Multimodal Transportation Pricing Alliance Design: Large-Scale Optimization for Rapid Gains

Kayla Cummings[1], Vikrant Vaze[2], Özlem Ergun[3], and Cynthia Barnhart[4]

[1]Massachusetts Institute of Technology
[2]Thayer School of Engineering at Dartmouth
[3]Northeastern University
[4]Massachusetts Institute of Technology

Keywords: Public transit, pricing, alliance design, mixed-integer non-convex optimization

1 INTRODUCTION AND CONTRIBUTIONS

Cities face critical challenges in their quest to improve urban mobility. Recent declines in transit ridership demonstrate the inability of transit's static infrastructure to accommodate rapidly evolving commuting patterns. Transportation Network Companies (TNCs) annually transport billions of passengers, but the majority of urban TNC patrons admit that they would have otherwise walked, biked, taken public transit, or not made the trip, coinciding with tens of millions in annual transit revenue losses, worsening congestion, higher emissions, lower navigability of cities, and reduced accessibility to affordable public transportation options. Mobility-on-demand (MOD) services have the potential to service *transit deserts* — low-density areas disconnected from public transit. However, cost presents a key barrier: MOD services administered by transit agencies incur the highest average per trip costs. High labor needs, outdated technology, and coordination difficulties lead to inefficient and expensive operations. The average TNC trip costs $13, a full $10 less than the agency-sponsored MOD trips. Outsourcing all 223 million on-demand transit trips to TNCs could hypothetically save billions of dollars for US transit agencies. Thus, pricing alliances between TNCs and transit agencies have the potential to improve service quality and coverage, while reducing costs and decreasing city-wide vehicle-miles traveled (VMT).

We propose a prescriptive pricing alliance to enable such alliances between transit agencies and ride-sharing operators. We formulate a fare setting model to maximize total benefits across the integrated network. We demonstrate how our framework could help operators navigate competing alliance objectives: (1) improving access to high-quality transportation options for underserved populations, (2) reducing emissions and congestion from single-occupancy vehicle trips, and (3) maintaining the financial well-being of participating operators and incentivizing profit-oriented operators to participate. A key technical challenge lies in capturing interdependence between fares and passenger choice. Our model integrates a discrete choice model of passengers' mode and route decisions based on prices and non-pricing attributes like travel times.

Our fare-setting model is a large-scale, mixed-integer, non-convex optimization problem — a challenging class of problems. Our first contribution is to design a two-stage decomposition in which the first-stage pricing decisions parameterize second-stage fare discounts and passengers' travel choices. The second stage becomes a more tractable mixed-integer linear optimization problem that can be solved with commercial solvers. To solve the first-stage model, we develop a new solution approach that combines tailored coordinate descent, parsimonious second-stage evaluations, and interpolations using Special Ordered Sets of type 2 (SOS2). We also develop acceleration techniques based on slanted coordinate traversal and search-direction randomization. This solution approach is our second technical contribution. It is applicable to any two-stage formulation with a low-dimensional, convex, continuous first-stage solution space and any computationally expensive black-box second stage. This solution approach has been found to significantly improve the outcomes, for passengers and operators, compared to those obtained with state-of-the-art benchmarks based on Bayesian Optimization (BO) (Liu *et al.*, 2019).

From a practical standpoint, for a large-scale case study in the Greater Boston area, our model sets fares in realistic ranges and with interpretable connections to alliance goals. An alliance with a greater focus on minimizing total VMT prefers flat fares rather than distance-varying fares to increase system utilization by long-distance commuters. On the other hand, alliances with a greater emphasis on increasing transit access set discounts with greater geographic variation to make alliance routes more attractive to heterogeneous populations. The clear alignment between operator goals and passenger choices achieved by our fare structures illustrates the value of modeling passenger choice. Analysis of our results also shows that the model is appropriately responsive to equity-oriented objectives. It sets lower fares for, and increases utilization by, low-income and long-distance commuters. Compared to noncooperative pricing, our fares and our tailored profit allocation mechanism together incentivize profit-oriented MOD operators not only to participate in the alliance but also to adopt the transit operator's priorities.

