Online Appendix

Appendix EC.1: Proofs

Proof of Lemma 1. First, we notice that the sum of the expected number of drivers in service and idle drivers equals the total number of drivers in the system, so we have $L^d_o(\lambda,\phi) = L - L^d_s(\lambda,\phi)$. Let $L^d_t(\lambda,\phi)$ be the expected number of riders on a trip under policy (λ,ϕ). Since the total number of riders in service includes both those waiting to be picked up and those currently on a trip, we have $L^d_s(\lambda,\phi) = L^r_p(\lambda,\phi) + L^d_t(\lambda,\phi)$. We consider an observation period T. Over this period, $L^d_t(\lambda,\phi)T$ is the total on-trip time for all riders. $L^d_t(\lambda,\phi)T/t_0$ is the expected number of riders served by the platform over period T. As $T\to\infty$, by the law of large numbers, the average on-trip time converges to t_0. Thus, the average number of riders served by the platform per time unit can be computed by $\lim_{T\to\infty} L^d_t(\lambda,\phi)T/(t_0T) = L^d_t(\lambda,\phi)/t_0$. On the other hand, if we increase the base fare p_0 by 1 unit, the objective $\tilde{\mathcal{R}}(p_0+1,\lambda,\phi)$ will increase by the average number of riders served per time unit according to equation (2). Thus, if the platform increases the base fare by $w^r_p t_0$, we have $\mathcal{R}(p_0+w^r_p t_0,\lambda,\phi) = \mathcal{R}(p_0,\lambda,\phi) + w^r_p t_0 L^d_t(\lambda,\phi)/t_0$. Consequently, equation (2) can be rewritten as

$$
\begin{aligned}
&\mathcal{R}(p_0,\lambda,\phi) - w^d_s L^d_s(\lambda,\phi) - w^r_q L^r_q(\lambda,\phi) - w^r_p L^r_p(\lambda,\phi) - w^d_o L^d_o(\lambda,\phi)\\
=&\mathcal{R}(p_0,\lambda,\phi) - w^d_s L^d_s(\lambda,\phi) - w^r_q L^r_q(\lambda,\phi) - w^r_p \left(L^d_s(\lambda,\phi) - L^d_t(\lambda,\phi)\right) - w^d_o (L - L^d_s(\lambda,\phi))\\
=&\mathcal{R}(p_0,\lambda,\phi) + \frac{w^r_p t_0 L^d_t(\lambda,\phi)}{t_0} - (w^d_s + w^r_p - w^d_o) L^d_s(\lambda,\phi) - w^r_q L^r_q(\lambda,\phi) - w^d_o L\\
=&\mathcal{R}(p_0+w^r_p t_0,\lambda,\phi) - (w^d_s + w^r_p - w^d_o) L^d_s(\lambda,\phi) - w^r_q L^r_q(\lambda,\phi) - w^d_o L,
\end{aligned}
$$

which proves the result. □

Proof of Lemma 2. We start with the proof of condition (1). Let (l',m') be a type 1 state. To show that (l,m) is also a type 1 state for any $l'\le l\le L-1$ and $1\le m\le m'$, it is sufficient to show that $(l'+1,m')$ and $(l',m'-1)$ are both type 1 state. Then condition (1) can be proved by recursively applying this result. We first show that $(l'+1,m')$ is a type 1 state. According to Assumption 1, we have

$$(l'+1)\mu_{l'+1,m'} - (l'+1)\mu_{l'+1,m'-1} \ge l'\mu_{l',m'} - l'\mu_{l',m'-1}, \tag{EC.1}$$

$$(l'+1)\mu_{l'+1,m'-1} - (l'+2)\mu_{l'+2,m'-1} \ge l'\mu_{l',m'-1} - (l'+1)\mu_{l'+1,m'-1}. \tag{EC.2}$$

Inequality (EC.1) is from condition (1) in Assumption 1, and inequality (EC.2) is from condition (2) in Assumption 1. Summing equation (EC.1) with equation (EC.2) on both sides yields

$$(l'+1)\mu_{l'+1,m'} - (l'+2)\mu_{l'+2,m'-1} \ge l'\mu_{l',m'} - (l'+1)\mu_{l'+1,m'-1} > 0, \tag{EC.3}$$

where the second inequality is due to the fact that (l', m') is a type 1 state which satisfies $l'\mu_{l',m'} > (l'+1)\mu_{l'+1,m'-1}$. Equation (EC.3) yields $(l'+1)\mu_{l'+1,m'} > (l'+2)\mu_{l'+2,m'-1}$, which shows that $(l'+1, m')$ is a type 1 state by definition. To show $(l', m'-1)$ is also a type 1 state, similarly, Assumption 1 gives

$$l'\mu_{l',m'-1} - l'\mu_{l',m'-2} \geq l'\mu_{l',m'} - l'\mu_{l',m'-1}, \tag{EC.4}$$

$$l'\mu_{l',m'-2} - (l'+1)\mu_{l'+1,m'-2} \geq l'\mu_{l',m'-1} - (l'+1)\mu_{l'+1,m'-1}. \tag{EC.5}$$

Summing equation (EC.4) with equation (EC.5) on both sides, we get

$$l'\mu_{l',m'-1} - (l'+1)\mu_{l'+1,m'-2} \geq l'\mu_{l',m'} - (l'+1)\mu_{l'+1,m'-1} > 0,$$

which shows that $(l', m'-1)$ is a type 1 state.

