Additional Notes on "Fair Fares for Vehicle Sharing Systems"

Hyemi Kim
Columbia University
hk3181@columbia.edu

First version: November 19, 2026

Abstract

These notes extend three results of Elmachtoub and Kim (2026) (hereafter, the paper). Section 1 characterizes the effect of permitting *a little* price differentiation when the platform starts from perfect price fairness. This is the endpoint opposite to the one the paper characterizes, and it admits five sign regimes rather than four, one of which is a Pareto improvement. Section 2 asks which of the paper's repositioning results survive when the platform operates more than one vehicle. With more than one vehicle, strong price fairness drives the platform to abandon the location whose riders are willing to pay less, while access fairness, as with one vehicle, makes one location gain and the other lose once vehicles are repositioned. Section 3 supplies a tractable relaxation, with a performance guarantee, for price-fair pricing on a general network, a problem the paper leaves open.

Notation. We follow the paper. Locations are indexed $i \in [N]$ and the fleet has $K \geq 1$ vehicles. On arc (i, j) customers arrive at rate Λ_{ij} and buy when their valuation exceeds the posted price p_{ij}, so the demand rate is $\lambda_{ij} = \Lambda_{ij}\bar{F}_{ij}(p_{ij})$, where $\bar{F}_{ij}$ is the tail of the valuation distribution. The stationary availability Π_i is the probability that a vehicle is available at i, and the realized flow on (i, j) is $\Pi_i\lambda_{ij}$. Revenue, consumer surplus per arrival at location i, and welfare are $\mathcal{R}$, S_i, and $\mathcal{W} = \mathcal{R} + \sum_i \Lambda_i S_i$. Fairness is indexed by $\alpha \in [0, 1]$, with $\alpha = 0$ imposing nothing and $\alpha = 1$ perfect fairness. In the two-location sections we drop the second index and write $\Lambda_i, b_i, p_i, \lambda_i$. Demand is linear there, i.e., valuations are uniform on $[0, b_i]$, so $\lambda_i = \Lambda_i(1 - p_i/b_i)^+$.

1 A little price differentiation

Uniform fares are common in shared mobility. Many bike-sharing systems charge the same rate everywhere, and microtransit services such as Via in Jersey City and Metro Micro in Los Angeles charge the same fare regardless of origin and destination. The paper asks what is lost when a platform that would like to differentiate is pushed toward such a uniform fare. This section asks the reverse question. Starting from a uniform fare, what happens when the platform is allowed to differentiate *a little*? Riders are usually assumed to dislike price differentiation, but is there always a loser, or can every party, the platform included, come out ahead?

In vehicle sharing, a price difference need not be read as exploiting information about customers. It can be an operational instrument. Citi Bike's *Bike Angels* program rewards riders who move a bicycle to a location where bicycles are scarce with ride credit, which lowers the effective price of the trips the platform wants. Proposition 1.1 identifies when such a discount makes everyone better off, when it merely redistributes, and when it is pure extraction.

1.1 Model

We use the paper's two-location benchmark: $N = 2$ locations, $K = 1$ vehicle, zero travel times, and linear demand, with $b_1, b_2 > 0$ and $\Lambda_1, \Lambda_2 > 0$. The vehicle waits at its location until a customer buys a trip to the other location, so it alternates between the two locations and a round trip takes $1/\lambda_1 + 1/\lambda_2$ time on average. The *round-trip rate* is

$$g(\boldsymbol{p}) = \frac{\lambda_1(\boldsymbol{p})\lambda_2(\boldsymbol{p})}{\lambda_1(\boldsymbol{p}) + \lambda_2(\boldsymbol{p})}, \qquad \boldsymbol{p} = (p_1, p_2) \in [0, b_1] \times [0, b_2],$$

with $g := 0$ at the corner $\boldsymbol{p} = (b_1, b_2)$, where both demand rates vanish, which makes g continuous on the box since $g \le \min\{\lambda_1, \lambda_2\}$. Trips leave each location at rate g, however many customers arrive there, so revenue, per-arrival consumer surplus at location i, and welfare are

$$\mathcal{R}(\boldsymbol{p}) = g(\boldsymbol{p})\,(p_1 + p_2), \qquad S_i(\boldsymbol{p}) = \frac{b_i \lambda_i}{2\Lambda_i^2}\, g(\boldsymbol{p}), \qquad \mathcal{W}(\boldsymbol{p}) = \mathcal{R}(\boldsymbol{p}) + \Lambda_1 S_1(\boldsymbol{p}) + \Lambda_2 S_2(\boldsymbol{p}). \tag{1.1}$$

Price fairness at level α is the constraint $|p_1 - p_2| \le (1 - \alpha)\Delta$, where $\Delta := |p_2^* - p_1^*|$ is the price gap the unconstrained platform would choose. The platform maximizes revenue subject to this constraint, and $\mathcal{W}$, S_1, and S_2 are read off at its choice $\boldsymbol{p}(\alpha)$. The choice is unique at every level by Lemma A.1(i) in Appendix A, which contains all proofs for this section, so the unconstrained optimum $\boldsymbol{p}^*$, Δ, and $\Phi(\alpha) := \Phi(\boldsymbol{p}(\alpha))$ are well defined for $\Phi \in \{\mathcal{R}, \mathcal{W}, S_1, S_2\}$. The object of study is the first-order effect of relaxing perfect fairness,

$$\mathcal{D}\Phi := \lim_{\varepsilon \downarrow 0} \frac{\Phi(1 - \varepsilon) - \Phi(1)}{\varepsilon}, \tag{1.2}$$

so that $\mathcal{D}\Phi > 0$ means that Φ rises when a little differentiation is permitted.

1.2 Result

Proposition 1.1 (A little price differentiation)**.** *Let $N = 2$, $K = 1$, demand be linear, and, without loss of generality, $b_2 \ge b_1$. Suppose $\Delta > 0$. Then:*

(i) Revenue rises: $\mathcal{D}\mathcal{R} > 0$.

(ii) One fare falls and the other rises. Location 1's fare is the one that falls if and only if

$$(3b_2 - b_1)\sqrt{\Lambda_2 b_1} > (3b_1 - b_2)\sqrt{\Lambda_1 b_2}. \tag{1.3}$$

(iii) Exactly one of the following five regimes holds, and each regime occurs for some parameters $(\Lambda_1, \Lambda_2, b_1, b_2)$ and persists under small changes to them:

(a)	$\mathcal{D}\mathcal{W} > 0$,	$\mathcal{D}S_1 > 0$,	$\mathcal{D}S_2 \ge 0$,
(b)	$\mathcal{D}\mathcal{W} \ge 0$,	$\mathcal{D}S_1 > 0$,	$\mathcal{D}S_2 < 0$,
(c)	$\mathcal{D}\mathcal{W} < 0$,	$\mathcal{D}S_1 \ge 0$,	$\mathcal{D}S_2 < 0$,
(d)	$\mathcal{D}\mathcal{W} < 0$,	$\mathcal{D}S_1 < 0$,	$\mathcal{D}S_2 < 0$,
(e)	$\mathcal{D}\mathcal{W} > 0$,	$\mathcal{D}S_1 < 0$,	$\mathcal{D}S_2 > 0$.

Location 1's fare falls in regimes (a)–(d), i.e., exactly when (1.3) *holds, and location 2's fare falls in regime (e).*

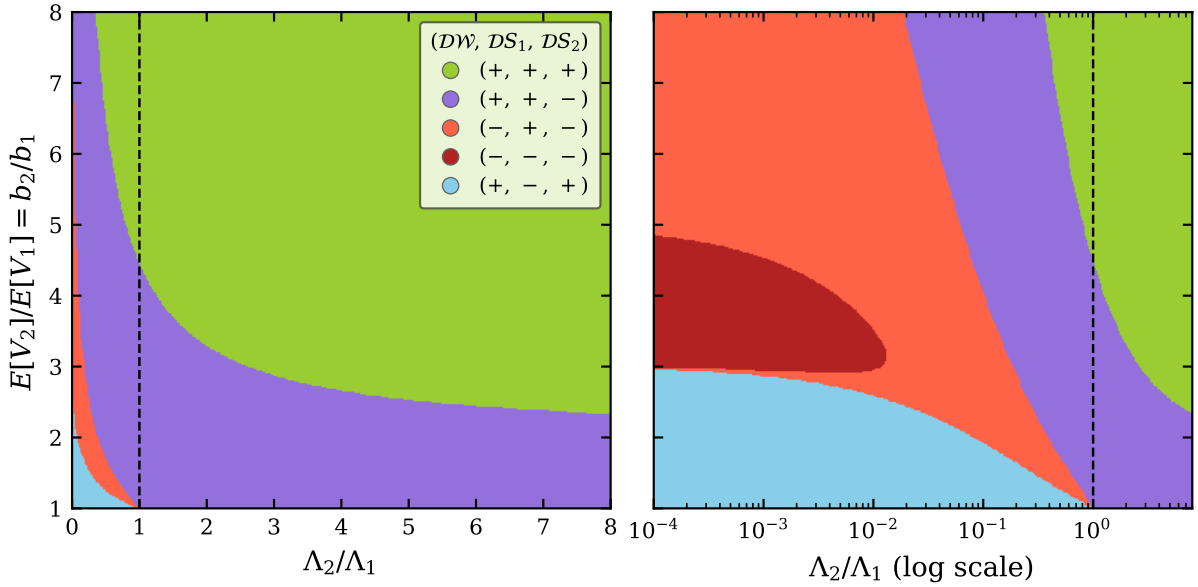


Figure 1: The five regimes of Proposition 1.1 in the plane $(\Lambda_2/\Lambda_1,\, b_2/b_1)$. Each point is coloured by its regime. The right panel repeats the map on a logarithmic horizontal axis, which is the only scale on which regime (d) is visible. The dashed line is $\Lambda_2/\Lambda_1 = 1$.

Part (ii) says which location gets the discount, namely the one where a discount shortens the vehicle's round trip. With equal valuations, (1.3) reads $\Lambda_2 > \Lambda_1$, so the discount goes to the quieter location, where the vehicle waits. This is the *Bike Angels* logic. When $b_2 \geq 3b_1$, location 1 is discounted whatever the traffic, because location 2's riders value the trip so far above the uniform fare that a surcharge barely deters them.

Price differentiation also need not hurt anyone. In regime (a), location 2's riders pay a higher fare and are still better off, because the vehicle returns to them more often, so perfect price fairness blocks a change that every rider and the platform would welcome. Figure 1 shows that this happens when $b_2 > 2b_1$ and location 2 has enough traffic.

2 Repositioning with multiple vehicles

This section studies the effect of price fairness and access fairness on the two-location network of the paper, $N = 2$, with repositioning and $K \geq 1$ vehicles. Proofs are in Appendix B.

2.1 Model

The setting is that of Section 1, now with $K \geq 1$ vehicles. The platform chooses demand rates $\lambda_i \in [0, \Lambda_i]$, i.e., prices $p_i(\lambda_i) = b_i(1 - \lambda_i/\Lambda_i)$, where any price $p_i \geq b_i$ gives $\lambda_i = 0$, and *repositioning rates* $\psi_i \geq 0$ at which idle vehicles are sent empty out of location i, at cost $c_i \geq 0$ per move, with $c_1 + c_2 > 0$. Let k be the number of vehicles at location 1 and write

$$\phi_1 := \lambda_1 + \psi_1, \qquad \phi_2 := \lambda_2 + \psi_2, \qquad \rho := \phi_2/\phi_1.$$

A vehicle leaves location 1 at rate ϕ_1 whenever $k \geq 1$ and leaves location 2 at rate ϕ_2 whenever $k \leq K-1$, so k is a birth–death chain with constant rates. With $\Pi_1 = \mathbb{P}(k \geq 1)$ and $\Pi_2 = \mathbb{P}(k \leq K-1)$,

$$\mathcal{R}_K(\lambda, \psi) = \Pi_1\big(\lambda_1 p_1(\lambda_1) - c_1\psi_1\big) + \Pi_2\big(\lambda_2 p_2(\lambda_2) - c_2\psi_2\big). \tag{2.1}$$

Write $Z_K(\rho) := \sum_{j=0}^{K} \rho^j$.

Lemma 2.1 (Availabilities)**.** *If $\phi_1, \phi_2 > 0$ then $\pi(k) = \rho^k / Z_K(\rho)$, so*

$$\Pi_1 = 1 - \frac{1}{Z_K(\rho)}, \qquad \Pi_2 = 1 - \frac{\rho^K}{Z_K(\rho)}.$$

Moreover:
(i) (flow balance) *$\phi_1\Pi_1 = \phi_2\Pi_2$, i.e., $\Pi_1 = \rho\Pi_2$.*
(ii) Π_1 increases from 0 to 1 and Π_2 decreases from 1 to 0 as ρ runs over $(0, \infty)$.
(iii) if $\phi_1 = 0$ or $\phi_2 = 0$ then $\mathcal{R}_K = 0$.

By Lemma 2.1 the availabilities depend on the policy only through ρ. Without repositioning $\rho = \lambda_2/\lambda_1$ is tied to the prices, and repositioning frees it, since sending vehicles empty out of location 1 raises ϕ_1 and lowers ρ, so Π_1 falls and Π_2 rises. Access and surplus are $A_i = \Pi_i(\rho)\lambda_i/\Lambda_i$ and $S_i = \Pi_i(\rho) b_i \lambda_i^2/(2\Lambda_i^2)$, so price fairness, which constrains $p_i(\lambda_i)$, and access fairness both depend on the policy only through (λ, ρ), as the next proposition requires.

Proposition 2.2 (One-sided repositioning)**.** *Fix $K \geq 1$ and any (possibly empty) set of fairness constraints on (λ, ρ). If (λ, ψ) is feasible with $\lambda \neq 0$ and $\psi_1\psi_2 > 0$, there is a feasible (λ, ψ') with the same ρ, hence the same Π_i, A_i, and S_i, with $\psi_1'\psi_2' = 0$ and strictly larger revenue. Every optimal solution therefore has $\psi_1^*\psi_2^* = 0$.*

This is the paper's result for $K = 1$. Moving vehicles empty in both directions only produces a ρ that moving them in one direction also produces, at a lower cost. The platform's problem is thus three-dimensional in $(\lambda_1, \lambda_2, \rho)$, with $\psi_1 = \lambda_2/\rho - \lambda_1$ and $\psi_2 = 0$ if $\rho \leq \lambda_2/\lambda_1$, and $\psi_1 = 0$ and $\psi_2 = \rho\lambda_1 - \lambda_2$ if $\rho \geq \lambda_2/\lambda_1$. Which direction is optimal depends on the parameters.

One-sided repositioning also simplifies the revenue. If vehicles pile up at location 1 and are sent back empty ($\psi_2 = 0$), flow balance gives $\Pi_1\psi_1 = \Pi_2\lambda_2 - \Pi_1\lambda_1$, i.e., empty moves $1 \to 2$ per unit time equal trips $2 \to 1$ minus trips $1 \to 2$, and substituting into (2.1) gives

$$\mathcal{R}_K = \lambda_1\big(p_1(\lambda_1) + c_1\big)\Pi_1(\rho) + \lambda_2\big(p_2(\lambda_2) - c_1\big)\Pi_2(\rho). \tag{2.2}$$

Each trip $1 \to 2$ saves one empty move and is worth $p_1 + c_1$, while each trip $2 \to 1$ creates one and is worth $p_2 - c_1$. Symmetrically, if $\psi_1 = 0$, the trips are worth $p_1 - c_2$ and $p_2 + c_2$. As in the paper, the platform may also send every vehicle that arrives at location 1 back at once, i.e., $\psi_1 = \infty$ and $\rho = 0$. Then $\Pi_1 = 0$, $\Pi_2 = 1$, and (2.2) gives $\mathcal{R}_K = \lambda_2\big(p_2(\lambda_2) - c_1\big)$. The same holds for location 2 with $\rho = \infty$.

2.2 Price fairness

Call a policy with $\lambda_1 = 0$ an *abandonment of location 1*, i.e., that location generates no trips, and vehicles landing there are only pushed back.

Proposition 2.3 (Price fairness and abandonment). *Fix $K \geq 1$, let $b_2 > c_1$ and $b_1 \leq \frac{1}{2}(b_2+c_1)$, and write $\mathcal{R}_K^*(\alpha)$ for the optimal revenue under price fairness at level α, as in Section 1.*

(a) Over all policies with $\lambda_1 = 0$, the largest revenue is $\overline{V} := \Lambda_2(b_2-c_1)^2/(4b_2)$. It is attained by $p_1 = p_2 = (b_2+c_1)/2$ with every vehicle that arrives at location 1 sent back at once ($\rho = 0$), and this policy is price fair at every α.

(b) $\mathcal{R}_K^*(\alpha) \geq \overline{V}$ for all α and $\mathcal{R}_K^*$ is nonincreasing, so $\mathcal{I}_K := \{\alpha : \mathcal{R}_K^*(\alpha) = \overline{V}\}$ is empty or an interval with right endpoint 1. On $\mathcal{I}_K$ the policy of (a) is optimal, and none of its prices, availabilities, accesses, or surpluses depends on α.

(c) If $K = 1$ and abandonment is optimal at $\alpha = 0$, then $\mathcal{I}_1 = [0,1]$. If $K \geq 2$, then $\mathcal{R}_K^*(0) > \overline{V}$, so $0 \notin \mathcal{I}_K$ and price fairness lowers revenue on $[0,1] \setminus \mathcal{I}_K$.

With location 1 abandoned the network is a single market, in which every paid trip $2 \to 1$ is followed by an empty move $1 \to 2$, so for every K the platform is a monopolist at location 2 with cost c_1 per ride. Since $p_2 \geq b_1$, posting $p_1 = p_2$ sells no rides at location 1, so this policy is fair for free.

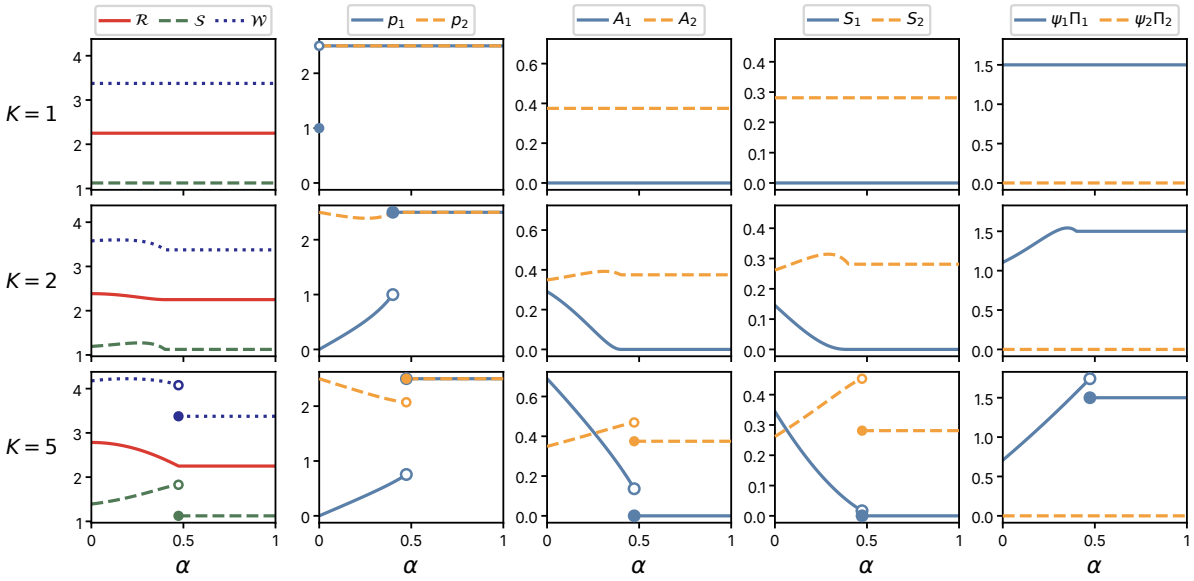


Figure 2: Price fairness with repositioning, $(\Lambda_1, \Lambda_2, b_1, b_2, c_1, c_2) = (1, 4, 1, 4, 1, 1)$, for $K = 1$, 2, and 5 (rows). Location 1 is abandoned at every α for $K = 1$, from $\alpha = 0.4$ for $K = 2$, and from $\alpha \approx 0.47$ for $K = 5$. A closed location 1 sells no rides at any price, and we set $p_1 = p_2$.

Figure 2 illustrates Proposition 2.3 with $\overline{V} = 2.25$. With one vehicle (top row), abandonment is optimal without fairness and stays optimal at every α, so price fairness changes nothing, as in the paper. With $K \geq 2$ vehicles one vehicle can sit at location 1 while another waits at location 2, so serving location 1 costs location 2 almost no availability ($\Pi_1(\rho) \approx \rho$ while $1 - \Pi_2(\rho) \approx \rho^K$ for small ρ), and without fairness the platform serves location 1 at a low price and lets its riders bring vehicles back to location 2. Price fairness pushes the two prices together, so revenue falls and access and surplus shift from location 1 to location 2, until the platform stops serving location 1 and posts the same price at both locations, from which point the constraint costs nothing more. The switch is gradual for $K = 2$ and abrupt for $K = 5$.

Either way, the prices end up equal, but location 1 has no service at all, and location 2 loses the lower price it had under weaker fairness, so strong price fairness ends with no service at the location whose riders are willing to pay less.

2.3 Access fairness

Access fairness requires $|A_1 - A_2| \leq (1-\alpha)\Delta$, where Δ is the gap at the unconstrained optimum. The platform can narrow the gap in two ways. It can serve fewer riders at both locations, or it can send vehicles empty toward the location with lower access, which raises access there and lowers it at the other location but costs c_i per move.

Proposition 2.4 (Access fairness)**.** *Fix $K \geq 1$.*

(a) (Without repositioning, both locations lose.) *Let $\Lambda_1 < \Lambda_2$, and suppose that the unconstrained optimum λ^* uses no repositioning and has $\lambda_i^* \in (0, \Lambda_i)$. Among policies without repositioning, the best one is $(1-\alpha)\lambda^*$. Along it the prices rise, $p_i = (1-\alpha)p_i^* + \alpha b_i$, the availabilities stay constant, and access, surplus, and revenue fall at both locations, to 0 at $\alpha = 1$.*

(b) (With repositioning, the location with lower access gains.) *Consider an interval of $\alpha < 1$ on which the optimum repositions, has $\lambda_i \in (0, \Lambda_i)$, and makes the access constraint bind. As α increases on it, access and surplus rise at the location with lower access and fall at the other location. At $K = 1$ the prices stay fixed and only ρ moves. For $K \geq 2$ the price at the location with lower access falls and the other price rises.*

(c) (Perfect access fairness.) *Under the assumptions of (a), the best revenue at $\alpha = 1$ is positive if and only if*

$$\Lambda_1 b_1 + \Lambda_2 b_2 > c_1(\Lambda_2 - \Lambda_1), \tag{2.3}$$

and then the optimum repositions for all α close to 1.

