Partial Controllability in Mean Field Incentive Dynamics

Takahiro Matsunaga[1] and Eiji Hato*[1]

[1]Department of Civil Engineering, The University of Tokyo, Japan

Short summary

Abstract

This study proposes a reward-based framework for public transit demand control under partial controllability, where only a subset of users can be directly influenced by incentives. The framework integrates a mean-field model of day-to-day behavioral adjustment with a budget-feasible reverse auction based on congestion externalities. Controllable and uncontrollable users are coupled through aggregate network congestion, while the operator allocates limited rewards using class-level externality weights computed from observed flow states. The proposed mechanism satisfies per-day strategy-proofness, individual rationality, and budget feasibility in the induced single-parameter environment. The resulting system is formulated as a feedback-perturbed population dynamic, in which uncontrolled users follow inertial logit revision and controllable users are additionally shifted by auction-induced reallocation. Numerical experiments show that the mechanism reduces peak-period congestion even when the controllable fraction is limited, while also revealing a trade-off between externality-based targeting and budget efficiency under heterogeneous weights.

Keywords: reverse auction; budget-feasible incentive mechanism; partial controllability; population dynamics; mean field control

1 Introduction

Congestion in transportation is a negative-externality problem: each user's route or departure time imposes congestion on others. In static settings, a standard first-best instrument is marginal cost pricing (Pigou (1920)), but are often difficult to implement because of informational and political constraints. A central issue is information asymmetry: operators cannot directly observe users' private costs of shifting strategies. Pricing-based measures are also often viewed as regressive, unpopular, and politically infeasible.

To address these challenges, this study proposes a reward-based demand control framework combining (i) a strategy-proof and budget-feasible mechanism, (ii) partial controllability, where only a subset of users can be directly influenced, and (iii) day-to-day user adaptation modeled as an asymmetric population game. Its core is a reverse auction that elicits private shift costs and allocates a limited budget to users who can reduce congestion most cost-effectively.

Existing studies remain fragmented: some elicit private information through VCG-type mechanisms (Vickrey (1961); Clarke (1971); Groves (1973)), others model aggregate adaptation through potential games (Sandholm (2002)) or trial-and-error pricing (Seo & Yin (2019)), while others infer responsiveness through predictive behavioral models (Matsunaga & Hato (2026)). Few studies integrate budget feasibility, private information, partial controllability, and day-to-day adaptation in one framework. Although Jiang et al. (2024) studies a reverse auction for Hong Kong's MTR system, it does not explicitly model day-to-day behavior or guarantee strategy-proofness and budget feasibility. Partial controllability is also common in practice, for example in mixed CAV–HV traffic control (Wu et al. (2024)). The main contribution of this study is to model the interaction between an incentive mechanism and an asymmetric population game with controllable and uncontrollable users. The following sections introduce the mean-field population game, formulate the reverse-auction policy as a perturbation to its baseline dynamics, analyze the mechanism properties, and examine the induced dynamics through simulation experiments.

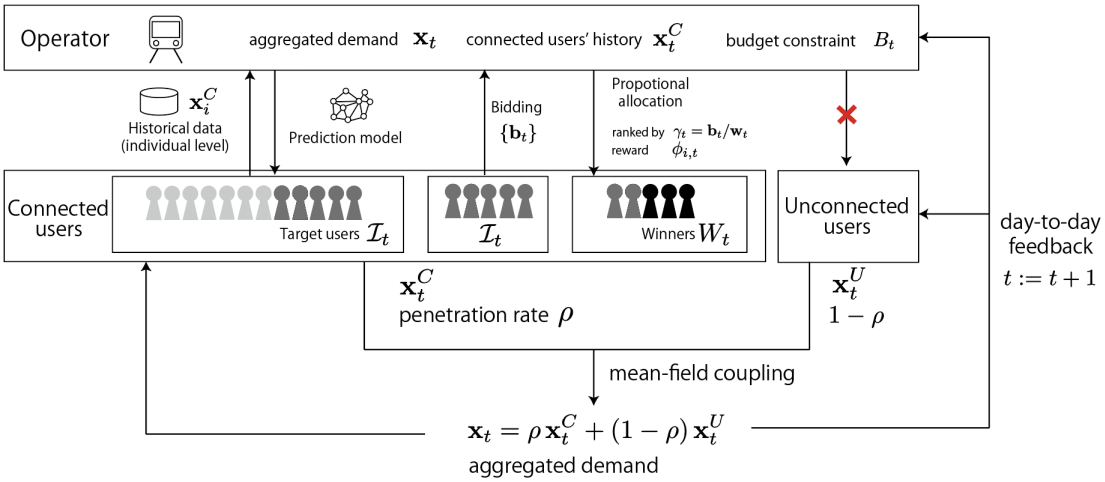


Figure 1: Mechanism framework

2 Mean-Field Population Game

Preliminaries

Consider a population of N users, each choosing a strategy s from a finite set $\mathcal{S}$. Let $\rho \in [0,1]$ be a penetration rate, the share of users directly reachable by the operator; these users form the controllable group C, while the remaining fraction $1-\rho$ forms the uncontrolled group U. Let $x_t^C, x_t^U \in \Delta^{|\mathcal{S}|}$ denote the population-level strategy shares of groups C and U, respectively, on day t, and write $x_{s,t}^C$ and $x_{s,t}^U$ for their s-th components.
This paper adopts a mean-field framework, where the population is represented by aggregate strategy distributions rather than individual-level states. Then the aggregate strategy share is given by the mean-field coupling of the two groups:

$$x_t = \rho\, x_t^C + (1-\rho)\, x_t^U. \tag{1}$$

Table 1 summarises the principal notation used throughout this paper, and Figure 2 illustrates the proposed framework.