So far, we have proved condition (1). Now we consider the case when (l', m') is a type 2 state. If there exists some type 1 state (l_1, m_1) such that $l_1 \leq l'$ and $m_1 \geq m'$, according to condition 1, we know (l', m') must be a type 1 state as well, which leads to contradiction. Hence, we know condition (2) holds. □

Proof of Theorem 1. We prove the theorem under three cases $c_d = c_r$, $c_d > c_r$, and $c_d < c_r$. In all cases, for an arbitrary dispatching policy, we construct a new policy that satisfies the optimality properties stated in the theorem. We then show that this constructed policy achieves a weakly higher objective value than the original policy.

Case 1: $c_d = c_r$. Let ϕ^* be the dispatching policy such that $\phi^*(l, m) = 0$ for all type 1 states and $\phi^*(l, m) = 1$ for all type 2 states and λ^* the optimal dynamic pricing policy under ϕ^*. We want to show $\tilde{\mathcal{R}}(\phi^*, \lambda^*) \geq \tilde{\mathcal{R}}(\phi, \lambda)$ for an arbitrary policy (ϕ, λ). We consider a coupling of two systems, where system 1 uses policy (ϕ^*, λ^*) and system 2 uses the policy (ϕ, λ). This coupling aligns all random events in both systems, including arrivals, trip completions, and dispatches. Instead of using the optimal pricing policy λ^* for system 1, we apply a dependent pricing policy λ^{D} that is dependent on system 2. Specifically, λ^{D} is constructed to keep the effective request rate in system 1 always the same as that in system 2, even if the two systems are in different states. In other words, for any state transition in system 2, the effective request rate in system 1 is immediately adjusted to the same as the new effective request rate in system 2. We know $\tilde{\mathcal{R}}(\phi^*, \lambda^*) \geq \tilde{\mathcal{R}}(\phi^*, \lambda^{\mathrm{D}})$. Thus, it is sufficient to show $\tilde{\mathcal{R}}(\phi^*, \lambda^{\mathrm{D}}) \geq \tilde{\mathcal{R}}(\phi, \lambda)$. We demonstrate this in the following.

Let $\{X_t\}_{t=1}^{\infty}$ and $\{Y_t\}_{t=1}^{\infty}$ denote the coupled Markov chain, where each step t corresponds to a transition in the joint process. At step t, the state of system 1 is given by $X_t = (l_t, m_t)$, and the state of system 2 is given by $Y_t = (l'_t, m'_t)$. Without loss of generality, we assume that both systems start with the same number of drivers in service and riders in the queue, that is, $X_1 = Y_1$. If a random event occurs in either system at step t, the states of systems 1 and 2 transition to X_{t+1} and Y_{t+1} after possible dispatchings, respectively, based on their dispatching policies. After that, the coupled chain proceeds to step $t+1$.

Note that any arrival events will occur in both chains at the same time in the coupled system because the effective request rate in system 1 is always set to be the same as that in system 2. However, a trip completion may occur in only one of the chains when the service completion rate in one system is larger than the other.

In the following proof, we show that under coupling, there are three possible "phases", which are mutually exclusive, for the pair (X_t, Y_t) at step t. The coupled process loops over these phases. Moreover, we will show that, at each phase, the penalty rate in system 1 is consistently less than or equal to one in system 2.

Phase 1. At step t, the process is in phase 1 if $X_t = Y_t = (l_t, m_t)$. Since both systems are in the same state and thus have the same service completion rate, there are two possible events for the coupled systems:

(1) a rider joining the queue for both systems;
(2) a trip completion for both systems.

For both events, we have two possible subcases described as follows. First, if the number of dispatched drivers under ϕ^* and ϕ is the same, then $X_{t+1} = Y_{t+1}$, and the process remains in phase 1 at step $t+1$.

Second, if the two policies dispatch a different number of drivers, the coupled states $X_{t+1} \neq Y_{t+1}$. Since the system will lie in the same anti-diagonal after dispatching, we know both systems lie in the same anti-diagonal at step $t+1$ under each event. Then, the process leaves phase 1 and enters phase 2 at step $t+1$.

Phase 2. At step t, the process is in phase 2 if $X_t = (l_t, m_t) \neq Y_t = (l'_t, m'_t)$ and $l_t + m_t = l'_t + m'_t$ (i.e., X_t and Y_t are on the same anti-diagonal line). By Lemma 2, the trip completion rate follows $l_t \mu_{l_t, m_t} \geq l'_t \mu_{l'_t, m'_t}$, which means system 1 has a larger trip completion rate than system 2 at the coupled state. Thus, it is possible that only system 1 experiences a trip completion at the next event. There are three possible events for the coupled systems:

(1) a rider joining the queue for both systems;
(2) a trip completion for both systems;
(3) a trip completion only for system 1 (due to larger trip completion rate in system 1).

For events (1) and (2), as the number of dispatched drivers does not affect the sum of l_{t+1} and m_{t+1}, both systems will lie on the same anti-diagonal line at step $t+1$. If $X_{t+1} \neq Y_{t+1}$, the process stays at phase 2. If $X_{t+1} = Y_{t+1}$, the process leaves phase 2 and re-enters phase 1.

For event (3), we have $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$, and the process leaves phase 2 and enters into phase 3.