Under the assumptions of (a), location 1, the location with fewer potential riders, has higher access whatever the willingness to pay, because without repositioning every vehicle that leaves a location must come back, so the same number of trips leaves each location per unit time, which is a larger share of the smaller market. Part (c) says when the platform can reach equal access without shutting down. With equal access, for every unit of access location 1 sends Λ_1 rides and location 2 sends Λ_2 rides, so $\Lambda_2 - \Lambda_1$ vehicles must return to location 2 empty at cost c_1 each, and (2.3) says that the fares, at most $\Lambda_1 b_1 + \Lambda_2 b_2$, can cover this cost.

Figure 3 illustrates Proposition 2.4 on the parameters of the paper, which satisfy the assumptions of (a) and (2.3). While the constraint is weak, the platform does not reposition and closes the gap by raising both prices, so riders at both locations lose, as in (a). Once the constraint is strong enough, and larger fleets reach that point earlier, the platform sends vehicles empty from location 1 to location 2, and from then on location 2, the location with lower access, gains (though not back to its level without fairness) and location 1 loses, as in (b), while revenue stays positive at $\alpha = 1$, as in (c). In this instance K changes only the channel, availability alone with one vehicle and prices as well with more. Who gains can depend on K, however. In the instance of Figure 2 the optimum repositions already at $\alpha = 0$, so riders at both locations never lose together, and the location with lower access, which is the one that gains, is location 1 for $K \leq 2$ and location 2 for $K = 5$.

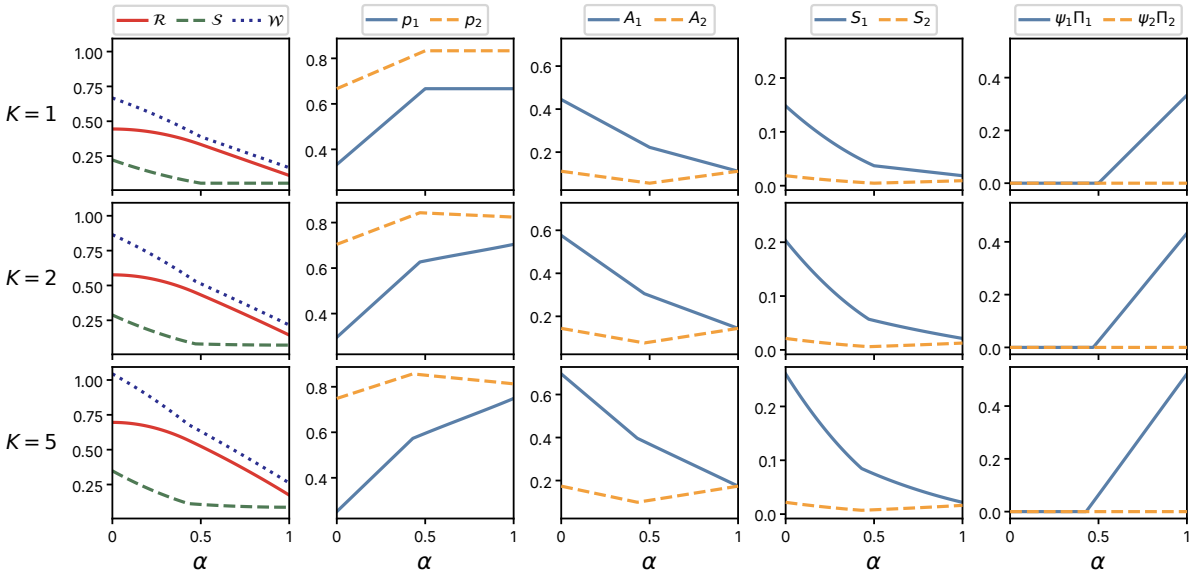


Figure 3: Access fairness with repositioning, $(\Lambda_1, \Lambda_2, b_1, b_2, c_1, c_2) = (1, 4, 1, 1, 1, 1)$, for $K = 1$, 2, and 5 (rows).

3 A price-fairness heuristic for general networks

This section returns to a general network with N locations and treats trip-based price fairness. The convex relaxation that the paper uses for access fairness does not apply here, since it balances the demand rates themselves at every location, which amounts to giving every location the same availability, and under price fairness this can be arbitrarily bad. We therefore keep the availability Π_i of each location as a decision variable, which gives the relaxation (3.1). Section 3.1 shows that this nonconvex problem becomes a one-dimensional family of convex programs once the highest normalized price $\bar{p}$ is fixed, Section 3.2 turns any feasible point of such a program into an α-fair policy with a revenue guarantee, Section 3.3 finds $\bar{p}$ by a grid search with a provable optimality gap, Section 3.4 illustrates the bounds, and Section 3.5 treats origin-based price fairness. All proofs are in Appendix C.

Setting. Let $\mathcal{A} \subseteq [N] \times [N]$ be the set of arcs, i.e., the pairs with $\Lambda_{ij} > 0$, where $i = j$ is allowed. Customers for arc (i, j) arrive at i as a Poisson process with rate Λ_{ij} and accept the static price p_{ij} with probability $\bar{F}_{ij}(p_{ij})$. An accepting customer who finds a vehicle available at i drives it to j, and every other customer is lost. Every trip from i to j, loaded or empty, takes an independent random time with mean $1/\mu_{ij} \in (0, \infty)$ and a distribution that depends only on (i, j). Fairness compares the normalized prices p_{ij}/d_{ij} for given $d_{ij} > 0$, such as trip distances, which lie in $[0, p_{\max}]$, so $p \in \mathcal{P} := \prod_{(i,j)\in\mathcal{A}} [0, d_{ij} p_{\max}]$. Let $b_{ij} := \sup\{p \geq 0 : \bar{F}_{ij}(p) > 0\} \in (0, \infty]$ be the maximum support of the valuation distribution on arc (i, j). A price $p_{ij} \geq b_{ij}$ *closes* the arc, since nobody buys, and it still enters the fairness constraint.

Repositioning. Every vehicle that completes a trip at location i, loaded or empty, is sent empty to a location $j \neq i$ with probability θ_{ij} and otherwise stays at i, where $\mathcal{V} := \{(i, j) \in [N] \times [N] : i \neq j\}$ and $\theta \in \Theta := \{\theta \in [0, 1]^{\mathcal{V}} : \sum_j \theta_{ij} \leq 1 \ \forall i\}$. Each repositioning move from i to j costs $c_{ij} \geq 0$. As in the supply redirection of Banerjee et al. (2022), whether a vehicle is

repositioned is decided at the moment it completes a trip, not while it is idle.

Revenue. For a static policy (p, θ) and an initial placement $\mathbf{s}$ of the K vehicles, at locations or at the start of trips, let

$$\mathcal{R}_K(p, \theta \mid \mathbf{s}) := \liminf_{t\to\infty} \frac{1}{t}\, \mathbb{E}_{\mathbf{s}}\big[\text{fares collected} - \text{repositioning costs incurred in } [0,t]\big].$$

Since the platform chooses where its vehicles start, let $\mathcal{R}_K(p,\theta) := \max_{\mathbf{s}} \mathcal{R}_K(p,\theta \mid \mathbf{s})$. Given $\alpha \in [0,1]$, let $W_\alpha := (1-\alpha)\Delta$, where $\Delta \in [0, p_{\max}]$ is a reference gap, such as the largest gap between the normalized prices of an unconstrained optimal policy. The trip-based α-fair prices, which we call α-fair throughout, and the best α-fair revenue are

$$\mathcal{P}_\alpha := \Big\{p \in \mathcal{P} : \ \Big|\frac{p_e}{d_e} - \frac{p_{e'}}{d_{e'}}\Big| \le W_\alpha \ \text{ for all } e, e' \in \mathcal{A}\Big\}, \qquad \mathcal{R}^*_K(\alpha) := \sup_{p\in\mathcal{P}_\alpha,\ \theta\in\Theta} \mathcal{R}_K(p,\theta).$$

Assumption 3.1.

(a) For every arc (i,j), $\bar{F}_{ij}$ is continuous and nonincreasing, strictly decreasing on $[0, b_{ij})$, and $\bar{F}_{ij}(0) > 0$.

(b) For every arc (i,j), the revenue curve $r_{ij}(q) := q\,\bar{F}_{ij}^{-1}(q)$ is concave on the set of purchase probabilities that admissible prices can produce, $\mathcal{Q}_{ij} := [\bar{F}_{ij}(d_{ij}p_{\max}), \bar{F}_{ij}(0)]$, where $\bar{F}_{ij}^{-1} : \mathcal{Q}_{ij} \to [0, \hat{p}_{ij}]$ is the inverse of $\bar{F}_{ij}$ on $[0,\hat{p}_{ij}]$, $\hat{p}_{ij} := \min\{b_{ij}, d_{ij}p_{\max}\}$.

Assumption 3.1(b) is the regularity condition of Myerson (1981) on the range of admissible prices, and holds whenever $\bar{F}_{ij}$ has a monotone hazard rate, in particular for linear, exponential, and logistic demand.

3.1 Reformulation as a family of convex programs

In the actual system the availabilities are not decisions. Once (p,θ) is fixed, the K vehicles move as a closed queueing network, and the probability $\Pi_i(p,\theta,K)$ that a vehicle is available at location i can be computed exactly (Section 3.2) but depends on (p,θ) in a complicated way, so maximizing $\mathcal{R}_K(p,\theta)$ over α-fair policies is a nonconvex problem. The following problem instead treats the availability Π_i of each location as a decision variable, constrained only by flow balance and by the fleet size, and y_{ij} as the rate of repositioning moves from i to j:

$$\begin{aligned} J^*_K(\alpha) := \max_{\substack{p\in\mathcal{P},\ \Pi\in[0,1]^N \\ y\in\mathbb{R}^{\mathcal{V}}_+}} \quad & \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\,\Pi_i\,\bar{F}_{ij}(p_{ij})\,p_{ij} - \sum_{(i,j)\in\mathcal{V}} c_{ij}\,y_{ij} \\ \text{s.t.} \quad & \sum_j \big(\Lambda_{ij}\bar{F}_{ij}(p_{ij})\Pi_i + y_{ij}\big) = \sum_j \big(\Lambda_{ji}\bar{F}_{ji}(p_{ji})\Pi_j + y_{ji}\big) \ \ \forall i \in [N], \\ & \sum_{(i,j)\in\mathcal{A}} \frac{\Lambda_{ij}\bar{F}_{ij}(p_{ij})\,\Pi_i}{\mu_{ij}} + \sum_{(i,j)\in\mathcal{V}} \frac{y_{ij}}{\mu_{ij}} \le K, \\ & p \in \mathcal{P}_\alpha. \end{aligned} \tag{3.1}$$

Every static policy produces availabilities and repositioning rates that satisfy these constraints, so (3.1) is a relaxation and $J^*_K(\alpha) \ge \mathcal{R}^*_K(\alpha)$ (Proposition C.4). It can, however, choose

availabilities that no system with K vehicles produces, such as a vehicle that is always available at two locations when $K = 1$, so $J_K^*(\alpha)$ can be much larger than $\mathcal{R}_K^*(\alpha)$ when K is small.

The fairness constraint involves all pairs of arcs. The first step replaces it by a window of fixed width whose position is a single scalar. A price vector $p \in \mathcal{P}$ is α-fair if and only if there is $\bar{p} \in \bar{\mathcal{P}}_\alpha := [W_\alpha,\ p_{\max}]$ such that

$$\bar{p} - W_\alpha \le p_{ij}/d_{ij} \le \bar{p} \qquad \text{for all } (i,j) \in \mathcal{A}. \tag{3.2}$$

Indeed, if such $\bar{p}$ exists, any two normalized prices differ by at most W_α. Conversely, if $\underline{m}$ is the smallest normalized price, then all normalized prices lie in $[\underline{m}, \underline{m} + W_\alpha] \cap [0, p_{\max}]$, so $\bar{p} := \min\{\underline{m} + W_\alpha,\ p_{\max}\}$ works, and $\bar{p} \ge W_\alpha$ because $W_\alpha \le p_{\max}$.

The second step lifts each arc to the pair (z_{ij}, Π_i) with $z_{ij} = \Pi_i \bar{F}_{ij}(p_{ij})$, i.e., z_{ij} is the fraction of arriving customers on (i,j) who buy and are served. The objective becomes the *perspective* of the revenue curve $r_{ij}(q) = q\,\bar{F}_{ij}^{-1}(q)$ of Assumption 3.1(b),

$$\tilde{r}_{ij}(z, \Pi) := \Pi\, r_{ij}(z/\Pi) \quad (\Pi > 0), \qquad \tilde{r}_{ij}(0,0) := 0,$$

and for $\bar{p} \in \bar{\mathcal{P}}_\alpha$ the *inner problem* is

$$\begin{aligned} V_\alpha(\bar{p}) := \max_{z \in \mathbb{R}^{\mathcal{A}},\ \Pi \in [0,1]^N,\ y \in \mathbb{R}^{\mathcal{V}}_+} \quad & \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\, \tilde{r}_{ij}(z_{ij}, \Pi_i) - \sum_{(i,j)\in\mathcal{V}} c_{ij}\, y_{ij} \\ \text{s.t.} \quad & \sum_j \big(\Lambda_{ij} z_{ij} + y_{ij}\big) = \sum_j \big(\Lambda_{ji} z_{ji} + y_{ji}\big) \qquad \forall i \in [N], \\ & \sum_{(i,j)\in\mathcal{A}} \frac{\Lambda_{ij} z_{ij}}{\mu_{ij}} + \sum_{(i,j)\in\mathcal{V}} \frac{y_{ij}}{\mu_{ij}} \le K, \\ & \bar{F}_{ij}(d_{ij}\bar{p})\, \Pi_i \le z_{ij} \le \bar{F}_{ij}\big(d_{ij}(\bar{p} - W_\alpha)\big)\, \Pi_i \quad \forall (i,j) \in \mathcal{A}. \end{aligned} \tag{3.3}$$

Proposition 3.2 (Reformulation)**.** *Under Assumption 3.1(a),*

$$J_K^*(\alpha) = \max_{\bar{p} \in \bar{\mathcal{P}}_\alpha} V_\alpha(\bar{p}),$$

and all maxima involved are attained. Moreover, for every $\bar{p} \in \bar{\mathcal{P}}_\alpha$ and every feasible (z, Π, y) of (3.3)*, the prices*

$$p_{ij} := \bar{F}_{ij}^{-1}\big(z_{ij}/\Pi_i\big) \ \text{ if } z_{ij} > 0, \qquad p_{ij} := d_{ij}\bar{p} \ \text{ if } z_{ij} = 0, \tag{3.4}$$

satisfy $p \in \mathcal{P}_\alpha$, and (p, Π, y) is feasible for (3.1) *with the same objective value. Under all of Assumption 3.1, problem* (3.3) *maximizes a concave, continuous function over a nonempty compact polyhedron for every $\bar{p} \in \bar{\mathcal{P}}_\alpha$, so $V_\alpha(\bar{p})$ is computed by a convex program.*

3.2 From a window to a policy

The policy. Fix $\bar{p} \in \bar{\mathcal{P}}_\alpha$ and let $(\hat{z}, \hat{\Pi}, \hat{y})$ be feasible for (3.3). Post the prices p^H that (3.4) assigns to $(\hat{z}, \hat{\Pi})$, and reposition with probabilities

$$\theta^H_{ij} := \hat{y}_{ij}/\eta_i \ \text{ if } \eta_i > 0, \qquad \theta^H_{ij} := 0 \ \text{ if } \eta_i = 0, \qquad \eta_i := \sum_j \big(\Lambda_{ji}\hat{z}_{ji} + \hat{y}_{ji}\big), \tag{3.5}$$

where η_i is the rate at which vehicles reach i under $(\hat{z}, \hat{\Pi}, \hat{y})$. By the balance constraint $\eta_i = \sum_j (\Lambda_{ij}\hat{z}_{ij} + \hat{y}_{ij}) \ge \sum_j \hat{y}_{ij}$, so $\theta^H \in \Theta$.

Its revenue. Since $\mathcal{R}_K$ is a long-run average, it is an expected rate under the stationary distribution π_K of the vehicles,

$$\mathcal{R}_K(p^H, \theta^H) = \sum_{(i,j)\in\mathcal{A}} p_{ij}^H \lambda_{ij}\, \pi_K(k_i \geq 1) - \sum_{(i,j)\in\mathcal{V}} c_{ij}\, \mu_{ij}\, \mathbb{E}_{\pi_K}\big[t_{ij}^{\mathrm{E}}\big], \tag{3.6}$$

where $\lambda_{ij} := \Lambda_{ij}\bar{F}_{ij}(p_{ij}^H)$ is the rate at which customers for arc (i,j) accept the price, so fares on (i,j) come in at rate $p_{ij}^H\lambda_{ij}$ whenever a vehicle waits at i, k_i is the number of vehicles waiting at location i, and t_{ij}^{E} is the number of vehicles repositioning from i to j, whose moves end, and hence start, at rate μ_{ij} per vehicle and cost c_{ij} each. Because θ^H is applied when a trip ends, the routing of each vehicle does not depend on where the other vehicles are, so π_K has the product form of classical closed networks (Gordon and Newell, 1967; Baskett et al., 1975), and the policy reproduces the feasible point up to one common *scaling factor* ω_K (Appendix C.2), in that $\pi_K(k_i \geq 1) = \omega_K\hat{\Pi}_i$ and $\mu_{ij}\mathbb{E}_{\pi_K}[t_{ij}^{\mathrm{E}}] = \omega_K\hat{y}_{ij}$, so by (3.6) the revenue is $\omega_K\hat{J}$, or at least $\omega_K\hat{J}$ if the vehicles use several separate parts of the network (Proposition C.4), where

$$\hat{J} := \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\tilde{r}_{ij}(\hat{z}_{ij}, \hat{\Pi}_i) - \sum_{(i,j)\in\mathcal{V}} c_{ij}\hat{y}_{ij}$$

is the objective value of the feasible point in (3.3). The factor ω_K is computed exactly by Buzen's recursion, and Lemma C.3 shows that it tends to one as the fleet grows. Applied at the best window, this gives the main result of this section. As for access fairness in the paper, the value $J_K^*(\alpha)$ of the relaxation bounds the best α-fair revenue from above, and the heuristic attains a known fraction of it.

Theorem 3.3 (Bounds on the best α-fair revenue)**.** *Let Assumption 3.1(a) hold and $K \geq 1$. Let $\bar{p}^*$ maximize V_α over $\bar{\mathcal{P}}_\alpha$. Then* (3.3) *at $\bar{p}^*$ has an optimal solution $(\hat{z}, \hat{\Pi}, \hat{y})$ with no directed cycle in the support of $\hat{y}$, and for every such solution $\hat{J} = J_K^*(\alpha)$ and*

$$\max\Big\{\frac{K}{N+K-1+\bar{\nu}},\ 1-\sqrt{\frac{N}{K}}\Big\}\, J_K^*(\alpha) \ \leq\ \omega_K J_K^*(\alpha) \ \leq\ \mathcal{R}_K(p^H,\theta^H) \ \leq\ \mathcal{R}_K^*(\alpha) \ \leq\ J_K^*(\alpha),$$

where

$$\bar{\nu} := \min\Big\{K,\ \sum_{(i,j)\in\mathcal{A}} \frac{\Lambda_{ij}}{\mu_{ij}} + \bar{\tau}\,\Lambda_{\mathrm{tot}}\Big\}, \qquad \Lambda_{\mathrm{tot}} := \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}, \qquad \bar{\tau} := (N-1)\max_{(i,j)\in\mathcal{V}} \mu_{ij}^{-1}.$$

Since $\bar{\nu} \leq \sum_{\mathcal{A}} \Lambda_{ij}/\mu_{ij} + \bar{\tau}\Lambda_{\mathrm{tot}}$ does not grow with K, with the network data held fixed the heuristic is within a factor $1 - O(1/K)$ of $\mathcal{R}_K^*(\alpha)$ as $K \to \infty$, and since the factor $1-\sqrt{N/K}$ involves no other data, it is also asymptotically optimal when the demand rates grow with the fleet, as long as $N/K \to 0$. Unlike the relaxation for access fairness, however, $J_K^*(\alpha)$ is not the value of a single convex program, since V_α need not be concave in $\bar{p}$, so $\bar{p}^*$ is not computed directly.

3.3 Choosing the window

We search for $\bar{p}^*$ over a grid $W_\alpha = \bar{p}_0 < \bar{p}_1 < \cdots < \bar{p}_n = p_{\max}$ and certify the result. A grid point $\hat{p}$ with the largest V_α gives the lower bound $V_\alpha(\hat{p}) \leq J_K^*(\alpha)$. For the upper bound,

let $U_\alpha(\bar{p}', \bar{p}'')$ be the value of (3.3) with the window constraint replaced by the wider window $[\bar{p}' - W_\alpha, \bar{p}'']$, i.e., by

$$\bar{F}_{ij}(d_{ij}\bar{p}'')\,\Pi_i \le z_{ij} \le \bar{F}_{ij}\big(d_{ij}(\bar{p}' - W_\alpha)\big)\,\Pi_i \qquad \forall (i,j) \in \mathcal{A}.$$

Its solution need not be α-fair, and it is used only for the bound. Since $[\bar{p}' - W_\alpha, \bar{p}'']$ contains the window of every $\bar{p} \in [\bar{p}', \bar{p}'']$ and each $\bar{F}_{ij}$ is nonincreasing, the feasible set of (3.3) at such $\bar{p}$ lies in that of the problem defining $U_\alpha(\bar{p}', \bar{p}'')$. Hence

$$V_\alpha(\bar{p}) \le U_\alpha(\bar{p}', \bar{p}'') \qquad \text{for every } \bar{p} \in [\bar{p}', \bar{p}''], \tag{3.7}$$

and in the same way U_α does not increase when the interval is replaced by a subinterval, with $U_\alpha(\bar{p}', \bar{p}') = V_\alpha(\bar{p}')$. Since the cells $[\bar{p}_{k-1}, \bar{p}_k]$ cover $\bar{\mathcal{P}}_\alpha$, one of them contains a maximizer of V_α, and (3.7) brackets $J_K^*(\alpha)$ within a provable optimality gap

$$V_\alpha(\hat{p}) \;\le\; J_K^*(\alpha) \;\le\; \max_{1 \le k \le n} U_\alpha(\bar{p}_{k-1}, \bar{p}_k).$$

If the gap is too large, the cells with $U_\alpha > V_\alpha(\hat{p})$ are split, which never raises their bounds, and the others are discarded, since they cannot contain a better window. Algorithm 1 summarizes the heuristic, and Figure 4 illustrates the search. Appendix C.3 states the search precisely, bounds the number of convex programs it solves, and shows that for every $\varepsilon > 0$ it stops after finitely many of them.