Cost structure and externality weights

Let $U_s(x_t)$ denote the generalized cost of strategy s under the aggregate state x_t. We decompose it into a congestion-dependent component $c_s(x_t)$ and a strategy-class-specific schedule cost ψ_s:

$$U_s(x_t) = c_s(x_t) + \psi_s, \tag{2}$$

$$c_s(x_t) = \sum_{a \in r(s)} c_a(f_a(x_t)), \tag{3}$$

$$f_a(x_t) = \sum_{s' \in \mathcal{S}} x_{s',t}\, \mathbf{1}_{a \in s'}, \tag{4}$$

where $r(s)$ is the set of links contained in strategy s, and $f_a(x_t)$ is the normalized flow on link a. We assume that ψ_s is exogenous and does not depend on the current aggregate state. Thus, schedule costs affect users' private evaluations but do not create interaction effects across users.
The generalized social cost is defined as the population average of generalized costs:

$$W(x_t) = \sum_{s \in \mathcal{S}} x_{s,t} \left(c_s(x_t) + \psi_s\right). \tag{5}$$

This can be decomposed into a congestion term and a schedule-cost term as

$$W(x_t) = \underbrace{\sum_{s \in \mathcal{S}} x_{s,t}\, c_s(x_t)}_{=:W^{\mathrm{cong}}(x_t)} + \underbrace{\sum_{s \in \mathcal{S}} x_{s,t}\, \psi_s}_{=:W^{\mathrm{sched}}(x_t)}.$$

Because ψ_s is class-specific and independent of x_t, it contributes to welfare but does not generate congestion externalities. Accordingly, the marginal congestion externality induced by an additional mass choosing strategy s is

$$\frac{\partial W}{\partial x_s} - \big(c_s(x_t) + \psi_s\big) = \frac{\partial W^{\text{cong}}}{\partial x_s} - c_s(x_t) = \sum_{s' \in \mathcal{S}} x_{s',t}\, \frac{\partial c_{s'}}{\partial x_{s,t}}. \tag{6}$$

We define the class-level externality weight as the per-unit marginal congestion externality under the current state x_t:

$$w_{s,t}(x_t) = \frac{\partial W}{\partial x_s} - \big(c_s(x_t) + \psi_s\big) = \sum_{s' \in \mathcal{S}} x_{s',t}\, \frac{\partial c_{s'}}{\partial x_{s,t}}. \tag{7}$$

Evolution dynamics of the system state

At the population level, we assume that, in the absence of intervention, the aggregate state follows an inertial logit revision dynamic. Let

$$u_s(x_t) := -U_s(x_t),$$

so that strategies with lower generalized cost have higher latent utility. The baseline evolution is then written as

$$x_{t+1} = \alpha\, x_t + (1-\alpha)\, \mathcal{G}\big(\mathbf{u}(x_t)\big), \tag{8}$$

$$\mathcal{G}_s(\mathbf{u}) = \frac{\exp(u_s)}{\sum_{k \in \mathcal{S}} \exp(u_k)}. \tag{9}$$

Here $\alpha \in [0,1)$ is the inertia parameter, and $\mathcal{G}(\mathbf{u}(x_t))$ gives the logit revision of the population under the current generalized costs.
The auction-based incentive mechanism perturbs this baseline evolution through the controllable group. Denoting this intervention-induced perturbation by $\mathcal{F}_t$, we write the aggregate dynamics in reduced form as

$$x_{t+1} = \alpha\, x_t + (1-\alpha)\, \mathcal{G}\big(\mathbf{u}(x_t)\big) + \rho \mathcal{F}_t. \tag{10}$$

The term $\mathcal{F}_t$ represents the net reallocation of mass induced by the incentive scheme within the controllable population. By construction, it is mass-preserving, i.e., $\sum_{s \in \mathcal{S}} \mathcal{F}_{t,s} = 0$.

Controlled and uncontrolled population dynamics

This subsection provides a microfoundation for the reduced-form dynamics in (10). The aggregate state evolves through inertial logit revision, together with the auction-induced reallocation within the controllable group.

Uncontrolled group dynamics In the absence of intervention, the uncontrolled group follows the baseline inertial logit revision:

$$x^U_{t+1,s} = \alpha\, x^U_{t,s} + (1-\alpha)\, \Lambda_s\big(\mathbf{u}(x_t)\big), \qquad s \in \mathcal{S}, \tag{11}$$

where $\alpha \in [0,1)$ is the inertia parameter and $\Lambda_s(\mathbf{u}(x_t))$ is the logit choice probability of strategy s under the current generalized costs, as already introduced in (8).

