limos lab for innovative mobility systems

TRISTAN XII, Okinawa, Japan, June 22, 2025

Participatory Traffic Control with Connected and Automated Vehicles for Network Efficiency

Yafeng Yin, Ph.D.
Donald Cleveland Collegiate Professor of Engineering
Donald Malloure Department Chair of Civil and Environmental Engineering
University of Michigan, Ann Arbor

Contributors

- **Minghui Wu**, Richard Lee, Xi Lin, and Jeff Scruggs, University of Michigan, Ann Arbor
- Zhibin Chen, New York University Shanghai
- Jerry Lynch, Duke University
- Meng Li, Tsinghua University

Outline

1 Participatory Traffic Control

2 CAVs as Demand Distributors

3 Recruiting CAVs for Participation

4 Controlling CAVs for Network Efficiency

5 Concluding Remarks

Participatory Traffic Control

- Utilizing connected and automated vehicles (CAVs) within the traffic stream as control actuators to manage traffic
- CAVs as traffic stream regulator
  - Wu et al. (2018), Stern et al. (2018) and Jin et al. (2018)
  - "Flow control will be possible via a few mobile actuators (less than 5%)" (Stern et al., 2018)

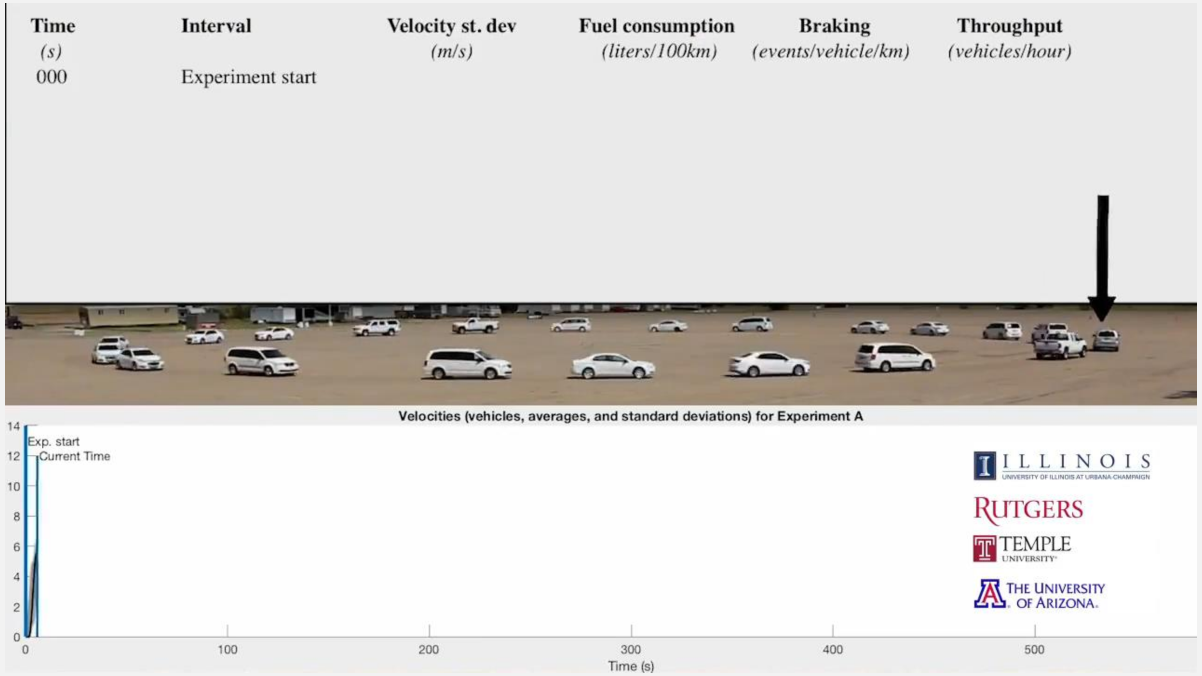

Source: https://flow-project.github.io/

Leveraging Relinquished Human Agency in Travel

- Human agency is the capacity for individuals to act independently and make their own decisions based on personal volition.
- This concept is highly relevant to travel, encompassing multiple facets such as the timing of travel, choice of destination, mode, and route, driving tasks, and adaptability.
- While traditionally personal vehicles have been marketed with cultural symbols of security, freedom, and control, the advent of driving automation requires travelers to increasingly surrender aspects of their agency.

Leveraging Relinquished Human Agency

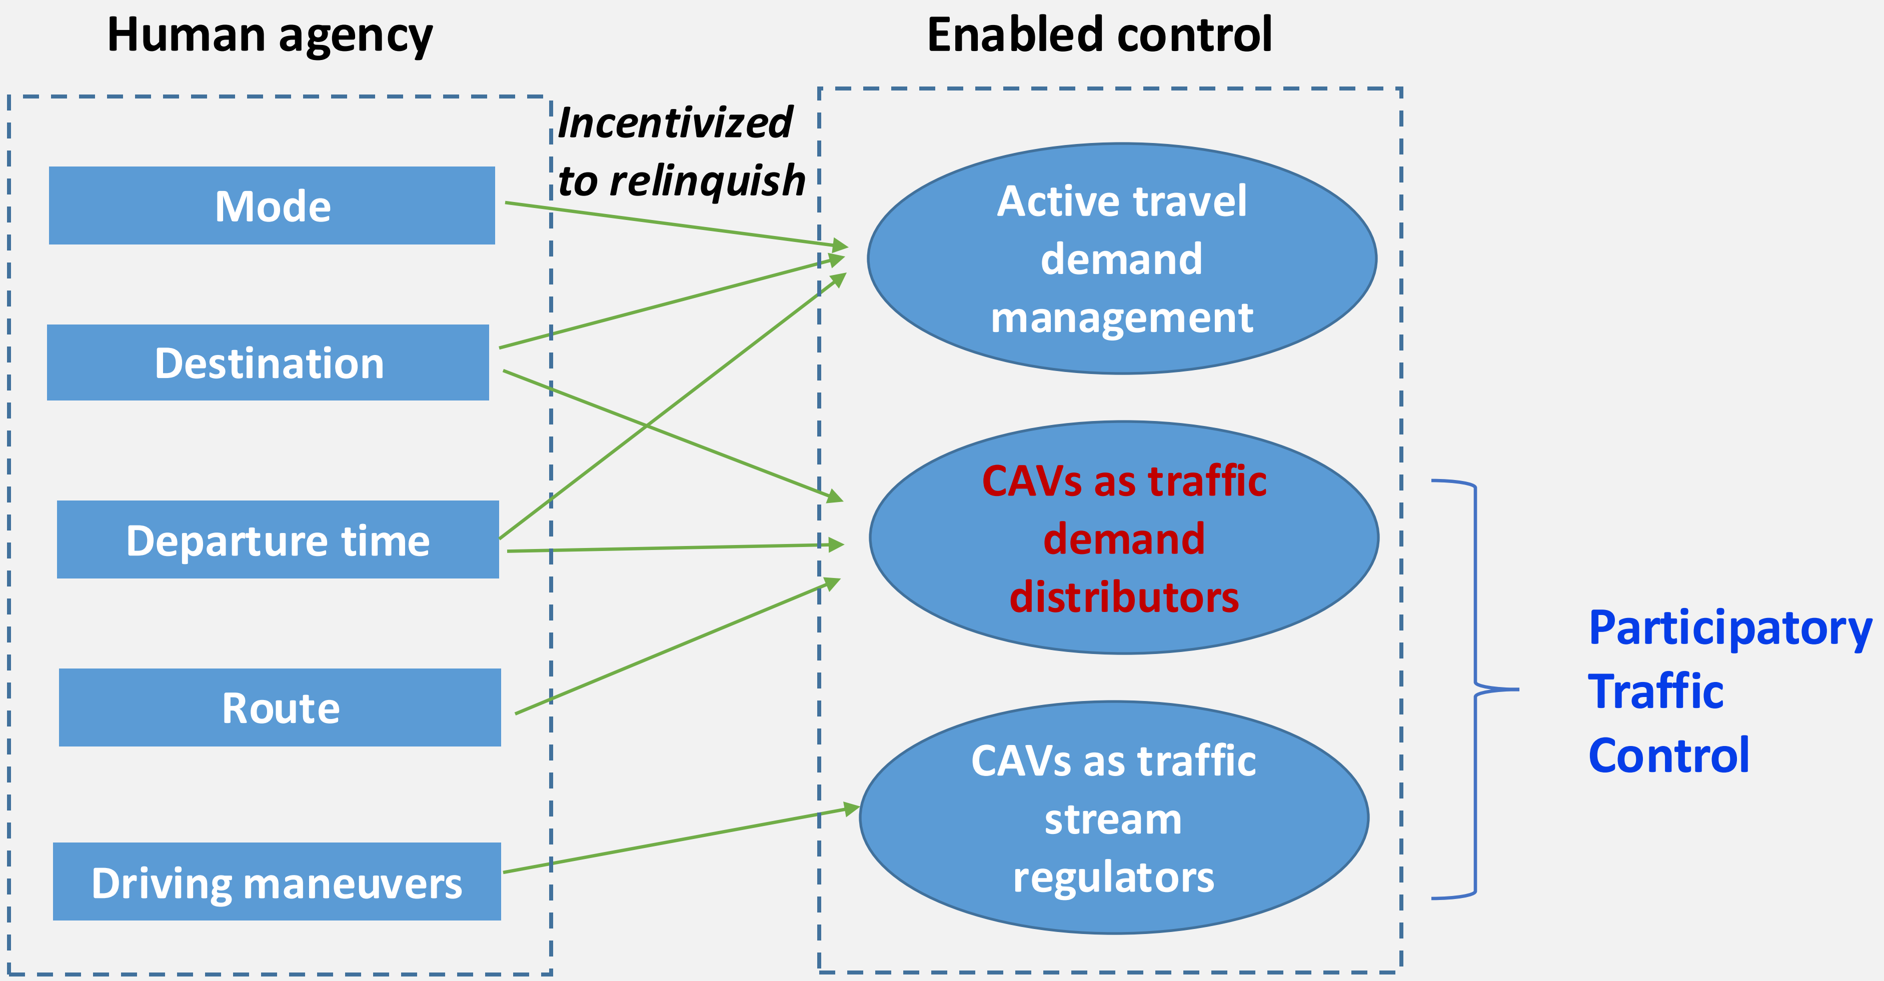

CAVs as Traffic Demand Distributors

- The working hypothesis is that, by controlling the departure time or route choice of a small number of CAVs, we can influence a larger number of uncontrolled vehicles' travel decisions to improve the overall system performance.

