Beyond FORCES: The Emerging Data Market and Sharing Economy

Lillian Ratliff and Shankar Sastry

University of California, Berkeley

June 5, 2016

Outline

From Action Webs to Resilient CPS

Resilient/High Confidence Networked Control

- Fault-tolerant networked control
  - Limits on stability, safety, & optimality
  - Scalable model predictive control
- Security & Resilient Control
  - Availability, Integrity, & Confidentiality
  - Graceful degradation
- Economic Incentives
  - Incentive Design for investing in security
  - Interdependent Risk Assessment & Cyber Insurance

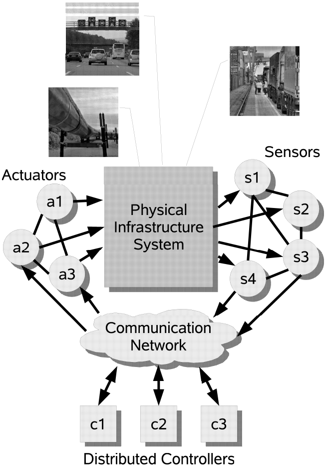

Societal Scale CPS

A complex collection of sensors, controllers, compute nodes, and actuators that work together to improve our daily lives

- From very small: Ubiquitous, Pervasive, Disappearing, Perceptive, Ambient
- To very large: Always Connectable, Reliable, Scalable, Adaptive, Flexible

Emerging Service Models

- Building energy management
- Automotive safety and control
- Management of metropolitan traffic flows
- Distributed health monitoring
- Smart Grid

Outline

Smart, Connected Infrastructure

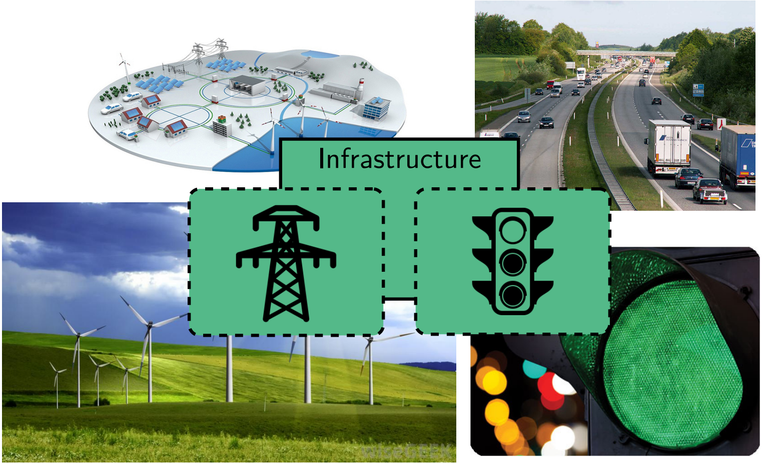

Operational Efficiency Informed by Usage Patterns

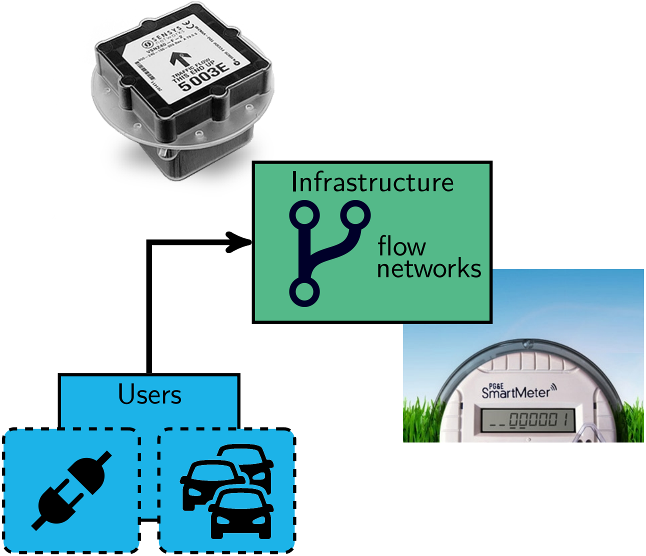

Data Disaggregation

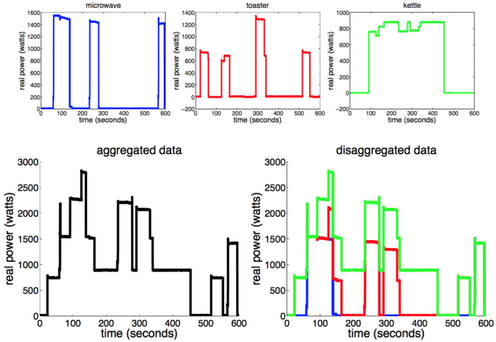

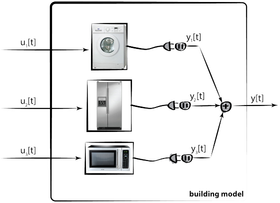

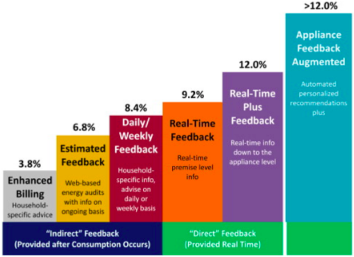

(a) improved load forecasting; (b) improved economic models

Modeling Urban Mobility Using Data

Data: parking transactions, meter locations, land use, economic indicators, etc.

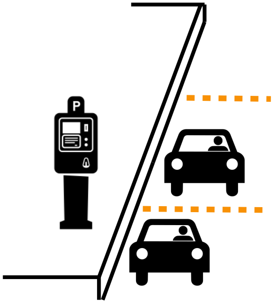

Modeling Urban Mobility Using Data

Data: parking transactions, meter locations, land use, economic indicators, etc.

↓

Features: arrival rates, geo-location, histogram of parking durations

↓

Machine Learning: weighted clustering

Learning

SDOT
Seattle Department of Transportation

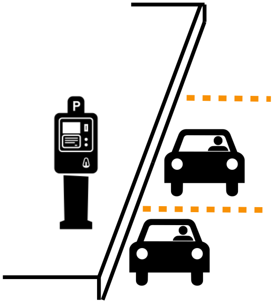

Modeling Urban Mobility Using Data

Data: parking transactions, meter locations, land use, economic indicators, etc.

Features: arrival rates, geo-location, histogram of parking durations

Machine Learning: weighted clustering

Learn natural neighborhoods
for more effective pricing

Downtown Seattle, 2013

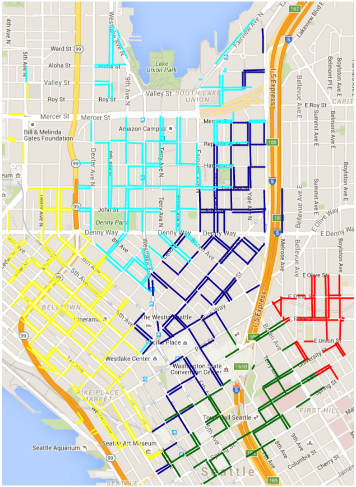

Modeling Urban Mobility Using Data

Data: parking transactions, meter locations, land use, economic indicators, etc.

