Dynamic Restaurant Ranking as Demand Control in On-Demand Meal Delivery Systems

Dongyang Xia*[1], Yousef Maknoon[2], Hai Wang[3], Niels Agatz[4], and Shadi Sharif Azadeh[1]

[1]Department of Transport & Planning, Delft University of Technology, Netherlands
[2]Department of Transport & Planning, Delft University of Technology, Netherlands
[3]School of Computing and Information Systems, Singapore Management University, Singapore
[4]Rotterdam School of Management, Erasmus University, Netherlands

Short summary

On-demand meal delivery (OMD) platforms influence customer choices by how restaurants are ranked in their app interface. Dynamically adjusting these rankings offers a powerful tool for real-time demand shaping. At the same time, ensuring efficient deliveries and fair courier workload distribution requires the platform to make coordinated decisions across both demand and supply sides, including restaurant ranking and courier routing. Since demand-shaping affects order distribution—which in turn impacts courier availability—jointly optimizing these decisions can enhance overall system performance. We propose a multi-stage optimization framework that determines restaurant rankings and courier routing decisions throughout the day while accounting for endogenous customer demand, congestion effects, the order bundling strategy, and the workload fairness of couriers. Equity is incorporated into dispatch through workload-based dominance constraints. We derive the structural properties that characterize when ranking adjustments improve a congestion-aware objective and use these properties to guide a learning-to-optimize algorithm. The proposed algorithm uses offline-trained machine learning models to predict whether a candidate solution will improve the objective and by how much. Dozens of candidate solutions are evaluated in parallel, without relying on commercial solvers. This enables real-time decision-making at scale. Additionally, equity-aware order-courier matching and couriers' routing decisions are constructed using a fast insertion-based algorithm with dominance rules that balance efficiency and courier workload fairness. Computational experiments on a city-scale meal delivery network covering the entire city of Delft, the Netherlands, show that dynamic ranking considerably improves system performance relative to fixed-ranking and non-learning benchmarks, particularly under constrained courier capacity. Our method produces higher-quality solutions within seconds, and achieves at least a 7.7% improvement in the platform's profit over the benchmark heuristic algorithm.
Keywords: On-demand delivery; Endogenous demand; Congestion effects; Equity-aware dispatch; Real-time optimization; Machine learning

1 INTRODUCTION

Urban logistics platforms increasingly rely on time-sensitive delivery, with on-demand meal delivery (OMD) a leading example. Despite rapid growth, platforms operate on thin margins and must balance consumers, restaurants, and couriers. With limited short-term flexibility in consumer prices and courier pay, platforms cannot quickly rebalance demand and supply through market levers. They therefore rely on operational controls to protect service levels. This pushes time pressure and workload variability onto couriers through higher assignment intensity, reduced slack, and more complex routes. This triggers sustainability and fairness concerns. Addressing these challenges requires coordinated control of customer demand and delivery routes in order to enhance operational efficiency while ensuring equitable workload assignment.

Prior work typically treats demand- and supply-side levers in isolation. Demand management (e.g., dynamic pricing and service-area adjustments) reshapes when and where orders arise (Ulmer et al., 2022; Bahrami et al., 2023; Yang et al., 2024), while supply management (e.g., courier routing) focuses on on-time fulfillment (Liu & Luo, 2023). This separation aids tractability but leaves critical gaps. Developed independently, these approaches miss feedback related to platform

policies, customer choices, and courier operations; often assume that demand is exogenous; and pay limited attention to fairness.

Motivated by these gaps, we propose a system-level optimization framework that coordinates demand control and supply control in real time. On the demand side, the platform uses dynamic restaurant ranking (i.e., the order of restaurants displayed on the app) to guide customer choices. On the supply side, it employs equity-based order-courier matching and courier routing strategies to manage resource allocation. The framework explicitly captures how ranking decisions affect demand realization, platform congestion, and workload equity among couriers. Our contributions are threefold. (i) We begin by deriving a set of general structural properties that characterize the interactions between restaurant rankings and key system outcomes that include platform profit, customer choice behavior, congestion effects, and workload equity. These properties offer theoretical insights into how ranking shapes both demand distribution and courier routes, and they inform the design of our modeling framework. (ii) We develop a multistage optimization model in which restaurant rankings and customer choice are jointly determined at each decision-making stage while accounting for congestion effects. Given these outcomes, the model then optimizes equity-based order-courier matching and courier routing decisions. (iii) We propose a learning-to-optimize algorithm. The algorithm efficiently evaluates dozens of candidate solutions in parallel within milliseconds without relying on commercial solvers, which enables real-time solution at scale.

