Dynamic Relocation in Ridesharing via Fixpoint Construction (Supplementary material)

Ian A. Kash[1] **Zhongkai Wen**[1] **Lenore D. Zuck**[1]

[1]University of Illinois Chicago, Chicago, IL, USA

A ADDITIONAL SIMULATION DETAILS

We use the dataset from Braverman et al. [2019], whose raw data is from the Di-Tech Challenge and covers 5-6 pm (evening rush hour) interval in each day between Jan 1-2016 until Jan 21-2016.

The transition matrix V induced from the raw data is in Table 1. The average travel times normalized for 10-min slots are in Table2. The demand rates (normalized for a fleet size of 1) are in Table 3.

For our unit-time setup, we adapt the non-uniform travel time by adding intermediary dummy regions for routes with longer travel time. So the travel times between real regions and dummy regions and between dummy regions are unit time. We accordingly adapt transition matrix V for dummy regions. For the 9-region Didi data, we use 1.5 as the normalized travel time between, which corresponds to 15 mins, thus divide the travel times in Table 1 by 1.5 and round to the nearest integer. This results in an extension of the 9-region problem to 48-regions when the dummy regions are included. Then we construct OA based on this 48 region problem.

B ADDITIONAL SIMULATIONS

We examine the performance of our dynamic relocation policy, referred to as CON here, relative to several baselines. In contrast to Sec. 7, these simulations use synthetic data with unit travel times to more closely match our theoretical results. We simulate a range of parameters, initial conditions, and policies and show that CON converges substantially faster, coming close to a lower bound. We also define two performance metrics, efficiency and availability, and show that our approach matches the baseline performance in terms of efficiency while being more flexibly able to target availability.

B.1 WHAT WE COMPARE

Consider the dynamics $\widehat{M}_{t+1} = \mathbb{1}F(\widehat{M}_t)$. These non-mass-conserving dynamics follow the drivers carrying passengers, discarding any who would relocate. It follows, from the relocation constraint of dynamic relocation, that $\widehat{M}_t \leq M_t$ for any dynamic relocation policy $\{M_t\}_{t=0}$. Thus, the lowest t for which $\widehat{M}_t \leq M^*$, if such a time exists, is a lower bound on the convergence time of any dynamic policy. We refer to this lower bound as LB.

We then compare our CON with LB and three policies. These are essentially the same as those from Sec. 7, with minor variations due to the uniform distances.

STA. This is a static policy that sets $\pi_t = \pi$ for all t. From Cor. 1 it follows that STA guarantees convergence to the fixpoint, yet, as we pointed out, it may do so slowly. Thus STA represents a baseline in the absence of a more sophisticated dynamic policy.

GDY. This is a greedy policy that distributes the relocating mass proportional to the unmet demand in each region with a one-step look ahead. That is, it takes $\pi_t[i] \propto [W\mathbb{1} - \mathbb{1}F(M_t)]_+$, which guarantees as many relocating drivers as possible will have a passenger at time $t+1$ while spreading them among the regions where they can be useful. As GDY does not depend on π, it may not converge to the fixpoint, but it does provide a meaningful baseline for other metrics based purely on the provision of service.

HMR. We adapt the dynamic policy of [Hosseini et al., 2021], which dispatches a single car at a time, to our setting. In particular, their algorithm computes a measure of which region will generate the most long-run service and sends the car there. Since the results of this computation do not change until a region is saturated, we adapt their policy by assigning relocation drivers to this region until (a) it becomes saturated or (b) the mass reaches the fixpoint mass of the region. While their policy is heuristic, with the inclusion of (b) it can

Accepted for the 38th Conference on Uncertainty in Artificial Intelligence (UAI 2022).

Region	1	2	3	4	5	6	7	8	9
1	0.230	0.297	0.372	0.004	0.026	0.029	0.009	0.018	0.015
2	0.044	0.655	0.146	0.005	0.079	0.038	0.018	0.005	0.011
3	0.165	0.291	0.288	0.007	0.054	0.126	0.017	0.025	0.027
4	0.0013	0.010	0.006	0.139	0.031	0.185	0.101	0.117	0.409
5	0.005	0.096	0.026	0.037	0.25	0.333	0.218	0.012	0.027
6	0.004	0.031	0.032	0.088	0.121	0.426	0.148	0.059	0.092
7	0.002	0.023	0.011	0.066	0.142	0.269	0.399	0.020	0.069
8	0.004	0.008	0.023	0.067	0.011	0.095	0.019	0.400	0.374
9	0.001	0.004	0.005	0.095	0.010	0.059	0.030	0.185	0.610

Table 1: Transition matrix V for 9 region Didi Data

Region	1	2	3	4	5	6	7	8	9
1	0.83	1.87	1.07	3.89	3.25	2.79	4.25	2.94	4.37
2	1.78	0.89	1.18	3.24	1.24	1.99	2.89	3.46	4.18
3	1.02	1.31	0.78	2.82	1.45	1.36	3.26	2.17	3.04
4	3.52	3.13	2.76	0.93	1.5	1.26	1.49	1.75	1.6
5	2.86	1.42	1.64	1.55	0.84	1.04	1.45	2.88	2.89
6	2.61	2.17	1.54	1.31	1.15	0.81	1.86	1.78	2.2
7	4.38	3.02	2.79	1.36	1.35	1.65	0.94	3.1	3
8	2.93	3.06	2.26	1.75	2.69	1.62	3.23	0.9	1.48
9	3.58	4.18	2.8	1.49	2.46	2.02	2.72	1.43	1.01

Table 2: Normalized travel times for 9 region Didi Data

often achieve convergence in our simulations.

