Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

The Opt-In Radius Auction: An Exact Two-Sided Hypergraph- VCG Mechanism for Multi-Passenger Ride-Sharing

Mario Verdyguer Cuevas
Universidad Politécnica de Madrid
Shailendra Kumar Tiwari
Manipal Academy of Higher Education
Sandhya Parasnath Dubey
sandhya.dubey@manipal.edu

Manipal Academy of Higher Education

Research Article

Keywords: Ride-sharing, mechanism design, Vickrey–Clarke–Groves mechanism, combinatorial optimization, hypergraph matching, intelligent transportation systems

Posted Date: November 12th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11409488/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

The Opt-In Radius Auction: An Exact Two-Sided Hypergraph-VCG Mechanism for Multi-Passenger Ride-Sharing

Mario Verdyguer Cuevas[1], Shailendra Kumar Tiwari[2], and Sandhya Parasnath Dubey[2]

[1]Escuela Técnica Superior de Ingenieros de Telecomunicación, Universidad Politécnica de Madrid, 28040 Madrid, Spain (ORCID: 0009-0001-4428-3538)

[2] Manipal Institute of Technology, Manipal Academy of Higher Education, Manipal 576104, India (ORCID: 0000-0002-8350-2676)

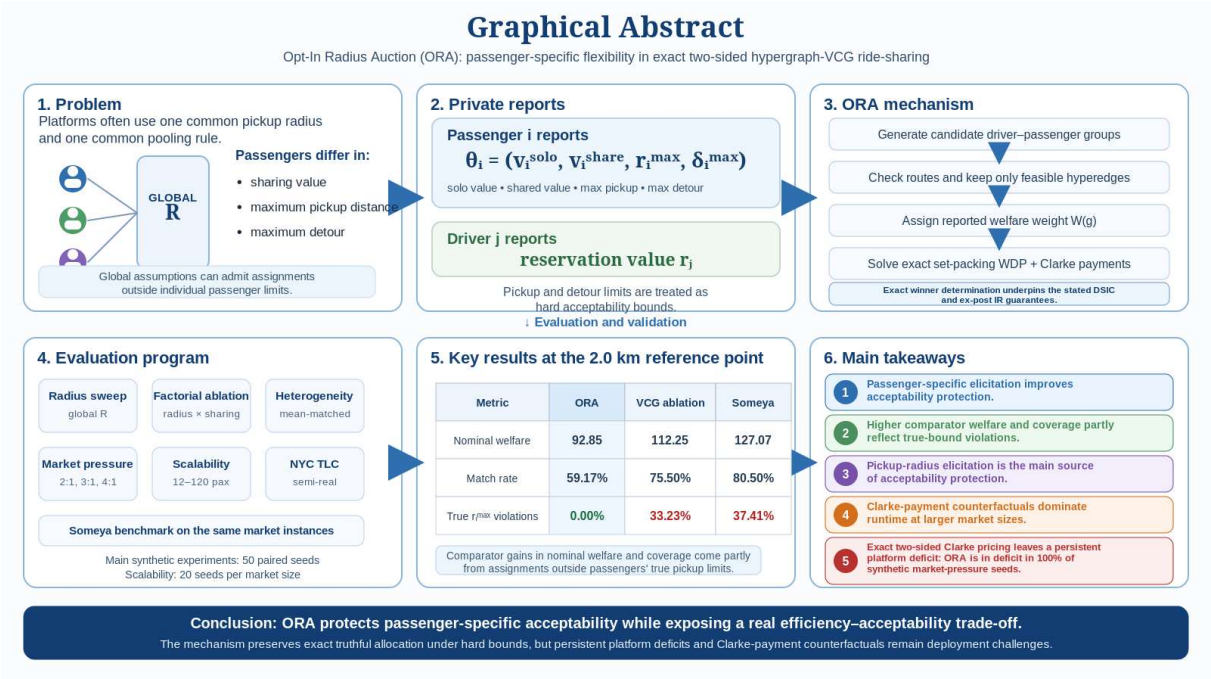
Corresponding author: Sandhya Parasnath Dubey (e-mail: sandhya.dubey@manipal.edu).

Abstract

Ride-sharing platforms usually apply a single pickup radius to every passenger and price pooled trips with one common rule, although passengers differ in how far they will travel to be picked up and in how much they value sharing. This paper proposes the Opt-In Radius Auction (ORA), an exact two-sided hypergraph Vickrey–Clarke–Groves (VCG) mechanism in which each passenger reports a solo valuation, a shared-ride valuation, a maximum pickup distance, and a maximum detour, and each driver reports a private reservation value. Under exact winner determination, quasi-linear utilities, and genuine hard acceptability bounds, ORA is dominant-strategy incentive compatible and ex-post individually rational on both sides; budget balance is not imposed and the platform deficit is reported as an outcome. The mechanism is evaluated on paired synthetic markets and on semi-synthetic Manhattan demand estimated from 15.15 million New York City Taxi and Limousine Commission trip records. At a 2.0 km global radius, a VCG ablation raises nominal welfare from 92.85 to 112.25 and the match rate from 59.17% to 75.50%, but serves 33.23% of its passengers beyond their true pickup limits, against 0% for ORA. The computational bottleneck at larger scale is the Clarke-payment counterfactuals rather than the allocation itself.

Keywords Ride-sharing, mechanism design, Vickrey–Clarke–Groves mechanism, combinatorial optimization, hypergraph matching, intelligent transportation systems.

Graphical abstract



1 Introduction

Multi-passenger ride-sharing requires a platform to make several coupled decisions within each clearing window: which requests to serve, which passengers to group, which driver to assign, which route to operate, and how transfers should be determined. These decisions affect operational efficiency, passenger waiting and detour, driver participation, and the platform budget. The problem is therefore not only one of routing or assignment. Passengers may value solo and shared service differently and may accept different pickup distances or detours, while drivers may have private participation costs. A mechanism for this setting must determine both an allocation and a transfer rule in the presence of strategic agents.

Many ride-sharing models simplify passenger flexibility through platform-wide parameters, such as a common pickup radius or a common pooling discount. These assumptions are convenient, but they suppress heterogeneity that can affect feasibility and welfare. Passengers requesting similar trips may still differ substantially in their willingness to share, in the pickup distance they accept, and in the additional in-vehicle travel they tolerate. Increasing a common pickup radius creates more feasible driver-passenger combinations and can improve nominal matching or welfare. However, some of these additional assignments may exceed individual passenger limits.

Evaluating only aggregate matching or nominal welfare can therefore conceal an acceptability problem.

Existing work addresses several components of this setting. Exact hypergraph Vickrey–Clarke–Groves (VCG) formulations combine multi-passenger matching with welfare-based pricing [1]. Two-sided auction mechanisms emphasize strategy-proofness, individual rationality, and stronger budget properties [2], while approximately efficient mechanisms trade exact welfare maximization for improved financial or computational guarantees [3]. Other studies focus on scalable assignment [4], strategic driver costs [5], or fairness-aware dispatch [6]. These contributions provide important tools, but passenger-specific sharing valuations, pickup limits, and detour limits are usually not elicited jointly as components of a strategic passenger type while the driver simultaneously reports a private participation cost.

The Opt-In Radius Auction (ORA) addresses this setting as a static, batch-cleared direct-revelation mechanism. Each passenger reports a solo valuation, a shared-ride valuation, a maximum pickup distance, and a maximum detour. Each driver reports a private reservation value, while per-kilometer operating cost is observable in the base model. ORA represents candidate allocations as driver-passenger-mode hyperedges and retains only groups for which a feasible route satisfies all reported pickup and detour constraints. It then solves an exact maximum-weight set-packing problem. Passenger payments and driver earnings are obtained from Clarke pivot counterfactuals. Under exact welfare maximization, quasi-linear utilities, and genuine hard passenger acceptability bounds, ORA is dominant-strategy incentive compatible (DSIC) and ex-post individually rational (IR) on both sides; platform budget balance is not guaranteed. What is new here is the combination: passenger pickup distance, detour tolerance, and sharing valuation are elicited jointly as private information alongside a private driver reservation value. Pickup and detour limits are enforced as hard feasibility constraints, solo and shared valuations enter the welfare objective, and exact Clarke pivot transfers determine payments on both sides. The mechanism is evaluated using an explicit acceptability metric on synthetic markets and semi-synthetic Manhattan demand estimated from 15.15 million trip records. Individual pieces of this appear in earlier work; the joint treatment does not.

The experiments are designed to isolate the effect of passenger-specific information rather than to compare mechanisms with different optimization objectives. ORA is

therefore evaluated against controlled ablations that preserve the same exact winner determination and payment structure while replacing individual pickup and sharing information with platform-wide assumptions. A factorial ablation separates the contribution of pickup-radius elicitation from that of sharing-value elicitation. Additional experiments study heterogeneity, market pressure, computational scalability, and a semi-synthetic validation derived from New York City Taxi and Limousine Commission (NYC TLC) demand. The results consistently show an efficiency-acceptability trade-off: global assumptions can increase nominal welfare and matching, but part of that gain is associated with assignments outside passengers' true pickup constraints.

The contributions of this paper are the following.

1) *Joint two-sided elicitation of flexibility.* Pickup distance, detour tolerance, and sharing valuation are elicited as components of a single strategic passenger type, while the driver simultaneously reports a private reservation value. Existing mechanisms elicit some of these, but not all of them jointly on both sides of the market.

2) *Hard acceptability bounds inside the mechanism.* Reported pickup and detour limits determine which hyperedges exist at all, not only how they are weighted. Feasibility, and not a penalty term, carries the passenger incentive result.

3) *Exact Clarke pricing on both sides.* Winner determination is solved exactly and transfers come from bilateral Clarke pivot counterfactuals, giving dominant-strategy incentive compatibility and ex-post individual rationality for passengers and drivers under the stated assumptions.

4) *An explicit acceptability metric.* Allocations produced under assumed passenger types are scored against the passengers' true types, which separates nominal welfare from welfare that is actually implementable. This makes visible a cost that aggregate matching statistics hide.

5) *A factorial ablation that separates the two elicited components.* Pickup-radius elicitation and sharing-value elicitation are varied independently over 50 paired seeds, showing that the acceptability protection comes almost entirely from the pickup component.

6) *Validation on semi-synthetic Manhattan markets.* Demand timing, origin–destination structure, and inter-zone distances are estimated from 15.15 million NYC TLC

high-volume for-hire vehicle trips, and the efficiency–acceptability trade-off persists outside the synthetic grid.

7) *Explicit reporting of the two costs of exactness.* The platform deficit and the runtime decomposition are reported as results rather than assumed away, identifying the Clarke counterfactuals as the computational bottleneck.

The remainder of this paper is organized as follows. Section 2 reviews incentive-compatible ride-sharing mechanisms, operational allocation methods, and fairness-oriented work, and positions ORA against them. Section 3 presents the market model, the type space, the feasibility and routing rules, exact winner determination, Clarke payments, and the incentive and budget properties. Section 4 describes the experimental design, including the comparators, the four experiments, the outcome measures, and the NYC TLC pipeline. Section 5 reports the results, Section 6 discusses what they imply for mechanism design and for platform operation, Section 7 states the limitations and the open questions, and Section 8 concludes.

2 Related Work

2.1 Incentive-Compatible Ride-Sharing Mechanisms

Mechanism-design approaches to ride-sharing differ mainly in how they balance allocation efficiency, incentive compatibility, individual rationality, and platform budget requirements. The closest structural reference for ORA is the hypergraph-VCG mechanism of Someya and Namerikawa [1]. Their formulation represents each feasible driver-passenger group as a hyperedge with a precomputed route cost and selects a maximum-weight set of non-overlapping groups. VCG pricing then provides passenger-side strategy-proofness and individual rationality under exact winner determination. ORA retains this hypergraph allocation structure but extends the reported information on both sides of the market. Passengers report individual sharing valuations and hard pickup and detour limits, while drivers report private reservation values.

