Dynamic Pricing Provides Robust Equilibria in Stochastic Ride-Sharing Networks: Online Companion

J. Massey Cashore* Peter I. Frazier† Éva Tardos‡

Cornell University

Abstract

This document is an online companion to our paper *Dynamic Pricing Provides Robust Equilibria in Stochastic Ridesharing Networks*. It includes additional proofs from the appendix, about concentration properties of the matching process and general convex analysis properties.

---

*Supported by an NSERC PGS D Fellowship

†Supported by AFOSR grant FA9550-19-1-0283

‡Supported by AFOSR grant FA9550-19-1-0183 and NSF grants CCF-1408673 and CCF-1563714

G Partial Derivatives of the State-Dependent Optimization Problem

In this section we prove a number of results about the fluid optimization problem.

G.1 Proof of Lemma 1

Lemma. *Consider the reward function $U_{(\ell,d,\omega_t)}$ associated with any route $(\ell,d) \in \mathcal{L}^2$ and any scenario ω_t. Assume the rider-value distribution $F_{(\ell,d,\omega_t)}$ satisfies Assumption 2. Then $U_{(\ell,d,\omega_t)}(g)$ is concave in g, is differentiable at every $g > 0$, and the derivative at each $g > 0$ satisfies:*

$$\frac{d}{dg} U_{(\ell,d,\omega_t)}(g) = P_{(\ell,d,\omega_t)}(g).$$

Moreover, the fluid optimization problem (6) has a concave objective function for any market state.

Proof. Let us write the utility function as $U(f) = U_{(\ell,d,\omega_t)}(f)$, First, observe that the utility function can be equivalently written as the following equation

$$U(g) = \bar{D}\mathbb{E}\left[V\mathbf{1}_{\left\{V \geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right\}}\right], \tag{59}$$

where we write $\bar{D}$ in place of $\bar{D}_{(\ell,d,\omega_t)}$, F to mean $F_{(\ell,d,\omega_t)}$, and $f \wedge \bar{D}$ to mean $\min(f, \bar{D})$. This characterization is justified by the following series of equalities:

$$\begin{aligned}
\mathbb{E}\left[V\mathbf{1}_{\left\{V\geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right\}}\right] &= \mathbb{E}\left[V \mid V \geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right]\mathbb{P}\left(V \geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right) \\
&= \mathbb{E}\left[V \mid V \geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right]\left(1-F\left(F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right)\right) \\
&= \mathbb{E}\left[V \mid V \geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right]\frac{g\wedge\bar{D}}{\bar{D}}.
\end{aligned}$$

Next, recall that if X is a uniform $[0,1]$ random variable then $F^{-1}(X)$ is a random variable with distribution function F. Using the fact that if X is uniform $[0,1]$ then so too is $1-X$, from the characterization (59) we have the following equalities:

$$\begin{aligned}
U(g) &= \bar{D}\int_0^1 F^{-1}(u)\mathbf{1}_{\left\{F^{-1}(u)\geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right\}}du \\
&= \bar{D}\int_0^1 F^{-1}(1-u)\mathbf{1}_{\left\{F^{-1}(1-u)\geq F^{-1}\left(1-\frac{g\wedge\bar{D}}{\bar{D}}\right)\right\}}du \\
&= \bar{D}\int_0^1 F^{-1}(1-u)\mathbf{1}_{\left\{\frac{g\wedge\bar{D}}{\bar{D}}\geq u\right\}}du \\
&= \bar{D}\int_0^{\frac{g\wedge\bar{D}}{\bar{D}}} F^{-1}(1-u)du.
\end{aligned}$$

Now, fix any $0 < g < \bar{D}$ and consider the above expression. The assumption $g < \bar{D}$ means that the integral upper bound $\frac{g\wedge\bar{D}}{\bar{D}}$ is simply $\frac{g}{\bar{D}}$, and from the assumption that F^{-1} satisfies Assumption 2 we know that $F^{-1}(\cdot)$ is continuous. Hence, by the fundamental theorem of calculus the function $U(g)$ is differentiable at $0 < g < \bar{D}$, and using the chain rule we compute the derivative to be:

$$\frac{d}{dg}U(g) = \bar{D}F^{-1}\left(1 - \frac{g}{\bar{D}}\right)\frac{1}{\bar{D}} = F^{-1}\left(1 - \frac{g}{\bar{D}}\right).$$

Further, note that $U(g)$ is constant for $g \geq \bar{D}$, from which we conclude that $\frac{d}{dg}U(g) = 0$ for $g > \bar{D}$. Finally, we must establish existence of the derivative at the point $g = \bar{D}$. To this end, it suffices to show that limit of the partial derivatives from below $g = \bar{D}$ and from above $g = \bar{D}$ are equal. This fact follows from the second assertion in Assumption 2 which states that $F^{-1}(0) = 0$:

$$\lim_{g \uparrow \bar{D}} \frac{d}{dg}U(g) = \lim_{g \uparrow \bar{D}} F^{-1}\left(1 - \frac{g}{\bar{D}}\right) = 0 = \lim_{g \downarrow \bar{D}} \frac{d}{dg}U(g).$$

Thus, we have established the derivative $\frac{d}{dg}U(g)$ exists for all $g > 0$ and is equal to $P_{(\ell,d,\omega_t)}(g)$. Concavity of $U(\cdot)$ follows from the observation that this derivative is non-increasing in g.

To show that (6) we has a concave objective we have to show that the remaining terms in $\mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ are also concave. Well, clearly the linear costs are concave. And finally, the add-passenger disutility cost function $-A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell)$ is convex, from the equation (10) and Lemma 20. ☐

Lemma 20. *Define*

$$f(x, y) = \begin{cases} \frac{y^2}{x} & \text{if } x > 0, \\ 0 & \text{if } x = 0. \end{cases}$$

Then f is convex over the domain $x, y \geq 0$, $y \leq x$.

Proof. Let (x_1, y_1) and (x_2, y_2) be two points in the domain of f, and let $\lambda \in [0, 1]$. We need to check

$$f(\lambda(x_1, y_1) + (1-\lambda)(x_2, y_2)) \leq \lambda f(x_1, y_1) + (1-\lambda)f(x_2, y_2).$$

Let's start with the case where x_1 and x_2 are both nonzero. In this case, the inequality we have to check is given by

$$\frac{(\lambda y_1 + (1-\lambda)y_2)^2}{\lambda x_1 + (1-\lambda)x_2} \leq \lambda\frac{y_1^2}{x_1} + (1-\lambda)\frac{y_2^2}{x_2}.$$

We will verify this inequality by applying the Cauchy-Schwarz inequality. Define the following values:

$$u_1 = \frac{\lambda y_1}{\sqrt{\lambda x_1}} \quad u_2 = \frac{(1-\lambda)y_2}{\sqrt{(1-\lambda)x_2}} \quad v_1 = \sqrt{\lambda x_1} \quad v_2 = \sqrt{(1-\lambda)x_2}.$$

The Cauchy-Schwarz inequality says $u^Tv \leq \|u\|\|v\|$. Observe the following equalities:

$$\begin{aligned} (u^Tv)^2 &= (\lambda y_1 + (1-\lambda)y_2)^2, \\ \|v\|^2 &= \lambda x_1 + (1-\lambda)x_2, \\ \|u\|^2 &= \lambda\frac{y_1^2}{x_1} + (1-\lambda)\frac{y_2^2}{x_2}. \end{aligned}$$

Rearranging C-S we have

$$\frac{(u^Tv)^2}{\|v\|^2} \leq \|u\|^2,$$

which implies the desired inequality holds, and hence f is convex whenever x_1 and x_2 are both nonzero.

The case where x_1 and x_2 are both zero is immediate. It remains to check the case where x_1 is nonzero and x_2 is 0. In this case, y_2 must also be zero, because of the constraint $y \leq x$. Therefore we have to verify

$$f(\lambda(x_1, y_1)) \leq \lambda f(x_1, y_1).$$

By inspection we see that $f(\lambda(x, y)) = \lambda f(x, y)$ is always satisfied, so f is convex in this case as well. □

G.2 State-Dependent Optimization Problem

For clarity we restate the state-dependent optimization problem below. Fix a time period t and a scenario ω_t. The state-dependent optimization problem depends on a supply-location vector $\mathbf{S} = (S_\ell : \ell \in \mathcal{L})$ where each component $S_\ell \geq 0$ specifies the volume of active drivers at location ℓ. Active drivers at a location ℓ consist of drivers who took a trip destined towards ℓ at the previous time period $t-1$, as well as new drivers who enter the market at location ℓ in the current time period t. The state-dependent optimization problem solves for the welfare-optimal trips in the current time period in the stochastic fluid model, given the market state specified by the scenario ω_t and the supply-location vector $\mathbf{S}$. We write $\Phi_{\omega_t}(\mathbf{S})$ to denote the value of the state-dependent optimization problem under the scenario ω_t as a function of the supply-location vector $\mathbf{S}$. The function $\Phi_{\omega_t}(\mathbf{S})$ is formally defined as the value of the following optimization problem:

$$\begin{aligned}
\Phi_{\omega_t}(\mathbf{S}) \equiv \quad & \sup_{\boldsymbol{f},\boldsymbol{g}} \quad \mathcal{U}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + \mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right] && (60)\\
& \text{subject to} \\
& \qquad f_{(\ell,d)} \geq 0 \qquad \forall (\ell, d) \in \mathcal{L}^2 && (61)\\
& \qquad g_{(\ell,d)} \geq 0 \qquad \forall (\ell, d) \in \mathcal{L}^2 && (62)\\
& \qquad f_{(\ell,d)} \geq g_{(\ell,d)} \qquad \forall (\ell, d) \in \mathcal{L}^2 && (63)\\
& \qquad \sum_{d \in \mathcal{L}} f_{(\ell,d)} = S_\ell \qquad \forall \ell \in \mathcal{L} && (64)
\end{aligned}$$

The decision variable $\boldsymbol{f} = (f_{(\ell,d)} : (\ell, d) \in \mathcal{L}^2)$ has components $f_{(\ell,d)}$ which specify the total trip volume along each route (ℓ, d). By total trip volume we mean$f_{(\ell,d)}$ specifies the sum of the relocation-trip volume and the dispatch trip-volume. The decision variable $\boldsymbol{g} = (g_{(\ell,d)} : (\ell, d) \in \mathcal{L}^2)$ has components $g_{(\ell,d)}$ which specify the dispatch trip volume along each route (ℓ, d).

The objective function is the sum of two functions: $\mathcal{U}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ specifies the welfare collected in the current time period t, and $\mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right]$ specifies the welfare to be collected in future time periods.

When $t = T$ is the final time period we just take $\mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right]$ to be 0. When $t < T$, we define $\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f})$ to be the supply-location vector arising at time $t+1$ under the trip volumes specified by $\boldsymbol{f}$ and the future scenario ω_{t+1}. The expectation $\mathbb{E}_{\omega_t}[\cdot]$ is taken over all time $t+1$ scenarios given the time t scenario ω_t. The function $\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f})$ follows the convention that upply-location vectors include new drivers who enter the market in the relevant time period. Let us write $S_{\omega_{t+1},\ell}(\boldsymbol{f})$ for the component of the supply-location vector $\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f})$ corresponding to location ℓ.

$S_{\omega_{t+1},\ell}(\boldsymbol{f})$ is defined formally by the following equation

$$S_{\omega_{t+1},\ell}(\boldsymbol{f}) = \frac{1}{k} M_{\omega_{t+1},\ell} + \sum_{o\in\mathcal{L}} f_{(o,\ell)}, \tag{65}$$

where $M_{\omega_{t+1},\ell}$ is the volume of new drivers who enter the market at location ℓ under the scenario ω_{t+1} and the sum is over all routes whose destination location is ℓ.

The utility collected in the current time period is the difference between the rider value we generate by serving dispatches and the disutility that drivers incur. Since we assume the price is a transfer from riders to drivers the price does not appear explicitly in the objective function. The function $\mathcal{U}_{\omega_t}(\boldsymbol{f},\boldsymbol{g})$ is formally defined by the following equation:

$$\mathcal{U}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) = \sum_{(\ell,d)\in\mathcal{L}^2} U_{(\ell,d,\omega_t)}(g_{(\ell,d)}) - \left(\sum_{(\ell,d)\in\mathcal{L}^2} c_{(\ell,d)} f_{(\ell,d)} + \sum_{\ell} A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell)\right). \tag{66}$$

The function $U_{(\ell,d,\omega_t)}(g_{(\ell,d)})$ specifies the total rider value generated as a function of dispatch-trip volume along the route (ℓ,d), the function $A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell)$ specifies the total add-passenger disutility incurred by drivers located at ℓ, as a function of the volume of dispatch trips originating from ℓ, $\boldsymbol{g}^T\mathbf{1}_\ell$, and the total volume of available drivers located at ℓ, $\boldsymbol{f}^T\mathbf{1}_\ell$.

We take $\mathbf{1}_\ell$ to be an indicator vector indexed by pairs of locations, where the value corresponding to each $(\ell',d)\in\mathcal{L}^2$ is 1 if $\ell'=\ell$ and 0 otherwise. With this convention, the quantities $\boldsymbol{g}^T\mathbf{1}_\ell$ and $\boldsymbol{f}^T\mathbf{1}_\ell$ specify the volume of dispatch trips originating from ℓ and the total volume of trips originating from ℓ, respectively:

$$\boldsymbol{g}^T\mathbf{1}_\ell = \sum_{d\in\mathcal{L}} g_{(\ell,d)},$$
$$\boldsymbol{f}^T\mathbf{1}_\ell = \sum_{d\in\mathcal{L}} f_{(\ell,d)}.$$

Assuming $\boldsymbol{f}$ satisfies the flow-conservation constraint (64), the total trip volume originating from ℓ is equal to the total volume of supply positioned at ℓ:

$$\boldsymbol{f}^T\mathbf{1}_\ell = S_\ell.$$

G.3 Optimality Conditions

We now derive the Lagrangian optimality conditions for the state-dependent optimization problem (60). For succinctness, we use the following notation for the objective function:

$$\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) = \mathcal{U}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) + \mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right]. \tag{67}$$

We begin by converting the optimization problem to a convex minimization problem where all inequality constraints have an upper bound of 0:

$$-\Phi_{\omega_t}(\mathbf{S}) \equiv \inf_{\boldsymbol{f},\boldsymbol{g}} \; -\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) \tag{68}$$

subject to

$$-f_{(\ell,d)} \leq 0 \qquad \forall(\ell,d)\in\mathcal{L}^2 \tag{69}$$
$$-g_{(\ell,d)} \leq 0 \qquad \forall(\ell,d)\in\mathcal{L}^2 \tag{70}$$
$$g_{(\ell,d)} - f_{(\ell,d)} \leq 0 \qquad \forall(\ell,d)\in\mathcal{L}^2 \tag{71}$$
$$\sum_{d\in\mathcal{L}} f_{(\ell,d)} - S_\ell = 0 \qquad \forall\ell\in\mathcal{L} \tag{72}$$

We associate dual variables $\alpha_{(\ell,d)}$, $\beta_{(\ell,d)}$, $\gamma_{(\ell,d)}$ and η_ℓ with each of the constraints (69), (70), (71), (72), respectively. We will write $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}$ to indicate the vector of dual variables.