2 PROBLEM SETTING

We address a real-time decision-making problem for OMD platforms, in which the platform jointly optimizes restaurant ranking, customer choices, order-courier matching, and courier routing while considering workload fairness. The service area (e.g., a city) is partitioned into a set of geographic districts $\mathcal{D}$, with each further subdivided into multiple service zones. The business day is discretized into a sequence of decision-making stages $\mathcal{S}$, each with a fixed duration δ (unit: min). The multistage decision-making problem involves four key stakeholders: the platform, customers, restaurants, and couriers. At each decision-making stage, the platform determines restaurant rankings based on carried-over courier availability, and demand is realized in response to these rankings. Although the platform does not control demand directly, the two interact because the ranking influences customer choices. These ranking and demand decisions together shape the congestion effects, which in turn inform subsequent dispatching and future ranking strategies. The congestion effects are evaluated based on demand realizations in terms of customer choices and available delivery capacity at each decision stage for each service zone.

3 METHODOLOGY

We formulate a multistage decision-making formulation that integrates customer choice behavior, congestion effects, and fairness considerations into a unified optimization framework. The proposed formulation supports a real-time decision-making framework. At the beginning of each stage, the developed flexible ranking optimization model determines restaurant rankings, which drives customer choices. The resulting demand realizations are then used in the equity-based order-courier matching and routing models, under an order-bundling strategy in which each courier may serve multiple orders per trip. Updated courier availability is passed forward as input to the next stage's ranking model. This feedback mechanism enables dynamic and adaptive coordination of demand and supply decisions throughout the operational horizon.

To solve this problem, we develop a real-time decision-making framework. The proposed solution method is repeated at every decision-making stage within a rolling horizon framework, which creates a feedback loop among the flexible ranking, equity-based order-courier matching, and courier routing decisions. The developed matching method considers both spatial proximity and the cumulative workload of couriers. By continuously feeding the latest information, including updated courier availability and newly arrived orders, into the current stage, the proposed solution framework adapts to real-time operational conditions. Specifically, at each stage, the proposed flexible ranking optimization model is first solved using a learning-to-optimize algorithm. The outputs are then passed to the order-courier matching and courier routing component, which is solved using the routing model in combination with the greedy insertion algorithm with the equity-based matching strategy. After these decisions have been made, the updated number of available couriers is carried

forward as input to the flexible ranking optimization model for the next decision-making stage.

4 RESULTS

To evaluate the performance and practical value of our mathematical formulations and real-time decision-making algorithms, we conduct numerical experiments based on realistic instances derived from data for the city of Delft, the Netherlands. The dataset incorporates all restaurants and the actual geographic layout of the city. Delft is partitioned into 18 regular hexagonal service zones, each with an edge length of 1 km. Each restaurant is mapped to a zone.

Finding 1. *Dynamically adjusting restaurant rankings yields substantial efficiency gains in both order fulfillment and OMD platform revenue, with the greatest benefits observed under ordered attraction and impatient customer scenarios.*
We assess the benefits of implementing a flexible restaurant ranking strategy by comparing it with a fixed-ranking benchmark. Table 1 reports the relative improvement in the number of fulfilled orders and platform revenue compared with the fixed-ranking benchmark across four instances, in which the parameter $a = 1$ indicates that customers have sufficient patience to view only the first page displayed, while $a = 100$ implies that customers browse all available pages, following the settings in Gao et al. (2021). In particular, we compare unordered (U) and ordered (O) attraction strategies. Under ordered attraction, the flexible ranking strategy yields substantial gains; these include over 60% more fulfilled orders and more than 24% improvement in revenue. These effects are especially pronounced when customers are less patient (i.e., $a = 1$). By contrast, when attraction scores are randomly ordered (U), the benefits of flexible ranking are smaller but still notable. These results demonstrate the value of dynamically adapting rankings to manage demand more effectively, particularly when customer preferences can be steered to align with platform objectives.

