Urban Congestion Relief Experiments Through Routing-App Interventions

Neha Arora[1] Alexandre Bayen[1,2] Theophile Cabannes[1] Kevin Chen[1]
Abdul Rahman Kreidieh[1] Yechen Li[1] Marc Nunkesser[1] Prem Ramaswami[1]
Andrew Tomkins[1] Eray Turkel[1] Shoshana Vasserman[1,3] Haizheng Zhang[1]

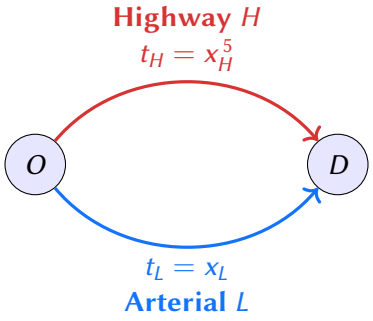
[1]Google Research [2]UC Berkeley [3]Stanford GSB

Nature Cities (2026)

Motivation

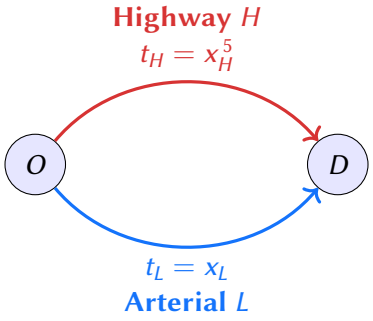
- Most routing apps (e.g., Google/Apple Maps) minimize travel time for the user only
- Increasingly, these apps consider other criteria too: tolls, ease of driving, etc.
- What about congestion externalities?

Motivation: individualistic routing vs the social optimum



- $x_H + x_L = 1$
- Highway: fast when empty, jams quickly
- Arterial: slower, degrades gradually

Motivation: individualistic routing vs the social optimum



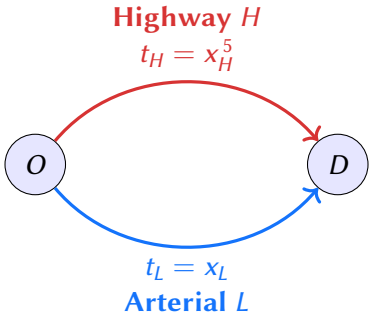
- $x_H + x_L = 1$
- Highway: fast when empty, jams quickly
- Arterial: slower, degrades gradually

Individual routing (most routing apps)

- Each driver takes the faster road
- NE: travel times equalize: $t_H = t_L$

$\rightsquigarrow \approx 3/4$ of drivers on H, $1/4$ on L

Motivation: individualistic routing vs the social optimum



- $x_H + x_L = 1$
- Highway: fast when empty, jams quickly
- Arterial: slower, degrades gradually

Individual routing (most routing apps)

- Each driver takes the faster road
- NE: travel times equalize: $t_H = t_L$

$\rightsquigarrow$ $\approx 3/4$ of drivers on H, $1/4$ on L

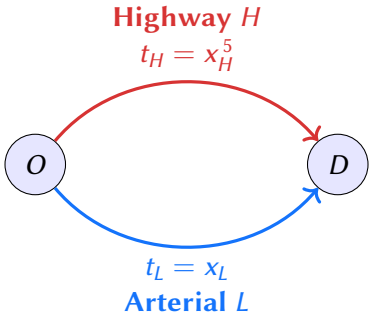
Optimal routing (social planner)

- Minimize average travel time:

$$\min_{x_H} \; x_H \, t_H(x_H) + (1 - x_H) \, t_L(1 - x_H)$$

$\rightsquigarrow$ Fewer drivers on H ($x_H \approx 0.65$), who each get there much faster

Motivation: individualistic routing vs the social optimum



- $x_H + x_L = 1$
- Highway: fast when empty, jams quickly
- Arterial: slower, degrades gradually

Individual routing (most routing apps)

- Each driver takes the faster road
- NE: travel times equalize: $t_H = t_L$

$\rightsquigarrow \approx 3/4$ of drivers on H, $1/4$ on L

Optimal routing (social planner)

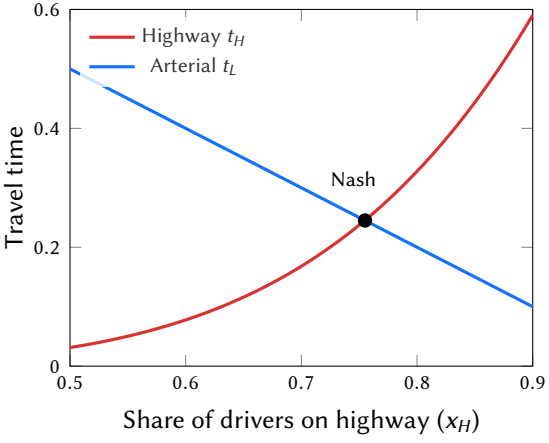
- Minimize average travel time:

$$\min_{x_H} \; x_H\, t_H(x_H) + (1 - x_H)\, t_L(1 - x_H)$$

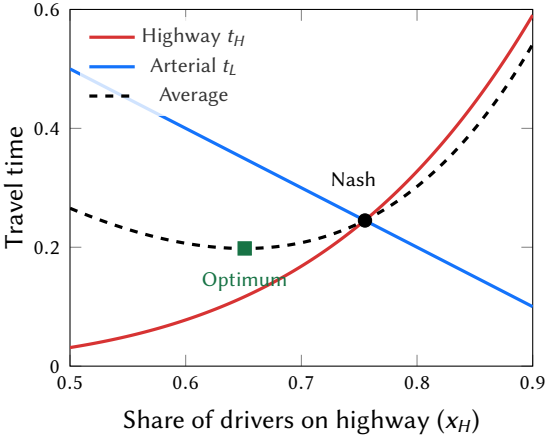
$\rightsquigarrow$ Fewer drivers on H ($x_H \approx 0.65$), who each get there much faster

Research Q: Can we get closer to optimal by nudging a few of the drivers on H to take L?