Controlled group dynamics The basic behavioral assumption is common across the two groups, and the incentive mechanism induces an additional reallocation of mass across strategies through rewarded shifts. Accordingly, the controlled-group dynamics can be written as

$$x^C_{t+1,s} = \alpha\, x^C_{t,s} + (1-\alpha)\, \Lambda_s\big(\mathbf{u}(x_t)\big) + \mathcal{F}_{t,s}, \qquad s \in \mathcal{S}, \tag{12}$$

where $\mathcal{F}_{t,s}$ denotes the net auction-induced inflow into strategy s and satisfies $\sum_{s \in \mathcal{S}} \mathcal{F}_{t,s} = 0$. Under the reverse-auction policy specified in Section 3, this perturbation is given by

$$x^C_{t+1,s} = \underbrace{\alpha\, x^C_{t,s} + (1-\alpha)\, \Lambda_s\big(\mathbf{u}(x_t)\big)}_{\text{baseline logit revision}} - \underbrace{\frac{1}{|C|} \sum_{i \in W_t :\, \bar{s}_{i,t} = s} p_{i,t}}_{\text{expected outflow from } s} + \underbrace{\frac{1}{|C|} \sum_{i \in W_t :\, s'_i = s} p_{i,t}}_{\text{expected inflow to } s}. \tag{13}$$

Table 1: Summary of notation.

Symbol	Description
Population and strategies	
N	Total number of users
$\mathcal{S}$	Set of strategies
$\rho \in [0,1]$	Penetration rate (share of controllable users)
C, U	Controllable and uncontrolled user groups
$x_{s,t}^C,\ x_{s,t}^U$	Share of group C (resp. U) choosing strategy s on day t
x_t	Aggregate strategy share ($= \rho\, x_t^C + (1-\rho)\, x_t^U$)
$\bar{s}_{i,t}$	Baseline strategy of user i on day t (without incentives)
$s'_{i,t}$	Post-shift (desired) strategy of user i on day t
Costs and externality	
$c_s(x_t)$	Congestion cost of strategy s under aggregate state x_t
ψ_s	Schedule cost of strategy class s (exogenous)
$U_s(x_t)$	Generalized cost ($= c_s(x_t) + \psi_s$)
$W(x_t)$	Social cost ($= \sum_s x_s\, U_s$)
$w_{s,t}$	Class-level externality weight of strategy s on day t
$\delta w_{i,t}$	Marginal congestion relief from user i's shift on day t
Auction mechanism	
B_t	Daily budget
$\theta_{i,t}$	True private shift cost of user i on day t
$b_{i,t}$	Bid submitted by user i on day t
$\phi_{i,t}$	Reward offered to user i on day t
$\mathcal{I}_t$	Invitation set on day t
W_t	Winner set on day t
$\gamma_{i,t}$	Efficiency ratio ($= b_{i,t}/w_{s(i),t}$)
Day-to-day dynamics	
$\alpha \in [0,1)$	Inertia parameter
$\Lambda_s(\mathbf{u})$	Softmax choice probability for strategy s
$p_{i,t}$	Compliance probability of winner i on day t

Here $p_{i,t} \in [0,1]$ is the probability that winner i complies with the requested shift on day t. In the deterministic benchmark, compliance is given by the threshold rule $p_{i,t} = \mathbf{1}\{\phi_{i,t} \geq \theta_i\}$. More generally, probabilistic compliance can be introduced through a binary logit response such as

$$p_{i,t} = \frac{1}{1 + \exp\{-\lambda(\phi_{i,t} - \theta_i)\}},$$

where $\lambda > 0$ governs the sharpness of the response. For those who lose the auction, we can set $\phi_{i,t} = 0$ and thus $p_{i,t} \simeq 0$, so they do not contribute to the reallocation.

Mean-field aggregate dynamics Using the mean-field coupling (1), the aggregate state evolves as

$$x_{t+1,s} = \alpha\, x_{t,s} + (1-\alpha)\, \Lambda_s\big(\mathbf{u}(x_t)\big) - \frac{\rho}{|C|} \sum_{i \in W_t:\, \bar{s}_{i,t} = s} p_{i,t} + \frac{\rho}{|C|} \sum_{i \in W_t:\, s'_i = s} p_{i,t}, \qquad s \in \mathcal{S}. \tag{14}$$

The first two terms are the endogenous day-to-day revision, while the last two terms are the auction-induced control input. Thus the mechanism enters the population game as an external feedback perturbation.

3 Day-to-Day Reverse Auction Mechanism

In the previous section, we introduced the mean-field population dynamics with perturbation generated by incentive intervention. This section describes the mechanism at the individual level, where users in the same class share the same externality weight w_s but differ in their private shift costs $\theta_{i,t}$.

Here we assume that the operator can reach the controllable group through individualised reward incentives delivered via their platforms like automated communication channels (e.g., smartphone app, email). The uncontrolled group, in contrast, adapts through the shared congestion environment, thus the total population dynamics will be controlled by the interaction between the mechanism and the inner population dynamics of the two groups.