↓

Features: blockface utilization

↓

Machine Learning: Hodge decomposition ranking

↓

Learn source and sink blockfaces to determine locations for validation studies

Belltown, 2015

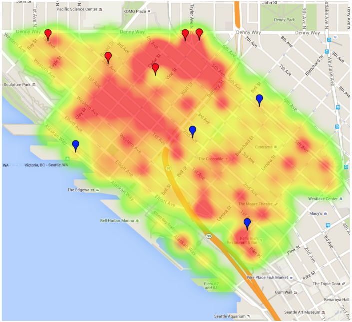

Outline

Closing the Loop — Integrating the User

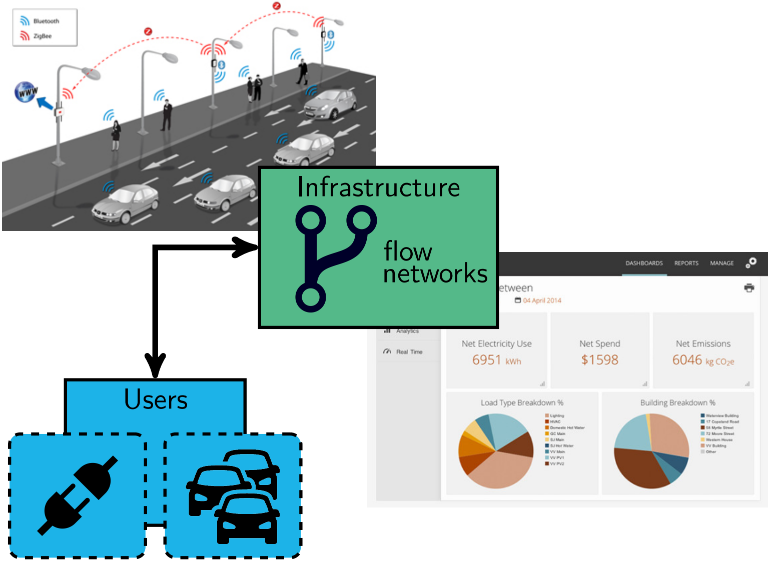

Quantifying the Efficiency–Privacy Tradeoff

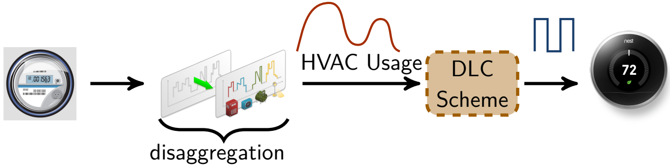

Efficiency-Privacy Tradeoff: DLC performance degrades as privacy–preserving metering is increased.

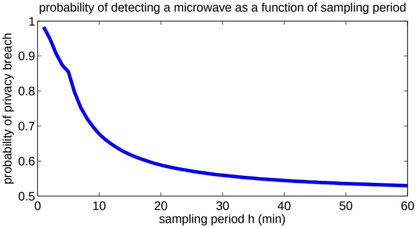

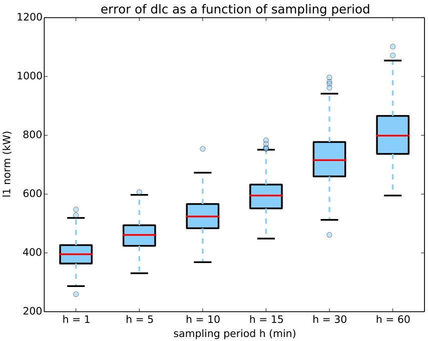

R. Dong, A. Cardenas, L. Ratliff, H. Ohlsson, S. Sastry. IEEE TSG, 2014 (under review, arxiv:1406.2568)

How do people value their data? — Privacy as Good

Contribution: Designed service contracts differentiated by value of data to balance efficiency-vulnerability tradeoff

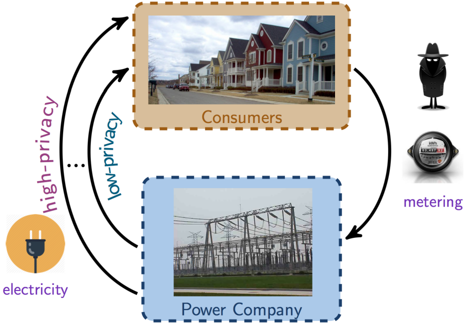

Results:
- Characterized contracts with privacy loss risk modeled using privacy metric and abstraction of loss.
- High-type free-rides ⇒ regulation to realize the social optimum.
- Privacy loss risk ⇒ incentive for investing in insurance.
- Designed insurance contracts for risk-averse utility company/ consumer.

Impact:
- Privacy loss risk motivates study of **security-insurance** investment.
- User valuations of data need to be factored in to improve efficiency.

L. Ratliff, *et al.* arxiv:1409.7926v3, 2014; L. Ratliff, *et al.* CDC 2014.

Queuing Game Framework for Urban Parking

Seattle Department of Transportation

Data: parking transactions, meter locations, land use, economic indicators, etc.

Features: arrival rates, geo-location, histogram of parking durations

Machine Learning: weighted clustering

Learning

System of Queues: M/G/n queue model for each cluster

Game Theory: supermarket game

Framework to Design Feedback

Supermarket Game & the Value of Information

Value of Information: expected reduction in expected waiting time due to a gain in information

\$ for Info: mean service time, arrival rate, expected occupancy, price, etc.

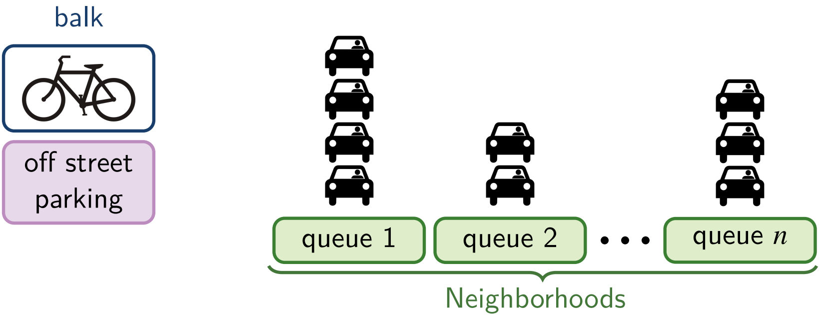

To Join or Not to Join

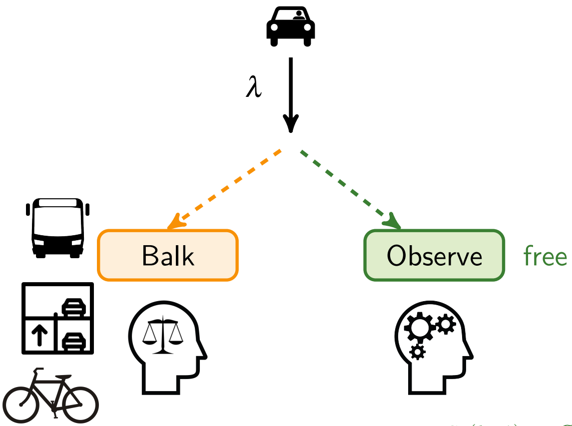

$U_b = 0$

(or outside opt)