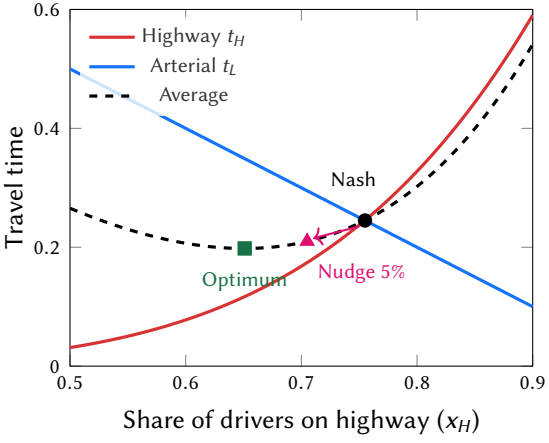
In this example, yes: a small nudge goes a long way



In this example, yes: a small nudge goes a long way



In this example, yes: a small nudge goes a long way

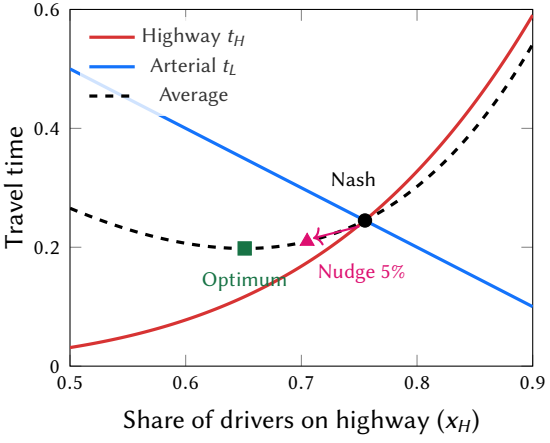


Move just 5% of drivers from H to L:

	Nash	Nudge	Optimum
x_H	0.755	0.705	0.651
t_H	0.245	0.174	0.116
t_L	0.245	0.295	0.349
Average	0.245	**0.210**	0.198

- In the example: 5% of drivers moved
  $\Rightarrow \approx$75% of gain from optimal routing

In this example, yes: a small nudge goes a long way



Move just 5% of drivers from H to L:

	Nash	Nudge	Optimum
x_H	0.755	0.705	0.651
t_H	0.245	0.174	0.116
t_L	0.245	0.295	0.349
Average	0.245	**0.210**	0.198

- In the example: 5% of drivers moved $\Rightarrow \approx$75% of gain from optimal routing
- Depends on the convexity of H and L
- Depends on reach, compliance, induced demand, etc.

This Project

▸ Related literature

Research questions

1. Can nudging a *small* share of trips off congested roads improve speeds in real life?
2. Or does it just move congestion somewhere else?
3. Why does it work?

This Project

Related literature

Research questions

1. Can nudging a *small* share of trips off congested roads improve speeds in real life?
2. Or does it just move congestion somewhere else?
3. Why does it work?

Approach

- Experiment on Google Maps: $\sim$1000 congested roads in 10 U.S. cities over 6 months
- Add small "cost" to routes through targeted roads rerouting $< 2\%$ of trips
- Switchback design: whole city alternates between treatment and control day to day
- Treatment effects on road speeds, trip durations, CO_2 emissions, and compliance

This Project

▸ Related literature

Research questions

1. Can nudging a *small* share of trips off congested roads improve speeds in real life?
   $\rightsquigarrow$ Yes. Speeds on targeted segments $\uparrow \approx 2\%$ (up to 4.6% in LA)
2. Or does it just move congestion somewhere else?
3. Why does it work?

Approach

- Experiment on Google Maps: $\sim$1000 congested roads in 10 U.S. cities over 6 months
- Add small "cost" to routes through targeted roads rerouting $< 2\%$ of trips
- Switchback design: whole city alternates between treatment and control day to day
- Treatment effects on road speeds, trip durations, CO_2 emissions, and compliance

This Project

▸ Related literature

Research questions

1. Can nudging a *small* share of trips off congested roads improve speeds in real life?
   ⇝ Yes. Speeds on targeted segments ↑≈ 2% (up to 4.6% in LA)
2. Or does it just move congestion somewhere else?
   ⇝ Net gains network-wide: speeds on all affected roads ↑ 0.35%; trip travel times ↓ 0.7%
3. Why does it work?

Approach

- Experiment on Google Maps: ∼1000 congested roads in 10 U.S. cities over 6 months
- Add small "cost" to routes through targeted roads rerouting $< 2\%$ of trips
- Switchback design: whole city alternates between treatment and control day to day
- Treatment effects on road speeds, trip durations, CO_2 emissions, and compliance

This Project

▸ Related literature

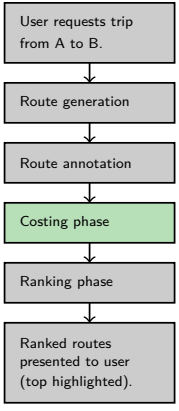
Research questions

1. Can nudging a *small* share of trips off congested roads improve speeds in real life?
   ⇝ Yes. Speeds on targeted segments ↑≈ 2% (up to 4.6% in LA)
2. Or does it just move congestion somewhere else?
   ⇝ Net gains network-wide: speeds on all affected roads ↑ 0.35%; trip travel times ↓ 0.7%
3. Why does it work?
   ⇝ Traffic pulled from central segments is spread thinly across many peripheral ones

Approach

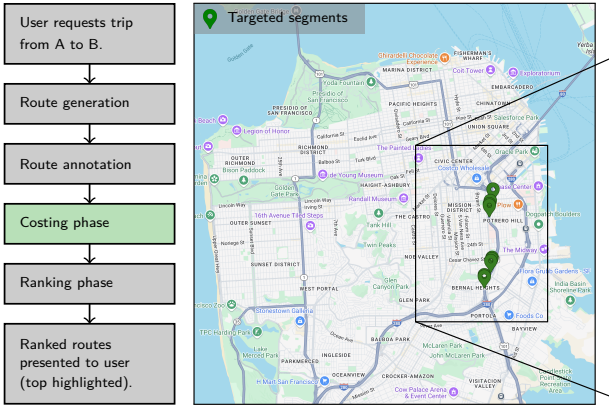
- Experiment on Google Maps: ∼1000 congested roads in 10 U.S. cities over 6 months
- Add small "cost" to routes through targeted roads rerouting $< 2\%$ of trips
- Switchback design: whole city alternates between treatment and control day to day
- Treatment effects on road speeds, trip durations, CO_2 emissions, and compliance

Intervention: a small penalty for traversing congested segments



- Routes through pre-selected **targeted segments** (green pins) get a small added cost

Intervention: a small penalty for traversing congested segments



- Routes through pre-selected **targeted segments** (green pins) get a small added cost