B.2 SIMULATION DESIGN

For each datapoint in each of our experiments we generate 40 different demand matrices W according the specified distribution and solve the resulting optimal allocation problem (Fig. 2) to determine π and M^* for each W. We then choose 20 different M_0 uniformly at random. We run each dynamic policy for 50 steps starting from each M_0 and report the results averaged across the choices of W and M_0. Thus each point is an average of 800 runs. We compare the performance of the various policies according to the following metrics, all of which have been considered in prior work on dynamic relocation [Braverman et al., 2019, Hosseini et al., 2021].

- Convergence Time. The number of steps of the dynamics until the ratio of current mass in a region to the fixpoint mass in that region is at most $1 + 10^{-6}$ for every region.
- Efficiency. The total number of full rides as a fraction of the total demand. In our stylized model this captures both the total value created for passengers and the revenue of the platform.
- Availability. The average efficiency on a per-region basis. That is, the ratio of full rides in a region to the demand in that region averaged across regions. High availability ensures some degree of fairness between regions

B.3 RESULTS

Fig. 1 shows the performance with all three of our metrics with 40 regions when the demand between each pair of regions is i.i.d. uniform, the supply of drivers is determined as a multiple of the total demand, and the policy is chosen by solving the optimization from Fig. 2. At high levels of supply of drivers relative to demand (above a ratio of about 1.1) CON, HMR, and GDY are all able to control the relocation of enough drivers to satisfy essentially all of the demand. Because so many drivers relocate, LB is essentially 1 and both CON and HMR achieve it. STA converges slowly and has worse efficiency and availability, showing the value of dynamic policies. As previously remarked, GDY does not in general converge to the fixpoint, and typically times out by reaching the 50 step limit. At very low ratios of supply to demand, the dynamics are dominated by the full rides. Thus LB is larger and even GDY converges. At intermediate ratios, CON converges substantially faster than the other approaches. Its performance is close to that of LB, but fails to match it because it only looks ahead a single step. The effects of CON on efficiency and availability relative to HMR

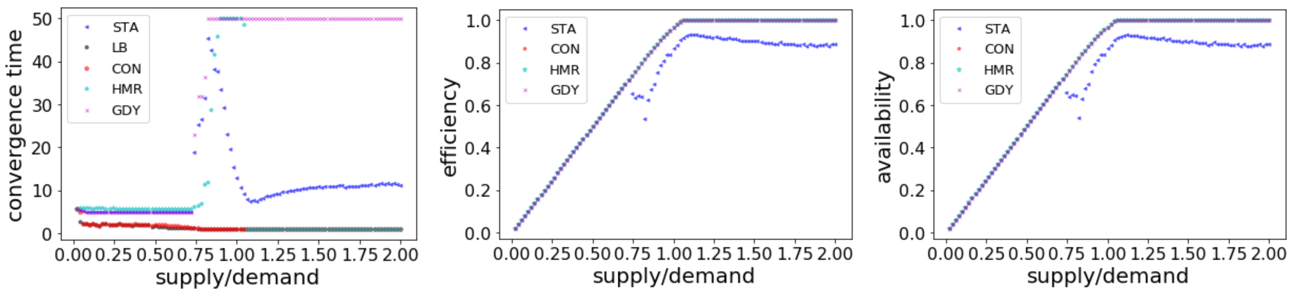

Figure 1: Performance with uniform demand; Efficiency objective

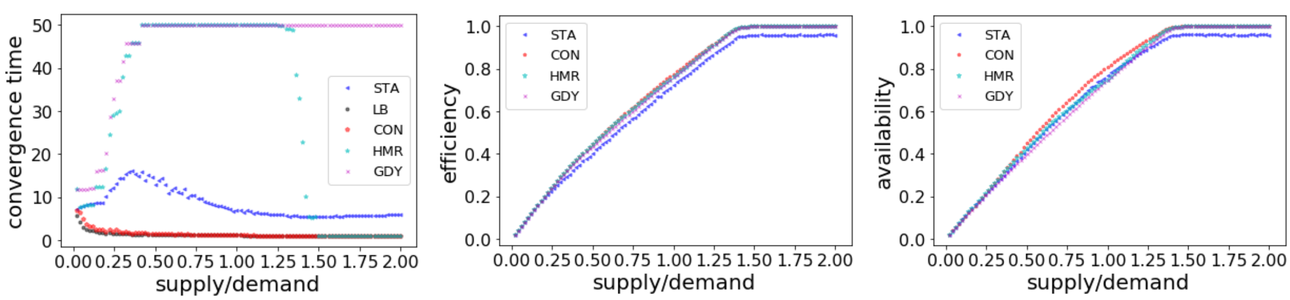

Figure 2: Performance with correlated demand; Half efficiency and half fairness objective