Since all we have done is changed the sign and direction of the objective function and algebraically rearranged the inequality constraint functions, the optimization problems (68) and (60) have the same set of optimal solutions.

We now obtain the Lagrangian function for the optimization problem (68):

$$\begin{aligned} L(\boldsymbol{f}, \boldsymbol{g}; \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}) = -\mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) &+ \sum_{(\ell,d)} \left[\gamma_{(\ell,d)}(g_{(\ell,d)} - f_{(\ell,d)}) - \alpha_{(\ell,d)} f_{(\ell,d)} - \beta_{(\ell,d)} g_{(\ell,d)}\right] \\ &+ \sum_{\ell} \eta_\ell \left(\sum_d f_{(\ell,d)} - S_\ell \right). \end{aligned}$$

Because all of the constraints for the problem (68) are linear, and the primal problem (60) has a finite optimal solution, we know strong duality holds. Therefore, a feasible solution $(\boldsymbol{f}, \boldsymbol{g})$ is optimal if and only if there exist feasible dual variables $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}$ for which the stationarity conditions and the complementary slackness conditions hold. For dual feasibility to hold the variables associated with inequality constraints must be nonnegative, that is the following inequalities must hold pointwise:

$$\boldsymbol{\alpha} \geq 0,\ \boldsymbol{\beta} \geq 0,\ \boldsymbol{\gamma} \geq 0.$$

The complementary slackness conditions are satisfied when the following equations hold for all origin-destination pairs $(\ell, d) \in \mathcal{L}^2$:

$$\alpha_{(\ell,d)} f_{(\ell,d)} = 0,\ \beta_{(\ell,d)} g_{(\ell,d)} = 0,\ \gamma_{(\ell,d)}(g_{(\ell,d)} - f_{(\ell,d)}) = 0,$$

that is the dual variables associated with inequality constraints must be 0 unless the corresponding inequality constraint is tight at the primal solution.

Finally, the stationarity conditions are satisfied when the primal solution $(\boldsymbol{f}, \boldsymbol{g})$ are a stationary point of the Lagrangian function when the dual variables are held fixed. Notice that when we hold the dual variables fixed the Lagrangian is a convex function of the primal solution, so a primal solution $(\boldsymbol{f}, \boldsymbol{g})$ is a stationary point if and only if 0 is a subgradient of the Lagrangian at $(\boldsymbol{f}, \boldsymbol{g})$. We use the notation $\partial L(\boldsymbol{f}, \boldsymbol{g}; \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta})$ to refer to the subgradient of the Lagrangian where the dual variables $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}$ are held fixed. The subgradient condition for a primal solution $(\boldsymbol{f}, \boldsymbol{g})$ to be a stationary point can thus be expressed as follows

$$0 \in \partial L(\boldsymbol{f}, \boldsymbol{g}; \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}).$$

We work in terms of the subgradient because the objective function $\mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ is not differentiable at coordinates where $f_{(\ell,d)} = 0$ or $g_{(\ell,d)} = 0$. However, we can use the following property (see Theorem 3.1.8 in the textbook by Borwein and Lewis) about general convex functions to obtain a stationarity condition in terms of the partial derivatives for the nonzero coordinates of $\boldsymbol{f}$ and $\boldsymbol{g}$:

Lemma. *Let $h : \mathbb{R}^m \to \mathbb{R}^n$ be a convex function and consider any point $x \in \mathbb{R}^m$ in its domain. Let $\partial h(x_0)$ be the subdifferential of h at x and assume the partial derivative $\frac{\partial}{\partial x_j} h(x)$ exists for some coordinate j. Then the jth component of every subgradient in the subdifferential of h at x is equal to the partial derivative of h at x. That is, for every $\phi \in \partial h(x)$ the equality $\phi_j = \frac{\partial}{\partial x_j} h(x)$ holds.*

Therefore, for any pair of locations (ℓ, d) where the objective function is differentiable with respect to $f_{(\ell,d)}$, the stationarity conditions require the following equality hold:

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta_\ell - \alpha_{(\ell,d)} - \gamma_{(\ell,d)}. \tag{73}$$

And, for any pair of locations (ℓ, d) where the objective function is differentiable with respect to $g_{(\ell,d)}$, the stationarity conditions require the following equality hold:

$$\frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \gamma_{(\ell,d)} - \beta_{(\ell,d)}. \tag{74}$$

G.4 Statement of Lemma 3

For the rest of this document we will focus on properties of the state-dependent optimization function (60). Recall the function $\Phi_{\omega_t}(\mathbf{S})$ gives the optimal value of the state-dependent optimization problem with respect to the time-scenario ω_t as a function of a supply-location vector $\mathbf{S}$. In this section we will show that the partial derivative $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S})$ exists for every location ℓ with a nonzero volume of drivers under $\mathbf{S}$.

First, let us introduce notation to refer to optimal primal and dual solutions of the state-dependent optimization problem. Let

$$F^*(\mathbf{S}) = \{(\boldsymbol{f}, \boldsymbol{g}) : \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \Phi_{\omega_t}(\mathbf{S}),\ (\boldsymbol{f}, \boldsymbol{g}) \text{ is feasible for (60) with respect to } \mathbf{S}\}$$

denote the set of primal optimal solutions as a function of the supply-location vector $\mathbf{S}$. Let

$$D^*(\mathbf{S}) = \{(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}) : \exists (\boldsymbol{f}^*, \boldsymbol{g}^*) \in F^*(\mathbf{S}) \text{ such that } 0 \in \partial L(\boldsymbol{f}^*, \boldsymbol{g}^*; \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}),\ \boldsymbol{\alpha} \geq 0, \boldsymbol{\beta} \geq 0, \boldsymbol{\gamma} \geq 0\}$$

denote the set of dual optimal solutions as a function of the supply-location vector $\mathbf{S}$.

Our main result in this section is the following Lemma, which characterizes important properties about partial derivatives of the state-dependent optimization function. Below is a restatement of Lemma 3.

Lemma. *Fix a time-scenario ω_t and let $\mathbf{S} = (S_\ell \geq 0 : \ell \in \mathcal{L})$ be any supply-location vector. Pick any location ℓ for which the volume of supply at ℓ is nonzero under $\mathbf{S}$, i.e. $S_\ell > 0$.*

1. *For the state-dependent optimization problem (60) with respect to $\mathbf{S}$ the value of any optimal dual variable associated with the flow conservation constraint (64) for location ℓ is unique. That is there exists a number η^*_ℓ such that $\eta_\ell = \eta^*_\ell$, where η_ℓ is the ℓth component of $\boldsymbol{\eta}$ for any optimal dual variables $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}) \in D^*(\mathbf{S})$.*
2. *The state dependent optimization function $\Phi_{\omega_t}(\cdot)$ is differentiable with respect to S_ℓ at the supply location vector $\mathbf{S}$. Moreover, the partial derivative is equal to the value of the optimal dual variable for the flow conservation constraint at location ℓ:*

$$\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}) = \eta^*_\ell.$$

3. *The partial derivative $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S})$ is continuous at $\mathbf{S}$.*

Below, in Section G.5 we prove Lemma 3.

G.5 Proof of Lemma 3

We prove Lemma 3 by backwards induction on the time t. For the rest of this section we hold fixed a time-scenario ω_t, a supply-location vector $\mathbf{S} = (S_\ell : \ell \in \mathcal{L})$, and we fix a location $\ell \in \mathcal{L}$ for which the volume of supply at ℓ under $\mathbf{S}$ is nonzero, i.e. $S_\ell > 0$.

Our backwards induction hypothesis states that the conclusion of Lemma 3 holds for all supply-location vectors at all time $t + 1$ scenarios. For clarity we formally state our backwards induction hypothesis in Assumption 3.

Assumption 3. *When $t = T$ is the final time period then we make no assumption. When $t < T$, let ω_{t+1} be any time $t + 1$ scenario, let $\mathbf{S}'$ be any supply location vector, and let ℓ' be any location for which the volume of supply at ℓ' under $\mathbf{S}'$ is nonzero. Then Lemma 3 parts 1, 2, and 3 are true, with respect to ω_{t+1}, $\mathbf{S}'$ and ℓ'.*

In the following subsections we prove parts 1, 2, and 3 of Lemma 3 assuming the backwards induction hypothesis.

G.5.1 Proof of Part 1

Lemma 3 Part 1 claims that the optimal dual variable associated with the flow-conservation constraint for location ℓ is unique. We prove this claim by invoking Lemma 21 which states that the Lagrangian optimality conditions for the state-dependent optimization problem hold between any pair of primal and dual optima.

Specifically, let $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}), (\boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \boldsymbol{\eta}') \in D^*(\mathbf{S})$ be any pair of dual optima and let $(\boldsymbol{f}, \boldsymbol{g}) \in F^*(\mathbf{S})$ be any primal optimum. Lemma 21 states that the stationarity conditions and complementary slackness conditions hold between the primal optimum $(\boldsymbol{f}, \boldsymbol{g})$ and both dual optima $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}), (\boldsymbol{\alpha}', \boldsymbol{\beta}', \boldsymbol{\gamma}', \boldsymbol{\eta}')$.

We first consider the stationarity optimality conditions. From the assumption that location ℓ has nonzero supply-volume under $\mathbf{S}$, there must be a destination $d \in \mathcal{L}$ for which a nonzero volume of drivers traverse from ℓ to d under any feasible solution. In particular, consider a location d for which the $f_{(\ell,d)}$ component of the optimal solution $(\boldsymbol{f}, \boldsymbol{g})$ is nonzero.

Observe that under the backwards induction hypothesis, the objective function $\mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ is differentiable with respect to $f_{(\ell,d)}$ at the primal optimum $(\boldsymbol{f}, \boldsymbol{g})$. Recall the objective function $\mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ is the sum of the current reward $\mathcal{U}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ and the future reward $\mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right]$. That the current reward $\mathcal{U}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$ is differentiable with respect to any nonzero component of $\boldsymbol{f}$ follows from Lemma 1. That the future reward is differentiable with respect to $f_{(\ell,d)}$ follows from the backward induction hypothesis. Specifically, for any time $t+1$ scenario ω_{t+1}, there will be nonzero supply-volume at location d under the resulting time $t + 1$ supply-location vector, since there is a nonzero volume of drivers driving from ℓ to d. Therefore the state-dependent optimization function $\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))$ is differentiable with respect to the volume of supply at location d. Therefore, it follows from the chain rule that the partial derivative of the future reward exists and can be written as follows:

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathbb{E}_{\omega_t}\left[\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right] = \mathbb{E}_{\omega_t}\left[\frac{\partial}{\partial S_d} \Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))\right].$$

Having established differentiability of the objective function with respect to the $f_{(\ell,d)}$ variable, let's return to the stationarity optimality conditions. It follows from equation (73) that the stationarity conditions imply the following equality:

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta_\ell - \alpha_{(\ell,d)} - \gamma_{(\ell,d)}.$$

The above equation gives us a useful characterization of the dual variable η_ℓ, whenever the supply-volume S_ℓ is greater than 0:

$$\begin{aligned}
\eta_\ell &= \frac{1}{S_\ell} \sum_{d \in \mathcal{L}} f_{(\ell,d)} \left[\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + \alpha_{(\ell,d)} + \gamma_{(\ell,d)} \right] \\
&= \frac{1}{S_\ell} \sum_{d \in \mathcal{L}} \left[f_{(\ell,d)} \frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + g_{(\ell,d)} \gamma_{(\ell,d)} \right] \\
&= \frac{1}{S_\ell} \sum_{d \in \mathcal{L}} \left[f_{(\ell,d)} \frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + g_{(\ell,d)} \frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) \right] \qquad (75)
\end{aligned}$$

For the remainder of this proof we consider two cases: in one case a nonzero volume of drivers traversing (ℓ, d) have no passenger, i.e. $g_{(\ell,d)} < f_{(\ell,d)}$; in the other case, we have all drivers traversing (ℓ, d) are carrying a passenger, i.e. $g_{(\ell,d)} = f_{(\ell,d)}$. In the first case, since the constraint $g_{(\ell,d)} \leq f_{(\ell,d)}$ is strict it follows from the complementary slackness conditions that dual variable associated with the constraint, $\gamma_{(\ell,d)}$, is 0. Similarly, since $f_{(\ell,d)}$ is nonzero, the dual variable associated with the nonnegativity constraint on $f_{(\ell,d)}$, that is $\alpha_{(\ell,d)}$, is 0. Therefore, the stationarity condition simplifies to the following:

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta_\ell.$$

Since Lemma 21 states the optimality conditions hold between any pair of primal and dual optima, we can apply the same line of reasoning to our other dual solution and conclude

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta'_\ell,$$

from which $\eta_\ell = \eta'_\ell$ follows.

In the second case where all drivers along (ℓ, d) have a passenger the dual variable $\gamma_{(\ell,d)}$ need not be 0, but equation (74) gives us the following characterization:

$$\frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \gamma_{(\ell,d)} - \beta_{(\ell,d)}.$$

Note that the objective function only depends on $\boldsymbol{g}$ for the current reward $\mathcal{U}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g})$, and this is differentiable with respect to any nonzero component $g_{(\ell,d)}$. Further, $\beta_{(\ell,d)}$ is the dual variable associated with the nonnegativity constraint on $g_{(\ell,d)}$, and from the assumption that $g_{(\ell,d)} = f_{(\ell,d)}$ and $f_{(\ell,d)} > 0$ the complementary slackness conditions imply that $\beta_{(\ell,d)}$ is 0. Therefore we obtain the equality

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + \frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta_\ell$$

and

$$\frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + \frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) = \eta'_\ell$$

from which $\eta'_\ell = \eta_\ell$ follows.

G.5.2 Proof of Part 2

We give a high-level outline of the proof for part 2. We start by using Lemma 22, which considers the *value function* associated with an optimization problem, which gives the optimal value of an optimization problem as a function of the constraint vector. Lemma 22 shows that the set of optimal dual variables for the optimization problem at a particular constraint vector is the same as the set of negative subgradients for the value function at that constraint vector.

The negative state-dependent optimization function $-\Phi_{\omega_t}(\mathbf{S})$ is similar to the value function considered by Lemma 22, except the state-dependent optimization problem has a mix of equality constraints and inequality constraints whereas the optimization problem considered in Lemma 22 only explicitly includes inequality constraints, and the supply-location vector $\mathbf{S}$ that $-\Phi_{\omega_t}(\mathbf{S})$ takes as an argument only varies the bounds for the equality constraints.