⇝ Trips with a similarly-costed alternative are steered away from those segments

Heuristic externality cost on targeted segments

- App ranks feasible routes $\mathcal{P}_i \in \mathcal{P}$ by cost:

$$C_i = \underbrace{f(\mathcal{P}_i)}_{\substack{\text{Predicted} \\ \text{travel time}}} + \underbrace{c(\mathcal{P}_i)}_{\substack{\text{Other costs} \\ \text{(tolls, etc.)}}}$$

Heuristic externality cost on targeted segments

- App ranks feasible routes $\mathcal{P}_i \in \mathcal{P}$ by cost:

$$C_i = \underbrace{f(\mathcal{P}_i)}_{\substack{\text{Predicted} \\ \text{travel time}}} + \underbrace{c(\mathcal{P}_i)}_{\substack{\text{Other costs} \\ \text{(tolls, etc.)}}}$$

- Add a fixed penalty ΔT for each targeted segment that is currently congested:

$$C_i' = f(\mathcal{P}_i) + c(\mathcal{P}_i) + \min\left(\sum_{x \in X_{\text{targeted}} \cap \mathcal{P}_i} \underbrace{\Delta T \cdot \mathbb{1}\{v(x,t) < v_{\text{congested}}(x)\}}_{\substack{\text{Only active when} \\ \text{segment is congested}}}, \underbrace{c_{\max}(\mathcal{P})}_{\text{Cap}} \right)$$

Heuristic externality cost on targeted segments

- App ranks feasible routes $\mathcal{P}_i \in \mathcal{P}$ by cost:

$$C_i = \underbrace{f(\mathcal{P}_i)}_{\substack{\text{Predicted} \\ \text{travel time}}} + \underbrace{c(\mathcal{P}_i)}_{\substack{\text{Other costs} \\ \text{(tolls, etc.)}}}$$

- Add a fixed penalty ΔT for each targeted segment that is currently congested:

$$C_i' = f(\mathcal{P}_i) + c(\mathcal{P}_i) + \min\left(\sum_{x \in X_{\text{targeted}} \cap \mathcal{P}_i} \underbrace{\Delta T \cdot \mathbb{1}\{v(x,t) < v_{\text{congested}}(x)\}}_{\substack{\text{Only active when} \\ \text{segment is congested}}}, \underbrace{c_{\max}(\mathcal{P})}_{\text{Cap}} \right)$$

- The cap keeps rerouted trips within a few percent of the fastest route
  - ⇝ On average, rerouted trips were assigned a route ≈ 30 seconds slower
  - ⇝ Under 2% of all trips received a different recommendation

Segment selection: ∼100 congested segments per city

- **Density at peak demand:** segments with highest 75th-percentile density
- **Bottlenecks:** speed drops relative to downstream segments (Chen et al., 2004)
- Road class varied across cities to compare freeways vs. arterials

City	Segs	Criterion	Road class
Atlanta	106	Peak density	Freeway
Boston	108	Peak density	Freeway
Chicago	81	Bottlenecks	Freeway + arterial
Los Angeles	120	Peak density	Freeway
Miami	100	Peak density	Arterial
New York	103	Peak density	Freeway
Philadelphia	101	Peak density	Freeway
Salt Lake City	94	Peak density	Arterial
San Francisco	102	Bottlenecks	US-101
Seattle	97	Peak density	I-5

Experimental design & data

- **Switchback design:** city-wide treatment vs. control on alternating days
  - Adjacent days give directly comparable hourly outcomes
  - Treatment def allows for spillovers on trips throughout the city

Experimental design & data

- **Switchback design:** city-wide treatment vs. control on alternating days
  - Adjacent days give directly comparable hourly outcomes
  - Treatment def allows for spillovers on trips throughout the city

- **Data:** aggregated + anonymized hourly stats from Google Maps
  - **Segment level** (6 months): distance-weighted road speeds + fuel consumption rates
    - X_{targeted}: the penalized segments themselves
    - X_{affected}: every segment on an original *or* rerouted path ($\approx 80\%$ of city traffic)

Experimental design & data

- **Switchback design:** city-wide treatment vs. control on alternating days
  - Adjacent days give directly comparable hourly outcomes
  - Treatment def allows for spillovers on trips throughout the city

- **Data:** aggregated + anonymized hourly stats from Google Maps
  - **Segment level** (6 months): distance-weighted road speeds + fuel consumption rates
    - X_{targeted}: the penalized segments themselves
    - X_{affected}: every segment on an original *or* rerouted path ($\approx 80\%$ of city traffic)
  - **Trip level** (2 months): average travel times of directly/indirectly affected trips

Experimental design & data

- **Switchback design:** city-wide treatment vs. control on alternating days
  - Adjacent days give directly comparable hourly outcomes
  - Treatment def allows for spillovers on trips throughout the city

- **Data:** aggregated + anonymized hourly stats from Google Maps
  - **Segment level** (6 months): distance-weighted road speeds + fuel consumption rates
    - X_{targeted}: the penalized segments themselves
    - X_{affected}: every segment on an original *or* rerouted path ($\approx 80\%$ of city traffic)
  - **Trip level** (2 months): average travel times of directly/indirectly affected trips

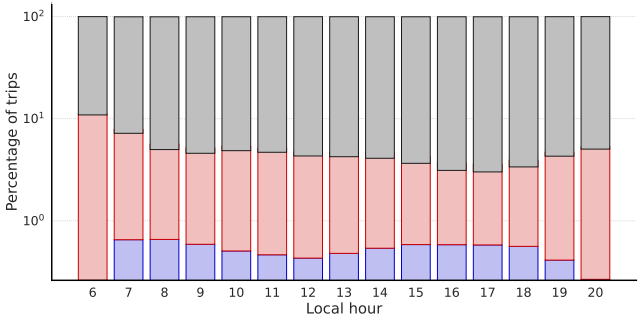
- **Estimator:** multi-level random effects model w/ city $\times$ hour effects (Gelman, 2006)
  - Main estimates with hierarchical Bayesian estimation + posterior inference
  - Robustness checks with additive mixed effects, MLE + placebo tests

▸ Estimation details

Who got rerouted? A small slice of trips

- Candidate trips had to traverse targeted segments
- Rerouted trips: 5–15% of those, $< 2\%$ of all trips
- Rerouted trips tend to be the *shorter* trips among those using targeted segments. These have close substitutes

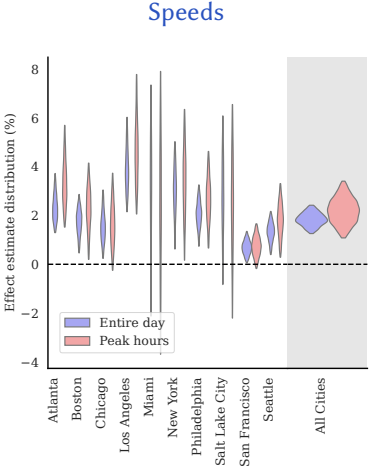
Atlanta: share of trips by hour



Note: Log scale. Blue: rerouted trips. Red: trips through targeted segments. Grey: other trips. Groups are not mutually exclusive.