$U_j = R - \frac{C_w(k+1)}{c\mu} - \frac{C_p}{\mu}$

$$R - \frac{C_w k}{\mu c} - \frac{C_p}{\mu} = 0$$

$$\downarrow$$

$$n_b = \left\lfloor \frac{R\mu c - C_p c}{C_w} \right\rfloor$$

(balking level)

Social Welfare the Effects of Balking & Price

$$U_{sw}(n) = \lambda \sum_{k=0}^{n-1} p_k(n)\left(R - \frac{C_w(k+1)}{\mu c} - \frac{C_p}{\mu}\right)$$

Thm: There exists n_{so} maximizing $U_{sw}(n)$ and $n_{so} \leqslant n_b$ so that $U_{sw}(n_b) \leqslant U_{sw}(n_{so})$.

Social Welfare the Effects of Balking & Price

$$U_{sw}(n) = \lambda \sum_{k=0}^{n-1} p_k(n) \left(R - \frac{C_w(k+1)}{\mu c} - \frac{C_p}{\mu} \right)$$

Thm: There exists n_{so} maximizing $U_{sw}(n)$ and $n_{so} \leqslant n_b$ so that $U_{sw}(n_b) \leqslant U_{sw}(n_{so})$.

Design price: $\hat{C}_p = C_p + \Delta C_p$

Prop: The pricing $\hat{C}_p$ that achieves the socially optimal balking level n_{so} is determined by $\alpha_{n_{so}} \leqslant \frac{\hat{C}_p}{\mu} \leqslant \alpha_{n_{so}-1}$ where $\alpha_k = R - \frac{C_w(k+1)}{\mu c}$.

Social Welfare the Effects of Balking & Price

$$U_{sw}(n) = \lambda \sum_{k=0}^{n-1} p_k(n)\left(R - \frac{C_w(k+1)}{\mu c} - \frac{C_p}{\mu}\right)$$

Thm: There exists n_{so} maximizing $U_{sw}(n)$ and $n_{so} \leqslant n_b$ so that $U_{sw}(n_b) \leqslant U_{sw}(n_{so})$.

Design price: $\hat{C}_p = C_p + \Delta C_p$

Prop: The pricing $\hat{C}_p$ that achieves the socially optimal balking level n_{so} is determined by $\alpha_{n_{so}} \leqslant \frac{\hat{C}_p}{\mu} \leqslant \alpha_{n_{so}-1}$ where $\alpha_k = R - \frac{C_w(k+1)}{\mu c}$.

congested limited balking: e.g. $n_{cl} = 20\%$ total volume

Prop: The pricing $\hat{C}_p$ that achieves a congestion limited balking level n_{cl} is determined by $\alpha_{n_{cl}} \leqslant \frac{\hat{C}_p}{\mu} \leqslant \alpha_{n_{cl}-1}$ where $\alpha_k = R - \frac{C_w(k+1)}{\mu c}$.

Social Welfare the Effects of Balking & Price

$$U_{sw}(n) = \lambda \sum_{k=0}^{n-1} p_k(n) \left(R - \frac{C_w(k+1)}{\mu c} - \frac{C_p}{\mu} \right)$$

Thm: There exists n_{so} maximizing $U_{sw}(n)$ and $n_{so} \leqslant n_b$ so that $U_{sw}(n_b) \leqslant U_{sw}(n_{so})$.

Design price: $\hat{C}_p = C_p + \Delta C_p$

Prop: The pricing $\hat{C}_p$ that achieves the socially optimal balking level n_{so} is determined by $\alpha_{n_{so}} \leqslant \frac{\hat{C}_p}{\mu} \leqslant \alpha_{n_{so}-1}$ where $\alpha_k = R - \frac{C_w(k+1)}{\mu c}$.

congested limited balking: e.g. $n_{cl} = 20\%$ total volume

Prop: The pricing $\hat{C}_p$ that achieves a congestion limited balking level n_{cl} is determined by $\alpha_{n_{cl}} \leqslant \frac{\hat{C}_p}{\mu} \leqslant \alpha_{n_{cl}-1}$ where $\alpha_k = R - \frac{C_w(k+1)}{\mu c}$.

Prop: If $n_{cl} \leqslant n_b$ (user selected), then $U_{sw}(n_b) \leqslant U_{sw}(n_{cl})$.

Prop: If $n_{cl} \leqslant n_{so}$, then $U_{sw}(n_{cl}) = U_{sw}(n_{so})$.

To Observe or Not to Observe

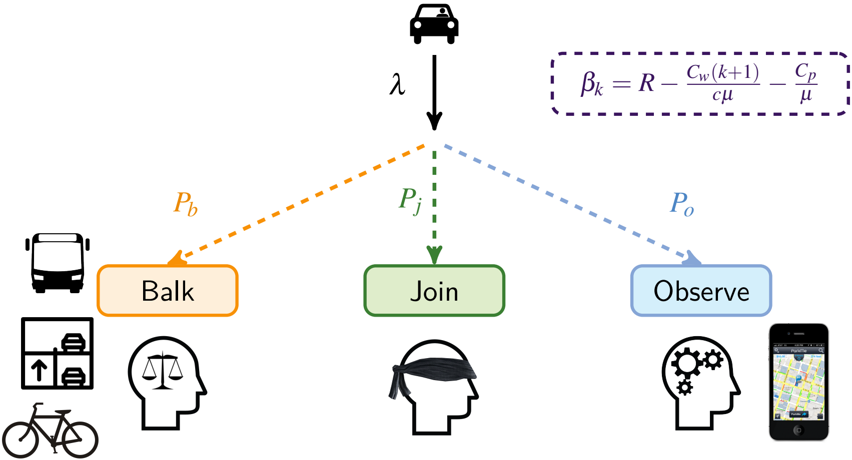

$U_b = 0$

(or outside opt)

$$U_j = \sum_{k=0}^{n-1} p_k \beta_k$$

$$U_o = \sum_{k=0}^{n_b-1} p_k \beta_k - C_o$$

Ratliff, et al. To Observe or Not to Observe: Queuing Game Framework for Urban Parking. IEEE CDC 2016 (submitted)