To use the result of Lemma 22 in the context of our state-dependent optimization function we first rewrite the state-depent optimization problem solely in terms of inequality constraints, where each equality constraint is replaced by two inequality constraints pointing in opposite directions. When the state-dependent optimization problem is written in this way it has the same structure as the optimization problem considered in Lemma 22, so we can associate a "value function" with the problem in the same manner as Lemma 22, and then the state-dependent optimization problem $-\Phi_{\omega_t}(\mathbf{S})$ is equivalent to this value function applied to a linear transformation of the supply-location vector $\mathbf{S}$.

Finally we invoke Lemma 23 which gives a version of the chain-rule that applies to subgradients. Invoking Lemma 23 tells us that a vector $\boldsymbol{\phi}$ is a subgradient of $(-\Phi_{\omega_t})(\mathbf{S})$ if and only iff $\boldsymbol{\phi} = -\boldsymbol{\eta}$, where $\boldsymbol{\eta}$ is the restriction of any optimal dual variable $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}) \in D^*(\mathbf{S})$ to the components associated with the flow-conservation equality constraints.

The conclusion of part 2 follows from the results of part 1, which states that there is a unique optimal dual variable for the flow-conservation constraint associated with location ℓ. It follows that there is a unique value for the ℓth component of any subgradient for the negative state-dependent optimization function evaluated at $\mathbf{S}$. We know that a function is differentiable at a point when the subderivative of that function at that point is unique . Therefore, the partial derivative of the negative state-dependent value function with respect to the ℓth component of the input vector exists and is equal to the negative dual variable associated with the ℓth flow-conservation constraint. Taking negatives on both sides of the equality, we conclude $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S})$ exists and is equal to η_ℓ^*, as claimed.

G.5.3 Proof of Part 3

The final result left to establish for Lemma 3 is that the state-dependent optimization function has continuous partial derivatives at any location where the supply-location vector is nonzero. We prove this result by showing that, for any sequence of supply-location vectors converging to $\mathbf{S}$, the corresponding sequence of partial derivatives with respect to location ℓ converges to the partial derivative evaluated at $\mathbf{S}$.

Formally, let $(\mathbf{S}_k : k = 1, 2, \dots)$ be a sequence of supply-location vectors converging to $\mathbf{S}$, and assume without loss of generality that the ℓth component of each iterate $\mathbf{S}_k$ is nonzero. Having already established parts 1 and 2 of Lemma 3, we know the following:

- For the state-dependent optimization problem with respect to each supply-location vector $\mathbf{S}_k$ there is a unique optimal dual variable associated with the flow-conservation constraint for location ℓ.
- The state-dependent optimization function evaluated at $\mathbf{S}_k$ is partially differentiable in the direction ℓ, and the value of the partial derivative is equal to the optimal dual variable for the location ℓ flow-conservation constraint.

Let $\eta_\ell^*(\mathbf{S}_k)$ denote the optimal dual variable for the location ℓ flow-conservation constraint with respect to $\mathbf{S}_k$ and let $\eta_\ell^*(\mathbf{S})$ denote the same optimal dual variable with respect to $\mathbf{S}$. We will show that the partial derivative of the state-dependent optimization function is continuous at $\mathbf{S}$ by showing that the sequence of optimal dual variables $(\eta_\ell^*(\mathbf{S}_k) : k = 1, 2, \dots)$ converges to $\eta_\ell^*(\mathbf{S})$, i.e.

$$\lim_{k\to\infty} \eta_\ell^*(\mathbf{S}_k) = \eta_\ell^*(\mathbf{S}). \tag{76}$$

Our approach for establishing the equality in equation (76) is to use the Lagrangian optimality conditions to obtain an equivalent expression in terms of primal solutions. To obtain this equivalent expression that works in the space of primal solutions, we construct a function that takes as input a primal optimal solution and produces the value of the optimal dual variable for location ℓ as the output. Define $\mathcal{S} = \{\mathbf{S}_k : k = 1, 2, \dots\} \cup \{\mathbf{S}\}$ to be the set of all supply-location vectors in our sequence and the limiting supply-location vector to which they converge, and define

$$\mathcal{F} = \bigcup_{\mathbf{S}'\in\mathcal{S}} F^*(\mathbf{S}')$$

to be the set of all primal solutions that are optimal for some supply-location vector in $\mathcal{S}$.

We write $E_\ell : \mathcal{F} \to \mathbb{R}$ to denote our function that recovers the optimal dual variable associated with location ℓ from a primal optimal solution. For a primal optimal solution $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in \mathcal{F}$, the exact definition of $E_\ell(\boldsymbol{f}^*, \boldsymbol{g}^*)$ will reflect the optimality conditions associated with a particular route (ℓ, d). The choice of the destination location d will depend on which components of $\boldsymbol{f}^*$ are nonzero. Specifically, order the locations in $\mathcal{L}$ as $d_1, d_2, \dots, d_n$ where $n = |\mathcal{L}|$, and let $i(\boldsymbol{f}^*) = i$ be the smallest index in $\{1, 2, \dots, n\}$ such that $f^*_{(\ell,d_i)}$ is nonzero. Note that by construction every supply-location vector in $\mathcal{S}$ has nonzero volume on location ℓ, so every primal optimal solution $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in \mathcal{F}$ always has at least one destination d for which $f^*_{(\ell,d)}$ is nonzero; in particular, the index $i(\boldsymbol{f}^*)$ is always well-defined.

Now, consider any sequence of optimal solutions $(\boldsymbol{f}_k^*, \boldsymbol{g}_k^*) \in F^*(\mathbf{S}_k)$ for $k \geq 1$, and observe the sequence $(\boldsymbol{f}_k^*, \boldsymbol{g}_k^*)$ is bounded, in particular there is a convergent subsequence $((\boldsymbol{f}_k^*(i), \boldsymbol{g}_k^*(i)) : i \geq 1)$. Let $\boldsymbol{f}^*, \boldsymbol{g}^*$ be the limit point of this subsequence, and observe that $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in F^*(\mathbf{S})$. Define the function $E_\ell(\boldsymbol{f}, \boldsymbol{g})$ as

$$E_\ell(\boldsymbol{f}, \boldsymbol{g}) = \frac{1}{S_\ell} \sum_{d\in\mathcal{L}} \left[f_{(\ell,d)} \frac{\partial}{\partial f_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) + g_{(\ell,d)} \frac{\partial}{\partial g_{(\ell,d)}} \mathcal{W}_{\omega_t}(\boldsymbol{f}, \boldsymbol{g}) \right]$$

From the equation (75), we know that $\eta_\ell^*(\mathbf{S}')$ is equal to $E_\ell(\boldsymbol{f}^*, \boldsymbol{g}^*)$, if $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in F^*(\mathbf{S}')$ for any $\mathbf{S}' \in \mathcal{F}$.

Observe that for i large enough, the convergent subsequence $((\boldsymbol{f}^*_{k(i)}, \boldsymbol{g}^*_{k(i)}) : i \geq 1)$ will be nonzero on the same components as $(\boldsymbol{f}^*, \boldsymbol{g}^*)$. Therefore

$$\lim_{i\to\infty} \eta_\ell^*(\mathbf{S}_{k(i)}) = \lim_{i\to\infty} E_\ell(\boldsymbol{f}^*_{k(i)}, \boldsymbol{g}^*_{k(i)}) = E_\ell(\boldsymbol{f}^*, \boldsymbol{g}^*) = \eta_\ell^*(\mathbf{S}).$$

Also observe that the sequence of dual variables $(\eta_\ell^*(\mathbf{S}_k) : k \geq 1)$ is the same sequence as $(E_\ell(\boldsymbol{f}_k^*, \boldsymbol{g}_k^*) : k \geq 1)$. From the above equation it follows that every limit point is equal to $\eta_\ell^*(\mathbf{S})$. Since the sequence $(\eta_\ell^*(\mathbf{S}_k) : k \geq 1)$ is bounded and since there is a single limit point, it follows the sequence converges: $\lim_{k\to\infty} \eta_\ell^*(\mathbf{S}_k) = \eta_\ell^*(\mathbf{S})$, finishing the proof.

G.5.4 Additional Lemma for Proof of Lemma 3

Lemma 21. *For any convex optimization problem the stationarity and complementary slackness conditions hold between any pair of primal and dual optima. In particular, for the state-dependent* *optimization problem (60) with respect to any scenario* ω_t *and supply location vector* $\mathbf{S}$*, if* $(\boldsymbol{f}, \boldsymbol{g}) \in F^*(\mathbf{S})$ *is any primal optimum and* $(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}) \in D^*(\mathbf{S})$ *is any dual optimum then the stationarity conditions hold, i.e.*

$$0 \in \partial L(\boldsymbol{f}, \boldsymbol{g}; \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\eta}),$$

and the complementary slackness conditions hold, i.e.

$$\alpha_{(\ell,d)} f_{(\ell,d)} = 0, \ \beta_{(\ell,d)} g_{(\ell,d)} = 0, \ \gamma_{(\ell,d)}(g_{(\ell,d)} - f_{(\ell,d)}) = 0.$$

For the following lemma, consider the optimization problem

$$\inf_{x \in \mathbb{R}^m} \left\{ f(x) \mid g(x) \leq 0 \right\}, \tag{77}$$

where $f : \mathbb{R}^m \to \mathbb{R}$ is our objective function and $g : \mathbb{R}^m \to \mathbb{R}^n$ is our constraint function. We assume that f and $g_1, \ldots, g_n$ are convex functions, where $g_j(x)$ is the jth component function of the multivariate constraint function g. The Lagrangian function $L : \mathbb{R}^m \times \mathbb{R}^n_+ \to \mathbb{R}$ is defined by

$$L(x, \lambda) = f(x) + \lambda^T g(x).$$

The dual function $\Gamma : \mathbb{R}^n_+ \to \mathbb{R}$ is defined by

$$\Gamma(\lambda) = \inf_{x \in \mathbb{R}^m} L(x; \lambda).$$

The *value function* associated with the mathematical program (77) describes how the optimal value changes as we perturb the constraint vector away from 0. Formally, it is a function $v : \mathbb{R}^n \to \mathbb{R}$ defined by the equation

$$v(b) = \inf_{x \in \mathbb{R}^m} \left\{ f(x) \mid g(x) \leq b \right\}. \tag{78}$$

The problem (77) is said to have zero duality gap when strong duality holds, i.e. when the primal optimum is equal to the dual optimum, as described by the following equation:

$$\inf_{x \in \mathbb{R}^m} \left\{ f(x) \mid g(x) \leq 0 \right\} = \sup_{\lambda \in \mathbb{R}^m_+} \Gamma(\lambda).$$

Any $\lambda^* \in \mathbb{R}^*_m$ which achieves the optimum on the right side of the above equation is said to be an optimal dual solution. The following lemma appears as Corollary 4.3.6 in **?**.

Lemma 22. *The mathematical program (77) has zero duality gap if and only if the value function* v *is lower semicontinuous at* 0*. In this case the set of dual optimal solutions is* $-\partial v(0)$*.*

In order to apply the result of Lemma 22 we also make use of the following result from **?**, which provides a chain rule for subdifferentials of convex functions composed with linear functions.

Lemma 23. *Let* $f : \mathbb{R}^m \to \mathbb{R}$ *be a convex function and let* $A \in \mathbb{R}^{m \times n}$ *be a matrix. Then the following equality is satisfied for* $x \in \mathbb{R}^n$*:*

$$\partial(f \circ A)(x) = A^T \partial f(Ax).$$

G.6 Proof of Lemma 4

We re-state Lemma 4 below.

Lemma. *Let* $\mathbf{S}$ *be a supply-location vector with nonnegative components and assume* $S_\ell = 0$ *for some location* ℓ*. Then the right-derivative* $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}^+)$ *is well-defined at* $\mathbf{S}$*. Moreover, the partial derivative function* $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S})$*, defined in (15), is continuous over the set* $\{\mathbf{S} \in \mathbb{R}^{\mathcal{L}} : S_\ell \geq 0 \forall \ell \in \mathcal{L}\}$*.*

Also, in the case where $S_\ell = 0$*, there exists an optimal dual solution such that the dual variable* η_ℓ *associated with the* ℓ*th flow-conservation constraint is equal to the right derivative* $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}^+)$*.*

We prove Lemma 4 by characterizing optimal primal solutions to the state-dependent optimization problem, in the regime where there is an infinitesimal volume of drivers at ℓ.

We start by defining the continuation utilities associated with a primal solution. Let ω_t be any scenario, and let $\boldsymbol{f} = (f_{(\ell,d)} \geq 0 : (\ell, d) \in \mathcal{L}^2)$ be any flow vector. For a destination $d \in \mathcal{L}$ define the continuation utility associated with d and $\boldsymbol{f}$ to be

$$U_d(\boldsymbol{f}) = \mathbb{E}\left[\frac{\partial}{\partial S_d}\Phi_{\omega_{t+1}}(\bar{\mathbf{S}}_{\omega_{t+1}}(\boldsymbol{f})) \mid \omega_t\right]. \tag{79}$$

The following lemma states that the continuation utilities associated with optimal solutions to the state-dependent optimization problem all take the same value.

Lemma 24. *Let* ω_t *be any scenario and let* $\mathbf{S}$ *be any feasible supply-location vector. Let* $(\boldsymbol{f}_i, \boldsymbol{g}_i) \in F^*_{\omega_t}(\mathbf{S})$*, for* $i = 1, 2$*, be any optimal solutions to the state-dependent optimization problem with respect to* $(\omega_t, \mathbf{S}_t)$*. Then the continuation utilities under* $\boldsymbol{f}_1$ *and* $\boldsymbol{f}_2$ *are the same, i.e.* $U_d(\boldsymbol{f}_1) = U_d(\boldsymbol{f}_2)$ *for any choice of destination* d*.*

Lemma 24 follows from the optimality conditions, and the fact that complementary slackness holds between any pair of primal and dual optima.

Our next Lemma characterizes optimal solutions for the state-dependent optimization problem in the regime where there is an infinitesimal volume of drivers at ℓ.