▸ Trip lengths by group

Targeted road segments: speeds ↑ ≈ 2%; fuel use rate ↓ 0.5–1%

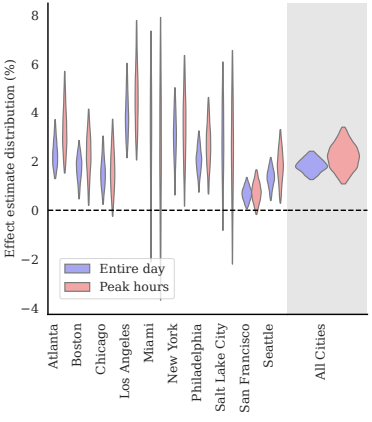


Note: Posterior distributions of percentage effects. All cities: +1.8% (entire day) and +2.1% (peak) on speeds, $P[\text{pos}] = 100\%$. Largest where freeways were targeted: Atlanta +3.3%, Los Angeles +4.6% (peak).

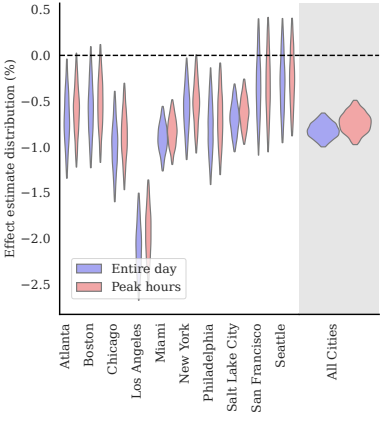
▸ Table

Targeted road segments: speeds ↑ ≈ 2%; fuel use rate ↓ 0.5–1%

Speeds



Fuel consumption rates

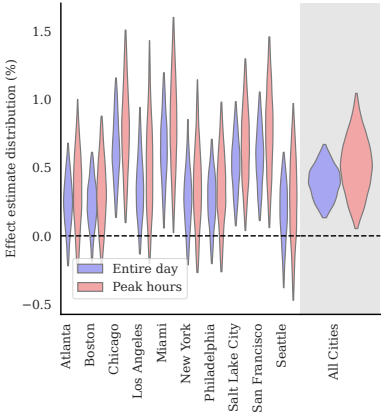


Note: Posterior distributions of percentage effects. All cities: +1.8% (entire day) and +2.1% (peak) on speeds, $P[\text{pos}] = 100\%$. Largest where freeways were targeted: Atlanta +3.3%, Los Angeles +4.6% (peak).

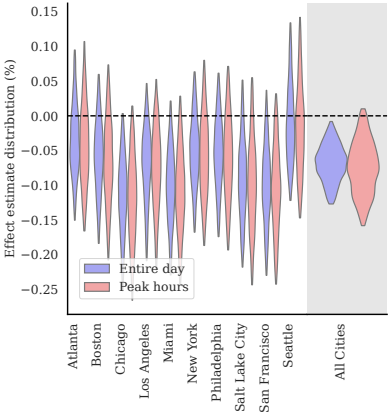
▸ Table

All affected road segments: speeds $\uparrow \approx 0.35 - 0.5\%$

Speeds



Fuel consumption rates



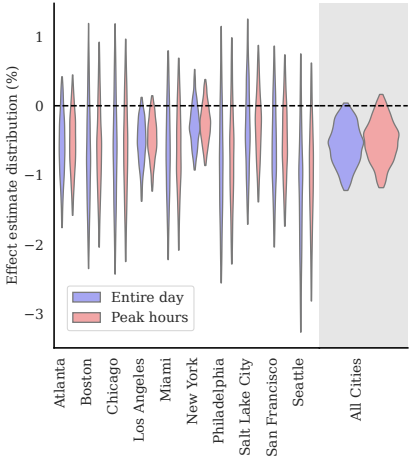
Note: All cities: speeds $+0.37\%$ (entire day) and $+0.49\%$ (peak). $P[\text{speeds} > 0] = 99.8\%$ and $P[\text{fuel} < 0] = 98.6\%$. The affected set covers $\approx 80\%$ of city traffic.

▸ Table

Trips through affected roads: travel times ↓ 0.7%

- Travel times for trips passing through any affected segment
- Median drops ≈ 0.5–1% across cities

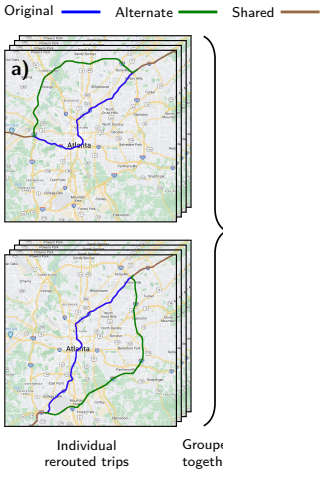
⇝ Suggests net network improvement



Note: Posterior distributions from the daily hierarchical model

▸ Environmental impact ▸ Robustness checks

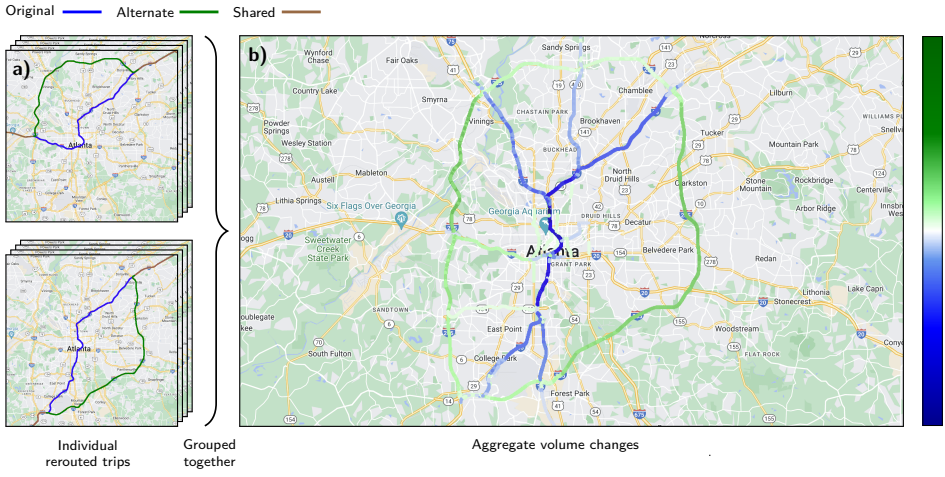
Mechanism: concentrated traffic is dispersed thinly



Note: Atlanta. Each rerouted trip: original path through central I-85 vs. alternate path. Overlap in brown.