Algorithm 1 Trip-based α-price-fair heuristic

Require: network $(\mathcal{A}, \Lambda, \bar{F}, \mu, c, d)$, fleet size K, price cap $p_{\max}$, window width $W_\alpha = (1-\alpha)\Delta$, grid $W_\alpha = \bar{p}_0 < \bar{p}_1 < \cdots < \bar{p}_n = p_{\max}$, tolerance $\varepsilon \ge 0$

Ensure: an α-fair static policy (p^H, θ^H) and bounds on $J_K^*(\alpha)$

For $k = 0, \ldots, n$, solve (3.3) at $\bar{p} = \bar{p}_k$, giving $V_\alpha(\bar{p}_k)$
For $k = 1, \ldots, n$, solve (3.3) with the wider window $[\bar{p}_{k-1} - W_\alpha, \bar{p}_k]$, giving $U_\alpha(\bar{p}_{k-1}, \bar{p}_k)$
$\hat{p} \leftarrow$ a grid point with the largest V_α, $\quad$ UB $\leftarrow \max_{1 \le k \le n} U_\alpha(\bar{p}_{k-1}, \bar{p}_k)$
while UB $- V_\alpha(\hat{p}) > \varepsilon$ and the computing budget is not exhausted **do**
    Split every cell with $U_\alpha > V_\alpha(\hat{p})$ at its midpoint and discard the other cells, evaluate V_α at the new grid points and U_α on the new cells, update $\hat{p}$, and set UB $\leftarrow \max\{V_\alpha(\hat{p}),\ \max U_\alpha$ over the current cells$\}$
end while
Let $(\hat{z}, \hat{\Pi}, \hat{y})$ be an optimal solution of (3.3) at $\bar{p} = \hat{p}$, and cancel the directed cycles of $\hat{y}$
Set p^H by (3.4) and θ^H by (3.5)
return (p^H, θ^H) and $V_\alpha(\hat{p}) \le J_K^*(\alpha) \le$ UB

Corollary 3.4 (Bounds from Algorithm 1)**.** *Let Assumption 3.1(a) hold and $K \ge 1$, and let (p^H, θ^H), $\hat{p}$, and* UB *be the output of Algorithm 1 with tolerance ε. Then $p^H \in \mathcal{P}_\alpha$ and $\theta^H \in \Theta$. If the algorithm stops with* UB $- V_\alpha(\hat{p}) \le \varepsilon$*, then*

$$\begin{aligned}\max\Big\{\frac{K}{N+K-1+\bar{\nu}},\ 1-\sqrt{\frac{N}{K}}\Big\}\,(\mathrm{UB}-\varepsilon) \;&\le\; \omega_K V_\alpha(\hat{p}) \;\le\; \mathcal{R}_K(p^H,\theta^H)\\ &\le\; \mathcal{R}_K^*(\alpha) \;\le\; J_K^*(\alpha) \;\le\; \mathrm{UB},\end{aligned}$$

where $\bar{\nu}$ is as in Theorem 3.3 and ω_K is the scaling factor of (p^H, θ^H). Otherwise the same chain holds with ε replaced by the optimality gap UB $- V_\alpha(\hat{p})$.

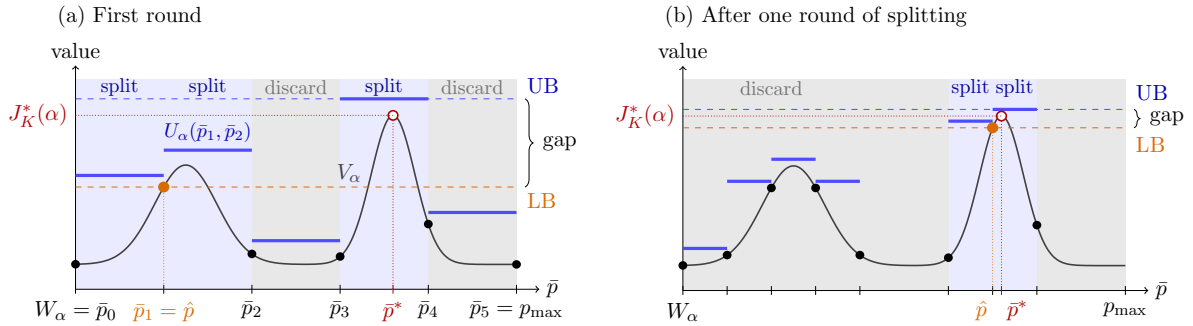


Figure 4: How Algorithm 1 searches over $\bar{p}$, on an illustrative V_α. (a) First round. Dots are the values of V_α at the grid points, and the best one, at $\hat{p}$, is the lower bound. Blue bars are the cell bounds U_α, and the largest one is the upper bound. The dashed lines are the two bounds, LB $= V_\alpha(\hat{p})$ and UB $= \max_k U_\alpha$, and the brace is the optimality gap. The hollow dot marks a maximizer $\bar{p}^*$, with value $J_K^*(\alpha)$. Gray cells have $U_\alpha \le V_\alpha(\hat{p})$ and are discarded. (b) After the cells marked "split" in (a) are split, both bounds move toward $J_K^*(\alpha)$.

Every term other than $\mathcal{R}_K^*(\alpha)$ and $J_K^*(\alpha)$ is computed, and $\mathcal{R}_K(p^H, \theta^H)$ is computed exactly from the product form. For every $\varepsilon > 0$ the algorithm stops with UB $- V_\alpha(\hat{p}) \le \varepsilon$ after finitely many convex programs if the budget allows (Appendix C.3), and then the heuristic keeps the guarantee of Theorem 3.3 with $J_K^*(\alpha)$ replaced by the computed value UB $- \varepsilon$, which is at most $J_K^*(\alpha)$.

3.4 Numerical illustration

We illustrate Corollary 3.4 on a network of $N = 3$ locations with all nine arcs, linear demand with a different maximum support on every arc, $d_{ij} = 1$, and $p_{\max} = 2$. The data and the computations are described in Appendix C.4. Figure 5 shows the revenue of the heuristic, i.e., the policy returned by Algorithm 1, the best static α-fair policy we could find by a global search over prices and repositioning probabilities, whose revenue is a lower estimate of $\mathcal{R}_K^*(\alpha)$, and the relaxation value $J_K^*(\alpha)$ with the upper bound UB of Algorithm 1. In every instance the algorithm certifies $J_K^*(\alpha)$ to within 10^{-4}, so the two coincide up to the width of the line.

Figure 5 shows two things. First, the heuristic is close to the best policy found except for the smallest fleets. At $\alpha = 1$ it earns 92% of the best revenue found for $K \le 2$ and more than 98% from $K = 4$ on, and at $K = 4$ it earns at least 98% of it for every α. Second, the gap between $J_K^*(\alpha)$ and the best policy found comes from the scaling factor ω_K, not from a weakness of the heuristic. At $\alpha = 1$ the best policy found earns only half of $J_K^*(\alpha)$ for $K = 3$ and 76% of it for $K = 8$, while the heuristic earns 96% and 99.7% of the best revenue found. The fleet constraint of (3.1) binds only for the smallest fleets, so $J_K^*(\alpha)$ barely changes from $K = 2$ on.

3.5 Origin-based price fairness

Origin-based price fairness bounds the gaps between the demand-weighted average normalized prices of the origins instead of the gaps between all normalized prices. Let $\mathcal{O} \subseteq [N]$ be the set of locations with at least one outgoing arc, let $T_{ij} := \Lambda_{ij} / \sum_k \Lambda_{ik}$ for $(i, j) \in \mathcal{A}$, and let $u_i(p) := \sum_j T_{ij}\, p_{ij}/d_{ij}$ for $i \in \mathcal{O}$. The origin-based α-fair prices and the best origin-based

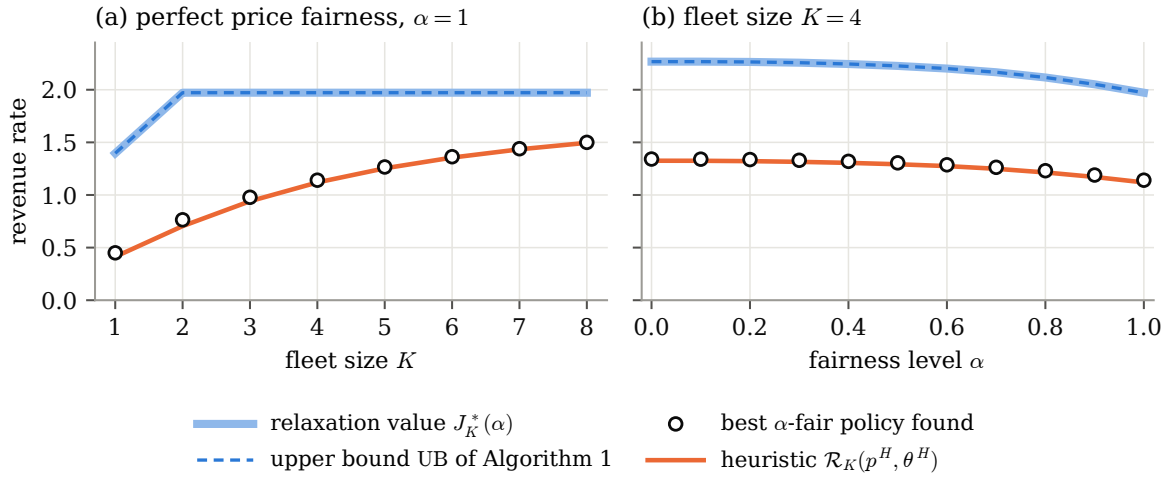


Figure 5: The relaxation value $J_K^*(\alpha)$, plotted as $V_\alpha(\hat{p})$, the upper bound UB of Algorithm 1, the revenue of the best α-fair policy found, and the exact revenue of the heuristic. (a) Perfect price fairness, $\alpha = 1$, as the fleet grows. (b) Fleet size $K = 4$, as the fairness level varies, with $W_\alpha = (1-\alpha)\Delta$ and $\Delta = 0.79$, the largest gap between the prices of the best policy found without a fairness constraint.

α-fair revenue are

$$\mathcal{P}_\alpha^{\mathrm{o}} := \big\{p \in \mathcal{P} : \ |u_i(p) - u_k(p)| \le W_\alpha \ \text{ for all } i,k \in \mathcal{O}\big\}, \qquad \mathcal{R}_K^{\mathrm{o}*}(\alpha) := \sup_{p \in \mathcal{P}_\alpha^{\mathrm{o}},\ \theta \in \Theta} \mathcal{R}_K(p,\theta).$$

This is a weaker requirement than trip-based price fairness, and a less natural one for riders. A rider pays the price of their own trip, not the demand-weighted average of the prices out of their origin, so two origins can have equal averages while individual trips from them are priced very differently, and the price of a closed arc enters the average although nobody pays it, so the platform can move an average with a price that has no effect on any rider. Every trip-based α-fair price vector is origin-based α-fair, since each average $u_i(p)$ lies in the window (3.2) together with the prices it averages, so the heuristic is also origin-based α-fair, but Corollary 3.4 compares it with $\mathcal{R}_K^*(\alpha)$, which can be smaller than $\mathcal{R}_K^{\mathrm{o}*}(\alpha)$.

The reformulation of Section 3.1 needs the fairness constraint to depend on the prices only through (z, Π), which fails here because the price of a closed arc enters an average. It holds for linear demand with prices capped at the maximum support, $p_{ij} \le \min\{b_{ij}, d_{ij}p_{\max}\}$. Then $p_{ij} = b_{ij}(1 - z_{ij}/\Pi_i)$, and fixing the largest average $\bar{u}$ makes the constraint linear,

$$(\bar{u} - W_\alpha)\,\Pi_i \ \le \ \sum_j T_{ij}\frac{b_{ij}}{d_{ij}}\big(\Pi_i - z_{ij}\big) \ \le \ \bar{u}\,\Pi_i \qquad \forall i \in \mathcal{O},$$

so the results of this section carry over with the same proofs. The cap rules out closing an arc with a price above b_{ij} to move an average, which is the mechanism behind the strategic market abandonment in the paper. Without the cap, or for other demand models, we know no tractable relaxation.

Appendices

A Proofs for Section 1

Throughout this appendix $b_2 \geq b_1$, as in Proposition 1.1.

Notation. The fairness constraint restricts only the difference between the two prices, so we work in the average price and the price gap, $s := \frac{1}{2}(p_1 + p_2)$ and $\gamma := p_2 - p_1$, in which the constraint is $|\gamma| \leq (1-\alpha)\Delta$ and perfect fairness is the line $\gamma = 0$. For $\Phi \in \{\mathcal{R}, \mathcal{W}, S_1, S_2\}$ we write $\widetilde{\Phi}(s,\gamma) := \Phi(s - \frac{\gamma}{2},\, s + \frac{\gamma}{2})$, with subscripts s, γ for partial derivatives. Where both demand rates are positive, the expected duration of a round trip is

$$H(\boldsymbol{p}) := \frac{1}{g(\boldsymbol{p})} = \frac{1}{\lambda_1} + \frac{1}{\lambda_2} = \frac{b_1}{\Lambda_1(b_1 - p_1)} + \frac{b_2}{\Lambda_2(b_2 - p_2)},$$

so $\mathcal{R}(\boldsymbol{p}) = (p_1 + p_2)/H(\boldsymbol{p})$, and H is strictly convex because each term is strictly convex in its own fare. At a common price $p < b_1$ the two demand rates satisfy $\lambda_1 + \lambda_2 = D(p)/(b_1 b_2)$ and $\lambda_1 \lambda_2 = \Lambda_1 \Lambda_2 (b_1 - p)(b_2 - p)/(b_1 b_2)$, where

$$D(p) := \Lambda_1 b_2 (b_1 - p) + \Lambda_2 b_1 (b_2 - p) > 0.$$

The proof has four steps. Lemma A.1 shows that the platform's choice is unique at every fairness level and that the common price P it charges at $\alpha = 1$ is interior, satisfies a strict second-order condition, and is in bijection with Λ_2/Λ_1, which lets us trade Λ_2 for P in the later algebra, and Lemma A.2 records the unconstrained optimum. Lemma A.3 shows that as the constraint is relaxed the optimum moves along a single path in the (s,γ) plane, so each $\mathcal{D}\Phi$ is a directional derivative along it. Lemma A.4 computes these derivatives and reduces their signs to those of three polynomials in (P, b_1, b_2). Lemma A.5 shows that the regions where the polynomials are positive are nested, which leaves exactly the five regimes.

Lemma A.1 (The platform's choice, and the common price)**.**

(i) Let C be a compact convex set of prices on which revenue is positive somewhere. Then $\mathcal{R}$ has a unique maximizer on C, with $p_i < b_i$ for both i. In particular $\boldsymbol{p}^*$, Δ, and $\Phi(\alpha)$ are well defined for every $\alpha \in [0,1]$.

(ii) Let $\varphi(p) := \widetilde{\mathcal{R}}(p,0)$ be revenue at the common price p. Then $\varphi(b_1) = 0$ and

$$\varphi(p) = \frac{2\Lambda_1\Lambda_2\, p\,(b_1 - p)(b_2 - p)}{D(p)} \qquad (0 \leq p < b_1),$$

and φ has a unique maximizer P on $[0, b_1]$, which is interior and satisfies $\widetilde{\mathcal{R}}_{ss}(P,0) = \varphi''(P) < 0$.

(iii) If $b_2 > b_1$, then P is the unique solution in $I := \big(\frac{b_1}{2}, \min\{b_1, \frac{b_2}{2}\}\big)$ of

$$\frac{\Lambda_2}{\Lambda_1} = h(P) := \frac{b_2(b_2 - 2P)(b_1 - P)^2}{b_1(2P - b_1)(b_2 - P)^2}, \tag{A.1}$$

and for fixed b_1, b_2 the map $\Lambda_2/\Lambda_1 \mapsto P$ is a strictly decreasing bijection from $(0,\infty)$ onto I, with inverse h.

Proof. (i). Revenue is continuous on the box and vanishes as soon as one λ_i does, so a maximizer on C exists and has $p_i < b_i$. For uniqueness, let $r > 0$ be the maximum of $\mathcal{R}$ on C. Since $H > 0$, $\mathcal{R}(\boldsymbol{p}) \le r$ if and only if $p_1 + p_2 - rH(\boldsymbol{p}) \le 0$, so $p_1 + p_2 - rH$ is nonpositive on C and vanishes exactly where $\mathcal{R} = r$, i.e., its maximizers are the maximizers of $\mathcal{R}$. Being linear minus strictly convex, it is strictly concave on the convex set of points of C where both rates are positive, which contains every maximizer, so it has at most one maximizer there. Finally, the box and its intersection with each slab $|p_1 - p_2| \le (1-\alpha)\Delta$ are compact and convex and contain the common prices (p, p) with $p \in (0, b_1)$, where revenue is positive, so this applies at every fairness level.

(ii). At $p_1 = p_2 = p \in [0, b_1)$, revenue is $\varphi(p) = 2p/H(p,p)$, and clearing denominators gives the closed form, while $\varphi(b_1) = 0$ because λ_1 vanishes there. The common prices (p, p) with $p \in [0, b_1]$ form a compact convex set on which revenue is positive somewhere, so by (i) φ has a unique maximizer P on $[0, b_1]$, which is interior because $\varphi(0) = \varphi(b_1) = 0$. Writing H' and H'' for the derivatives of $p \mapsto H(p,p)$,

$$\varphi'(p) = \frac{2\,(H - pH')}{H^2}, \qquad (H - pH')' = -pH'',$$

so $\varphi'(P) = 0$ means $H(P,P) = PH'(P)$. Differentiating φ' once more, the term containing $H - pH'$ vanishes at P, which leaves

$$\varphi''(P) = \frac{2\,(H - pH')'}{H^2}\Big|_{p=P} = -\frac{2P\,H''(P)}{H(P,P)^2} < 0,$$

since $H''(p) = 2b_1/\big(\Lambda_1(b_1 - p)^3\big) + 2b_2/\big(\Lambda_2(b_2 - p)^3\big) > 0$. Since $\varphi(p) = \widetilde{\mathcal{R}}(p, 0)$, we have $\widetilde{\mathcal{R}}_{ss}(P, 0) = \varphi''(P) < 0$.

(iii). The first-order condition $H(P,P) = PH'(P)$ from (ii) reads

$$\frac{b_1}{\Lambda_1} \cdot \frac{b_1 - 2P}{(b_1 - P)^2} + \frac{b_2}{\Lambda_2} \cdot \frac{b_2 - 2P}{(b_2 - P)^2} = 0.$$

The two terms cannot both vanish, since $b_1 \ne b_2$, so they have opposite signs. As $b_1 < b_2$, this forces $b_1 - 2P < 0 < b_2 - 2P$, and with $P < b_1$ we get $P \in I$. Solving for Λ_2/Λ_1 gives (A.1), where h is the ratio of $b_2(b_2 - 2P)/(b_2 - P)^2$ to $b_1(2P - b_1)/(b_1 - P)^2$. On I the first is positive with derivative $-2b_2P/(b_2 - P)^3 < 0$ and the second is positive with derivative $2b_1P/(b_1 - P)^3 > 0$, so h is continuously differentiable with $h' < 0$. As $P \downarrow \frac{b_1}{2}$ the second vanishes while the first stays positive, so $h \to +\infty$. At the right end, either $\frac{b_2}{2} < b_1$ and the first tends to 0 while the second stays bounded away from zero, or $\frac{b_2}{2} \ge b_1$ and the second tends to $+\infty$ while the first stays bounded, since $b_2 - P \ge b_2 - b_1 > 0$. Either way $h \to 0^+$, so h is a strictly decreasing bijection of I onto $(0, \infty)$. Hence (A.1) has exactly one solution in I for each value of Λ_2/Λ_1, namely P, and the map $\Lambda_2/\Lambda_1 \mapsto P$ is h^{-1}. □

Lemma A.2 (The unconstrained optimum, cf. Elmachtoub and Kim, 2026)**.** *Let* $\vartheta := \sqrt{\Lambda_1 b_2/(\Lambda_2 b_1)}$. *If* $b_1 = b_2$ *or* $\Lambda_2/\Lambda_1 \le 4b_1b_2/(b_2 - b_1)^2$, *the unconstrained optimum is the point*

$$p_1^* = \frac{b_1 - b_2 + 2b_1\vartheta}{2(1+\vartheta)}, \qquad p_2^* - p_1^* = \frac{(3b_2 - b_1) - \vartheta(3b_1 - b_2)}{2(1+\vartheta)}, \tag{A.2}$$

at which both demand rates are positive. Otherwise it is a corner point with $p_1^* = 0 < p_2^*$. *In both cases* $p_2^* - p_1^*$ *has the sign of* $(3b_2 - b_1)\sqrt{\Lambda_2 b_1} - (3b_1 - b_2)\sqrt{\Lambda_1 b_2}$.

Proof. Where both demand rates are positive, $\partial\mathcal{R}/\partial p_1 = g - (p_1+p_2)\big(\frac{\lambda_2}{\lambda_1+\lambda_2}\big)^2\frac{\Lambda_1}{b_1}$ and $\partial\mathcal{R}/\partial p_2 = g-(p_1+p_2)\big(\frac{\lambda_1}{\lambda_1+\lambda_2}\big)^2\frac{\Lambda_2}{b_2}$. Setting both to zero gives $\lambda_1(\lambda_1+\lambda_2) = (p_1+p_2)\lambda_2\Lambda_1/b_1$ and $\lambda_2(\lambda_1+\lambda_2) = (p_1+p_2)\lambda_1\Lambda_2/b_2$, whose ratio is $(\lambda_2/\lambda_1)^2 = \Lambda_2 b_1/(\Lambda_1 b_2) = \vartheta^{-2}$. Substituting $\lambda_2 = \lambda_1/\vartheta$, $p_1 = b_1(1-\lambda_1/\Lambda_1)$, and $p_2 = b_2(1-\lambda_2/\Lambda_2)$ into the first gives $\lambda_1 = \Lambda_1(b_1+b_2)/\big(2b_1(1+\vartheta)\big)$ and $p_1+p_2 = \frac{b_1+b_2}{2}$, as in the paper, hence (A.2). So any critical point of $\mathcal{R}$ with positive demand rates is the point (A.2), which lies in the box exactly when $p_1^* \geq 0$, i.e., $\Lambda_2/\Lambda_1 \leq 4b_1b_2/(b_2-b_1)^2$ (one checks $p_1^* < b_1$ and $0 < p_2^* < b_2$ always). When it does, it is the maximizer, because $\boldsymbol{p} \mapsto p_1+p_2-\mathcal{R}(\boldsymbol{p}^*)H(\boldsymbol{p})$ is strictly concave with zero gradient at $\boldsymbol{p}^*$, hence nonpositive wherever both rates are positive, i.e., $\mathcal{R} \leq \mathcal{R}(\boldsymbol{p}^*)$ there, while $\mathcal{R} = 0 \leq \mathcal{R}(\boldsymbol{p}^*)$ elsewhere. When it does not, the maximizer, unique with $p_i < b_i$ by Lemma A.1(i), lies on the boundary $p_1 = 0$ or $p_2 = 0$ of the box. The second is impossible, because along $p_2 = 0$

$$\frac{\partial\mathcal{R}}{\partial p_2} = \frac{\lambda_1\Lambda_2}{\lambda_1+\Lambda_2}\Big[1-\frac{p_1\lambda_1}{b_2(\lambda_1+\Lambda_2)}\Big] > 0, \qquad \text{since } \frac{p_1\lambda_1}{b_2(\lambda_1+\Lambda_2)} < \frac{p_1}{b_2} < 1.$$

So a corner optimum has $p_1^* = 0 < p_2^*$. For the sign of the gap, multiplying the numerator in (A.2) by $\sqrt{\Lambda_2 b_1} > 0$ shows that at the interior point $p_2^* - p_1^*$ has the sign of $(3b_2-b_1)\sqrt{\Lambda_2 b_1} - (3b_1-b_2)\sqrt{\Lambda_1 b_2}$. At a corner optimum the gap is $p_2^* > 0$, and the expression is positive too, automatically if $b_2 \geq 3b_1$. If $b_1 < b_2 < 3b_1$, both sides of $(3b_2-b_1)\sqrt{\Lambda_2 b_1} > (3b_1-b_2)\sqrt{\Lambda_1 b_2}$ are positive, so after squaring it is equivalent to $\Lambda_2/\Lambda_1 > b_2(3b_1-b_2)^2/\big(b_1(3b_2-b_1)^2\big)$, and this follows from $\Lambda_2/\Lambda_1 > 4b_1b_2/(b_2-b_1)^2$ because $4b_1^2(3b_2-b_1)^2-(3b_1-b_2)^2(b_2-b_1)^2 = (b_1+b_2)^2(-b_2^2+10b_1b_2-5b_1^2) \geq 0$ for $b_1 < b_2 < 3b_1$, the quadratic in b_2 having roots $(5\pm2\sqrt{5})b_1$. □

At $\alpha = 1-\varepsilon$ the gap is confined to $[-\varepsilon\Delta, \varepsilon\Delta]$. The next lemma shows that the constraint binds, i.e., the platform opens the gap as far as allowed, in the direction in which revenue increases.