Assumptions

The objective of the operator in this auction mechanism is to select a subset of users (winner set W_t) to maximize congestion relief under a limited budget B_t, which is given as follows:

$$W_t = \arg\max_{A \subseteq C} \sum_{i \in A} \delta w_{i,t} \quad \text{subject to} \quad \sum_{i \in W_t} \phi_{i,t} \leq B_t, \tag{15}$$

$$\delta w_{i,t} = w_{\bar{s}(i),t}(\hat{x}_t) - w_{s'(i),t}(x'_t), \tag{16}$$

where $\delta w_{i,t}$ denotes the marginal congestion relief associated with user i's shift on day t. Specifically, it is measured as the difference between the class-level externality weight under the user's baseline strategy $\bar{s}(i)$ and that under the post-shift strategy $s'(i)$. These two weights are evaluated at the predicted aggregate state $\hat{x}_t$ and the expected post-shift state $x't$, respectively. $\phi i, t$ represents the reward offered to user i.
Here, to simplify the problem, we introduce the two assumptions below.

Assumption 1 First, we only consider the contribution of leaving the baseline strategy when calculating the marginal relief $\delta w_{i,t}$. This is reasonable in the context of peak-period congestion relief, where operators want to focus on shifting users away from highly congested peak-period strategies. Under this assumption, operators only need to predict the baseline $\bar{s}_{i,t}$ and its externality $w_{\bar{s}(i),t}$ at the aggregate state $\hat{x}_t$ without the incentive scheme to calculate (15). Since the aggregate state $\hat{x}_t$ is now given by the evolutional dynamics of the system as defined in (10), operators can predict $\hat{x}_t$ from the observed state x_{t-1} and predicted baseline dynamics under $\mathcal{F} = 0$.

Assumption 2 Second, we also assume that operators can predict controllable users' baseline strategy from historical data. In a context of daily morning commute, this kind of prediction will work reasonably well, as users' baseline strategies are relatively stable. With this assumption, operators can evaluate the externality weights $w_{\bar{s}(i),t}$ for each user i in the controllable group C. This setting, almost the same as the work by Jiang et al. (2024), will require additional discussion on predictive modeling and estimation, but the details of the prediction method are outside the scope of this paper.

Mechanism structure

(i) Invitation based on predicted baseline strategies

Under Assumption 2, operators can target users with high externality weights based on predicted baseline strategies $\bar{s}_{i,t}$, and thus can invite only those users to the auction. More specifically, operators define a set of "undesired" peak-period strategies $\mathcal{S}^- \subseteq \mathcal{S}$, and invite users whose predicted baseline strategy $\bar{s}_{i,t}$ belongs to this set, forming invited users set $\mathcal{I}_t$. This targeting (invitation) is important because if users with low externality are invited, they will win the auction with low bids but yield little congestion relief, causing loss of social welfare.

(ii) Bidding by invited users

Each invited user $i \in \mathcal{I}_t$ submits a bid $b_{i,t} \geq 0$, representing the minimum reward they would accept to shift from their baseline. Users who prefer not to participate do not submit a bid. In reporting their bids, truthful bidding is an essential assumption for the mechanism to work properly; its incentive compatibility is verified in the mechanism-properties subsection below. Under truthful bidding, the bid $b_{i,t}$ directly reflects the user's private shift cost $\theta_{i,t}$.

(iii) Winner selection under additive contribution approximation

With predicted marginal externality weights $w_{\bar{s}(i),t}$ and bid $b_{i,t}$ for each user i, now operators can choose winners based on the principle of cost-effectiveness $\gamma_i = b_{i,t}/\delta w_{i,t}$. This winner selection and reward allocation is implemented based on a budget-feasible mechanism proposed by Singer (2010) (Algorithm 1). Those in winner set W_t will be offered rewards $\phi_{i,t}$, and the operator expects them to shift from their baseline strategy $\bar{s}_{i,t}$ to another strategy $s'_{i,t} \notin \mathcal{S}^-$.
Here, the objective (15) treats marginal contributions $\delta w_{i,t}$ as additive across winners. In a congestion-relief problem, welfare gains are generally non-additive because of the diminishing

Algorithm 1 Budget-Feasible Allocation Rule under Additive Contribution Approximation (Singer (2010))

Require: Invited set $\mathcal{I}_t$; bids $\{b_i\}_{i\in\mathcal{I}_t}$; class-level externality weights $\{w_{s(i)}\}_{i\in\mathcal{I}_t}$; budget B_t
Ensure: Winner set W_t; rewards $\{\phi_i\}_{i\in W_t}$

1: Compute efficiency ratios: $\gamma_i \leftarrow b_i/w_{s(i)}$ for each $i \in \mathcal{I}_t$ ▷ individual bid / class-level weight
2: Sort invited users by γ_i in ascending order, and relabel them as $(1), \ldots, (|\mathcal{I}_t|)$ so that $\gamma_{(1)} \le \gamma_{(2)} \le \cdots \le \gamma_{(|\mathcal{I}_t|)}$
3: Find threshold index:

$$k^* \leftarrow \max\left\{ k \in \{1,\ldots,|\mathcal{I}_t|\} \;\middle|\; \gamma_{(k)} \le \frac{B_t}{\sum_{j=1}^{k} w_{(j)}} \right\}$$