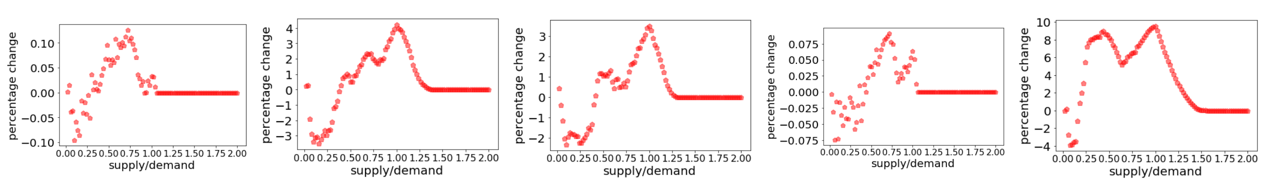

(a) (b) (c) (d) (e)

Figure 3: Relative availability of CON to HMR

Region	1	2	3	4	5	6	7	8	9
Rate	0.0131	0.0624	0.0381	0.0652	0.0870	0.1178	0.0762	0.1438	0.2751

Table 3: Demand rate for 9 region Didi Data

and GDY are minimal.

In Fig. 2, we change the setup of the experiment: (1) We introduce correlation between the demands from various regions by first choosing the total demand in each region i.i.d. uniform and then assigning that demand to each destination proportional to the total demand in the destination raised to a random exponent determined for each source region independently and uniformly from $[0, 10]$. (2) We change the objective of the optimization from maximizing $\mathbf{F}$ to maximizing $0.5\mathbf{F} + 0.5 \min_i F\mathbb{1}[i]$. This shifts the objective from purely efficiency into a hybrid of efficiency and a (somewhat minimal) fairness criterion of maximizing welfare of the least-served region. The overall shape of the results is similar to Fig. 1, with a small improvement in availability relative to HMR and GDY. Because the policy aims at fairness and not just efficiency, CON has a tendency to do so as well. In contrast, HMR and GDY are inherently efficiency-focused, yet we still match their performance on that metric which achieving the improved availability.

Fig. 3 focuses specifically on the percentage change in availability of CON relative to HMR. Subfigures (a) and (b) correspond to Fig. 1 and Fig. 2 respectively. In (a) CON typically performs better than HMR, though the effect is very small, substantially less than 1%. In contrast (b) shows a much larger effect, a 5% improvement in performance in a meaningful range where supply is somewhat less than demand.[1] Correlated demand alone (c) or the fairness objective alone (d) do not show this large benefit. Making the demand perfectly correlated, by making it exactly proportional to the total demand in the destination, shows benefits of nearly 10%. Overall, these results show that in the more plausible ranges of supply and demand CON can more effectively achieve a non-efficiency objective and are suggestive of the size of the benefit being driven by the correlation of the demand pattern.

Finally, we provide the results of several additional experiments. These are:

- An experiment that examines how our results for the correlated demand pattern and hybrid objective depend on the number of regions (using a ratio of 1 between supply and demand) and shows that the benefits do not significantly depend on the number of regions except for very small values of r (Fig. 4).
- An experiment that examines how correlation level of demand patterns affects the relative availability of CON compared to HMR, and shows that the relative availability of CON is positively related with the correlation level of demand patterns. The higher the correlation level of demand patterns, the better relative performance of CON over HMR (Fig. 5).
- An experiment with a third objective, directly optimizing for availability, which leads to similar results as our fairness objective (Fig. 6).
- An experiment in the spirit of our second experiment on the DiDi data that runs each policy for 100 steps with demand pattern changing every 5 steps, to represent changes in demand over the course of a day (Fig. 7).

To summarize the results of our experiments, we have seen that CON consistently converges substantially faster than other policies and has a performance that is often close to or matching LB. The effect of this on efficiency relative to the other policies is, however, quite small. When targeting a policy that puts weight on fairness rather than just efficiency, CON leads to economically meaningful improvements in availability, showings its ability to target a wider range of objectives than previous approaches.

C OMITTED PROOFS

PROOF OF LEM. 3

If suffices to show that for $M' = M + q\pi$, $M' = next(M')$. A region $i \in U_M$ is in Z_π by assumption, hence $M'[i] = M[i]$ and, since $i \in Z_\pi$, $out(M')[j, i] = out(M)[j, i]$ for every $j \notin U_M$. Thus $next(M)[i] = next(M')[i]$.

For a region i not in U_M, $M' = M + q\pi[i]$. While non-U_M regions contribute no additional flow into i, each other region j contributes $q\pi[j]\pi[i]$ additional flow into i. Since $\sum_{j \notin U_M} \pi[j]\pi[i] = \pi[i]$, hence the additional flow into i is $q\sigma[i]$. $\bowtie$

[1] CON performs a bit worse than HMR when the supply of drivers is sharply limited and the efficiency-focused approach appears beneficial. However, this is not the regime ridesharing platforms strive to operate in.