On-Street vs. Off-Street Parking Example

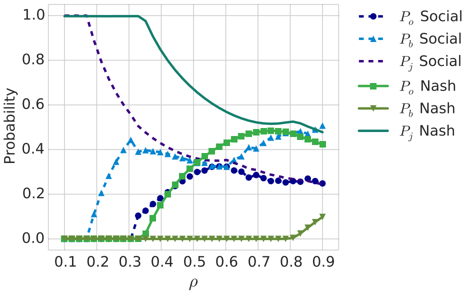

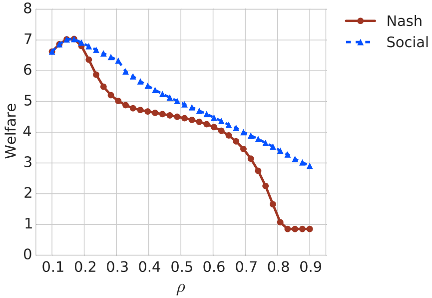

$c = 30, \mu = \frac{1}{120}$ (service rate)

$\lambda \in [0.025, 0.225]$ (arrival rate)

$\rho = \frac{\lambda}{c\mu}$ (traffic intensity)

As $\rho \to 1$, it seems $|P_j - P_o| \to \varepsilon$

$P_b^* = 0$ for $\rho \in [0, 0.8]$ while $P_b^{so} > 0$ for $\rho > 0.15$

Nash: less people use off-street parking $\Rightarrow$ spend more time **circling**

Congestion vs. Occupancy

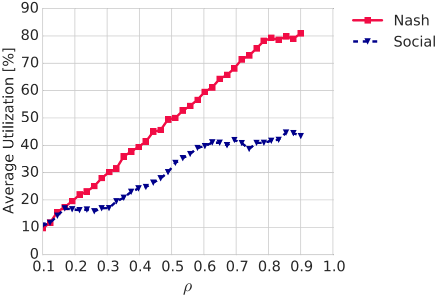

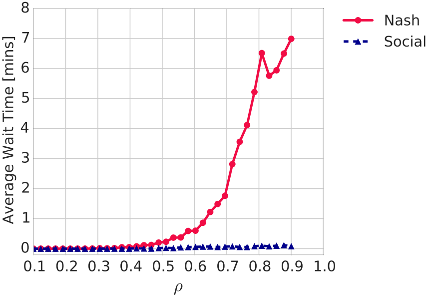

drivers arrive, if block is full they choose a new block uniformly at random and there is a travel time along links

$C_p = \$0.05/\text{min}$, $C_o = \$3.85$

$C_w = \$1.5/\text{min}$, $C_{off} = \$0.9/\text{min}$

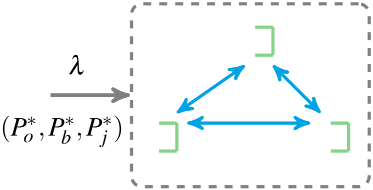

Ratliff, et al. To Observe or Not to Observe: Queuing Game Framework for Urban Parking. IEEE CDC 2016 (submitted)

Outline

Emerging Data Market – User Models

Emerging Data Market – Closing the Loop

Emerging Data Market – Sharing Economy

Conclusions and Future Research

Data as a Commodity

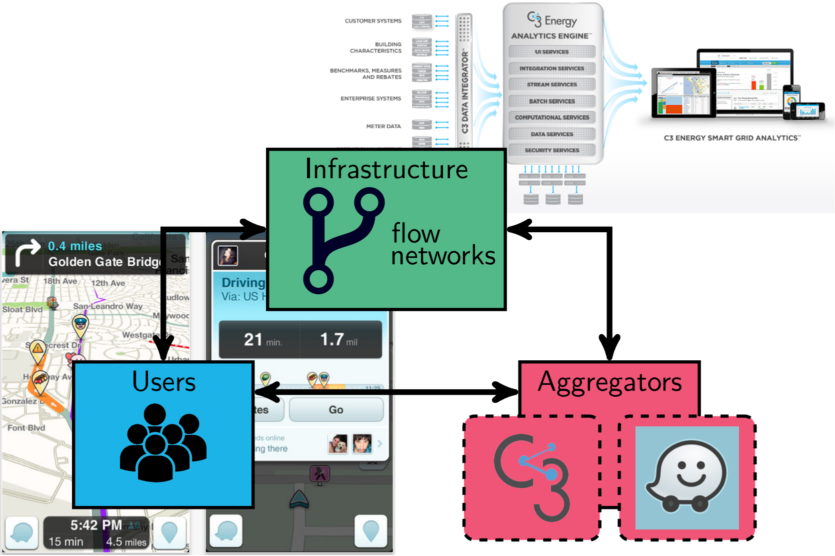

Shared Economy & Platform Markets

A smart infrastructure empowered by the Internet of Things (IoT) has at its core an ecosystem consisting of a *shared economy*.

