Addendum: on the uniqueness assumption

Why the assumption is needed. Without uniqueness of the fluid optimum, Robinson's stability bound (Hoffman, 1952; Robinson, 1977) places some offline solution near the fluid solution set, but not necessarily near the particular $\bar{x}$ computed by the algorithm. The entry $\bar{x}_{jl}$ that is large then need not have a counterpart X_{jl} that is large, and the implication underlying the information-loss bound can fail. Under primal uniqueness[3], the bound specializes to the prescribed $\bar{x}$ and yields the coordinatewise estimate $\bar{x}_{jl} \geq t^{\frac{1}{2}+\epsilon} \Rightarrow X_{jl} \geq 1/\eta$ used in Lemma 7. The assumption is therefore necessary for this proof rather than cosmetic.

The following simple example, by Yun-Tung Kuo and Yehua Wei, shows that without some condition the regret bound of Theorem 3 may fail: take $d = n = 1$, prices $f_1 = 2,\ f_2 = 1$ with $q_1 = \frac{1}{2},\ q_2 = 1$, and non-binding initial budget $B = T$. Because $f_1 q_1 = f_2 q_2 = 1$, both static-price policies are DP optimal with expected value T. The fluid $P[B, q, T]$ has, as well, value T and multiple optima. OFFLINE, in contrast, has the value $\mathbb{E}\big[T \cdot \max\{f_1 Q_1(T),\ f_2\}\big] = T + \Omega(\sqrt{T})$.

On reachability. We require uniqueness for every feasible $\bar{\mathrm{b}}$. The proof in fact uses uniqueness only at the residual budgets $\bar{\mathrm{b}}^t$ that Algorithm 3* actually visits. Constant regret is maintained, for instance, if with probability $1 - \mathcal{O}(1/T)$ the algorithm visits during its run only states with a unique fluid solution.

The uniqueness hypothesis is non-vacuous: the single-resource case. We show that the uniqueness assumption holds for a natural family of instances. Take a single resource, $d = 1$, with $A_j = 1$ for all types $j \in [n]$, so that the fluid LP is

$$\max_{x \geq 0} \sum_{j,\ell} f_{j\ell} q_{j\ell}\, x_{j\ell} \quad \text{s.t.} \quad \sum_{j,\ell} q_{j\ell}\, x_{j\ell} \leq b, \qquad \sum_{\ell \in \mathcal{L}_j} x_{j\ell} = p_j \ \ \forall j,$$

with $p \in \Delta_n := \{p \in \mathbb{R}^n_{\geq 0} : \sum_j p_j = 1\}$. For a type j and two of its prices $\ell \neq \ell'$ with $q_{j\ell} \neq q_{j\ell'}$, let

$$\mu_{j,\ell\ell'} = \frac{f_{j\ell'} q_{j\ell'} - f_{j\ell} q_{j\ell}}{q_{j\ell'} - q_{j\ell}}$$

be the capacity price at which type j is indifferent between ℓ and ℓ'. We impose:

(A1) For every type j, the values $\{f_{j\ell} q_{j\ell}\}_\ell$ are pairwise distinct, and the indifference values $\mu_{j,\ell\ell'}$ are pairwise distinct over price pairs.

[3] Robinson's bound applies to the primal–dual pair: for every offline optimum (X, U) it produces a nearby fluid optimum (x^0, u^0). When the fluid primal optimum is unique, every such (x^0, u^0) has the same primal part $\bar{x}$, so the bound delivers $\|X - \bar{x}\|_2 \leq \frac{1}{\zeta}\epsilon$ even when the fluid dual is not unique.

(A2) No indifference value is shared by two types: $\mu_{j,\ell\ell'} \neq \mu_{j',mm'}$ whenever $j \neq j'$.

Proposition. *Under* (A1)–(A2)*, with $d = 1$ and $A_j = 1$ for all j, the fluid LP has a unique optimal solution for every $p \in \Delta_n$ and every $b > 0$ for which it is feasible.*

PROOF. Fix feasible $b > 0$ and $p \in \Delta_n$ (wlog $p_j > 0$). Suppose x and x' are both optimal. Let $(\mu^\star, \nu)$ be an optimal dual solution. By complementary slackness, for any price ℓ active in x or x':

$$\nu_j = f_{j\ell} q_{j\ell} - \mu^\star q_{j\ell}.$$

If two prices $\ell \neq \ell'$ of type j satisfy this, then $q_{j\ell} \neq q_{j\ell'}$ (by (A1)), so $\mu^\star = \mu_{j,\ell\ell'}$. By (A1)–(A2), at most one type j_0 can have two eligible prices at $\mu^\star$; all other types have a unique eligible price $\ell^\star(j)$, so x and x' agree on $j \neq j_0$. If no such j_0 exists, $x = x'$. Otherwise, x and x' can only differ in how type j_0 splits p_{j_0} between its two eligible prices ℓ, ℓ'. Since $\mu^\star = \mu_{j_0,\ell\ell'} > 0$ (at $\mu = 0$ the tightness condition reads $\nu_j = f_{j\ell} q_{j\ell}$, which by (A1) holds for at most one ℓ per type), the capacity constraint binds. Because x and x' consume the same amount on every type other than j_0, both leave the same residual b_0 for j_0, and both satisfy

$$x_{j_0\ell} + x_{j_0\ell'} = p_{j_0}, \quad q_{j_0\ell} x_{j_0\ell} + q_{j_0\ell'} x_{j_0\ell'} = b_0.$$

The system has a unique solution (determinant $q_{j_0\ell'} - q_{j_0\ell} \neq 0$), so $x = x'$. ■

Uniqueness does not preclude degeneracy. Uniqueness and primal degeneracy are independent. Take $d = 1$ and $n = 2$, and let $\ell^\star(j) = \arg\max_\ell f_{j\ell} q_{j\ell}$ (unique by (A1)). At $b = q_{1\ell^\star(1)} p_1 + q_{2\ell^\star(2)} p_2$, the unique optimum has $x_{j,\ell^\star(j)} = p_j$ for $j = 1, 2$ and all other variables equal to 0. The capacity constraint is binding, so the capacity slack variable is 0. Since there are three equality constraints (two demand constraints and one capacity constraint), every basis contains three basic variables. However, only the two variables $x_{1,\ell^\star(1)}$ and $x_{2,\ell^\star(2)}$ are positive. Therefore at least one basic variable is 0, and the unique optimal basic feasible solution is degenerate.