Passenger preferences have also been incorporated explicitly in earlier mechanism-design models. Bian, Liu, and Bai [7] study on-demand first-mile ride-sharing in which passengers report personalized arrival deadlines, maximum willingness to pay, and detour tolerances. Their VCG-based mechanism combines these mobility preferences

with baseline price control and is designed to satisfy incentive compatibility and individual rationality while avoiding unreasonably low prices and carrier deficits. This work is particularly relevant to ORA because feasibility depends partly on passenger-specific requirements rather than only on platform-wide parameters. The setting is different, however: it focuses on first-mile service to transit hubs and does not model the same two-sided multi-passenger hypergraph allocation with strategic driver reservation values.

Two-sided ride-sharing mechanisms make a different trade-off. Yu, Huang, and Yue [2] propose a strategy-proof multi-party carpooling double auction that provides two-sided strategy-proofness, individual rationality, and budget balance, while its efficiency guarantee is asymptotic rather than exact. McAfee's double-auction mechanism [8] similarly obtains truthfulness and budget balance by sacrificing the marginal trade, and approximately efficient two-sided combinatorial auctions [3] provide strong incentive and budget properties without requiring exact welfare maximization. This tension is not specific to ride-sharing. Myerson and Satterthwaite [9] show that no bilateral trade mechanism can be simultaneously efficient, incentive compatible, individually rational, and budget balanced under their standard assumptions, and Parkes, Kalagnanam, and Eso [10] study how VCG-based payments in exchanges can be adjusted toward budget balance at the cost of exact incentive properties. ORA takes the exact-efficiency side of this trade-off: it uses Groves transfers and reports the resulting platform deficit rather than modifying the allocation or payment rule to impose budget balance.

Other studies concentrate on strategic behavior on the driver side. Shi et al. [5], for example, model heterogeneous private driver costs and propose a second-pricing mechanism that combines incentive-compatible driver reporting with matching and route planning. Their mechanism provides incentive compatibility, individual rationality, and budget-balance properties, but passengers are not strategic agents in the same sense as in ORA. Dynamic market design provides another point of comparison. Ma, Fang, and Parkes [11] develop a spatio-temporal pricing mechanism in which accepting platform dispatches forms a subgame-perfect equilibrium for drivers. In their complete-information dynamic model, the equilibrium is welfare-optimal, individually rational, envy-free, budget balanced, and core-selecting. ORA addresses a different setting: it is static and batch-cleared, but seeks dominant-strategy truthfulness for

reported passenger and driver types and supports multi-passenger allocations with passenger-specific hard feasibility constraints.

Taken together, these mechanisms show that incentive-compatible ride-sharing can be approached through VCG pricing, double auctions, second-pricing rules, or dynamic equilibrium pricing. ORA does not attempt to dominate these alternatives on every mechanism-design property. Its narrower contribution is to combine exact multi-passenger hypergraph allocation with two-sided private information while allowing passenger-specific pickup, detour, and sharing preferences to affect both the value and the feasibility of an allocation.

2.2 Operational and Scalable Ride-Sharing Allocation

A separate line of research focuses on the operational problem of assigning riders and vehicles at scale. Alonso-Mora et al. [4] introduced a high-capacity ride-sharing framework that first constructs feasible rider-vehicle groups and then solves an integer program for real-time fleet assignment. This structure is closely related to ORA at the optimization level: both enumerate feasible multi-passenger groups before selecting a compatible set of assignments. The main difference is that the Alonso-Mora framework treats passengers and vehicles as non-strategic inputs and does not model private valuations or mechanism payments.

Large-scale dispatch has also been studied through learning and network-aware approaches. Qin et al. [12] use reinforcement learning to incorporate long-term value into ride-hailing dispatch, while Zhou and Roncoli [13] account for the interaction between vehicle assignment, routing, and endogenous congestion. These methods address dynamic operation and computational scale more directly than ORA, but truthful reporting of passenger preferences and driver participation costs is outside their scope.

Recent dynamic ride-sharing research also emphasizes that operational performance depends strongly on user flexibility. Müller et al. [14] study dynamic ride-sharing with an agent-based MATSim model and examine how service outcomes change with assumptions about matching, waiting, and departure-time flexibility. Their results reinforce the operational importance of user-side flexibility, although preferences enter a simulation and matching framework rather than an incentive-compatible allocation mechanism.

ORA addresses a different problem. Its purpose is not to outperform city-scale dispatch systems in runtime, but to examine what changes when passenger flexibility and driver participation costs are private information and the allocation is required to satisfy exact mechanism-design guarantees. This choice necessarily limits computational scale. Experiment D therefore measures the cost of exact group enumeration, winner determination, and Clarke-payment counterfactuals rather than presenting ORA as a replacement for large-scale operational dispatch algorithms.

2.3 Fairness and Participant Outcomes

Fairness-oriented ride-pooling research usually focuses on the distribution of service or income rather than on truthful multi-passenger allocation. Raman, Shah, and Dickerson [6] incorporate fairness objectives into ride-pooling and study income redistribution among drivers. Saxena, Zhang, and Shahabi [15] instead examine disparate pricing in ride-hailing and propose pricing rules intended to improve affordability for disadvantaged areas. These studies show that efficiency alone is not sufficient to describe participant outcomes, but neither addresses the same two-sided private-information problem considered by ORA.

Kamar and Horvitz [16] provide an earlier connection between ridesharing and VCG-based payments. Their system uses approximate ride assignments together with VCG-style transfers, which creates both computational and incentive limitations because exact welfare maximization is not preserved. They also identify the possibility of a platform loss under VCG payments. ORA takes a different approach: winner determination remains exact and the resulting platform deficit is reported explicitly rather than removed through a payment adjustment that would alter the mechanism's incentive properties.

Empirical evidence supports treating sharing and waiting tolerance as heterogeneous individual attributes. Stated-preference studies find substantial variation across travelers in the willingness to share a vehicle with strangers and in the sensitivity to waiting and in-vehicle time [17], [18], while preferences for shared services also differ systematically with trip purpose and traveler characteristics [19]. These studies do not imply that pickup and detour limits must be modeled as hard constraints, but they motivate the assumption in ORA that sharing valuation and pickup tolerance are passenger-specific quantities rather than platform constants.

2.4 Positioning of ORA

The studies above address several parts of the ride-sharing mechanism problem, but with different priorities. Someya and Namerikawa [1] provide exact hypergraph allocation with VCG-style passenger pricing; Yu, Huang, and Yue [2] focus on two-sided strategy-proofness and budget balance; Alonso-Mora et al. [4] address large-scale multi-passenger assignment; Bian, Liu, and Bai [7] incorporate individualized mobility requirements in a VCG-based first-mile setting; and Ma, Fang, and Parkes [11] study dynamic pricing and driver incentives. Empirical work on pooled mobility also shows that passengers differ in their willingness to share and in the value they place on waiting and travel-time changes [17]–[19].

ORA combines a narrower set of these elements in one static, batch-cleared mechanism. Each passenger reports a solo valuation, a shared-ride valuation, a maximum pickup distance, and a maximum detour, while each driver reports a private reservation value. The pickup and detour reports are treated as hard acceptability constraints, so they determine whether a hyperedge is feasible; the solo and shared valuations determine the welfare weight of feasible service. The resulting allocation is selected by exact maximum-weight set packing and transfers are computed with two-sided Clarke pivot payments. This is not a claim that ORA dominates every mechanism on every dimension: budget-balanced auctions provide stronger budget guarantees, operational approaches scale to much larger markets, and fairness-focused models address dimensions outside this allocation objective.

3 Model and Mechanism

3.1 Market Setting

ORA considers a static ride-sharing market cleared in fixed batch windows. In each window, a finite passenger set N submits trip requests and a finite driver set M is available to serve them. The platform acts as the mechanism designer: it collects reports, constructs the feasible allocation set, selects the reported-welfare-maximizing allocation, and determines transfers. Platform budget balance is evaluated as an outcome of the mechanism rather than included in the allocation objective.

Each passenger has an origin and destination, and each driver has a current location and vehicle capacity. ORA is a direct-revelation mechanism, so strategic behavior is

represented through the private information reported before clearing. Driver repositioning across windows and responses to future prices are outside the static model. Fig. 1 summarizes the clearing pipeline, from reported types to the transfers computed for each participant.

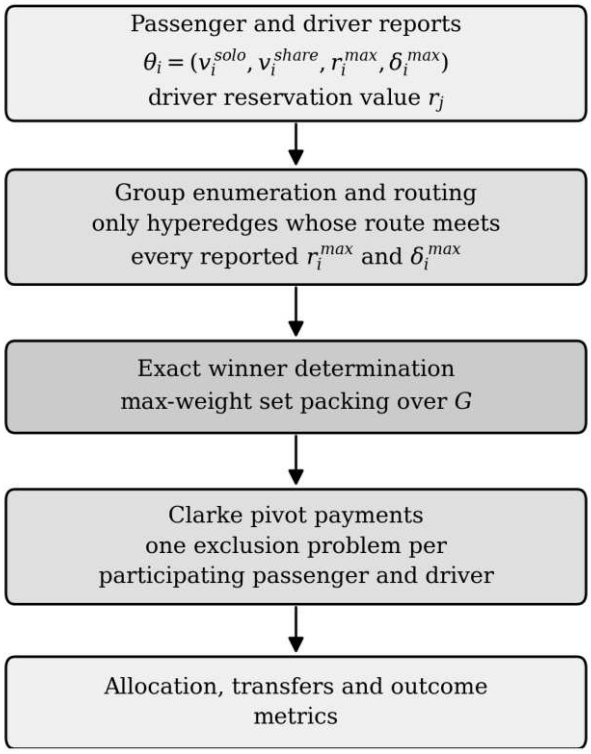


Fig. 1 Overview of the ORA clearing pipeline. Reported passenger and driver types determine which groups are feasible, exact winner determination selects the allocation, and Clarke pivot counterfactuals set the transfers

3.2 Passenger and Driver Types

Each passenger i ∈ N has private type

$$\theta_i = (v_i^{solo}, v_i^{share}, r_i^{max}, \delta_i^{max}), \quad (1)$$

where v_i^{solo} is the value of solo service, v_i^{share} is the value of shared service, r_i^{max} is the maximum acceptable pickup distance, and δ_i^{max} is the maximum acceptable in-vehicle detour relative to the direct trip. The valuations satisfy $v_i^{share} \leq v_i^{solo}$, and sharing tolerance is defined as

$$\tau_i = v_i^{share} / v_i^{solo}, \quad 0 < \tau_i \leq 1. \quad (2)$$

The central modeling choice is that pickup limits, detour limits, and sharing valuations are passenger-specific rather than common platform parameters. Pickup and detour bounds are treated as hard acceptability constraints. An allocation that violates a

passenger's true bound is therefore non-implementable for that passenger. This hard-constraint interpretation is required for the passenger incentive result below.

Let $A_i(\theta_i)$ denote the allocations in which passenger i is unmatched or, if served, satisfies her true pickup and detour bounds. The incentive argument uses the true utility

$$u_i(a, p_i;\ \theta_i) = v_i(a) - p_i,\ a \in A_i(\theta_i), \tag{3}$$
$$u_i = 0 \text{ if } i \text{ is unmatched},$$
$$u_i = -\infty \text{ if } a \notin A_i(\theta_i).$$

The $-\infty$ term in (3) is a formal representation of non-implementability, not a calibrated monetary disutility.

Each driver j ∈ M has a current location, vehicle capacity, observable marginal operating cost c_j, and private reservation value ρ_j. The reservation value represents the driver's opportunity cost of participation in the current clearing window. The economic cost of an assigned route is the observable route operating cost plus the driver's reservation value. Table 1 collects the notation used in the rest of the paper.