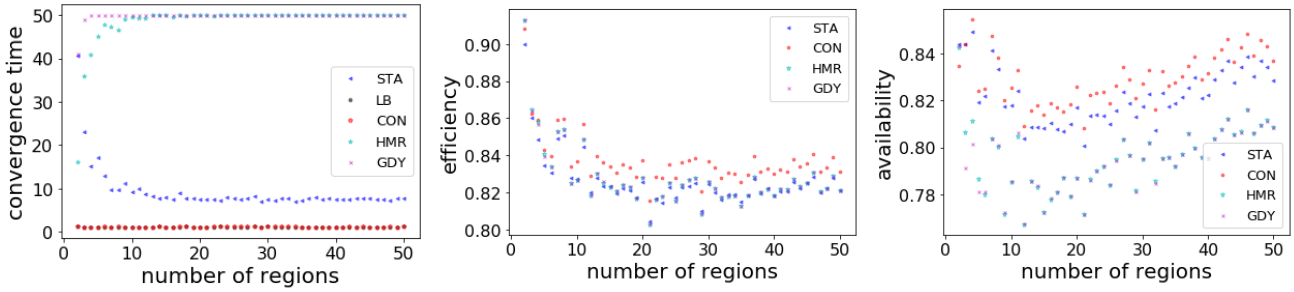

Figure 4: Performance varying number of regions

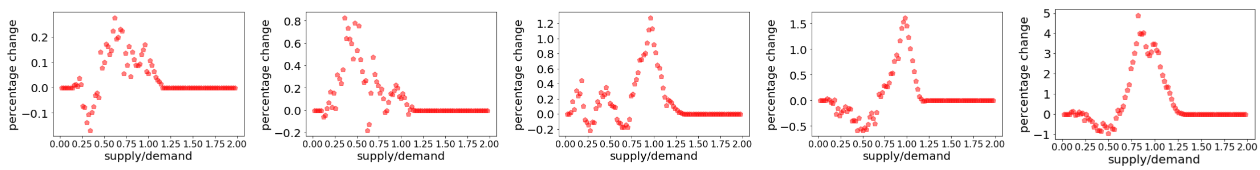

(a) correlation 0.05 (b) correlation 0.25 (c) correlation 0.45 (d) correlation 0.75 (e) correlation 0.95

Figure 5: Relative availability of CON to HMR for demand patterns of different correlation levels

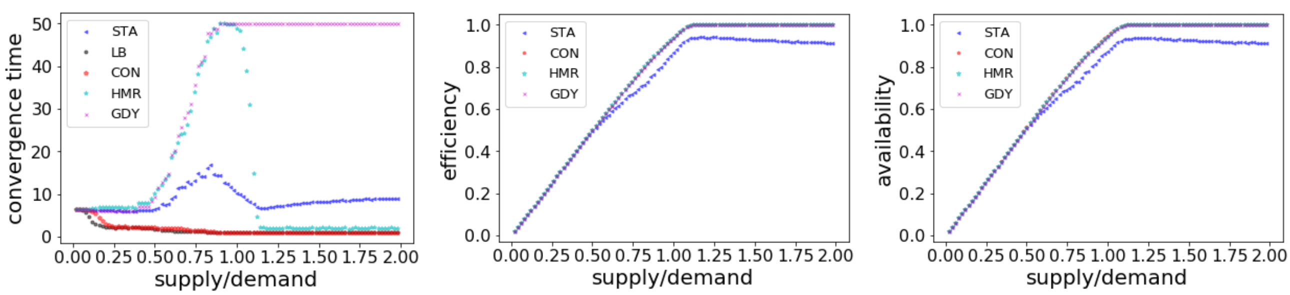

Figure 6: Performance with correlated demand; availability objective

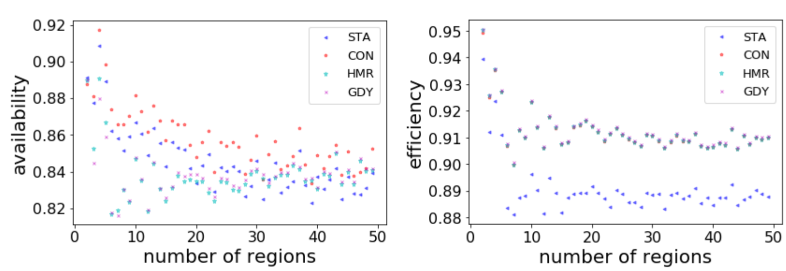

Figure 7: Performance varying demand

PROOF OF TH. 1

The proof is by showing that the following is an inductive invariant[2] of the while loop:

$$\varphi\colon\ M = next(M)\ \wedge\ \boldsymbol{M} = q - m \\ \wedge\ |U_M| \le i\ \wedge\ U_M \not\subseteq Z_\pi \to Erg(T_M)$$

Termination of the loop trivially follows since i, being an upper bound of the size of U_M, cannot go below 0.