4: Set winner set: $W_t \leftarrow \{(1), \ldots, (k^*)\}$
5: **for** each $i \in W_t$ **do**
6: Assign proportional reward: $\phi_i \leftarrow \min\left\{ \frac{w_{s(i)}}{\sum_{j\in W_t} w_{s(j)}} B_t,\ \gamma_{(k^*+1)} w_{s(i)} \right\}$
7: **end for**
8: **return** $W_t,\ \{\phi_i\}_{i\in W_t}$

marginal relief as more users shift. Hence the exact contribution of each user is set-dependent and cannot be represented by a single ex-ante scalar without specifying an ordering or a selected set. However, in the proposed approach, the operator must rank cost-effectiveness γ_i and determine payments ϕ_i before these interaction effects are realized. Therfoere, we ignore this diminishing marginal relief and treat $\delta w_{i,t}$ as additive scalars.

(iv) Ex-post reward

Winners who *actually* changed their travel strategy (verified through records) receive their promised reward $\phi_{i,t}$; those who do not shift receive nothing. As shown in next subsection, the mechanism satisfies strategy-proofness (SP) and individual rationality (IR).

Mechanism properties

Here we show that the proposed mechanism simultaneously satisfies strategy-proofness (SP), individual rationality (IR), and budget feasibility (BF), on the assumption of additive class-level externality weights. This proposition is mainly based on Myerson's theorem, which shows that single-parameter mechanism with monotonicity and threshold payment be truthful. Our model is also single-parameter, because shift cost θ is only parameter for each user. However, as mentioned in the previous subsection, we treat marginal contributions as additive values, which is an approximation and does not hold in the original congestion-relief problem and Singer's setting.

Proposition (Per-day SP, IR, BF)**.** *Algorithm 1 simultaneously satisfies **SP**, **IR**, and **BF** under the additive contribution approximation.*

Proof. **Strategy-proofness; SP.** For a fixed day t, Algorithm 1 ranks invited users by the efficiency ratio

$$\gamma_i = \frac{b_{i,t}}{w_{s(i),t}},$$

where $w_{s(i),t}$ is fixed from the viewpoint of the daily bid-reporting game. Hence each user's private information (θ_i) enters only through the scalar bid $b_{i,t}$, so the mechanism is a single-parameter reverse auction. A lower bid decreases γ_i, weakly improves the user's position in the ranking, and therefore weakly improves the chance of being selected. Thus the allocation rule is monotone in the standard single-parameter sense.
For each winning user i, the payment is given by

$$\phi_{i,t} = \min\left\{ \frac{w_{s(i),t}}{\sum_{j\in W_t} w_{s(j),t}} B_t,\ \gamma_{(k^*+1)}\, w_{s(i),t} \right\}, \tag{17}$$

which is the threshold payment associated with the ranking rule induced by γ_i. Therefore, by Myerson's characterization for single-parameter domains Myerson (1981), monotonicity together

Parameter	Value	Description
N	1600	Total number of users
μ	2.0	Schedule preference for peak travel
α	0.75	Inertia parameter
β	1.0	Congestion sensitivity
κ	2.0	Linear congestion coefficient

Table 2: Common parameters for the numerical experiments.

with threshold payments implies strategy-proofness with respect to bid reporting in the daily auction.

Individual rationality; IR. From ranking condition in the winner selection step, we have

$$\frac{b_i}{w_{s(i)}} \;\le\; \min\left\{\frac{B_t}{\sum_{j\in W_t} w_{s(j)}},\; \gamma_{(k^*+1)}\right\}. \tag{18}$$

which can be rearranged to

$$b_i \;\le\; \min\left\{\frac{w_{s(i)}}{\sum_{j\in W_t} w_{s(j)}}\cdot B_t,\; \gamma_{(k^*+1)} w_{s(i)}\right\}. \tag{19}$$

From (17), we have $b_i \le \phi_i$, meaning that the reward offered to user i exceeds their bid b_i, which is the minimum compensation they require to shift under the condition of SP.

Budget feasibility; BF. Budget feasibility follows directly from

$$\sum_{i\in W_t} \phi_i \le \sum_{i\in W_t} \frac{w_i}{\sum_{j\in W_t} w_j} B_t = B_t.$$

□

Here, this SP holds at the per-day level: truthful bidding is optimal within each day's auction, conditional on the current state and weights. However, since the auction result and the following behavior shift change future aggregate states through (10), these dynamics might cause cross-day manipulation in users' bidding process. A full analysis of repeated-game incentive compatibility is left for future work.