Table 1 Notation used in the model

Symbol	Meaning
N, M	Passenger and driver sets; indices i, j
θ_i	Private type of passenger i
v_i^{solo}, v_i^{share}	Value of solo and of shared service for passenger i
τ_i	Sharing tolerance, v_i^{share} / v_i^{solo}
r_i^{max}, δ_i^{max}	Maximum acceptable pickup distance and detour
$\hat{x}$	Reported value of a private quantity x
c_j, ρ_j	Operating cost per km (observable) and reservation value (private) of driver j
$v_i(a)$	True value obtained by passenger i under allocation a
$W_{-i}(a)$	Reported welfare of allocation a evaluated without the value of passenger i
F_{-i}	Allocations feasible in the market from which passenger i has been removed
g, G	Candidate group (hyperedge) and set of feasible groups
$P(g)$, $j(g)$, $m(g)$	Passengers, driver and service mode of group g

Symbol	Meaning
$C(g)$	Operating cost of the route serving group g
d_i^{pickup}, d_i^{ride}, d_i^{direct}	Distance before boarding, in vehicle, and direct origin–destination; the first two are evaluated on the route of group g
$W(g)$	Reported welfare of group g
W^{nom}, W^{adj}, C^{acc}	Nominal welfare, adjusted welfare, acceptability cost
x_g	Binary variable selecting group g
a, a^*	Allocation and welfare-maximizing allocation
H_i	Maximum welfare of remaining agents when i is excluded
p_i, e_j	Clarke payment of passenger i; Clarke earnings of driver j
u_i, u_j	Realized utility of passenger i and driver j
B, D	Platform budget balance and required subsidy

3.3 Feasible Groups and Routing

ORA represents the allocation space as a hypergraph. A candidate group g consists of one driver j(g), a passenger subset P(g) ⊆ N, and a service mode m(g) ∈ {SOLO, SHARED}. In the implementation, a one-passenger group is SOLO and a group of two or more passengers is SHARED.

A group is feasible only if vehicle capacity is respected and a valid route satisfies every reported pickup and detour constraint. Concretely, for each i ∈ P(g), the route must satisfy

$$d_i^{\mathrm{pickup}}(g) \le \hat{r}_i^{\max}, \tag{4}$$

$$d_i^{\mathrm{ride}}(g) - d_i^{\mathrm{direct}} \le \hat{\delta}_i^{\max}. \tag{5}$$

Here $d_i^{pickup}(g)$ is route distance before passenger i boards, $d_i^{ride}(g)$ is distance travelled while i is in the vehicle, and d_i^{direct} is the direct origin-to-destination distance. Routing therefore determines feasibility before winner determination. Only groups whose realized routes satisfy all reported bounds enter the allocation problem. Passenger flexibility consequently affects the feasible allocation set, not only the weight assigned to an allocation. A feasible allocation contains no passenger or driver more than once. The empty allocation is always feasible and provides the outside option used in the individual-rationality arguments.

3.4 Welfare and Winner Determination

For passenger i, the value reported for a candidate group g depends on the service mode. If m(g)=SOLO, the mechanism uses $\hat{v}_i^{solo}$; if m(g)=SHARED, it uses $\hat{v}_i^{share}$. Let C(g) denote the operating cost of the route and let $\hat{\rho}_{j(g)}$ be the reported reservation value of the driver assigned to that group. The reported welfare of group g is therefore

$$W(g) = \sum_{i \in P(g)} \hat{v}_i(g) - C(g) - \hat{\rho}_{j(g)}. \quad (6)$$

The empirical evaluation distinguishes between nominal and adjusted welfare. W^{nom}(a) is the true valuation-minus-cost total of the allocation produced by a mechanism, even when a comparison arm assigns a passenger to a ride that exceeds her true pickup or detour limits. To make those violations visible, C^{acc}(a) records the total acceptability penalty associated with the allocation, and adjusted welfare is defined as

$$W^{adj}(a) = W^{nom}(a) - C^{acc}(a). \quad (7)$$

The adjusted quantity in (7) is used only as an evaluation diagnostic. It is not the objective optimized by ORA, and subtracting an acceptability penalty does not make an allocation outside a passenger's hard bounds feasible. With truthful reports, ORA only selects allocations that satisfy those bounds, so C^{acc}(a)=0 and consequently W^{adj}(a)=W^{nom}(a).

Let $x_g \in \{0,1\}$ indicate whether candidate group g is selected. ORA chooses the allocation by solving the winner-determination problem (WDP)

$$\max_x \sum_{g \in G} W(g)\, x_g \quad (8)$$

$$\text{s.t. } \sum_{g:\, i \in P(g)} x_g \le 1, \ \forall i \in N,$$

$$\sum_{g:\, j(g)=j} x_g \le 1, \ \forall j \in M,$$

$$x_g \in \{0,1\}, \ \forall g \in G.$$

Thus, (8) is an exact maximum-weight set-packing problem over the feasible hyperedges. Each passenger and each driver can appear in at most one selected group. This follows the standard winner-determination structure used in combinatorial auctions [20], [21]. Solving the problem exactly is important for the mechanism itself, not only for computational accuracy, because the incentive-compatibility results assume that reported welfare is maximized over the full feasible allocation set.

3.5 Clarke Payments and Utilities

Payments follow the Clarke pivot rule [22]–[24]. For each participating agent, the winner-determination problem is solved again with that agent excluded. The transfer equals the externality that the agent imposes on the remaining market. The main allocation and every exclusion counterfactual use the same welfare definition.

For passenger i, H_i is the maximum welfare available to the remaining agents when i is removed. Write $W_{-i}(a)$ for the reported welfare of allocation a evaluated without the value of passenger i, and F_{-i} for the allocations that are feasible in the market from which i has been removed. The payment from the passenger to the platform is

$$p_i = H_i - W_{-i}(a*), \tag{9}$$
$$H_i = \max_{a \in F-i} W_{-i}(a).$$

Passenger utility is then

$$u_i = v_i(a*) - p_i. \tag{10}$$

On the driver side, positive earnings represent payments from the platform to the driver. If driver j operates group g_j and receives Clarke earnings e_j, realized utility is

$$u_j = e_j - C(g_j) - \rho_j. \tag{11}$$

An idle driver receives the outside option of zero.

3.6 Incentive Compatibility and Individual Rationality

Proposition 1 — Passenger incentive compatibility. Suppose that the winner-determination problem is solved exactly, passenger utilities are quasi-linear, and r_i^{max} and δ_i^{max} represent hard limits that passenger i is not willing to exceed. Under these assumptions, reporting the true passenger type is a weakly dominant strategy.

Consider passenger i while keeping the reports of all other agents fixed. In the Clarke payment, the exclusion term H_i is determined by the welfare that can be obtained when passenger i is removed from the market, so it does not depend on her own report. For any allocation that respects her true pickup and detour limits, her utility is therefore equal to her true value for the allocated ride plus the welfare generated for the other agents, minus the constant H_i. When passenger i reports truthfully, the mechanism considers exactly the allocations she is willing to accept and uses her true valuations for solo and shared rides. Since the allocation is chosen by exact welfare maximization, the mechanism selects the best outcome for this expression among those feasible

allocations. Reporting a smaller pickup radius or detour limit can only remove outcomes that would actually have been acceptable to her. Reporting larger limits may make additional allocations available, but these cannot improve her true utility: if the resulting allocation still satisfies her real limits, it was already feasible under truthful reporting; if it exceeds either real limit, it is not an acceptable outcome and has utility $-\infty$ under the hard-constraint assumption. The usual Groves argument applies to misreports of v_i^{solo} and v_i^{share}. Therefore, neither changing one component nor jointly misreporting several components can provide passenger i with a higher true utility than truthful reporting.

Passenger ex-post individual rationality follows from the fact that remaining unmatched is always a feasible alternative and gives utility zero. With truthful reports and exact welfare maximization, the mechanism therefore cannot choose an acceptable allocation that leaves the passenger worse off than not being served.

Proposition 2 — Driver incentive compatibility and individual rationality. Under exact welfare maximization and quasi-linear utility, truthful reporting of the driver reservation value is a dominant strategy. The reservation value enters reported welfare as a negative participation cost, while the exclusion welfare of the remaining market does not depend on the driver's own report. The standard Groves argument therefore applies. The allocation in which the driver is present but idle reproduces the outside option available in the exclusion market, establishing ex-post individual rationality.

These guarantees depend critically on the modeling assumptions. If passengers are willing to exchange additional pickup distance or detour for a lower price, then pickup and detour should enter a continuous valuation function rather than define hard feasibility. The DSIC result in this paper is therefore conditional on genuine hard passenger acceptability bounds.

3.7 Budget Balance

The platform budget is defined as

$$B = \sum_i p_i - \sum_j e_j, \tag{12}$$

where passenger payments enter positively and driver earnings negatively. When B<0, the payments collected from passengers are not enough to cover driver earnings, so the difference has to be subsidized by the platform. The required subsidy is

$$D = \max\{-B, 0\}. \tag{13}$$

ORA does not require the budget to balance in every market. This follows from the mechanism-design choice made in this paper: the mechanism keeps exact reported-welfare maximization, incentive compatibility on both sides, and ex-post individual rationality instead of modifying the allocation or transfers to guarantee a non-negative platform balance. For this reason, the budget deficit is treated as an outcome of the mechanism and measured directly in the experiments. No additional threshold or redistribution rule is introduced, since doing so would change the payment structure and could affect the theoretical guarantees established for ORA.

4 Experimental Design

4.1 Common Simulation Setting

The experiments are designed to measure what changes when passenger-specific information is replaced by common platform assumptions while the rest of the mechanism is kept unchanged. Unless an experiment states otherwise, the baseline setting uses a 6 × 6 Manhattan grid with 12 passengers, four drivers, vehicle capacity three, and 50 paired random seeds. For each seed, every mechanism is evaluated on the same generated market. This pairing is important because differences between the results can then be attributed to the information available to the mechanism, rather than to different passenger or driver realizations. Table 2 lists the baseline setting.

The demand used in these experiments is synthetic. Passenger origins and destinations are sampled from the network, and their valuations and flexibility limits are drawn from predefined distributions. In the baseline configuration, the mean solo valuation is 20 monetary units, the mean sharing tolerance is 0.75, and the mean maximum pickup distance is 1.5 km. Drivers have a mean reservation value of 5 monetary units, while operating cost is assumed to be observable and fixed at 0.8 monetary units per kilometer. These values are used as simulation parameters and should not be interpreted as estimates calibrated from a commercial ride-hailing platform.

Table 2 Baseline simulation settings

Element	Baseline setting
Network	6 × 6 Manhattan grid
Baseline market	12 passengers; 4 drivers
Vehicle capacity	3
Random seeds	50 paired seeds
Passenger means	v^{solo}=20; τ=0.75; r^{max}=1.5 km
Driver economics	reservation=5; operating cost=0.8 per km

All experimental arms are evaluated through the same measurement layer. When a comparator allocates using assumed passenger types, the realized allocation is scored against the passengers' true generated types. In particular, a common-radius mechanism is not treated as acceptable merely because the platform imposed that radius; realized pickup distance is evaluated against each passenger's true maximum pickup distance.

4.2 External Benchmark and Radius Sweep

The first comparison evaluates ORA against a controlled VCG ablation and a reimplementation of the Someya-Namerikawa hypergraph mechanism [1]. The VCG ablation preserves ORA's exact winner determination, routing, driver model, and Clarke-payment calculation. It differs only in information: passenger-specific pickup radii and shared-ride valuations are replaced by a common radius R and the common sharing rule $v_i^{share} = 0.75 v_i^{solo}$. Passenger-specific detour bounds remain unchanged. This design isolates the two passenger-specific components studied by ORA without simultaneously changing the optimization procedure.