Upon entering the loop for the first time, M is the all-0 vector and $m = q$ hence $M = next(M)$, $\boldsymbol{M} = q - m$, $|U_M| = r = i$, and $Erg(T_M)$, which establishes the base case.

Assume now a new iteration that starts when φ holds. Obviously, $m > 0$. We now distinguish between two cases, the first where $U_M \subseteq Z_\pi$, which in fact can be true only before the last iteration, and the case where $U_M \not\subseteq Z_\pi$ where the loop may terminate or enter a new iteration with a lower i.

In the first case, since σ is set to π and Δq is set to m, according to Lem. 3, after line 12 is executed, $M = next(M)$, $\boldsymbol{M} = q$ and $m = 0$. Hence the new M is the desired one and the loop terminates.

In the second case where there are unsaturated regions not in Z_π, the induction hypothesis implies that $Erg(T_M)$. Line 9 calls an *external* procedure to compute, in σ, the unique fixpoint of T_M, and Line 10 assigns to $\Delta q (> 0)$ the maximum available mass that can be distributed among the regions so to avoid an unsaturated becoming oversaturated. Hence, after Δq is distributed according to σ, at least one additional region becomes saturated. From Lem. 2 it follows that $M + \Delta q\sigma$, whose total mass is $\boldsymbol{M} + \Delta q$, is fixpoint mass distribution. Line 12 updates M and, implicitly, U_M. Thus φ, with the updated variables, holds also at the end of the iteration. If the new $m = 0$, then φ (over the new variables) imply that $M = next(M)$ and $\boldsymbol{M} = q$. Otherwise, another iteration is executed. $\bowtie$

PROOF OF LEM. 4

Consider a region i. From the definition of $next$ and the monotonicity of out in M (for every i and j) it follows that $next(M_0)[i] = next(M_1)[i]$ iff $out(M_0)[j,i] = out(M_1)[j,i]$ for every region j. It thus suffices to prove that for every region j such that $M_0[j] < M_1[j]$,

$$out(M_0)[j,i] = out(M_1)[j,i] \quad \text{iff} \quad i \in Z_\pi \text{ and } j \notin J$$

For $\Rightarrow$, assume for contraposition that $j \in J$. Since $M_0[j] < M_1[j]$, $V[j,i](\min(M_1[j], W\mathbb{1}[j]) - M_0[j]))$ of j's additional outflow is directed towards i implying that $out(M_0)[j,i] \neq out(M_1)[j,i]$. If, however, $j \notin J$, j's additional outflow is distributed according to π, and since $i \notin Z_\pi$, $\pi[i]$ fraction of it towards i implying $out(M_0)[j,i] \neq out(M_1)[j,i]$.

For $\Leftarrow$, since $j \notin J$ the additional outflow is directing according to π. Since $i \in Z_\pi$, $out(M_0)[j,i] = out(M_1)[j,i]$. $\bowtie$

[2] An *inductive invariant* is an assertion that holds upon entry to the loop, and if it holds at the beginning of the loop then it holds when the loop is next re-entered.

PROOF OF TH. 2

From Th. 1, such a fixpoint exists. It suffices to show that the fixpoint is unique.

Assume, by way of contradiction, that there are two distinct such fixpoints, say $\widehat{M_0}$ and $\widehat{M_1}$. Let $\widehat{M} = \min(\widehat{M_0}, \widehat{M_1})$. Since $\widehat{M_0}$ and $\widehat{M_1}$ are distinct fixpoints in S_q, $\widehat{M}$ is not in S_q.

Let

$$\mathcal{L} = \{M \in (\mathbb{R}_{\ge 0} \cup \{\infty\})^r : \widehat{M} \le M\}$$

With the vector (point-wise) $\le$ as a preorder, and vector (point-wise) $\min$ (resp. $\max$) as meet (resp. join), $\mathcal{L}$ is a complete lattice.

While $\mathcal{L}$ is complete lattice, it may not be closed under $next$. We thus define $aux = \max(M, next(M))$ as an auxiliary function under which $\mathcal{L}$ is closed. From the Knaster-Tarski theorem it follows that aux has a set of fixpoints in $\mathcal{L}$, and that the set of aux's fixpoints in $\mathcal{L}$ is a complete lattice.

Both $\widehat{M_0}$ and $\widehat{M_1}$ are fixpoints of $next$ in $S_q \subseteq \mathcal{L}$, as well as fixpoints of aux in $\mathcal{L}$. Since $\widehat{M} = \widehat{M_0} \sqcap \widehat{M_1}$ is the minimum element of $\mathcal{L}$, it is the least fixpoint of aux in $\mathcal{L}$.