4 Numerical experiments

Setup

In the numerical experiments, each strategy $s \in \mathcal{S}$ corresponds to a space-time path on a transportation network $\mathcal{G} = (\mathcal{N}, \mathcal{A})$. Let $\pi_s = N_s/N$ denote the population share of class s. Compliance is assumed deterministic ($\phi_{i,t} \ge \theta_i$ implies shift). Since the auction mechanism targets congestion relief without accounting for users' schedule costs caused by post-shift strategy, the planner's objective is to minimise congestion cost alone, without schedule delay costs. The logit social optimum $\mathbf{x}^*$ is therefore defined as the minimiser of the congestion cost regularised by the logit entropy:

$$\mathbf{x}^* = \arg\min_{\mathbf{x}} \left\{ W^{\text{cong}}(f(\mathbf{x})) - \beta^{-1} \sum_s \pi_s H(\mathbf{x}_s) \right\}, \tag{20}$$

where $H(\mathbf{x}_s) = -\sum_k x_{s,k} \ln x_{s,k}$ is the entropy of the strategy distribution within class s. The entropy term ensures consistency with bounded rationality in the logit choice model.

Table 2 summarises common parameters.

Experiment 1: Case of single-link binary choice

Here we use a single-link scalar model with one population class and binary choice ($\mathcal{S} = \{\text{peak}, \text{off-peak}\}$), and we set externality weights homogeneous ($w_i = 1$ for all $i \in N$). We fix $\rho = 0.20$, $B = 0.7$ and compare four policies over 60 days. We also set private shift costs as $\theta_i \sim \text{Uniform}[0, 1.0]$.

The proposed mechanism closes 8% of the NE–SO gap in W ($= (W^{\mathrm{NE}} - \bar{W})/(W^{\mathrm{NE}} - W^{\mathrm{SO}})$, Table 3). Uniform reward distributes the budget equally among all invited users, yielding per-person rewards too small to cover any θ_i; random selection pays each winner's bid in random order until the budget is exhausted, producing marginal reduction. Proposed method outperforms both by concentrating the budget on the cheapest users, and also it shows clear convergence to a new stationary state with lower W and peak share (Figure 2).

Policy	W	NE–SO gap closed
No intervention	1.085	0%
Uniform reward	1.085	0%
Random selection	1.080	0%
Proposed (SP + BF)	**1.006**	**8%**
Logit SO	0.136	100%

Table 3: Stationary $\bar{W}$ and NE–SO gap closed under different policies (Experiment 1).

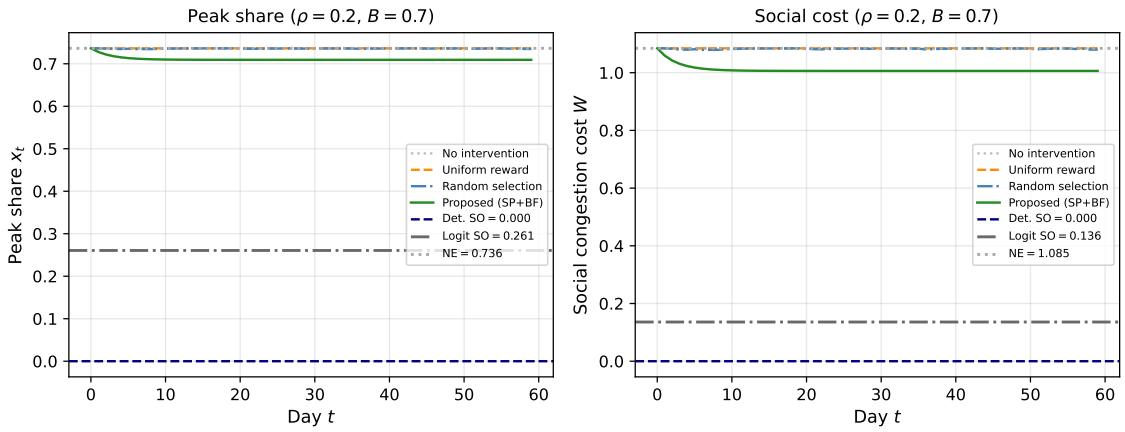


Figure 2: Convergence of peak share and social cost under different policies (Experiment 1).

Experiment 2: Case of multi-slot corridor network

We extend the model to a linear corridor network with four origin nodes $s_0, \ldots, s_3$ connected in series to a CBD destination (Figure 3). Each link has travel time $\tau = 0.5$ time units, so the total travel time from station s to the CBD is $(4 - s)\tau$; users from s_0 spend 2.0 time units in the network while those from s_3 spend only 0.5. The departure time is discretised into 6 slots at 30-minute intervals from 6:00 to 8:30. All users share a common desired arrival time $t^* = 8{:}00$ at the CBD, and schedule delay is asymmetric: $\beta_E = 1.0$ (early), $\beta_L = 2.0$ (late). Population is CBD-concentrated, with $N_s = (200, 300, 400, 700)$ for stations $s_0, \ldots, s_3$. The peak arrival window at the CBD is defined as [7:30, 8:30]; the set of departure slots classified as peak varies by origin due to the differing travel times.