Lemma A.3 (The constraint binds, and the expansion)**.** *Suppose $\Delta > 0$ and $\widetilde{\mathcal{R}}_\gamma(P,0) \neq 0$, and set $\sigma := \operatorname{sign}\widetilde{\mathcal{R}}_\gamma(P,0)$. Then for all small $\varepsilon > 0$ the problem at $\alpha = 1-\varepsilon$ has a unique optimum $(s_\varepsilon, \gamma_\varepsilon)$, the constraint binds there, $\gamma_\varepsilon = \sigma\varepsilon\Delta$, and $\widetilde{\mathcal{R}}_s(s_\varepsilon,\gamma_\varepsilon) = 0$. More precisely, $s_\varepsilon = s(\gamma_\varepsilon)$, where $s(\gamma)$ denotes the best average fare at gap γ, i.e., the C^1 solution near P of $\widetilde{\mathcal{R}}_s(s(\gamma),\gamma) = 0$ with $s(0) = P$. Consequently, for every $\Phi \in \{\mathcal{R}, \mathcal{W}, S_1, S_2\}$,*

$$\mathcal{D}\Phi = \sigma\Delta\, G_\Phi, \qquad G_\Phi := \big(\widetilde{\Phi}_\gamma + \widetilde{\Phi}_s\, s'(0)\big)\big|_{(P,0)}, \qquad s'(0) = -\frac{\widetilde{\mathcal{R}}_{s\gamma}(P,0)}{\widetilde{\mathcal{R}}_{ss}(P,0)}, \tag{A.3}$$

and in particular $\mathcal{D}\mathcal{R} = \Delta\,|\widetilde{\mathcal{R}}_\gamma(P,0)| > 0$.

Proof. Localization. The point $(P,0)$ is feasible at every ε and carries revenue $\varphi(P) > 0$, so no optimum has a vanishing demand rate, and the feasible set is compact, so an optimum exists. Let $L := \{(s,0) : 0 \leq s \leq b_1\}$, on which $\widetilde{\mathcal{R}} = \varphi$ has the unique maximizer $(P,0)$ by Lemma A.1(ii). If $\varepsilon_n \downarrow 0$ and z_n are optima at those levels, any limit point $\bar{z}$ of (z_n) lies in L and satisfies $\widetilde{\mathcal{R}}(\bar{z}) \geq \widetilde{\mathcal{R}}(x)$ for all $x \in L$, since every point of L is feasible at every level, hence $\bar{z} = (P,0)$. So for small ε every optimum lies in any prescribed neighbourhood U of $(P,0)$, and we work inside U, where $\widetilde{\mathcal{R}}$ is smooth and, shrinking U, $\widetilde{\mathcal{R}}_{ss} < 0$ by Lemma A.1(ii).

The inner problem. Since $\widetilde{\mathcal{R}}_s(P,0) = \varphi'(P) = 0$ and $\widetilde{\mathcal{R}}_{ss}(P,0) < 0$, the implicit function theorem gives $\varepsilon_0 > 0$ and a unique C^1 map $\gamma \mapsto s(\gamma)$ on $(-\varepsilon_0, \varepsilon_0)$ with $s(0) = P$ and

$\widetilde{\mathcal{R}}_s(s(\gamma),\gamma)=0$, with $s'(0)$ as in (A.3). Strict concavity of $s\mapsto\widetilde{\mathcal{R}}(s,\gamma)$ on U makes $s(\gamma)$ the unique maximizer for each gap γ.

Binding. Let $V(\gamma):=\widetilde{\mathcal{R}}(s(\gamma),\gamma)$. Then $V'(\gamma)=\widetilde{\mathcal{R}}_\gamma(s(\gamma),\gamma)$ because $\widetilde{\mathcal{R}}_s(s(\gamma),\gamma)=0$, so $V'(0)=\widetilde{\mathcal{R}}_\gamma(P,0)\neq 0$ has sign σ, and so does V' near 0. On $|\gamma|\le\varepsilon\Delta<\varepsilon_0$ the strictly monotone V is maximized only at $\gamma_\varepsilon=\sigma\varepsilon\Delta$, which gives uniqueness and binding.

Expansion. With $\gamma(\varepsilon):=\sigma\varepsilon\Delta$ and $\Psi(\varepsilon):=\widetilde{\Phi}(s(\gamma(\varepsilon)),\gamma(\varepsilon))=\Phi(1-\varepsilon)$, the chain rule gives $\Psi'(0)=\sigma\Delta\,(\widetilde{\Phi}_\gamma+\widetilde{\Phi}_s\,s'(0))|_{(P,0)}=\sigma\Delta\,G_\Phi$, which is (A.3) by (1.2). For $\Phi=\mathcal{R}$, $\widetilde{\mathcal{R}}_s(P,0)=0$ gives $G_{\mathcal{R}}=\widetilde{\mathcal{R}}_\gamma(P,0)$ and $\mathcal{DR}=\sigma\Delta\widetilde{\mathcal{R}}_\gamma(P,0)=\Delta|\widetilde{\mathcal{R}}_\gamma(P,0)|$. □

The coefficients G_Φ are rational in $(P,b_1,b_2,\Lambda_1,\Lambda_2)$ with P defined only implicitly, so we substitute $\Lambda_2=\Lambda_1 h(P)$ from (A.1), which leaves rational functions of (P,b_1,b_2,Λ_1) whose signs are carried by three polynomials in (P,b_1,b_2).

Lemma A.4 (Closed forms and signs of the coefficients)**.** *Assume $b_2>b_1$, and let $P\in I$ be as in Lemma A.1(iii). Write*

$$\begin{aligned} B&:=b_1+b_2-3P, \qquad \widehat{S}_1:=b_1(b_1+b_2)(b_2-P)-8P^2B,\\ \widehat{S}_2&:=b_2(b_1+b_2)(b_1-P)-8P^2B, \qquad \widehat{W}:=(b_1+b_2)^2(b_1-P)(b_2-P)-16P^3B. \end{aligned} \tag{A.4}$$

(i) $\operatorname{sign}\widetilde{\mathcal{R}}_\gamma(P,0)=\operatorname{sign}(b_1+b_2-4P)=:\sigma$, *and σ is the sign of* $(3b_2-b_1)\sqrt{\Lambda_2 b_1}-(3b_1-b_2)\sqrt{\Lambda_1 b_2}$, *as in* (1.3), *hence of $p_2^*-p_1^*$ by Lemma A.2. In particular $\sigma=0$ only if $b_2<3b_1$.*

(ii) With $\Lambda_2=\Lambda_1 h(P)$ from (A.1), *the coefficients G_Φ of* (A.3) *are*

$$\begin{aligned} G_{\mathcal{R}}&=\frac{\Lambda_1(b_1-P)^2(b_2-2P)\,(b_1+b_2-4P)}{P\,b_1(b_2-b_1)^2}, \quad G_{\mathcal{W}}=\frac{\Lambda_1(b_1-P)^2(b_2-2P)}{8P^3b_1(b_2-b_1)^2B}\,\widehat{W},\\ G_{S_1}&=\frac{(b_1-P)^3(b_2-2P)}{8P^3b_1(b_2-b_1)^2B}\,\widehat{S}_1, \qquad\qquad G_{S_2}=\frac{(b_2-P)^3(2P-b_1)}{8P^3\,b_2\,(b_2-b_1)^2B}\,\widehat{S}_2. \end{aligned} \tag{A.5}$$

All four prefactors are strictly positive, so $\operatorname{sign}\mathcal{DW}=\sigma\operatorname{sign}\widehat{W}$, $\operatorname{sign}\mathcal{D}S_1=\sigma\operatorname{sign}\widehat{S}_1$, *and* $\operatorname{sign}\mathcal{D}S_2=\sigma\operatorname{sign}\widehat{S}_2$.

(iii) The three polynomials satisfy

$$\widehat{W}=8P^2(b_1+b_2-4P)B+(b_1-P)\widehat{S}_1+(b_2-P)\widehat{S}_2, \tag{A.6a}$$

$$\widehat{S}_1-\widehat{S}_2=P(b_1+b_2)(b_2-b_1)>0. \tag{A.6b}$$

(iv) $-\frac{1}{2}<s'(0)<\frac{1}{2}$. *Hence along the optimal path $dp_1/d\gamma=s'(0)-\frac{1}{2}<0<s'(0)+\frac{1}{2}=dp_2/d\gamma$, i.e., one price falls and the other rises.*

Proof. (i). Since $p_1+p_2=2s$, $\widetilde{\mathcal{R}}_\gamma=2s\,g_\gamma$, and since raising γ lowers p_1 and raises p_2, each by half as much, $\partial_\gamma\lambda_1=\frac{\Lambda_1}{2b_1}$ and $\partial_\gamma\lambda_2=-\frac{\Lambda_2}{2b_2}$, so

$$g_\gamma=\frac{\lambda_2^2\,\partial_\gamma\lambda_1+\lambda_1^2\,\partial_\gamma\lambda_2}{(\lambda_1+\lambda_2)^2}, \qquad \text{so} \qquad \widetilde{\mathcal{R}}_\gamma(p,0)=\frac{\Lambda_1\Lambda_2\,p\,E(p)}{D(p)^2},$$

where $E(p):=\Lambda_2 b_1(b_2-p)^2-\Lambda_1 b_2(b_1-p)^2$. Substituting $\Lambda_2=\Lambda_1 h(P)$,

$$E(P)=\Lambda_1 b_2(b_1-P)^2\,\frac{b_1+b_2-4P}{2P-b_1}, \qquad D(P)=\frac{\Lambda_1 b_2\,P\,(b_1-P)(b_2-b_1)}{(2P-b_1)(b_2-P)}, \tag{A.7}$$

the second because $(2P-b_1)(b_2-P)+(b_2-2P)(b_1-P) = P(b_2-b_1)$. On I we have $2P-b_1>0$, so $\operatorname{sign}\widetilde{\mathcal{R}}_\gamma(P,0) = \operatorname{sign}(b_1+b_2-4P)$. It remains to compare this sign with that of the expression in (1.3). If $b_2 \ge 3b_1$, then $b_1+b_2-4P > b_2-3b_1 \ge 0$ since $P<b_1$, and $(3b_1-b_2)\sqrt{\Lambda_1 b_2} \le 0 < (3b_2-b_1)\sqrt{\Lambda_2 b_1}$, so both are $+1$. If $b_1 < b_2 < 3b_1$, then $\frac{b_1+b_2}{4} \in I$ and h is strictly decreasing, so $b_1+b_2-4P>0$ holds exactly when $\Lambda_2/\Lambda_1 > h(\frac{b_1+b_2}{4}) = b_2(3b_1-b_2)^2/\big(b_1(3b_2-b_1)^2\big)$, while $(3b_2-b_1)\sqrt{\Lambda_2 b_1} > (3b_1-b_2)\sqrt{\Lambda_1 b_2}$ has both sides positive and squares to the same inequality, and likewise with equality. So the two signs agree, and the last claim follows from Lemma A.2.

(ii). Differentiating (1.1) and substituting $\Lambda_2 = \Lambda_1 h(P)$,

$$\widetilde{\mathcal{R}}_{ss}(P,0) = -\frac{4\Lambda_1(b_2-2P)(b_1-P)\,B}{P\,b_1(b_2-P)(b_2-b_1)}, \qquad \widetilde{\mathcal{R}}_{s\gamma}(P,0) = \frac{\Lambda_1(b_2-2P)(b_1-P)\,M(P)}{P^2\,b_1(b_2-P)(b_2-b_1)^2},$$

where $M(P) := 12P^3 - 11P^2(b_1+b_2) + 3P(b_1+b_2)^2 - b_1b_2(b_1+b_2)$, so that

$$s'(0) = \frac{M(P)}{4P\,(b_2-b_1)\,B}. \tag{A.8}$$

For the surplus terms, at $(P,0)$ and on (A.1),

$$\widetilde{S}_{1,s} = -\frac{(b_2-2P)(b_1-P)^2}{2P^2(b_2-b_1)}, \qquad \widetilde{S}_{1,\gamma} = \frac{(b_2-2P)(b_1-P)^2\,(4P^2-6Pb_1+b_1^2+b_1b_2)}{4P^2b_1(b_2-b_1)^2},$$
$$\widetilde{S}_{2,s} = -\frac{(2P-b_1)(b_2-P)^2}{2P^2(b_2-b_1)}, \qquad \widetilde{S}_{2,\gamma} = \frac{(2P-b_1)(b_2-P)^2\,(4P^2-6Pb_2+b_1b_2+b_2^2)}{4P^2\,b_2\,(b_2-b_1)^2}.$$

Inserting (A.7) into $\widetilde{\mathcal{R}}_\gamma = \Lambda_1\Lambda_2 PE/D^2$ with $\Lambda_2 = \Lambda_1 h(P)$ gives the $G_{\mathcal{R}}$ entry of (A.5), and $G_{S_i} = \widetilde{S}_{i,\gamma} + \widetilde{S}_{i,s}s'(0)$ with (A.8), collected over the common denominators $8P^3b_1(b_2-b_1)^2B$ and $8P^3b_2(b_2-b_1)^2B$, gives G_{S_1} and G_{S_2}. Finally $G_{\mathcal{W}} = G_{\mathcal{R}} + \Lambda_1 G_{S_1} + \Lambda_2 G_{S_2}$ by linearity, and with $\Lambda_2 = \Lambda_1 h(P)$ the three entries collected over the common denominator $8P^3b_1(b_2-b_1)^2B$ have numerator $\Lambda_1(b_1-P)^2(b_2-2P)\,\widehat{W}$ with $\widehat{W}$ as in (A.4), which is the $G_{\mathcal{W}}$ entry of (A.5). Dividing the same identity by $\Lambda_1(b_1-P)^2(b_2-2P)/[8P^3b_1(b_2-b_1)^2B]$ turns its four terms into $\widehat{W}$, $8P^2(b_1+b_2-4P)B$, $(b_1-P)\widehat{S}_1$, and $(b_2-P)\widehat{S}_2$, which is (A.6a). For the prefactors, on I we have $P>0$, $b_1-P>0$, $b_2-P>0$, $b_2-2P>0$, $2P-b_1>0$, $b_2>b_1$, $\Lambda_1>0$, and $B = (b_1-P)+(b_2-2P) > 0$.

(iii). Identity (A.6a) was shown in (ii). For (A.6b), $\widehat{S}_1 - \widehat{S}_2 = (b_1+b_2)\big[b_1(b_2-P) - b_2(b_1-P)\big] = P(b_1+b_2)(b_2-b_1)$.

(iv). From (A.8),

$$\tfrac{1}{2} - s'(0) = \frac{(b_1-P)\,Q_1(P)}{4P(b_2-b_1)B}, \qquad s'(0)+\tfrac{1}{2} = \frac{(b_2-P)\,Q_2(P)}{4P(b_2-b_1)B},$$

with $Q_1(P) := 12P^2 - 5P(b_1+b_2) + b_2(b_1+b_2)$ and $Q_2(P) := -12P^2 + 5P(b_1+b_2) - b_1(b_1+b_2)$, so it suffices to show $Q_1>0$ and $Q_2>0$ on I. Expanding around $P = \frac{b_1}{2}$,

$$Q_1(P) = 12\big(P-\tfrac{b_1}{2}\big)^2 + (7b_1-5b_2)\big(P-\tfrac{b_1}{2}\big) + \tfrac{(2b_2-b_1)(b_2-b_1)}{2}.$$

If $b_2 \le \frac{7}{5}b_1$ all three terms are nonnegative and the last is positive. If $b_2 > \frac{7}{5}b_1$ the discriminant in $P-\frac{b_1}{2}$ is $(7b_1-5b_2)^2 - 24(2b_2-b_1)(b_2-b_1) = -(b_1+b_2)(23b_2-25b_1) < 0$. Either way $Q_1>0$ for $P > \frac{b_1}{2}$. The function Q_2 is concave in P with $Q_2(\frac{b_1}{2}) = \frac{3}{2}b_1(b_2-b_1) > 0$, and at the right end of I either $Q_2(b_1) = 4b_1(b_2-2b_1) \ge 0$ (when $b_2 \ge 2b_1$) or $Q_2(\frac{b_2}{2}) = \frac{1}{2}(b_2-b_1)(2b_1-b_2) > 0$ (when $b_1 < b_2 < 2b_1$). A concave function that is positive at one end of an interval and nonnegative at the other is positive inside, so $Q_2>0$ on I. □

It remains to show which sign patterns the three polynomials of (A.4) can take. Throughout, (P, b_1, b_2) ranges over

$$\Omega := \big\{(P, b_1, b_2) : \ \tfrac{b_1}{2} < P < b_1, \ b_2 > 2P\big\},$$

which contains every (P, b_1, b_2) with $b_2 > b_1$ and $P \in I$. The next lemma shows that the sets where the polynomials are positive are nested, $\{\widehat{S}_2 > 0\} \subset \{\widehat{W} > 0\} \subset \{\widehat{S}_1 > 0\}$.

Lemma A.5 (The nesting)**.** *For $(P, b_1, b_2) \in \Omega$:*
(i) If $b_1 + b_2 \le 4P$ then $\widehat{W} < 0$, $\widehat{S}_2 < 0$, and $\widehat{S}_1 > 0$.
(ii) If $b_1 + b_2 > 4P$ and $\widehat{S}_2 \ge 0$ then $\widehat{W} > 0$ and $\widehat{S}_1 > 0$.
(iii) If $\widehat{S}_1 \le 0$ then $\widehat{W} < 0$. Equivalently, $\widehat{W} \ge 0$ implies $\widehat{S}_1 > 0$.

Proof. A preliminary bound. We first show that $\widehat{S}_1 \le 0$ at a point of Ω forces $P > \frac{8}{9}b_1$ and $b_2 < P + 8P^2/b_1 < 9b_1$ there. As a polynomial in b_2, $\widehat{S}_1 = b_1 b_2^2 + (b_1^2 - b_1 P - 8P^2)\, b_2 - (b_1^2 P + 8 b_1 P^2 - 24P^3)$ is an upward-opening parabola with minimizer $b_2^* = (8P^2 + b_1 P - b_1^2)/(2b_1)$. For fixed P and b_1, Ω lets b_2 range over $(2P, \infty)$, and at the left end $\widehat{S}_1|_{b_2 = 2P} = P(2P - b_1)(4P - b_1) > 0$ because $P > \frac{b_1}{2}$. So $\widehat{S}_1 \le 0$ at some $b_2 > 2P$ requires two conditions. Condition 1 is that the minimizer lies to the right of $2P$, since otherwise $\widehat{S}_1$ only increases on $b_2 > 2P$. This reads $8P^2 - 3b_1 P - b_1^2 > 0$, i.e., $P > \frac{3+\sqrt{41}}{16} b_1$, and fails for $P \le \frac{7}{12} b_1$ since $\frac{7}{12} < \frac{3+\sqrt{41}}{16}$. Condition 2 is that the minimum value is at most zero, i.e., the discriminant is nonnegative, $\delta_1(P) \ge 0$, where

$$\delta_1(P) := (b_1^2 - b_1 P - 8P^2)^2 + 4b_1(b_1^2 P + 8 b_1 P^2 - 24P^3) = 64P^4 - 80 b_1 P^3 + 17 b_1^2 P^2 + 2 b_1^3 P + b_1^4.$$

This fails for $\frac{7}{12} b_1 \le P \le \frac{8}{9} b_1$. Indeed, $\delta_1''(P) = 768P^2 - 480 b_1 P + 34 b_1^2$ is an upward parabola in P with lowest point at $P = \frac{5}{16} b_1$, so it increases for $P \ge \frac{5}{16} b_1$, and it equals $\frac{46}{3} b_1^2 > 0$ at $P = \frac{7}{12} b_1$, so δ_1 is convex on $[\frac{7}{12} b_1, b_1]$. Moreover $\delta_1(\frac{7}{12} b_1) = -\frac{671}{1296} b_1^4 < 0$ and $\delta_1(\frac{8}{9} b_1) = -\frac{143}{6561} b_1^4 < 0$, and a convex function negative at both ends of an interval is negative on it. Both conditions can therefore hold only if $P > \frac{8}{9} b_1$. For the bound on b_2, by (A.4), $\widehat{S}_1 \le 0$ reads $b_1(b_1 + b_2)(b_2 - P) \le 8P^2 B < 8P^2(b_1 + b_2)$, using $B = b_1 + b_2 - 3P < b_1 + b_2$. Dividing by $b_1(b_1 + b_2) > 0$ gives $b_2 - P < 8P^2/b_1$, and since $P < b_1$ on Ω, $P + 8P^2/b_1 < 9b_1$.

(i). Fix P and b_1, and let b_2 range over $(2P, 4P - b_1]$. Both $\widehat{S}_2$ and $\widehat{W}$ are convex in b_2 there. Expanding (A.4), $\widehat{S}_2 = (b_1 - P) b_2^2 + \big(b_1(b_1 - P) - 8P^2\big) b_2 - 8P^2(b_1 - 3P)$ has leading coefficient $b_1 - P > 0$, and in $\widehat{W}$ the term $-16P^3 B$ is linear in b_2 while $(b_1 + b_2)^2 (b_2 - P)$ has second derivative $6b_2 + 4b_1 - 2P$, so $\partial^2 \widehat{W}/\partial b_2^2 = (b_1 - P)(6b_2 + 4b_1 - 2P) > 0$. A convex function is maximized at an endpoint, and

$$\widehat{W}\big|_{b_2 = 2P} = -P(b_1 - P)(2P - b_1)(b_1 + 6P), \qquad \widehat{W}\big|_{b_2 = 4P - b_1} = -16P^2(2P - b_1)^2,$$
$$\widehat{S}_2\big|_{b_2 = 2P} = -2P(b_1 - P)(2P - b_1), \qquad \widehat{S}_2\big|_{b_2 = 4P - b_1} = -4P(2P - b_1)(3P - b_1)$$

are all negative for $\frac{b_1}{2} < P < b_1$. For $\widehat{S}_1$, if $P \le \frac{8}{9} b_1$ the preliminary bound gives $\widehat{S}_1 > 0$. If $P > \frac{8}{9} b_1$, then $b_2^* - (4P - b_1) = (8P^2 - 7b_1 P + b_1^2)/(2b_1) > 0$, since $8P^2 - 7b_1 P + b_1^2$ is increasing for $P > \frac{7}{16} b_1$ and equals $\frac{89}{81} b_1^2$ at $P = \frac{8}{9} b_1$, so $\widehat{S}_1$ is decreasing in b_2 on $(2P, 4P - b_1]$ and bounded below by $\widehat{S}_1|_{b_2 = 4P - b_1} = 4P(2P - b_1)(b_1 - P) > 0$.