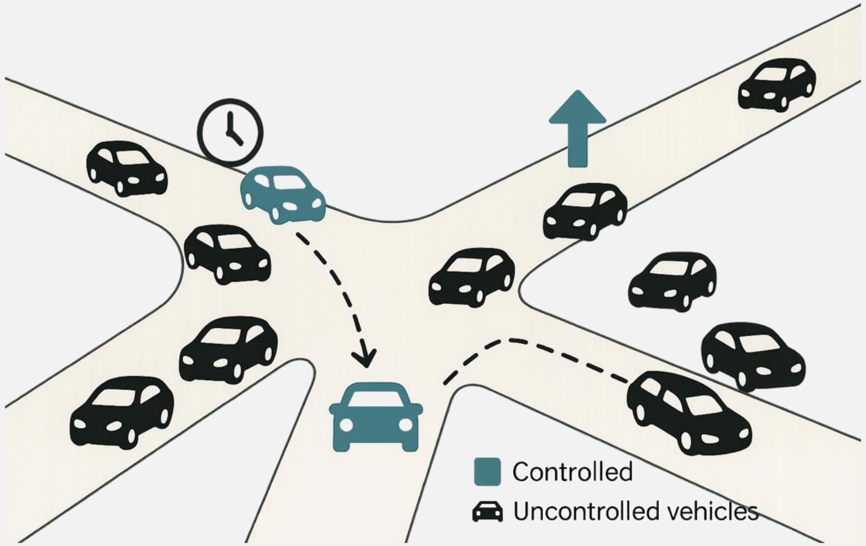

What percentage of traffic do we need to control?

Let's examine the minimum percentage of traffic that needs to be controlled to replicate system optimum in a traffic network

Wardrop's Principles of Traffic Assignment*

JOHN WARDROP
(1922–1989)

Feature	First Principle (User Equilibrium)	Second Principle (System Optimum)
Behavior Type	Selfish (user optimal)	Altruistic or centrally controlled
Stability	Nash Equilibrium	Requires control to be stable
Realism	More realistic for decentralized systems	Ideal goal for planning or control
Total System Cost	Higher than or equal to SO	Minimal possible total travel time
Formulation	$\min z = \sum_a \int_0^{v_a} t_a(\omega)d\omega$ s.t. $\sum_{p\in P^w} h_p = q^w \quad \forall w \in W$ $h_p \geq 0 \quad \forall p \in P$ $v_a = \sum_{p\in P} h_p \delta_a^p \quad \forall a \in A$ $q^w \geq 0 \quad \forall w \in W$	$\min z = \sum_a t_a(v_a)v_a$ s.t. $\sum_{p\in P^w} h_p = q^w \quad \forall w \in W$ $h_p \geq 0 \quad \forall p \in P$ $v_a = \sum_{p\in P} h_p \delta_a^p \quad \forall a \in A$ $q^w \geq 0 \quad \forall w \in W$

** Wardrop J. (1952) Some Theoretical Aspects of Road Traffic Research. Proceedings, Institution of Civil Engineers II(1), pp.325-378.*

Controlling CAVs to Replicate System Optimum

- What is the minimum percentage of traffic that needs to be controlled?
  - Cooperative AVs (CAVs): System optimum routing principle
  - Uncontrolled Vehicles (UVs): User optimal routing principle

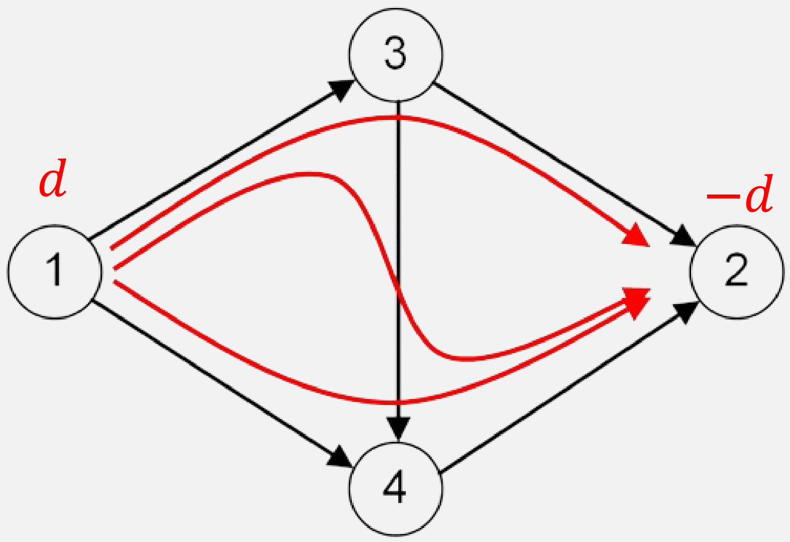

Braess network

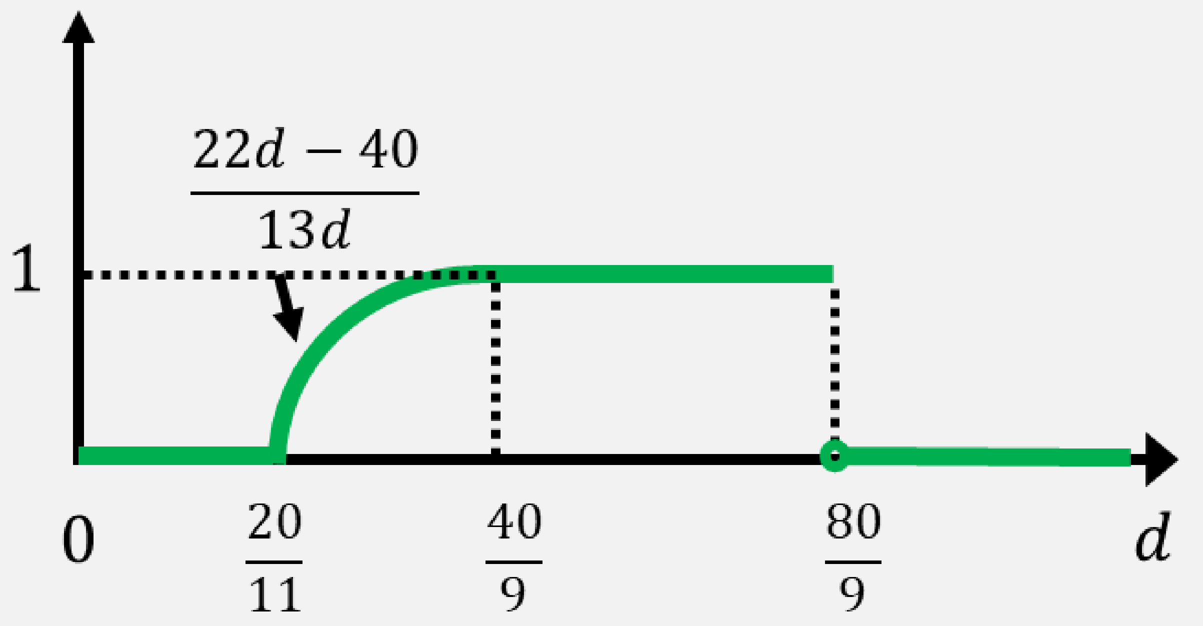

Minimum Control Ratio (MCR) Formulation

lab for innovative mobility systems

$$\min_{\boldsymbol{f}^U,\boldsymbol{f}^C,\boldsymbol{d}^U,\boldsymbol{d}^C,\boldsymbol{v}^U,\boldsymbol{v}^C} \sum_{w\in W} d_w^C$$

$s.t.$

$$v_a^U = \sum_{w\in W}\sum_{r\in \hat{R}_w} f_r^U \delta_{a,r} \qquad \forall a \in A \qquad (1)$$

Set of paths with minimum travel time

$$\sum_{r\in \hat{R}_w} f_r^U = d_w^U \qquad \forall w \in W \qquad (2)$$

$$f_r^U \geq 0 \qquad \forall r \in \boxed{\hat{R}_w}, w \in W \qquad (3)$$

$$v_a^C = \sum_{w\in W}\sum_{r\in \bar{R}_w} f_r^C \delta_{a,r} \qquad \forall a \in A \qquad (4)$$

Set of paths with minimum marginal travel time

$$\sum_{r\in \bar{R}_w} f_r^C = d_w^C \qquad \forall w \in W \qquad (5)$$

$$f_r^C \geq 0 \qquad \forall r \in \boxed{\bar{R}_w}, w \in W \qquad (6)$$

$$v_a^U + v_a^C = \bar{v}_a \qquad \forall a \in A \qquad (7)$$

$$d_w^U + d_w^C = \bar{d}_w \qquad \forall w \in W \qquad (8)$$

MCR in Real-World Networks*

Network	Zones	Nodes	Links	MCR (%)
Anaheim	38	416	914	**21**
Barcelona	110	1020	2522	**36**
Berlin-Mitte-Prenzlauer Berg-Friedrichshain-Center	98	975	2184	**14**
Eastern-Massachusetts	74	74	258	**20**
Sioux Falls	24	24	76	**14**
Terrassa	55	1609	3264	**57**
Winnipeg	147	1052	2836	**42**

Chen, Z., Lin, X, **Yin, Y. and Li, M. (2020) Path controlling of automated vehicles for system optimum on transportation networks with heterogeneous traffic stream. Transportation Research Part C: Emerging Technologies, 110, 312-329.*

Recruiting CAVs for Participation

- Drivers delegating control to machines does not necessarily mean they are willing to sacrifice personal benefits for system-level improvements.
- Incentives may be necessary to encourage CAV participation in traffic control schemes. These incentives should compensate drivers for their loss of agency and potential increases in travel costs.
- So, how can we design effective incentive schemes to recruit CAVs for implementing participatory traffic control?

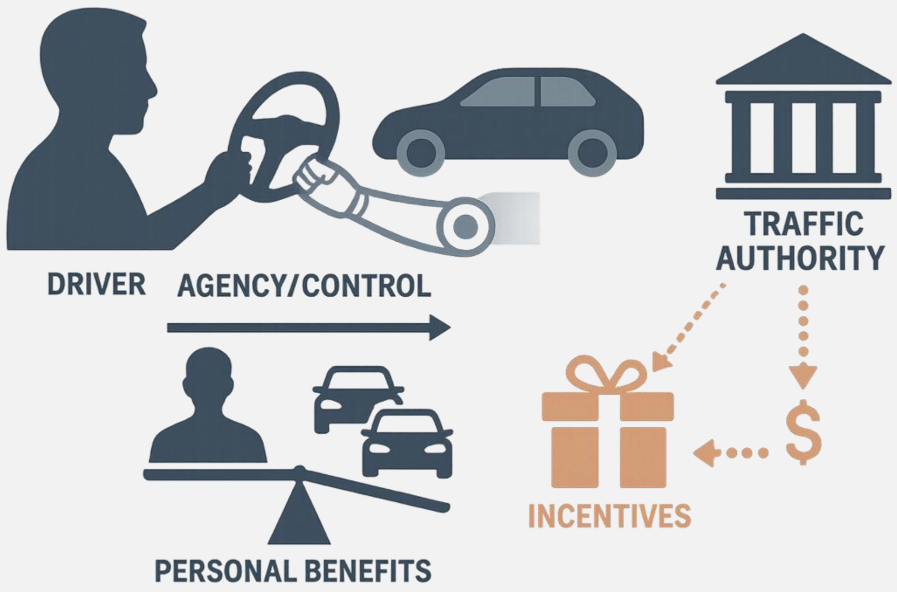

Recruiting CAVs for Participation

Morning Commute Setting (Vickrey, 1969)

