this covers $k = mq$ with $a = m - 1$, giving $B_{mq,q} = f/2 = \frac{1}{2} - \frac{1}{2m}$. If $\rho > 0$ then $\lceil n/q \rceil = a + 1$ and $2\lceil n/(2q) \rceil$ equals $a + 2$ for even a and $a + 1$ for odd a; substituting into $M_{k,q} \leq aq + \rho + 1$ gives the criterion in the equivalent form $a + \rho + 1 \geq q$ (a even) or $a + \rho + 2 \geq q$ (a odd). ☐

Theorem 6.10 (the two-scale variant)

Proof. Since $\sum_i h_i = 0$, the weighted full residual is

$$T = \sum_i d_i h_i = \sum_i h_i + \sum_i (d_i - 1)h_i = -\tfrac{1}{3}h_0. \tag{19}$$

Lower bound. Fix any support and credits, and let j be an omitted coordinate. The weighted identity prefix strictly before j and the weighted reverse prefix strictly after j have errors summing to $T - d_j h_j = T + d_j f$, so

$$B \;\geq\; \frac{T + d_j f}{2}.$$

First note that $h_0 \leq q/k = 1/m$ *unconditionally*: if coordinate 0 is selected then $h_0 = q/k - c_0 \leq q/k$, and if it is omitted then $h_0 = -f \leq 0 < q/k$. By (19) this gives $T = -\frac{1}{3}h_0 \geq -\frac{1}{3m}$ in every case.

Two cases. If some omitted j has $d_j = 1$, then $B \geq \frac{T+f}{2} \geq \frac{f}{2} - \frac{1}{6m}$. If instead every omitted coordinate has $d_j \neq 1$, then — since coordinate 0 is the only one with $d_0 \neq 1$ — the star omits exactly $j = 0$ and $q - 1 = 1$, i.e. $q = 2$; then $h_0 = -f$, so $T = f/3$ by (19) and

$$B \;\geq\; \frac{f/3 + (2/3)f}{2} = \frac{f}{2},$$

which is even stronger. In both cases $B \geq \frac{f}{2} - \frac{1}{6m} = \frac{1}{2} - \frac{1}{2m} - \frac{1}{6m}$.

Upper bound. Use the sawtooth star of Theorem 6.9 and place the low-demand coordinate at a *selected* coordinate whose residual h_i is zero. Such a coordinate exists: the minimal segment allocation uses $M_{k,q}$ of the $n+1$ selected coordinates, and the surplus $n + 1 - M_{k,q}$ is m when m is odd and $m - 1$ when m is even, hence positive for $m \geq 2$; surplus coordinates receive $h_i = 0$ in the greedy filling. Multiplying a zero-residual coordinate by $2/3$ changes no prefix error, so the same allocation still has radius at most $f/2$. (For $m = 1$ we have $f = 0$ and both bounds are trivial.) ☐

G Verification artifacts

The repository accompanying this paper[1] contains, for each numbered claim, a program that rebuilds the object from its primitive data and recomputes the claimed quantity; a machine-readable certificate; and, where the claim is a search, the complete scope record. The verification summarized in Section 9 is a second, independent layer written from the statements in this paper.

[1] https://github.com/snikolenko/unsplittable-flows