The common radius varies from 0.5 km to 4.0 km in 0.25 km increments. ORA is solved once per seed because its allocation does not depend on the platform-wide radius, whereas the VCG ablation is re-solved at every radius. ORA outcomes are repeated across the radius axis only for comparison and visualization.

The Someya reimplementation is evaluated on the same market instances. It retains exact maximum-weight hypergraph allocation but differs in feasibility and participant modeling: capacity is enforced, waiting and detour enter as soft penalties rather than passenger-specific hard bounds, and drivers are fixed-rate executors rather than

strategic agents with private reservation values. It is also independent of the global-radius sweep.

Before serving as an external comparator, the reimplementation is behaviorally checked against the Standard-VCG counterpart described in [1]. Across the same 50 paired synthetic markets, group generation, routing, exact set packing, and Clarke pricing are fixed, while waiting- and detour-disutility coefficients are set to zero in the comparator. The purpose is qualitative validation of comparative behavior rather than numerical reproduction of the original Yokohama experiment, because the original trip instances and preference parameters are not available.

4.3 Experiment A: Factorial Ablation

Experiment A uses a 2 × 2 factorial design to separate the contribution of passenger-specific pickup radii from that of passenger-specific shared-ride valuations. The pickup factor compares individual r_i^{max} with a global radius of 2.0 km. The sharing-value factor compares individual v_i^{share} with the common rule $v_i^{share} = 0.75 v_i^{solo}$. The four arms are listed in Table 3.

Table 3 Factorial-ablation arms

Arm	Pickup radius	Sharing value
ORA	Individual	Individual
No radius elicitation	Global 2.0 km	Individual
No pooling elicitation	Individual	Global 0.75
Full ablation	Global 2.0 km	Global 0.75

All four arms use the same market for each of the 50 seeds and the same exact WDP and Clarke-payment implementation. Allocations produced under assumed passenger types are evaluated against the true generated passenger population. Factorial contrasts are computed at seed level before aggregation, preserving the paired design and allowing the two elicited components and their interaction to be evaluated separately.

For any outcome Y, the radius main effect is

$$\tfrac{1}{2}[(Y_{\mathrm{ORA}} - Y_{\mathrm{NoRadius}}) + (Y_{\mathrm{NoPooling}} - Y_{\mathrm{Full}})], \tag{14}$$

the sharing-value main effect is

$$½[(Y_{\mathrm{ORA}} - Y_{\mathrm{NoPooling}}) + (Y_{\mathrm{NoRadius}} - Y_{\mathrm{Full}})], \tag{15}$$

and the interaction is

$$Y_{\mathrm{ORA}} - Y_{\mathrm{NoRadius}} - Y_{\mathrm{NoPooling}} + Y_{\mathrm{Full}}. \tag{16}$$

4.4 Experiment B: Heterogeneity Robustness

Experiment B tests whether the value of individual elicitation arises from passenger heterogeneity rather than from an unfavorable choice of a platform-wide constant. Pickup-radius dispersion is varied while the population mean remains 1.5 km. At each level, ORA is compared with a no-radius-elicitation arm whose global radius is the realized mean $r_i^{\max}$ for that seed and heterogeneity level. This removes mean differences and isolates dispersion.

$$\sigma_r \in \{0, 0.25, 0.50, 0.75\}\,\mathrm{km}. \tag{17}$$

A second sensitivity analysis varies the standard deviation of sharing tolerance while holding its mean at 0.75. ORA is compared with a no-pooling-elicitation arm whose common sharing discount equals the realized mean τ_i for the same seed and heterogeneity level. The homogeneous case is an internal control: when the relevant characteristic has zero dispersion and the comparator receives the same realized mean, the individual and global descriptions coincide.

$$\sigma_\tau \in \{0, 0.075, 0.150, 0.225\}. \tag{18}$$

4.5 Experiment C: Market Pressure

Experiment C tests whether the contribution of passenger-specific information changes as demand increases relative to supply. The driver count is fixed at four and vehicle capacity remains three. Passenger populations of 8, 12, and 16 correspond to passenger-driver ratios of 2:1, 3:1, and 4:1. The three pressure levels are listed in Table 4.

Table 4 Market-pressure design

Passengers	Drivers	Passenger/driver	Demand/seat capacity
8	4	2:1	0.67
12	4	3:1	1.00
16	4	4:1	1.33

The four factorial arms from Experiment A are evaluated at each pressure level with the same global assumptions of a 2.0 km radius and a 0.75 common sharing discount. Markets are nested within random seeds. The 12-passenger market is the baseline Experiment A instance; the low-pressure market retains the first eight passengers, while the high-pressure market adds four independently generated passengers without redrawing the original drivers. This design changes demand pressure while holding the driver draw fixed.

Factorial effects are calculated separately at each pressure level and then compared within seed. The resulting trends distinguish effects that remain approximately constant from effects that change as the market becomes more supply-constrained.

4.6 Experiment D: Computational Scalability

Experiment D examines how the computational cost of exact ORA grows as the market becomes larger. The benchmark uses a 10 × 10 Manhattan grid, vehicle capacity three, and a fixed passenger-to-driver ratio of 3:1. For each market size, 20 random seeds are evaluated. The tested instances range from 12 to 120 passengers, with the number of drivers increased proportionally to keep the same market pressure.

For every instance, several parts of the computation are measured separately. These include the time required to enumerate feasible groups, the number of candidate and feasible groups, the runtime of the main winner-determination problem, the total clearing time including the Clarke counterfactuals, the additional payment-computation overhead, and the complete end-to-end runtime. This decomposition is necessary because solving the allocation problem once represents only part of the cost of exact ORA. Computing Clarke payments requires the mechanism to solve additional exclusion problems for the participating agents, so reporting only the main WDP runtime would give an incomplete picture of computational performance.

Runtime results are summarized using both means with 95% confidence intervals and medians with interquartile ranges. The median-based summaries are particularly useful because some runs are noticeably slower than others, producing right-skewed runtime distributions. To describe how runtime changes with market size, log-log fits of the form (19) are estimated using the median values at each size. The resulting

exponent β is used only as an empirical description of growth over the tested range and should not be interpreted as a formal computational-complexity bound.

$$y = \alpha\, n^{\beta}. \tag{19}$$

4.7 Outcome Measures and Statistical Reporting

A common metrics layer is used wherever a quantity is well defined. Efficiency is reported through nominal realized welfare W^{nom}, match rate, and pooling rate. Service quality includes mean and 95th-percentile waiting time among served passengers and mean in-vehicle detour. Acceptability-adjusted welfare, adjusted passenger utility, acceptability cost, and true-bound violation rates are reported separately so that nominal economic value is not conflated with implementability.

Passenger acceptability is measured as the fraction of served passengers whose realized pickup distance exceeds their true $r_i^{\max}$,

$$V = \left|\left\{i \in S : d_i^{\text{pickup}} > r_i^{\max}\right\}\right| / |S|, \tag{20}$$

where S is the set of served passengers. Passenger IR violations, nominal passenger utility, and acceptability-adjusted passenger utility are also reported. Adjusted utility is a robustness diagnostic and does not replace the hard-feasibility interpretation used in the incentive proof.

Driver outcomes include IR violations, allocated-driver utility, earnings inequality, and platform budget balance. The subsidy requirement D is defined in (13); deficit frequency is the fraction of seeds with B<0, and deficit per served passenger divides D by the number of matched passengers. Gini and Jain indices summarize distributional outcomes but are interpreted jointly with match rate and waiting-time measures rather than as standalone performance criteria.

For the main synthetic experiments, uncertainty is summarized using 95% confidence-interval half-widths across the 50 paired seeds. Factorial effects and sensitivity trends are formed at seed level before aggregation. Headline 2.0 km ORA-versus-VCG comparisons additionally report paired t-tests, Cohen's d_z, and a 10,000-resample paired bootstrap of the seed-level difference.

$$\text{CI}_{95} = 1.96\, s / \sqrt{n}. \tag{21}$$

4.8 Semi-synthetic NYC TLC Validation

The final experiment replaces the synthetic geography and uniform demand source with public TLC High Volume For-Hire Vehicle (HVFHV) trip records [25]. The TLC data provide observed geography, trip-distance structure, and demand timing, while passenger valuations, sharing tolerances, pickup and detour limits, and driver economics remain simulated. The validation therefore remains semi-synthetic rather than fully calibrated.

TLC HVFHV data have also been used in peer-reviewed ride-sharing research: Storch, Timme, and Schröder [26] analyzed more than 250 million NYC requests from 2019 using TLC origin-destination zones and shared-request information.

The sample contains Manhattan-to-Manhattan trips from January through March 2025. After filtering to the 69 Manhattan taxi zones, the cached pool contains 15,146,072 observed trips. The synthetic benchmark's market size, capacity, mechanism arms, radius sweep, and 50 paired seeds are retained. Passenger origin-destination pairs are drawn from observed trip records, and driver starting zones follow the empirical drop-off distribution.

Travel is represented by a directed taxi-zone distance matrix. For an ordered origin-destination pair with at least 20 observed trips, the matrix uses the median recorded trip distance; thinner pairs use centroid distance multiplied by the preprocessing detour factor. This estimator covers 87.75% of directed non-diagonal zone pairs and 99.99% of observed trip mass. It is a zone-level empirical distance representation rather than a full road-network routing model.

The purpose of this experiment is a robustness test, not a claim of complete behavioral calibration. It asks whether the synthetic efficiency-acceptability trade-off persists when regular-grid geography and uniform demand are replaced by observed Manhattan trip patterns while the controlled preference model is held fixed.

5 Results

5.1 ORA versus the Global-Radius Benchmark and Someya

The radius sweep shows a consistent trade-off between nominal allocation performance and passenger-specific acceptability. ORA achieves mean realized

welfare of 92.85 and serves 59.17% of requesting passengers. Because pickup limits are enforced during feasibility and routing, the measured pickup-radius violation rate is zero throughout the experiment. Increasing the common radius available to the VCG ablation enlarges its feasible set and raises both nominal welfare and coverage, while also increasing the fraction of served passengers whose true pickup limits are exceeded. Fig. 2 shows nominal welfare and match rate across the sweep.

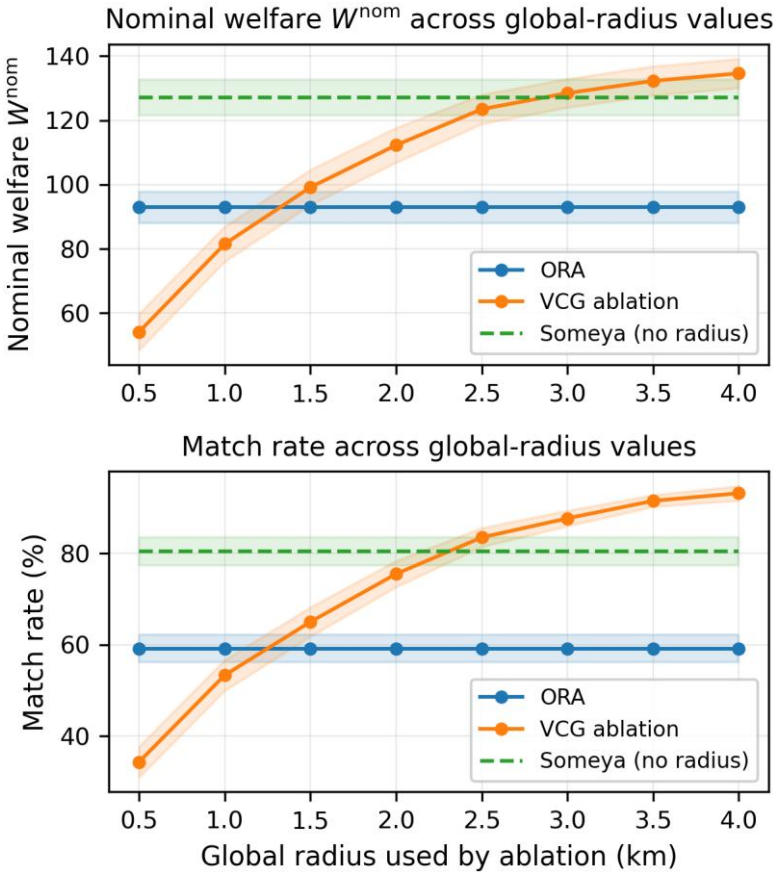


Fig. 2 Nominal welfare and match rate across the global-radius sweep. ORA and Someya are radius-independent references

At a 2.0 km global radius, the VCG ablation achieves mean realized welfare of 112.25 compared with 92.85 for ORA and serves 75.50% rather than 59.17% of passengers. However, 33.23% of the passengers served by the VCG ablation are picked up beyond their true r_i^{max}. The ablation also has a higher pooling rate (92.40% versus 72.87%), a longer mean served-passenger wait (3.45 versus 2.66 minutes), and a larger mean detour (0.531 versus 0.442 km).