Define $I_0 = \{i : \widehat{M_0}[i] < \widehat{M_1}[i]\}$ and $I_1 = \{i : \widehat{M_0}[i] > \widehat{M_0}[i]\}$, that is, for every $i \in I_0$, $\widehat{M}[i] = \widehat{M_0}[i]$, for every $i \in I_1$, $\widehat{M}[i] = \widehat{M_1}[i]$, and for every $i \notin I_0 \cup I_1$, $\widehat{M}[i] = \widehat{M_0}[i] = \widehat{M_1}[i]$. Recall that Z_π is the set $\{j : \pi[j] = 0\}$. Obviously, I_0 and I_1 are disjoint and neither is empty. By construction, $\widehat{M_0}, \widehat{M_1} \gneq \widehat{M}$, and all three are fixpoints of $next$. It then follows from Lem. 4 that for every $i \in I_0 \cup I_1$, $\widehat{M}[i]$ is saturated in $\widehat{M}$ and therefore in both $\widehat{M_0}$ and $\widehat{M_1}$. It then follows from Lem. 4 that Z_π is pairwise disjoint from both I_0 and I_1.

If Z_π is empty, it follows from Lem. 4 that $\widehat{M}$ is strictly less than both $\widehat{M_0}$ and $\widehat{M_1}$, which is obviously a contradiction. Assume therefore that Z_π is nonempty.

Now consider some $i \in I_0$. Since region i is saturated in both $\widehat{M_0}$ and $\widehat{M}$, it follows from Lem. 4 that for every $j \notin Z_\pi$, $next(\widehat{M})[j] < next(\widehat{M_0})[j]$. Since I_0 is disjoint from Z_π, it follows that $\widehat{M}[i] < \widehat{M_0}[i]$, contradicting the assumption that $i \in I_0$. $\bowtie$

PROOF OF LEM. 5

Assume $\langle F_0, E_0\rangle$, i_0, and j_0 as in the lemma statement. We construct $\langle F, E\rangle$ that is also feasible for q, but with $\boldsymbol{F} > \boldsymbol{F_0}$, thus showing that $\langle F_0, E_0\rangle$ is nonoptimal.

Let x be the vector induced by $\min(E_0[i_0, j_0], (W - F_0)\mathbb{1}[i_0])$ given $V[i_0]$.

Define F_x be such that $F_x[i_0] = x$ and 0 elsewhere, and E_x such that:

$$E_x[i,j] \;=\; \begin{cases} x[i] & i \neq i_0, j = j_0 \\ \boldsymbol{x} - x[i_0] & i = j_0 \neq i_0 = j \\ 0 & \text{otherwise} \end{cases}$$

We claim that $\langle F_x, E_x\rangle$ feasible for $2\boldsymbol{x}-x[i_0]$ if $i_0 = j_0$ and $3\boldsymbol{x}-2x[i_0]$ otherwise: The D- and P- constraints are satisfied by construction. As for the F-constraint, every region i, $i \neq i_0, j_0$, receives $x[i]$ full cars and sends $x[i]$ empty cars. Region i_0 sends $\boldsymbol{x}$ full cars and receives $x[i_0]$ cars full from itself, and (whether or not $i_0 = j_0$), $\boldsymbol{x} - x[i_0]$ empty cars from other regions. Finally, when $j_0 \neq i_0$, region j_0 receives $x[j_0]$ full cars from i_0 and $x[i]$ empty cars from every region $i, i \neq i_0$, thus it receives $\boldsymbol{x}-x[i_0]+x[j_0]$ cars in total, which is exactly the number of cars it sends ($\boldsymbol{x} - x[i_0]$ empty, to i_0 and $x[j_0]$ empty to j_0).

Let $F_1 = F_0 + F_x$, and $E_1 = E_0 + E_x$. The choice of x guarantees that $F_1 \leq W$, hence it follows from Ob. 2 that $\langle F_1, E_1\rangle$ is feasible for $q + 2\boldsymbol{x} - x[i_0]$ (or $q + 3\boldsymbol{x} - 2[i_0]$ if $i_0 \neq j_0$). Let $y = \boldsymbol{x} - x[i_0]$ ($y = 2\boldsymbol{x} - 2x[i_0]$). Obviously, $y > 0$. If $i_0 = j_0$, then $E_0[i_0, i_0] \geq \boldsymbol{x} > y$. If $i_0 \neq j_0$, then $E_0[i_0, j_0] \geq \boldsymbol{x} > y/2$ and $E_x[j_0, i_0] = \boldsymbol{x} - x[i_0] \geq y/2$. It then follows from Ob. 2 that $\langle F_2, E_2\rangle$ where $F_2 = F_1$ and E_2 is just like E with y subtracted from $E[i_0, i_0]$ (in the case of $i_0 = j_0$) or $y/2$ subtracted from both $E_1[i_0, j_0]$ and $E_1[j_0, i_0]$ (in the case of $i_0 \neq i_j$) is feasible for $q + \boldsymbol{x}$.

Let $c = q/(q+\boldsymbol{x})$, $F = c \cdot F_2$ and $E = c \cdot E_2$. From Ob. 1 it follows that $\langle F, E\rangle$ is feasible for q. Moreover,

$$\boldsymbol{F} = c(\boldsymbol{F_2}) = c(\boldsymbol{F_0} + \boldsymbol{x}) = \frac{q\boldsymbol{F_0} + q\boldsymbol{x}}{q + \boldsymbol{x}} > \boldsymbol{F_0}$$

where the last inequality follows from $q > \boldsymbol{F_0}$ (which holds since $\boldsymbol{E_0} > 0$.)