Phase 3. At step t, the process is in phase 3 if $l_t + m_t < l'_t + m'_t$. In this phase, we cannot determine which state, X_t or Y_t, has the higher trip completion rate. Thus, there are four possible events for the coupled systems:

(1) a rider joining the queue for both systems;
(2) a trip completion for both systems;
(3) a trip completion only for system 1 (if system 1 has a larger trip completion rate);
(4) a trip completion only for system 2 (if system 2 has a larger trip completion rate).

When events (1), (2), and (3) occur, the process stays at phase 3 since one can easily show that $l_{t+1}+m_{t+1} < l'_{t+1}+m'_{t+1}$ holds.

When event (4) occurs, there are two possible sub-cases: $l'_{t+1}+m'_{t+1} = l_{t+1}+m_{t+1}$ and $l'_{t+1}+m'_{t+1} > l_{t+1}+m_{t+1}$. If $l'_{t+1}+m'_{t+1} = l_{t+1}+m_{t+1}$, then the process leaves phase 3 and re-enters either phase 1 (when $X_{t+1}=Y_{t+1}$) or phase 2 (when $X_{t+1}\neq Y_{t+1}$). Otherwise, if $l'_{t+1}+m'_{t+1} > l_{t+1}+m_{t+1}$, the process stays at phase 3.

In all phases above, since the effective request rate for both systems are kept the same, the revenue rate for both systems are the same as well. However, under coupling, the penalty rate in system 1 is consistently less or equal to that in system 2 because of weakly smaller $l+m$ values. Hence, we have $\tilde{\mathcal{R}}(\phi^*,\lambda^*)\geq\tilde{\mathcal{R}}(\phi^*,\lambda^{\mathrm{D}})\geq\tilde{\mathcal{R}}(\phi,\lambda)$, which proves the result. The first inequality holds because there always exists an optimal Markovian policy (see Theorem 8.1.2 of Puterman 1994).

Case 2: $c_d > c_r$. The proof idea here is similar to the first case. We show that there exists a policy (ϕ^D,λ^D) such that $\tilde{\mathcal{R}}(\phi^{\mathrm{D}},\lambda^{\mathrm{D}})\geq\tilde{\mathcal{R}}(\phi,\lambda)$ holds for any policy (ϕ,λ). Construct ϕ^{D} such that $\phi^{\mathrm{D}}(l,m)=0$ for all type 1 states (l,m). We consider a coupling of two systems, where system 1 uses policy $(\phi^{\mathrm{D}},\lambda^{\mathrm{D}})$ and system 2 uses policy (ϕ,λ). Similar to the case of $c_d=c_r$, λ^{D} here is defined to keep the effective request rate in system 1 always the same as that in system 2. Again, let $\{X_t\}_{t=1}^{\infty}$ and $\{Y_t\}_{t=1}^{\infty}$ denote the coupled Markov chain, where each step t corresponds to a transition in the joint process. In the following, we describe the coupling processes and the constructed dispatching policy ϕ^{D} in each phase.

Phase 1. At step t, the process is in phase 1 if $X_t=Y_t$. Let $X_t=(l_t,m_t)$ denote the corresponding state of system 1 at step t. There are two possible events for the coupled systems, which are the same as those in case 1, phase 1. In this phase, we set $\phi^{\mathrm{D}}(l,m)=\phi(l,m)$, which is the dispatching policy in system 2, for all type 2 states (l,m). We now consider next state $Y_{t+1}=(l'_{t+1},m'_{t+1})$ for system 2. Let P^* be the corresponding zigzag path under ϕ^*, the optimal dispatching policy when $c_d=c_r$. There are two sub-cases for Y_{t+1}:

- Y_{t+1} is a type 2 state or $Y_{t+1}\in P^*$. Then by the construction of ϕ^{D}, both systems dispatch the same number of drivers at step t, so we have $X_{t+1}=Y_{t+1}$. The process remains in phase 1.
- Y_{t+1} is a type 1 state and $Y_{t+1}\notin P^*$. the number of dispatched drivers differs between the two policies, so we have $X_{t+1}\neq Y_{t+1}$. Since $\phi^{\mathrm{D}}(l,m)=0$ for all type 1 states (l,m), we know $X_{t+1}\in P^*$. Moreover, X_{t+1} and Y_{t+1} still lie on the same anti-diagonal, and $l_{t+1}<l'_{t+1}$. The process leaves phase 1 and enters into phase 2.

Phase 2. At step t, the process is in phase 2 if $l_t<l'_t$, $m_t>m'_t$, $l_t+m_t=l'_t+m'_t$, and $X_t\in P^*$. In other words, Y_t lies on the same anti-diagonal as X_t but is positioned to the lower left of X_t. In this phase, we set $\phi^{\mathrm{D}}=\phi^*$, which is the optimal dispatching policy when $c_d=c_r$. The coupled systems face the same three events defined in phase 2, case 1 ($c_d=c_r$). Again, we use the same label (1), (2), and (3) described in phase 2, case 1 to represent each event.

When event (1) or (2) occurs, if $X_{t+1} \neq Y_{t+1}$, we know $l_{t+1} < l'_{t+1}$ as ϕ^* makes at most one dispatch after an event occurs. Thus, the process stays at phase 2. If $X_{t+1} = Y_{t+1}$, the process leaves phase 2 and re-enters into phase 1.

When event (3) occurs, we have $l_{t+1} < l'_{t+1}$ and $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$, and the process leaves phase 2 and enters phase 3.