The Someya reimplementation reaches the highest nominal welfare and match rate at this reference point, 127.07 and 80.50%, respectively. These outcomes are

accompanied by a 37.41% pickup-radius violation rate, a mean served-passenger wait of 4.09 minutes, and a mean detour of 0.644 km. ORA therefore does not dominate the comparators on unconstrained welfare or coverage. The comparators obtain additional nominal value partly by admitting assignments that ORA excludes as incompatible with passenger-specific pickup limits.

Paired seed-level inference confirms that the synthetic 2.0 km trade-off is not driven by a few market realizations. Relative to the VCG ablation, ORA changes nominal welfare by -19.40 ± 3.80 units ($p < 10^{-12}$, $d_z = -1.42$; paired-bootstrap 95% CI [-23.06, -15.74]), match rate by -16.33 ± 2.99 percentage points ($p < 10^{-13}$, $d_z = -1.52$; bootstrap CI [-19.33, -13.33]), and true pickup-radius violations by -33.23 ± 3.93 percentage points ($p < 10^{-20}$, $d_z = -2.35$; bootstrap CI [-37.21, -29.34]). ORA also raises adjusted passenger utility by 1.86 ± 1.24 units ($p = 0.005$, $d_z = 0.41$) and reduces passenger-IR violations by 12.02 ± 4.12 percentage points ($p < 10^{-6}$, $d_z = -0.81$). Table 5 summarizes the three mechanisms at the 2.0 km reference point, and Fig. 3 shows the corresponding true pickup-radius violations across the whole sweep.

Table 5 ORA, VCG ablation, and Someya at the 2.0 km reference point

Metric	ORA	VCG ablation	Someya
Nominal welfare	92.85	112.25	127.07
Match rate	59.17%	75.50%	80.50%
Pooling rate	72.87%	92.40%	90.72%
Mean wait, served	2.66 min	3.45 min	4.09 min
Mean detour	0.442 km	0.531 km	0.644 km
True r^{max} violations	0.00%	33.23%	37.41%
Adjusted passenger utility	9.43	7.56	—
Passenger IR violations	0.00%	12.02%	—

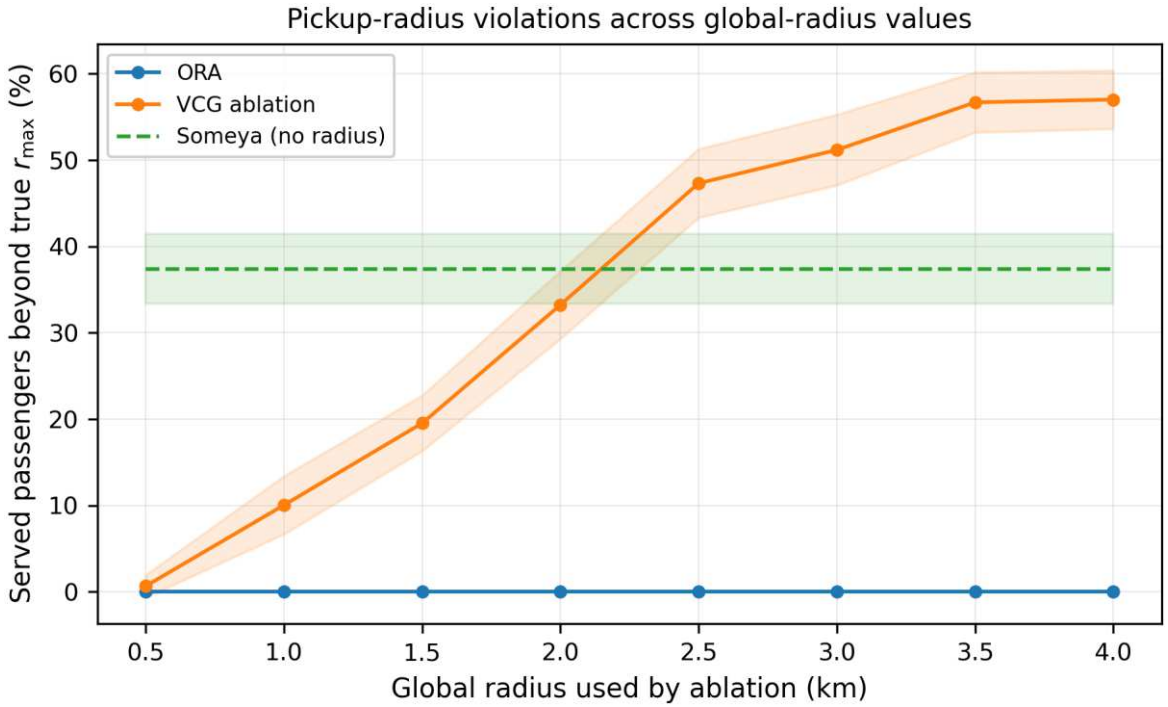


Fig. 3 True pickup-radius violations across the global-radius sweep. ORA remains at zero while the global-radius VCG ablation worsens as the assumed radius increases

Acceptability-adjusted diagnostics lead to the same interpretation. ORA has $W^{nom}=W^{adj}=92.85$ because its acceptability cost is zero. The VCG ablation decreases from 112.25 nominal welfare to 101.27 after an acceptability cost of 10.98, while Someya decreases from 127.07 to 111.69 after a cost of 15.38. These adjusted values quantify the nominal welfare associated with service outside stated flexibility; they do not make those allocations feasible under the hard-bound model.

Distributional indices at the 2.0 km reference point are similar across mechanisms. For waiting times over all requesters, ORA has Gini/Jain values of 0.419/0.608, versus 0.418/0.602 for the VCG ablation and 0.427/0.609 for Someya. Driver-earnings Gini/Jain values are 0.118/0.937 for ORA and 0.101/0.949 for the VCG ablation; they are not reported for Someya because its driver model does not use comparable strategic earnings. These indices are read jointly with match rate and service levels rather than as standalone fairness measures.

5.2 Factorial Ablation: Separating Radius and Sharing-Value Elicitation

Experiment A separates the effects of the two passenger-specific components. Removing only individual pickup radii increases welfare from 92.85 to 114.29 and match rate from 59.17% to 74.50%, but produces a 33.64% pickup-radius violation

rate, a 7.23% passenger-IR violation rate, and an acceptability cost of 9.60. Removing only passenger-specific sharing values produces much smaller aggregate changes: welfare is 91.31, match rate is 60.83%, and pickup-radius violations remain zero. Passenger-IR violations nevertheless rise to 3.78%. Table 6 reports the headline factorial results.

Table 6 Factorial-ablation headline results

Arm	Nom. welfare	Match	r^{max} viol.	Passenger IR	Adj. utility	Accept. cost
ORA	92.85	59.17%	0.00%	0.00%	9.43	0.00
No radius elicitation	114.29	74.50%	33.64%	7.23%	7.77	9.60
No pooling elicitation	91.31	60.83%	0.00%	3.78%	9.09	0.00
Full ablation	112.25	75.50%	33.23%	12.02%	7.56	10.98

Paired factorial contrasts support these distinct roles. Individual pickup-radius elicitation changes realized welfare by -21.19 ± 3.77 units and match rate by -15.00 ± 2.91 percentage points relative to a global radius, while reducing pickup-radius violations by 33.43 ± 3.92 percentage points and acceptability cost by 10.29 ± 2.85. Mean acceptability-adjusted passenger utility increases by 1.59 ± 1.11.

Sharing-value elicitation has a smaller effect on aggregate allocation but a distinct preference-related role. Its paired main effect adds 1.79 ± 0.52 welfare units and reduces passenger-IR violations by 4.29 ± 1.97 percentage points. The effect on adjusted passenger utility is 0.27 ± 0.36, whose interval includes zero, and the headline interaction intervals also include zero. The factorial evidence is therefore consistent with an approximately separable interpretation: pickup bounds primarily determine feasibility and acceptability, while sharing valuations affect selection among feasible groups.

5.3 Heterogeneity: When Does Individual Elicitation Matter?

Experiment B evaluates whether these differences arise from heterogeneity rather than from an arbitrary common parameter. In the zero-heterogeneity control, ORA and the mean-matched global comparator coincide exactly. As pickup-radius dispersion increases, the advantage of individual elicitation in acceptability becomes larger.

At a target pickup-radius standard deviation of 0.75 km, ORA reduces pickup-radius violations by 22.66 percentage points and passenger-IR violations by 10.52 percentage points relative to the mean-matched comparator. Mean acceptability-adjusted passenger utility is 4.84 units higher and acceptability cost is 20.01 units lower. Across the tested levels, each additional 0.25 km of target pickup-radius dispersion changes ORA's relative effect by -7.54 ± 1.60 percentage points for pickup-radius violations, -3.56 ± 1.13 for passenger-IR violations, +1.56 ± 0.57 for adjusted passenger utility, and -6.32 ± 2.48 for acceptability cost. The welfare trend of -0.15 ± 1.34 includes zero.

Sharing-preference heterogeneity produces a different pattern. For each additional 0.075 of target sharing-tolerance dispersion, individual sharing-value elicitation adds 1.01 ± 0.33 to relative welfare and reduces passenger-IR violations by 2.30 ± 0.92 percentage points. At the largest tested dispersion, ORA has 3.00 more welfare units and 6.56 percentage points fewer passenger-IR violations than the mean-matched common-share comparator. These controls show that the observed differences are linked to passenger heterogeneity itself rather than to a poorly chosen global mean.

5.4 Market Pressure

Experiment C examines whether the component-level effects change as demand rises relative to supply. The pickup-radius main effect remains strongly protective against true-bound violations at all three passenger-driver ratios: 32.14 ± 5.34 percentage points at 2:1, 33.43 ± 3.92 at 3:1, and 34.82 ± 3.67 at 4:1. Passenger-IR effects remain protective across the same settings.

The clearest pressure-dependent change appears in the welfare cost of individual pickup constraints. The radius main effect on welfare is -15.27 ± 3.43 at 2:1, -21.19 ± 3.77 at 3:1, and -24.03 ± 3.92 at 4:1, corresponding to a within-seed trend of -4.38 ± 1.78 welfare units per additional passenger per driver. At the same time, the acceptability-cost reduction becomes larger and the match-rate penalty becomes less negative.

Sharing-value elicitation also becomes more consequential as demand increases. Its welfare advantage rises by 0.38 ± 0.31 units per additional passenger per driver, and its reduction in passenger-IR violations strengthens by 1.95 ± 1.17 percentage points per step. No interaction trend has a paired 95% interval excluding zero. Market

pressure therefore changes the magnitude of the two main effects without providing evidence of a qualitatively different interaction mechanism over the tested range.

5.5 Computational Scalability

Exact ORA becomes more computationally demanding as market size increases. Median end-to-end clearing time is approximately 0.30 seconds with 12 passengers and four drivers, 6.36 seconds with 60 passengers and 20 drivers, 40.57 seconds with 96 passengers and 32 drivers, and 154.70 seconds with 120 passengers and 40 drivers. Table 7 reports the full benchmark, and Fig. 4 decomposes the runtime by stage.