Learning & Optimization with Strategic Sources

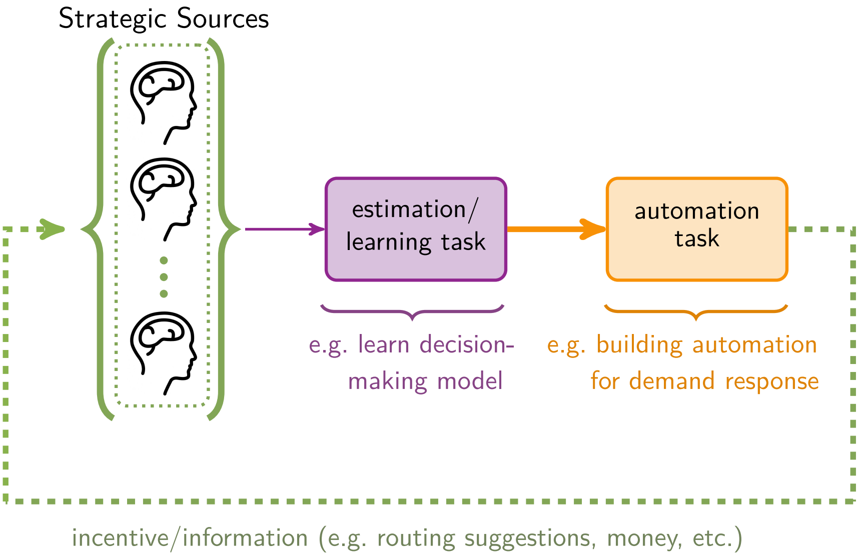

Abstraction of the Adaptive Incentive Design Problem

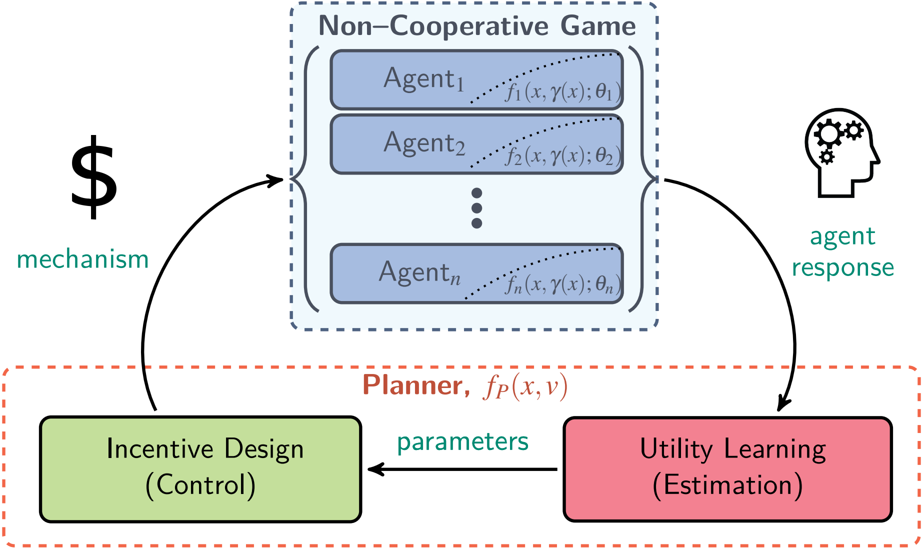

Closing–the–Loop via Adaptive Incentive Design

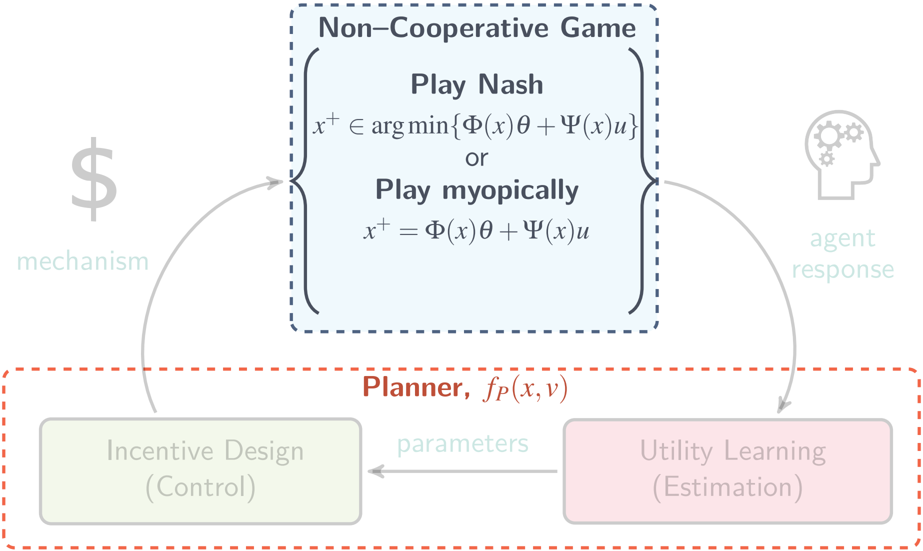

Closing–the–Loop via Adaptive Incentive Design

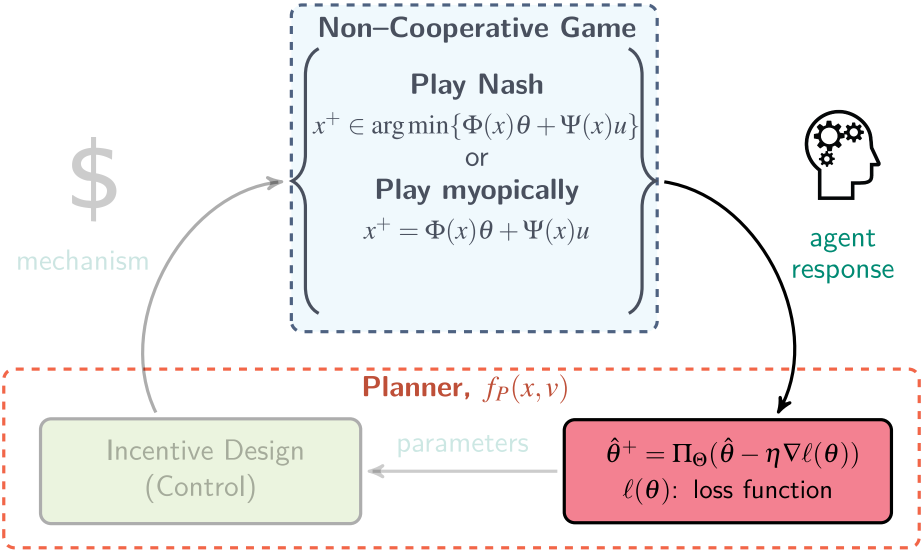

Closing–the–Loop via Adaptive Incentive Design

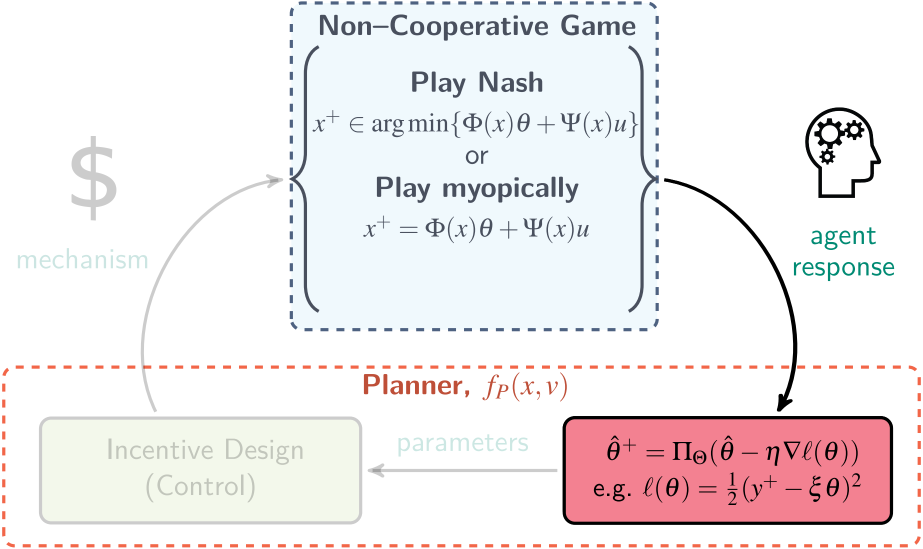

Closing–the–Loop via Adaptive Incentive Design

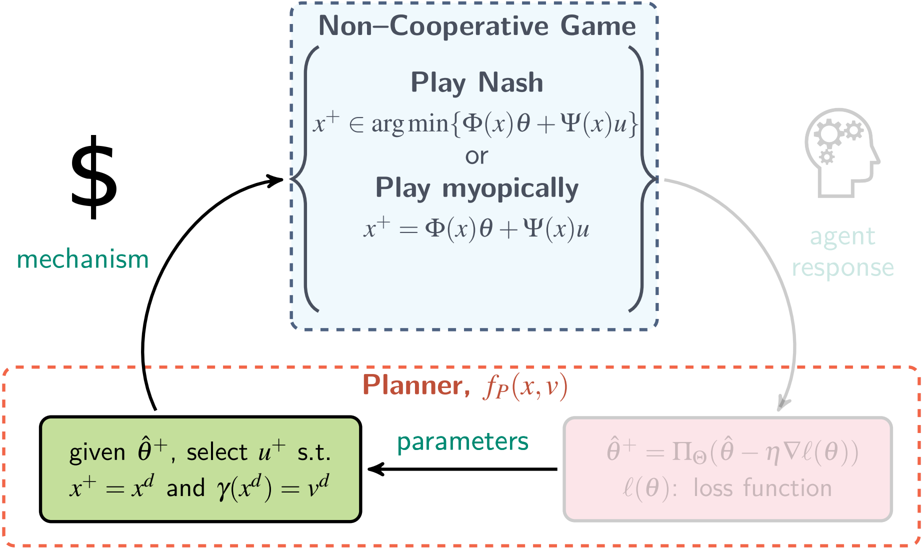

Closing–the–Loop via Adaptive Incentive Design

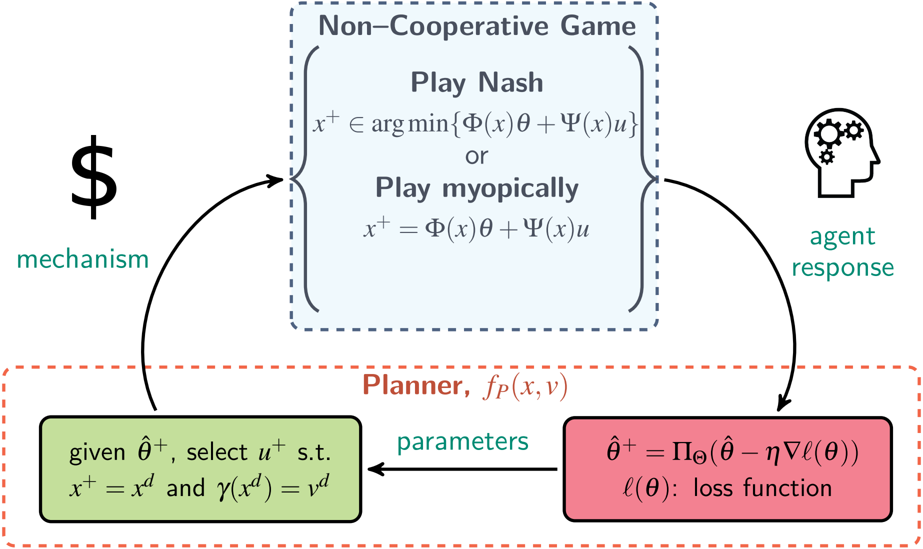

Convergence Results — Adaptive Control/Online Learning

A1: For all agent preferences θ, there exists incentive parameters s.t. the agents play the desired strategy.

Thm: Suppose that the algorithm is *persistently exciting* and *stable* $(c_2 I \leqslant \xi^T \xi \leqslant c_1 I)$, then $\hat{\theta} \to \theta$.

parameter convergence

Convergence Results — Adaptive Control/Online Learning

A1: For all agent preferences θ, there exists incentive parameters s.t. the agents play the desired strategy.

Thm: Suppose that the algorithm is *persistently exciting* and *stable* ($c_2 I \leqslant \xi^T \xi \leqslant c_1 I$), then $\hat{\theta} \to \theta$.

parameter convergence

Thm: Agent strategies and incentive parameters converge to the desired values ($x_i \to x_i^d$ and $u_i \to u_i^d$).

myopic

Convergence Results — Adaptive Control/Online Learning

A1: For all agent preferences θ, there exists incentive parameters s.t. the agents play the desired strategy.

Thm: Suppose that the algorithm is *persistently exciting* and *stable* ($c_2 I \leqslant \xi^T \xi \leqslant c_1 I$), then $\hat{\theta} \to \theta$.

parameter convergence

Thm: Agent strategies and incentive parameters converge to the desired values ($x_i \to x_i^d$ and $u_i \to u_i^d$).

