Information Design in Non-atomic Games: Computation, Repeated Setting and Experiment

Workshop on Algorithmic Game Theory, Mechanism Design, and Learning
Politecnico di Torino
8 Nov 2022

Ketan Savla
ksavla@usc.edu

University of Southern California

Information Design in Non-atomic Games: Computation, Repeated Setting and Experiment

Workshop on Algorithmic Game Theory, Mechanism Design, and Learning
Politecnico di Torino
8 Nov 2022

Ketan Savla
ksavla@usc.edu

University of Southern California

Joint work with: Yixian Zhu

Funding: NSF, USDoT

Routing Decisions in Urban Traffic Networks

- resource sharing among non-atomic agents
- uncertain and unpredictable environment

Routing Decisions in Urban Traffic Networks

- resource sharing among non-atomic agents
- uncertain and unpredictable environment

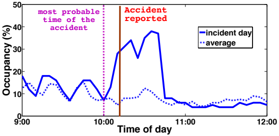

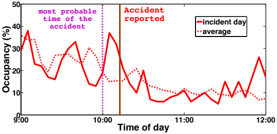

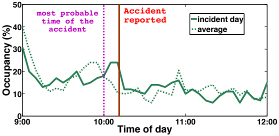

Routing Decisions in Urban Traffic Networks

- resource sharing among non-atomic agents
- uncertain and unpredictable environment
- repeated interactions

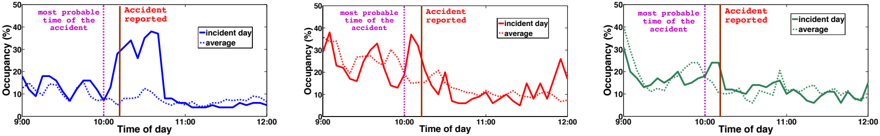

Routing Decisions in Urban Traffic Networks

- resource sharing among non-atomic agents
- uncertain and unpredictable environment
- repeated interactions

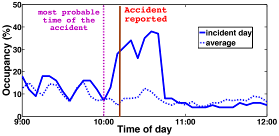

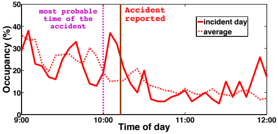

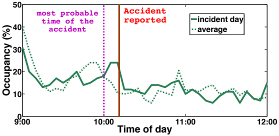

Information Design Setting

- Full info not optimal [AMMO:18]

[AMMO:18] Acemoglu, Makhdoumi, *et al*, "Informational Braess paradox: The effect of information on traffic congestion", *Operations Research*, 2018

Information Design Setting

- Full info not optimal [AMMO:18]

$$\min_{\text{signal}} \; E[\text{cost}(\text{signal})]$$

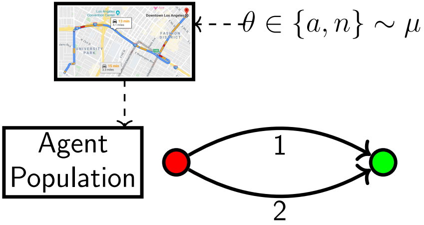

Information Design Setting

- Full info not optimal [AMMO:18]

$$\boxed{\min_{\text{signal}} \; E[\text{cost}(\text{signal})]}$$

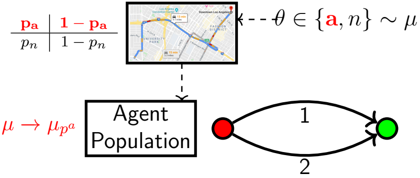

Information Design Setting

- Full info not optimal [AMMO:18]

$$\boxed{\min_{\text{signal}} \ E[\text{cost}(\text{signal})]}$$

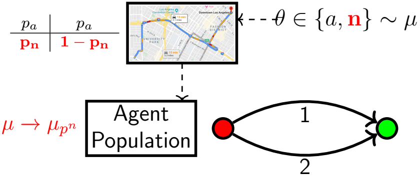

Information Design Setting

- Full info not optimal [AMMO:18]

$$\boxed{\min_{\text{signal}} \; E[\text{cost(signal)}]}$$

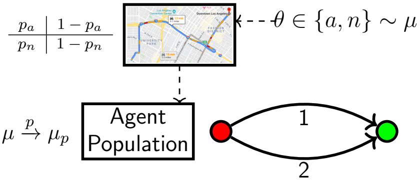

Information Design Setting

- Full info not optimal [AMMO:18]