▸ Other cities

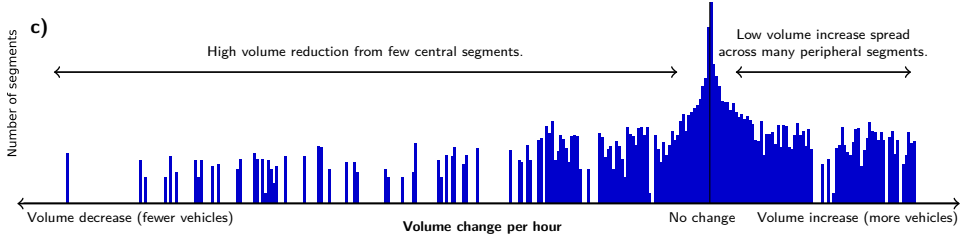
Mechanism: concentrated traffic is dispersed thinly



Note: Net change in hourly volume if all rerouted trips took the alternate vs. original path. Blue: fewer vehicles. Green: more vehicles.

▸ Other cities

Large cuts on a few segments, small additions on many



Note: Atlanta. Histogram of net hourly volume changes by segment if all rerouted trips took the alternate vs. original path. Similar patterns in all 10 cities (SI Fig. 7).

Why dispersion can help

- Penalties sit on a few fixed, high-demand segments...
- ...but trips through them have many different origins and destinations
  - Each trip's best alternative uses different roads

  ⇝ Diverted traffic is spread over many segments

Why dispersion can help

- Penalties sit on a few fixed, high-demand segments...
- ...but trips through them have many different origins and destinations
  - Each trip's best alternative uses different roads

  ⇝ Diverted traffic is spread over many segments
- Congestion is convex
  - Penalty only switches on when a segment is very congested
  - Alternatives are close substitutes of similar road class with spare capacity

Why dispersion can help

- Penalties sit on a few fixed, high-demand segments...
- ...but trips through them have many different origins and destinations
  - Each trip's best alternative uses different roads

  ⇝ Diverted traffic is spread over many segments
- Congestion is convex
  - Penalty only switches on when a segment is very congested
  - Alternatives are close substitutes of similar road class with spare capacity

⇝ Low-cost penalties can even out effective traffic across the network

Open questions

- **How far from optimal?**
  - We used fixed penalties on historically congested segments
  - Dynamic, optimized penalties that account for feedback loops could do better

Open questions

- **How far from optimal?**
  - We used fixed penalties on historically congested segments
  - Dynamic, optimized penalties that account for feedback loops could do better
- **Induced demand?**
  - Faster roads might draw in drivers who used to avoid them (Volker and Handy, 2021)
  - In our setting, the gains are nearly invisible to any single driver:
    - $\approx 0.25\%$ of avg trip time; $\approx 1/40$ of day-to-day SD
  - A bigger, more visible intervention might induce an equilibrium response

Open questions

- **How far from optimal?**
  - We used fixed penalties on historically congested segments
  - Dynamic, optimized penalties that account for feedback loops could do better
- **Induced demand?**
  - Faster roads might draw in drivers who used to avoid them (Volker and Handy, 2021)
  - In our setting, the gains are nearly invisible to any single driver:
    - $\approx 0.25\%$ of avg trip time; $\approx 1/40$ of day-to-day SD
  - A bigger, more visible intervention might induce an equilibrium response
- **Beyond Maps**
  - Our experiment demonstrates the inefficiency of baseline traffic conditions
  - An alternative to nudging drivers is to coordinate fleets

  $\rightsquigarrow$ Suggests (potential) road network improvement as AV adoption expands

Conclusion

- Rerouted $< 2\%$ of Google Maps trips from $\sim$1000 congested segments in 10 cities
  - Rerouted trips were only $\approx$ 30 seconds slower on average
- Targeted segments: speeds $\uparrow \approx 2\%$, fuel use $\downarrow$ 0.5–1%
- Modest network-wide gains
  - Affected segments: speeds $\uparrow$ 0.35% (0.5% at peak). Trip travel times $\downarrow$ 0.7%
  - Traffic pulled off a few central roads spreads thinly across many peripheral roads

$\rightsquigarrow$ Small, network-aware nudges from routing apps can measurably reduce congestion

Thank you!

svass@stanford.edu

References

Tamer Başar and Geert Jan Olsder. *Dynamic Noncooperative Game Theory, 2nd Edition*. Society for Industrial and Applied Mathematics, 1998. doi: 10.1137/1.9781611971132.

Iavor Bojinov, David Simchi-Levi, and Jinglong Zhao. Design and analysis of switchback experiments. *Management Science*, 2022.

Audrey Boruvka, Daniel Almirall, Katie Witkiewitz, and Susan A Murphy. Assessing time-varying causal effect moderation in mobile health. *Journal of the American Statistical Association*, 113(523):1112–1121, 2018.

Aaron Brooker, Jeffrey Gonder, Lijuan Wang, Eric Wood, Sean Lopp, and Laurie Ramroth. Fastsim: A model to estimate vehicle efficiency, cost and performance. Technical report, SAE Technical Paper, 2015.

Chao Chen, Alexander Skabardonis, and Pravin Varaiya. Systematic identification of freeway bottlenecks. *Transportation Research Record*, 1867(1):46–52, 2004.

Yi-Chang Chiu, Jon Bottom, Michael Mahut, Alex Paz, Ramachandran Balakrishna, Travis Waller, and Jim Hicks. *Dynamic Traffic Assignment: A Primer*. Number June. Transportation Research Board, Washington, D.C., 7 2011. ISBN 978-0-309-43529-1. doi: 10.17226/22872. URL https://www.nap.edu/catalog/22872.

Andrew Gelman. Multilevel (hierarchical) modeling: what it can and cannot do. *Technometrics*, 48(3): 432–435, 2006.