Lemma 25. *Let* ω_t *be any scenario and let* $\mathbf{S}$ *be any feasible supply-location vector. Assume* $S_\ell = 0$ *for some location* ℓ*. Let* $(\mathbf{S}_n)_{n=1}^\infty$ *be a sequence of feasible supply-location vectors converging to* $\mathbf{S}$*, such that each element of the sequence has a nonzero volume of drivers at* ℓ*, i.e.* $S^n_\ell > 0$ *for all* n*. Let* $(\boldsymbol{f}^*_n, \boldsymbol{g}^*_n) \in F^*_{\omega_t}(\mathbf{S}_n)$ *be an optimal primal solution for each* n*, and define*

$$\boldsymbol{f}^n_\ell = \frac{1}{S^n_\ell}\left(f^*_{n,(\ell,d)} : d \in \mathcal{L}\right) \quad \boldsymbol{g}^n_\ell = \frac{1}{S^n_\ell}\left(g^*_{n,(\ell,d)} : d \in \mathcal{L}\right)$$

to be the restriction of $(\boldsymbol{f}^*_n, \boldsymbol{g}^*_n)$ *to components that correspond to trips originating from* ℓ*, divided by the volume of drivers at* ℓ *under the* n*th iterate in the sequence. Then every limit point of the sequence* $(\boldsymbol{f}_{\ell,n}, \boldsymbol{g}_{\ell,n})_{n=1}^\infty$ *is an optimal solution to the following optimization problem:*

$$\sup \quad \sum_{d\in\mathcal{L}} V_d g_d + \sum_{d\in\mathcal{L}}(U_d - c_{(\ell,d)})f_d - A\left(\sum_{d\in\mathcal{L}} g_d, \sum_{d\in\mathcal{L}} f_d\right) \tag{80}$$

$$\textit{such that} \quad 0 \leq g_d \leq f_d \ \ \forall d \in \mathcal{L}, \tag{81}$$

$$\sum_{d\in\mathcal{L}} f_d = 1. \tag{82}$$

In the above, V_d is to the maximum rider value held by riders requesting a trip from ℓ to d under ω_t, U_d is the continuation utility associated with each destination $d \in \mathcal{L}$ under an optimal solution for the limiting supply-location vector $\mathbf{S}$, and $A(\cdot,\cdot)$ is the add-passenger disutility cost function $A(g,f) = \frac{C}{2}\frac{g^2}{f}$.

Moreover, the optimization problem (80) has a unique optimal dual variable η_ℓ associated with the constraint (82), and the value of this dual variable is equal to the right-derivative limit $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}^+)$.

Proof. We give a high-level outline for the proof of Lemma 25. First, a backwards induction argument lets us assume that the objective function of the original state-dependent optimization problem $\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g})$ has continuous right derivatives on the boundary of the feasible region (the backwards induction assumption applies to the future-period reward function $\Phi_{\omega_{t+1}}(\mathbf{S}_{\omega_{t+1}}(\boldsymbol{f}))$ which appears as a summand in the objective $\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g})$).

Next, by considering convergent subsequences, we can assume without loss of generality that the sequence of optimal solutions $(\boldsymbol{f}_n^*, \boldsymbol{g}_n^*)$, $n = 1, 2, \ldots$, converges to some limit $(\boldsymbol{f}^*, \boldsymbol{g}^*)$, and by continuity of the objective function it follows that the limit point is optimal with respect to the limiting supply-location vector $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in F^*_{\omega_t}(\mathbf{S})$.

Also by considering convergent subsequences, we can assume without loss of generality that the scaled sequence of points

$$\boldsymbol{f}_\ell^n = \frac{1}{S_\ell^n}\left(f^*_{n,(\ell,d)} : d \in \mathcal{L}\right) \quad \boldsymbol{g}_\ell^n = \frac{1}{S_\ell^n}\left(g^*_{n,(\ell,d)} : d \in \mathcal{L}\right)$$

converges to some limit $(\bar{\boldsymbol{f}}_\ell, \bar{\boldsymbol{g}}_\ell)$, and by virtue of the scaling it follows that this limit is a feasible solution for the linearized optimization problem (80).

Next, we consider the optimality conditions associated with each iterate in our sequence of supply-location vectors. Let η_ℓ^n be the dual variable associated with the ℓth flow-conservation constraint, for the nth supply-location vector in our sequence. Since every iterate in our sequence has nonzero volume of drivers at ℓ, the dual variable η_ℓ^n is unique. Further, the characterization (75) of this dual variable yields the following expression:

$$\eta_\ell^n = \frac{1}{S_\ell^n}\sum_{d\in\mathcal{L}}\left[f^n_{(\ell,d)}\frac{\partial}{\partial f_{(\ell,d)}}\mathcal{W}_{\omega_t}(\boldsymbol{f}_n^*, \boldsymbol{g}_n^*) + g^n_{(\ell,d)}\frac{\partial}{\partial g_{(\ell,d)}}\mathcal{W}_{\omega_t}(\boldsymbol{f}_n^*, \boldsymbol{g}_n^*)\right].$$

Notice that the partial derivatives of the objective function have the following expressions:

$$\frac{\partial}{\partial f_{(\ell,d)}}\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) = -c_{(\ell,d)} + \frac{\partial}{\partial f_{(\ell,d)}}\mathcal{U}^{>t}_{\omega_t}(\boldsymbol{f}) - \frac{\partial}{\partial f_{(\ell,d)}}A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell),$$

and

$$\frac{\partial}{\partial g_{(\ell,d)}}\mathcal{W}_{\omega_t}(\boldsymbol{f},\boldsymbol{g}) = \frac{d}{dg}U_{(\ell,d,\omega_t)}(g_{(\ell,d)}) - \frac{\partial}{\partial g_{(\ell,d)}}A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell).$$

Also, notice that the partial derivatives of the add-passenger disutility function have the following expressions:

$$\frac{\partial}{\partial f_{(\ell,d)}}A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell) = -\frac{C}{2}\left(\frac{\boldsymbol{g}^T\mathbf{1}_\ell}{\boldsymbol{f}^T\mathbf{1}_\ell}\right)^2 \quad \text{and} \quad \frac{\partial}{\partial g_{(\ell,d)}}A(\boldsymbol{g}^T\mathbf{1}_\ell, \boldsymbol{f}^T\mathbf{1}_\ell) = C\frac{\boldsymbol{g}^T\mathbf{1}_\ell}{\boldsymbol{f}^T\mathbf{1}_\ell}.$$

Therefore, the partial derivatives of $A(\cdot,\cdot)$ are invariant to both of its arguments being scaled by the same multiple. In particular, we have the equality

$$\frac{\partial}{\partial f_{(\ell,d)}}A(\boldsymbol{g}_n^{*T}\mathbf{1}_\ell, \boldsymbol{f}_n^{*T}\mathbf{1}_\ell) = \frac{\partial}{\partial f_{(\ell,d)}}A(\frac{\boldsymbol{g}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}, \frac{\boldsymbol{f}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}) \text{ and } \frac{\partial}{\partial g_{(\ell,d)}}A(\boldsymbol{g}_n^{*T}\mathbf{1}_\ell, \boldsymbol{f}_n^{*T}\mathbf{1}_\ell) = \frac{\partial}{\partial g_{(\ell,d)}}A(\frac{\boldsymbol{g}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}, \frac{\boldsymbol{f}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}).$$

We can rewrite our expression for the optimal dual variable η_ℓ^n as follows:

$$\begin{aligned}\eta_\ell^n = &\sum_{d\in\mathcal{L}} \frac{f_{(\ell,d)}^n}{S_\ell^n}\left(\frac{\partial}{\partial f_{(\ell,d)}}\mathcal{U}_{\omega_t}^{>t}(\boldsymbol{f}_n^*) - c_{(\ell,d)} - \frac{\partial}{\partial f_{(\ell,d)}}A(\frac{\boldsymbol{g}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}, \frac{\boldsymbol{f}_n^{*T}\mathbf{1}_\ell}{S_\ell^n})\right)\\ &+\sum_{d\in\mathcal{L}} \frac{f_{(\ell,d)}^n}{S_\ell^n}\left(\frac{d}{dg}U_{(\ell,d,\omega_t)}(g_{(\ell,d)}^n) - \frac{\partial}{\partial g_{(\ell,d)}}A(\frac{\boldsymbol{g}_n^{*T}\mathbf{1}_\ell}{S_\ell^n}, \frac{\boldsymbol{f}_n^{*T}\mathbf{1}_\ell}{S_\ell^n})\right)\end{aligned}$$

We know that $(\boldsymbol{f}_n^*, \boldsymbol{g}_n^*)$ converges as $n \to \infty$, as does $\frac{f_{(\ell,d)}^n}{S_\ell^n}$ and $\frac{g_{(\ell,d)}^n}{S_\ell^n}$. Therefore the sequence of optimal dual variables converges to the following limit.

$$\begin{aligned}\lim_{n\to\infty}\eta_\ell^n = &\sum_{d\in\mathcal{L}} \bar{f}_{(\ell,d)}\left(\frac{\partial}{\partial f_{(\ell,d)}}\mathcal{U}_{\omega_t}^{>t}(\boldsymbol{f}^*) - c_{(\ell,d)} - \frac{\partial}{\partial f_{(\ell,d)}}A(\bar{\boldsymbol{g}}^T\mathbf{1}_\ell, \bar{\boldsymbol{f}}^T\mathbf{1}_\ell)\right)\\ &+\sum_{d\in\mathcal{L}} \bar{g}_{(\ell,d)}\left(\frac{d}{dg}U_{(\ell,d,\omega_t)}(g_{(\ell,d)}^*) - \frac{\partial}{\partial g_{(\ell,d)}}A(\bar{\boldsymbol{g}}^T\mathbf{1}_\ell, \bar{\boldsymbol{f}}^T\mathbf{1}_\ell)\right).\end{aligned}$$

However, since $\boldsymbol{g}^*$ is feasible for the limiting supply-location vector $\mathbf{S}$, we know

$$\frac{d}{dg}U_{(\ell,d,\omega_t)}(g_{(\ell,d)}^*) = \frac{d}{dg}U_{(\ell,d,\omega_t)}(0) = V_d,$$

where V_d is the maximum rider value for riders requesting from ℓ to d under ω_t. Also, since $(\boldsymbol{f}^*, \boldsymbol{g}^*)$ is an optimal solution with respect to $\mathbf{S}$, we have the partial derivative $\frac{\partial}{\partial f_{(\ell,d)}}\mathcal{U}_{\omega_t}^{>t}(\boldsymbol{f}^*)$ is equal to the optimal continuation utility U_d associated with d. Therefore, the limit of the sequence of dual variables, $\eta_\ell = \lim_{n\to\infty}\eta_\ell^n$, is equal to

$$\begin{aligned}\eta_\ell = &\sum_{d\in\mathcal{L}} \bar{f}_{(\ell,d)}\left(U_d - c_{(\ell,d)} - \frac{\partial}{\partial f_{(\ell,d)}}A(\bar{\boldsymbol{g}}^T\mathbf{1}_\ell, \bar{\boldsymbol{f}}^T\mathbf{1}_\ell)\right)\\ &+\sum_{d\in\mathcal{L}} \bar{g}_{(\ell,d)}\left(V_d - \frac{\partial}{\partial g_{(\ell,d)}}A(\bar{\boldsymbol{g}}^T\mathbf{1}_\ell, \bar{\boldsymbol{f}}^T\mathbf{1}_\ell)\right).\end{aligned}$$

Optimality of $(\bar{\boldsymbol{f}}, \bar{\boldsymbol{g}})$ and η_ℓ follow by using the above characterization of η_ℓ to show that $(\bar{\boldsymbol{f}}, \bar{\boldsymbol{g}})$ satisfies optimality conditions for the optimization problem (80). Finally, uniqueness of the dual variable η_ℓ follows from the same argument we used in Part one (appendix G.5.1) of the proof for Lemma 3. ☐

Lemma 25 shows that the right-hand derivative limit $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}^+)$ is well-defined when $S_\ell = 0$, and that every sequence of partial derivatives $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}_n)$ with $S_\ell^n > 0$ converges to the same limit. To finish showing that the partial derivative function To finish proving Lemma 4 it suffices to show that sequences of right-hand derivatives $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}_n^+)$ converge, for supply-location vectors $\mathbf{S}_n^+$ on the boundary of the feasible space, i.e. with $S_\ell^n = 0$.

Lemma 26 follows from analyzing the optimization problem (80) using the same logic as part three of our proof of Lemma 3.

Lemma 26. *Let $(\mathbf{S}_n)_{n=1}^\infty$ be a sequence of supply location vectors which converge to $\mathbf{S}$, all of which have 0 driver volume at ℓ, i.e. $S_\ell^n = 0$ for all n. Then $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S}_n^+)$ converges to $\frac{\partial}{\partial S_\ell}\Phi_{\omega_t}(\mathbf{S})$ as $n \to \infty$.*

H Matching Process Details

We view the matching process as a generic procedure for allocating trips to available drivers. In general, we assume there is a stochastic matching process for the two-level model as well as a deterministic matching process for the fluid model. Our analysis holds for any matching process that satisfies three properties, stated informally below:

Assumption 4.

1. *The random trip-volumes produced by the stochastic two-level model matching process converge to their corresponding deterministic fluid trip-volumes as the population size parameter grows to infinity.*
2. *In the deterministic fluid matching process, the only way for the trip volume produced by the matching process along a route to be smaller than the optimal trip volume along that route is if the drivers are using an acceptance threshold smaller than the optimal acceptance threshold.*
3. *In the stochastic two-level model matching process, conditioning on the action taken by a single driver has negligible effect on the overall distribution of aggregate trip counts in the limit as the population size grows to infinity. Specifically, we assume there exists a sequence* $(\beta_k)_{k=1}^{\infty}$ *converging to 0 as* $k \to \infty$ *such that the conditional distribution* $\mathbb{P}(\mathbf{S}_{t+1} \mid a_i^t)$ *is at most* β_k *different from the unconditional distribution* $\mathbb{P}(\mathbf{S}_{t+1})$*, i.e.* $|\mathbb{P}(\mathbf{S}_{t+1} \mid a_i^t) - \mathbb{P}(\mathbf{S}_{t+1})| \leq \beta_k$*. We assume the sequence* $(\beta_k)_{k=1}^{\infty}$ *works for all initial states* $(\omega_t, \mathbf{S}_t)$ *and all driver strategy profiles.*

For completeness, we define one example of a matching process that the platform can use, and show that it satisfies properties 1 and 2 listed above. We conjecture that this process also satisfies 3, but have not yet verified this.

H.1 Example Matching Process Definition

The SSP matching process definition differs slightly between the fluid model and the two-level model, because granular rider and driver decisions which affect the dynamics of the matching process are stochastic in the two level model but deterministic in the fluid model.

In both cases the SSP matching process makes use of a subroutine which takes a collection of drivers and a collection of riders all heading towards the same destination, and allocates dispatches towards that destination until either no drivers or riders remain.