$$\min_{\text{signal}} \; E[\text{cost}(\text{signal})]$$

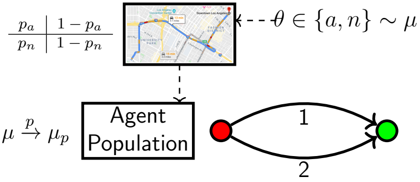

Signal Types

- private vs. public
- direct vs. indirect

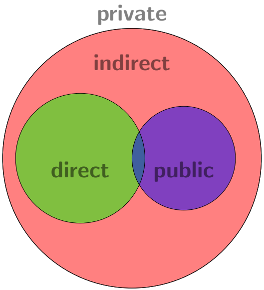

Randomization

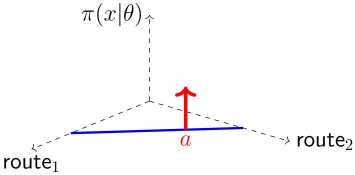

Randomization

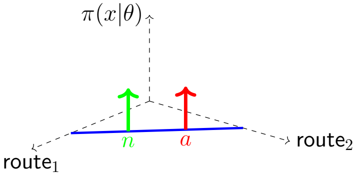

Randomization

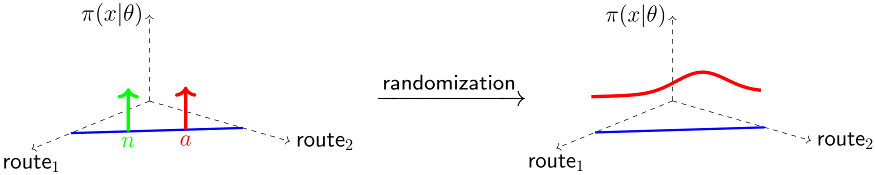

Randomization

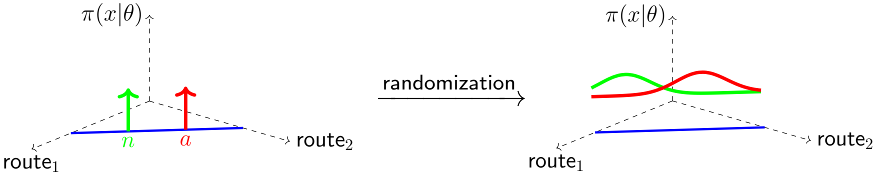

Randomization

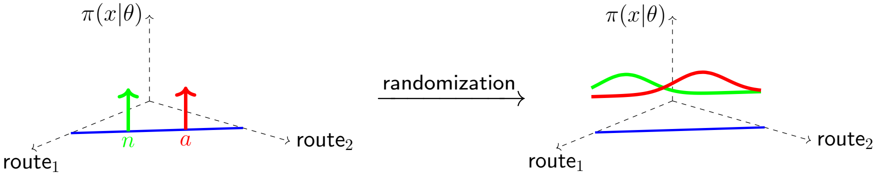

$$\mu \sim \theta \xrightarrow{\pi} x \rightarrow \text{posterior}_i(x,\theta) \propto x_i \, \pi(x|\theta) \, \mu(\theta)$$

- μ and π are public knowledge

Information Design

$$\boxed{\min_\pi \sum_{\theta,i} \int x_i \, \ell_i^\theta(x_i) \, \pi(x|\theta) \, \mu(\theta) \, \mathrm{d}x \quad \text{s.t. obedience constraint}}$$

- ℓ: link latency
- Existing works limited to stylized settings: [DKR:17], [TT:19]

[DKR:17]: S. Das, E. Kamenica, and R. Mirka, "Reducing congestion through information design", *Allerton 2017*
[TT:19]: H. Tavafoghi, D. Teneketzis, "Strategic information provision in routing games", 2019

Information Design

$$\min_\pi \sum_{\theta,i} \int x_i\, \ell_i^\theta(x_i)\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \quad \text{s.t. obedience constraint}$$

- ℓ: link latency
- Existing works limited to stylized settings: [DKR:17], [TT:19]

Obedience Constraint

Recommendation i is *obeyed* if:

$$\sum_\theta \int \left(\ell_i^\theta(x_i) - \ell_j^\theta(x_j)\right) \underbrace{\text{posterior}_i(x,\theta)}_{\propto\, x_i\, \pi(x|\theta)\, \mu(\theta)} \mathrm{d}x \le 0 \quad \forall\, j$$

Information Design

$$\min_\pi \sum_{\theta,i} \int x_i \, \ell_i^\theta(x_i) \, \pi(x|\theta) \, \mu(\theta) \, \mathrm{d}x \quad \text{s.t. obedience constraint}$$

- ℓ: link latency
- Existing works limited to stylized settings: [DKR:17], [TT:19]

Obedience Constraint

Recommendation i is *obeyed* if:

$$\sum_\theta \int \left(\ell_i^\theta(x_i) - \ell_j^\theta(x_j)\right) \underbrace{\text{posterior}_i(x,\theta)}_{\propto x_i \, \pi(x|\theta) \, \mu(\theta)} \mathrm{d}x \leq 0 \quad \forall j$$

infinite dimensional (non-convex) optimization!