Table 7 Computational scalability benchmark

Passengers	Drivers	Groups	Enum. time	Single WDP	Full clear	Total
12	4	10	0.000 s	0.032 s	0.295 s	0.295 s
60	20	484	0.155 s	0.105 s	6.224 s	6.365 s
96	32	1,818	1.208 s	0.355 s	39.612 s	40.566 s
120	40	3,593	2.429 s	1.112 s	152.001 s	154.702 s

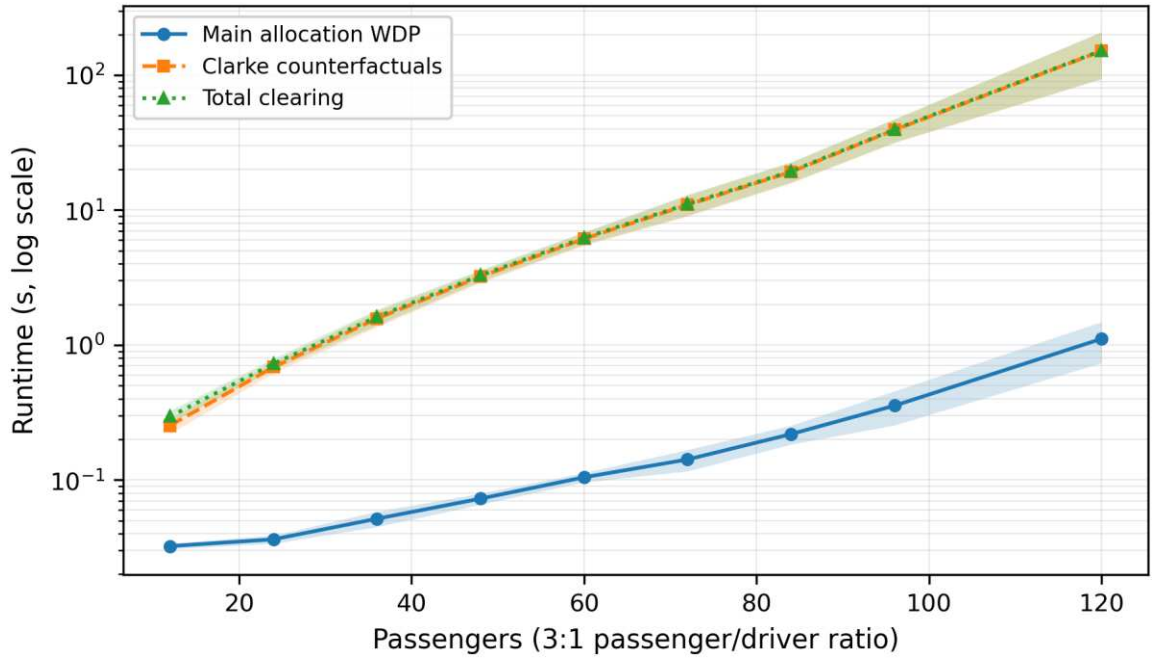


Fig. 4 Runtime decomposition for exact ORA as market size increases at a fixed 3:1 passenger-driver ratio. Curves show median runtimes with interquartile ranges over 20 seeds; the vertical axis is logarithmic

At 120 passengers, the median instance contains 3,593 feasible groups and group enumeration takes 2.43 seconds. The main exact WDP requires 1.11 seconds in the median run, whereas full clearing requires 152.00 seconds because Clarke payments require repeated agent-exclusion optimizations. The median number of exclusion solves is 142, and payment overhead represents approximately 98.0% of mean total runtime at this size. The computational bottleneck is therefore the counterfactual payment calculation rather than the main allocation solve.

Descriptive log-log fits are consistent with this decomposition. The generated-group count has a fitted exponent of 2.58 (R^2 = 0.995), the main WDP runtime 1.40 (R^2 = 0.797), full clearing 2.60 (R^2 = 0.928), and total runtime 2.61 (R^2 = 0.930). These exponents describe empirical growth over the nine tested market sizes rather than theoretical complexity bounds.

5.6 Semi-synthetic Validation on NYC TLC Demand

The NYC TLC experiment produces a more difficult spatial environment for ORA while preserving the same qualitative trade-off. At the 2.0 km reference point, ORA achieves mean nominal welfare of 27.99 and a match rate of 18.83%. The VCG ablation reaches 39.58 welfare and 27.50% matching, while Someya reaches 78.49 welfare and 49.17% matching. ORA remains at 0% true pickup-radius violations, compared with 50.77% for the VCG ablation and 77.95% for Someya. Fig. 5 shows nominal welfare and match rate in the semi-synthetic markets.

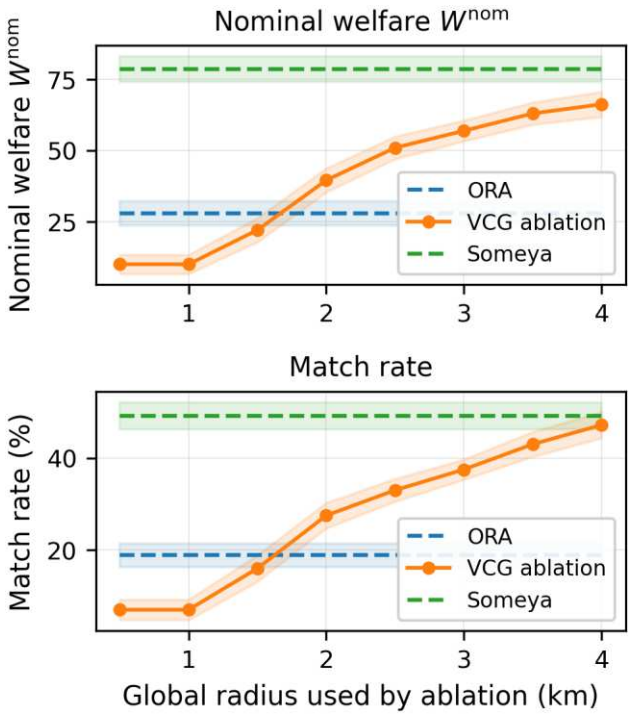


Fig. 5 Semi-synthetic NYC TLC validation: nominal welfare and match rate across the global-radius sweep. Shaded bands show 95% confidence intervals over 50 paired seeds

The radius sweep shows the same structural pattern as the synthetic experiment, but more sharply. Widening the ablation radius expands nominal welfare and matching, whereas ORA stays fixed because its feasible set is built from passenger-specific bounds rather than the platform radius. At 2.0 km, the paired ORA-minus-VCG difference is -11.59 ± 3.66 welfare units (paired t-test $p < 10^{-6}$; Cohen's $d_z = -0.88$) and -8.67 ± 2.38 percentage points in match rate ($p < 10^{-8}$; $d_z = -1.01$). Fig. 6 shows the corresponding true pickup-radius violations.

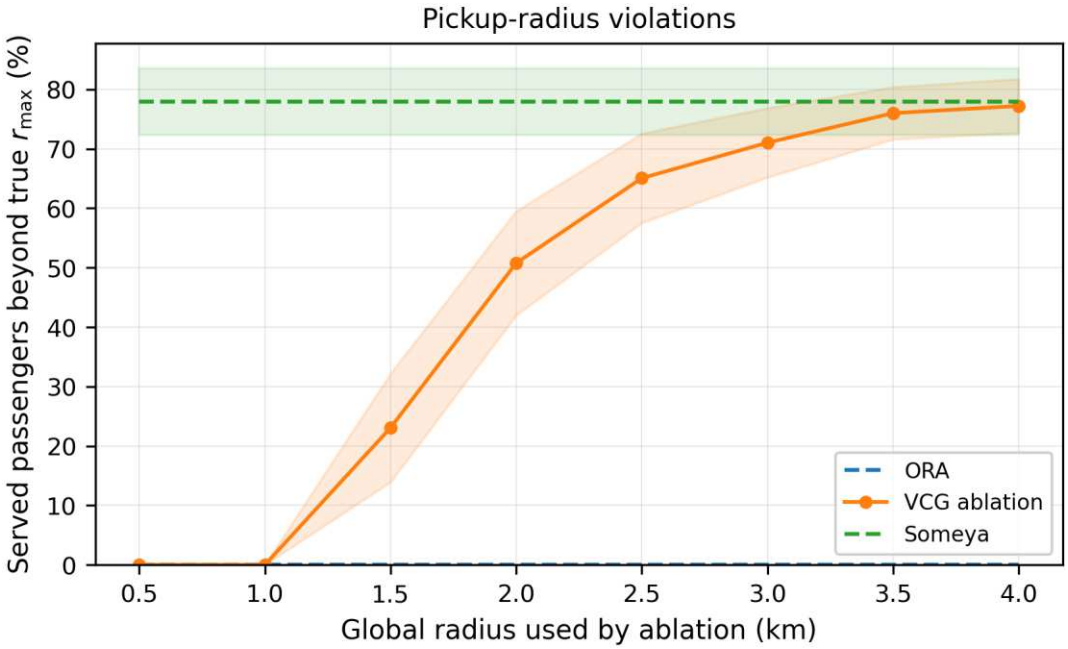


Fig. 6 Semi-synthetic NYC TLC validation: true pickup-radius violations across the global-radius sweep. ORA remains at zero while violations rise rapidly for the global-radius VCG ablation

Acceptability moves in the opposite direction. At 2.0 km, ORA cuts true pickup-radius violations by 50.77 ± 8.77 percentage points relative to the VCG ablation ($p < 10^{-14}$; d_z = -1.60). It also raises mean acceptability-adjusted passenger utility by 4.19 ± 2.11 units (p = 0.0003; d_z = 0.55) and cuts passenger-IR violations by 12.63 ± 5.77 percentage points ($p < 0.001$; d_z = -0.61). On the welfare diagnostic, ORA remains at W^{adj} = 27.99, while the VCG ablation falls from 39.58 nominal welfare to 30.97 after its 8.61 acceptability cost; Someya falls from 78.49 to 36.71 after an acceptability cost of 41.78. W^{adj} is a diagnostic rather than a replacement for true feasibility under hard bounds.

The semi-synthetic experiment therefore supports the synthetic conclusion under observed Manhattan demand and empirically estimated zone-pair distances. ORA serves fewer passengers, but passenger-specific information continues to prevent assignments outside true pickup limits and produces better acceptability-adjusted passenger outcomes than the global-radius benchmark.

5.7 Platform Budget Analysis

ORA does not balance the platform budget in the synthetic market-pressure experiment. Across the 50 paired seeds at each pressure level, all ORA markets are

in deficit. Mean subsidy requirements are 37.07 ± 5.92, 49.15 ± 7.02, and 65.72 ± 10.16 monetary units at passenger-driver ratios of 2:1, 3:1, and 4:1. Deficit per served passenger is considerably more stable at 6.90 ± 0.89, 6.80 ± 0.80, and 7.11 ± 0.98. The increase in total deficit therefore mainly reflects higher market activity rather than a large increase in subsidy per matched passenger.

The controlled VCG ablation is also frequently in deficit. Its mean deficits are 36.54 ± 5.19, 56.72 ± 9.14, and 64.22 ± 9.76 across the three pressure levels, with deficit frequencies of 100%, 98%, and 100%. This comparison indicates that the subsidy requirement cannot be attributed solely to ORA's passenger-specific pickup constraints; it is also associated with the two-sided Clarke transfer structure shared by both exact arms. Table 8 collects the budget diagnostics for ORA, and Fig. 7 shows the platform budget across pressure levels.