References

Andrew Gelman and Jennifer Hill. *Data analysis using regression and multilevel/hierarchical models*. Cambridge university press, 2006.

Jacob Holden, Nicholas Reinicke, and Jeffrey Cappellucci. Routee: A vehicle energy consumption prediction engine. *Society of Automotive Engineers Technical Paper Series*, 2(NREL/JA-5400-78089), 2020.

Alexander Keimer and Alexandre Bayen. Routing on traffic networks incorporating past memory up to real-time information on the network state. *Annual Review of Control, Robotics, and Autonomous Systems*, 3(1):151–172, 2020.

W. Krichene, J. D. Reilly, S. Amin, and A. M. Bayen. Stackelberg routing on parallel networks with horizontal queues. *IEEE Transactions on Automatic Control*, 59(3):714–727, March 2014. ISSN 0018-9286. doi: 10.1109/TAC.2013.2289709.

Elizabeth O Lillie, Bradley Patay, Joel Diamant, Brian Issell, Eric J Topol, and Nicholas J Schork. The n-of-1 clinical trial: the ultimate strategy for individualizing medicine? *Personalized medicine*, 8(2):161–173, 2011.

M. Patriksson. *The traffic assignment problem: models and methods*. Courier Dover Publications, 2015.

Bin Ran and David Boyce. *Modeling Dynamic Transportation Networks*. Springer Berlin, Heidelberg, 01 1996. ISBN 978-3-642-80232-4. doi: 10.1007/978-3-642-80230-0.

David Ratkowsky, Richard Alldredge, and Marc A Evans. *Cross-over experiments: design, analysis and application*, volume 135. CRC Press, 1992.

References

Tim Roughgarden. *Selfish routing and the price of anarchy*. MIT Press, 2005.

Jamey M.B. Volker and Susan L Handy. The induced travel calculator and its applications, 2021.

Related literature

- **Selfish vs. system-optimal routing, in theory**

  Başar and Olsder (1998); Roughgarden (2005); Patriksson (2015); Krichene et al. (2014); Keimer and Bayen (2020)

  ⇝ Field evidence that marginal, network-aware routing improves a real network

◂ Back

Related literature

- **Selfish vs. system-optimal routing, in theory**

  Başar and Olsder (1998); Roughgarden (2005); Patriksson (2015); Krichene et al. (2014); Keimer and Bayen (2020)

  ⇝ Field evidence that marginal, network-aware routing improves a real network

- **Dynamic traffic assignment**

  Ran and Boyce (1996); Chiu et al. (2011)

  ⇝ Penalties based on *historical* congestion: far simpler than DTA, still effective

◂ Back

Related literature

- **Selfish vs. system-optimal routing, in theory**

  Başar and Olsder (1998); Roughgarden (2005); Patriksson (2015); Krichene et al. (2014); Keimer and Bayen (2020)

  ⇝ Field evidence that marginal, network-aware routing improves a real network

- **Dynamic traffic assignment**

  Ran and Boyce (1996); Chiu et al. (2011)

  ⇝ Penalties based on *historical* congestion: far simpler than DTA, still effective

- **Switchback / crossover experiments**

  Ratkowsky et al. (1992); Lillie et al. (2011); Boruvka et al. (2018); Bojinov et al. (2022)

  ⇝ Switchback design for a city-wide traffic intervention

◂ Back

Estimation: hierarchical Bayesian model

- **Challenge:** only 10 cities, noisy day-to-day traffic, many city × hour effects of interest
- **Approach:** hierarchical (partial pooling) model (Gelman, 2006; Gelman and Hill, 2006)

$$y_{it} \sim N(\mu_{it} + \tau_{it}, \sigma_y), \qquad \tau_{it} \sim N(\tau_i + \tau_t, \sigma_{it}), \qquad \tau_i,\ \tau_t \sim N(\tau, \cdot)$$

  - City × hour baselines μ_{it}. Effects pooled across cities and (correlated) hours
  - Separate global effect τ for freeway-targeted vs. other cities
  - Posterior probabilities are easy to read and reduce multiple-testing concerns
- **Robustness:**
  - Mixed-effects linear model with randomization-based p-values (1,000 placebo assignments)
  - Drop outlier days (holidays, severe weather). Results with these days kept are similar

◂ Back ▸ Model hierarchy ▸ Mixed model & placebo ▸ With outliers

Environmental impact: ∼1,000+ tons CO_2e per city per year

- Fuel use from NREL's FASTSim-based RouteE model (Brooker et al., 2015; Holden et al., 2020)
- Scale up from affected segments (weekdays, 7am–8pm) to a year, assuming a fixed penetration rate
- Median savings exceed 1,000 tCO_2e/yr in 8 of 10 cities
- Caveat: city-level 90% intervals include zero

	Savings (tCO_2e/yr)	
City	Median	90% interval
Atlanta	1,991	[−4,716, 7,537]
Boston	1,499	[−1,411, 4,980]
Chicago	4,297	[−656, 9,548]
Los Angeles	4,511	[−2,369, 13,956]
Miami	1,817	[−485, 4,124]
New York	2,514	[−3,820, 8,587]
Philadelphia	1,369	[−1,522, 4,453]
Salt Lake City	658	[−415, 1,577]
San Francisco	2,285	[−992, 5,248]
Seattle	504	[−3,921, 3,458]

Note: Positive numbers are savings. Intervals are the 5th–95th posterior percentiles.

Back

Environmental impact: ∼1,000+ tons CO_2e per city per year

- Fuel use from NREL's FASTSim-based RouteE model (Brooker et al., 2015; Holden et al., 2020)
- Scale up from affected segments (weekdays, 7am–8pm) to a year, assuming a fixed penetration rate
- Median savings exceed 1,000 tCO_2e/yr in 8 of 10 cities
- Caveat: city-level 90% intervals include zero

	Savings (tCO_2e/yr)	
City	Median	90% interval
Atlanta	1,991	[−4,716, 7,537]
Boston	1,499	[−1,411, 4,980]
Chicago	4,297	[−656, 9,548]
Los Angeles	4,511	[−2,369, 13,956]
Miami	1,817	[−485, 4,124]
New York	2,514	[−3,820, 8,587]
Philadelphia	1,369	[−1,522, 4,453]
Salt Lake City	658	[−415, 1,577]
San Francisco	2,285	[−992, 5,248]
Seattle	504	[−3,921, 3,458]

Note: Positive numbers are savings. Intervals are the 5th–95th posterior percentiles.