Support Discretization

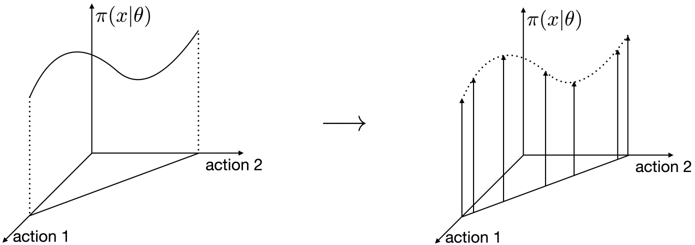

Support Discretization

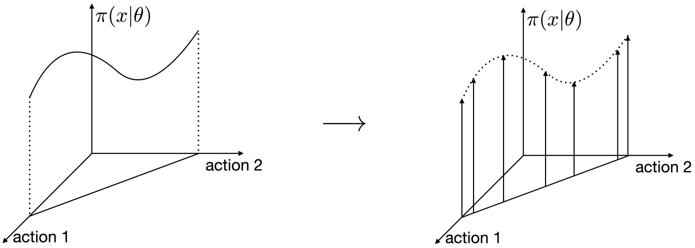

$$\sum_{\theta} \int \left(\ell_i^\theta(x_i) - \ell_j^\theta(x_j)\right) x_i \, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \quad \longrightarrow \quad \sum_{k,\,\theta} \left(\ell_i^\theta(x_i^{(k)}) - \ell_j^\theta(x_j^{(k)})\right) x_i^{(k)} \, \pi(x^{(k)}|\theta)\, \mu(\theta)$$

Support Discretization

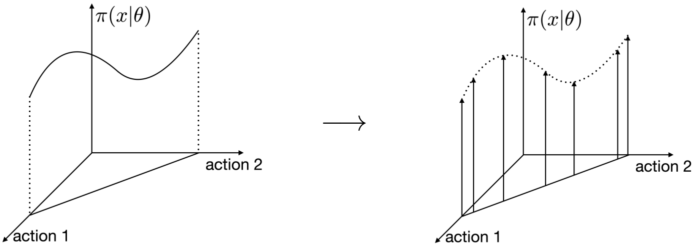

$$\sum_{\theta} \int \left(\ell_i^\theta(x_i) - \ell_j^\theta(x_j)\right) x_i \, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \quad \longrightarrow \quad \sum_{k,\,\theta} \left(\ell_i^\theta(x_i^{(k)}) - \ell_j^\theta(x_j^{(k)})\right) x_i^{(k)} \, \pi(x^{(k)}|\theta)\, \mu(\theta)$$

finite dimensional but non-convex

Equivalence under Finite Discretization

Optimize over Grid Points

$$\min_{\pi, \{x^{(k)}\}} \sum_{\theta,i,k} x_i^{(k)} \, \ell_i^\theta(x_i^{(k)}) \, \pi(x^{(k)}|\theta) \, \mu(\theta)$$

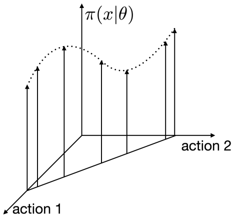

Equivalence under Finite Discretization

Optimize over Grid Points

$$\min_{\pi, \{x^{(k)}\}} \sum_{\theta,i,k} x_i^{(k)} \, \ell_i^\theta(x_i^{(k)}) \, \pi(x^{(k)}|\theta) \, \mu(\theta)$$

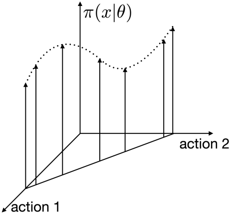

$+$

- polynomial ℓ

Example

BPR function is 4th order

Equivalence under Finite Discretization

Optimize over Grid Points

$$\min_{\pi, \{x^{(k)}\}} \sum_{\theta, i, k} x_i^{(k)} \, \ell_i^{\theta}(x_i^{(k)}) \, \pi(x^{(k)}|\theta) \, \mu(\theta)$$

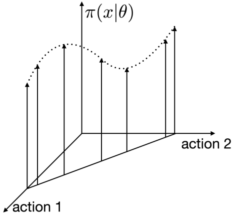

$+$

- polynomial ℓ

Example

BPR function is 4th order

$\downarrow$

equivalent poly optimization if # grid pts $\geq \binom{\text{poly order} + \text{num links}}{\text{poly order} + 1}$

Proof Sketch: Cubature Formula

Tchakaloff Thm

- $\exists \{x^{(k)}\}$ s.t.

$$\int \text{poly}(x)\,\pi(x)\,\mathrm{d}x = \sum_k \text{poly}(x^{(k)})\,\pi(x^{(k)}) \quad \forall \text{ poly}$$

Proof Sketch: Cubature Formula

Tchakaloff Thm

- $\exists \{x^{(k)}\}$ s.t.

$$\int \text{poly}(x)\,\pi(x)\,\mathrm{d}x = \sum_k \text{poly}(x^{(k)})\,\pi(x^{(k)}) \quad \forall \text{ poly}$$

Problem of Moments

$$\min_\pi \sum_{\theta,i} \int \underbrace{x_i\,\ell_i^\theta(x_i)}_{\text{moments}}\,\pi(x|\theta)\,\mu(\theta)\,\mathrm{d}x$$

$$\text{s.t. } \sum_\theta \int \left(\ell_i^\theta(x_i) - \underbrace{\ell_j^\theta(x_j)\right) x_i}_{\text{moments}}\,\pi(x|\theta)\,\mu(\theta)\,\mathrm{d}x \leq 0 \quad \forall j$$

Information Design as Polynomial Optimization

optimize over sufficiently many grid pts + polynomial ℓ

$$\downarrow$$

polynomial optimization (non-convex)

Information Design as Polynomial Optimization

optimize over sufficiently many grid pts + polynomial ℓ

$$\downarrow$$

polynomial optimization (non-convex)

- arbitrarily tight lower bound by hierarchy of semi-definite relaxations [L01]
- low hierarchies sufficient in practice
  - e.g., first hierarchy is tight for 2 links and affine ℓ

[L01]: J. B. Lasserre, "Global optimization with polynomials and the problem of moments", *SIAM Journal on Optimization*, 2001.