Phase 3. At step t, the process is in phase 3 if $l_t < l'_t$ and $l_t + m_t < l'_t + m'_t$. There are four possible events for the coupled systems, which are exactly the same as those in phase 3, case 1. We use event (1), (2), (3), and (4) to describe each event. We set dispatching policy ϕ^{D} to be dependent on the event that occurred.

When event (1), (2), or (3) occur, the two systems keep their anti-diagonal relationship, that is, $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$. Under these events, we set $\phi^{\mathrm{D}}(l, m) = 0$ for all type 2 states (l, m), which means system 1 will not make a dispatch under these events. Then $l_{t+1} < l'_{t+1}$ and $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$ hold due to the following facts:

- For event (1), we have $m_{t+1} = m_t + 1 > m'_t + 1 \geq m'_{t+1}$ and $l_{t+1} = l_t < l'_t \leq l'_{t+1}$.
- For events (2) and (3), we have $m_{t+1} = m_t > m'_t \geq m'_{t+1}$ and $l_{t+1} = l_t - 1 < l'_t - 1 \leq l'_{t+1}$.

Thus, the process will stay at phase 3 since all conditions for phase 3 are satisfied.

When event (4) occurs, we have $l_{t+1} + m_{t+1} = l_t + m_t$ and $l'_{t+1} + m'_{t+1} = l'_t + m'_t - 1$. We construct ϕ^{D} according to the following two sub-cases.

- $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$. We set $\phi^{\mathrm{D}}(l, m) = 0$ for all type 2 states (l, m) as well. Then we know $m_{t+1} = m_t > m'_t \geq m'_{t+1}$, and thus $l_{t+1} < l'_{t+1}$ holds. As a result, the process will stay at phase 3.
- $l_t + m_t = l'_{t+1} + m'_{t+1}$. We set $\phi^{\mathrm{D}} = \phi^*$, the optimal dispatching policy when $c_d = c_r$. It is easy to verify that the process will enter either phase 1 if $l_{t+1} = l'_{t+1}$ or phase 2 if $l_{t+1} < l'_{t+1}$.

In all phases above, the revenue rates are identical for both systems, while the penalty rate in system 1 is always less than or equal to that in system 2 at coupled states because system 1 has weakly smaller $l + m$ and l values. Hence, again the objective function satisfies $\tilde{\mathcal{R}}(\phi^*, \lambda^*) \geq \tilde{\mathcal{R}}(\phi^{\mathrm{D}}, \lambda^{\mathrm{D}}) \geq \tilde{\mathcal{R}}(\phi, \lambda)$, which proves the result.

Case 3: $c_d < c_r$. Similar to case 2, we show there exists ϕ^{D} and λ^{D} such that $\tilde{\mathcal{R}}(\phi^{\mathrm{D}}, \lambda^{\mathrm{D}}) \geq \tilde{\mathcal{R}}(\phi, \lambda)$ for an arbitrary policy (ϕ, λ). We set λ^{D} to keep the effective request rate in system 1 always the same as that in system 2. We set ϕ^{D} such that $\phi^{\mathrm{D}}(l, m) = 1$ for all type 2 states (l, m). For type 1 states, ϕ^{D} depends on system 2, and will be defined later. Again, let $\{X_t\}_{t=1}^{\infty}$ and $\{Y_t\}_{t=1}^{\infty}$ denote the coupled Markov chain, where each step t corresponds to a transition in the joint process. By the partial construction of ϕ^{D}, we know X_t is a type 1 state for all t. There are four possible phases in this case, which are illustrated below.

Phase 1. At step t, the process is in phase 1 if $X_t = Y_t$. In this phase, there are two possible cases for Y_{t+1} and the construction of ϕ^{D} depends on whether Y_{t+1} is a type 1 state or type 2 state.

- Y_{t+1} is a type 1 state. We set $\phi^{\mathrm{D}} = \phi$ for all type 1 states. Then the next coupled states satisfy $X_{t+1} = Y_{t+1}$, and the process remains in phase 1.
- Y_{t+1} is a type 2 state. We set $\phi^{\mathrm{D}} = 0$ for all type 1 states (i.e., $\phi^{\mathrm{D}} = \phi^*$). The next coupled states satisfy $X_{t+1} \neq Y_{t+1}$ and $X_{t+1} \in P^*$ since $\phi^{\mathrm{D}}(l,m) = 1$ for all type 2 states. Since the state always stay on the same anti-diagonal after a dispatching action, we have $l_{t+1} + m_{t+1} = l'_{t+1} + m'_{t+1}$ and $m_{t+1} < m'_{t+1}$. Then the process leaves phase 1 and enters phase 2.

Phase 2. At step t, the process is in phase 2 if $m_t < m'_t$, $l_t > l'_t$, $l_t + m_t = l'_t + m'_t$, and $X_t \in P^*$. Thus, Y_t must be a type 2 state. There are three possible events for the coupled systems, which are the same as those in phase 2, case 1. Again, there are two possible cases for Y_{t+1} and the construction of ϕ^{D} depends on whether Y_{t+1} is a type 1 state or type 2 state.