Definition 4. *The single destination dispatch subroutine in the fluid model is a procedure that takes as input a rider volume* $\bar{R}$*, driver volume* $\bar{M}$*, and a single disutility threshold* x*. The output is a number* $\bar{G}(\bar{R}, \bar{M}, x)$ *specifying the volume of dispatches that were accepted, and a number* $\bar{U}(\bar{R}, \bar{M}, x)$ *specifying the volume of drivers who were not allocated a dispatch in the process. The function definitions are stated below.*

$$\bar{G}(\bar{R}, \bar{M}, x) = \min\left(Z(\bar{R}, x), \bar{M}\right) \frac{x}{C}, \tag{83}$$

$$\bar{U}(\bar{R}, \bar{M}, x) = \bar{M} - \min\left(Z(\bar{R}, x), \bar{M}\right). \tag{84}$$

(Recall $Z(\bar{R}, x)$ *is the volume of drivers who need to be allocated a dispatch in order to see* $\bar{R}$ *accepted dispatches given a disutility threshold* x*).*

In the two level model it is a procedure that takes as input a number of riders R*, a number of drivers* M*, and a choice of disutility thresholds* $x_1, \ldots, x_M$ *for each driver. The procedure allocates dispatches to riders until all rides have been served or no drivers remain. The output is a number of drivers* G *who accepted a dispatch, and a set* $U \subseteq [M]$ *of driver labels who were not allocated a dispatch in the process. The stochastic dynamics governing* G *and* U *are stated in Algorithm 1.*

Algorithm 1: The Single Destination Dispatch Subroutine in the Two Level Model

```
1. Input: A number of dispatch requests $R$, a number of drivers $M$, a choice of disutility
   threshold $x_i$ for each $i = 1, 2, \ldots, M$.
2. Randomly permute the driver labels: select a permutation $\pi : [M] \to [M]$ uniformly at
   random and define new labels $j = \pi(i)$.
3. Initialize $G \leftarrow 0$, $U \leftarrow \{1, 2, \ldots, M\}$.
4. For $j = 1, 2, \ldots, M$:
     • Allocate a dispatch to driver $j$.
     • Sample the accept/reject decision $\delta_j \sim \mathrm{Ber}(x_j/C)$.
     • Record the decision: $R \leftarrow R - \delta_j$, $G \leftarrow G + \delta_j$.
     • Remove $j$ from $U$: $U \leftarrow U \setminus \{j\}$.
     • If $R = 0$: go to step 5.
5. Return $G$, $U$.
```

The SSP matching process, in both the fluid model and the two level model, uses the single destination dispatch subroutine in two separate stages. In the first stage, drivers are subdivided into groups, where there is one group for each destination, and group sizes are determined by the dispatch volumes and the disutility threshold from the optimal solution. The single destination subroutine is then used to allocate dispatches for each destination to drivers in the group associated with that destination. This is the first stage of the matching process. If any drivers remain undispatched after the first stage, the second stage goes through the dispatch destinations one by one and it uses the single destination dispatch subroutine to allocate all remaining demand for that destination to all remaining drivers.

Notice that the disutility threshold associated with the optimal solution $(\boldsymbol{f}^*, \boldsymbol{g}^*)$ is the same for every destination d:

$$x_\ell^* = C \frac{\boldsymbol{g}^{*T} \mathbf{1}_\ell}{\boldsymbol{f}^{*T} \mathbf{1}_\ell}. \tag{85}$$

Let

$$Z_{(\ell,d)} = Z(g^*_{(\ell,d)}, x^*_\ell) = g^*_{(\ell,d)} \frac{x^*_\ell}{C} \tag{86}$$

be the volume of drivers we need to allocate a dispatch toward d in order to see $g^*_{(\ell,d)}$ accepted dispatches under the threshold x^*_ℓ. Observe

$$\sum_{d \in \mathcal{L}} Z_{(\ell,d)} = \sum_{d \in \mathcal{L}} Z(g^*_{(\ell,d)}, x^*_\ell) = \sum_{d \in \mathcal{L}} g^*_{(\ell,d)} \frac{C}{x^*_\ell} = \boldsymbol{f}^{*T} \mathbf{1}_\ell = S_\ell,$$

so the fractions $Z_{(\ell,d)}/S_\ell$ sum to 1 over all d. These fractions are used to determine the partition sizes in the first stage of the matching process.

Definition 5. *The matching process in the fluid model takes as input a market state* $(\omega_t, \mathbf{S}_t)$*, a volume of requests* $\bar{R}_d$ *for each destination* d*, a location* ℓ*, a volume of drivers* $\bar{M}_\ell$*, and a disutility threshold vector* $\boldsymbol{x}_\ell = (x_{(\ell,d)} : d \in \mathcal{L})$*. It proceeds in two stages:*

Algorithm 2: The Matching Process in the Two Level Model

```
1. Input: A location $\ell$, a number of dispatch requests $R_{(\ell,d)}$ for each destination $d$, a number of drivers $M_\ell$, an add-passenger threshold vector $\boldsymbol{x}_i = (x^i_{(\ell,d)} : d \in \mathcal{L})$ for each driver $i = 1, 2, \ldots, M_\ell$, the fluid optimal actions $(\boldsymbol{f}^*, \boldsymbol{g}^*, \boldsymbol{x}^*)$.
2. Stage one:
    (a) Compute partition sizes $Z^*_{(\ell,d)}$ for each destination $d$, using equation (86) with the fluid optimal trip volume $g^*_{(\ell,d)}$ and threshold $x^*_\ell$.
    (b) Partition the $M_\ell$ drivers into groups of size $M_{(\ell,d)}$, where each $M_{(\ell,d)}$ is rounded up or down from $Z^*_{(\ell,d)}$.
    (c) Use the single-destination dispatch subroutine to allocate the $R_{(\ell,d)}$ dispatch requests to the $M_{(\ell,d)}$ drivers, for each destination $d$.
    (d) Record the output from the single-destination dispatch subroutine: Let $G^{(1)}_{(\ell,d)}$ be the number of dispatch trips accepted and $U^{(1)}_{(\ell,d)}$ the number of drivers who were not allocated a trip.
3. Stage two:
    (a) Let $R^{(2)}_{(\ell,d)} = R_{(\ell,d)} - G^{(1)}_{(\ell,d)}$ be the number of riders who have not been matched to a driver at the end of the first stage.
    (b) Let $M^{(2)}_\ell = \sum_d U^{(1)}_{(\ell,d)}$ be the number of drivers who were not allocated a dispatch at the end of the first stage.
    (c) Pick an ordering of the destinations $d_1, d_2, \ldots, d_{|\mathcal{L}|}$. For each destination $d = d_1, d_2, \ldots, d_{|\mathcal{L}|}$:
        • Use the single-destination dispatch subroutine to allocate the $R^{(2)}_{(\ell,d)}$ dispatch requests to the $M^{(2)}_\ell$ remaining drivers.
        • Record the output from the single-destination dispatch subroutine: let $G^{(2)}_{(\ell,d)}$ be the number of dispatch trips accepted, and $U^{(2)}_{(\ell,d)}$ be the number of drivers who remain undispatched.
        • Update the number of remaining drivers: set $M^{(2)}_\ell = U^{(2)}_{(\ell,d)}$.
```

1. *In the first stage, it partitions the* M_ℓ *drivers into groups of size* $\bar{M}_{(\ell,d)} = \frac{Z_{(\ell,d)}}{S_\ell}\bar{M}_\ell$ *for each* d*. It uses the single destination dispatch subroutine for each* d *to allocate the* $\bar{R}_d$ *dispatches to the* $\bar{M}_{(\ell,d)}$ *drivers. The first stage produces* $\bar{G}_d^{(1)} = \bar{G}(\bar{R}_d, \bar{M}_{(\ell,d)}, x_{(\ell,d)})$ *accepted dispatches towards each destination* d*, and* $\bar{U}_d = \bar{U}(\bar{R}_d, \bar{M}_{(\ell,d)}, x_{(\ell,d)})$ *drivers remain unallocated from each group* d*.*

2. *In the second stage the initial volume of drivers who were not allocated in the first stage is equal to* $\bar{U}^{(0)} = \sum_{d\in\mathcal{L}} \bar{U}_d$*. The matching process orders the locations* $d_1, \ldots, d_L$ *and it goes through the destinations and uses the single destination dispatch subroutine to allocate all remaining dispatches to the pool of unallocated drivers, until either all demand has been served or all drivers have been allocated. Specifically, for each destination* $i = 1, 2, \ldots, L$*, it runs the single destination dispatch subroutine on* $\bar{R}_d - \bar{G}_d^{(1)}$*,* $\bar{U}^{(i-1)}$*,* $x_{(\ell,d_i)}$*, and it records* $\bar{G}^{(2)} = \bar{G}(\bar{R}_d - \bar{G}_d^{(1)}, \bar{U}^{(i-1)}, x_{(\ell,d_i)})$ *accepted dispatches and* $\bar{U}^{(i)} = \bar{U}(\bar{R}_d - \bar{G}_d^{(1)}, \bar{U}^{(i-1)}, x_{(\ell,d_i)})$ *remaining unallocated drivers.*

The matching process in the two level model takes as input a market state $(\omega_t, \mathbf{S}_t)$*, a number of requests* R_d *for each destination* d*, a location* ℓ*, a number of drivers* M_ℓ*, and a disutility threshold vector* $\boldsymbol{x}_i = (x_d^i : d \in \mathcal{L})$ *for each driver* $i = 1, 2, \ldots, M_\ell$*. It proceeds in two stages:*

1. *In the first stage, it partitions the* M_ℓ *drivers into groups of size* $M_{(\ell,d)}$*, which are either rounded up or down from* $\frac{Z_{(\ell,d)}}{S_\ell}M_\ell$*, for each* d*. The allocation of drivers to groups happens uniformly at random. It uses the single destination dispatch subroutine for each* d *to allocate the* R_d *dispatches to the* $M_{(\ell,d)}$ *drivers. The first stage produces* $G_d^{(1)}$ *accepted dispatches towards each destination* d*, and* U_d *is the set of driver indices which remain unallocated from each group* d*.*

2. *In the second stage the initial volume of drivers who were not allocated in the first stage is equal to* $U^{(0)} = \cup_{d\in\mathcal{L}} U_d$*. The matching process orders the locations* $d_1, \ldots, d_L$ *and it goes through the destinations and uses the single destination dispatch subroutine to allocates all remaining dispatches to the pool of unallocated drivers, until either all demand has been served or all drivers have been allocated. Specifically, for each destination* $i = 1, 2, \ldots, L$*, it runs the single destination dispatch subroutine on* $R_d - G_d^{(1)}$*,* $U^{(i-1)}$*, and the thresholds* x_d^i *for* $i \in U^{(i-1)}$*. The output is* $G^{(2)}$ *accepted dispatches and* $U^{(i)}$ *is the index set of unallocated drivers.*

Lemma 27. *Let* $(\omega_t, \mathbf{S}_t)$ *be any market state, let* $(\boldsymbol{f}^*, \boldsymbol{g}^*) \in F^*_{\omega_t}(\mathbf{S}_t)$ *be a solution to the fluid optimization problem. Consider the fluid matching process which allocates all dispatch demand* $(g^*_{(\ell,d)} : d \in \mathcal{L})$ *originating from* ℓ *to all* S_ℓ *drivers positioned at* ℓ*. Let* $\boldsymbol{x} = (x_{(\ell,d)} : d \in \mathcal{L})$ *be any disutility threshold vector used by drivers at* ℓ *and let* $\boldsymbol{g}_\ell = (g_{(\ell,d)} : d \in \mathcal{L})$ *be the output of the matching process. Then* $g_{(\ell,d)} < g^*_{(\ell,d)}$ *implies* $x_{(\ell,d)} < x^*_\ell$*.*

H.2 Matching Process Concentration Properties

In this section of the appendix we establish that the matching process satisfies good concentration properties as the population size goes to infinity. For ease of use in the analysis of our main algorithm, we establish concentration inequalities that hold uniformly across all relevant market states. Informally, the relevant market states are those in which approximately every agent uses approximately the same add-passenger disutilities. In addition, we require that the total number of drivers is no larger than a multiple of the population-size parameter.

To describe the relevant market states to which our concentration inequalities apply, fix a time period t and scenario ω_t. Let $\mathcal{M}$ be an index set of all active drivers in the marketplace, and let $\mathcal{M}_\ell$ be the subset of drivers who are positioned at each location ℓ. For a driver $i \in \mathcal{M}_\ell$, we use $\boldsymbol{x}_i = (x^i_{(\ell,d)} : d \in \mathcal{L})$ to denote the add-passenger disutility threshold vector selected by driver i, and we use $r_i \in \mathcal{L}$ to denote the relocation destination selected by driver i, which is the destination towards which driver i will drive empty if they do not accept a dispatch trip. We use the term *driver-state* to mean the specification of add-passenger disutility threshold vector $\boldsymbol{x}_i$ and relocation destination r_i for each driver $i \in \mathcal{M}_\ell$ for each location ℓ. At a location ℓ, we will use $\boldsymbol{r}_\ell = (r_{(\ell,d)} : d \in \mathcal{L})$ to mean the distribution of relocation destinations used by drivers positioned at ℓ. Each component $r_{(\ell,d)}$ is the probability a randomly selected driver from ℓ would have selected d as their relocation destination:

$$r_{(\ell,d)} = \frac{\sum_{i \in \mathcal{M}_\ell} \mathbf{1}\{r_i = d\}}{|\mathcal{M}_\ell|}.$$

To simplify notation we will use $\mathcal{M}$ to refer to the set of driver indices, as well as their choice of disutility threshold and relocation destinations.

For each location ℓ let $\boldsymbol{x}_\ell = (x_{(\ell,d)} : d \in \mathcal{L})$ denote a common add-passenger threshold vector, potentially used by drivers at ℓ. For an error term $\epsilon > 0$, define

$$\mathcal{M}_\ell(\epsilon, \boldsymbol{x}_\ell) = \{i \in \mathcal{M}_\ell : \|\boldsymbol{x}_i - \boldsymbol{x}_\ell\|_\infty < \epsilon\}$$

to be the subset of drivers positioned at ℓ whose threshold vector is no more than ϵ away from $\boldsymbol{x}_\ell$ in any component. Let $\boldsymbol{x} = (\boldsymbol{x}_\ell : \ell \in \mathcal{L})$ denote a common threshold vector for each location. For $\epsilon > 0$, define

$$\mathcal{M}(\epsilon, \boldsymbol{x}) = \bigcup_{\ell \in \mathcal{L}} \mathcal{M}_\ell(\epsilon, \boldsymbol{x}_\ell).$$

Definition 6. *Let γ be any constant, and let $(\epsilon_k : k \geq 1)$ and $(\delta_k : k \geq 1)$ be nonnegative sequences which converge to 0 as $k \to \infty$. For any population-size $k \geq 1$ and any driver state $\mathcal{M}$, we say $\mathcal{M}$ is permissible with respect to $\gamma, \epsilon_k, \delta_k$ if the following conditions are satisfied:*

1. *The total number of drivers is no larger than γk, i.e $|\mathcal{M}| \leq \gamma k$.*
2. *There exists a common disutility threshold vector $\boldsymbol{x} = (\boldsymbol{x}_\ell : \ell \in \mathcal{L})$ such that the number of drivers who use a disutility threshold vector that is further than ϵ_k from $\boldsymbol{x}$ is vanishingly small, relative to k:*
$$\frac{|\mathcal{M} \setminus \mathcal{M}(\epsilon_k, \boldsymbol{x})|}{k} \leq \delta_k.$$

The concentration inequalities we provide in this section show that, when the matching process is applied to a driver state that is permissible with respect to $\gamma, \epsilon_k, \delta_k$, then with high probability, the difference between the stochastic output of the matching process and the corresponding fluid output is small.