myopic

Thm: $\left[x^d\right.$ non-degenerate differential Nash ($\det D\omega(x^d) \neq 0$) and $\left.\|\hat{\theta} - \theta\| < \varepsilon\right] \implies \exists$ N.E. x^* close to x^d.

Nash

Convergence Results — Adaptive Control/Online Learning

A1: For all agent preferences θ, there exists incentive parameters s.t. the agents play the desired strategy.

Thm: Suppose that the algorithm is *persistently exciting* and *stable* ($c_2 I \leqslant \xi^T \xi \leqslant c_1 I$), then $\hat{\theta} \to \theta$.

parameter convergence

Thm: Agent strategies and incentive parameters converge to the desired values ($x_i \to x_i^d$ and $u_i \to u_i^d$).

myopic

Thm: $\left[x^d\right.$ non-degenerate differential Nash ($\det D\omega(x^d) \neq 0$) and $\left.\|\hat{\theta} - \theta\| < \varepsilon\right] \implies \exists$ N.E. x^* close to x^d.

Thm: If x^* is stable ($D\omega(x^*) > 0$), then users will converge to x^* under tâtonnement (gradient play).

Nash

Convergence Results — Adaptive Control/Online Learning

A1: For all agent preferences θ, there exists incentive parameters s.t. the agents play the desired strategy.

Thm: Suppose that the algorithm is *persistently exciting* and *stable* ($c_2 I \leqslant \xi^T \xi \leqslant c_1 I$), then $\hat{\theta} \to \theta$.

parameter convergence

Thm: Agent strategies and incentive parameters converge to the desired values ($x_i \to x_i^d$ and $u_i \to u_i^d$).

myopic

Thm: $\left[x^d\right.$ non-degenerate differential Nash ($\det D\omega(x^d) \neq 0$) and $\left.\|\hat{\theta} - \theta\| < \varepsilon\right] \implies \exists$ N.E. x^* close to x^d.

Thm: If x^* is stable ($D\omega(x^*) > 0$), then users will converge to x^* under tâtonnement (gradient play).

Nash

Noise: Under classical assumptions, convergence with noise!