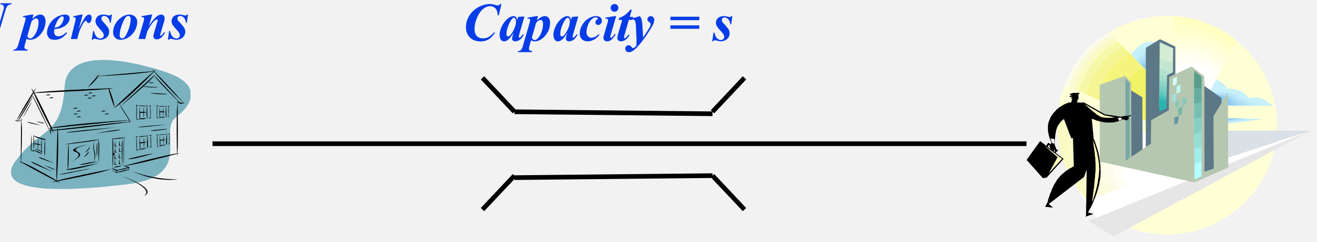

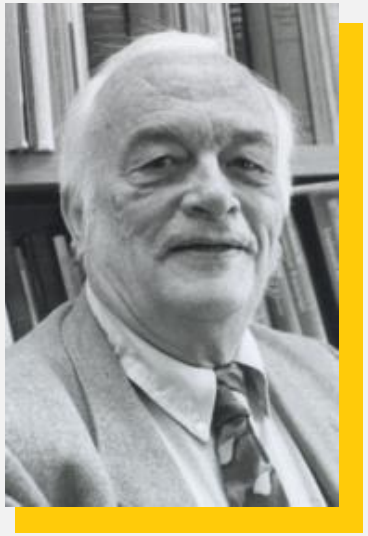

WILLIAM VICKREY
(1914-1996; Nobel Laureate, 1996)

Early/Late Arrival Penalty

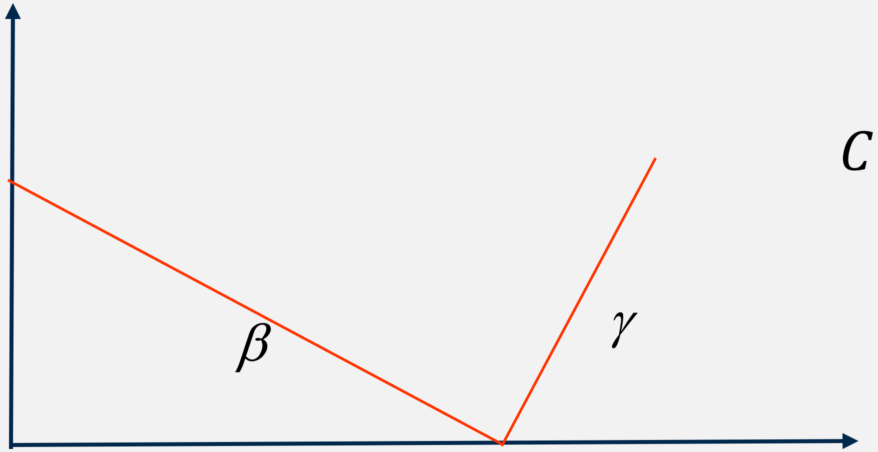

$$C(t) = \begin{cases} \alpha T(t) + \beta(t^* - t - T(t)) & t + T(t) \le t^* \\ \alpha T(t) + \gamma(t + T(t) - t^*) & t + T(t) \ge t^* \end{cases}$$