Extension to Heterogeneous Agents

- only a fraction participate in information design

Extension to Heterogeneous Agents

- only a fraction participate in information design

$$\min_{\pi, y} \sum_{\theta, i} \int (x_i + y_i)\, \ell_i^\theta (x_i + y_i)\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x$$

$$\text{s.t.} \quad \sum_{\theta} \int \left(\ell_i^\theta (x_i + y_i) - \ell_j^\theta (x_j + y_j) \right) x_i\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \le 0 \quad \forall\, i, j \quad \text{(obedience)}$$

Extension to Heterogeneous Agents

- only a fraction participate in information design

$$\min_{\pi,y} \sum_{\theta,i} \int (x_i+y_i)\, \ell_i^\theta(x_i+y_i)\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x$$

$$\text{s.t.} \sum_{\theta} \int \left(\ell_i^\theta(x_i+y_i) - \ell_j^\theta(x_j+y_j)\right) x_i\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \le 0 \quad \forall\, i,j \quad \text{(obedience)}$$

$$\sum_{\theta} \int \left(\ell_i^\theta(x_i+y_i) - \ell_j^\theta(x_j+y_j)\right) y_i\, \pi(x|\theta)\, \mu(\theta)\, \mathrm{d}x \le 0 \quad \forall\, i,j \quad \text{(Nash)}$$

Public vs Private Signals

- Private: optimal social cost non-increasing with participation rate

Public vs Private Signals

- Private: optimal social cost non-increasing with participation rate
- Public: optimal social cost may increase with participation rate

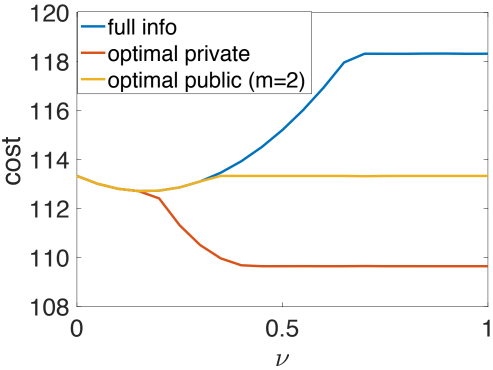

Towards Repeated Setting

Obedience involves agents computing $\sum_\theta \int \ell_i(x_i) \underbrace{x_i \, \pi(x|\theta) \, \mu(\theta)}_{\text{posterior}} \, \mathrm{d}x$

Towards Repeated Setting

Obedience involves agents computing $\sum_\theta \int \ell_i(x_i) \underbrace{x_i \, \pi(x|\theta) \, \mu(\theta)}_{\text{posterior}} \, \mathrm{d}x$

obedience constraint

- Bayesian calculation ✗
- long evaluation phase ✗
- requires knowledge of π, μ ✗

Towards Repeated Setting

Obedience involves agents computing $\sum_\theta \int \ell_i(x_i) \underbrace{x_i \, \pi(x|\theta) \, \mu(\theta)}_{\text{posterior}} \, \mathrm{d}x$

obedience constraint		"empirical" obedience?
Bayesian calculation ✗		myopic decision ✓
long evaluation phase ✗	$\rightarrow$	dynamic obedience ✓
requires knowledge of π, μ ✗		no knowledge of π, μ ✓

A Dynamic Model of Obedience

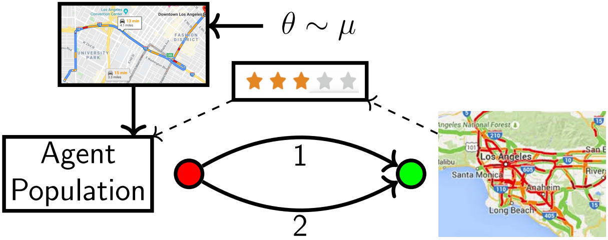

A Dynamic Model of Obedience

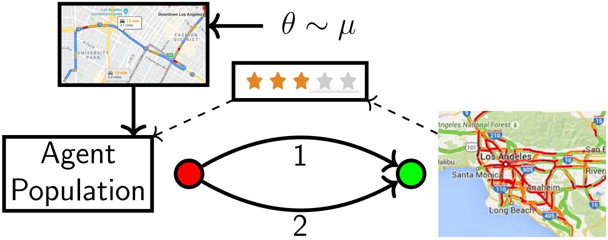

Participating Agents

- action space $\equiv$ {obey, do not obey}
- $\Pr(\text{obey}) \sim \underbrace{\sum_{\text{agents, stages}}}_{\text{socio-temporal}} \underbrace{\ell_{\text{reco}} - \ell_{\text{alternate}}}_{\text{regret}}$