- Y_{t+1} is a type 1 state. We set $\phi^{\mathrm{D}} = \phi$. Only event (1) or (2) can occur here because Y_t is a type 2 state and Y_{t+1} is a type 1 state, so there must be a state transition in system 2. Thus, we have $l_{t+1} + m_{t+1} = l'_{t+1} + m'_{t+1}$. Moreover, we know $X_{t+1} = Y_{t+1}$ because both systems have the same dispatching policy at type 1 states and Y_{t+1} becomes a type 1 state now. Consequently, the process leaves phase 2 and enters phase 1 again.
- Y_{t+1} is a type 2 state. We set $\phi^{\mathrm{D}} = \phi^*$, the optimal dispatching policy under $c_d = c_r$. Then $X_{t+1} \in P^*$ is a type 1 state. When event (1) or (2) occurs, we have $l_{t+1} + m_{t+1} = l'_{t+1} + m'_{t+1}$ because dispatching does not change the anti-diagonal of the state. Since Y_{t+1} is a type 2 state and X_{t+1} must be a type 1 state, we know $m_{t+1} < m'_{t+1}$ and $l_{t+1} > l'_{t+1}$ hold. Thus, the process stays in phase 2.

  When event (3) occurs, we have $l_{t+1} + m_{t+1} = l_t + m_t - 1 < l'_t + m'_t = l'_{t+1} + m'_{t+1}$ Moreover, m_{t+1} either equals m_t (not dispatch) or $m_t - 1$ (dispatch), which is strictly less than m'_t. The process then enters phase 3.

Phase 3. At step t, the process is in phase 3 if $m_t < m'_t$, $l_t + m_t < l'_t + m'_t$, and $X_t \in P^*$. There are four possible events for the coupled systems, which are exactly the same as those in phase 3, case 1. We construct ϕ^{D} based on whichever of m_t and m'_{t+1} is larger.

If $m_t > m'_{t+1}$, at all type 1 states, we set ϕ^{D} to keep dispatching until $m_{t+1} = m'_{t+1}$. Note that event (3) is not possible to occur under $m_t > m'_{t+1}$. In specific, we set ϕ^{D} according to the following under events (1), (2), and (4):

(a) Under event (1), for every $k = 0, \cdots, \min\{L - l_t, m_t - m'_{t+1} + 1\}$, we set $\phi^{\mathrm{D}}(l_t + k, m_t + 1 - k) = 1$;
(b) Under event (2), for every $k = 0, \cdots, \min\{L - l_t + 1, m_t - m'_{t+1}\}$, we set $\phi^{\mathrm{D}}(l_t - 1 + k, m_t - k) = 1$;
(c) Under event (4), for every $k = 0, \cdots, \min\{L - l_t, m_t - m'_{t+1}\}$, we set $\phi^{\mathrm{D}}(l_t - 1 + k, m_t - k) = 1$.

For any type 1 state (l,m) not covered above, we set $\phi^{\mathrm{D}}(l,m) = 0$.

On the other hand, if $m_t \leq m'_{t+1}$, we set $\phi^{\mathrm{D}} = \phi^*$, the optimal dispatching policy under $c_d = c_r$.

Now we discuss what happens for systems 1 and 2 under each event. When events (1) and (2) occur, there are three possible cases, which are $m_t > m'_{t+1}$, $m_t = m'_{t+1}$, and $m_t < m'_{t+1}$.

- If $m_t > m'_{t+1}$ due to a dispatching in system 2, we make dispatches in system 1 to state $m_{t+1} = m'_{t+1}$. Since both systems have an arrival or a trip completion, we have $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$. Then, the process enters into phase 4.
- If $m_{t+1} = m'_{t+1}$, we have $\phi^{\mathrm{D}} = \phi^*$, so we know $X_{t+1} \in P^*$ and $l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$. Thus, the process enters phase 4.
- If $m_{t+1} = m'_{t+1}$, we have $\phi^{\mathrm{D}} = \phi^*$, the process stays in phase 3.

When event (3) occurs, the state Y_{t+1} is the same as Y_t, so we have $m_t < m'_t = m'_{t+1}$. As a result, ϕ^{D} is set as ϕ^*. Consequently, $m_{t+1} \le m_t < m'_t = m'_{t+1}$ and $X_{t+1} \in P^*$ hold. The process stays at phase 3. When event (4) occurs, we have $l_t + m_t = l_{t+1} + m_{t+1}$. There are two sub-cases, depending on whether the two systems are on the same anti-diagonal at $t + 1$.

- $l_t + m_t = l_{t+1} + m_{t+1} = l'_{t+1} + m'_{t+1}$, that is, the two systems are on the same anti-diagonal. Since $m_{t+1} \le m_t$ holds, we know $m_{t+1} \le m'_{t+1}$ holds according to our constructed ϕ^{D}. If $m'_{t+1} = m_{t+1}$, the process re-enters phase 1. If $m_{t+1} < m'_{t+1}$, we know ϕ^{D} is constructed as ϕ^*. Thus, the process re-enters phase 2.
- $l_t + m_t = l_{t+1} + m_{t+1} < l'_{t+1} + m'_{t+1}$. If $m_t < m'_{t+1}$, we set $\phi^{\mathrm{D}} = \phi^*$, so the state in system 1 does not change (i.e., $l_{t+1} = l_t$ and $m_{t+1} = m_t$). Thus, the process stays at phase 3. If $m_t > m'_{t+1}$, according to constructed ϕ^{D}, we have $m_{t+1} = m'_{t+1}$. The process enters phase 4. If $m_t = m'_{t+1}$, we set $\phi^{\mathrm{D}} = \phi^*$, and the process also enters phase 4 due to $m_{t+1} = m_t = m'_{t+1}$.