Consequently, $\langle F, E\rangle$ is feasible for q and $\boldsymbol{F} > \boldsymbol{F_0}$ so that $\langle F_0, E_0\rangle$ is not optimal. $\bowtie$

PROOF OF LEM. 6

Let $c = \min_{i,j}(V[i,j])$. If $\mathbf{E}(\mathbf{M_t}) > 0$, then $[\nu(q) - \mathbb{1}F(M_t)]_+$ is monotone increasing in $\nu(q)$ on an open interval including q_{t+1}. Subsequently we show that

$$\sum_i [\nu(q_t + c\Delta(t)) - \mathbb{1}F(M_t)]_+[i] \leq \sum_i [[\nu(q_{t+1}) - \mathbb{1}F(M_t)]]_+[i]$$

Monotonicity then implies that $\nu(q_{t+1}) \geq \nu(q_t + c\Delta(t)$, and therefore $q_{t+1} - q_t > c\Delta(t)$.

If $\mathbf{E}(\mathbf{M_t}) = 0$, there is no relocation and we cannot apply the previous argument because we only know that $[\nu(q) - \mathbb{1}F(M_t))]_+$ is monotone non-decreasing. Subsequently we show that then $M_{t+1}[i] - \nu(q_t)[i] \geq c\Delta(t)$. Thus $M_{t+1} = \nu(q_t) + (M_{t+1} - \nu(q_t)) \geq \nu(q_t) + c\Delta(t)\mathbb{1} \geq \nu(q_t + c\Delta(t))$ and $q_{t+1} \geq q_t + c\Delta(t)$. $\bowtie$

PROOF OF LEM. 6 FOR THE CASE THAT $\mathbf{E}(\mathbf{M_t}) > 0$

$$\begin{aligned}
&\textstyle\sum_i [\nu(q_t + c\Delta(t)) - \mathbb{1}F(M_t)]_+[i] \\
&\overset{(1)}{\leq} \textstyle\sum_i [(\nu(q_t) + c\Delta(t)\mathbb{1} - \mathbb{1}F(M_t))]_+[i] \\
&\overset{(2)}{=} \textstyle\sum_i [\mathbb{1}F(\nu(q_t)) + \mathbb{1}E(\nu(q_t)) + c\Delta(t)\mathbb{1} - \mathbb{1}F(M_t)]_+[i] \\
&\overset{(3)}{=} \textstyle\sum_i (\!\!\!\;[\mathbb{1}F(\nu(q_t)) + \mathbb{1}E(\nu(q_t)) + c(\mathbf{F}(\mathbf{M_t}) - \mathbf{F}(\nu(\mathbf{q_t})))\mathbb{1} + \\
&\qquad c(\mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\nu(\mathbf{q_t})))\mathbb{1} - \mathbb{1}F(M_t)]\!\!\!\;)_+[i] \\
&\overset{(4)}{\leq} \textstyle\sum_i (\!\!\!\;[\mathbb{1}F(\nu(q_t)) + \mathbb{1}E(\nu(q_t)) + \mathbb{1}F(M_t) - \mathbb{1}F(\nu(q_t)) + \\
&\qquad c(\mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\nu(\mathbf{q_t})))\mathbb{1} - \mathbb{1}F(M_t))]\!\!\!\;)_+[i] \\
&\overset{(5)}{=} \textstyle\sum_i [\mathbb{1}E(\nu(q_t)) + c(\mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\nu(\mathbf{q_t})))\mathbb{1}]_+[i] \\
&\overset{(6)}{=} \mathbf{E}(\nu(\mathbf{q_t})) + c(\mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\nu(\mathbf{q_t}))) \\
&\overset{(7)}{<} \mathbf{E}(\mathbf{M_t}) \\
&\overset{(8)}{=} \textstyle\sum_i [\nu(q_{t+1}) - \mathbb{1}F(M_t)]_+[i]
\end{aligned}$$

(1) follows from the monotonicity of ν; (2) follows since $\nu(q_t)$ is a fixpoint; (3) follows from the expansion $\Delta(t) = \mathbf{M_t} - q_t$, followed mass conservation; (4) follows by the definition of c, $\mathbb{1}F(M_t)$ and $\mathbb{1}F(\nu(q_t))$; In particular, $(\mathbb{1}F(M_t) - \mathbb{1}F(\nu(q_t)))[i] = \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(Q_t))\mathbb{1}[j])V[j,i] \geq \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(Q_t))\mathbb{1}[j]) \min_{k,\ell} V[k,l] = c(\mathbf{F}(\mathbf{M_t}) - \mathbf{F}(\nu(\mathbf{q_t})))$; (5) follows by cancellation; (6) follows because $\mathbb{1}E(\nu(q_t)) \geq \mathbb{0}$, and $M_t \geq \nu(q_t)$ and so $E(M_t) \geq E(\nu(q_t))$; (7) follows because $c < 1$ and $E(M_t) \geq E(\nu(q_t))$; (8) is by definition of q_{t+1}.