(ii). By (A.6b), $\widehat{S}_1 = \widehat{S}_2 + P(b_1+b_2)(b_2-b_1) > 0$. In (A.6a) the first term is positive because $b_1+b_2 > 4P$ and $B = (b_1-P)+(b_2-2P) > 0$, the second is positive because $b_1 > P$ and $\widehat{S}_1 > 0$, and the third is nonnegative because $b_2 > P$ and $\widehat{S}_2 \geq 0$. Hence $\widehat{W} > 0$.

(iii). Suppose $\widehat{S}_1 \leq 0$. Since the first term of (A.6a) can be positive while the other two are nonpositive, we instead write $\widehat{W}$ in terms of $\widehat{S}_1$ alone. By (A.4), $b_1(b_1+b_2)(b_2-P) = \widehat{S}_1 + 8P^2B$, so

$$\begin{aligned} b_1\widehat{W} &= (b_1+b_2)(b_1-P)\big(\widehat{S}_1 + 8P^2B\big) - 16b_1P^3B \\ &= (b_1+b_2)(b_1-P)\,\widehat{S}_1 + 8P^2B\big[(b_1+b_2)(b_1-P) - 2b_1P\big]. \end{aligned}$$

The first term is nonpositive, since $(b_1+b_2)(b_1-P) > 0$ and $\widehat{S}_1 \leq 0$. For the second, $B > 0$, and by the preliminary bound $P > \frac{8}{9}b_1$ and $b_2 < 9b_1$, so $(b_1+b_2)(b_1-P) < 10b_1 \cdot \frac{1}{9}b_1 < \frac{16}{9}b_1^2 < 2b_1P$, i.e., the bracket is negative. Hence $b_1\widehat{W} < 0$. □

Proof of Proposition 1.1. If $\Delta = 0$ the constraint forces $\gamma = 0$ at every α, so the optimum, unique by Lemma A.1(i), does not move. Assume $\Delta > 0$. We treat $b_2 > b_1$ first and $b_2 = b_1$ at the end.

Direction of the split ($b_2 > b_1$). By Lemma A.4(i), $\sigma = \operatorname{sign}(p_2^* - p_1^*) \neq 0$ because $\Delta > 0$, so $\widetilde{\mathcal{R}}_\gamma(P,0) \neq 0$ and Lemma A.3 gives $\mathcal{D}\Phi = \sigma\Delta G_\Phi$ for each Φ. Part (i) of the proposition is $\mathcal{D}\mathcal{R} > 0$ from that lemma, and part (ii) is Lemma A.4(i) and (iv).

The case $\sigma = -1$. Then $b_1+b_2 < 4P$ and Lemma A.5(i) gives $(\operatorname{sign}\widehat{W}, \operatorname{sign}\widehat{S}_1, \operatorname{sign}\widehat{S}_2) = (-,+,-)$, all strict, so by the sign rule of Lemma A.4(ii) the pattern is $(+,-,+)$, regime (e), with no vanishing entry.

The case $\sigma = +1$. Then $b_1+b_2 > 4P$, and Lemma A.5(ii),(iii) say that $\{\widehat{S}_2 > 0\} \subset \{\widehat{W} > 0\} \subset \{\widehat{S}_1 > 0\}$. A parameter point sits in exactly one of four layers of this nest:

	condition	$\operatorname{sign}\widehat{W}$	$\operatorname{sign}\widehat{S}_1$	$\operatorname{sign}\widehat{S}_2$
(a)	$\widehat{S}_2 \geq 0$	$+$	$+$	≥ 0
(b)	$\widehat{S}_2 < 0 \leq \widehat{W}$	≥ 0	$+$	$-$
(c)	$\widehat{W} < 0 \leq \widehat{S}_1$	$-$	≥ 0	$-$
(d)	$\widehat{S}_1 < 0$	$-$	$-$	$-$

In (a) the two strict entries are Lemma A.5(ii). In (b) the strict $\widehat{S}_1$ is Lemma A.5(iii). In (c), $\widehat{S}_2 \geq 0$ is excluded by Lemma A.5(ii). In (d), $\widehat{W} < 0$ is Lemma A.5(iii) and $\widehat{S}_2 < \widehat{S}_1 < 0$ by (A.6b). The sign rule of Lemma A.4(ii) also holds where a polynomial vanishes, because each prefactor is strictly positive, so multiplying by $\sigma = +1$ turns the four rows into regimes (a)–(d) of the proposition, weak inequalities included. The five regimes are mutually exclusive by their inequalities, which proves part (iii) except for existence.

Existence (part (iii)). The five points below are given for arbitrary $b_1 > 0$ and $\Lambda_1 > 0$. Each has $b_2 > b_1$ and $P \in I$, and Λ_2 is recovered from P through (A.1). Each realizes the regime named with all three entries strict, and a strict pattern persists on a neighbourhood in $(\Lambda_1, \Lambda_2, b_1, b_2)$, because σ and the three polynomials are continuous in (P, b_1, b_2) and P is continuous in $(\Lambda_1, \Lambda_2, b_1, b_2)$ by the implicit function theorem applied to (A.1), since h is

continuously differentiable with $h' < 0$ on I:

$$\begin{array}{lllll} \text{(a)} & P = \frac{13}{25}b_1, & b_2 = 4b_1, & \Lambda_2 = \frac{4736}{841}\Lambda_1, \\ \text{(b)} & P = \frac{3}{5}b_1, & b_2 = 2b_1, & \Lambda_2 = \frac{32}{49}\Lambda_1, \\ \text{(c)} & P = \frac{9}{10}b_1, & b_2 = 5b_1, & \Lambda_2 = \frac{20}{1681}\Lambda_1, \\ \text{(d)} & P = \frac{99}{100}b_1, & b_2 = 4b_1, & \Lambda_2 = \frac{404}{4439449}\Lambda_1, \\ \text{(e)} & P = \frac{7}{10}b_1, & b_2 = \frac{3}{2}b_1, & \Lambda_2 = \frac{27}{512}\Lambda_1. \end{array}$$

(At point (a) the unconstrained optimum is at the corner $p_1^* = 0$, which Lemma A.2 covers.)

The case $b_2 = b_1 =: b$. Then $D(p) = (\Lambda_1 + \Lambda_2)b(b-p)$, so the closed form of Lemma A.1(ii) becomes $\varphi(p) = 2\Lambda_1\Lambda_2\, p(b-p)/\big((\Lambda_1+\Lambda_2)b\big)$, which is maximized at $P = \frac{b}{2}$. There $b_1+b_2-4P = 0$ and all three polynomials vanish, so their signs say nothing and we evaluate (A.3) directly at $(\frac{b}{2}, 0)$ from (1.1), without substituting for Λ_2. This gives

$$s'(0) = \frac{\Lambda_2 - \Lambda_1}{4(\Lambda_1 + \Lambda_2)}, \qquad G_{\mathcal{R}} = G_{\mathcal{W}} = \frac{\Lambda_1\Lambda_2(\Lambda_2 - \Lambda_1)}{2(\Lambda_1 + \Lambda_2)^2},$$
$$G_{S_1} = \frac{\Lambda_2}{8(\Lambda_1 + \Lambda_2)}, \qquad G_{S_2} = -\frac{\Lambda_1}{8(\Lambda_1 + \Lambda_2)},$$

with the second-order condition $\widetilde{\mathcal{R}}_{ss}(P, 0) < 0$ from Lemma A.1(ii), so Lemma A.3 applies. Thus $\sigma = \operatorname{sign} \widetilde{\mathcal{R}}_\gamma(\frac{b}{2}, 0) = \operatorname{sign}(\Lambda_2 - \Lambda_1)$. By Lemma A.2 the unconstrained optimum is (A.2), which specializes to $p_2^* - p_1^* = b(\sqrt{\Lambda_2} - \sqrt{\Lambda_1})/(\sqrt{\Lambda_2} + \sqrt{\Lambda_1})$, so $\sigma = \operatorname{sign}(p_2^* - p_1^*)$ again, and $\Delta > 0$ forces $\Lambda_1 \neq \Lambda_2$. Then $\mathcal{DR} = \mathcal{DW} > 0$, $|s'(0)| < \frac{1}{4}$, so part (ii) holds because (1.3) reads $\Lambda_2 > \Lambda_1$ when $b_1 = b_2$, and the pattern is $(+, +, -)$ when $\Lambda_2 > \Lambda_1$ and $(+, -, +)$ when $\Lambda_2 < \Lambda_1$, i.e., regimes (b) and (e), with nothing vanishing. No point falls outside the five regimes in either case.

Which fare falls (part (iii)). By part (ii), location 1's fare falls exactly when $\sigma = +1$, i.e., exactly when (1.3) holds. When $b_2 > b_1$, $\sigma = -1$ gives regime (e) and $\sigma = +1$ gives one of (a)–(d), and when $b_2 = b_1$, $\sigma = -1$ gives (e) and $\sigma = +1$ gives (b). So location 1's fare falls in regimes (a)–(d) and location 2's in regime (e).

Dependence on ratios. Finally, the regime depends on the parameters only through b_2/b_1 and Λ_2/Λ_1, as Figure 1 presumes. When $b_2 > b_1$, scaling (P, b_1, b_2) by $1/b_1$ multiplies $b_1 + b_2 - 4P$, $\widehat{W}$, $\widehat{S}_1$, and $\widehat{S}_2$ by positive constants and leaves h in (A.1) unchanged, so P/b_1 solves (A.1) with $(1, b_2/b_1)$ in place of (b_1, b_2) and depends only on the two ratios. When $b_2 = b_1$, the regime is decided by the sign of $\Lambda_2 - \Lambda_1$. □

B Proofs for Section 2

Proof of Lemma 2.1. Detailed balance $\pi(k-1)\phi_2 = \pi(k)\phi_1$ gives $\pi(k) = \pi(0)\rho^k$, and summing $\pi(k)\phi_1 = \pi(k-1)\phi_2$ over $k = 1, \ldots, K$ gives (i). For (ii), $\pi(0) = 1/\sum_j \rho^j$ decreases from 1 to 0 and $\pi(K) = 1/\sum_m \rho^{-m}$ increases from 0 to 1. For (iii), if $\phi_1 = 0$ all mass sits at $k = K$, so $\Pi_2 = 0$ and $\lambda_1 = \psi_1 = 0$. The case $\phi_2 = 0$ is symmetric. □

Proof of Proposition 2.2. Since $\lambda \neq 0$ and $\psi_1\psi_2 > 0$, both ϕ_1 and ϕ_2 are positive. Put $t := \min\{\psi_1, \psi_2/\rho\} > 0$ and $\psi_1' := \psi_1 - t$, $\psi_2' := \psi_2 - \rho t$. Then $\phi_1' = \phi_1 - t$ and $\phi_2' = \rho\phi_1'$, so $\rho' = \rho$ provided $\phi_1' > 0$, and $\phi_1' = 0$ would force $t = \psi_1$, $\lambda_1 = 0$, hence $\phi_2' = 0$ and

$\lambda_2 = 0$, contradicting $\lambda \neq 0$. So the availabilities and every fairness constraint are unchanged, $\min\{\psi_1', \psi_2'\} = 0$, and by Lemma 2.1(i),

$$\mathcal{R}_K(\lambda, \psi') - \mathcal{R}_K(\lambda, \psi) = c_1 \Pi_1 t + c_2 \Pi_2 \rho t = (c_1 + c_2)\Pi_1 t > 0.$$

If $\lambda = 0$ then $\mathcal{R}_K \leq 0$ with equality at $\psi = 0$. □

Notation. In the rest of this appendix write

$$a_1(\lambda_1) := \lambda_1\big(p_1(\lambda_1) + c_1\big), \qquad a_2(\lambda_2) := \lambda_2\big(p_2(\lambda_2) - c_1\big)$$

for the two coefficients in (2.2), so that $\mathcal{R}_K = a_1(\lambda_1)\Pi_1(\rho) + a_2(\lambda_2)\Pi_2(\rho)$ on the branch $\psi_2 = 0$, $0 < \rho \leq \lambda_2/\lambda_1$. The function $a_2(\lambda_2) = \lambda_2(b_2 - c_1) - b_2\lambda_2^2/\Lambda_2$ is strictly concave with maximizer $\lambda_2^A := \Lambda_2(b_2 - c_1)^+/(2b_2)$ over $[0, \Lambda_2]$, i.e., $p_2 = (b_2 + c_1)/2$ when $b_2 > c_1$, and then $a_2(\lambda_2^A) = \overline{V}$.

Lemma B.1 (The exchange rate between availabilities)**.** *Write $Z_K'(\rho) = \sum_{j=0}^{K-1}(j+1)\rho^j$ for the derivative of Z_K. Then*

$$\kappa_K(\rho) := -\frac{\Pi_2'(\rho)}{\Pi_1'(\rho)} = \rho^{2K-2}\,\frac{Z_K'(1/\rho)}{Z_K'(\rho)}. \tag{B.1}$$

For $K = 1$, $\kappa_1 \equiv 1$. For $K \geq 2$, κ_K is strictly increasing on $(0, \infty)$.

Proof. From $\Pi_1 = 1 - 1/Z_K$ and $\Pi_2 = 1 - \rho^K/Z_K$, $\Pi_1' = Z_K'/Z_K^2$ and $\Pi_2' = -\rho^{K-1}(KZ_K - \rho Z_K')/Z_K^2$. Since $KZ_K(\rho) - \rho Z_K'(\rho) = \sum_{j=0}^{K}(K-j)\rho^j = \rho^{K-1}\sum_{i=0}^{K-1}(i+1)\rho^{-i} = \rho^{K-1}Z_K'(1/\rho)$, this gives (B.1). For $K = 1$, $Z_1' \equiv 1$ and $\kappa_1 \equiv 1$.

Let $K \geq 2$. Write $m(\rho) := \rho Z_K''(\rho)/Z_K'(\rho) = \sum_{j<K} j(j+1)\rho^j / \sum_{j<K}(j+1)\rho^j$, a weighted average of $j \in \{0, \ldots, K-1\}$ with strictly positive weights, so $0 < m(\rho) < K - 1$ for every $\rho \in (0, \infty)$. Differentiating $\log \kappa_K(\rho) = (2K-2)\log\rho + \log Z_K'(1/\rho) - \log Z_K'(\rho)$,

$$\rho\,\frac{d}{d\rho}\log\kappa_K(\rho) = (2K-2) - m(1/\rho) - m(\rho) > 0,$$

so κ_K is strictly increasing. □

Proof of Proposition 2.3. (a) With $\lambda_1 = 0$ and $\rho > 0$, flow balance reads $\Pi_1\psi_1 = \Pi_2(\lambda_2 + \psi_2)$, so from (2.1) and $\Pi_2 \leq 1$,

$$\mathcal{R}_K = \Pi_2\big[a_2(\lambda_2) - (c_1 + c_2)\psi_2\big] \leq \Pi_2\, a_2(\lambda_2) \leq \max\{0, a_2(\lambda_2)\} \leq \overline{V}.$$

With $\rho = 0$ the revenue is $a_2(\lambda_2)$, which attains $\overline{V}$ at $\lambda_2 = \lambda_2^A$, i.e., $p_2 = (b_2 + c_1)/2$. Since $p_2 \geq b_1$, posting $p_1 = p_2$ gives $\lambda_1 = 0$ and $|p_1 - p_2| = 0$, so the policy is price fair at every α.

(b) By (a) the abandonment policy is feasible at every α and earns $\overline{V}$, so $\mathcal{R}_K^*(\alpha) \geq \overline{V}$, and $\mathcal{R}_K^*$ is nonincreasing because the feasible set shrinks in α. If $\mathcal{R}_K^*(\alpha_0) = \overline{V}$ then $\overline{V} \leq \mathcal{R}_K^*(\alpha) \leq \mathcal{R}_K^*(\alpha_0) = \overline{V}$ for $\alpha \geq \alpha_0$, so $\mathcal{I}_K$ is empty or an interval with right endpoint 1. On $\mathcal{I}_K$ the value equals $\overline{V}$, the revenue of the policy of (a), which does not involve α.

(c) If $K = 1$ and abandonment is optimal at $\alpha = 0$, then $\mathcal{R}_1^*(0) = \overline{V}$, so $0 \in \mathcal{I}_1$ and $\mathcal{I}_1 = [0, 1]$. Let $K \geq 2$. Since $Z_K(\rho) = 1 + \rho + O(\rho^2)$, we have $\pi(0) = 1 - \rho + O(\rho^2)$ and $\pi(K) = \rho^K(1 + O(\rho))$, so $\Pi_1(\rho) = \rho + O(\rho^2)$ and $\Pi_2(\rho) = 1 - \rho^K + O(\rho^{K+1})$. Take $\lambda_2 = \lambda_2^A$, any $\lambda_1 \in (0, \Lambda_1)$, so that $a_1(\lambda_1) > 0$, and repositioning from location 1 to location 2. Since $K \geq 2$, (2.2) gives $\mathcal{R}_K - \overline{V} = a_1(\lambda_1)\rho + O(\rho^2) > 0$ for small $\rho > 0$. Hence $\mathcal{R}_K^*(0) > \overline{V}$, so $0 \notin \mathcal{I}_K$. □

Proof of Proposition 2.4. (a) Without repositioning write $g_K(\lambda) := \Pi_1(\lambda_2/\lambda_1)\lambda_1$ for the throughput, with $g_K := 0$ when $\lambda_1\lambda_2 = 0$, and $g_K^* := g_K(\lambda^*)$. For $K = 1$ this is the round-trip rate g of Section 1. It is continuous, positive when $\lambda_1, \lambda_2 > 0$, and homogeneous of degree one, since scaling λ leaves ρ unchanged. Revenue is $g_K(\lambda)(p_1 + p_2) = g_K(\lambda)(b_1 + b_2) - g_K(\lambda)\chi(\lambda)$ with $\chi(\lambda) := b_1\lambda_1/\Lambda_1 + b_2\lambda_2/\Lambda_2$, and by flow balance the access gap is $A_1 - A_2 = g_K(\lambda)(1/\Lambda_1 - 1/\Lambda_2) > 0$, so $\Delta = g_K^*(1/\Lambda_1 - 1/\Lambda_2)$ and the access constraint reads $g_K(\lambda) \le (1-\alpha)g_K^*$. For $t > 0$ the policies with throughput t are the points $t\varpi$ with $g_K(\varpi) = 1$, and their revenue is $t(b_1+b_2) - t^2\chi(\varpi)$. On the set $\{g_K = 1\}$ we have $\varpi_1, \varpi_2 \ge 1$, because $g_K \le \min\{\lambda_1, \lambda_2\}$, and $\chi \to \infty$ as $|\varpi| \to \infty$, so $\chi^* := \min\{\chi(\varpi) : g_K(\varpi) = 1\}$ is attained, at some ϖ^*. The best revenue at throughput t is therefore $t(b_1 + b_2) - t^2\chi^*$, increasing in t on $[0, (b_1 + b_2)/(2\chi^*)]$. The unconstrained optimum is interior by assumption, so $g_K^* = (b_1+b_2)/(2\chi^*)$ and $\lambda^* = g_K^*\varpi^*$. For $\alpha \in [0, 1]$ the constrained optimum is $t = (1-\alpha)g_K^*$, i.e., $\lambda(\alpha) = (1-\alpha)\lambda^*$. Along it ρ and the availabilities are constant, $p_i(\lambda_i(\alpha)) = (1-\alpha)p_i^* + \alpha b_i$, so A_i and S_i fall, and the revenue $(1-\alpha)g_K^*\big[(1-\alpha)(p_1^* + p_2^*) + \alpha(b_1 + b_2)\big]$ falls to 0.

(b) By Proposition 2.2 the optimum repositions in one direction. Suppose it sends vehicles out of location 1, i.e., $\psi_2 = 0 < \psi_1$, the other case following by exchanging the locations. Then $\rho < \lambda_2/\lambda_1$, and by (2.2) the optimum is a local maximizer of $a_1(\lambda_1)\Pi_1(\rho) + a_2(\lambda_2)\Pi_2(\rho)$ subject to $\sigma(A_1 - A_2) \le (1-\alpha)\Delta$, where $\sigma = 1$ if $A_1 > A_2$ and $\sigma = -1$ otherwise. Since the constraint binds and $(1-\alpha)\Delta > 0$, σ is constant on the interval. Since λ is interior and $\rho < \lambda_2/\lambda_1$, the first-order conditions hold with a multiplier $\zeta \ge 0$. Put $\beta := \sigma\zeta$ and $L = a_1\Pi_1 + a_2\Pi_2 - \beta\big[\frac{\lambda_1}{\Lambda_1}\Pi_1 - \frac{\lambda_2}{\Lambda_2}\Pi_2\big]$. The conditions $\partial_{\lambda_i} L = 0$, with $a_1' = b_1 + c_1 - 2b_1\lambda_1/\Lambda_1$ and $a_2' = b_2 - c_1 - 2b_2\lambda_2/\Lambda_2$, give the first two equations of

$$\lambda_1 = \frac{\Lambda_1(b_1 + c_1) - \beta}{2b_1}, \qquad \lambda_2 = \frac{\Lambda_2(b_2 - c_1) + \beta}{2b_2}, \qquad \frac{b_1\Lambda_2\,\lambda_1^2}{b_2\Lambda_1\,\lambda_2^2} = \kappa_K(\rho). \tag{B.2}$$

Then $\partial_\rho L = \Pi_1'\Xi_1 + \Pi_2'\Xi_2$ with $\Xi_1 = \lambda_1(p_1 + c_1 - \beta/\Lambda_1) = b_1\lambda_1^2/\Lambda_1$ and $\Xi_2 = \lambda_2(p_2 - c_1 + \beta/\Lambda_2) = b_2\lambda_2^2/\Lambda_2$ by the first two equations, and setting it to zero gives the third, with κ_K as in (B.1).

Let $K \ge 2$. As β increases, λ_1 falls and λ_2 rises by the first two equations, so the left side of the third falls, and ρ falls since κ_K is strictly increasing (Lemma B.1). By Lemma 2.1(ii), Π_1 falls and Π_2 rises, so $A_1 = \Pi_1\lambda_1/\Lambda_1$ and $S_1 = \Pi_1 b_1\lambda_1^2/(2\Lambda_1^2)$ fall, A_2 and S_2 rise, p_1 rises, and p_2 falls. In particular $A_1 - A_2$ is strictly decreasing in β, so the binding constraint $\sigma(A_1 - A_2) = (1-\alpha)\Delta$ determines β as a function of α. If $\sigma = 1$, the right side falls as α increases, so β rises, and location 1, the location with higher access, loses access and surplus and its price rises, while location 2 gains and its price falls. If $\sigma = -1$, β falls and the roles are exchanged.