$$\gamma > \alpha > \beta$$

$$T(t) = \frac{Q(t)}{s}$$

Equilibrium Departure Time Choice

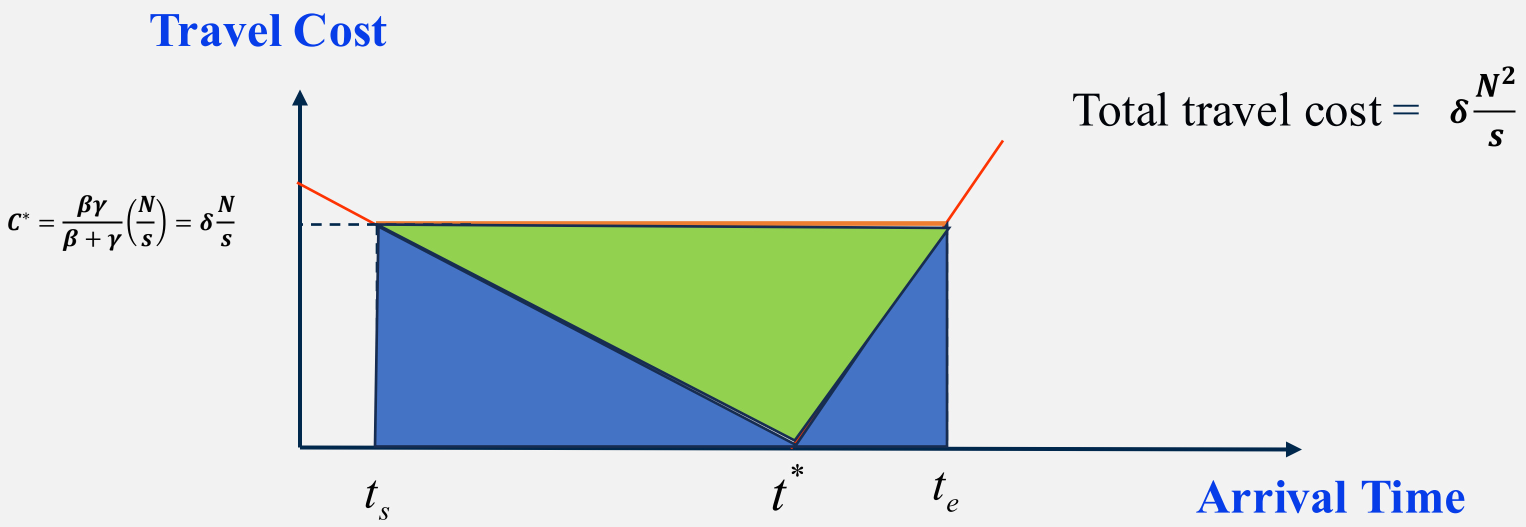

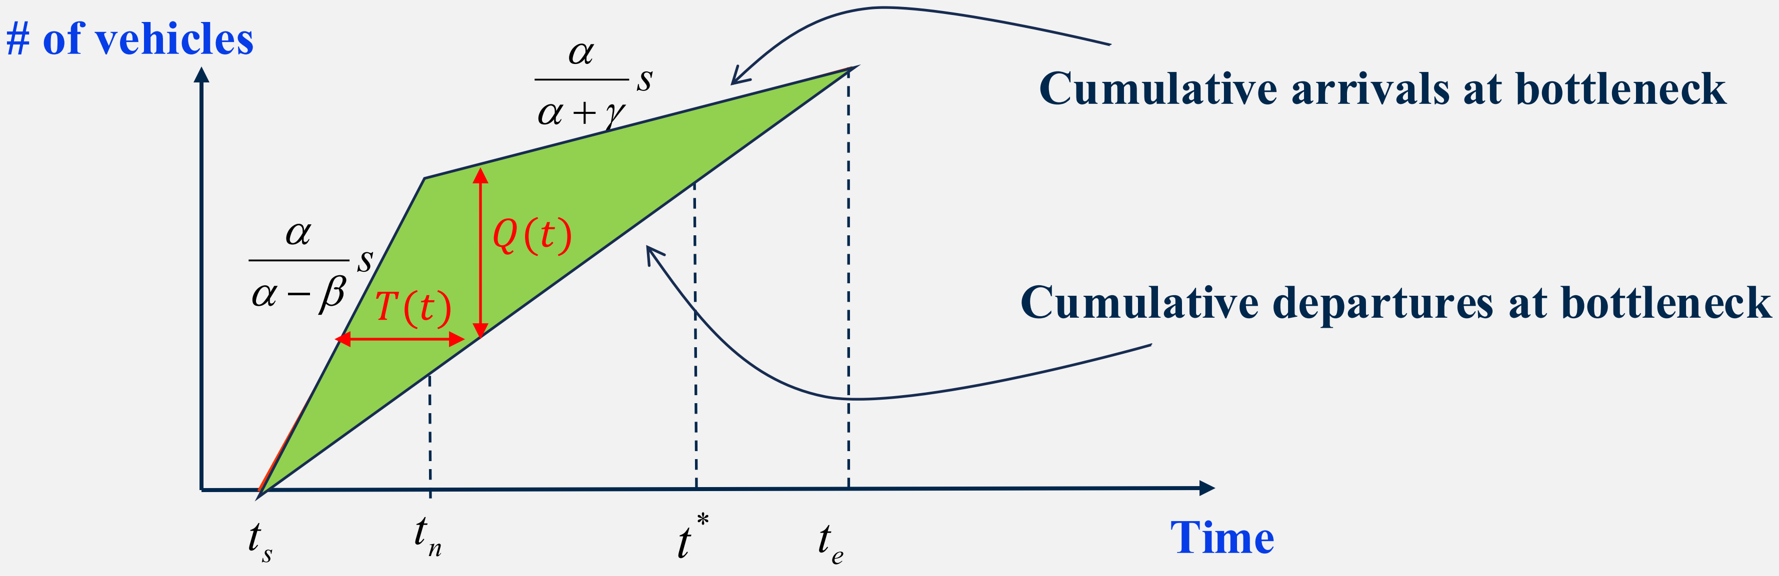

Queueing Diagram at Bottleneck

Equilibrium vs Optimum

Equilibrium

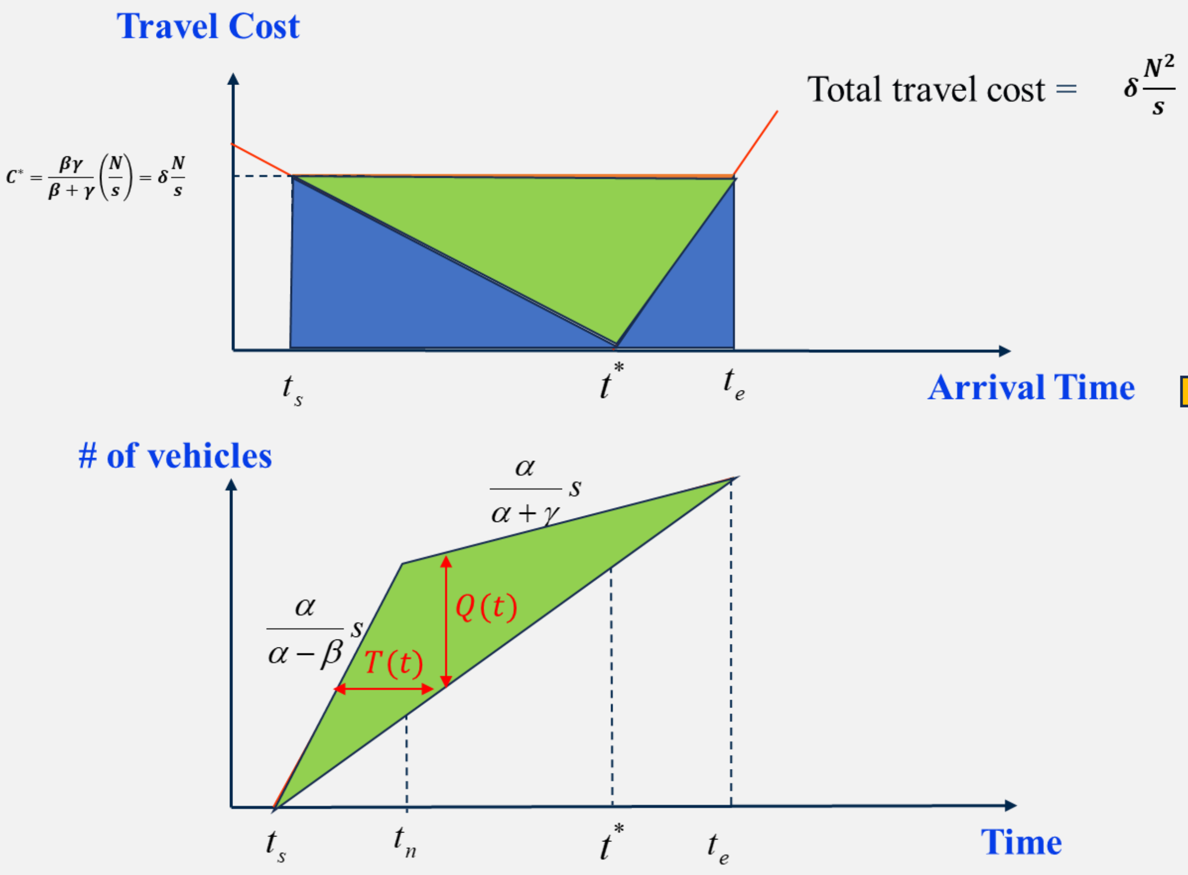

Queueing Diagram at Bottleneck

Optimum

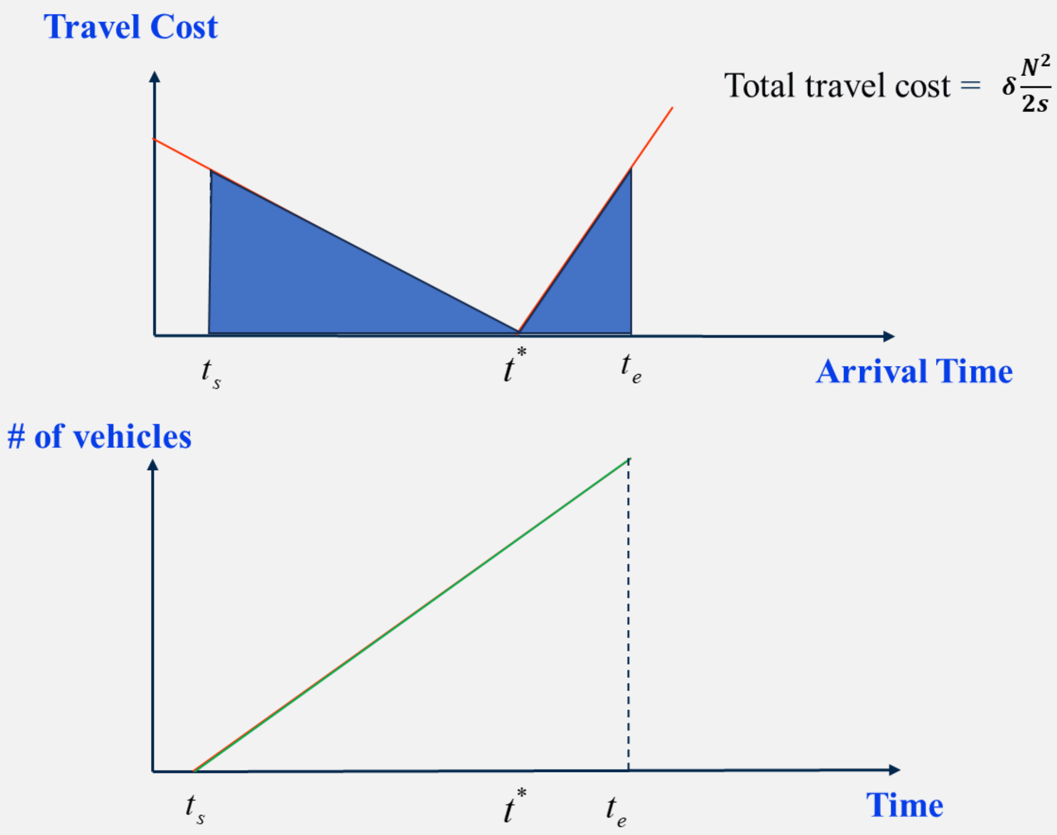

Queueing Diagram at Bottleneck

Recruiting CAVs for Participatory Traffic Control

Group 1: Providing incentive p in exchange for both departure time and maneuver control

- xN recruited users
- Heterogeneous cost of agency $c_d \in [d, D]$ with PDF $f(c_d)$, CDF $F(c_d)$, and inverse CDF F^{-1}

Group 2: Providing incentive q in exchange for driving maneuvers only

- yN recruited users
- Cost of agency $c_m = \rho c_d$ with $\rho < 1$

Participatory Traffic Control

- Departure time: placing Group 1 at the shoulders of the departure window
- Maneuver: uniformly mixing Group 2 in the traffic stream with UVs
- The bottleneck capacity $s(w)$ is an increasing function of the mixing rate of CAVs w due to the smoothing effect. Denote $s(0) = s_m < s(1) = s_M = rs_m$, $r(> 1)$ is the maximum capacity improvement ratio

Equilibrium Solutions under Control

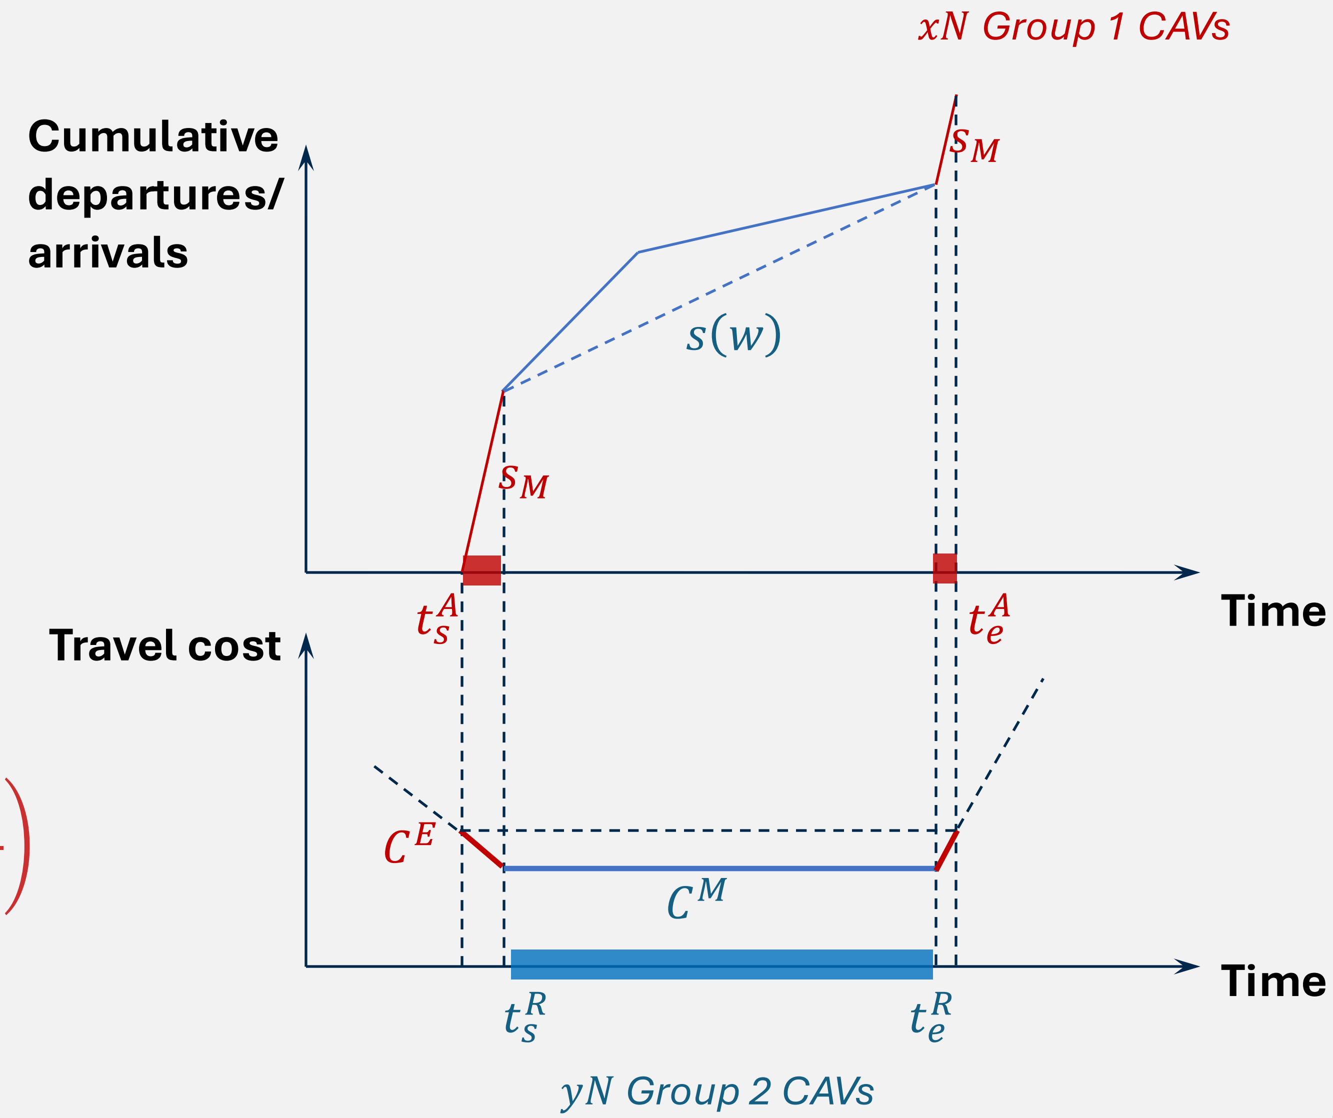

Average travel cost of Group 1 recruits

$$C_E = \delta N \left(\frac{1-x}{s(w)} + \frac{x}{2s_M} \right)$$

Equilibrium travel cost of remaining commuters

$$C_M = \delta \frac{N(1-x)}{s(w)}$$

First-Best Design

$$\min_{x,y} \delta N^2 \underbrace{\left(\frac{1-x}{s\left(\frac{y}{1-x}\right)}+\frac{x^2}{2s_M}\right)}_{\textit{Total travel cost}} + \underbrace{\theta N \int_d^{F^{-1}\left(\frac{x}{\theta}\right)} cf(c)dc + \rho\theta N \int_{F^{-1}\left(\frac{x}{\theta}\right)}^{F^{-1}\left(\frac{x+y}{\theta}\right)} cf(c)dc}_{\textit{Total cost of agency}} + \underbrace{\kappa(x+y)N}_{\textit{Operation cost}}$$

$$s.t. \qquad x+y \le \theta, x, y \ge 0$$

Specification of $s(w)$ and F:

- The smoothing effect is marginal increasing

$$s(w) = \frac{s_M}{-(r-1)w+r}$$

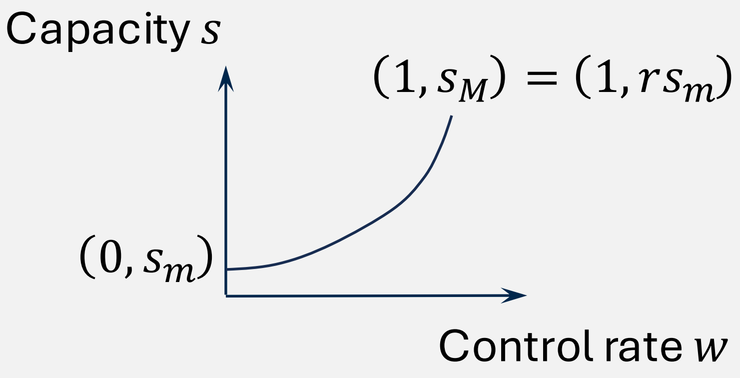

- Uniformly distributed c_d

Analytical Solutions

Conditions		x^*	y^*
$B < 0$		$clamp_{[0,\theta]}(C)$	0
$B \in [0, \theta]$	$B < A$	$clamp_{[0,\theta]}(C)$	0
	$B \geq A$	$clamp_{[0,\theta]}(A)$	$B - clamp_{[0,\theta]}(A)$
$B \geq \theta$		$clamp_{[0,\theta]}(A)$	$\theta - clamp_{[0,\theta]}(A)$

Unconstrained recruitment target for Group 1

Unconstrained total recruitment target

Adjusted recruitment target for Group 1 with Group 2 not being activated

where $A = \frac{\theta\delta N - (1-\rho) r s_m \theta \frac{Dd}{D-d}}{\theta\delta N + (1-\rho) r D s_m}$, $B = \frac{\left(r\left(1-\frac{s_m\kappa}{\delta N}\right)-1\right)\theta\delta N}{\rho r D s_m} - \theta\frac{d}{D-d}$, $C = \frac{r\left(1-\frac{s_m\kappa}{\delta N}\right)\theta\delta N - r s_m \theta\frac{Dd}{D-d}}{\theta\delta N + r D s_m}$

$$clamp_{[0,\theta]}(X) = \min\{\max\{0, X\}, \theta\}.$$

Optimal incentives:

$$p^* = \left(\frac{D-d}{\theta} + \frac{\delta N}{2 r s_m}\right) x^* + \frac{\rho(D-d)}{\theta} y^* + d \qquad q^* = \frac{\rho(D-d)}{\theta}(x^* + y^*) + \rho d$$

Managerial Insights – Economic Feasibility

- The incentive scheme is always active if there exist drivers whose cost of agency is sufficiently small (e.g., the lower bound of the cost $d = 0$) and the unit operation cost is sufficiently low compared to the equilibrium travel cost ($\kappa \ll \delta \frac{N}{s_m}$)
  - No recruitment arises when $B < 0, C < 0$. Since $r > 1$, these inequalities cannot hold simultaneously if both conditions are met
- Empirical studies indicate that more than 60% of people are willing to take prosocial actions when there is no personal cost. This suggests a high likelihood of the existence of drivers with sufficiently low or zero agency costs.
- **Recruiting CAVs for participatory traffic control likely makes economic sense for rush hour congestion management….**

Managerial Insights – Economic Feasibility

- **…. especially in large cities!**
  - The number of recruits increases as *N* increases

$$B = \frac{\left(r\left(1 - \frac{s_m \kappa}{\delta N}\right) - 1\right)\theta \delta N}{\rho r D s_m} - \theta \frac{d}{D - d}$$

  - In high-demand scenarios, travel cost dominates, and participatory control offers greater leverage to reduce the total cost

$$\min_{x,y} \underbrace{\delta N^2 \left(\frac{1-x}{s\left(\frac{y}{1-x}\right)} + \frac{x^2}{2 s_M}\right)}_{\textit{Total travel cost}} + \underbrace{\theta N \int_{d}^{F^{-1}\left(\frac{x}{\theta}\right)} c f(c) dc + \rho \theta N \int_{F^{-1}\left(\frac{x}{\theta}\right)}^{F^{-1}\left(\frac{x+y}{\theta}\right)} c f(c) dc}_{\textit{Total cost of agency}} + \underbrace{\kappa (x+y) N}_{\textit{Operation cost}}$$

Managerial Insights – Hiring Strategy

- **We should generally prioritize recruiting Group 1, even though they are more costly to hire**.
  - This preference is driven by their marginally increasing smoothing effect and the presence of drivers with low agency costs.
  - When $d = 0$, we always have Group 1 recruits. Positioning them at the shoulders, where the mixing rate is 100%, yields the greatest benefits.

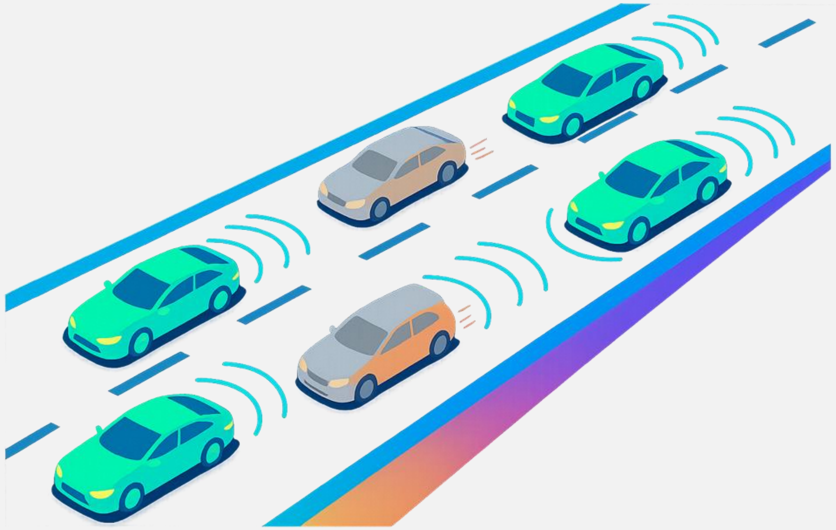

Managerial Insights – Hiring Strategy

- **The recruitment focus shifts toward Group 2 as the smoothing effect becomes stronger**
  - A higher r leads to more recruits

$$B = \frac{\left(r\left(1 - \frac{s_m \kappa}{\delta N}\right) - 1\right)\theta \delta N}{\rho r D s_m} - \theta \frac{d}{D - d}$$

  - On the other hand, a higher r leads to fewer recruits in Group 1, implying that the share of Group 1 in the recruits drops

$$A = \frac{\theta \delta N - (1-\rho) r s_m \theta \frac{Dd}{D - d}}{\theta \delta N + (1-\rho) r D s_m}$$

  - When r increases, recruiting Group 2 becomes more cost-effective, as they are less costly but become more effective

Numerical Validation

Scenarios		Smoothing effect: Marginal increasing	Smoothing effect: Linear	Smoothing effect: Marginal decreasing
Cost of agency	Uniform	(a)	(b)	(c)
Cost of agency	Truncated normal	(d)	(e)	(f)

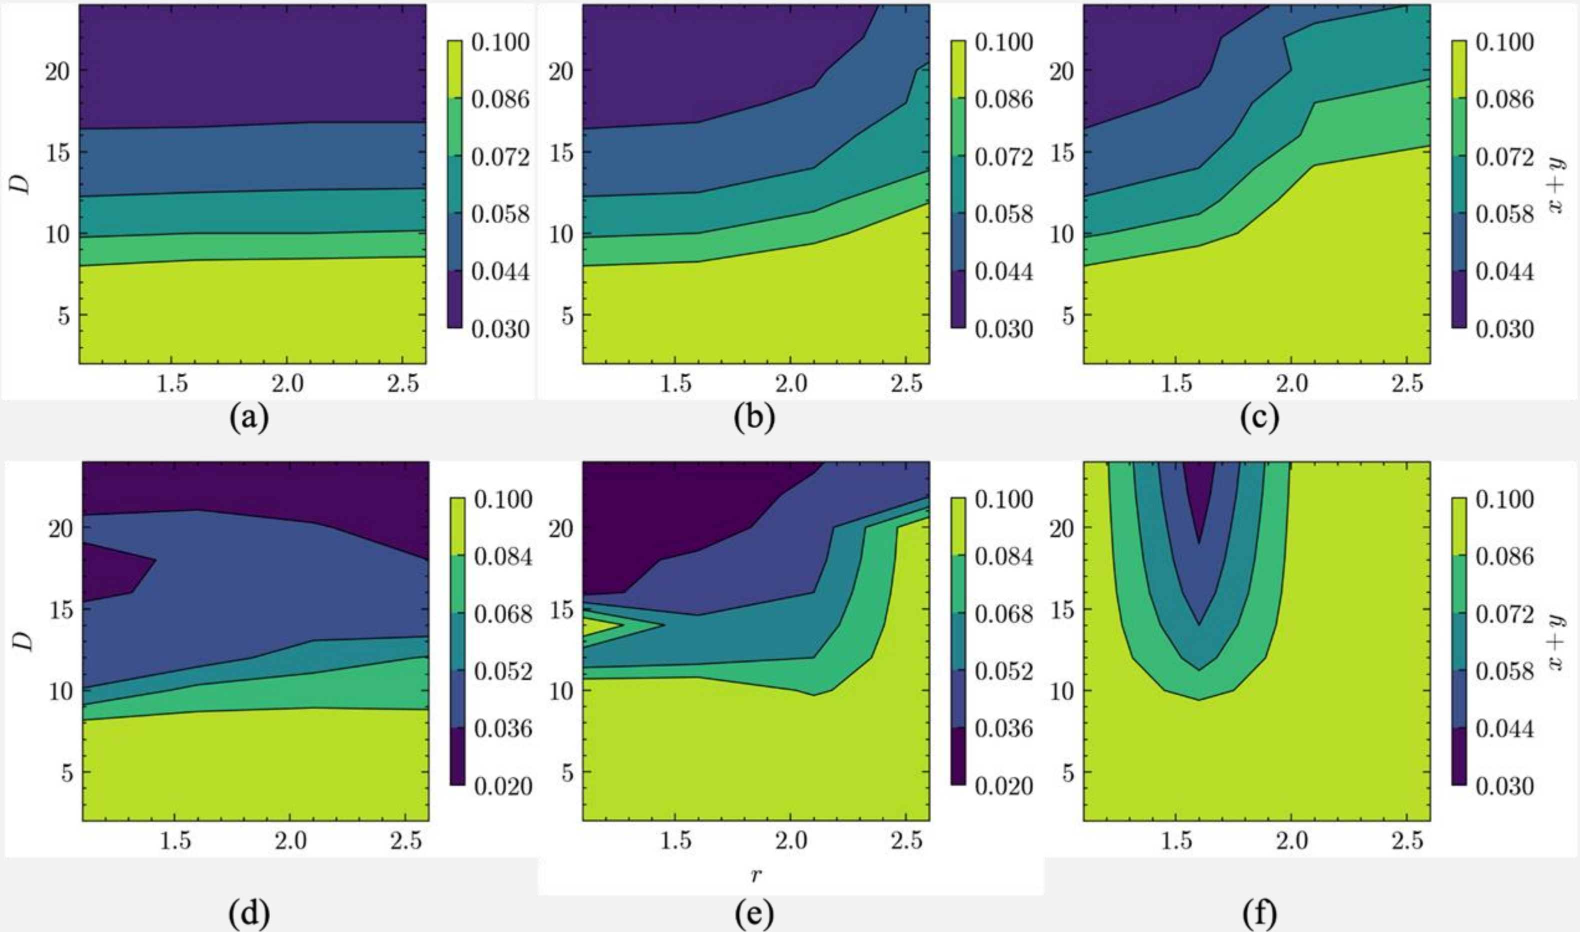

At $d = 0$, $\rho = 0.9$ and 10% AV market penetration, the recruitment level stays above 2% across all cases when varying D and r