PROOF OF LEM. 6 FOR THE CASE THAT $\mathbf{E}(\mathbf{M_t}) = 0$.

$$\begin{aligned} M_{t+1}[i] - \nu(q_t)[i] &= \\ &\overset{(1)}{=} \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(q_t))\mathbb{1}[j])V[j,i] \\ &\overset{(2)}{\geq} \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(q_t))\mathbb{1}[j])c \\ &\overset{(3)}{=} c\Delta(t) \end{aligned}$$

where (1) and (3) hold because of lack of relocation. Thus $M_{t+1} = \nu(q_t) + (M_{t+1} - \nu(q_t)) \geq \nu(q_t) + c\Delta(t)\mathbb{1} \geq \nu(q_t + c\Delta(t))$ and $q_{t+1} \geq q_t + c\Delta(t)$.

PROOF OF COR. 1

Let $c = \min(\min_{\{i|\pi[i]>0\}}(\pi[i]), \min_{i,j}(V[i,j]))$. As we are now following π at each step, let $M_{t+1} = next_\pi(M_t)$ and $q_t = \nu^{-1}(M_t)$.

If $\pi[i] > 0$ then:

$$\begin{aligned} M_{t+1}[i] - \nu(q_t)[i] & \\ &\overset{(1)}{=} \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(q_t))\mathbb{1}[j])V[j,i] + \\ &\qquad \sum_j (E(M_t)\mathbb{1}[j] - E(\nu(q_t))\mathbb{1}[j])\pi[i] \\ &\overset{(2)}{\geq} \sum_j (F(M_t)\mathbb{1}[j] - F(\nu(q_t))\mathbb{1}[j])c + \\ &\qquad \sum_j (E(M_t)\mathbb{1}[j] - E(\nu(q_t))\mathbb{1}[j])c \\ &\overset{(3)}{=} c\Delta(t) \end{aligned}$$

If $\pi[i] = 0$ the same conclusion holds by omitting the terms involving $\mathbf{E}$. Thus $M_{t+1} = \nu(q_t) + (M_{t+1} - \nu(q_t)) \geq \nu(q_t) + c\Delta(t)\mathbb{1} \geq \nu(q_t + c\Delta(t))$ and M_t converges at least linearly to M^*. $\bowtie$

PROOF OF LEM. 7

$$\begin{aligned} \Delta^E(t) &\overset{(1)}{=} \mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\nu(q_t)) \\ &\overset{(2)}{=} \sum_i [(\nu(q_{t+1}) - \mathbb{1}F(M_t)]_+[i] - \mathbf{E}(\nu(q_t)) \\ &\overset{(3)}{=} \sum_i [\nu(q_{t+1}) - \nu(q_t) + \nu(q_t) - \mathbb{1}F(M_t)]_+[i] \\ &\qquad - \mathbf{E}(\nu(q_t) \\ &\overset{(4)}{=} \sum_i [\![(\nu(q_{t+1}) - \nu(q_t) + \mathbb{1}F(\nu(q_t)) + \mathbb{1}E(\nu(q_t)) \\ &\qquad - \mathbb{1}F(M_t)]\!]_+[i] - \mathbf{E}(\nu(q_t)) \\ &\overset{(5)}{\leq} \boldsymbol{\nu}(\mathbf{q_{t+1}}) - \boldsymbol{\nu}(\mathbf{q_t})) + \mathbf{E}(\nu(\mathbf{q_t})) - \mathbf{E}(\nu(\mathbf{q_t})) \\ &\overset{(6)}{=} q_{t+1} - q_t \end{aligned}$$

(1) follows from the definition of $\Delta^E(t)$; (2) follows from definition of q_{t+1}; (3) adds and subtracts $\nu(q_t)$; (4) follows because $\nu(q_t)$ is a fixpoint; (5) follows because $\nu(q_{t+1}) \geq \nu(q_t)$ and $\mathbb{1}F(\nu(q_t)) \leq \mathbb{1}F(M_t)$; (6) follows because q_t is by definition the total mass of $\nu(q_t)$. $\bowtie$

PROOF OF OB. 3

$$\begin{aligned} &\sum_t (\mathbf{E}(\mathbf{M_t}) - \mathbf{E}(\mathbf{M}^*)) \\ &\overset{(1)}{=} \sum_t (\mathbf{E}(\nu(\mathbf{q_t})) + \Delta^E(t) - \mathbf{E}(\mathbf{M}^*)) \\ &\overset{(2)}{\leq} \sum_t \Delta^E(t) \overset{(3)}{\leq} \sum_t (q_{t+1} - q_t) \overset{(4)}{\leq} \mathbf{M}^* \end{aligned}$$

(1) follows by the definition of $\Delta^E(t)$; (2) follows since $\nu(q_t) \leq M^*$, hence $\mathbf{E}(\nu(\mathbf{q_t})) \leq \mathbf{E}(\mathbf{M}^*)$; (3) follows from Lem. 7 and(4) line follows since the q_ts are the masses of the $\nu(q_t)$s. $\bowtie$