A Dynamic Model of Obedience

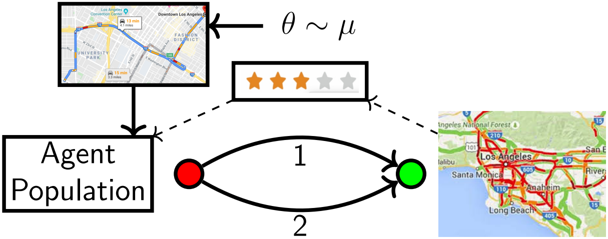

Participating Agents

- action space ≡ {obey, do not obey}
- $\Pr(\text{obey}) \sim \underbrace{\sum_{\text{agents, stages}}}_{\text{socio-temporal}} \underbrace{\ell_{\text{reco}} - \ell_{\text{alternate}}}_{\text{regret}}$

alternate

$$\text{reco} \begin{bmatrix} 0 & * & \dots & * \\ * & 0 & \dots & * \\ \vdots & \vdots & \vdots & * \\ * & * & \dots & 0 \end{bmatrix}$$

A Dynamic Model of Obedience

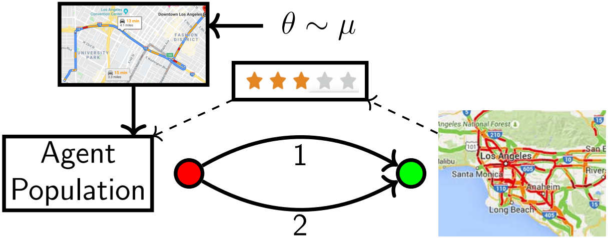

Non-participating Agents

- strategy space $\equiv \{1, 2\}$
- best response to $\Pr(\text{obey})$

A Dynamic Model of Obedience

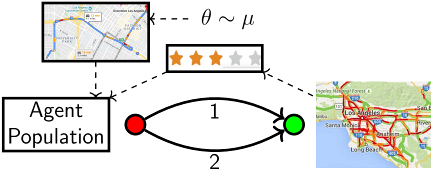

Participating Agents: ν fraction

- $\underbrace{\Pr(\text{obey})}_{q} \propto$  $\xleftarrow{\text{projection}}$ (socio-temp) avg regret

A Dynamic Model of Obedience

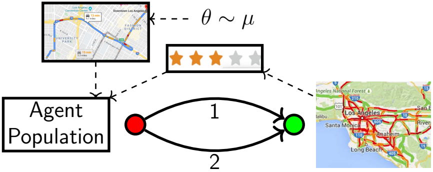

Participating Agents: ν fraction

- $\underbrace{\mathrm{Pr}(\text{obey})}_{q} \propto$  $\overset{\text{projection}}{\longleftarrow}$ (socio-temp) avg regret

$$\underbrace{x(q)}_{\text{actual link flows}} = \nu\, q\, \underbrace{\pi(.|\theta)}_{\text{reco link flows}} + \nu\,(1-q)\, \underbrace{P^T}_{\text{non-obedient routing}}\, \pi(.|\theta)$$

A Dynamic Model of Obedience

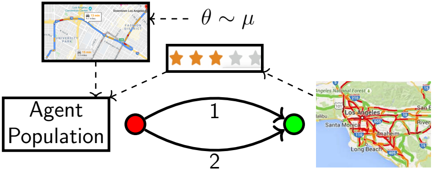

Participating Agents: ν fraction

- $\underbrace{\mathrm{Pr}(\text{obey})}_{q} \propto$  $\xleftarrow{\text{projection}}$ (socio-temp) avg regret

$$\underbrace{x(q)}_{\text{actual link flows}} = \nu\, q\, \underbrace{\pi(.|\theta)}_{\text{reco link flows}} + \nu\,(1-q)\, \underbrace{P^T}_{\text{non-obedient routing}}\, \pi(.|\theta)$$

Non-participating Agents: $1-\nu$ fraction

- $y(q)$: Bayes Nash flow for given q

Model Summary

Key Assumptions

- $\Pr(\text{obey}) \propto$ 
- avg regret $\xrightarrow{\text{projection}}$ 
- $\exists$ non-obedient routing matrix P

Model Summary

Key Assumptions

- $\Pr(\text{obey}) \propto$ 
- avg regret $\xrightarrow{\text{projection}}$ 
- $\exists$ non-obedient routing matrix P

Contrast with Literature: [HM:00,FV:97,...]

- finite agents
- action space $\equiv$ {route 1, route 2, ...}
- agent's action based on *personal* regret associated with action pairs

[HM:00]: S. Hart and A. Mas-Colell, "A simple adaptive procedure leading to correlated equilibrium", 2000.

[FV:97]: D. P. Foster and R. V. Vohra, "Calibrated learning and correlated equilibrium", 1997.