Second-Best Solution

Scenarios		Smoothing effect		
		Marginal increasing	Linear	Marginal decreasing
Cost of agency	Uniform	(a)	(b)	(c)
	Truncated normal	(d)	(e)	(f)

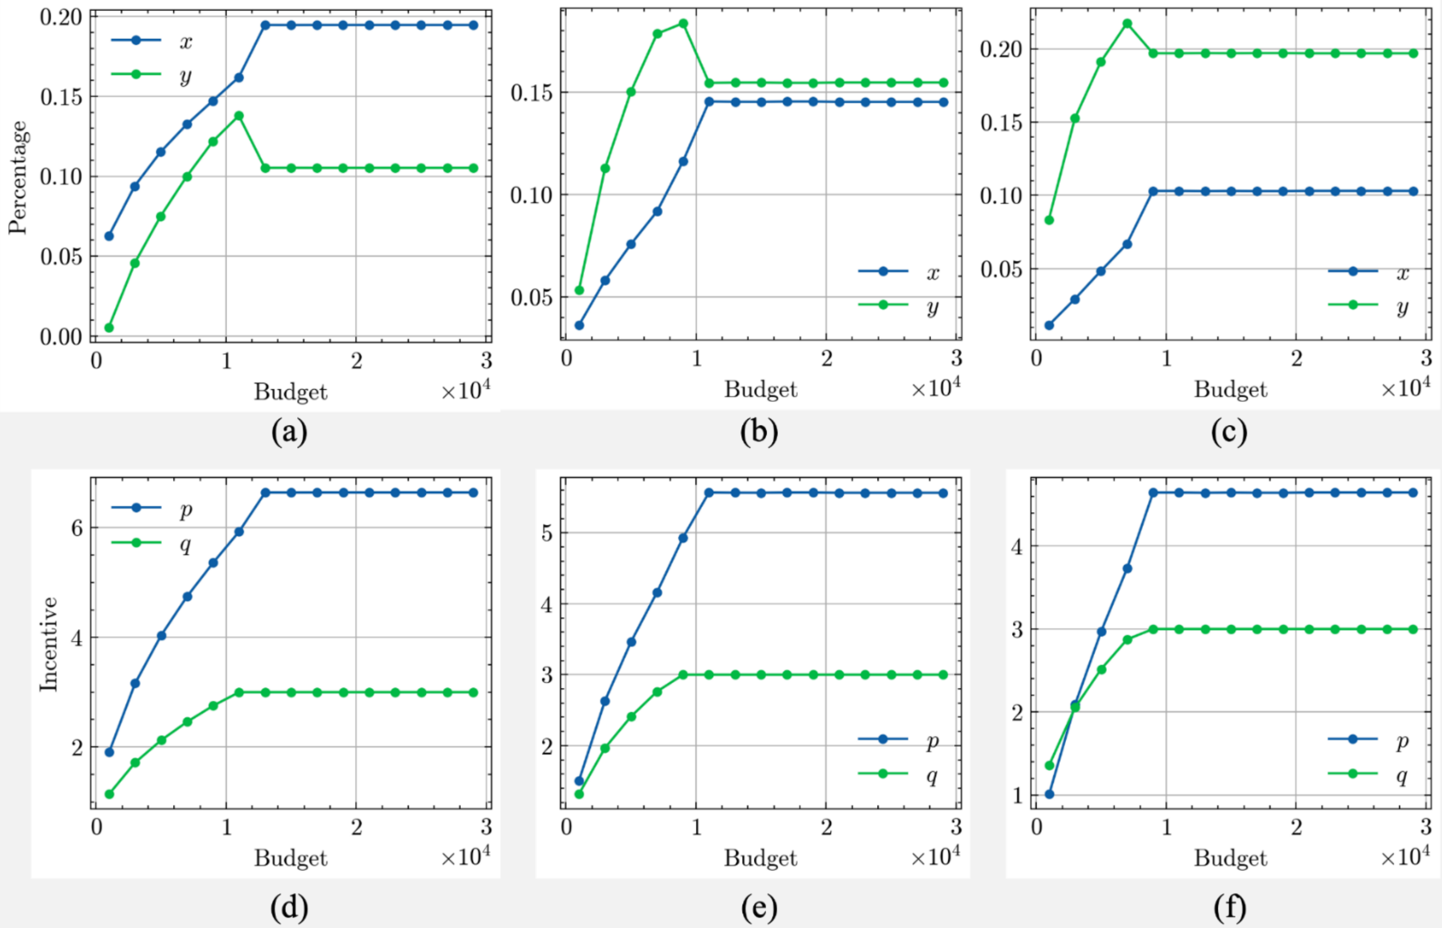

At $d = 2, D = 10, \rho = 0.3$, $r = 1.6$, and 30% AV market penetration, both groups see higher recruitment and incentives as the budget increases

Controlling CAVs for Network Efficiency

Problem Setting

- We consider a general network where the routes or/and departure times of a fraction of CAVs can be controlled to reduce cumulative system travel cost over time.
- Uncontrolled vehicles (UVs) behave selfishly, adjusting their choices day-to-day following a certain dynamics like the one proposed by Smith (1984), though this process is not necessarily known to the traffic authority.
- We thus aim for a model-free, distributed control policy that lets CAVs act based on local information.

MIKE SMITH
(2007 INFORMS Robert Herman Lifetime Achievement Awardee)

Finite-Agent Optimal Control Problem

Total discounted cost over time

$$\min_{\pi} J(\pi) = E\left[\sum_{t=0}^{\infty} \gamma^t C(\mu_t^N, \nu_t^M)\right]$$

$$s.t.\ \ \nu_{t+1}^M = q(\mu_t^N, \nu_t^M)$$ Transition kernel for uncontrolled vehicles

$$a_t^i \sim \pi(\cdot | x_t^i, \mu_t^N, \nu_t^M) \quad i \in [N]$$ Sample action from CAV routing policy

$$x_{t+1}^i \sim p(\cdot | x_t^i, a_t^i, \mu_t^N, \nu_t^M) \quad i \in [N]$$ Transition kernel for CAVs

Proposition

If the transition kernels p and q, as well as the system cost C are Lipschitz continuous w.r.t μ and ν, there always exists an optimal policy

Finding such a policy becomes intractable as the number of agents grows large*

*Yang, Y., R. Luo, M. Li, M. Zhou, W. Zhang, and J. Wang (2018) Mean field multi-agent reinforcement learning. Proceedings of the 35th International Conference on Machine Learning, PMLR 80:5571-5580.

Mean-Field Approximation

Let $N, M \to \infty$

Uncontrolled vehicles (UV) (modeled by a Markov process)

State: the empirical distribution $\nu_t^N \to \nu_t \in \mathcal{P}(\mathcal{X})$, a mean-field travel choice distribution under the law of large numbers

Transition $\nu_{t+1} \sim q(\cdot \mid \nu_t, \mu_t)$

CAVs (modeled by a Markov decision process)

State $x_t \in \mathcal{X}$ (current travel choice), a realization of a random variable whose distribution matches the MF distribution μ_t

Action $a_t \in \mathcal{X}$ (choosing next travel choice) sample from policy $\pi(\cdot \mid x_t, \mu_t, \nu_t)$

Transition $x_{t+1} \sim p(\cdot \mid x_t, a_t, \mu_t, \nu_t)$

- (μ_t, ν_t) as the state (population state) population cost $c(\mu_t, \nu_t)$ (system travel cost)
- Let $\pi_t(\mu_t, \nu_t) = \pi_t(\cdot \mid \cdot, \mu_t, \nu_t)$, and represent the state-action joint distribution across the CAVs as the population action $h_t = \mu_t \otimes \pi_t(\mu_t, \nu_t)$ (CAV assignments)
- The action is sampled from a new policy $h_t \sim \hat{\pi}(\mu_t, \nu_t)$
- Transition $\mu_{t+1} = T(\nu_t, \mu_t, h_t), \nu_{t+1} \sim q(\nu_t, h_t)$ (population response)

where $T(\nu_t, \mu_t, \mu_t) = \sum_x [\sum_a p(\cdot|x, a, \mu_t, \nu_t)\, \pi_t(a|x, \mu_t, \nu_t) \mu_t(x)]$

Mean-Field Approximation

	Finite Agent Control	Mean-Field Control
State	Current route choice of each CAV Route choice distribution of UVs	Current route choice distributions of two groups (representative agents)
Action	Assigned route choice of each CAV	Joint distribution of current and assigned route choices for all CAVs (population action)
Transition	Each CAV follows assignment UVs follow day-to-day dynamics	CAVs follow a population kernel UVs follow day-to-day dynamics
Cost	Total cost of all vehicles	Total cost of all vehicles

Single-Agent Mean-Field Control

$$\min_{\hat{\pi}} J(\hat{\pi}) = E\left[\sum_{t=0}^{\infty} \gamma^t c(\mu_t, \nu_t)\right]$$

$$s.t.\ \nu_{t+1} = q(\mu_t, \nu_t)$$ Transition kernel for uncontrolled vehicles

$$h_t \sim \hat{\pi}(\cdot|\mu_t, \nu_t)$$ Sample action from the population policy

$$\mu_{t+1} \sim T(\cdot|\mu_t, \nu_t, h_t)$$ Transition kernel for CAVs

Proposition*

If the transition kernels p and q, as well as the system cost c are Lipschitz continuous w.r.t μ and ν, there always exists an optimal policy for MFC

**Cui, K., C. Fabian, Tahir, A., and H. Koeppl (2024) Major-Minor Mean Field Multi-Agent Reinforcement Learning, Proceedings of the 41st International Conference on Machine Learning, PMLR 235:9603-9632.*

RL-Based Distributed Policy *

Input: Initialize population policy $\hat{\pi}^{\omega}$

for iterations $t = 1,2,\ldots$ **do**

Sample population action $h_t \sim \hat{\pi}^{\omega}(\cdot \mid \mu_t, \nu_t)$
Broadcast h_t

for CAV agent $i = 1,2,\ldots,N$ **do**

Retrieve individual policy π_t from h_t
Sample and execute action $a_t^i \sim \pi_t(\cdot \mid x_t^i)$

end for

Observe cost C_t, next population state μ_{t+1}, ν_{t+1}

Update policy $\hat{\pi}^{\omega}$ (with RL algorithms)

end for

*Wu, M., Wang, M., **Yin, Y.** and Lynch, J. (2024) Leveraging connected and automated vehicles for participatory traffic control. Transportation Research Part C, Vol. 166, 104757.*