We now describe what we mean by the fluid outcome associated with a particular driver state. Given population-size parameter $k \geq 1$, a driver state $\mathcal{M}$ determines the supply-location vector $\mathbf{S}_t$ by, in each component ℓ, taking the ratio between the total number of drivers at ℓ and k:

$$S_\ell = \frac{|\mathcal{M}_\ell|}{k}.$$

The market state $(\omega_t, \mathbf{S}_t)$ then determines the prices, $P_{(\ell,d)}$ for $(\ell,d) \in \mathcal{L}^2$, set by the SSP mechanism. The prices then determine the expected number of riders who request a dispatch:

$$\bar{R}_{(\ell,d)} = \mathbb{E}[D^k_{(\ell,d)}](1 - F_{(\ell,d)}(P_{(\ell,d)})),$$

where $D^k_{(\ell,d)}$ is the (stochastic) number of riders who are potentially interested in a dispatch from ℓ to d.

Definition 7. *Fix a population-size parameter k and let $\mathcal{M}$ be a driver-state that is permissible with respect to parameters $(\gamma, \epsilon_k, \delta_k)$. Let $\boldsymbol{x}$ be the common disutility threshold vector used by approximately all drivers in $\mathcal{M}$ (which exists from the second condition in the definition of permissible driver state, Definition 6). Let $\mathbf{r} = (\boldsymbol{r}_\ell : \ell \in \mathcal{L})$ denote the relocation distributions used by the population of drivers across each location.*

The fluid outcome associated with $\mathcal{M}$, k, are, for each route (ℓ,d), the dispatch trip volumes $\bar{G}_{(\ell,d)}$ and total trip volumes $\bar{F}_{(\ell,d)}$, which result from using the fluid matching process to allocate the dispatch demand volumes $\bar{R}_{(\ell,d)}$, along each route (ℓ,d), assuming $|\mathcal{M}_\ell|$ drivers are positioned at each location ℓ, and the drivers at ℓ use disutility threshold vector $\boldsymbol{x}_\ell$ and relocation-trip distribution $\boldsymbol{r}_\ell$.

We also define $\bar{f}_{(\ell,d)} = \frac{\bar{F}_{(\ell,d)}}{k}$ and $\bar{g}_{(\ell,d)} = \frac{\bar{G}_{(\ell,d)}}{k}$ to be the fluid outcome, normalized by the population size k.

To summarize, a permissible driver state $\mathcal{M}$ and a population-size parameter k induce both a deterministic fluid trip specification, denoted by dispatch trip volumes $\bar{G}_{(\ell,d)}$ and total trip volumes $\bar{F}_{(\ell,d)}$, for each route (ℓ,d), and stochastic trip specifications, denoted by dispatch trip volumes $G_{(\ell,d)}$ and total trip volumes $F_{(\ell,d)}$, for each (ℓ,d). The stochastic procedure governing $G_{(\ell,d)}$ and $F_{(\ell,d)}$ is described in Algorithm 2, and the deterministic procedure governing $\bar{G}_{(\ell,d)}$ and $\bar{F}_{(\ell,d)}$ is described in Definition 5. We will use $\mathbf{F}$, $\boldsymbol{G}$, $\bar{\mathbf{F}}$, $\bar{\boldsymbol{G}}$, to mean the corresponding vectors of trip counts (the vectors are indexed by routes $(\ell,d) \in \mathcal{L}^2$).

Also, define $H_{(\ell,d)} = F_{(\ell,d)} - G_{(\ell,d)}$ to mean the (stochastic) total number of relocation trips along (ℓ,d), and define $\bar{H}_{(\ell,d)} = \bar{F}_{(\ell,d)} - \bar{G}_{(\ell,d)}$ to mean the deterministic fluid number of relocation trips along (ℓ,d). Let $\boldsymbol{H}$ be the vector with components $H_{(\ell,d)}$ for each $(\ell,d) \in \mathcal{L}^2$ and let $\bar{\boldsymbol{H}}$ be the vector with components $\bar{H}_{(\ell,d)}$ for each $(\ell,d) \in \mathcal{L}^2$.

For each $k \geq 1$, let $\mathcal{D}_k$ be the set of driver states $\mathcal{M}$ that are permissible with respect to $(\gamma, \epsilon_k, \delta_k)$ when the population-size parameter is k. The main concentration lemma that we prove in this section is stated below:

Lemma 28. *There exist nonnegative sequences $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$, both converging to 0 as $k \to \infty$, such that the following equation is true for every k:*

$$\sup_{\mathcal{M} \in \mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}\left(\|\boldsymbol{H} - \bar{\boldsymbol{H}}\|_1 + \|\boldsymbol{G} - \bar{\boldsymbol{G}}\|_1\right) \leq \alpha_k\right) \geq 1 - q_k. \tag{87}$$

In the above equation, it is understood that the trip specifications $\boldsymbol{H}$, $\boldsymbol{G}$, $\bar{\boldsymbol{H}}$, $\bar{\boldsymbol{G}}$ are those which arise from the driver state $\mathcal{M}$ and the population-size parameter k.

Proof. We give a brief summary of the proof of Lemma 28, the details of which are contained in the Lemmas below.

Observe that the outcome of the matching process, i.e. the vector $\boldsymbol{G}$, is the sum of two vectors $\boldsymbol{G} = \boldsymbol{G}_1 + \boldsymbol{G}_2$ where $\boldsymbol{G}_1$ encodes the output from the first stage of the matching process and $\boldsymbol{G}_2$ encodes the output from the second stage. We analyze the convergence of $\boldsymbol{G}_1$ and $\boldsymbol{G}_2$ separately.

Lemma 31 provides an asymptotic concentration result for the output of the single-destination dispatch procedure with deterministic inputs, which we use to deduce that $\boldsymbol{G}_1$ converges asymptotically to $\bar{\boldsymbol{G}}_1$. Next, Lemma 32 provides asymptotic concentration for the single-destination dispatch procedure with inputs that have small stochastic perturbations. These small stochastic perturbations correspond to the second stage of the matching process, and are used to show that $\boldsymbol{G}_2$ converges to $\bar{\boldsymbol{G}}_2$.

Finally, Lemma 33 shows that the remaining undispatched drivers, i.e. the trips encoded by $\boldsymbol{H}$, converge to the deterministic fluid approximation $\bar{\boldsymbol{H}}$. □

Before proving Lemma 28, we provide a number of Lemmas that help us analyze the different components of the matching process. The following Lemma follows from standard concenteration inequalities for sub-Gaussian random variables.

Lemma 29. *Let $\mathcal{I}$ be an arbitrary index set and let $\gamma > 0$ be a constant. For each $k \geq 1$ and $i \in \mathcal{I}$, let $X_{k,i}$ be a Binomial random variable and let $R_{k,i}$ be a constant no larger than γk. Let $Z_{k,i} = \min(R_{k,i}, X_{k,i})$ and let $\bar{Z}_{k,i} = \min(R_{k,i}, \mathbb{E}[X_{k,i}])$. Then there exists concentration functions $\epsilon(k)$, $q(k)$ such that*

$$\sup_{i \in \mathcal{I}} \mathbb{P}\left(|Z_{k,i} - \bar{Z}_{k,i}| \geq \epsilon(k)\right) \leq q(k). \tag{88}$$

Lemma 30. *Moreover, for each k and i let $Y_{k,i}$ have a negative binomial distribution, let $M_{k,i}$ be a constant no larger than γk. Define $Z_{k,i} = \min(Y_{k,i}, M_{k,i})$ and $\bar{Z}_{k,i} = \min(\mathbb{E}[Y_{k,i}], M_{k,i})$. Then there exist concentration functions $\epsilon(k)$, $q(k)$ such that*

$$\sup_{i \in \mathcal{I}} \mathbb{P}\left(|Z_{k,i} - \bar{Z}_{k,i}| \geq \epsilon(k)\right) \leq q(k). \tag{89}$$

Next, we analyze asymptotic convergence of the single destination dispatch subroutine.

For a population-size parameter value k, consider the single destination dispatch subroutine with R riders and M drivers, both of which are smaller than γk. Assume that, except for a subset of size at most $k\delta_k$ drivers, each driver has probability of accepting a dispatch no more than ϵ_k away from some constant p.

Let G and U be random variables counting the number of accepted dispatches, and the number of undispatched drivers, respectively. Let $\bar{G}$ and $\bar{U}$ be the volume of accepted dispatches and undispatched drivers from the fluid matching process, with R riders, M drivers, and acceptance probability p.

Lemma 31. *There exists sequences of nonnegative numbers, $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$, both of which converge to 0 as $k \to \infty$, such that the following statement is true for every k:*

$$\sup_{M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]} \mathbb{P}\left(\frac{1}{k}|G - \bar{G}| \leq \alpha_k\right) \geq 1 - q_k \tag{90}$$

$$\sup_{M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]} \mathbb{P}\left(\frac{1}{k}|U - \bar{U}| \leq \alpha_k\right) \geq 1 - q_k. \tag{91}$$

In the above equations, it is understood that G and $\bar{G}$ are the stochastic and fluid number of accepted dispatches from the single-destination dispatch subroutine with M drivers, R riders, and, except for a subset of size at most $k\delta_k$ drivers, drivers use an acceptance probability within ϵ_k of p. Similarly, it is understood that U and $\bar{U}$ are the stochastic and fluid number of remaining undispatched drivers.

We defer the proof of Lemma 31 to Appendix H.3.

Our next Lemma proves asymptotic convergence when the single-destination dispatch subroutine is called twice, where the number of riders and drivers remaining unmatched in the first call are used as input for the second call to the procedure. We consider a situation where the input parameters for the single-destination dispatch subroutine are stochastically perturbed by a random variable which satisfies asymptotic concentration properties. We show the conclusion of Lemma 31 still hold despite this stochastic perturbation.

Specifically, for each $k \geq 1$, let $\mathcal{I}_k$ be a set of tuples of random variables $(X, Y) \in \mathcal{I}_k$, with deterministic fluid approximations $(\bar{X}, \bar{Y})$, such that the following concentration property is satisfied:

$$\sup_{(X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}\left(|X - \bar{X}| + |Y - \bar{Y}|\right) \leq \alpha_k\right) \geq 1 - q_k, \tag{92}$$

$$\tag{93}$$

where $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$ are nonnegative sequences which converge to 0 as $k \to \infty$.

Our next Lemma analyzes convergence of the single-destination dispatch subroutine when the initial number of drivers and riders are perturbed by subtracting X and Y. Let G and U be the stochastic output from when the single-destination dispatch subroutine when $M - X$ is the initial number of drivers and $R - Y$ is the initial number of drivers. Let $\bar{G}$ and $\bar{U}$ be the fluid number of drivers when the initial driver volume is $M - \bar{X}$ and the initial rider volume is $R - \bar{Y}$. We show that G and U converge asymptotically to $\bar{G}$ and $\bar{U}$, assuming that X and Y satisfy the concentration property (92).

Lemma 32. *For each $k \geq 1$, let $\mathcal{I}_k$ be a set of tuples of nonnegative random variables $(X, Y) \in \mathcal{I}_k$, with deterministic fluid approximations $(\bar{X}, \bar{Y})$, which satisfy the asymptotic concentration property (92). Let (M, R, p) be any constants satisfying $M \leq \gamma k$, $R \leq \gamma k$, and $p \in [0, 1]$. Let G be the number of dispatches and U the number of remaining drivers, when the single-destination dispatch subroutine (1) is used to allocate $R - Y$ dispatch requests to $M - X$ drivers, assuming that, except for a subset of size at most $k\delta_k$ drivers, drivers use an acceptance probability within ϵ_k of p. Let $\bar{G}$ and $\bar{U}$ be the output of the fluid subroutine when $R - \bar{Y}$ riders are allocated to $M - \bar{X}$ drivers. Then G converges asymptotically to $\bar{G}$ and U converges asymptotically to $\bar{U}$, in the sense that the following equation holds for all $k \geq 1$:*

$$\sup_{M\leq\gamma k,\ R\leq\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|G - \bar{G}| \leq \beta_k\right) \geq 1 - p_k \tag{94}$$

$$\sup_{M\leq\gamma k,\ R\leq\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|U - \bar{U}| \leq \beta_k\right) \geq 1 - p_k, \tag{95}$$

where $(\beta_k : k \geq 1)$ and $(p_k : k \geq 1)$ are nonnegative sequences which converge to 0 as $k \to \infty$.

Our final intermediate Lemma analyzes the asymptotic convergence of the relocation trips taken by the drivers. Recall that, in the fluid matching process, the volume of relocation trips towards each location is proportional to the volume of drivers who selected that destination as their relocation destination. Specifically, the fluid volume of relocation trips from ℓ to d is determined by the following equation:

$$\bar{H}_{(\ell,d)} = \left(M_\ell - \sum_{d'} \bar{G}_{(\ell,d')}\right)\left(\frac{\sum_{i\in\mathcal{M}_\ell} \mathbf{1}\left\{r_i = d\right\}}{|\mathcal{M}_\ell|}\right).$$

The factor on the left, i.e. $M_\ell - \sum_{d'} \bar{G}_{(\ell,d')}$, counts the volume of supply that does not serve a dispatch in the fluid model (recall $M_\ell = |\mathcal{M}_\ell|$ is the unnormalized volume of drivers at ℓ), and the factor on the right, $\frac{\sum_{i\in\mathcal{M}_\ell}\mathbf{1}\{r_i=d\}}{|\mathcal{M}_\ell|}$, counts the proportion of drivers positioned at ℓ who choose d as their relocation destination.

The stochastic number of relocation trips along each route (ℓ, d) is defined as the total number of undispatched drivers positioned at ℓ who chose d as their relocation destination. For a driver state $\mathcal{M}$, let $\mathcal{M}_\ell^R \subseteq \mathcal{M}_\ell$ be the (stochastic) subset of drivers who take a relocation trip. The number of relocation trips $H_{(\ell,d)}$ is defined by

$$H_{(\ell,d)} = \sum_{i\in\mathcal{M}_\ell^R} \mathbf{1}\left\{r_i = d\right\}.$$

Let us also use the notation $\bar{H}_\ell$ to mean the total volume of fluid relocation trips, and H_ℓ to mean the stochastic total number of relocation trips in the two level model:

$$\bar{H}_\ell = \sum_d \bar{H}_{(\ell,d)} \quad\text{and}\quad H_\ell = \sum_d H_{(\ell,d)}.$$

The following Lemma shows that the distribution of relocation trips converges asymptotically to the fluid distribution of relocation trips.