Congestion cost on each link-time pair (a, τ) is linear: $c_a(f_{a,\tau}) = \kappa\, f_{a,\tau}$ with $\kappa = 2.0$ uniform across all links. Because users from distant origins traverse more links, their route congestion cost $\sum_{a \in r_s} \kappa\, f_{a,\tau_a}$ and, consequently, their externality weight w_s (eq. (7)) are structurally larger. The weights are recomputed daily from the observed flow state. Private shift costs are drawn as $\theta_i \sim \mathrm{Uniform}[0, 1.5]$; note that this experiment uses a wider support than the previous experiment ([0, 1.0]) to reflect the larger schedule delay costs in the multi-slot setting.

Figure 4 shows the convergence of W for the cases with/without considering heterogeneity of externality weights. The converged social cost W is slightly lower when considering heterogeneity, which is expected because targeting high-externality users should improve cost-effectiveness.

Figure 5 compares results of with/without this heterogeneity under different settings of budget constraints and penetration rate. The illustrated resluts are reasonable: decreasing of converged social costs as budget and penetration increase. The green region illustrates a clear improvement in congestion relief by considering heterogeneity, but there is not an obvious relationship with budget or penetration rate. We can also see that as budget and penetration rate increase, the converged

social cost gets closer to the logit SO line. As same as in Experiment 1, there are large gaps between the converged scores and those of SO, but it depends on the setting of hyperparameters. Figure 6 focuses on the overpayment problem caused by the proportional payment rule (Algorithm 1). It shows that around 50% of the payment is information rent, implying total payment is roughly twice the true compensation cost. This overpayment is a structural consequence of the algorithm's design to ensure strategy-proofness and individual rationality: the reward ϕ_i must exceed the user's private cost θ_i to incentivise truthful bidding, but the proportional rule inflates rewards for high-w_i users. It also illustrates the number of winners per day, which increases almost linearly with budget, as expected. This trade-off between mechanism properties and efficiency is a fundamental aspect of the design: while the proportional payment rule ensures SP and IR, it can lead to significant overpayment. Also, looking at the number of winners, the mechanism under this setting is very selective, resulting in only about 30 users in the total of 1600 users being selected as winners per day. This is one possible reason why the improvement in W is limited compared to the logit SO benchmark.

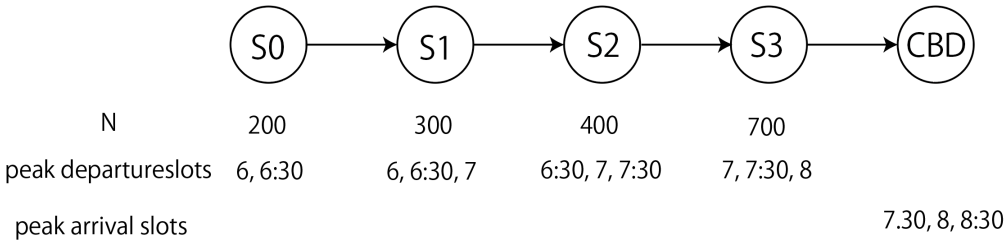


Figure 3: Linear corridor with four origin nodes to CBD (for Experiment 2).

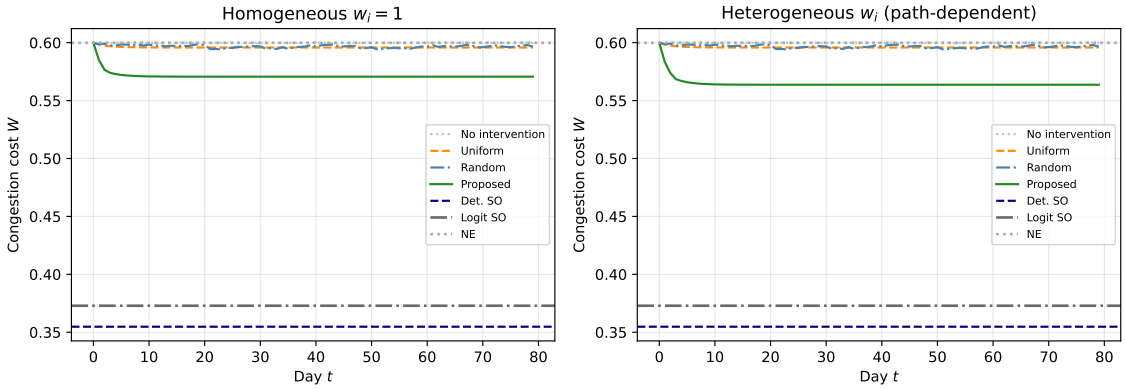


Figure 4: W trajectory for homogeneous vs. heterogeneous w_i ($\rho = 0.30$, $B = 1.5$).