Table 8 Platform budget diagnostics for ORA (50 seeds)

Setting	Mean deficit	Def./served	Freq.
Pressure 2:1	37.07±5.92	6.90±0.89	100%
Pressure 3:1	49.15±7.02	6.80±0.80	100%
Pressure 4:1	65.72±10.16	7.11±0.98	100%
Radius σ=0	47.72±7.45	5.98±0.80	100%
Radius σ=0.75	49.24±6.72	6.96±0.83	100%
Sharing σ=0	50.49±7.04	6.76±0.76	100%
Sharing σ=0.225	48.75±7.46	6.77±0.87	100%
NYC TLC ORA	18.82±3.57	7.88±1.19	96%

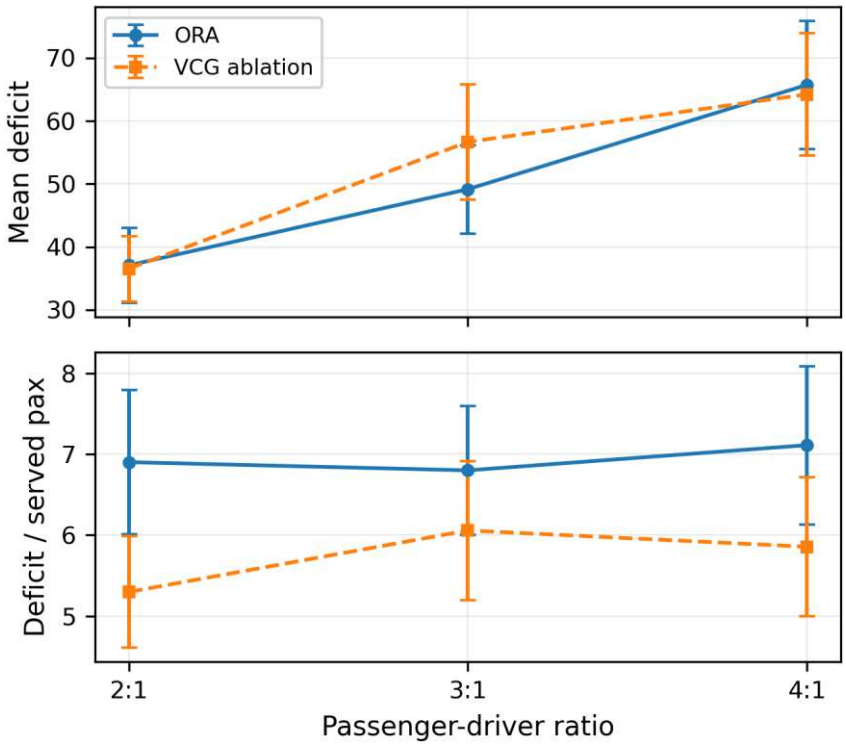


Fig. 7 Platform budget under market pressure. Positive values are subsidy requirements; error bars are 95% confidence intervals over 50 paired seeds. The upper panel reports mean deficit per market and the lower panel deficit per served passenger

Budget effects from heterogeneity are mixed. Greater pickup-radius dispersion increases ORA's deficit per served passenger from 5.98 ± 0.80 at σ_r=0 to 6.96 ± 0.83 at σ_r=0.75 km, while sharing-value dispersion produces little corresponding change. In the semi-synthetic TLC markets, ORA has a mean deficit of 18.82 ± 3.57, a 96% deficit frequency, and 7.88 ± 1.19 deficit units per served passenger. The VCG ablation records 23.70 ± 3.74, 98%, and 6.82 ± 0.79, respectively. These monetary units are not calibrated to observed fares, but the results show that the platform subsidy requirement is persistent rather than confined to the synthetic grid.

5.8 Validation of the Someya Reimplementation

The behavioral validation confirms that the preference-aware Someya reimplementation differs materially from its Standard-VCG counterpart. Each seed contains 1,192 candidate hyperedges before feasibility filtering, since the Someya arm does not impose passenger-specific pickup bounds; removing waiting and detour disutility changes between 1,132 and 1,144 group weights, and the selected allocation changes in all 50 paired markets.

The comparative outcomes reproduce the expected service-versus-passenger protection trade-off. Preference-aware Someya serves 80.50% of requests, versus 100.00% under Standard VCG, with average occupancy of 2.415 versus 3.000 and 9.66 versus 12.00 matched passengers per market. Mean native passenger utility is 9.06 versus 8.41 (paired difference +0.65; $p = 0.257$; Cohen's $d_z = 0.16$), so the mean-utility direction is favorable but not statistically distinguishable from zero in these markets. The lower tail shows a clearer difference: mean minimum passenger utility rises from -14.59 to 2.00 (difference +16.59 ± 7.76; $p = 1.17 \times 10^{-4}$), while native passenger-IR violations fall from 12.67% under Standard VCG to 0.00% under the preference-aware mechanism (difference -12.67 ± 3.21 percentage points; $p = 4.78 \times 10^{-10}$; $d_z = -1.09$).

These results provide behavioral validation rather than numerical replication of [1]. On the common markets used here, the preference-aware benchmark exhibits the expected reduction in service and occupancy together with stronger passenger-side protection. Exact replication of the original Yokohama experiment remains outside the scope of this study because the original trip instances and reported preference parameters are not available.

6 Discussion

6.1 Interpreting the Welfare-Acceptability Trade-off

The central empirical result is a trade-off rather than unconditional dominance. Increasing a platform-wide pickup radius expands the comparator's feasible set and can improve nominal welfare and matching, but the same change also increases service beyond passengers' true pickup limits. ORA excludes such assignments by construction. A similar pattern appears in the Someya benchmark, which achieves high nominal allocation performance without passenger-specific hard pickup constraints.

This distinction affects the interpretation of efficiency. ORA evaluates feasibility using individual hard acceptability bounds, whereas comparison mechanisms can obtain additional nominal welfare from assignments outside a passenger's true feasible set. The nominal-welfare difference is therefore not equivalent to an implementable efficiency gain. Acceptability-adjusted welfare and utility quantify this distinction for

ORA and the VCG ablation, while Someya's native utility and IR are reported separately under its own soft-disutility convention.

6.2 What the Two Elicited Passenger Components Contribute

The factorial ablation separates two effects that are combined in the overall ORA-versus-ablation comparison. Passenger-specific pickup distance primarily determines feasibility and acceptability. At the reference setting, individual pickup limits reduce pickup-radius violations by 33.43 percentage points and remove most measured acceptability cost, while also reducing nominal welfare and matching because some groups admitted by a larger global radius are excluded.

Individual shared-ride valuation affects a different part of the mechanism. It changes the feasible set only modestly and instead influences selection among already feasible groups. Its main effects are a smaller welfare improvement and fewer passenger-IR violations, while the estimated change in acceptability-adjusted passenger utility is less precise. The interaction estimates do not provide evidence of strong reinforcement between the two components. The pickup bound therefore determines which assignments are acceptable, while the solo-versus-shared valuation determines how acceptable service modes are valued.

6.3 Heterogeneity and Market Pressure

The mean-matched heterogeneity experiment provides an internal check on this interpretation. When the relevant passenger characteristic is homogeneous and the comparator receives the same realized mean, the mechanisms coincide. As dispersion increases, their outcomes diverge. Pickup-radius heterogeneity strengthens ORA's advantages in pickup compliance, passenger IR, adjusted utility, and acceptability cost, while the welfare trend remains statistically indistinguishable from zero. Sharing-value heterogeneity increasingly improves welfare and passenger IR under individual elicitation.

Market pressure changes the magnitude of these effects without changing their interpretation. As the passenger-driver ratio increases, the nominal welfare cost of individual pickup constraints becomes larger, while the associated acceptability protection also strengthens. Sharing-value elicitation becomes more important for welfare and passenger IR. No interaction trend is statistically distinguishable from zero

over the tested range, so higher demand pressure amplifies the main effects rather than producing evidence of a different interaction structure.

6.4 Mechanism-Design and Computational Implications

ORA occupies a specific point in the mechanism-design trade-off. Relative to Someya [1], it retains exact hypergraph winner determination while making passenger acceptability explicit and treating driver reservation values as private information. Relative to budget-balanced double auctions [2], [8] and approximately efficient two-sided mechanisms [3], ORA retains exact reported-welfare maximization and Groves incentives but accepts possible platform deficits. It should therefore be interpreted as an exact truthful benchmark for passenger-specific flexibility, not as a mechanism that dominates alternatives in budget balance or large-scale real-time operation.

The scalability experiment identifies the principal computational bottleneck. At 120 passengers, the median main WDP solve remains close to one second, whereas full clearing exceeds two and a half minutes because Clarke payments require repeated agent-exclusion solves; the allocation solve accounts for well under one percent of clearing time. Future computational work should therefore focus on payment computation, including parallel exclusion solves, solver warm starts, reuse of shared model structure, and exactness-preserving pruning. Any approximation to allocation or payment computation would need to be evaluated against the incentive guarantees that motivate the mechanism.

6.5 What the NYC TLC Validation Adds

The semi-synthetic experiment improves spatial external validity without becoming a behavioral calibration exercise. The 69-zone Manhattan network and observed HVFHV origin-destination patterns replace two strong simplifications of the synthetic benchmark: regular-grid geography and uniform spatial demand. Persistence of the efficiency-acceptability trade-off under this change indicates that the synthetic result is not generated only by the geometry of the 6 × 6 grid.

The lower ORA match rate in the TLC experiment also illustrates the operational cost of hard passenger-specific feasibility. Observed Manhattan demand produces origin-destination and driver-location configurations that are less favorable than the uniform synthetic setting, so individual bounds remove a larger fraction of candidate groups.

This reinforces the need to distinguish nominal welfare from acceptability-adjusted diagnostics.

6.6 Platform Budget Implications

The budget results make the mechanism-design trade-off explicit. Exact two-sided Clarke transfers preserve the allocation and incentive structure studied here, but the platform subsidizes nearly all cleared markets. Total subsidy increases with market pressure, while ORA's subsidy per served passenger remains close to seven monetary units in the main synthetic settings. Pickup-radius heterogeneity modestly increases the burden, whereas sharing-value heterogeneity produces little corresponding trend.

The controlled VCG ablation is also almost always in deficit, so the budget problem cannot be attributed solely to ORA's individual passenger constraints. A more cautious interpretation is that the deficit is associated with the two-sided Groves transfer design used by both exact mechanisms. A deployable platform would therefore require an external subsidy or a modified transfer rule, and any such modification should be evaluated explicitly for the efficiency, DSIC, and individual-rationality properties it changes.

7 Limitations and Future Work

These results should be interpreted within the limits of the current model and experiments. First, the controlled comparisons rely on synthetic private types: passenger valuations, flexibility bounds, and driver reservation values are generated rather than calibrated to observed willingness-to-pay or acceptance behavior. The NYC TLC validation replaces the synthetic geography and uniform demand with 15.15 million observed Manhattan HVFHV trips and empirical directed trip distances, but it remains semi-synthetic because strategic types are simulated and the zone-pair distance estimator has not been evaluated on a held-out trip partition. The paired design therefore supports controlled mechanism comparisons and the TLC experiment improves spatial external validity, but neither establishes that the reported effect sizes transfer unchanged to a commercial platform.

Second, the Someya comparison is a same-instance reimplementation rather than a numerical replication of the original Yokohama experiment. Behavioral validation against its Standard-VCG counterpart confirms that the preference-aware

implementation changes allocation outcomes in the expected direction, but the original trip instances and preference parameters are not available.

Third, ORA is static and batch cleared. Driver repositioning, online arrivals, future demand, endogenous congestion, and continuation values are outside the model. These elements are central to dynamic ride-hailing systems, and extending ORA across clearing windows would require a new dynamic formulation rather than repeated application of the current static mechanism.

Fourth, passenger incentive compatibility depends on interpreting maximum pickup distance and maximum detour as genuine hard acceptability constraints. Under this interpretation, an outcome beyond a true bound is non-implementable and cannot be compensated by a lower fare. If passengers instead trade additional pickup distance or detour against price continuously, these dimensions should enter the utility function as soft preferences and the current DSIC proof would require revision.