L. Ratliff, Thesis, UC Berkeley, 2015; L. Ratliff, et al., IEEE TAC, 2015; L. Ratliff, et al., SICON, 2015.

Social Game — Robust Utility Learning

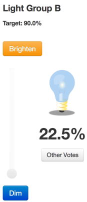

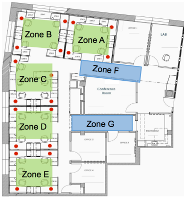

Strategic Sources

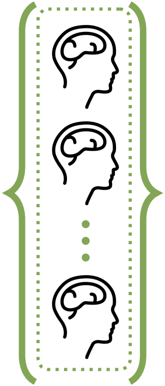

estimation/ learning task

automation task

e.g. learn decision-making model

e.g. building automation for demand response

incentive/information (e.g. routing suggestions, money, etc.)

Robust Utility Learning: Bootstrapping, Boosting, Bagging

1. Use equilibrium conditions for differential Nash equilibria, to construct a Constrained Generalized Least Squares problem in agent parameters.

$Y = X\theta + \varepsilon$

$\text{cov}(\varepsilon|X) = G \succ 0$

$G^{-1/2}Y = G^{-1/2}X\theta + G^{-1/2}\varepsilon$

G is unknown.
We impose structure;
let's us learn coalitions!

$\theta_{GLS}, \hat{G}$

2. *Wild bootstrapping* and bagging:
we have little data and we expect bias.

Generate N pseudo-datasets;
fit N *weak* FGLS estimators

$\tilde{Y} = X\theta_{GLS} + \hat{G}^{1/2}\varepsilon$

$\theta_{\text{Bagged}} = \frac{1}{N}\sum_{k=1}^{N} \theta_{FGLS}^{k}$

Robust Utility Learning — Results

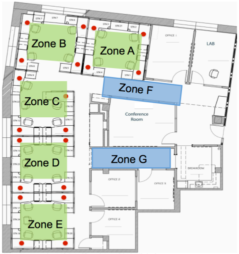

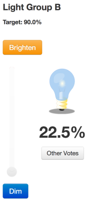

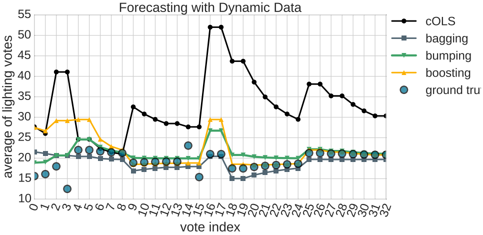

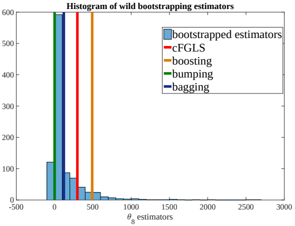

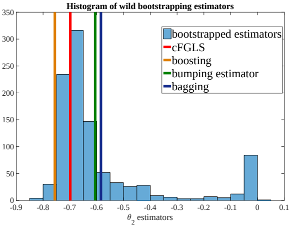

Mixture of Utilities — Myopic Decision-Making

$$f(x) = \phi(x) + \sum_{j=1}^{M} \pi_j(x, \xi_j)\psi(x) = \phi(x) + \sum_{j=1}^{M} \frac{\exp(-\xi_j^\top x)}{\sum_{k=1}^{M} \exp(-\xi_k^\top x)} \theta_j \psi(x)$$

θ_1 θ_2 θ_3

I. Konstantakopolous, et al. IEEE CDC 2016; L. Ratliff, et al., Allerton 2014

Mixture of Utilities — Myopic Decision-Making

$$f(x) = \phi(x) + \sum_{j=1}^{M} \pi_j(x, \xi_j)\psi(x) = \phi(x) + \sum_{j=1}^{M} \frac{\exp(-\xi_j^\top x)}{\sum_{k=1}^{M} \exp(-\xi_k^\top x)} \theta_j \psi(x)$$

θ_1 θ_2 θ_3

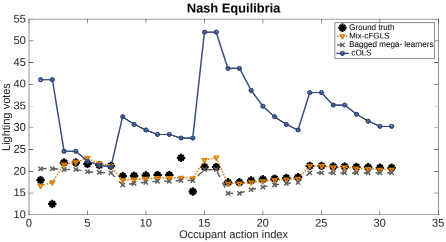

I. Konstantakopolous, et al. IEEE CDC 2016; L. Ratliff, et al., Allerton 2014

Mixture of Utilities — Myopic Decision-Making

$$f(x) = \phi(x) + \sum_{j=1}^{M} \pi_j(x, \xi_j)\psi(x) = \phi(x) + \sum_{j=1}^{M} \frac{\exp(-\xi_j^\top x)}{\sum_{k=1}^{M}\exp(-\xi_k^\top x)}\theta_j\psi(x)$$

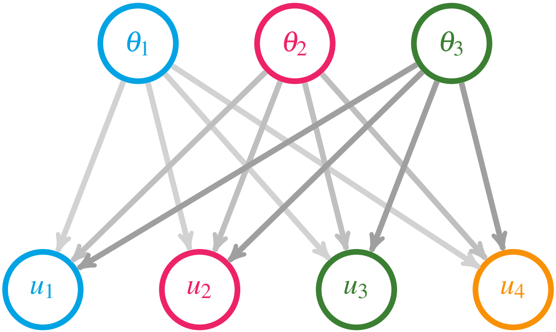

$$\psi(x) = \sum_{\ell=1}^{N} \frac{\exp(u_\ell(x))}{\sum_{k=1}^{N}\exp(u_k(x))}$$ Luce-Shepard

I. Konstantakopolous, et al. IEEE CDC 2016; L. Ratliff, et al., Allerton 2014

Outline

Open Data Initiatives and Data Sharing Mechanisms

Many cities are adopting open data policies in which all data collected by municipal service providers is made available.

At the same time, third-party companies are emerging on the scene to provide services to cities (e.g. platform-as-a-service). In addition, companies often want to share data with researchers.

Can we generate mechanisms that deconflict open data policies with intellectual property protection?