Phase 4. At step t, the process is in phase 4 if $m_t = m'_t$, $l_t < l'_t$, and $l_t + m_t < l'_t + m'_t$ hold. There are four possible events for the coupled system, which are exactly the same as those in phase 3, case 1. We construct ϕ^{D} according to the state type of Y_{t+1}. If Y_{t+1} is a type 1 state, we use the same construction of ϕ^{D} as one in phase 3 under $m_t > m'_{t+1}$. That is, at all type 1 states, ϕ^{D} is set to keep dispatching until $m_{t+1} \le m'_{t+1}$ (i.e., dispatching the same number of drivers in both systems if system 1 is still in a type 1 state before taking the action). Otherwise, if Y_{t+1} is a type 2 state, we set $\phi^{\mathrm{D}} = \phi^*$.

Now we discuss how states change under each event. When events (1) and (2) occur, if Y_{t+1} is a type 1 state, we have two subcases:

- $(l_t, m_t + 1)$ is a type 1 state under event (1) or $(l_t - 1, m_t)$ is a type 1 state under event (2). In this case, both systems dispatch the same number of drivers. Thus, we have $m_{t+1} = m'_{t+1}$ and the process stays in phase 4.
- $(l_t, m_t + 1)$ is a type 2 state under event (1) or $(l_t - 1, m_t)$ is a type 2 state under event (2). If system 2 dispatches at least one driver, then according to the construction of ϕ^{D}, both systems dispatch the same number of drivers as well, so the process stays in phase 4. If system 2 does not dispatch any drivers, system 1 will dispatch exactly one driver since we require $\phi^{\mathrm{D}}(l, m) = 1$ for all type 2 states. Therefore, it follows $m_{t+1} = m_t - 1 = m'_t - 1 = m'_{t+1} - 1 < m'_{t+1}$ and $X_{t+1} \in P^*$ after dispatching, so the process enters phase 3.

Otherwise, if Y_{t+1} is a type 2 state under event (1) or (2), there is no dispatching in system 2 at step t since Y_t is a type 1 state. By the zigzag property, we know both $(l_t - 1, m_t)$ and $(l_t, m_t + 1)$ are type 2 states, so system 1 will make a dispatch. Then we have $m_{t+1} < m'_{t+1}$, and $X_{t+1} \in P^*$. Thus, the process re-enters phase 3.

When event (3) occurs, Y_{t+1} is the same as Y_t, so we have $m'_{t+1} = m'_t = m_t$. If $(l_t - 1, m_t)$ is a type 1 state, we have $m_{t+1} = m_t = m'_t = m'_{t+1}$, so the process stays at phase 4. On the other hand, if $(l_t - 1, m_t)$ is a type 2 state, system 1 will make a dispatch and go into a type 1 state. Thus, we have $m_{t+1} < m_t = m'_t = m'_{t+1}$ and $X_{t+1} \in P^*$, so the process re-enters phase 3.

When event (4) occurs, Y_{t+1} must be a type 1 state. This is because $X_t = (l_t, m_t)$ is a type 1 state, and $l_t \le l'_{t+1}$ and $m_t \ge m'_{t+1}$ hold, so by Lemma 2, $Y_{t+1} = (l'_{t+1}, m'_{t+1})$ is also a type 1 state. Then, according to the construction of ϕ^{D}, we have $m_{t+1} = m'_{t+1}$ since the number of dispatches under the event is the same in both system 1 and system 2. Thus, if $l_t + m_t < l'_{t+1} + m'_{t+1}$ holds, the process stays at phase 4. Otherwise, if $l_t + m_t = l'_{t+1} + m'_{t+1}$, the next states $X_{t+1} = Y_{t+1}$, so the process returns to phase 1.

In all phases, the revenue rate is identical for both systems, while the penalty rate in system 1 is always less than or equal to that in system 2 at coupled states because system 1 has weakly smaller $l + m$ and m values. Hence, again the objective function satisfies $\tilde{\mathcal{R}}(\phi^*, \lambda^*) \ge \tilde{\mathcal{R}}(\phi^{\mathrm{D}}, \lambda^{\mathrm{D}}) \ge \tilde{\mathcal{R}}(\phi, \lambda)$, which proves the result. □

Proof of Theorem 2. By Theorem 1, when $c_d = c_r$, there exists an optimal dispatching policy ϕ^* that satisfies $\phi^*(l, m) = 0$ for all type 1 states and $\phi^*(l, m) = 1$ for all type 2 states. Let $P^* = \big((l_1, m_1), (l_2, m_2), \cdots\big)$ be the corresponding zigzag path under ϕ^* and $P^*_j := \big((l_1, m_1), \cdots, (l_j, m_j)\big)$ be a sub-path of P^* that terminates at state $(l_j, m_j) \in P^*$. In the following, we show that when $c_d = c_r$, P^*_j yields an objective based on optimal static pricing that is greater than or equal to any other zigzag paths terminating at state (l_j, m_j) when we compare two paths in Algorithm 1 (i.e., $P_{l_j, m_j} = P^*_j$ for $j = 1, 2, \cdots$, where $P_{l,m}$ stores the best zigzag path terminated at (l, m) when Algorithm 1 terminates).