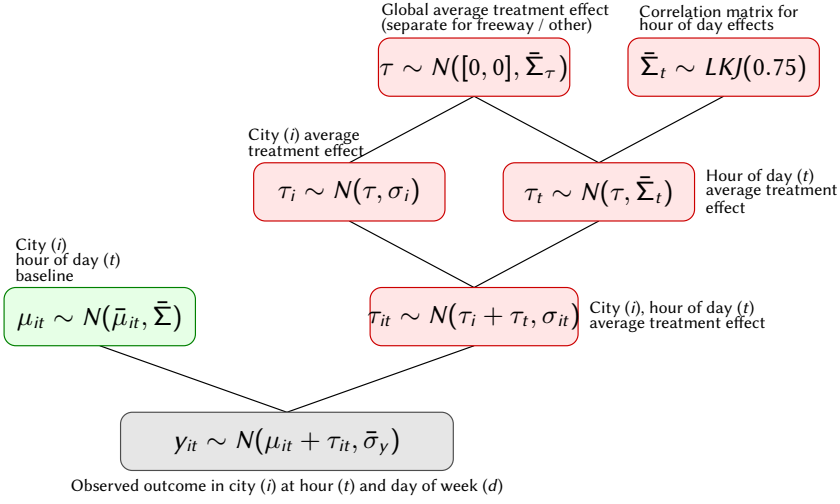
Back

Robustness

- **Frequentist mixed-effects model + placebo tests** ▸ Details
  - No partial pooling, so less precise. Same qualitative results
- **Keeping outlier days** ▸ Results
  - Effects point the same way and are stronger in some cases
- **Did rerouted drivers follow their new routes?** ▸ Compliance
  - No significant change in trip completion or route deviation rates between treatment and control
- **Is it just differences in demand across days?** ▸ Speeds vs. flows
  - Fitted speed–volume line for treatment days sits above control days across trip volumes

◂ Back

Hourly Bayesian model hierarchy



Note: Green: control-period baseline hierarchy. Red: treatment effect hierarchy. Priors on some variance components omitted. The trip travel time analysis uses an analogous daily (day-of-week) hierarchy.

Back

Mixed-effects linear model & placebo tests

- For each outcome y in city c, hour h, date t:

$$y_{c,h,t} = \beta_{\text{day}(t)} + \beta_{\text{pk}(h)} + \beta_{\text{day}(t),\text{pk}(h)} + \beta_{\text{sb}(h,t)} + \beta_{\text{sb}(h,t),\text{pk}(h)} + \eta_c + \sum_{k=1}^{3} \eta_{h^k} h^k + \epsilon_{c,h,t}$$

  - sb: switchback status (treatment/control). pk: peak hours. η_c: city random effect

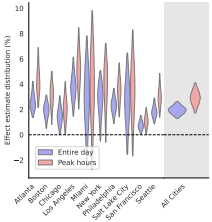
- **Inference by randomization:**
  - Randomly relabel each switchback interval as treatment or control
  - Re-estimate 1,000 times to get the null distribution of effects
  - p-value = share of placebo runs with a larger effect than the true assignment

- No partial pooling, so less precise, but no concern that pooling "borrows" effects across cities

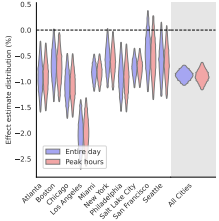
Back Back

Results including outlier days

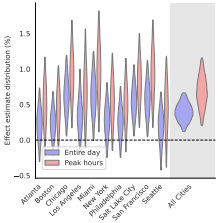
(a) Targeted speeds



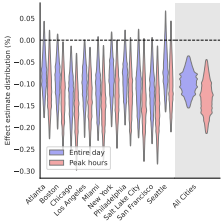
(b) Targeted fuel



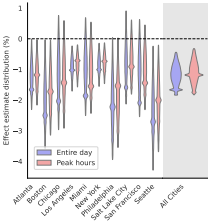
(c) Affected speeds



(d) Affected fuel

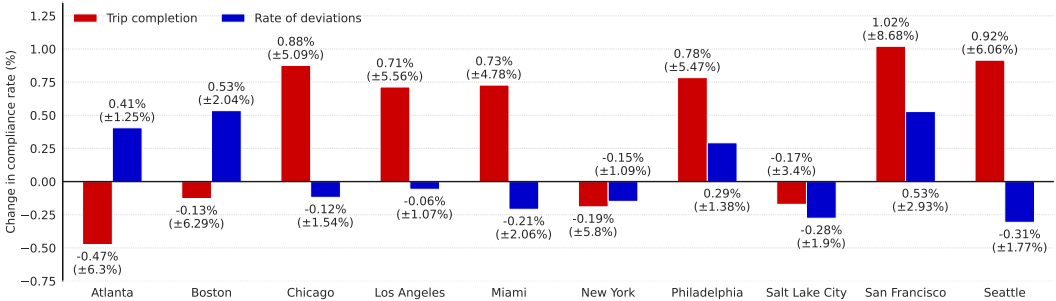


(e) Trip travel times



Note: Same analysis as in the main results, keeping days that the outlier filter would drop (mostly U.S. holidays and severe weather). Effects point the same

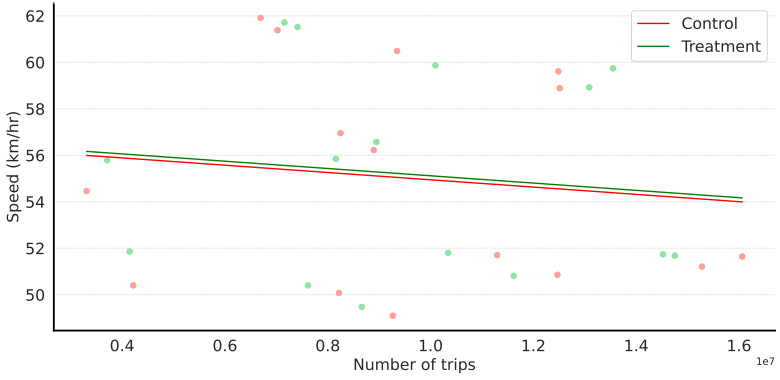
Trip guidance compliance



Note: Difference in trip completion and route deviation rates between treatment and control periods. Values in brackets are standard deviations across treatment and control intervals. None are significant.

Back

Speeds vs. trip flows



Note: Measured impact of the intervention as a function of observed trip flows. Each point is a day. Lines are linear fits by treatment status.