2 METHODOLOGY

$$\max_{\boldsymbol{x},\boldsymbol{\beta},\boldsymbol{\mu},\Lambda} \quad \text{Weighted Sum of Passenger Utility, Operator Profit, and VMT Reduction} \tag{1}$$

$$\begin{aligned}
\text{s.t.} \quad & \text{Route Fare} = \text{Base} + \text{Markup} \times \text{Distance} && \forall \text{ Discount-Ineligible Routes} && (2)\\
& \text{Route Fare} = (1 - \Lambda \boldsymbol{x}) \times \text{Base} + \text{Markup} \times \text{Distance} && \forall \text{ Discount-Eligible Routes} && (3)\\
& \text{Route Share} = \frac{e^{\text{Route Utility } - \text{ Fare}}}{e^{\text{OO Utility}} + \sum_{\text{Routes}} e^{\text{Route Util. } - \text{ Fare}}} && \forall \text{ Routes, Time Periods} && (4)\\
& \text{OO Share} = \frac{e^{\text{OO Utility}}}{e^{\text{OO Utility}} + \sum_{\text{Routes}} e^{\text{Route Util. } - \text{ Fare}}} && \forall \text{ Markets, Time Periods} && (5)\\
& \text{Lower bounds} \leq \boldsymbol{\beta}, \boldsymbol{\mu}, \Lambda \leq \text{Upper bounds}, \boldsymbol{x} \text{ binary} && \forall \text{ Operators} && (6)
\end{aligned}$$

The fare-setting model (1)-(6) jointly sets the fare parameters (base ($\boldsymbol{\beta}$) and distance-based markups ($\boldsymbol{\mu}$)), discount level (Λ), and binary discount activation decisions ($\boldsymbol{x}$) across the integrated network to maximize a weighted sum of passenger benefits, operator profits, and the negative of vehicle miles traveled (VMT). Discounts are applied on selected routes (Constraints (2) and (3)) and utility-maximizing passengers choose between the various available routes and the outside option (OO) according to a multinomial logit model (Constraints (4) and (5)). The fare parameters and discount level obey bounds and the discount activation decisions are binary (Constraints (6)). It is a non-convex mixed-integer nonlinear optimization problem. We decompose the formulation to tractably obtain high-quality solutions for practically-sized problems (with tens of thousands of variables and hundreds of thousands of constraints). The decomposition approach separates the two main sources of model complexity, namely, non-convexity and the presence of discrete variables. The first stage of the decomposition handles the non-convexity but relegates discrete variables to the second stage. First-stage pricing decisions parameterize second-stage discount activations and passenger choice, which makes the first-stage solution space a small-dimensional, continuous, convex region. The second stage can then be formulated as a more tractable mixed-integer linear optimization problem (MILOP).

2.1 Solution Approach

The first-stage objective function's gradients are inaccessible, which eliminates gradient-based approaches. Low-dimensional, convex first-stage space is amenable to *coordinate descent*, which takes turns fixing all fare parameters except one and greedily optimizing along the free dimension. But a one-dimensional search is also difficult for a search space comprising a continuous spectrum of optimal MILOP solutions. Although the second-stage problem solves fast enough to be a useful tool, it is slow enough to warrant a judicious selection of first-stage valuation points. Our tailored coordinate descent scans each search direction by solving an auxiliary model that approximates the objective function in that search direction with Special Ordered Sets of type 2 (SOS2). It terminates when no improvements are found in the coordinate directions. We call this new procedure as *SOS2 Coordinate Descent (SOS2-CD)*. We also develop three acceleration strategies that build upon SOS2-CD. The first one randomizes search direction ordering. The second one exploits the fact that the SOS2 approximation is valid along *any search direction* through the first-stage problem's solution space, not just those parallel to coordinate axes. Natural search direction candidates are those which jointly vary an operator's base fare and markup while holding other parameters constant. Given a pair of dimensions, this strategy randomly selects the spanning dimension and then selects the line's slope in this 2D plane uniformly at random from the set of lines that intersect the current solution and span the selected dimensions. Finally, it drops anchors at evenly spaced points along the chosen line and obtains the next candidate solution maximizing the auxiliary objective function. SOS2-CD over such slanted directions unlocks directions that navigate trade-offs between high-base-low-markups vs. low-base-high-markups. Finally, we use BO warm starts to mitigate SOS2-CD's sensitivity to random initializations.

2.2 Profit Allocation Mechanism

When considering a pricing alliance, an operator assesses whether cooperation would improve their prioritized system-wide metrics compared to noncooperation. The MOD operator maximizes its own profit, while the transit agency may maximize a linear combination of multiple system-wide metrics. By entering a pricing alliance, the transit agency is guaranteed to fare no worse than that under the noncooperative setting. However, the MOD operator's participation depends on whether the alliance participation increases its profit. We design a profit allocation mechanism that guarantees the MOD operator's participation. Because the alliance's priorities may include benefits to passengers and/or to the rest of society (as reduced VMT), the alliance may earn less profit than the combined profits of the operators in the noncooperative setting. Under the profit allocation mechanism, the transit operator guarantees that the profit-oriented MOD operator earns at least as much as it would have earned outside of the alliance.