Table 1: Improvements in the number of fulfilled orders and revenue achieved by the flexible-ranking strategy compared with the fixed-ranking strategy.

Instance (Att., # of couriers, a)	Increase in fulfilled orders (%)	Increase in revenue (%)
(O, 140, 1)	60.4	30.4
(O, 140, 100)	68.7	24.6
(U, 140, 1)	13.7	4.2
(U, 140, 100)	9.9	3.0

Finding 2. *Incorporating congestion effects—defined as the ratio of order quantity to driver capacity—in the flexible ranking strategy balances spatial demand distribution, alleviates delivery pressure in highly congested zones, and enhances overall network efficiency.*
We further evaluate the effectiveness of the proposed flexible ranking strategy that incorporates congestion effects. Table 2 summarizes the results. Across all instances, the strategy reduces average click-to-door time by 7.7%-26.9% but lowers the platform's revenue by 0.6%-7.2%, which reveals a trade-off in which efficiency gains usually dominate. The benefits are robust across attraction models, patience levels, and fleet sizes, and are largest in constrained settings.
Managerial insights. Our analyses of the results yield four insights. First, allowing flexible ranking dominates fixed benchmarks. Second, incorporating real-time congestion smooths demand across zones and cuts click-to-door time, especially with limited fleets, and thus improves service quality without heavy resource expansion. Third, including fairness in matching and routing is beneficial, but overemphasis can be costly. Lastly, integrating learning into solution methods materially helps real-time operations. The proposed learning-to-optimize algorithm considerably improves the objective under tight time limits. Detailed computational results and managerial insights will be presented at the conference, including analyses of structural properties, formal analyses, performance of our algorithm, and results across various instances.

References

Bahrami, S., Nourinejad, M., Yin, Y., & Wang, H. (2023). The three-sided market of on-demand

Table 2: Comparison of average click-to-door time and platform revenue with and without incorporating congestion effects.

Instance ($\|\mathcal{C}\|$, a)	Ave. Ctd (min)	Ave. Ctdc (min)	Dev^c (%)	Rev. (€)	Revc. (€)	Dev^r (%)
(150, 1)	27.4	25.3	**7.7**	16,084.2	14,934.1	**-7.2**
(140, 1)	30.8	27.7	**10.1**	15,900.0	15,600.0	**-1.7**
(130, 1)	38.9	35.4	**9.0**	15,592.7	15,496.2	**-0.6**
(150, 100)	30.0	27.1	**9.7**	16,639.7	15974.1	**-4.0**
(140, 100)	34.4	30.2	**12.2**	16,621.8	15,800.3	**-4.9**
(130, 100)	39.0	35.8	**8.2**	16,365.3	15,678.2	**-4.2**

Abbreviations: $|\mathcal{C}|$ = # of couriers; Ave. Ctd = Average click-to-door time without congestion; a = Customers' patience; Ave. Ctdc = Average click-to-door time with congestion; Dev^c = Relative deviation in click-to-door time; Rev. = Revenue without congestion; Revc. = Revenue without with congestion; Dev^r = Relative deviation in revenue.

delivery. *Transportation Research Part E: Logistics and Transportation Review*, *179*, 103313.

Gao, P., Ma, Y., Chen, N., Gallego, G., Li, A., Rusmevichientong, P., & Topaloglu, H. (2021). Assortment optimization and pricing under the multinomial logit model with impatient customers: Sequential recommendation and selection. *Operations research*, *69*(5), 1509–1532.

Liu, S., & Luo, Z. (2023). On-demand delivery from stores: Dynamic dispatching and routing with random demand. *Manufacturing & Service Operations Management*, *25*(2), 595-612.

Ulmer, M. W., Erera, A., & Savelsbergh, M. (2022). Dynamic service area sizing in urban delivery. *OR Spectrum*, *44*, 763–793.

Yang, J., Lau, H. C., & Wang, H. (2024). Optimization of customer service and driver dispatch areas for on-demand food delivery. *Transportation Research Part C: Emerging Technologies*, *165*, 104653.