Fifth, the base driver model treats marginal operating cost per kilometer as observable and only the reservation value as private. Real drivers may have heterogeneous private fuel, time, vehicle, and opportunity costs. Extending the driver type space would increase realism but also the dimensionality of the mechanism and the empirical elicitation requirements.

Sixth, ORA does not impose platform budget balance. The budget analysis shows that the limitation is persistent: ORA is in deficit in 100% of the synthetic market-pressure seeds and 96% of the semi-synthetic TLC seeds, with a deficit of approximately 6.8-7.9 monetary units per served passenger across the main settings. This deficit is not corrected by a threshold or redistribution rule so that the exact Groves structure remains transparent. Future work should compare this benchmark with budget-balanced or approximately efficient alternatives and quantify the efficiency or incentive guarantees exchanged for financial sustainability.

Finally, fairness is evaluative rather than allocative in the present experiments. Gini and Jain indices, waiting-time outcomes, driver earnings, and passenger outcomes are measured after clearing and do not modify the base WDP. A history-based passenger-priority extension has been designed separately but has not been integrated into the allocation and payment counterfactuals tested here. Incorporating such a term,

validating the corresponding payment rule, and evaluating it in multi-window simulations with persistent passenger identities remain directions for future work.

8 Conclusion

This paper presents the Opt-In Radius Auction (ORA), an exact two-sided hypergraph-VCG mechanism for multi-passenger ride-sharing. Its novelty is not the hypergraph allocation itself but the strategic information embedded in it: passenger-specific pickup and detour limits enter as hard feasibility constraints, solo and shared valuations enter the welfare objective, and drivers report private reservation values. Transfers on both sides are determined using exact Clarke pivot payments. When the winner-determination problem is solved exactly, utilities are quasi-linear, and passenger pickup and detour limits are treated as genuine hard constraints, the mechanism provides the incentive-compatibility and ex-post individual-rationality properties established in this paper.

The experimental results show that using passenger-specific information involves a clear trade-off. In both the controlled synthetic experiments and the semi-synthetic NYC TLC validation, mechanisms using a common pickup radius generally obtain higher nominal welfare and serve more passengers. Part of this improvement, however, comes from allocations that exceed the actual pickup limits of some passengers. The factorial experiments show that individual pickup radii are mainly responsible for preventing these unacceptable assignments. Individual sharing valuations have a smaller but separate role, mainly affecting preference-sensitive welfare and passenger individual rationality. The heterogeneity experiments further show that these differences become relevant when passenger preferences and constraints vary across the population.

These results do not imply that ORA improves every aspect of ride-sharing allocation. Respecting individual constraints can reduce nominal coverage, and the mechanism still faces two important limitations of exact Groves mechanisms. The budget experiments show that platform subsidies are frequently required, while the scalability analysis shows that most of the computational cost comes from the repeated Clarke-payment counterfactuals rather than from solving the main winner-determination problem. Future work should therefore focus on better behavioral calibration and dynamic validation, as well as budget-aware transfer rules, richer models of driver

private information, and more efficient methods for computing exact or carefully controlled approximate payments.

Acknowledgements

During the preparation of this manuscript, the authors used AI-assisted tools, specifically OpenAI ChatGPT and Anthropic Claude. Their use was limited to language revision and copy-editing of author-written text, formatting support for tables and figures, and assistance with the implementation and debugging of parts of the simulation and analysis code released with this paper. The research content of the manuscript was not generated by these tools: the design of the mechanism, the propositions and their proofs in Section 3, the experimental design in Section 4, and the interpretation of the results in Sections 5–8 originate with the authors. All code, simulation runs, numerical results, statistical analyses, figures, tables, and references were independently reviewed, executed, and verified by the authors, who take full responsibility for the content and integrity of the published work.

Declarations

Funding

This work received no external funding.

Code and data availability

The ORA implementation, synthetic market generators, experiment scripts, fixed random seeds, NYC TLC preprocessing and filtering code, and analysis scripts used to produce the reported results are available at https://github.com/iaeste-internship/opt-in-radius-auction. The raw TLC High Volume For-Hire Vehicle trip records are publicly distributed by the New York City Taxi and Limousine Commission [25] and are not redistributed in the repository. The synthetic experiments can be reproduced directly from the repository, while reproduction of the semi-synthetic NYC TLC experiment additionally requires the corresponding public TLC trip files.

Competing interests: The authors declare no competing interests.

Ethics approval and consent to participate/consent for publication: Not applicable

Author contributions

M.V.C. and S.P.D. conceived the study and designed the mechanism. M.V.C. developed the formal model, implemented the mechanism and the simulation framework, carried out the experiments and the statistical analysis, prepared the figures and tables, and wrote the first draft. S.K.T. contributed to the methodology and validated the analysis. S.P.D. supervised the work, validated the mechanism design and the experimental protocol, administered the project, and revised the manuscript. All authors read and approved the final manuscript.

References

[1] R. Someya and T. Namerikawa, “Mechanism design-based matching and pricing for multi-passenger ride-sharing,” SICE J. Control, Meas., Syst. Integr., vol. 19, no. 1, Art. no. 2680718, 2026, https://doi.org/10.1080/18824889.2026.2680718

[2] H. Yu, M. Huang, and X. Yue, “Sharing the shared rides: Multi-party carpooling supported strategy-proof double auctions,” Prod. Oper. Manag., vol. 33, no. 7, pp. 1569–1590, 2024, https://doi.org/10.1177/10591478241252746

[3] R. Colini-Baldeschi, P. W. Goldberg, B. de Keijzer, S. Leonardi, T. Roughgarden, and S. Turchetta, “Approximately efficient two-sided combinatorial auctions,” ACM Trans. Econ. Comput., vol. 8, no. 1, pp. 1–29, 2020, https://doi.org/10.1145/3381523

[4] J. Alonso-Mora, S. Samaranayake, A. Wallar, E. Frazzoli, and D. Rus, “On-demand high-capacity ride-sharing via dynamic trip-vehicle assignment,” Proc. Natl. Acad. Sci. USA, vol. 114, no. 3, pp. 462–467, 2017, https://doi.org/10.1073/pnas.1611675114

[5] B. Shi, L. Fu, Z. Cao, and L. Zhu, “A second-pricing based incentive-compatible mechanism for matching and pricing in ride-sharing,” Expert Syst. Appl., vol. 248, Art. no. 123377, 2024, https://doi.org/10.1016/j.eswa.2024.123377

[6] N. Raman, S. Shah, and J. P. Dickerson, “Data-driven methods for balancing fairness and efficiency in ride-pooling,” in Proc. 30th Int. Joint Conf. Artif. Intell. (IJCAI), 2021, pp. 363–369, https://doi.org/10.24963/ijcai.2021/51

[7] Z. Bian, X. Liu, and Y. Bai, “Mechanism design for on-demand first-mile ridesharing,” Transp. Res. B, Methodol., vol. 138, pp. 77–117, 2020, https://doi.org/10.1016/j.trb.2020.03.011

[8] R. P. McAfee, “A dominant strategy double auction,” J. Econ. Theory, vol. 56, no. 2, pp. 434–450, 1992, https://doi.org/10.1016/0022-0531(92)90091-U

[9] R. B. Myerson and M. A. Satterthwaite, “Efficient mechanisms for bilateral trading,” J. Econ. Theory, vol. 29, no. 2, pp. 265–281, 1983, https://doi.org/10.1016/0022-0531(83)90048-0

[10] D. C. Parkes, J. Kalagnanam, and M. Eso, “Achieving budget-balance with Vickrey-based payment schemes in exchanges,” in Proc. 17th Int. Joint Conf. Artif. Intell. (IJCAI), 2001, pp. 1161–1168.

[11] H. Ma, F. Fang, and D. C. Parkes, “Spatio-temporal pricing for ridesharing platforms,” Oper. Res., vol. 70, no. 2, pp. 1025–1041, 2022, https://doi.org/10.1287/opre.2021.2178

[12] Z. T. Qin, X. Tang, Y. Jiao, F. Zhang, Z. Xu, H. Zhu, and J. Ye, “Ride-hailing order dispatching at DiDi via reinforcement learning,” INFORMS J. Appl. Anal., vol. 50, no. 5, pp. 272–286, 2020, https://doi.org/10.1287/inte.2020.1047

[13] Z. Zhou and C. Roncoli, “A scalable vehicle assignment and routing strategy for real-time on-demand ridesharing considering endogenous congestion,” Transp. Res. C, Emerg. Technol., vol. 139, Art. no. 103658, 2022, https://doi.org/10.1016/j.trc.2022.103658

[14] J. Müller, E. Nassar, M. Straub, and A. Tsui Moreno, “Exploring the dynamics of dynamic ride-sharing: Insights from a sensitivity analysis with an agent-based simulation,” Transportation, vol. 53, pp. 2279–2300, 2026, https://doi.org/10.1007/s11116-024-10564-8

[15] N. A. Saxena, W. Zhang, and C. Shahabi, “Unveiling and mitigating bias in ride-hailing pricing for equitable policy making,” AI Ethics, vol. 5, no. 2, pp. 1549–1560, 2025, https://doi.org/10.1007/s43681-024-00498-3

[16] E. Kamar and E. Horvitz, “Collaboration and shared plans in the open world: Studies of ridesharing,” in Proc. 21st Int. Joint Conf. Artif. Intell. (IJCAI), 2009, pp. 187–194.

[17] P. S. Lavieri and C. R. Bhat, “Modeling individuals’ willingness to share trips with strangers in an autonomous vehicle future,” Transp. Res. A, Policy Pract., vol. 124, pp. 242–261, 2019, https://doi.org/10.1016/j.tra.2019.03.009

[18] M. J. Alonso-González, N. van Oort, O. Cats, S. Hoogendoorn-Lanser, and S. Hoogendoorn, “Value of time and reliability for urban pooled on-demand services,” Transp. Res. C, Emerg. Technol., vol. 115, Art. no. 102621, 2020, https://doi.org/10.1016/j.trc.2020.102621

[19] R. Krueger, T. H. Rashidi, and J. M. Rose, “Preferences for shared autonomous vehicles,” Transp. Res. C, Emerg. Technol., vol. 69, pp. 343–355, 2016, https://doi.org/10.1016/j.trc.2016.06.015

[20] T. Sandholm, “Algorithm for optimal winner determination in combinatorial auctions,” Artif. Intell., vol. 135, no. 1–2, pp. 1–54, 2002, https://doi.org/10.1016/S0004-3702(01)00159-X

[21] P. Cramton, Y. Shoham, and R. Steinberg, Eds., Combinatorial Auctions. Cambridge, MA, USA: MIT Press, 2006.

[22] W. Vickrey, “Counterspeculation, auctions, and competitive sealed tenders,” J. Finance, vol. 16, no. 1, pp. 8–37, 1961, https://doi.org/10.1111/j.1540-6261.1961.tb02789.x

[23] E. H. Clarke, “Multipart pricing of public goods,” Public Choice, vol. 11, pp. 17–33, 1971, https://doi.org/10.1007/BF01726210

[24] T. Groves, “Incentives in teams,” Econometrica, vol. 41, no. 4, pp. 617–631, 1973, https://doi.org/10.2307/1914085

[25] New York City Taxi and Limousine Commission, “TLC trip record data,” High Volume For-Hire Vehicle trip records and taxi-zone lookup data. Accessed: Aug. 2026. [Online]. Available: https://www.nyc.gov/site/tlc/about/tlc-trip-record-data.page

[26] D.-M. Storch, M. Timme, and M. Schröder, “Incentive-driven transition to high ride-sharing adoption,” Nat. Commun., vol. 12, Art. no. 3003, 2021, https://doi.org/10.1038/s41467-021-23287-6