Lemma 33. *Suppose that the total number of relocation trips converges to the fluid volume of relocation trips, as $k \to \infty$, for all admissible driver states. That is, assume there exists sequences $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$ such that $\lim_{k\to\infty}\alpha_k = \lim_{k\to\infty} q_k = 0$, for which the following inequality holds for every k and every ℓ*

$$\sup_{\mathcal{M}\in\mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}|H_\ell - \bar{H}_\ell| \leq \alpha_k\right) \geq 1 - q_k.$$

Then the relocation trip volumes along each individual route converge to their fluid approximations, i.e. there exist sequences $(\beta_k : k \geq 1)$ and $(p_k : k \geq 1)$ such that $\lim_{k\to\infty}\beta_k = \lim_{k\to\infty} p_k = 0$, for which the following inequality holds for every k:

$$\sup_{\mathcal{M}\in\mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}\|\boldsymbol{H} - \bar{\boldsymbol{H}}\|_1 \leq \beta_k\right) \geq 1 - p_k.$$

Note that in the above equation, $\boldsymbol{H}$ and $\bar{\boldsymbol{H}}$ are vectors with components corresponding to $H_{(\ell,d)}$ and $\bar{H}_{(\ell,d)}$ for each route (ℓ, d). Therefore, the difference inside the probability is equal to the following:

$$\|\boldsymbol{H} - \bar{\boldsymbol{H}}\|_1 = \sum_{(\ell,d)\in\mathcal{L}^2} |H_{(\ell,d)} - \bar{H}_{(\ell,d)}|.$$

We defer the proof of Lemma 33 to Appendix H.5

H.3 Proof of Lemma 31

We restate Lemma 31 below.

Lemma. *There exists sequences of nonnegative numbers, $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$, both of which converge to 0 as $k \to \infty$, such that the following statement is true for every k:*

$$\sup_{M\leq\gamma k,\ R\leq\gamma k,\ p\in[0,1]} \mathbb{P}\left(\frac{1}{k}|G - \bar{G}| \leq \alpha_k\right) \geq 1 - q_k \tag{96}$$

$$\sup_{M\leq\gamma k,\ R\leq\gamma k,\ p\in[0,1]} \mathbb{P}\left(\frac{1}{k}|U - \bar{U}| \leq \alpha_k\right) \geq 1 - q_k. \tag{97}$$

In the above equations, it is understood that G and $\bar{G}$ are the stochastic and fluid number of accepted dispatches from the single-destination dispatch subroutine with M drivers, R riders, and, except for a subset of size at most $k\delta_k$ drivers, drivers use an acceptance probability within ϵ_k of p. Similarly, it is understood that U and $\bar{U}$ are the stochastic and fluid number of remaining undispatched drivers.

Proof. We start with a proof of equation (96), which states that G converges to $\bar{G}$ as $k \to \infty$.

Fix $k \geq 1$ and any $M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]$. Write the number of drivers M as $M = M' + M''$, such that M' is the number of drivers whose acceptance probability is within ϵ_k of p, and M'' is the number of drivers whose acceptance probability is further from p than ϵ_k. By assumption we have $M'' \leq \delta_k k$.

Consider the following modification of the single-destination dispatch subroutine parameters, which is designed to slightly underestimate the total number of dispatches produced the matching process. Assume that the dispatches are only allocated to the M' drivers whose acceptance probability is within ϵ_k of p, and assume that all M' drivers exactly use acceptance threshold $p - \epsilon_k$. Let the number of riders stay R. Let G_L be the number of accepted dispatches from this version of the single-destination dispatch subroutine. Also, define

$$\bar{G}_L = \min(M'(p - \epsilon_k), R)$$

to be the fluid output from this version of the single-destination dispatch subroutine.

Also consider the following modification, which is desgined to slightly overestimate the total number of dispatches. Assume that all the M'' drivers agree to serve a dispatch before the single-destination dispatch subroutine is called, so that the remaining number of riders is $R' = R - \min(R, M'')$ and the remaining number of drivers is M'. Also assume that all M' drivers exactly use threshold value $p + \epsilon_k$. Let G_U be the number of accepted dispatches from this process, i.e. G_U is equal to $\min(R, M'')$ plus the stochastic number of dispatches that occur when R' riders are matched to M' drivers using the single-destination dispatch subroutine, assuming all M' drivers have acceptance probability exactly equal to $p + \epsilon_k$. Define

$$\bar{G}_U = \min(R, M'') + \min(R', M'(p + \epsilon_k))$$

to be the fluid output from this version of the dispatch subroutine.

Notice that when all drivers use the same acceptance probability, the resulting number of dispatch trips is equal in distribution to the minimum of the number of drivers and a Binomial distribution parameterized by the number of drivers and the common acceptance probability. Therefore, by Lemma 29, we have the following bounds:

$$\sup_{M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]} \mathbb{P}\left(\frac{1}{k}|G_U - \bar{G}_U| \leq \alpha'_k\right) \geq 1 - q'_k$$

$$\sup_{M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]} \mathbb{P}\left(\frac{1}{k}|G_L - \bar{G}_L| \leq \alpha'_k\right) \geq 1 - q'_k,$$

where $(\alpha'_k : k \geq 1)$ and $(q'_k : k \geq 1)$ are sequences that converge to 0 as $k \to \infty$. Also, Notice that for any parameter values $M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1]$, we have the upper bound

$$\bar{G}_U - \bar{G}_L \leq M'' + M' 2\epsilon_k \leq k(\delta_k + 2\gamma\epsilon_k). \tag{98}$$

By construction, we have that G stochastically dominates G_L, and G_U stochastically dominates G. That is, for any $g \geq 0$, we have the following:

$$\mathbb{P}(G \leq g) \leq \mathbb{P}(G_L \leq g),$$

and

$$\mathbb{P}(G_U \leq g) \leq \mathbb{P}(G \leq g).$$

Therefore we obtain the following bounds, for any $\epsilon > 0$:

$$\begin{aligned}\mathbb{P}\left(|G - \bar{G}| \geq \epsilon\right) &\leq \mathbb{P}\left(G - \bar{G} \geq \epsilon\right) + \mathbb{P}\left(\bar{G} - G \geq \epsilon\right) \\ &\leq \mathbb{P}\left(G_U - \bar{G} \geq \epsilon\right) + \mathbb{P}\left(\bar{G} - G_L \geq \epsilon\right) \\ &\leq \mathbb{P}\left(|G_U - \bar{G}_U| \geq \epsilon + \bar{G} - \bar{G}_U\right) \\ &\quad + \mathbb{P}\left(|\bar{G}_L - G_L| \geq \epsilon + \bar{G}_L - \bar{G}\right). \end{aligned} \tag{99}$$

Finally, define

$$\alpha_k = \alpha'_k + (\delta_k + 2\gamma\epsilon_k).$$

By equation (99) we have the bound

$$\begin{aligned}\mathbb{P}\left(\frac{1}{k}|G - \bar{G}| \geq \alpha_k\right) &\leq \mathbb{P}\left(|G_U - \bar{G}_U| \geq k\alpha_k + \bar{G} - \bar{G}_U\right) \\ &\quad + \mathbb{P}\left(|\bar{G}_L - G_L| \geq k\alpha_k + \bar{G}_L - \bar{G}\right). \end{aligned} \tag{100}$$

Now, from equation (98), we have

$$k\alpha_k + \bar{G} - \bar{G}_U \geq k\alpha'_k,$$

and similarly

$$k\alpha_k + \bar{G}_L - \bar{G} \geq k\alpha'_k.$$

Therefore, continuing from (100), we have

$$\begin{aligned}\mathbb{P}\left(\frac{1}{k}|G - \bar{G}| \geq \alpha_k\right) &\leq \mathbb{P}\left(|G_U - \bar{G}_U| \geq k\alpha'_k\right) + \mathbb{P}\left(|\bar{G}_L - G_L| \geq k\alpha'_k\right) \\ &\leq q'_k + q'_k. \end{aligned} \tag{101}$$

Taking $q_k = 2q'_k$ finishes the proof.

The proof of U converging to $\bar{U}$ is analogous to the above argument. We first define U_L and U_U to mean the random number of undispatched drivers assuming all drivers use the acceptance probability $p - \epsilon_k$ and $p + \epsilon_k$, respectively, and we observe show that U_U stochastically dominates U which in turn stochastically dominates U_L. Stochastic dominance lets us bound the convergence of U in terms of the convergence of U_U and U_L. We then observe that each U_L and U_U is equal in distribution to the minimum of a constant and a negative binomial distribution, so Lemma 30 gives us large-population convergence.

□

H.4 Proof of Lemma 32

To prove Lemma 32 we first state and prove the following Lemma.

Lemma 34. *For each $k \geq 1$, let $\mathcal{Z}_k$ be a set of random variables $Z \in \mathcal{Z}_k$ with deterministic fluid approximations $\bar{Z}$. Assume that Z converges asymptotically to $\bar{Z}$, in the sense that there exists sequences $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$, both converging to 0 as $k \to \infty$, such that the following holds for every $k \geq 1$:*

$$\sup_{Z \in \mathcal{Z}_k} \mathbb{P}\left(\frac{1}{k}|Z - \bar{Z}| \leq \alpha_k\right) \geq 1 - q_k.$$

Let f be a Lipschitz continuous function with Lipschitz constant L. Then $f(Z)$ converges asymptotically to $f(\bar{Z})$, in the sense that the following equation holds:

$$\sup_{Z\in\mathcal{Z}_k} \mathbb{P}\left(\frac{1}{k}|f(Z)-f(\bar{Z})| \le L\alpha_k\right) \ge 1-q_k.$$

Proof. Observe that, if $\frac{1}{k}|Z-\bar{Z}| \le \alpha_k$ is true, then we have

$$\frac{1}{k}|f(Z)-f(\bar{Z})| \le \frac{1}{k}L|Z-\bar{Z}| \le L\alpha_k.$$

Therefore we have

$$\sup_{Z\in\mathcal{Z}_k} \mathbb{P}\left(\frac{1}{k}|f(Z)-f(\bar{Z})| \le L\alpha_k\right) \ge \sup_{Z\in\mathcal{Z}_k} \mathbb{P}\left(\frac{1}{k}|Z-\bar{Z}| \le \alpha_k\right) \ge 1-q_k,$$

as claimed.

□

We restate Lemma 32 below.

Lemma. *For each $k \ge 1$, let $\mathcal{I}_k$ be a set of tuples of nonnegative random variables $(X, Y) \in \mathcal{I}_k$, with deterministic fluid approximations $(\bar{X}, \bar{Y})$, which satisfy the asymptotic concentration property (92). Let (M, R, p) be any constants satisfying $M \le \gamma k$, $R \le \gamma k$, and $p \in [0,1]$. Let G be the number of dispatches and U the number of remaining drivers, when the single-destination dispatch subroutine (1) is used to allocate $R-Y$ dispatch requests to $M-X$ drivers, assuming that, except for a subset of size at most $k\delta_k$ drivers, drivers use an acceptance probability within ϵ_k of p. Let $\bar{G}$ and $\bar{U}$ be the output of the fluid subroutine when $R-\bar{Y}$ riders are allocated to $M-\bar{X}$ drivers. Then G converges asymptotically to $\bar{G}$ and U converges asymptotically to $\bar{U}$, in the sense that the following equation holds for all $k \ge 1$:*

$$\sup_{M\le\gamma k,\ R\le\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|G-\bar{G}| \le \beta_k\right) \ge 1-p_k \tag{102}$$

$$\sup_{M\le\gamma k,\ R\le\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|U-\bar{U}| \le \beta_k\right) \ge 1-p_k, \tag{103}$$

where $(\beta_k : k \ge 1)$ and $(p_k : k \ge 1)$ are nonnegative sequences which converge to 0 as $k \to \infty$.

Proof. Consider the single-destination dispatch subroutine when the population-size parameter is $k \ge 1$, with $M \le \gamma k$ drivers, $R \le \gamma k$ riders, $p \in [0,1]$ common acceptance probability, and let $(X, Y) \in \mathcal{I}_k$. Let G' and U' be the output of the fluid matching process with $R-Y$ riders and $M-X$ drivers. By Lemma 31 we have that G and U converge asymptotically to G' and U', in the sense that there exists $(\alpha'_k : k \ge 1)$, $(q'_k : k \ge 1)$, both converging to 0 as $k \to \infty$, such that the following holds for every k:

$$\sup_{M\le\gamma k,\ R\le\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|G-G'| \le \alpha'_k\right) \ge 1-q'_k$$

$$\sup_{M\le\gamma k,\ R\le\gamma k,\ p\in[0,1],\ (X,Y)\in\mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|U-U'| \le \alpha'_k\right) \ge 1-q'_k.$$

We now claim that G' and U' converge asymptotically to $\bar{G}$ and $\bar{U}$. Recall the output of the fluid single-destination dispatch subroutine is defined (Definition 4) in terms of deterministic functions $\bar{G}(\cdot)$ and $\bar{U}(\cdot)$, so we have

$$\begin{aligned} G' &= \bar{G}(R-Y, M-X, p) \quad & \bar{G} &= \bar{G}(R-\bar{Y}, M-\bar{X}, p) \\ U' &= \bar{U}(R-Y, M-X, p) \quad & \bar{U} &= \bar{U}(R-\bar{Y}, M-\bar{X}, p). \end{aligned}$$

Observe that the functions $\bar{G}(\cdot)$ and $\bar{U}(\cdot)$ are both Lipschitz continuous, so by Lemma 34 we have G' converges asymptotically to $\bar{G}$, i.e. the following equation holds for every $k \geq 1$

$$\sup_{M \leq \gamma k,\ R \leq \gamma k,\ p \in [0,1],\ (X,Y) \in \mathcal{I}_k} \mathbb{P}\left(\frac{1}{k}|G' - \bar{G}| \leq L\alpha_k\right) \geq 1 - q_k,$$

where α_k and q_k are the error term and probability term from the convergence of (X, Y) to $(\bar{X}, \bar{Y})$ (see equation (92)), and L is the Lipschitz constant for the function $\bar{G}(\cdot)$. Taking $\beta_k = \alpha'_k + L\alpha_k$ and $p_k = q'_k + q_k$ finishes the proof. □

H.5 Proof of Lemma 33

We restate Lemma 33 below.

Lemma. *Suppose that the total number of relocation trips converges to the fluid volume of relocation trips, as $k \to \infty$, for all admissible driver states. That is, assume there exists sequences $(\alpha_k : k \geq 1)$ and $(q_k : k \geq 1)$ such that* $\lim_{k\to\infty} \alpha_k = \lim_{k\to\infty} q_k = 0$*, for which the following inequality holds for every k and every ℓ*

$$\sup_{\mathcal{M} \in \mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}|H_\ell - \bar{H}_\ell| \leq \alpha_k\right) \geq 1 - q_k.$$

Then the relocation trip volumes along each individual route converge to their fluid approximations, i.e. there exist sequences $(\beta_k : k \geq 1)$ and $(p_k : k \geq 1)$ such that $\lim_{k\to\infty} \beta_k = \lim_{k\to\infty} p_k = 0$*, for which the following inequality holds for every k:*

$$\sup_{\mathcal{M} \in \mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}\|\boldsymbol{H} - \bar{\boldsymbol{H}}\|_1 \leq \beta_k\right) \geq 1 - p_k.$$

Proof. It suffices to prove the following is true for each route (ℓ, d):

$$\sup_{\mathcal{M} \in \mathcal{D}_k} \mathbb{P}\left(\frac{1}{k}\|H_{(\ell,d)} - \bar{H}_{(\ell,d)}\|_1 \leq \beta_k\right) \geq 1 - p_k.$$

If the above inequality holds for each route (ℓ, d), then the claimed inequality follows from a union bound over all locations.