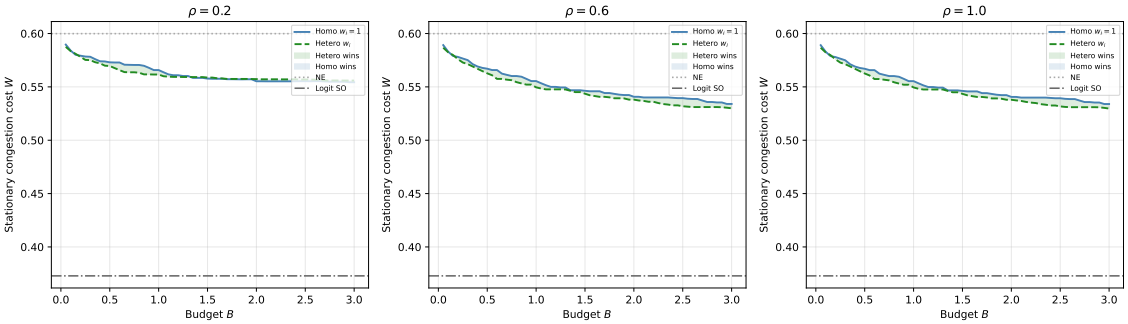


Figure 5: Stationary congestion cost W under homogeneous ($w_i = 1$) and heterogeneous w_i as a function of budget B, for $\rho \in \{0.2, 0.4, 0.6\}$. Green (blue) shading indicates regions where heterogeneous (homogeneous) weights achieve lower W.

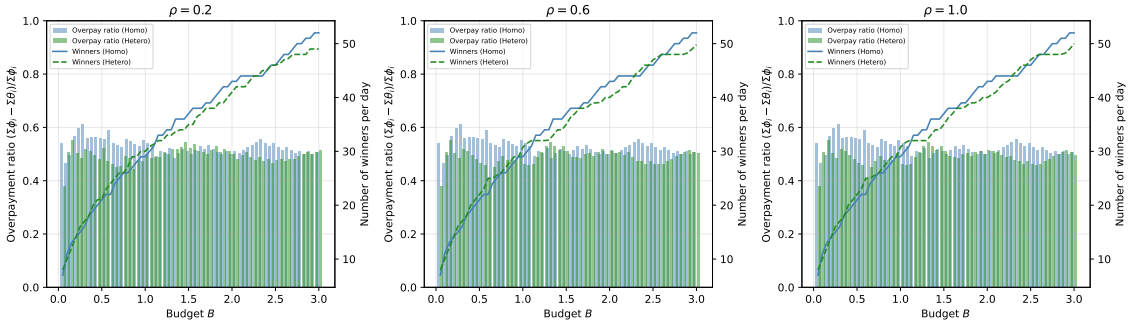


Figure 6: Overpayment ratio (bars, left axis) and number of winners per day (lines, right axis) under homogeneous and heterogeneous w_i, as a function of budget B.

5 Discussion and Conclusions

This paper proposed a day-to-day reverse auction mechanism for congestion control under partial controllability. Under additive congestion assumptions, the mechanism satisfies SP, IR, and BF through an externality-weighted Singer rule. Numerical experiments suggest that even limited controllability can move the system toward the social optimum through feedback to the uncontrolled group.

Several extensions remain. The payment rule should be refined to reduce overpayment. Marginal contribution should also account for the social cost of the destination strategy, not only that of leaving the baseline. In repeated settings, additional devices may be needed to deter strategic bidding across days. Since the present case study relies on simplified numerical assumptions, validation in more realistic settings is essential.

A further issue is the submodularity of marginal contributions. Here they are treated additively to determine the auction outcome ex ante. In reality, congestion relief is generally set-dependent and exhibits diminishing returns. Incorporating this structure may better approximate the original welfare problem and improve efficiency.

Acknowledgements

This research was performed by the commissioned research fund provided by F-REI (JPFR 26-05-07-01).

References

Clarke, E. H. (1971). Multipart pricing of public goods. *Public choice*, 17–33.

Groves, T. (1973). Incentives in teams. *Econometrica: Journal of the Econometric Society*, 617–631.

Jiang, W., Koutsopoulos, H. N., & Ma, Z. (2024). A reverse auction-based individualized incentive system for transit mobility management. *IEEE Transactions on Intelligent Transportation Systems*.

Matsunaga, T., & Hato, E. (2026, January). Dynamic congestion control in metropolitan railways using generative language models: Approach by bid-free reverse auctions. In *Transportation research board annual meeting (trb 2025)*. Washington, D.C., United States.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of operations research*, *6*(1), 58–73.

Pigou, A. C. (1920). *The economics of welfare* (First ed.). London: Macmillan.

Sandholm, W. H. (2002, 07). Evolutionary implementation and congestion pricing. *The Review of Economic Studies*, *69*(3), 667-689. Retrieved from `https://doi.org/10.1111/1467-937X.t01-1-00026` doi: 10.1111/1467-937X.t01-1-00026

Seo, T., & Yin, Y. (2019). Optimal pricing for departure time choice problems with unknown preference and demand: Trial-and-error approach. In *Transportation research board 98th annual meeting.*

Singer, Y. (2010). Budget feasible mechanisms. In *2010 ieee 51st annual symposium on foundations of computer science* (pp. 765–774).

Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of finance*, *16*(1), 8–37.

Wu, M., Wang, B., Yin, Y., & Lynch, J. P. (2024). Participatory traffic control: Leveraging connected and automated vehicles to enhance network efficiency. *Transportation Research Part C: Emerging Technologies*, *166*, 104757. Retrieved from `https://www.sciencedirect.com/science/article/pii/S0968090X2400278X` doi: https://doi.org/10.1016/j.trc.2024.104757