We use induction to show this. Recall that in Algorithm 1, we evaluate the objective values of the two predecessor paths leading into state (l_j, m_j): one arriving from above and the other from the left. For $k = 1, 2, \cdots$, suppose that path P^*_k follows $P_{l_k, m_k} = P^*_k$. This holds trivially for $k = 1$. We want to show $P_{l_{k+1}, m_{k+1}} = P^*_{k+1}$ holds as well. We have the following two cases.

(1) $l_{k+1} = l_k$ and $m_{k+1} = m_k + 1$, that is, $P_{\text{left}} = P^*_{k+1}$ when Algorithm 1 is at (l_{k+1}, m_{k+1}). We show that $\mathcal{R}_{\text{left}} \ge \mathcal{R}_{\text{above}}$. Similar to the proof idea of Theorem 1, we construct a coupling between two systems: system 1 follows path P_{left} and system 2 follows P_{above}. Both systems adopt the same static pricing, and we align their cutoff states so that the cutoff state in system 1 lies on the same anti-diagonal (same $l + m$ value) as the cutoff state in system 2. Under this coupling, the revenue rates are identical when neither system is at its cutoff state. Moreover, for states at the same anti-diagonal, the trip completion rate in system 1 is greater than or equal to that in system 2, so its penalty rate is always no greater than

that of system 2 under coupling. In addition, system 1 never reaches its cutoff state if system 2 is not at its cutoff state. Together, these facts imply that the objective of system 1 is at least as large as that of system 2, i.e., $\mathcal{R}_{\text{left}} \geq \mathcal{R}_{\text{above}}$. If the objectives under the two paths satisfy $\mathcal{R}_{\text{left}} > \mathcal{R}_{\text{above}}$, according to Algorithm 1, we have $P_{l_{k+1},m_{k+1}} = P^*_{k+1}$. Otherwise, if we have $R_{\text{left}} = R_{\text{above}}$, since $(l_{k+1}-1, m_{k+1})$ must be a type 2 state as it is not on the optimal zigzag path P^*, according to Algorithm 1, we have $P_{l_{k+1},m_{k+1}} = P^*_{k+1}$ as well.

(2) $l_{k+1} = l_k + 1$ and $m_{k+1} = m_k$, that is, $P_{\text{above}} = P^*_{k+1}$. Under a similar coupling argument, one can show that $\mathcal{R}_{\text{above}} \geq \mathcal{R}_{\text{left}}$. If the objectives satisfy $\mathcal{R}_{\text{above}} > \mathcal{R}_{\text{left}}$, we have $P_{l_{k+1},m_{k+1}} = P^*_{k+1}$. If the objectives satisfy $\mathcal{R}_{\text{left}} = \mathcal{R}_{\text{above}}$, then $(l_{k+1}-1, m_{k+1})$ must be a type 1 state. Thus, according to Algorithm 1, we have $P^*_{k+1} = P_{l_{k+1},m_{k+1}}$.

Therefore, by induction, Algorithm 1 will set $P_{l_j,m_j} = P^*$ for every $j = 1, 2, \cdots$. Thus, Algorithm 1 yields the same optimal dispatching policy described in Theorem 1. Moreover, the corresponding pricing policy λ^* converges to the optimal solution under value iteration, so the pricing policy is optimal as well. □

Proof of Proposition 1. Consider the threshold $M \geq L\bar{\mu}(p_0 + p_{\max}t_0)/c_r$. First, we notice that for any $l \in \mathcal{L}$ and any $m \geq M$, we have

$$p_0 + p_{\max}t_0 \leq \frac{c_d l + c_r m}{l\mu_{l,m}}. \tag{EC.6}$$

To see why, we start with the chosen threshold condition:

$$M \geq \frac{L\bar{\mu}(p_0 + p_{\max}t_0)}{c_r} \implies \frac{c_r M}{L\bar{\mu}} \geq (p_0 + p_{\max}t_0). \tag{EC.7}$$

Since $m \geq M$ and $\bar{\mu} \geq \mu_{l,m} > 0$, we have

$$\frac{c_d l + c_r m}{l\mu_{l,m}} \geq \frac{c_r M}{l\bar{\mu}} \geq \frac{c_r M}{L\bar{\mu}} \geq (p_0 + p_{\max}t_0) \tag{EC.8}$$

for any $l \leq L$. Now, consider extending the path $P = \big((l_1, m_1), \cdots, (l_I, m_I)\big)$ to some state (l_{I+1}, m_{I+1}) where $m_{I+1} \geq M$. Let the extended path be $P' := \big((l_1, m_1), \cdots, (l_{I+1}, m_{I+1})\big)$. We show that the objective under the optimal dynamic pricing for P' is the same as that for path P. To illustrate this, we examine a cycle of recurrence of the process that starts and ends at (l_I, m_I). Within this cycle, if the first event is a trip completion, then the expected return in this cycle will be the same as that of the original path P starting at (l_I, m_I). On the other hand, if the first event is an effective arrival, the state transitions to (l_{I+1}, m_{I+1}). Then the expected time for the Markov chain to return to state (l_I, m_I) is $1/(l_{I+1}\mu_{l_{I+1},m_{I+1}})$. During this period, the expected accumulated penalty is $(c_d l_{I+1} + c_r m_{I+1})/(l_{I+1}\mu_{l_{I+1},m_{I+1}})$. However, the expected revenue coming from a single arrival is at most $p_0 + p_{\max}t_0$, which is less than or equal to the penalty by equation (EC.6). Thus optimal dynamic pricing on P' will not admit any new rider arrival at state (l_I, m_I).