3 RESULTS

Table 1 compares performance for three time budgets and 50 trials each and provides statistics on the number of search trajectories. All versions of our approach under all time budgets outperform the BO benchmark on average, and especially so for the versions with acceleration strategies. The larger time budgets allow for accelerated SOS2-CD to offer robust performance. In general, the SOS2-CD trajectories with multidimensional search (that is, SOS2-CD-MD and SOS2-CD-MD-R) converge more quickly, allowing more trajectories to be computed in a given time limit.

3.1 Insights from Practical Case Study

Our results confirm that our model yields interpretable output with prices in realistic ranges. The minimum, average and maximum real-world route prices in the service region are \$4, \$10.23, and \$54.46, respectively, while those given by our model are in the range \$0-\$60.37. We compare the allied system ridership across priority regimes and towns under fares corresponding to base (i.e. income-agnostic) and refined (i.e. income-aware) choice models. In a passenger-oriented regime, the income-aware model increases the system ridership by 13% on average for towns with below-average median income compared to an average increase of less than 2% for towns with above-average median income. Although ridership increases across the board due to generally lower fares, the greatest increases occur in the two towns with the lowest median household

Table 1 – *Objective function statistics (surplus compared to the 1-hour BO performance).*

Time budget	Algorithm	Trajectories			Objective (Thousand USD)		
		Min	Avg	Max	Min	Avg	Max
1 hour	BO	-	-	-	−28.0	0*	18.6
	SOS2-CD	2.0	3.8	5.0	−103.8	13.5	24.1
	SOS2-CD-MD	2.0	4.0	7.0	**−5.7**	**19.3**	24.1
	SOS2-CD-R	2.0	3.8	5.0	−77.6	17.1	**24.4**
	SOS2-CD-MD-R	2.0	4.2	6.0	−20.5	19.0	23.7
6 hours	BO	-	-	-	7.7	17.9	23.9
	SOS2-CD	16.0	20.1	25.0	−2.7	22.3	24.3
	SOS2-CD-MD	15.0	21.3	28.0	22.0	23.3	24.1
	SOS2-CD-R	13.0	19.7	24.0	**22.4**	**23.8**	24.4
	SOS2-CD-MD-R	15.0	21.4	26.0	21.1	23.3	24.3
12 hours	BO	-	-	-	18.7	21.9	23.9
	SOS2-CD	32.0	39.6	50.0	22.0	23.7	24.3
	SOS2-CD-MD	32.0	42.4	54.0	22.8	23.7	24.2
	SOS2-CD-R	26.0	38.2	47.0	**23.5**	**24.0**	24.4
	SOS2-CD-MD-R	31.0	42.7	52.0	22.2	23.7	24.4

incomes. The lower middle-income bracket sees the next-highest ridership increase. In the profit-oriented regime, ridership increases across all towns as higher profit can be achieved with lower fares and higher volume. The above-average income towns gain 10% in average ridership, while the below-average income towns gain 37%. When prioritizing VMT, we obtain significantly larger ridership increases for longer distance passengers. In particular, the five towns farthest from the inner city are the only ones with ridership increases. Their average increase is ∼25% while the remaining 10 towns see an average 7% ridership drop. Thus, the VMT-focused regime increases access to the allied network for commuters who are farther from the inner city.

3.2 Quantifying the Value of Cooperation

The profit allocation mechanism ensures that the MOD operator receives its noncooperative earnings, despite the lower MOD fares that the alliance sets to achieve lower VMT or higher passenger benefits. This in turn reduces the transit's profit allocation as high VMT or low passenger benefits are increasingly penalized. On the other hand, the transit operator always strictly improves its objective of optimizing total system-wide performance. As a result, the MOD operator interestingly finds it in its interest to adopt transit's priorities as transit increasingly diverges from profit maximization. In other words, *the profit-maximizing MOD operator would not prefer a profit-maximizing alliance.* In fact, the MOD operator would benefit most from total altruism on the transit side (an exclusive focus on passenger utility or VMT reduction). Passengers also benefit from strictly lower prices and increased system utilization.

4 DISCUSSION

We contribute a pricing alliance design framework to enable incentive-aligned collaboration between transit agencies and MOD operators. It captures the interdependent decisions of passengers and operators. We develop a tractable two-stage formulation equivalent to the original mixed-integer non-convex model, and solve it with a new tailored SOS2 coordinate descent approach. Our approach consistently and significantly outperforms benchmarks, generating daily system-wide benefits worth tens of thousands of dollars. Our framework aligns profit-oriented MOD operators with transit goals of passenger utility and VMT reduction. Ultimately, cooperative pricing results in win-win-win outcomes for passengers, MOD operators, and transit agencies.

References

Liu, Yang, Bansal, Prateek, Daziano, Ricardo, & Samaranayake, Samitha. 2019. A framework to integrate mode choice in the design of mobility-on-demand systems. *Transp. Research Part C*, **105**, 648 – 665.