Distributed departure-time policy with Vickrey's bottleneck setting

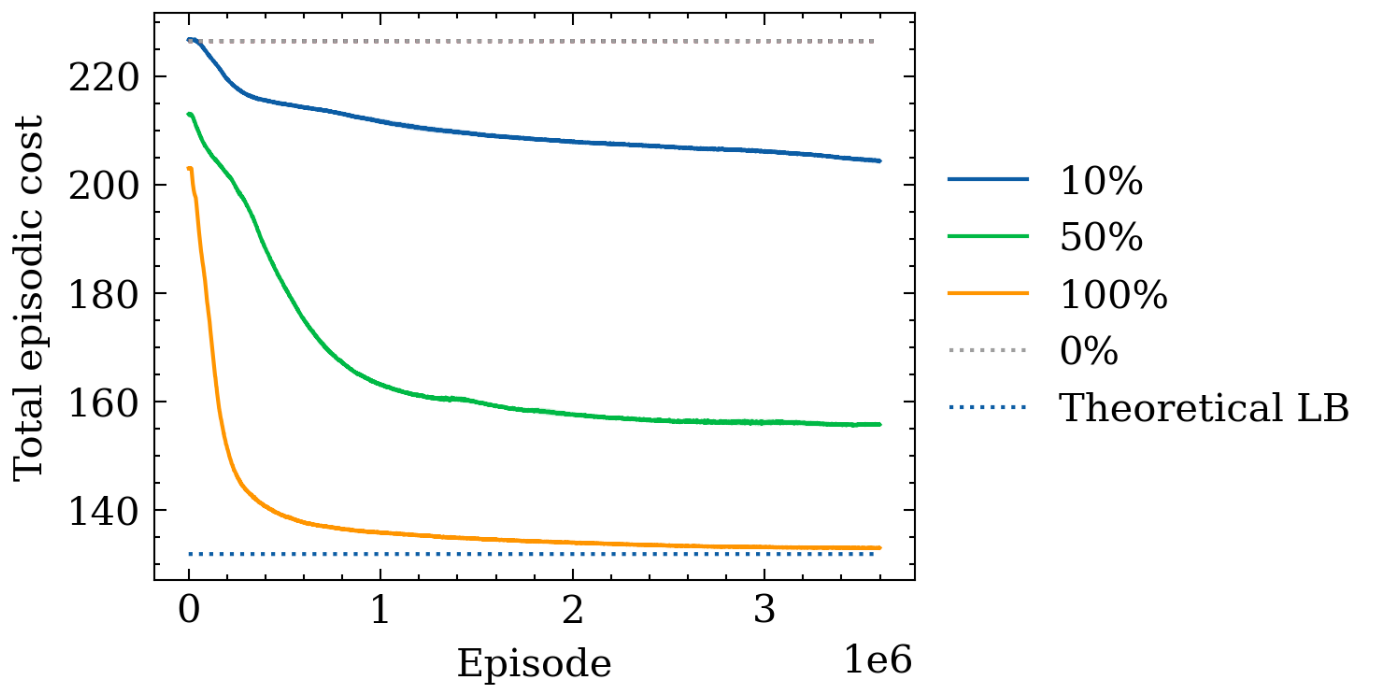

Trained policy performance under 100% penetration.

	Nguyen–Dupuis	Sioux Falls	Bottleneck	Bathtub	Joint choices
System cost	81.18	16.34	133.09	105.58	47.92
LB	81.04	13.98	131.82	103.57	47.76
Relative error(%)	0.17	16.88	0.96	1.94	0.34

Numerical Experiments (Cont'd)

Implementation of the resulting policy: 100% controllable, random initialization, average of three trials

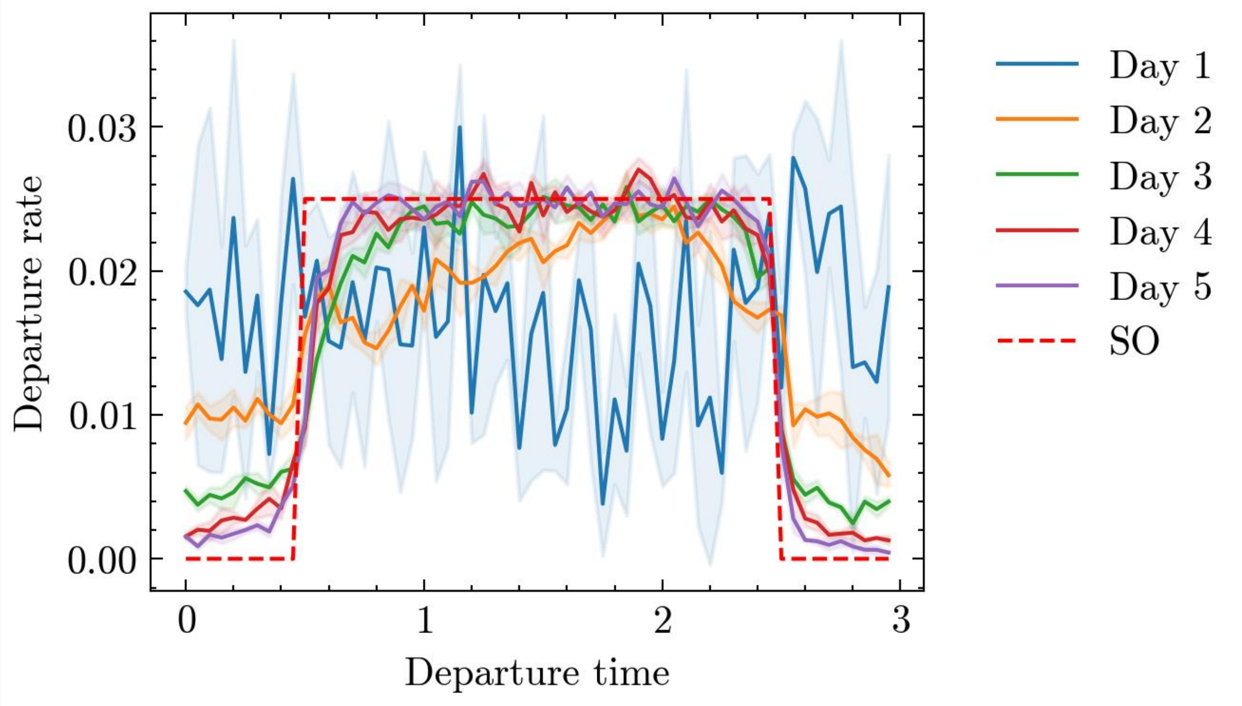

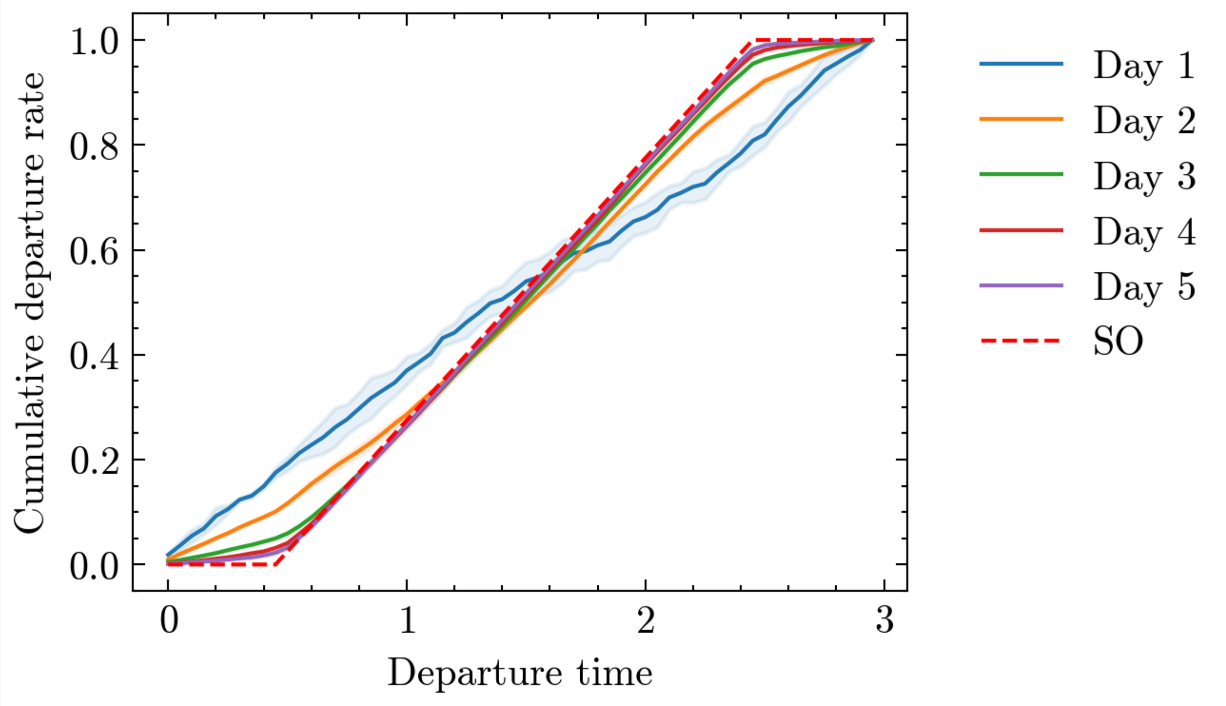

Concluding Remarks

Key Takeaways

- Advancements in connectivity and automation will make travelers more willing to give up personal control over their travel choices.
- CAVs with relinquished agency can serve as “traffic demand distributors,” coordinating routes and departure times to enhance system performance.
- Controlling ~15–20% of traffic can be effective, but required share varies by context.
- Incentives help recruit CAVs by compensating drivers for their reduced control and any added travel costs.
- Implementing such schemes is likely economically sensible for managing rush hour congestion in large cities.
- Mean-field reinforcement learning effectively trains distributed control policies.
- **Participatory traffic control could be a promising addition to our congestion management toolbox.**

Synergistic Integration with Other Instruments*

For example, participatory traffic control, when integrated with path-based pricing, can create synergies that offset each method's weaknesses.