Clearly, the same argument works for optimal static pricing under path comparison in Algorithm 1: when $m > M$, extending the path from state $(l-1, m)$ or $(l, m-1)$ to state (l, m) will not improve the objective under any static pricing. Therefore, for $m > M$ and $l \in \mathcal{L}$, the objective recorded at each state $R_{l,m}$ satisfies $R_{l,m} = \max\{R_{l,m-1}, R_{l-1,m}\}$. As a consequence, for $m > M$, the following equation holds:

$$\begin{aligned}
R_{L,m} &= \max\{R_{L,m-1}, R_{L-1,m}\} \\
&= \max\{R_{L,m-1}, R_{L-1,m-1}, R_{L-2,m}\} \\
&= \max\{R_{L,m-1}, R_{L-1,m-1}, \cdots, R_{1,m}\} \\
&= \max_{l \in \{1,\cdots,L\}} \{R_{l,m-1}\} \\
&= \max_{l \in \{1,\cdots,L\}} \{R_{l,M}\} = R_{L,M}.
\end{aligned}$$

Thus, extending a path to any state $(l,m), m > M$ cannot produce a better solution. □

Proof of Proposition 2. Given the estimated parameters $\hat{\alpha}_1, \hat{\alpha}_2$ and $\hat{C}$, the predicted service rate at state (l,m) can be written as

$$\hat{\mu}_{l,m} = \frac{1}{\hat{C}(m+1)^{\hat{\alpha}_2}(L-l+1)^{\hat{\alpha}_1} + t_0}.$$

Since the fitted coefficient $\hat{\alpha}_1, \hat{\alpha}_2 < 0$, we know $\hat{\mu}_{l,m}$ is strictly increasing in m and strictly decreasing in l. To verify Assumption 1, we need to show that the differences in the trip completion rate, $l\mu_{l,m+1} - l\mu_{l,m}$ and $l\mu_{l,m} - (l+1)\mu_{l+1,m}$, are both non-increasing in m and non-decreasing in l. Instead of directly proving this, we establish a stronger result: the partial differences $\mu_{l,m+1} - \mu_{l,m}$ and $\mu_{l,m} - \mu_{l+1,m}$ are non-increasing in m and non-decreasing in l. We first perform a variable transformation on l by replacing it with $o = L - l$, where o represents the number of idle drivers. Under this transformation, with a slight abuse of notation, we denote by the service rate $\hat{\mu}(o,m)$. We further extend the domain of o and m to continuous values, with $o \in [0, L]$ and $m \in [0, +\infty]$. In specific, $\hat{\mu}(o,m)$ can be written as

$$\hat{\mu}(o,m) = \frac{1}{\hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1} + t_0}.$$

Since $\hat{\mu}(o,m)$ is continuously differentiable, it suffices to show that its Hessian matrix is element-wise non-positive. More specifically, we need to confirm that:

$$\frac{\partial^2 \hat{\mu}}{\partial o \partial m} \le 0, \quad \frac{\partial^2 \hat{\mu}}{\partial o^2} \le 0, \quad \text{and} \quad \frac{\partial^2 \hat{\mu}}{\partial m^2} \le 0.$$

Due to the symmetry in o and m, it is enough to verify $\frac{\partial^2 \hat{\mu}}{\partial o^2} \le 0$ and $\frac{\partial^2 \hat{\mu}}{\partial o \partial m} \le 0$.

$$\frac{\partial^2 \hat{\mu}}{\partial o \partial m} = -\frac{\hat{C}\hat{\alpha}_2\hat{\alpha}_1 \left(t_0 - \hat{C}(o+1)^{\hat{\alpha}_1}(m+1)^{\hat{\alpha}_2}\right)(m+1)^{\hat{\alpha}_2 - 1}(o+1)^{\hat{\alpha}_1 - 1}}{\left(t_0 + \hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1}\right)^3}, \tag{EC.9}$$

$$\frac{\partial^2 \hat{\mu}}{\partial o^2} = -\frac{\hat{C}\hat{\alpha}_1 \left(t_0(\hat{\alpha}_1 - 1) - \hat{C}(1+\hat{\alpha}_1)(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1}\right)(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1 - 2}}{\left(t_0 + \hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1}\right)^3}. \tag{EC.10}$$

Since $\hat{\alpha}_1 < 0$, $\hat{\alpha}_2 < 0$, and $t_0 - \hat{C}(o+1)^{\hat{\alpha}_1}(m+1)^{\hat{\alpha}_2} \geq t_0 - \hat{C} \geq 0$, it follows from equation (EC.9) that $\frac{\partial^2 \hat{\mu}}{\partial o \partial m} \leq 0$. Next, for the second derivative with respect to o, we use the fact that $\hat{\alpha}_1 < 0$ and we have:

$$
\begin{aligned}
& t_0(\hat{\alpha}_1 - 1) - \hat{C}(1+\hat{\alpha}_1)(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1} \\
= & - t_0 - \hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1} + \hat{\alpha}_1 \left(t_0 - \hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1} \right) \\
\leq & \hat{\alpha}_1 \left(t_0 - \hat{C}(m+1)^{\hat{\alpha}_2}(o+1)^{\hat{\alpha}_1} \right) \\
\leq & 0,
\end{aligned}
$$

which implies that $\frac{\partial^2 \hat{\mu}}{\partial o^2} \leq 0$ from equation (EC.10). This completes the proof. □