Experiment Procedure

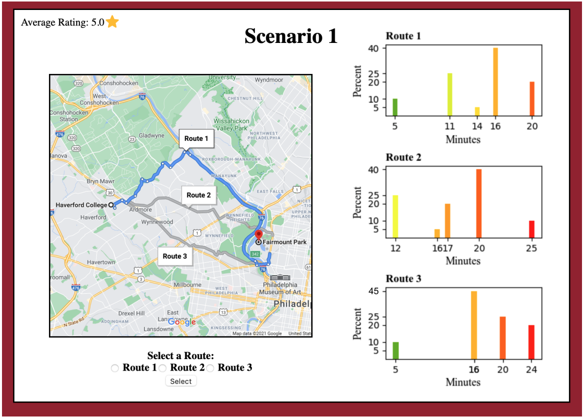

Experiment Procedure

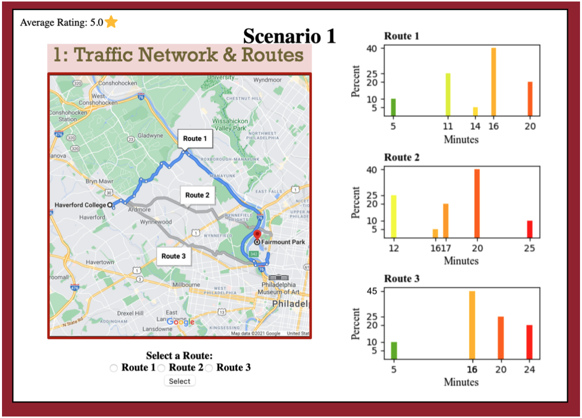

Experiment Procedure

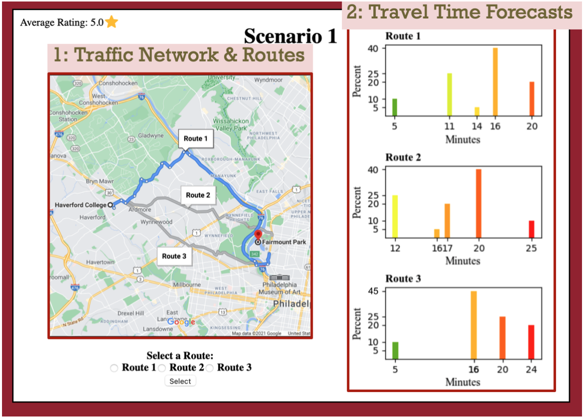

Experiment Procedure

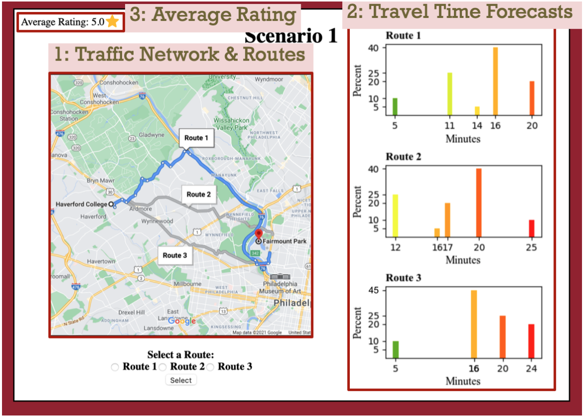

Experiment Procedure

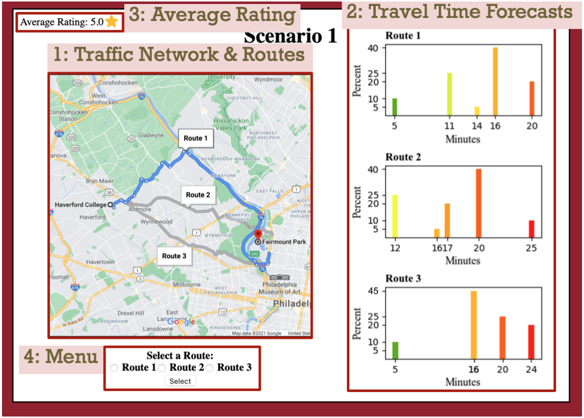

Experiment Procedure

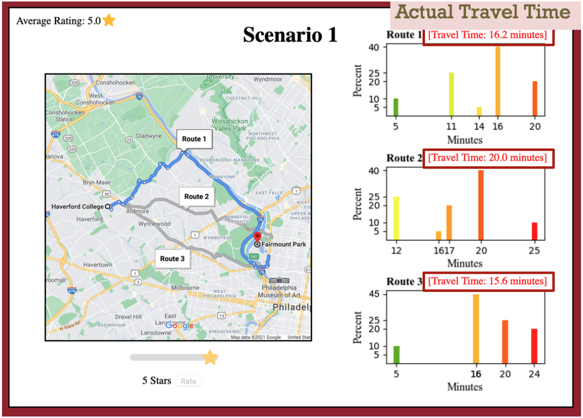

Experiment Procedure

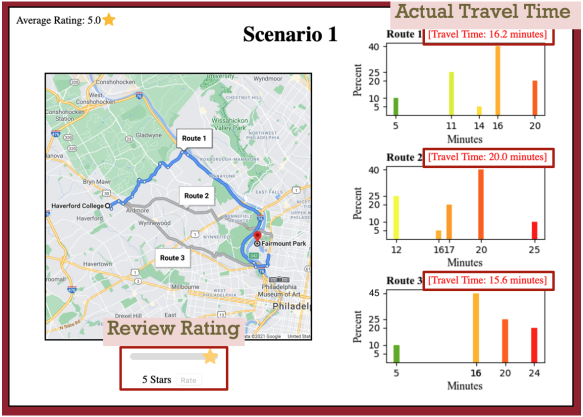

- $\nu = 1$
- 100 scenarios/participant $\times$ 33 participants

Connecting Experiment to Theory

- $\underbrace{\star\star\star\star\star}_{\text{rating}}(t+1) = \frac{t}{t+1}\text{rating}(t) + \frac{1}{t+1} \cdot \text{avg historical reviews}$
- $x = \text{rating} \cdot \pi(.|\theta) + (1 - \text{rating})\hat{P}^T\pi(.|\theta)$

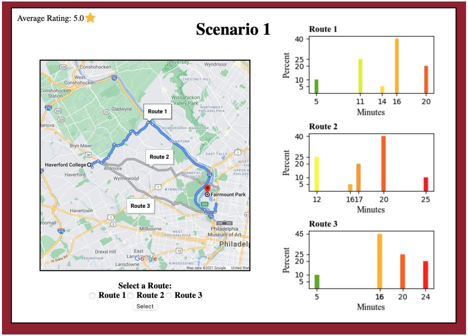

Connecting Experiment to Theory

-  $\underbrace{}_{\text{rating}}(t+1) = \frac{t}{t+1}\text{rating}(t) + \frac{1}{t+1} \cdot \text{avg historical reviews}$
- $x = \text{rating} \cdot \pi(.|\theta) + (1 - \text{rating})\hat{P}^T\pi(.|\theta)$

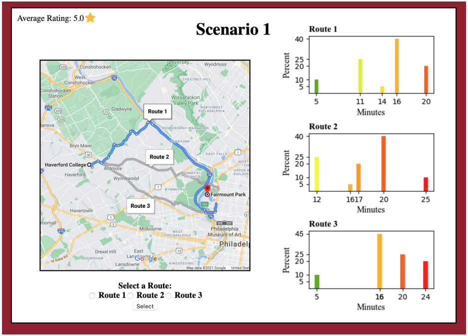

Hypotheses