Fix $k \geq 1$, an admissible driver state $\mathcal{M} \in \mathcal{D}_k$, and a location ℓ. Let $\mathcal{M}_\ell^R \subseteq \mathcal{M}$ be the stochastic subset of drivers who serve a relocation trip. By definition, we have $|\mathcal{M}_\ell^R|$ is the random variable H_ℓ, and by assumption we have that H_ℓ converges to a deterministic fluid approximation $\bar{H}_\ell$.

The probability any individual driver falls in this subset depends on the exact threshold vector that the driver has selected, as well as their realized add-passenger disutility. In particular, if two drivers use the exact same threshold vector, they have the same probability of going non-dispatched. If two drivers use approximately the same threshold, and we condition on the event that their sampled add-passenger disutilities are bounded away from the region where the different thresholds would lead to different decisions, then drivers still have the same probability of going non-dispatched.

Let $\boldsymbol{x}_\ell$ be the common disutility threshold vector which is approximately used by approximately all drivers at ℓ. Define

$$\mathcal{X}(\epsilon_k, \boldsymbol{x}_\ell) = [0, C] \setminus \left(\bigcup_{d \in \mathcal{L}} [x_{(\ell,d)} - \epsilon_k, x_{(\ell,d)} + \epsilon_k] \right)$$

to be the subset of feasible disutility thresholds $[0, C]$ where a small band $[x_{(\ell,d)} - \epsilon_k, x_{(\ell,d)} + \epsilon_k]$ centered at each threshold $x_{(\ell,d)}$ is removed.

Define

$$\mathcal{M}'_\ell = \mathcal{M}_\ell(\epsilon_k, \boldsymbol{x}_\ell) \cap \{i \in \mathcal{M}_\ell : X_i \in \mathcal{X}(\epsilon_k, \boldsymbol{x}_\ell)\}$$

to be the subset of drivers who approximately use the threshold $\boldsymbol{x}_\ell$ and whose sampled disutilities lie in $\mathcal{X}(\epsilon_k, \boldsymbol{x}_\ell)$. Notice that the cardinality $|\mathcal{M}'_\ell|$ has Binomial distribution with parameters $|\mathcal{M}_\ell(\epsilon_k, \boldsymbol{x}_\ell)|$ and $1 - 2|\mathcal{L}|\epsilon_k$, so $|\mathcal{M}'_\ell|$ converges to $|\mathcal{M}_\ell|$ as $k \to \infty$.

Define

$$H'_{(\ell,d)} = \left| \{i \in \mathcal{M}_\ell : r_i = d\} \cap \mathcal{M}'_\ell \cap \mathcal{M}^R_\ell \right|.$$

Notice that the number of trips from ℓ to d is bounded by

$$|\mathcal{M}_\ell \setminus \mathcal{M}'_\ell| + H'_{(\ell,d)} \geq H_{(\ell,d)} \geq H'_{(\ell,d)}.$$

Finally, notice that the distribution of $H'_{(\ell,d)}$ is equivalent to sampling $|\mathcal{M}^R_\ell|$ balls, without replacement, from a bag with $|\mathcal{M}'_\ell|$ balls, where each ball is associated with a destination, and counting how many balls are associated with the destination d. Concentration inequalities for sampling without replacement show that $H'_{(\ell,d)}$ converges asymptotically to $\bar{H}_{(\ell,d)}$, and this is sufficient to prove that $H_{(\ell,d)}$ converges to $\bar{H}_{(\ell,d)}$ asymptotically, because the difference between $H'_{(\ell,d)}$ and $H_{(\ell,d)}$ vanishes asymptotically.

□

I General Convex Analysis Properties

In this section we obtain useful convex analysis properties. We change notation from the rest of the paper, and consider the following generic convex optimization problem

$$\inf \left(f(x) \mid g_i(x) \leq 0, i = 1, 2, \ldots, m, x \in \mathbb{R}^n \right), \tag{104}$$

where $f, g_1, \ldots, g_m$ are convex functions from $\mathbb{R}^n$ to $\mathbb{R}$. We assume our convex program (104) satisfies the conditions described in Assumption 5.

Assumption 5. *Assume the following conditions hold:*

1. $f, g_1, \ldots, g_m$ *are continuously differentiable at every point* $x \in \mathbb{R}^n$.
2. *The feasible region* $\{x \mid g_i(x) \leq 0, i = 1, 2, \ldots, m, x \in \mathbb{R}^n\}$ *is a bounded compact set. In particular, there is a constant* γ *such that the feasible region lies in* $\bar{B}(0, \gamma)$*, i.e. the closed ball centered at* 0 *with radius* γ.
3. *The gradients* $\nabla f(x), \nabla g_1(x), \ldots, \nabla g_m(x)$ *have norm smaller than some constant* C *for all* $x \in \mathbb{R}^n$.

The Lagrangian associated with the optimization problem (104) is the function $L : \mathbb{R}^n \times \mathbb{R}^m_+ \to \mathbb{R} \cup \{+\infty\}$ defined by

$$L(x; \lambda) = f(x) + \lambda^T g(x), \tag{105}$$

where $g(x)$ is the vector in $\mathbb{R}^m$ with $g_i(x)$ as its ith component. Note the min-max theorem states the following relads :

$$\inf_{x \in \mathbb{R}^n} \sup_{\lambda \in \mathbb{R}^m_+} L(x; \lambda) = \sup_{\lambda \in \mathbb{R}^m_+} \inf_{x \in \mathbb{R}^n} L(x; \lambda). \tag{106}$$

Definition 8. *Let $x \in \mathbb{R}^n$. A nonnegative vector $\lambda \in \mathbb{R}^m_+$ is said to be a Lagrange multiplier vector for x if it satisfies the following conditions:*

- *Complementary slackness: $\lambda_i g_i(x) = 0$ holds for every $i = 1, 2, \ldots, m$.*
- *Stationarity: $\nabla_x L(x; \lambda) = 0$. In other words, x is a global minimizer of $L(\cdot; \lambda)$.*

We say λ is an ϵ-approximate Lagrange multiplier vector for x if it satisfies the following conditions:

- *Approximate complementary slackness: $|\lambda_i g_i(x)| \leq \epsilon$ holds for every $i = 1, 2, \ldots, m$.*
- *Approximate stationarity: $\|\nabla_x L(x; \lambda)\|_2 \leq \epsilon$.*

For the next Lemma, we consider an optimal solution x^* and a Lagrange multiplier vector λ^* for x^*. Part of our proof is concerned with the function mapping $x \in \mathbb{R}^n$ to the gradient of the Lagrangian $\nabla_x L(x; \lambda^*)$. In particular, we care about how large the gradient can vary when evaluated at two points that are close to one another. To reason about this maximum perturbation effect, for $\delta > 0$ define

$$\kappa(\delta) = \max \left(\|\nabla_x L(x; \lambda^*) - \nabla_x L(y; \lambda^*)\|_2 \mid x, y \in \bar{B}(0, \gamma), \|x - y\|_2 \leq \delta \right) \tag{107}$$

which gives the maximum norm of the difference between the gradients of any two points in $\bar{B}(0, \gamma)$ whose distance from each other is at most δ. From Assumption 5 we know the gradient $\nabla_x L(x; \lambda^*)$ is continuous in x, and we know that every continuous function over a compact set is uniformly continuous, so it follows that the gradient function, restricted to the closed ball $\bar{B}(0, \gamma)$, is uniformly continuous. Therefore, the maximum perturbation $\kappa(\delta)$ goes to 0 as $\delta \to 0$.

Lemma 35. *Let $\bar{x}$ be a feasible solution for (104) and let $\bar{\lambda} \in \mathbb{R}^m_+$. If $\bar{\lambda}$ is a Lagrange multiplier vector for $\bar{x}$ then $\bar{x}$ is an optimal solution for (104). If $\bar{\lambda}$ is an ϵ-approximate Lagrange multiplier vector for $\bar{x}$ then $\bar{x}$ is an ϵ'-optimal solution for (104), where $\epsilon' = \epsilon(n + 2\gamma)$.*

Proof. We focus on the case where $\bar{\lambda}$ is an ϵ-approximate Lagrange multiplier vector for $\bar{x}$. Let x^* and λ^* be optimal primal and dual variables. Observe the following chain of inequalities:

$$f(x^*) = L(x^*; \lambda^*) \geq L(x^*; \bar{\lambda}) \geq L(\bar{x}; \bar{\lambda}) + \nabla_x L(\bar{x}; \bar{\lambda})^T (x^* - \bar{x}),$$

where the first bound follows from optimality of λ^* and the second line follows by convexity of $L(\cdot; \bar{\lambda})$. Therefore we have the upper bound

$$f(\bar{x}) \leq f(x^*) + \|\nabla_x L(\bar{x}; \bar{\lambda})\|_2 \|(x^* - \bar{x})\|_2 + |\bar{\lambda}^T g(\bar{x})|.$$

By approximate complementary slackness we have $|\bar{\lambda}^T g(\bar{x})| \leq n\epsilon$, by approximate stationarity we have $\|\nabla_x L(\bar{x}; \bar{\lambda})\|_2 \leq \epsilon$, and by feasibility of x^* and $\bar{x}$ we have $\|(x^* - \bar{x})\|_2 \leq 2\gamma$. Therefore we obtain the bound

$$f(\bar{x}) \leq f(x^*) + \epsilon(n + 2\gamma) = f(x^*) + \epsilon'$$

as claimed. □

Lemma 36. *Let x^* be an optimal solution for (104) and let λ^* be a Lagrange multiplier vector for x^*. Let $\bar{x}$ be a feasible solution for (104). If $\bar{x}$ is an optimal solution then λ^* is a Lagrange multiplier vector for $\bar{x}$. If $\bar{x}$ is an ϵ-optimal solution, then λ^* is an ϵ'-approximate Lagrange multiplier vector, where $\epsilon' = \sqrt{\sqrt{\epsilon} + C\kappa(C\sqrt{\epsilon})}$.*

Proof. Observe the following chain of inequalities:

$$\begin{aligned} f(x^*) &= L(x^*; \lambda^*) \\ &= \min_{x \in \mathbb{R}^n} L(x; \lambda^*) \\ &\leq L(\bar{x}; \lambda^*) && (108) \\ &= f(\bar{x}) + \lambda^{*T} g(\bar{x}) \\ &\leq f(\bar{x}) && (109) \end{aligned}$$

The first two lines follow from the complementary slackness and stationarity conditions which hold between x^* and λ^*. Line (109) follows from feasibility of $\bar{x}$ and nonnegativity of λ^*.

We first consider the case where $\bar{x}$ is an optimal solution. In this case, we have $f(\bar{x}) = f(x^*)$ so every line in the above chain of inequalities holds with equality. In particular, line (108) holding with equality shows that $\bar{x}$ is a global minimizer of $L(\cdot; \lambda^*)$, so the stationarity condition $\nabla_x L(\bar{x}; \lambda^*)$ holds, and line (109) holding with equality shows that complementary slackness holds between $\bar{x}$ and λ^*. Therefore λ^* is a Lagrange multiplier vector for $\bar{x}$.

Next, we consider the case wehre $\bar{x}$ is an ϵ-optimal solution. In this case, we have $f(\bar{x}) \leq f(x^*) + \epsilon$, so ϵ is an upper bound on the difference between any two consecutive terms in the above chain of inequalities. Applying this upper bound to line (109) we conclude that $|\lambda^{*T} g(\bar{x})| \leq \epsilon$. Since every term $\lambda_i^* g_i(\bar{x})$ is nonpositive, it follows that $|\lambda_i^* g_i(\bar{x})| \leq |\lambda^{*T} g(\bar{x})| \leq \epsilon \leq \epsilon'$, so approximate complementary slackness holds between λ^* and $\bar{x}$.

We show in Lemma 37 that $\bar{x}$ satisfies approximate stationarity with respect to λ^*. We use the result in Lemma 37 by taking $h(x)$ to mean $L(x; \lambda^*)$. The result of Lemma 37 then provides the bound

$$\|\nabla_x L(\bar{x}; \lambda^*)\|_2^2 \leq \sqrt{\epsilon} + C\kappa(C\sqrt{\epsilon}).$$

Therefore, $\|\nabla_x L(\bar{x}; \lambda^*)\|_2 \leq \epsilon'$, establishing that λ^* is an ϵ'-approximate Lagrange multiplier vector for $\bar{x}$. □

Lemma 37. *Let $h : \mathbb{R}^n \to \mathbb{R}$ be a continuously differentiable convex function. Let x^* be a global minimizer of h, let $\bar{x}$ satisfy $h(\bar{x}) < h(x^*) + \epsilon$, and assume that $\bar{x}$ and x^* both lie in $\bar{B}(0, \gamma)$, i.e. the closed ball centered at 0 with radius γ. For $\delta > 0$, define*

$$\kappa(\delta) = \max\big(\|\nabla h(x) - \nabla h(y)\|_2 \mid x, y \in \bar{B}(0, \gamma), \|x - y\|_2 \leq \delta\big)$$

to be the maximum norm of the difference between the gradients of any two points in $\bar{B}(0, \gamma)$ whose distance from each other is at most δ. Assume C is a constant upper bound on $\|\nabla h(x)\|_2$. Then

$$\|\nabla h(\bar{x})\|_2^2 \leq \sqrt{\epsilon} + C\kappa(C\sqrt{\epsilon}). \tag{110}$$

Proof. Define

$$x' = \bar{x} - t\nabla h(\bar{x}).$$

for some $t > 0$. From convexity of h we have

$$h(x') + \langle \nabla h(x'), \bar{x} - x' \rangle \leq h(\bar{x}).$$

Rearranginge the above, and using the definition of x',

$$t\langle \nabla h(x'), \nabla h(\bar{x})\rangle \le h(\bar{x}) - h(x') \le \epsilon.$$

Now consider the following chain of inequalities:

$$\begin{aligned}
\|\nabla h(\bar{x})\|_2^2 &= \langle \nabla h(\bar{x}), \nabla h(\bar{x})\rangle \\
&= \langle \nabla h(x'), \nabla h(\bar{x})\rangle + \langle \nabla h(\bar{x}) - \nabla h(x'), \nabla h(\bar{x})\rangle \\
&\le \frac{\epsilon}{t} + \langle \nabla h(\bar{x}) - \nabla h(x'), \nabla h(\bar{x})\rangle \\
&\le \frac{\epsilon}{t} + \|\nabla h(\bar{x}) - \nabla h(x')\|_2 \|\nabla h(\bar{x})\|_2 \\
&\le \frac{\epsilon}{t} + \|\nabla h(\bar{x}) - \nabla h(x')\|_2 C \\
&\le \frac{\epsilon}{t} + C\kappa(\|\bar{x} - x'\|_2) \\
&\le \frac{\epsilon}{t} + C\kappa(\|t\nabla h(\bar{x})\|_2) \\
&\le \frac{\epsilon}{t} + C\kappa(Ct).
\end{aligned}$$

Setting $t = \sqrt{\epsilon}$, we obtain the claimed bound (110). □