Let $K = 1$. Then $\kappa_1 \equiv 1$ by Lemma B.1, and the three equations of (B.2) determine β, hence λ and the prices, without reference to α. The gap $A_1 - A_2 = \Pi_1(\rho)\lambda_1/\Lambda_1 - \Pi_2(\rho)\lambda_2/\Lambda_2$ is strictly increasing in ρ, so the binding constraint makes ρ fall as α increases if $\sigma = 1$, and rise if $\sigma = -1$. With λ fixed, access and surplus at each location move with its availability, so again the location with higher access loses and the other gains.

(c) Consider a policy with $\psi_2 = 0$, $\lambda_1, \lambda_2 > 0$, and equal access, $\Pi_1\lambda_1/\Lambda_1 = \Pi_2\lambda_2/\Lambda_2$. By Lemma 2.1(i) this fixes $\rho = \Pi_1/\Pi_2 = \lambda_2\Lambda_1/(\lambda_1\Lambda_2) < \lambda_2/\lambda_1$, so the policy lies on the branch, with $\Pi_1\psi_1 = \Pi_2\lambda_2 - \Pi_1\lambda_1 = \Pi_1\lambda_1(\Lambda_2/\Lambda_1 - 1) \ge 0$. Substituting both relations into (2.1),

$$\mathcal{R}_K = \frac{\Pi_1\lambda_1}{\Lambda_1}\Big[\Lambda_1 p_1 + \Lambda_2 p_2 - c_1(\Lambda_2 - \Lambda_1)\Big].$$

Under (2.3), letting $\lambda_1, \lambda_2 \downarrow 0$ sends $p_i \uparrow b_i$, so the bracket is positive for small $\lambda_1, \lambda_2 > 0$, and $\Pi_1 \lambda_1 > 0$, hence $\mathcal{R}_K^*(1) > 0$. Conversely, if (2.3) fails, the bracket is nonpositive for every $p_i \le b_i$, so every such policy earns at most 0. On the other branch, $\psi_1 = 0$, flow balance and equal access give $\Pi_1 \lambda_1 = \Pi_2(\lambda_2 + \psi_2) \ge \Pi_2 \lambda_2 = \Pi_1 \lambda_1 \Lambda_2 / \Lambda_1$, which forces $\lambda_1 = 0$ and then $\lambda_2 = 0$ since $\Lambda_2 > \Lambda_1$. Policies with $\phi_1 = 0$ or $\phi_2 = 0$ earn 0 by Lemma 2.1(iii). A policy with $\rho = 0$ has $A_1 = 0$, so equal access forces $\lambda_2 = 0$ and revenue 0, and similarly for $\rho = \infty$. Hence $\mathcal{R}_K^*(1) = 0$. Finally, under (2.3), $\mathcal{R}_K^*$ is nonincreasing in α, so $\mathcal{R}_K^*(\alpha) \ge \mathcal{R}_K^*(1) > 0$, while the best revenue without repositioning tends to 0 by (a). So for α close to 1 the optimum repositions. □

C Proofs for Section 3

C.1 Proofs for Section 3.1

By Assumption 3.1(a), $\bar{F}_{ij}$ is a continuous strictly decreasing bijection from $[0, \hat{p}_{ij}]$ onto $\mathcal{Q}_{ij}$, so $\bar{F}_{ij}^{-1}$ and r_{ij} are well defined and continuous on $\mathcal{Q}_{ij}$, with $\bar{F}_{ij}^{-1}(0) = b_{ij}$ and $r_{ij}(0) = 0$ when some admissible price closes the arc. The perspective $\tilde{r}_{ij}$ is defined on $\mathcal{D}_{ij} := \{(z, \Pi) : 0 \le \Pi \le 1,\ \bar{F}_{ij}(d_{ij} p_{\max})\Pi \le z \le \bar{F}_{ij}(0)\Pi\}$. Because $0 \le \bar{F}_{ij} \le 1$, the window constraint already implies $0 \le z_{ij} \le \Pi_i$ and $(z_{ij}, \Pi_i) \in \mathcal{D}_{ij}$, so these need not be imposed separately. When $d_{ij}\bar{p} \ge b_{ij}$, the window constraint reads $0 \le z_{ij} \le \bar{F}_{ij}(d_{ij}(\bar{p} - W_\alpha))\Pi_i$, and the program may close arc (i, j) by setting $z_{ij} = 0$.

Proof of Proposition 3.2. Step 1 shows that, whenever $z_{ij} = \Pi_i \bar{F}_{ij}(p_{ij})$ on every arc, the two problems have the same balance and fleet constraints and the same objective value, so only the fairness constraint remains, which is $p \in \mathcal{P}_\alpha$ in (3.1) and the window constraint in (3.3). Steps 2 and 3 convert feasible points in the two directions, Step 4 shows that the maxima are attained, and Step 5 shows convexity.

Step 1: matching constraints and objectives. Let $p \in \mathcal{P}$, $\Pi \in [0,1]^N$, $y \in \mathbb{R}_+^{\mathcal{V}}$, and let $z_{ij} = \Pi_i \bar{F}_{ij}(p_{ij})$ for all $(i, j) \in \mathcal{A}$. Then $(z_{ij}, \Pi_i) \in \mathcal{D}_{ij}$, since $\bar{F}_{ij}(d_{ij} p_{\max}) \le \bar{F}_{ij}(p_{ij}) \le \bar{F}_{ij}(0)$. The balance and fleet constraints of (3.1) at (p, Π, y) are those of (3.3) at (z, Π, y), since both are written in terms of $\Pi_i \bar{F}_{ij}(p_{ij}) = z_{ij}$ and y, and the repositioning costs are the same in both objectives. It remains to show that $\tilde{r}_{ij}(z_{ij}, \Pi_i) = \Pi_i \bar{F}_{ij}(p_{ij})\, p_{ij}$ on every arc. If $\Pi_i > 0$ and $\bar{F}_{ij}(p_{ij}) > 0$, then $p_{ij} < b_{ij}$, so p_{ij} lies in $[0, \hat{p}_{ij}]$, the interval on which $\bar{F}_{ij}$ is strictly decreasing and $\bar{F}_{ij}^{-1}$ is defined, hence $\tilde{r}_{ij}(z_{ij}, \Pi_i) = \Pi_i\, r_{ij}(\bar{F}_{ij}(p_{ij})) = \Pi_i \bar{F}_{ij}(p_{ij})\, \bar{F}_{ij}^{-1}(\bar{F}_{ij}(p_{ij})) = \Pi_i \bar{F}_{ij}(p_{ij})\, p_{ij}$. If $\Pi_i > 0$ and $\bar{F}_{ij}(p_{ij}) = 0$, then $z_{ij} = 0$ and both sides are zero, since $\tilde{r}_{ij}(0, \Pi_i) = \Pi_i\, r_{ij}(0) = 0$. If $\Pi_i = 0$, then $z_{ij} = 0$ and both sides are zero, since $\tilde{r}_{ij}(0, 0) = 0$.

Step 2: from (3.1) *to* (3.3). Let (p, Π, y) be feasible for (3.1). By (3.2) there is $\bar{p} \in \bar{\mathcal{P}}_\alpha$ with $d_{ij}(\bar{p} - W_\alpha) \le p_{ij} \le d_{ij}\bar{p}$ on every arc. Set $z_{ij} := \Pi_i \bar{F}_{ij}(p_{ij})$. Since $\bar{F}_{ij}$ is nonincreasing, $\bar{F}_{ij}(d_{ij}\bar{p}) \le \bar{F}_{ij}(p_{ij}) \le \bar{F}_{ij}(d_{ij}(\bar{p} - W_\alpha))$, and multiplying by $\Pi_i \ge 0$ gives the window constraint of (3.3). By Step 1, (z, Π, y) also satisfies the balance and fleet constraints of (3.3) and has the same objective value as (p, Π, y). Hence the objective of (3.1) at (p, Π, y) is at most $V_\alpha(\bar{p})$, and $J_K^*(\alpha) \le \sup_{\bar{p}} V_\alpha(\bar{p})$.

Step 3: from (3.3) *to* (3.1). Let $\bar{p} \in \bar{\mathcal{P}}_\alpha$, let (z, Π, y) be feasible for (3.3), and define p by (3.4). We show that on every arc $d_{ij}(\bar{p} - W_\alpha) \le p_{ij} \le d_{ij}\bar{p}$ and $\Pi_i \bar{F}_{ij}(p_{ij}) = z_{ij}$. If $z_{ij} > 0$, then $\Pi_i > 0$ and, by the window constraint, $q := z_{ij}/\Pi_i$ satisfies $\bar{F}_{ij}(d_{ij}\bar{p}) \le q \le \bar{F}_{ij}(d_{ij}(\bar{p} - W_\alpha))$, so $q \in \mathcal{Q}_{ij}$ and $q > 0$. Then $p_{ij} = \bar{F}_{ij}^{-1}(q)$ satisfies $\bar{F}_{ij}(p_{ij}) = q$, i.e., $\Pi_i \bar{F}_{ij}(p_{ij}) = z_{ij}$, and

$\bar{F}_{ij}(d_{ij}\bar{p}) \le \bar{F}_{ij}(p_{ij}) \le \bar{F}_{ij}(d_{ij}(\bar{p} - W_\alpha))$ with $\bar{F}_{ij}(p_{ij}) > 0$. Since $\bar{F}_{ij}$ is nonincreasing, and strictly decreasing where it is positive, this forces $d_{ij}(\bar{p} - W_\alpha) \le p_{ij} \le d_{ij}\bar{p}$. If $z_{ij} = 0$, then $p_{ij} = d_{ij}\bar{p}$ lies in the window, and the window constraint gives $0 \le \Pi_i \bar{F}_{ij}(d_{ij}\bar{p}) \le z_{ij} = 0$, so $\Pi_i \bar{F}_{ij}(p_{ij}) = z_{ij}$. Every normalized price therefore lies in $[\bar{p} - W_\alpha, \bar{p}] \subseteq [0, p_{\max}]$, so $p \in \mathcal{P}$ and, by (3.2), $p \in \mathcal{P}_\alpha$. By Step 1, (p, Π, y) is feasible for (3.1) and has the same objective value as (z, Π, y). This is the statement about (3.4), and it gives $V_\alpha(\bar{p}) \le J_K^*(\alpha)$ for every $\bar{p}$.

Step 4: attainment. Steps 2 and 3 show that $J_K^*(\alpha)$ equals the supremum of $V_\alpha(\bar{p})$ over $\bar{p} \in \bar{\mathcal{P}}_\alpha$. First, $J_K^*(\alpha)$ is attained, since the feasible set of (3.1) is compact (the fleet constraint bounds y) and contains $(p, 0, 0)$ for every $p \in \mathcal{P}_\alpha$, and the objective is continuous. Second, $V_\alpha(\bar{p})$ is attained because (3.3) maximizes a continuous function over a nonempty compact set. The set is cut out by linear constraints, contains $(0, 0, 0)$, and is bounded by $\Pi \in [0, 1]^N$, $0 \le z_{ij} \le \Pi_i$, and the fleet constraint. Each $\tilde{r}_{ij}$ is continuous on $\mathcal{D}_{ij}$, where $\Pi > 0$ because r_{ij} is continuous on the compact interval $\mathcal{Q}_{ij}$, and at $(0, 0)$, where z/Π is undefined, because $|\tilde{r}_{ij}(z, \Pi)| \le \Pi \max_{\mathcal{Q}_{ij}} |r_{ij}| \to 0$. Third, applying Step 2 to an optimal solution of (3.1) gives $\bar{p}^*$ with $J_K^*(\alpha) \le V_\alpha(\bar{p}^*)$, so the maximum over $\bar{p}$ is attained at $\bar{p}^*$.

Step 5: convexity. Fix $\bar{p} \in \bar{\mathcal{P}}_\alpha$. Step 4 shows that the feasible set of (3.3) is a nonempty compact polyhedron and that its objective is continuous, so it remains to show that the objective is concave. The objective is a nonnegative combination of the functions $\tilde{r}_{ij}(z_{ij}, \Pi_i)$ plus a linear term in y, so it suffices that each $\tilde{r}_{ij}$ is concave on the convex polyhedron $\mathcal{D}_{ij}$. If two points of $\mathcal{D}_{ij}$ both have $\Pi > 0$, so does every point between them, and there $\tilde{r}_{ij}$ is the perspective of the concave function r_{ij}, hence jointly concave (Boyd and Vandenberghe, 2004, §3.2.6). The only point of $\mathcal{D}_{ij}$ with $\Pi = 0$ is $(0, 0)$, since $\bar{F}_{ij}(d_{ij}p_{\max})\Pi \le z \le \bar{F}_{ij}(0)\Pi$ on $\mathcal{D}_{ij}$. On the segment from $(0, 0)$ to a point $(z, \Pi) \in \mathcal{D}_{ij}$ with $\Pi > 0$, the ratio z/Π does not change, so for $t \in [0, 1]$

$$\tilde{r}_{ij}(tz, t\Pi) = t\Pi\, r_{ij}(z/\Pi) = t\,\tilde{r}_{ij}(z, \Pi) = t\,\tilde{r}_{ij}(z, \Pi) + (1 - t)\,\tilde{r}_{ij}(0, 0),$$

the last step because $\tilde{r}_{ij}(0, 0) = 0$. This is the concavity inequality on that segment, with equality. Hence $\tilde{r}_{ij}$ is concave on all of $\mathcal{D}_{ij}$, including $(0, 0)$. □

C.2 Proofs for Section 3.2

The general product form. We recall the product form of a closed network with single-server and infinite-server nodes, and then apply it to the network of Section 3.2. For finite disjoint sets $\mathcal{N}_1$ and $\mathcal{N}_\infty$ of single-server and infinite-server nodes, weights $w \in (0, \infty)^{\mathcal{N}_1}$ and $\varrho \in (0, \infty)^{\mathcal{N}_\infty}$, and integers $m \ge 0$, let $\mathcal{X}_m := \{(k, t) \in \mathbb{Z}_+^{\mathcal{N}_1} \times \mathbb{Z}_+^{\mathcal{N}_\infty} : \sum_n k_n + \sum_e t_e = m\}$ be the set of states with m vehicles, where k_n and t_e are the numbers of vehicles at the nodes, and

$$\pi_m(k, t) := \frac{1}{Z_m} \prod_{n \in \mathcal{N}_1} w_n^{k_n} \prod_{e \in \mathcal{N}_\infty} \frac{\varrho_e^{t_e}}{t_e!}, \qquad Z_m := \sum_{(k,t) \in \mathcal{X}_m} \prod_{n \in \mathcal{N}_1} w_n^{k_n} \prod_{e \in \mathcal{N}_\infty} \frac{\varrho_e^{t_e}}{t_e!}.$$

If $\mathcal{N}_1 \cup \mathcal{N}_\infty \ne \emptyset$ then $Z_m > 0$, because a state with all m vehicles at one node has positive weight, and $Z_0 = 1$. Write $\mathbb{E}_m$ for the expectation under π_m, (X, T) for the coordinates, and set

$$\omega_K := \frac{Z_{K-1}}{Z_K}, \qquad \nu := \sum_{e \in \mathcal{N}_\infty} \varrho_e.$$

Grouping the states by the number j of vehicles at single-server nodes, and using the multinomial theorem for the infinite-server nodes, gives

$$Z_m = \sum_{j=0}^{m} \ell_j \, \frac{\nu^{m-j}}{(m-j)!}, \qquad \ell_j := \sum_{k \in \mathbb{Z}_+^{\mathcal{N}_1} : \sum_n k_n = j} \; \prod_{n \in \mathcal{N}_1} w_n^{k_n}. \tag{C.1}$$

Buzen's recursion computes $\ell_0, \dots, \ell_K$ by adding one single-server node at a time. With $\mathcal{N}_1 = \{n_1, \dots, n_L\}$, let $\ell_j^{(0)} := \mathbf{1}\{j = 0\}$ and $\ell_j^{(l)} := \ell_j^{(l-1)} + w_{n_l} \ell_{j-1}^{(l)}$ for $0 \le j \le K$ and $1 \le l \le L$, with $\ell_{-1}^{(l)} := 0$. Then $\ell_j^{(L)} = \ell_j$, because $\sum_{j \ge 0} \ell_j x^j = \prod_{n \in \mathcal{N}_1} (1 - w_n x)^{-1}$ as formal power series and dividing a series by $1 - wx$ is exactly this recursion. So $\ell_0, \dots, \ell_K$, and with them Z_{K-1}, Z_K, and ω_K, are computed exactly in $O(|\mathcal{N}_1| K)$ operations.

Lemma C.1 (Product form, Gordon and Newell, 1967; Baskett et al., 1975; Reiser and Lavenberg, 1980; Kelly, 1979)**.**

(i) For every $K \ge 1$, $n \in \mathcal{N}_1$, $e \in \mathcal{N}_\infty$, and $1 \le j \le K$,

$$\pi_K(X_n \ge j) = \frac{w_n^j Z_{K-j}}{Z_K}, \qquad \mathbb{E}_K[T_e] = \omega_K \varrho_e, \qquad \mathbb{E}_K[X_n] = \omega_K w_n \big(1 + \mathbb{E}_{K-1}[X_n]\big),$$

and in particular $\pi_K(X_n \ge 1) = \omega_K w_n$.

(ii) Let K vehicles move among the nodes $\mathcal{N}_1 \cup \mathcal{N}_\infty$ as follows. A node $n \in \mathcal{N}_1$ releases a vehicle at rate $\lambda_n > 0$ while it holds at least one, a node $e \in \mathcal{N}_\infty$ releases each of its vehicles at rate $\mu_e > 0$, and a released vehicle moves to node n' with probability $R_{nn'}$, where R is an irreducible stochastic matrix with zero diagonal. Let $v \in (0, \infty)^{\mathcal{N}_1 \cup \mathcal{N}_\infty}$ solve $v_{n'} = \sum_n v_n R_{nn'}$ for all n', and set $w_n := v_n / \lambda_n$ and $\varrho_e := v_e / \mu_e$. Then the chain of node occupancies is irreducible on $\mathcal{X}_K$ with stationary distribution π_K, and every node n releases vehicles at long-run rate $\omega_K v_n$.

Part (ii) takes the times spent at infinite-server nodes to be exponential. If instead each visit to $e \in \mathcal{N}_\infty$ lasts an independent random time with mean $1/\mu_e$, the node occupancies still have the stationary distribution π_K, the long-run averages are its expectations, and every node n still releases vehicles at long-run rate $\omega_K v_n$, because infinite-server nodes are insensitive to the distribution beyond its mean (Baskett et al., 1975; Barbour, 1976).

Proof. (i) Let $\mathbf{1}_n$ denote the unit vector of node n. The map $\mathbf{s} \mapsto \mathbf{s} - \mathbf{1}_n$ is a bijection from $\{\mathbf{s} \in \mathcal{X}_K : k_n \ge 1\}$ onto $\mathcal{X}_{K-1}$ that divides the weight by w_n, so $\pi_K(X_n \ge 1) = w_n Z_{K-1} / Z_K$, and iterating gives $\pi_K(X_n \ge j) = w_n^j Z_{K-j} / Z_K$. Hence

$$\mathbb{E}_K[X_n] = \sum_{j=1}^{K} \frac{w_n^j Z_{K-j}}{Z_K} = \frac{w_n Z_{K-1}}{Z_K} \Big(1 + \sum_{j=1}^{K-1} \frac{w_n^j Z_{K-1-j}}{Z_{K-1}}\Big) = \omega_K w_n \big(1 + \mathbb{E}_{K-1}[X_n]\big).$$

Similarly $\mathbf{s} \mapsto \mathbf{s} - \mathbf{1}_e$ maps $\{t_e \ge 1\}$ onto $\mathcal{X}_{K-1}$, and $t_e \varrho_e^{t_e} / t_e! = \varrho_e \cdot \varrho_e^{t_e - 1} / (t_e - 1)!$ gives $\mathbb{E}_K[T_e] = \varrho_e Z_{K-1} / Z_K$.

(ii) The chain is irreducible because every vehicle can be moved along a path of positive routing probabilities, one vehicle at a time. For stationarity, fix $\mathbf{s} \in \mathcal{X}_K$ and a node n that holds a vehicle in $\mathbf{s}$. Put $\varsigma_n(\mathbf{s}) := w_n$ if $n \in \mathcal{N}_1$ and $\varsigma_n(\mathbf{s}) := \varrho_n / t_n$ if $n \in \mathcal{N}_\infty$. For $n' \ne n$,

the state $\mathbf{s} - \mathbf{1}_n + \mathbf{1}_{n'}$ moves to $\mathbf{s}$ when n' releases a vehicle that goes to n. In that state n' releases at rate $\lambda_{n'}$ if $n' \in \mathcal{N}_1$ and $\mu_{n'}(t_{n'}+1)$ if $n' \in \mathcal{N}_\infty$, and in both cases $\pi_K(\mathbf{s} - \mathbf{1}_n + \mathbf{1}_{n'}) \times (\text{release rate of } n') = \pi_K(\mathbf{s})\, v_{n'}/\varsigma_n(\mathbf{s})$. So the rate into $\mathbf{s}$ by an arrival at n is $\pi_K(\mathbf{s}) \sum_{n'} v_{n'} R_{n'n}/\varsigma_n(\mathbf{s}) = \pi_K(\mathbf{s})\, v_n/\varsigma_n(\mathbf{s})$, which is $\pi_K(\mathbf{s})\lambda_n$ if $n \in \mathcal{N}_1$ and $\pi_K(\mathbf{s})\mu_n t_n$ if $n \in \mathcal{N}_\infty$, i.e., the rate out of $\mathbf{s}$ by a release from n. Every transition moves one vehicle from one node to another, so summing over n gives global balance. Finally, by part (i), a node $n \in \mathcal{N}_1$ releases at rate $\lambda_n \pi_K(X_n \geq 1) = \omega_K v_n$ and a node $e \in \mathcal{N}_\infty$ at rate $\mu_e \mathbb{E}_K[T_e] = \omega_K v_e$. □

The implemented network. Under (p^H, θ^H) the vehicles move through a closed network with three kinds of nodes, the locations, where vehicles wait, and the loaded and the empty trips, written $(i,j)^{\mathrm{L}}$ and $(i,j)^{\mathrm{E}}$. A vehicle waiting at i leaves when a customer accepts, i.e., at rate $\lambda_i := \sum_j \lambda_{ij}$, and goes to j with probability λ_{ij}/λ_i. A vehicle that finishes a trip at j, loaded or empty, starts an empty trip to $k \neq j$ with probability θ^H_{jk}, and otherwise waits at j. These routing probabilities do not depend on where the other vehicles are, as Lemma C.1(ii) requires. This is why repositioning is decided when a trip ends, since a rule that moves waiting vehicles according to the current inventories would break this. Let

$$v_i := \sum_j \Lambda_{ij} \hat{z}_{ij} \quad (\text{location } i), \qquad v^{\mathrm{L}}_{ij} := \Lambda_{ij} \hat{z}_{ij} \quad (\text{loaded trip}), \qquad v^{\mathrm{E}}_{ij} := \hat{y}_{ij} \quad (\text{empty trip}) \tag{C.2}$$

be the rates at which vehicles enter the nodes in the feasible point, called the visit rates, and let $\mathcal{N}^+$ be the set of nodes with $v > 0$, i.e., the nodes that the vehicles use in the feasible point. We apply the product form with the locations and the trips in $\mathcal{N}^+$ as $\mathcal{N}_1$ and $\mathcal{N}_\infty$, $w_i = \hat{\Pi}_i$, and $\varrho = v/\mu$, which the next lemma justifies. Then $|\mathcal{N}_1| \leq N$ and ν equals the *workload* of the feasible point,

$$\hat{\nu} := \sum_{(i,j)\in\mathcal{A}} \frac{\Lambda_{ij}\hat{z}_{ij}}{\mu_{ij}} + \sum_{(i,j)\in\mathcal{V}} \frac{\hat{y}_{ij}}{\mu_{ij}},$$

i.e., the expected number of vehicles in transit in it. The weight of location i is then $\hat{\Pi}_i^{k_i}$, its availability in the feasible point raised to the number of vehicles waiting there, and this is the point of the prices (3.4), which make the ratio v_i/λ_i of the rate at which vehicles enter location i to the rate at which it releases one equal to $\hat{\Pi}_i$. The weight of a trip is $(v/\mu)^t/t!$, with the factorial because the vehicles on a trip move independently, whereas a location releases its vehicles one at a time.