- $\Pr(\text{obey}) \propto \text{rating}$?
- convergence of $\hat{P}$?

Connecting Experiment to Theory

- $\underbrace{\text{★★★☆☆}}_{\text{rating}}(t+1) = \frac{t}{t+1}\text{rating}(t) + \frac{1}{t+1} \cdot \text{avg historical reviews}$
- $x = \text{rating} \cdot \pi(.|\theta) + (1 - \text{rating})\hat{P}^T \pi(.|\theta)$

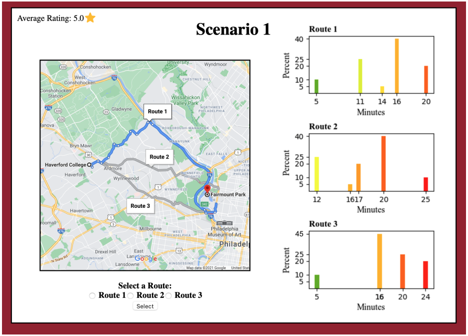

Hypotheses

- $\Pr(\text{obey}) \propto \text{rating}$?
- convergence of $\hat{P}$?
- $\text{rating} \propto -$ avg regret?
- $\text{rating}(t) \to$ ★★★★★ ?

Obedience vs Displayed Rating

$$\boxed{\Pr(\text{obey}) \propto \text{rating ?}}$$

Obedience vs Displayed Rating

$\Pr(\text{obey}) \propto \text{rating}$?

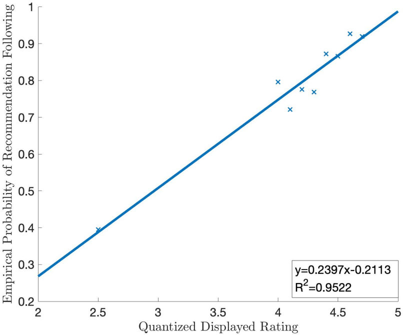

strong positive correlation ($R^2 = 0.9522$)

Long Run Behavior of $\hat{P}$

$\hat{P}$ converges with increasing s?

Long Run Behavior of $\hat{P}$

$\hat{P}$ converges with increasing s?

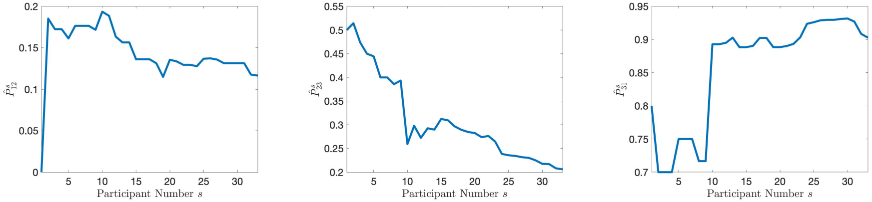

Displayed Rating vs Average Regret

$$\boxed{\text{rating} \propto -\text{avg regret ?}}$$

Displayed Rating vs Average Regret

rating $\propto$ −avg regret ?