SIDE WALK LABS

Swiftly

SDOT

Seattle Department of Transportation

Privacy Guarantees for Data Exchange at Scale

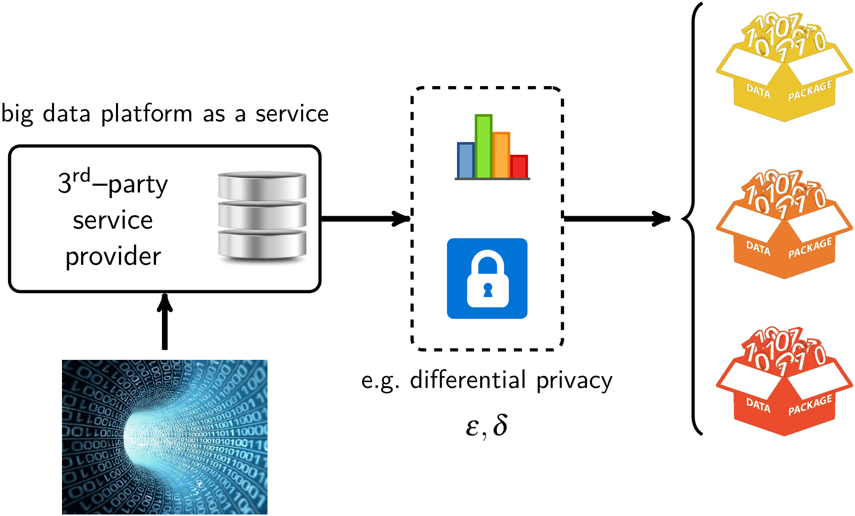

Old School Regulation

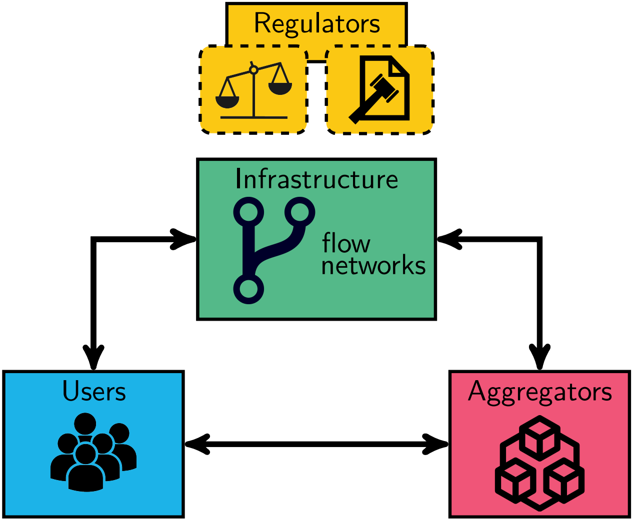

Breakin' the Law, Breakin' the Law,. . .

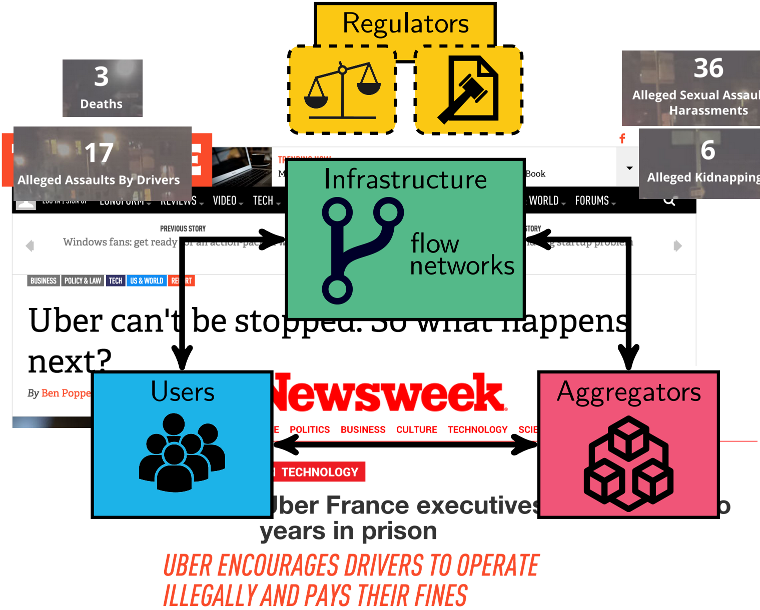

Technologically-Aware Regulation and Policies

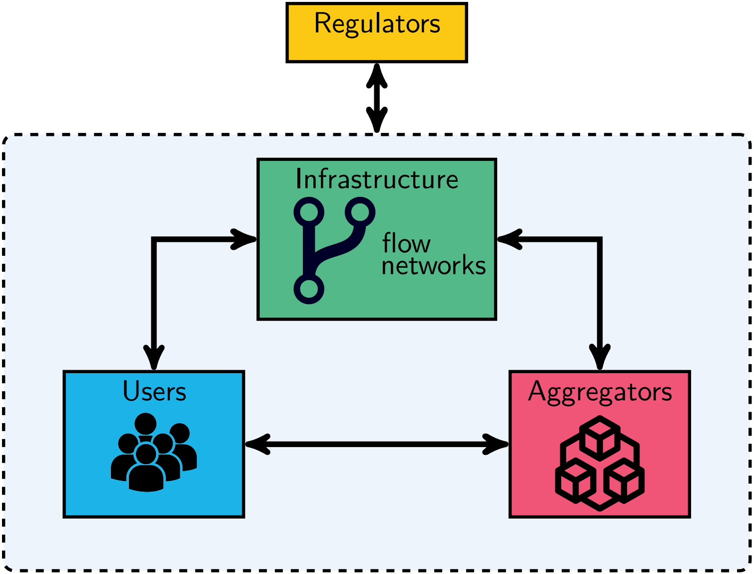

Infrastructure Evolution

The shared economy will require service providers to evolve in order to provide **improved services** that are **competitive in the new marketplace**

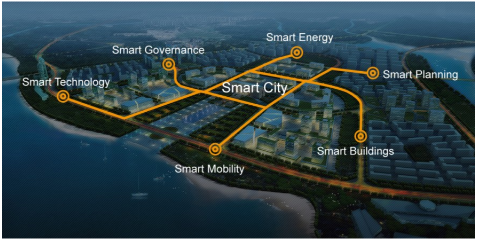

- Companies emerging that capitalizing on streaming data.
- Forcing existing infrastructure systems to modify their operational model in order to survive.

Not Just Existing Infrastructure: New infrastructure systems are emerging! (e.g. UAVs+UTM monitoring health of road, water, power networks)

Towards a Theory of RC + EI

Issues Addressed

- Incentive Design
- Disaggregation and Fundamental Privacy Bounds
- Privacy Aware Contract Design: Free Riding and Adverse Selection
- Value of Information in Urban Mobility
- Adaptive Incentive Mechanisms

Next Steps

- Modeling New Market Mechanisms
- Integrating DM models into RC
- Incentivize investments in security, privacy

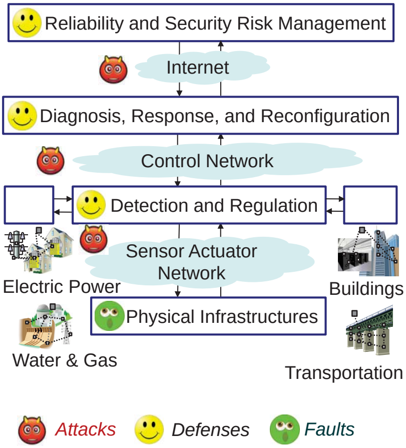

Thank you for your attention. Questions?

Shankar Sastry
sastry@coe.berkeley.edu