Lemma C.2 (Implemented network)**.** *Let $(\hat{z}, \hat{\Pi}, \hat{y})$ be feasible for* (3.3)*. Then*

(a) the rates (C.2) *solve the traffic equations $v_{n'} = \sum_n v_n R_{nn'}$ of the routing under (p^H, θ^H),*

(b) no vehicle leaves $\mathcal{N}^+$, $\mathcal{N}^+$ is the disjoint union of finitely many closed communicating classes of the routing, and v restricted to each class solves the traffic equations of that class,

(c) every location $i \in \mathcal{N}^+$ has $\hat{\Pi}_i > 0$ and $v_i/\lambda_i = \hat{\Pi}_i$.

Proof. Let $\hat{\Pi}_i > 0$. For an arc (i,j) with $\hat{z}_{ij} > 0$, $\bar{F}_{ij}(p^H_{ij}) = \hat{z}_{ij}/\hat{\Pi}_i$ by (3.4). For an arc with $\hat{z}_{ij} = 0$, the window constraint gives $\bar{F}_{ij}(d_{ij}\bar{p})\hat{\Pi}_i \leq 0$, so $\bar{F}_{ij}(d_{ij}\bar{p}) = 0$ and $p^H_{ij} = d_{ij}\bar{p}$ closes the arc. In both cases $\lambda_{ij} = \Lambda_{ij}\hat{z}_{ij}/\hat{\Pi}_i$, hence $\lambda_i = v_i/\hat{\Pi}_i$. A location with $v_i > 0$ has some $\hat{z}_{ij} > 0$, hence $\hat{\Pi}_i > 0$, which proves (c).

(a) A loaded trip (i,j) receives vehicles only from location i, at rate $v_i\lambda_{ij}/\lambda_i = \hat{\Pi}_i\lambda_{ij} = \Lambda_{ij}\hat{z}_{ij}$ if $\hat{\Pi}_i > 0$, and $0 = \Lambda_{ij}\hat{z}_{ij}$ otherwise. Trips into j deliver vehicles at the total rate η_j. An empty trip (j,k) receives $\eta_j\theta^H_{jk} = \hat{y}_{jk}$, and location j receives $\eta_j(1-\sum_k\theta^H_{jk}) = \eta_j - \sum_k\hat{y}_{jk} = \sum_k\Lambda_{jk}\hat{z}_{jk} = v_j$ by the balance constraint (both sides vanish if $\eta_j = 0$).

(b) If $v_n > 0$ and $R_{nn'} > 0$ then $v_{n'} \geq v_nR_{nn'} > 0$ by (a), so no vehicle leaves $\mathcal{N}^+$. Restricted to $\mathcal{N}^+$, the routing is a finite Markov chain with the strictly positive invariant measure v. An invariant measure vanishes on transient states, so every node of $\mathcal{N}^+$ is recurrent and $\mathcal{N}^+$ splits into closed communicating classes. No node outside a class routes into it, so v restricted to each class solves the traffic equations of that class. □

The size of the scaling factor. The scaling factor depends on the network only through the weights, and the following lemma bounds it by the number of locations, the fleet size, and the workload of the feasible point.

Lemma C.3 (Lower bounds on ω_K)**.** *If $\mathcal{N}^+ \neq \emptyset$, the ratios $\omega_m := Z_{m-1}/Z_m$, $m \geq 1$, are nondecreasing in m, and*

$$\omega_K \;\geq\; \max\Big\{\frac{K}{N+K-1+\hat{\nu}},\; 1-\sqrt{\frac{N}{K}}\Big\}.$$

The second bound does not depend on $\hat{\nu}$, which matters when the fleet constraint of (3.3) binds, so that $\hat{\nu} = K$ and the first bound is close to $\frac{1}{2}$.

Proof. By Lemma C.2(c), $w_i = \hat{\Pi}_i \in (0,1]$, there are at most N locations, so $|\mathcal{N}_1| \leq N$, and $\nu = \hat{\nu} \leq K$ by the fleet constraint of (3.3). This is the only place where the constraint $\Pi \leq 1$ is used, and it is essential.

Monotonicity. By (C.1), $(Z_m)_{m\geq 0}$ is the convolution of the geometric sequences $(w_n^j)_{j\geq 0}$, $n \in \mathcal{N}_1$, and, if $\mathcal{N}_\infty \neq \emptyset$, the sequence $(\nu^j/j!)_{j\geq 0}$. All of these are positive and log-concave, the last because $(\nu^j/j!)^2 \geq (\nu^{j-1}/(j-1)!)\,(\nu^{j+1}/(j+1)!)$. The coefficients of a product of polynomials with positive log-concave coefficients are log-concave (Hoggar, 1974), and truncating every sequence at index $m+1$ leaves $Z_0,\dots,Z_{m+1}$ unchanged. Hence $Z_m^2 \geq Z_{m-1}Z_{m+1}$ for every $m \geq 1$, i.e., $\omega_m \leq \omega_{m+1}$.

First bound. Every state in $\mathcal{X}_K$ holds K vehicles, so by Lemma C.1(i)

$$K = \sum_n \mathbb{E}_K[X_n] + \sum_e \mathbb{E}_K[T_e] = \omega_K\Big(\sum_n w_n + \sum_n w_n\,\mathbb{E}_{K-1}[X_n] + \nu\Big) \leq \omega_K\,(N+K-1+\hat{\nu}),$$

because $w_n \leq 1$, $\sum_n \mathbb{E}_{K-1}[X_n] \leq K-1$, $|\mathcal{N}_1| \leq N$, and $\nu = \hat{\nu}$.

Second bound. If $\omega_K \geq 1$ there is nothing to prove, so let $\omega_K < 1$. By Lemma C.1(i), monotonicity, and $w_n \leq 1$,

$$\pi_K(X_n \geq j) = \frac{w_n^j Z_{K-j}}{Z_K} = w_n^j\,\omega_K\omega_{K-1}\cdots\omega_{K-j+1} \leq \omega_K^j, \qquad 1 \leq j \leq K.$$

Hence $\mathbb{E}_K[X_n] = \sum_{j=1}^K \pi_K(X_n \geq j) \leq \omega_K/(1-\omega_K)$, and with $\mathbb{E}_K[T_e] = \omega_K\varrho_e$ and $\nu \leq K$,

$$K = \sum_n \mathbb{E}_K[X_n] + \sum_e \mathbb{E}_K[T_e] \leq \frac{|\mathcal{N}_1|\,\omega_K}{1-\omega_K} + \omega_K K.$$

So $K(1-\omega_K)^2 \leq |\mathcal{N}_1|\,\omega_K \leq N$, which is the claim. □

Proposition C.4 (From a feasible point to a policy)**.** *Let Assumption 3.1(a) hold and $K \geq 1$. For every $\bar{p} \in \bar{\mathcal{P}}_\alpha$ and every feasible point of* (3.3)*, the policy (p^H, θ^H) is α-fair, i.e., $p^H \in \mathcal{P}_\alpha$ and $\theta^H \in \Theta$, and*

$$\omega_K \hat{J} \;\leq\; \mathcal{R}_K(p^H, \theta^H) \;\leq\; \mathcal{R}_K^*(\alpha) \;\leq\; J_K^*(\alpha),$$

where $\omega_K \hat{J}$ is read as 0 if $\mathcal{N}^+ = \emptyset$. The first inequality is an equality when $\mathcal{N}^+$ is a single closed class and all vehicles start in $\mathcal{N}^+$.

If $\hat{J} \geq 0$, Lemma C.3 turns the first inequality into the a priori bound $\max\{K/(N+K-1+\hat{\nu}),\, 1-\sqrt{N/K}\}\,\hat{J} \leq \mathcal{R}_K(p^H, \theta^H)$, and since the proposition holds for any feasible point of (3.3), an approximate solution of the convex program is enough.

Proof. Last inequality. Let $p \in \mathcal{P}_\alpha$, $\theta \in \Theta$, and an initial state $\mathbf{s}$ be given, and write $\mathbb{E}$ for $\mathbb{E}_{\mathbf{s}}$. For $t \geq 0$ let $L_{ij}(t)$, $(i,j) \in \mathcal{A}$, and $Y_{ij}(t)$, $(i,j) \in \mathcal{V}$, be the numbers of loaded and empty trips from i to j started in $(0,t]$, let $X_i(t)$ be the number of vehicles available at i, and let $Q_{ij}(t)$ be the number of vehicles in transit from i to j, loaded or empty, at time t. Set $L_{ij} := 0$ for $(i,j) \notin \mathcal{A}$ and $Y_{ii} := 0$.

Loaded trips. Customers for (i,j) arrive by a Poisson process, independently of the past, and each one starts a loaded trip exactly when it accepts the price and a vehicle is available at i. Hence $\mathbb{E}[L_{ij}(t)] = \Lambda_{ij}\bar{F}_{ij}(p_{ij}) \int_0^t \mathbb{P}(X_i(u) \geq 1)\, du$.

Balance. Every trip into i that ends delivers a vehicle that is repositioned at once or stays, and every loaded trip out of i takes an available vehicle. Counting the vehicles that reach and leave i in $(0,t]$ therefore gives

$$\sum_j \big(L_{ij}(t) + Y_{ij}(t)\big) - \sum_j \big(L_{ji}(t) + Y_{ji}(t)\big) = K_i^{\mathbf{s}} - X_i(t) - \sum_j Q_{ji}(t),$$

where $K_i^{\mathbf{s}}$ is the number of vehicles that $\mathbf{s}$ places at i or on trips into i. The right side lies in $[-K, K]$.

Fleet. Fix $t_0 > 0$ and let $m_{ij}(t_0) := \mathbb{E}[\min\{\xi_{ij}, t_0\}] > 0$, where ξ_{ij} has the trip-time distribution of (i,j). A trip from i to j started in $(0,t]$ is in transit for at least $\min\{\xi, t_0\}$ time units in $[0, t+t_0]$, where ξ is its duration, and ξ is drawn independently of the start time. Hence $m_{ij}(t_0)\,\mathbb{E}[L_{ij}(t) + Y_{ij}(t)] \leq \mathbb{E}\int_0^{t+t_0} Q_{ij}(u)\, du$, and the right sides sum to at most $K(t+t_0)$ over all pairs, since at most K vehicles are in transit at any time. In particular $\mathbb{E}[Y_{ij}(t)]/t$ is bounded for $t \geq 1$.

Limits. Choose $t_n \to \infty$ along which $t_n^{-1}\mathbb{E}[\text{fares} - \text{costs in } [0,t_n]]$ converges to $\mathcal{R}_K(p, \theta \mid \mathbf{s})$ and, passing to a subsequence, $t_n^{-1}\int_0^{t_n} \mathbb{P}(X_i(u) \geq 1)\, du \to \Pi_i \in [0,1]$ for all i and $t_n^{-1}\mathbb{E}[Y_{ij}(t_n)] \to y_{ij} \geq 0$ for all $(i,j) \in \mathcal{V}$. Then $t_n^{-1}\mathbb{E}[L_{ij}(t_n)] \to \Lambda_{ij}\bar{F}_{ij}(p_{ij})\Pi_i$. Dividing the balance identity by t_n and letting $n \to \infty$ gives the balance constraint of (3.1) at (p, Π, y). Dividing the fleet inequalities by t_n, summing over all pairs, and letting $n \to \infty$ gives $\sum_{\mathcal{A}} m_{ij}(t_0)\Lambda_{ij}\bar{F}_{ij}(p_{ij})\Pi_i + \sum_{\mathcal{V}} m_{ij}(t_0) y_{ij} \leq K$, and letting $t_0 \to \infty$, so that $m_{ij}(t_0) \uparrow 1/\mu_{ij}$ by monotone convergence, gives the fleet constraint. Fares and repositioning costs in $[0,t]$ are $\sum_{\mathcal{A}} p_{ij} L_{ij}(t)$ and $\sum_{\mathcal{V}} c_{ij} Y_{ij}(t)$, so $\mathcal{R}_K(p, \theta \mid \mathbf{s}) = \sum_{\mathcal{A}} \Lambda_{ij}\bar{F}_{ij}(p_{ij}) p_{ij} \Pi_i - \sum_{\mathcal{V}} c_{ij} y_{ij}$, the objective of (3.1) at the feasible point (p, Π, y). Hence $\mathcal{R}_K(p, \theta \mid \mathbf{s}) \leq J_K^*(\alpha)$, and taking the maximum over $\mathbf{s}$ and the supremum over (p, θ) gives the claim.

First inequality. $p^H \in \mathcal{P}_\alpha$ by Proposition 3.2, and $\theta^H \in \Theta$ was shown after (3.5). If $\mathcal{N}^+ = \emptyset$ then $\hat{z} = 0$ and $\hat{y} = 0$, so $\hat{J} = 0$, $\theta^H = 0$ and $\mathcal{R}_K(p^H, \theta^H) \geq 0$. Otherwise, by Lemma C.2, the closed classes C of $\mathcal{N}^+$ do not interact. Start from a state with K_C vehicles in each class C,

$\sum_C K_C = K$. By Lemma C.2(b), v restricted to each class solves the traffic equations of that class, so Lemma C.1(ii), with the insensitivity stated after Lemma C.1, applies to each class separately, and the long-run averages are expectations under the product over the classes C of the product-form distributions of K_C vehicles in C.

Now consider π_K on all of $\mathcal{N}^+$, with $w_i = \hat{\Pi}_i$ and $\varrho = v/\mu$ (Lemma C.2(c)). Its weight factors over the classes, so conditionally on the class counts (K_C) it is exactly the long-run law just described. Fares are collected at rate $\sum_{\mathcal{A}} p^H_{ij}\lambda_{ij}\mathbf{1}\{X_i \geq 1\}$ and repositioning moves from i to j start at rate θ^H_{ij} times the rate at which trips into i end. Averaging over (K_C) under π_K and using Lemma C.1, the expected fare rate is $\omega_K \sum_{\mathcal{A}} p^H_{ij}\lambda_{ij}\hat{\Pi}_i$ and the expected repositioning rate on (i,j) is $\theta^H_{ij}\omega_K\eta_i = \omega_K\hat{y}_{ij}$. Since $p^H_{ij}\lambda_{ij}\hat{\Pi}_i = p^H_{ij}\Lambda_{ij}\hat{z}_{ij} = \Lambda_{ij}\tilde{r}_{ij}(\hat{z}_{ij}, \hat{\Pi}_i)$ (both sides vanish when $\hat{z}_{ij} = 0$), this gives

$$\sum_{(K_C)} \pi_K\big((K_C)\big)\,\big(\text{long-run revenue from counts } (K_C)\big) = \omega_K \hat{J}.$$

The best initial state does at least as well as this average, so $\mathcal{R}_K(p^H, \theta^H) \geq \omega_K\hat{J}$, with equality for every initial state in $\mathcal{N}^+$ when $\mathcal{N}^+$ is a single class, since then there is only one count. The middle inequality holds because $p^H \in \mathcal{P}_\alpha$ and $\theta^H \in \Theta$. □

Lemma C.5 (Acyclic repositioning)**.** *Let $\bar{p} \in \bar{\mathcal{P}}_\alpha$. Cancelling the directed cycles of $\hat{y}$ turns any optimal solution of* (3.3) *into an optimal solution with no directed cycle in the support of $\hat{y}$, and every feasible point of* (3.3) *with no directed cycle in the support of $\hat{y}$ has $\hat{\nu} \leq \bar{\nu}$.*

Proof. Subtracting the smallest flow around a directed cycle of $\hat{y}$ keeps feasibility and does not lower the objective, since $c \geq 0$, and finitely many such steps remove all directed cycles. For the bound, the fleet constraint gives $\hat{\nu} \leq K$. Since $\hat{y}$ has no directed cycle, it decomposes into flows $f_{\mathfrak{p}} \geq 0$ on simple directed paths $\mathfrak{p}$, each running from a location where $\hat{y}$ has net outflow to one where it has net inflow (Ahuja et al., 1993, Thm. 3.5). By the balance constraint the net outflow of $\hat{y}$ at i equals $\sum_j \Lambda_{ji}\hat{z}_{ji} - \sum_j \Lambda_{ij}\hat{z}_{ij}$, so $\sum_{\mathfrak{p}} f_{\mathfrak{p}} \leq \sum_i \sum_j \Lambda_{ji}\hat{z}_{ji} \leq \Lambda_{\text{tot}}$, using $\hat{z}_{ji} \leq \hat{\Pi}_j \leq 1$. Each path has at most $N-1$ pairs, so $\sum_{\mathcal{V}} \hat{y}_{ij}/\mu_{ij} = \sum_{\mathfrak{p}} f_{\mathfrak{p}} \sum_{(i,j)\in\mathfrak{p}} \mu_{ij}^{-1} \leq \bar{\tau}\Lambda_{\text{tot}}$, and $\sum_{\mathcal{A}} \Lambda_{ij}\hat{z}_{ij}/\mu_{ij} \leq \sum_{\mathcal{A}} \Lambda_{ij}/\mu_{ij}$. Hence $\hat{\nu} \leq \bar{\nu}$. □

Proof of Theorem 3.3. An optimal solution of (3.3) at $\bar{p}^*$ exists by Proposition 3.2, and Lemma C.5 gives one with no directed cycle in the support of $\hat{y}$. For every such solution, $\hat{J} = V_\alpha(\bar{p}^*) = J^*_K(\alpha) \geq 0$ by Proposition 3.2, since $(0,0,0)$ is feasible, and $\hat{\nu} \leq \bar{\nu}$ by Lemma C.5. Proposition C.4 gives the last three inequalities. If $\mathcal{N}^+ = \emptyset$ then $J^*_K(\alpha) = \hat{J} = 0$, so every term in the chain is 0, with $\omega_K\hat{J}$ read as in Proposition C.4. Otherwise Lemma C.3 gives the first inequality, since its factor decreases in $\hat{\nu}$. □

C.3 Proofs for Section 3.3

This subsection states the search of Algorithm 1 precisely, bounds the number of convex programs it solves, shows that it stops for every $\varepsilon > 0$, proves Corollary 3.4, and gives the cost of the remaining steps. Throughout, Assumption 3.1(a) holds, and for $W_\alpha \leq a \leq b \leq p_{\max}$ we write $U_\alpha(a,b)$ for the value of (3.3) with the window constraint replaced by $\bar{F}_{ij}(d_{ij}b)\,\Pi_i \leq z_{ij} \leq \bar{F}_{ij}\big(d_{ij}(a - W_\alpha)\big)\,\Pi_i$ for all $(i,j) \in \mathcal{A}$, so that $U_\alpha(a,a) = V_\alpha(a)$. Since $0 \leq a - W_\alpha$ and $b \leq p_{\max}$, every ratio z_{ij}/Π_i allowed by this window lies in $\mathcal{Q}_{ij}$, so, as in Step 4 of the proof of

Proposition 3.2, this program maximizes a continuous function over a nonempty compact set and its maximum is attained.

Short cells. For $(i,j) \in \mathcal{A}$ and $h, t \geq 0$ let

$$\mathcal{F}_{ij}(h) := \sup\bigl\{|\bar{F}_{ij}(u) - \bar{F}_{ij}(v)| : \ u, v \in [0, d_{ij}p_{\max}],\ |u - v| \leq d_{ij}h\bigr\},$$
$$\mathcal{H}_{ij}(t) := \sup\bigl\{|r_{ij}(q) - r_{ij}(q')| : \ q, q' \in \mathcal{Q}_{ij},\ |q - q'| \leq t\bigr\}$$

be the moduli of continuity of $\bar{F}_{ij}$ on $[0, d_{ij}p_{\max}]$ and of r_{ij} on $\mathcal{Q}_{ij}$. Both are nondecreasing, and both tend to 0 as $h \downarrow 0$ and $t \downarrow 0$, because $\bar{F}_{ij}$ is continuous on the compact interval $[0, d_{ij}p_{\max}]$ by Assumption 3.1(a) and r_{ij} is continuous on the compact interval $\mathcal{Q}_{ij}$ (Appendix C.1). Let

$$\bar{J} := \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij} \max_{q\in\mathcal{Q}_{ij}} |r_{ij}(q)| + K \sum_{(i,j)\in\mathcal{V}} c_{ij}\mu_{ij}, \qquad \tau_{ij} := \frac{1}{\mu_{ij}} + \frac{\mathbf{1}\{i \neq j\}}{\mu_{ji}},$$

$$\mathcal{E}(h) := \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\Bigl[\mathcal{H}_{ij}\bigl(\mathcal{F}_{ij}(h)\bigr) + \Bigl(c_{ji}\mathbf{1}\{i \neq j\} + \frac{\bar{J}}{K}\,\tau_{ij}\Bigr)\mathcal{F}_{ij}(h)\Bigr].$$

Then $\mathcal{E}$ is nondecreasing and $\mathcal{E}(h) \to 0$ as $h \downarrow 0$, and $\mathcal{E}$ depends only on the data.

Lemma C.6 (Short cells)**.** *For all $W_\alpha \leq a \leq b \leq p_{\max}$,*

$$V_\alpha(a) \ \leq \ U_\alpha(a,b) \ \leq \ V_\alpha(a) + \mathcal{E}(b-a).$$

Proof. The first inequality is (3.7). For the second, write $f(z, \Pi, y) := \sum_{\mathcal{A}} \Lambda_{ij}\tilde{r}_{ij}(z_{ij}, \Pi_i) - \sum_{\mathcal{V}} c_{ij}y_{ij}$ for the objective of (3.3), and note two facts. First, f is positively homogeneous of degree one, since $\tilde{r}_{ij}(sz, s\Pi) = s\,\tilde{r}_{ij}(z, \Pi)$ for $s \geq 0$. Second, $|f(z, \Pi, y)| \leq \bar{J}$ at every point feasible for some $U_\alpha(a,b)$, since $|\tilde{r}_{ij}(z_{ij}, \Pi_i)| = \Pi_i |r_{ij}(z_{ij}/\Pi_i)| \leq \max_{\mathcal{Q}_{ij}} |r_{ij}|$ when $\Pi_i > 0$, $\tilde{r}_{ij}(0,0) = 0$, and the fleet constraint gives $y_{ij} \leq K\mu_{ij}$.