Back

Targeted segments: posterior summaries

		Speed (%)				Fuel rate (%)	
City	Period	5th	Median	95th	P[pos]	Median	P[neg]
Atlanta	Peak	+1.34	+3.30	+5.88	100%	−0.63	97.4%
Boston	Peak	+0.06	+2.20	+4.00	97.8%	−0.49	93.3%
Chicago	Peak	−0.10	+1.61	+3.73	96.8%	−0.85	99.7%
Los Angeles	Peak	+1.86	+4.56	+7.82	100%	−1.89	100%
Miami	Peak	−4.21	+2.56	+7.59	78.9%	−0.82	100%
New York	Peak	−0.80	+3.24	+6.36	94.5%	−0.53	97.7%
Philadelphia	Peak	+0.64	+2.70	+4.78	99.5%	−0.66	98.3%
Salt Lake City	Peak	−2.83	+2.24	+6.28	82.2%	−0.62	100%
San Francisco	Peak	−0.21	+0.71	+1.65	93.8%	−0.36	81.4%
Seattle	Peak	−0.02	+1.80	+3.28	97.4%	−0.25	77.4%
All cities	Peak	+1.00	+2.13	+3.27	100%	−0.74	100%
All cities	Entire day	+1.20	+1.77	+2.38	100%	−0.82	100%

Note: Posterior percentiles from the hourly hierarchical model. Peak hours are 9:00–11:00 and 16:00–18:00. Entire-day estimates for each city are in SI Table 1.

Back

Affected segments: posterior summaries

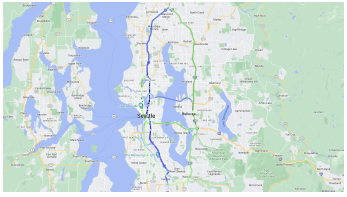
		Speed (%)				Fuel rate (%)	
City	Period	5th	Median	95th	P[pos]	Median	P[neg]
Atlanta	Peak	−0.34	+0.30	+0.95	82.9%	−0.04	74.7%
Boston	Peak	−0.31	+0.27	+0.86	82.9%	−0.06	83.4%
Chicago	Peak	+0.10	+0.77	+1.48	98.7%	−0.11	95.1%
Los Angeles	Peak	−0.26	+0.45	+1.26	88.9%	−0.07	87.9%
Miami	Peak	+0.10	+0.83	+1.64	98.7%	−0.10	93.3%
New York	Peak	−0.34	+0.32	+0.99	83.5%	−0.06	80.8%
Philadelphia	Peak	−0.30	+0.32	+0.95	84.7%	−0.06	81.0%
Salt Lake City	Peak	+0.09	+0.68	+1.29	98.9%	−0.09	90.7%
San Francisco	Peak	+0.08	+0.75	+1.45	98.4%	−0.10	92.3%
Seattle	Peak	−0.63	+0.22	+0.87	72.6%	−0.02	62.5%
All cities	Peak	+0.06	+0.49	+0.92	98.8%	−0.07	95.1%
All cities	Entire day	+0.12	+0.37	+0.64	99.8%	−0.07	98.6%

Note: Posterior percentiles from the hourly hierarchical model. Peak hours are 9:00–11:00 and 16:00–18:00. Entire-day estimates for each city are in SI Table 2.

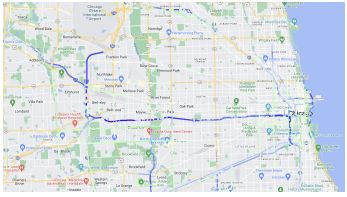
Back

Traffic dispersion in other cities

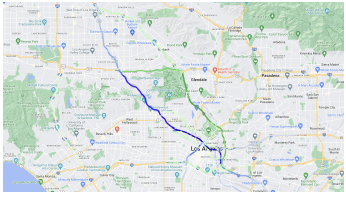
(a) Seattle



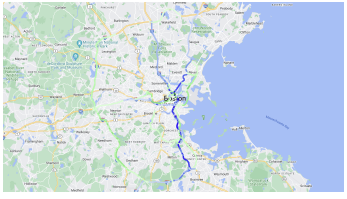
(b) Chicago



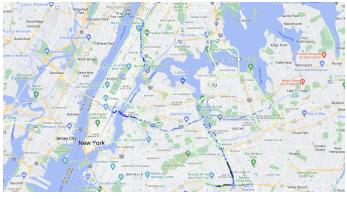
(c) Los Angeles



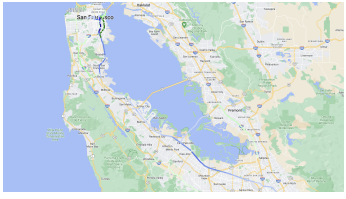
(d) Boston



(e) New York



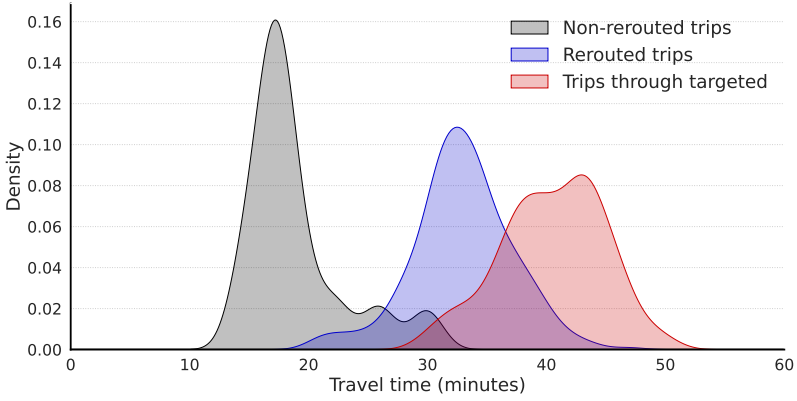
(f) San Francisco



Note: Projected volume change per segment per hour, normalized to the largest change in each city. Blue: fewer vehicles. Green: more vehicles. Imagery ©2026, Map data ©2026 Google.

Back

Trip travel times by subgroup (Atlanta)



Note: Trips through targeted segments are longer than typical trips. Among them, rerouted trips tend to be shorter, because those trips have attractive alternatives.

Back

Energy model

- Physics-based FASTSim simulator (NREL), validated on dynamometer data
- Too costly to run on every trip, so an ML surrogate (as in RouteE) is fit on segment speed, grade, and length
- Convex curve: fuel efficiency is best at intermediate speeds