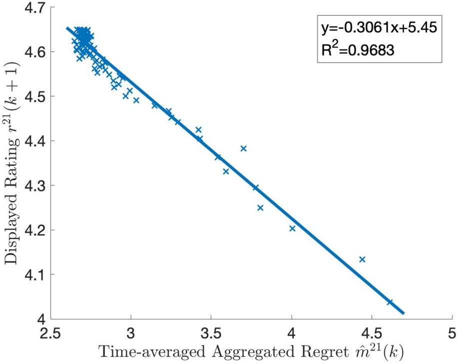

Strong negative correlation ($R^2 = 0.9683$) for an unbiased individual

Convergence

For all obedient π, $q(t) \to 1$ a.s. ★★★★★

Convergence

For all obedient π, $q(t) \to 1$ a.s. ★★★★★

Proof Sketch

- $\underline{\nu = 1}$: avg regret $\to \underbrace{\sum_\theta \mu(\theta)\, \pi^T(x|\theta)\, (I-P) \overbrace{\ell^*}^{\text{all obey}}}_{\leq 0\text{: obedience constraint}} - \underbrace{\cdots}_{\geq 0\text{: quadratic form}}$

Convergence

For all obedient π, $q(t) \to 1$ a.s. ★★★★★

Proof Sketch

- $\underline{\nu = 1}$: avg regret $\to \underbrace{\sum_{\theta} \mu(\theta)\, \pi^T(x|\theta)\,(I-P)\,\overbrace{\ell^*}^{\text{all obey}}}_{\leq 0\text{: obedience constraint}} - \underbrace{\cdots}_{\geq 0\text{: quadratic form}}$
- $\underline{\nu < 1}$:
  - $\exists$ avg regret subsequence $\to \leq 0$
  - deviation of parent sequence is sufficiently small asymptotically

Long Run Behavior of the Displayed Rating

$$\boxed{\text{rating}(t) \to ★★★★★\ ?}$$

Long Run Behavior of the Displayed Rating

$$\text{rating}(t) \rightarrow ★★★★★\ ?$$

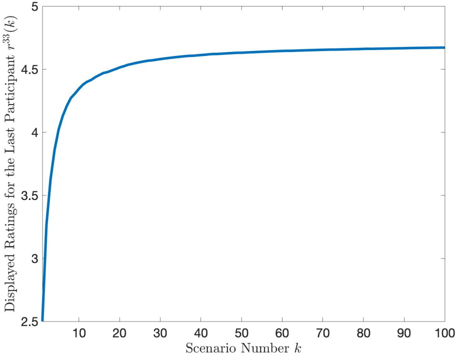

Long Run Behavior of $\Pr(\text{obey})$

$$\boxed{\Pr(\text{obey}) \to 1 \ ?}$$

Long Run Behavior of $\Pr(\text{obey})$

$$\Pr(\text{obey}) \to 1 \ ?$$

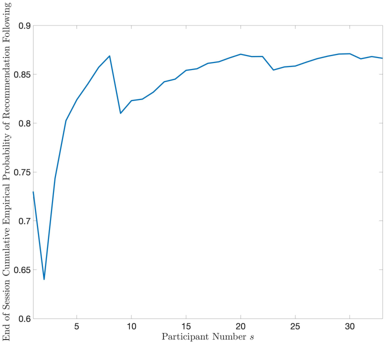

- $\Pr(\text{obey}) \to 0.87$
- consistent with $\Pr(\text{obey}) \to 0.91$ from the regression model

Summary

Key Takeaways

- Information design for non-atomic games as *finite* optimization
- A learning model for correlated equilibrium in non-atomic games
- Empirical evidence and analytical convergence

Summary

Key Takeaways

- Information design for non-atomic games as *finite* optimization
- A learning model for correlated equilibrium in non-atomic games
- Empirical evidence and analytical convergence

Ongoing Work

- Approximation algorithms
- Learning model for non participating agents
- Other decision-making settings: e.g., scheduling

Relevant Publications

- Y. Zhu, K. Savla, "*Information Design in Non-atomic Routing Games with Partial Participation: Computation & Properties*", IEEE Transactions on Control of Network Systems, 2022.
- Y. Zhu, K. Savla, "*Convergence in a Repeated Non-atomic Routing Game with Partial Signaling*", `arxiv` 2022.
- Y. Zhu, K. Savla, "*An Experimental Study on Learning Correlated Equilibrium in Routing Games*", `arxiv` 2022.

Relevant Publications

- Y. Zhu, K. Savla, "*Information Design in Non-atomic Routing Games with Partial Participation: Computation & Properties*", IEEE Transactions on Control of Network Systems, 2022.
- Y. Zhu, K. Savla, "*Convergence in a Repeated Non-atomic Routing Game with Partial Signaling*", `arxiv` 2022.
- Y. Zhu, K. Savla, "*An Experimental Study on Learning Correlated Equilibrium in Routing Games*", `arxiv` 2022.

`https://viterbi-web.usc.edu/~ksavla/publications.html`

`ksavla@usc.edu`