Let (z, Π, y) attain $U_\alpha(a,b)$ and put $h := b - a$. For each arc $(i,j) \in \mathcal{A}$ let $\delta_{ij} := \max\{0,\ \bar{F}_{ij}(d_{ij}a)\Pi_i - z_{ij}\}$. Since $z_{ij} \geq \bar{F}_{ij}(d_{ij}b)\Pi_i$ and $\Pi_i \leq 1$,

$$\delta_{ij} \leq \bigl(\bar{F}_{ij}(d_{ij}a) - \bar{F}_{ij}(d_{ij}b)\bigr)\Pi_i \leq \mathcal{F}_{ij}(h)\,\Pi_i \leq \mathcal{F}_{ij}(h).$$

Raise each z_{ij} to $z'_{ij} := z_{ij} + \delta_{ij}$, keep Π, and compensate with empty moves, $y'_{ji} := y_{ji} + \Lambda_{ij}\delta_{ij}$ for every $(i,j) \in \mathcal{A}$ with $i \neq j$, leaving the other entries of y unchanged. The point (z', Π, y') satisfies the window constraint of $V_\alpha(a)$. Indeed, $z'_{ij} \geq \bar{F}_{ij}(d_{ij}a)\Pi_i$ by construction, and $z'_{ij} \leq \bar{F}_{ij}(d_{ij}(a - W_\alpha))\Pi_i$, because either $\delta_{ij} = 0$ and $z'_{ij} = z_{ij}$, or $z'_{ij} = \bar{F}_{ij}(d_{ij}a)\Pi_i \leq \bar{F}_{ij}(d_{ij}(a - W_\alpha))\Pi_i$ since $\bar{F}_{ij}$ is nonincreasing. It also satisfies the balance constraints. Raising z_{ij} with $i \neq j$ adds $\Lambda_{ij}\delta_{ij}$ to the outflow of i and to the inflow of j, and the empty move $j \to i$ adds the same amount to the outflow of j and to the inflow of i, while raising z_{ii} adds the same amount to both sides at i. Its fleet usage exceeds that of (z, Π, y), which is at most K, by

$$\Delta_F := \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\delta_{ij}\,\tau_{ij} \leq \sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\tau_{ij}\mathcal{F}_{ij}(h).$$

Let $s := K/(K + \Delta_F) \in (0,1]$. Scaling by s preserves the balance and window constraints, which are homogeneous in (z, Π, y) and in (z, Π), keeps $\Pi \in [0,1]^N$ and $y \geq 0$, and brings the

fleet usage to at most $s(K+\Delta_F) = K$. Hence $s\,(z', \Pi, y')$ is feasible for (3.3) at $\bar{p} = a$, and $V_\alpha(a) \geq f\big(s(z', \Pi, y')\big) = s\, f(z', \Pi, y')$.

It remains to compare $f(z', \Pi, y')$ with $f(z, \Pi, y) = U_\alpha(a, b)$. If $\Pi_i > 0$, both z_{ij}/Π_i and z'_{ij}/Π_i lie in $\mathcal{Q}_{ij}$ and differ by $\delta_{ij}/\Pi_i \leq \mathcal{F}_{ij}(h)$, so $|\tilde{r}_{ij}(z'_{ij}, \Pi_i) - \tilde{r}_{ij}(z_{ij}, \Pi_i)| = \Pi_i |r_{ij}(z'_{ij}/\Pi_i) - r_{ij}(z_{ij}/\Pi_i)| \leq \mathcal{H}_{ij}(\mathcal{F}_{ij}(h))$. If $\Pi_i = 0$, then $z_{ij} = z'_{ij} = 0$. The added empty moves cost at most $\sum_{i \neq j} c_{ji} \Lambda_{ij} \mathcal{F}_{ij}(h)$. Hence

$$f(z', \Pi, y') \;\geq\; U_\alpha(a, b) - E_1, \qquad E_1 := \sum_{(i,j) \in \mathcal{A}} \Lambda_{ij} \big[\mathcal{H}_{ij}(\mathcal{F}_{ij}(h)) + c_{ji} \mathbf{1}\{i \neq j\} \mathcal{F}_{ij}(h) \big].$$

Finally, $s\, f(z', \Pi, y') - U_\alpha(a, b) = s\big(f(z', \Pi, y') - U_\alpha(a, b)\big) - (1 - s)\, U_\alpha(a, b) \geq -E_1 - (1 - s)\bar{J}$, and $1 - s = \Delta_F/(K + \Delta_F) \leq \Delta_F/K$. So $V_\alpha(a) \geq U_\alpha(a, b) - E_1 - (\bar{J}/K) \sum_{\mathcal{A}} \Lambda_{ij} \tau_{ij} \mathcal{F}_{ij}(h) = U_\alpha(a, b) - \mathcal{E}(h)$. □

The search. Algorithm 1 runs in rounds $r = 0, 1, 2, \ldots$. At the start of round r it holds a finite set $\Gamma_r \subseteq \bar{\mathcal{P}}_\alpha$ of grid points with their values $V_\alpha(\bar{p})$, a finite set $\mathcal{C}_r$ of active cells, i.e., intervals $[a, b]$ with $a < b$ and $a, b \in \Gamma_r$, with their values $U_\alpha(a, b)$, a point $\hat{p}_r \in \arg\max_{\bar{p} \in \Gamma_r} V_\alpha(\bar{p})$, and

$$\mathrm{UB}_r := \max\Big\{ V_\alpha(\hat{p}_r),\; \max_{[a,b] \in \mathcal{C}_r} U_\alpha(a, b) \Big\},$$

with the inner maximum read as $-\infty$ when $\mathcal{C}_r = \emptyset$. Round 0 starts from $\Gamma_0 = \{\bar{p}_0, \ldots, \bar{p}_n\}$ and $\mathcal{C}_0 = \{[\bar{p}_{k-1}, \bar{p}_k] : 1 \leq k \leq n\}$. In round r, if $\mathrm{UB}_r - V_\alpha(\hat{p}_r) \leq \varepsilon$, or if the round would exceed the computing budget, the algorithm stops and returns $\hat{p} := \hat{p}_r$ and $\mathrm{UB} := \mathrm{UB}_r$. Otherwise it splits each cell $[a, b] \in \mathcal{C}_r$ with $U_\alpha(a, b) > V_\alpha(\hat{p}_r)$ at its midpoint $m = (a + b)/2$, discards the other cells of $\mathcal{C}_r$, evaluates $V_\alpha(m)$, $U_\alpha(a, m)$, and $U_\alpha(m, b)$ for each split cell, and sets $\Gamma_{r+1} := \Gamma_r \cup \{\text{midpoints}\}$ and $\mathcal{C}_{r+1} := \{\text{halves of the split cells}\}$. Write $\mathcal{S}_r \subseteq \mathcal{C}_r$ for the set of cells split in round r, and $h := \max_{1 \leq k \leq n} (\bar{p}_k - \bar{p}_{k-1})$ for the mesh of the initial grid.

Proposition C.7 (The search)**.** *Let $\varepsilon \geq 0$.*

(a) In every round r, $V_\alpha(\hat{p}_r) \leq J^*_K(\alpha) \leq \mathrm{UB}_r$. Moreover $V_\alpha(\hat{p}_r)$ is nondecreasing and UB_r nonincreasing in r.

(b) Rounds $0, \ldots, r-1$ together with the initial grid solve $2n+1+3\sum_{q<r} |\mathcal{S}_q|$ convex programs, and $|\mathcal{S}_q| \leq |\mathcal{C}_q| \leq n 2^q$.

(c) In every round r, every cell in $\mathcal{C}_r$ has length at most $h 2^{-r}$, and $\mathrm{UB}_r - V_\alpha(\hat{p}_r) \leq \mathcal{E}(h 2^{-r})$.

(d) If $\varepsilon > 0$ and the budget is unlimited, the algorithm stops at the latest in round R, the smallest integer with $\mathcal{E}(h 2^{-R}) \leq \varepsilon$, after at most $2n + 1 + 3n(2^R - 1)$ convex programs, and then $\mathrm{UB} - V_\alpha(\hat{p}) \leq \varepsilon$.

Proof. (a) Since $\Gamma_r \subseteq \Gamma_{r+1} \subseteq \bar{\mathcal{P}}_\alpha$, $V_\alpha(\hat{p}_r)$ is nondecreasing and at most $J^*_K(\alpha) = \max_{\bar{\mathcal{P}}_\alpha} V_\alpha$ (Proposition 3.2). For the upper bound we show by induction on r that every $\bar{p} \in \bar{\mathcal{P}}_\alpha$ either lies in a cell of $\mathcal{C}_r$ or satisfies $V_\alpha(\bar{p}) \leq V_\alpha(\hat{p}_r)$. For $r = 0$ the cells of $\mathcal{C}_0$ cover $\bar{\mathcal{P}}_\alpha$. Suppose the claim holds for r, and let $\bar{p}$ lie in $[a, b] \in \mathcal{C}_r$. If the cell is split, $\bar{p}$ lies in one of its halves, which belong to $\mathcal{C}_{r+1}$. If it is discarded, then (3.7) gives $V_\alpha(\bar{p}) \leq U_\alpha(a, b) \leq V_\alpha(\hat{p}_r) \leq V_\alpha(\hat{p}_{r+1})$. If instead $V_\alpha(\bar{p}) \leq V_\alpha(\hat{p}_r)$, then $V_\alpha(\bar{p}) \leq V_\alpha(\hat{p}_{r+1})$. This proves the claim. In round r, a

point $\bar{p}$ in a cell $[a,b] \in \mathcal{C}_r$ has $V_\alpha(\bar{p}) \le U_\alpha(a,b) \le \mathrm{UB}_r$ by (3.7), and any other point has $V_\alpha(\bar{p}) \le V_\alpha(\hat{p}_r) \le \mathrm{UB}_r$. Applying this to a maximizer of V_α gives $J_K^*(\alpha) \le \mathrm{UB}_r$. Finally, every cell of $\mathcal{C}_{r+1}$ is a subinterval of a cell of $\mathcal{C}_r$, so its value U_α is at most that of the larger cell, as noted after (3.7), and $V_\alpha(\hat{p}_{r+1}) \le J_K^*(\alpha) \le \mathrm{UB}_r$. Hence $\mathrm{UB}_{r+1} \le \mathrm{UB}_r$.

(b) The initial grid needs V_α at $n+1$ points and U_α on n cells. Splitting a cell needs V_α at its midpoint and U_α on its two halves, i.e., three programs. Each cell of $\mathcal{C}_q$ has at most two halves in $\mathcal{C}_{q+1}$, so $|\mathcal{C}_q| \le n2^q$, and $\mathcal{S}_q \subseteq \mathcal{C}_q$.

(c) The cells of $\mathcal{C}_0$ have length at most h, and halving preserves the bound $h2^{-r}$ from one round to the next. Let $[a,b] \in \mathcal{C}_r$. Since $a \in \Gamma_r$, Lemma C.6 and the monotonicity of $\mathcal{E}$ give $U_\alpha(a,b) \le V_\alpha(a) + \mathcal{E}(b-a) \le V_\alpha(\hat{p}_r) + \mathcal{E}(h2^{-r})$. Taking the maximum over $\mathcal{C}_r$, together with $V_\alpha(\hat{p}_r)$, gives $\mathrm{UB}_r \le V_\alpha(\hat{p}_r) + \mathcal{E}(h2^{-r})$.

(d) Since $\mathcal{E}(h2^{-r}) \to 0$ as $r \to \infty$, the integer R exists, and by (c) the stopping test $\mathrm{UB}_r - V_\alpha(\hat{p}_r) \le \varepsilon$ holds at the latest in round R. Rounds $0, \dots, R-1$ split at most $\sum_{q<R} n2^q = n(2^R - 1)$ cells, so (b) bounds the number of programs. □

Thus ε is an input, the tolerance, while the optimality gap $\mathrm{UB} - V_\alpha(\hat{p})$ is an output. The gap is at most ε when the algorithm stops by the test, and it may exceed ε when the budget stops it, in which case the bracket $V_\alpha(\hat{p}) \le J_K^*(\alpha) \le \mathrm{UB}$ remains valid by (a). Part (c) also shows that the gap tends to 0 as the budget grows, for every $\varepsilon \ge 0$.

Proof of Corollary 3.4. By Proposition C.4, $p^H \in \mathcal{P}_\alpha$ and $\theta^H \in \Theta$. Algorithm 1 cancels the directed cycles of $\hat{y}$, so by Lemma C.5 it uses an optimal solution of (3.3) at $\hat{p}$ with no directed cycle in the support of $\hat{y}$, for which $\hat{J} = V_\alpha(\hat{p}) \ge 0$ by Proposition 3.2, since $(0,0,0)$ is feasible, and $\hat{\nu} \le \bar{\nu}$. Proposition C.4 gives the second, third, and fourth inequalities, and Proposition C.7(a) gives the last.

For the first inequality, $\mathrm{UB} - \varepsilon \le V_\alpha(\hat{p})$ when the algorithm stops with $\mathrm{UB} - V_\alpha(\hat{p}) \le \varepsilon$, and $\mathrm{UB} - \varepsilon = V_\alpha(\hat{p})$ when ε is replaced by $\mathrm{UB} - V_\alpha(\hat{p})$. The factor $\max\{K/(N+K-1+\bar{\nu}),\, 1-\sqrt{N/K}\}$ is positive, since $K \ge 1$, so in both cases the left end is at most the factor times $V_\alpha(\hat{p})$. If $\mathcal{N}^+ = \emptyset$, then $V_\alpha(\hat{p}) = \hat{J} = 0$ by the proof of Proposition C.4, so the left end is at most 0, which is $\omega_K V_\alpha(\hat{p})$ read as in that proposition. Otherwise Lemma C.3 gives $\omega_K \ge \max\{K/(N+K-1+\hat{\nu}),\, 1-\sqrt{N/K}\}$, which is at least the factor because the bound of Lemma C.3 decreases in $\hat{\nu}$ and $\hat{\nu} \le \bar{\nu}$, and multiplying by $V_\alpha(\hat{p}) \ge 0$ gives the first inequality. □

The cost of each step. Under all of Assumption 3.1, each $V_\alpha(\bar{p})$ and each $U_\alpha(a,b)$ is a convex program (Proposition 3.2) with $|\mathcal{A}| + N + N(N-1)$ variables, N balance equations, one fleet inequality, $2|\mathcal{A}|$ window inequalities, and the bounds $0 \le \Pi \le 1$ and $y \ge 0$. For linear demand it is a second-order cone program.

Lemma C.8 (Conic form under linear demand)**.** *Let $\bar{F}_{ij}(p) = (1 - p/b_{ij})^+$ for every $(i,j) \in \mathcal{A}$. Then $V_\alpha(\bar{p})$ equals the largest value of*

$$\sum_{(i,j)\in\mathcal{A}} \Lambda_{ij}\, b_{ij}\, (z_{ij} - t_{ij}) - \sum_{(i,j)\in\mathcal{V}} c_{ij}\, y_{ij}$$

over the points (z,t,Π,y) that satisfy the constraints of (3.3) and $z_{ij}^2 \le t_{ij}\Pi_i$, $t_{ij} \ge 0$ for all $(i,j) \in \mathcal{A}$, and likewise for $U_\alpha(a,b)$. Since $z^2 \le t\Pi$ with $t, \Pi \ge 0$ is the rotated second-order

cone $\|(2z,\, t-\Pi)\|_2 \le t+\Pi$*, this is a second-order cone program with* $2|\mathcal{A}|+N+N(N-1)$ *variables and* $|\mathcal{A}|$ *three-dimensional cones.*

Proof. On $\mathcal{Q}_{ij}$, $\bar{F}_{ij}^{-1}(q) = b_{ij}(1-q)$, so $r_{ij}(q) = b_{ij}q(1-q)$ and $\tilde{r}_{ij}(z,\Pi) = b_{ij}(z - z^2/\Pi)$ for $\Pi > 0$. Fix a point (z,Π,y) feasible for (3.3) and an arc (i,j). If $\Pi_i > 0$, the constraints on t_{ij} read $t_{ij} \ge z_{ij}^2/\Pi_i$, and since $b_{ij} > 0$ the largest value of $b_{ij}(z_{ij} - t_{ij})$ is $b_{ij}(z_{ij} - z_{ij}^2/\Pi_i) = \tilde{r}_{ij}(z_{ij},\Pi_i)$. If $\Pi_i = 0$, the window constraint forces $z_{ij} = 0$, every $t_{ij} \ge 0$ is allowed, and the largest value is $0 = \tilde{r}_{ij}(0,0)$. So maximizing over t first recovers the objective of (3.3). The same argument applies to $U_\alpha(a,b)$, whose window also lies in $[0, p_{\max}]$. □

After the search, Algorithm 1 needs one optimal solution $(\hat{z}, \hat{\Pi}, \hat{y})$ at $\hat{p}$, which the evaluation of $V_\alpha(\hat{p})$ already provides. Cancelling the directed cycles of $\hat{y}$ (Lemma C.5) takes at most $N(N-1)$ steps, since each step sets the flow on at least one pair in $\mathcal{V}$ to zero and makes no zero flow positive, and each directed cycle is found by depth-first search in $O(N^2)$ operations. The prices (3.4) and the probabilities (3.5) are closed forms. The exact revenue $\omega_K \hat{J}$ needs $\hat{J}$, which is the objective value of the solution, and $\omega_K = Z_{K-1}/Z_K$, which the recursion after (C.1) gives from $\hat{\Pi}$ and $\hat{\nu}$ in $O(NK)$ operations. No Markov chain is solved, so beyond the convex programs of the search the computation takes $O(N^4 + NK)$ operations.

C.4 Experimental details for Section 3.4

The instance has $N = 3$ locations and all nine arcs, including the three trips that return to their origin, with

$$\Lambda = \begin{pmatrix} 0.5 & 1.2 & 0.5 \\ 0.8 & 0.5 & 1.0 \\ 0.4 & 1.5 & 0.5 \end{pmatrix}, \qquad \mu = \begin{pmatrix} 3.0 & 2.0 & 1.5 \\ 2.0 & 3.0 & 2.5 \\ 1.5 & 2.5 & 3.0 \end{pmatrix}, \qquad c = \begin{pmatrix} \cdot & 0.15 & 0.25 \\ 0.15 & \cdot & 0.10 \\ 0.25 & 0.10 & \cdot \end{pmatrix}.$$

Demand is linear, $\bar{F}_{ij}(p) = (1 - p/b_{ij})^+$, with a different b_{ij} on every arc,

$$b = \begin{pmatrix} 1.0 & 2.0 & 1.7 \\ 1.3 & 0.8 & 1.1 \\ 0.9 & 1.5 & 0.7 \end{pmatrix}.$$

We take $d_{ij} = 1$ and $p_{\max} = 2$. Every revenue reported in Figure 5 is computed exactly from the product form of Lemma C.1, with the vehicles placed in the closed class, or split among the closed classes, that earns most, and this evaluator was checked against a brute-force continuous-time Markov chain with exponential trip times on random policies with $K \in \{1, 2\}$.

Upper bound and heuristic. Algorithm 1 is run with an initial grid of 21 equally spaced points of $[W_\alpha, p_{\max}]$, the tolerance $\varepsilon = 10^{-4}$, and a budget of 2000 convex programs, each solved with the conic solver CLARABEL through `cvxpy`, and a round of splitting is carried out only if it fits in the budget. In every instance of Figure 5 the algorithm stops with $\mathrm{UB} - V_\alpha(\hat{p}) < 10^{-4}$ within the budget, after between 41 and 941 convex programs. The heuristic (p^H, θ^H) is the policy that the algorithm returns, and its exact revenue agrees with $\omega_K \hat{J}$, where ω_K is computed from (C.1) with $w_i = \hat{\Pi}_i$ and $\nu = \hat{\nu}$.

Benchmark. A static α-fair policy is encoded as a vector in $[0,1]^{1+|\mathcal{A}|+|\mathcal{V}|}$, namely the position of $\bar{p}$ in $[W_\alpha, p_{\max}]$, the position of each price in the window $[\bar{p} - W_\alpha, \bar{p}]$, and repositioning weights normalized so that $\sum_j \theta_{ij} \le 1$. The best policy is sought by differential evolution

(SciPy) with two random seeds, a population of 12 times the dimension, and at most 200 generations, followed by L-BFGS-B polishing. The heuristic (p^H, θ^H) is inserted into the initial population, so the benchmark is never below the heuristic, and in panel (b) the best unconstrained policy is inserted as well. In panel (b), $\Delta = 0.79$ is the gap between the largest and the smallest price of the best policy found without a fairness constraint by the same search (three seeds, population 15 times the dimension, 300 generations).

Certified check at $\alpha = 1$. For $K \leq 8$ we also solved, with the global solver SCIP (Achterberg, 2009), an exact mixed-integer formulation of the problem over the policies whose routing forms a single closed class. In it, the product form is imposed through mean value analysis, a binary variable per arc records whether the arc is open, and the prices on open arcs lie in the window. With a time limit of 600 seconds and a relative gap of 10^{-4}, SCIP proved optimality for $K \leq 2$ and stopped at the gap limit for $3 \leq K \leq 8$, each within three minutes. In every case the value found by the global search is within 0.01% of the upper bound of SCIP, i.e., it is optimal among those policies up to the tolerance of the solver.

References

Tobias Achterberg. SCIP: Solving constraint integer programs. *Mathematical Programming Computation*, 1(1):1–41, 2009.

Ravindra K. Ahuja, Thomas L. Magnanti, and James B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Englewood Cliffs, NJ, 1993.

Siddhartha Banerjee, Daniel Freund, and Thodoris Lykouris. Pricing and optimization in shared vehicle systems: An approximation framework. *Operations Research*, 70(3):1783–1805, 2022.

Andrew D. Barbour. Networks of queues and the method of stages. *Advances in Applied Probability*, 8(3):584–591, 1976.

Forest Baskett, K. Mani Chandy, Richard R. Muntz, and Fernando G. Palacios. Open, closed, and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22 (2):248–260, 1975.

Stephen Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.

Adam N. Elmachtoub and Hyemi Kim. Fair fares for vehicle sharing systems. *Operations Research*, 2026. doi: 10.1287/opre.2024.1043.

William J. Gordon and Gordon F. Newell. Closed queuing systems with exponential servers. *Operations Research*, 15(2):254–265, 1967.

Stuart G. Hoggar. Chromatic polynomials and logarithmic concavity. *Journal of Combinatorial Theory, Series B*, 16(3):248–254, 1974.

Frank P. Kelly. *Reversibility and Stochastic Networks*. Wiley, Chichester, UK, 1979.

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

Martin Reiser and Stephen S. Lavenberg. Mean-value analysis of closed multichain queuing networks. *Journal of the ACM*, 27(2):313–322, 1980.