Network	Zones	Nodes	Links	MCR (%)	**Zero – Revenue MCR (%)**
Anaheim	38	416	914	**21**	**0.73**
Barcelona	110	1020	2522	**36**	**0.32**
Berlin-Mitte-Prenzlauer Berg-Friedrichshain-Center	98	975	2184	**14**	**0.06**
Eastern-Massachusetts	74	74	258	**20**	**0.48**
Sioux Falls	24	24	76	**14**	**0.70**
Terrassa	55	1609	3264	**57**	**23**
Winnipeg	147	1052	2836	**42**	**0.6**

Chen, Z., Lin, X, **Yin, Y. and Li, M. (2020) Path controlling of automated vehicles for system optimum on transportation networks with heterogeneous traffic stream. Transportation Research Part C: Emerging Technologies, 110, 312-329.*

Leveraging δ-dissipativity for stabilizing discrete-time mixed-traffic dynamic systems

Dissipation Inequality

$$S(x_{k+1}, p_k) \leq \sigma(x_{k+1}, p_k) + s(\delta x_k, \delta p_k)$$

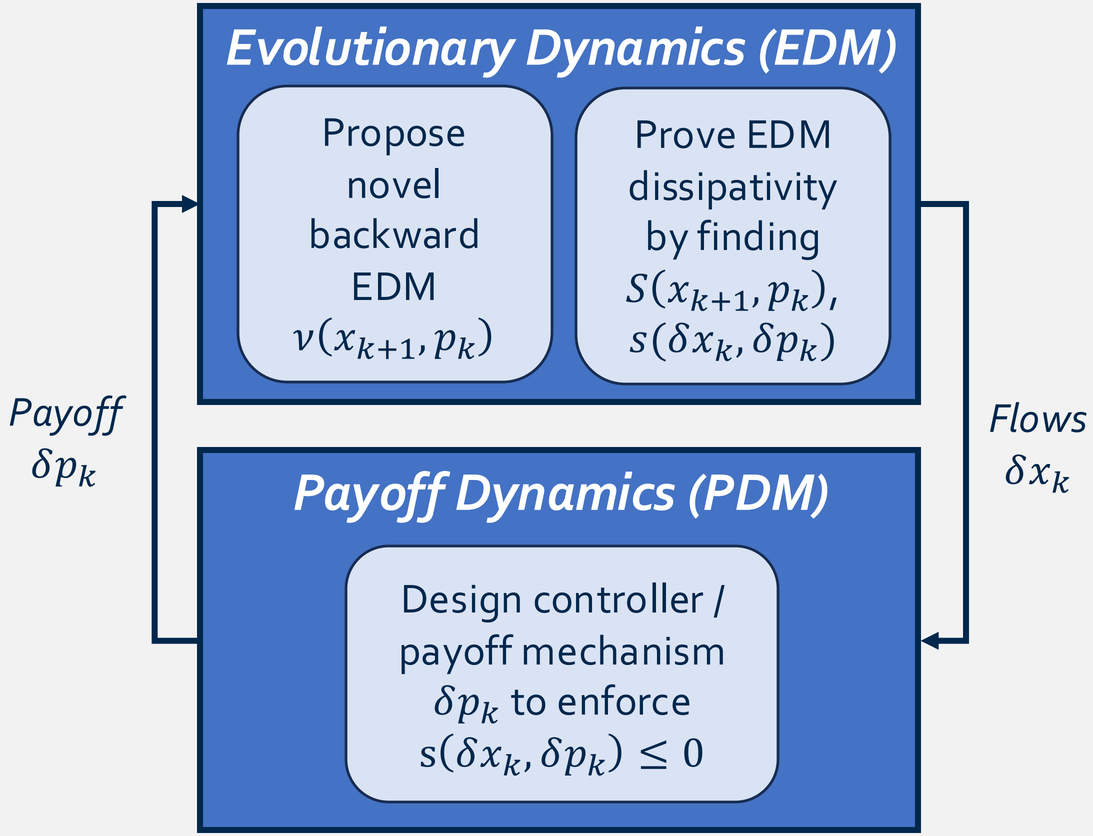

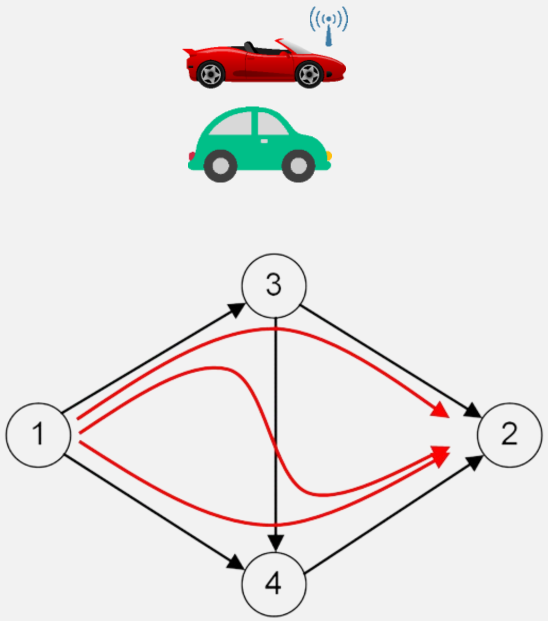

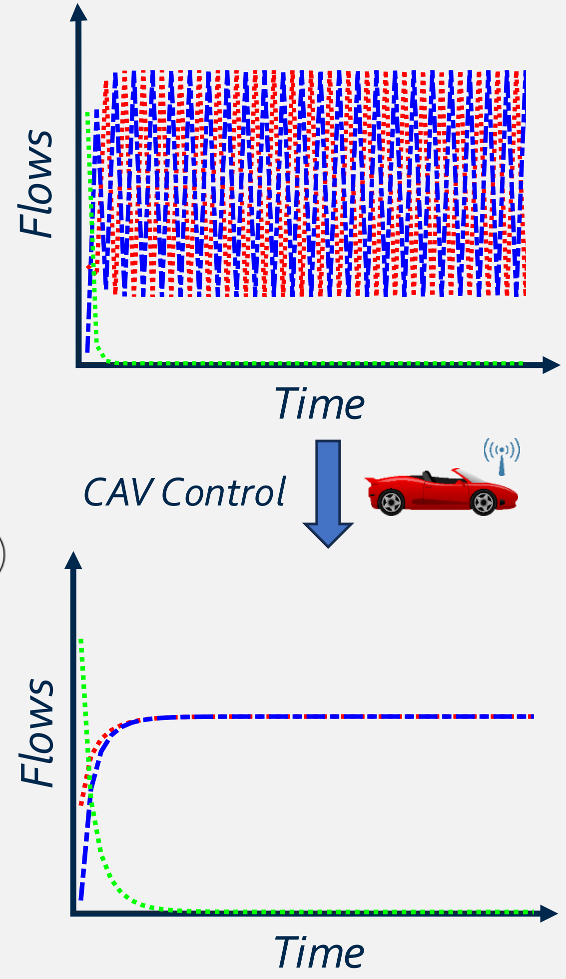

* Lee, R., Scruggs, J., and **Yin, Y.** (2025) Discrete-Time Stabilization of Nash Equilibrium for Mixed Traffic Routing. 2025 American Control Conference, July 8-10, 2025, Denver, CO, USA

Thank You

Email: yafeng@umich.edu

Yafeng Yin, Ph.D.
Donald Cleveland Collegiate Professor of Engineering
Donald Malloure Department Chair of Civil and Environmental Engineering
University of Michigan, Ann Arbor

Numerical Experiments (Cont'd)

Distributed routing policy:

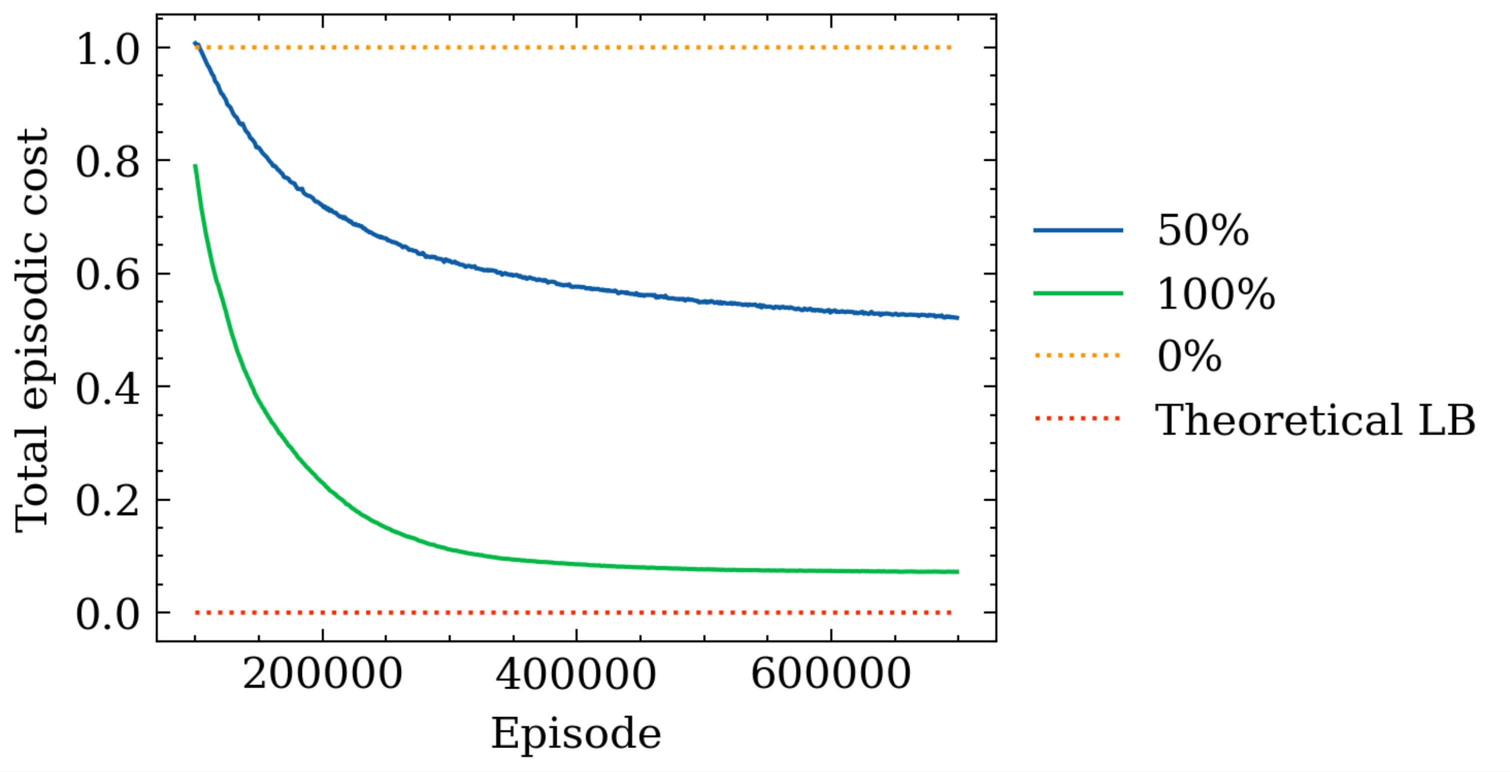

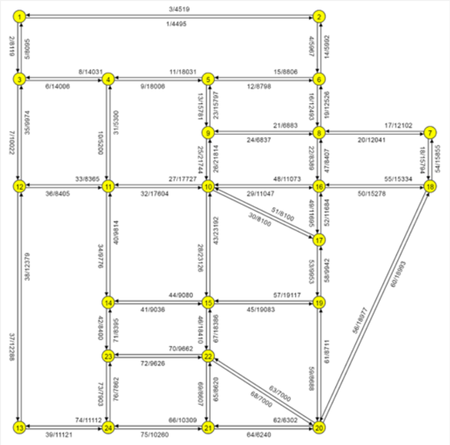

Sioux Falls Network