

New Vistas and Opportunities

Societal Scale CPS Systems: Data Markets and Incentive Design

Shankar Sastry

University of California, Berkeley

Joint work Lillian Ratliff, Roy Dong, Henrik Ohlsson (C3 Energy) and Alvaro Cardenas (UT Dallas))

NSF FORCES Review: 5th November 2015



Action Webs to Resilient CPS

- Action Webs & Societal CPS

Incentive Design for Societal CPS

- Incentive Design Framework

- Learning Utilities of Players

- Non-Intrusive Load Monitoring

Emerging Data Markets

- Privacy Metrics

- Privacy Aware Incentive Design

Conclusions and Future Research

Action Webs to Resilient CPS

Action Webs & Societal CPS

Incentive Design for Societal CPS

Incentive Design Framework

Learning Utilities of Players

Non-Intrusive Load Monitoring

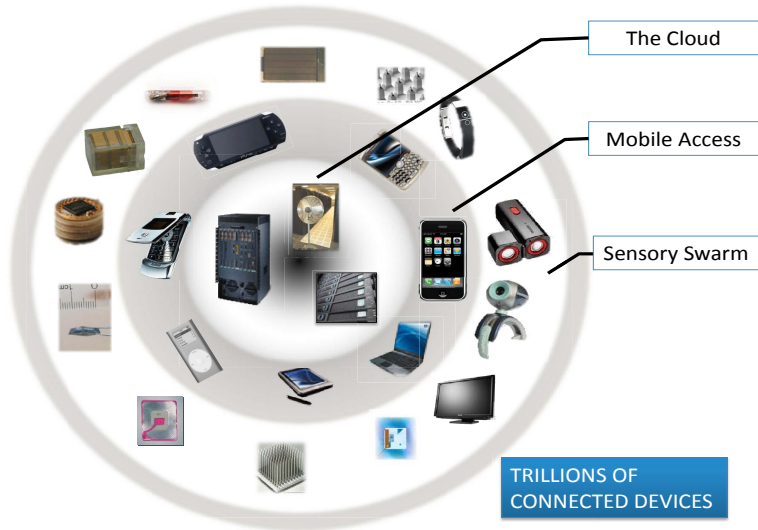
Emerging Data Markets

Privacy Metrics

Privacy Aware Incentive Design

Conclusions and Future Research

The swarm at the edge of the cloud




Source: J. Rabaey [ASPDAC'08]

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

Economic Impact

Electricity Grid:

- Utilities are currently utilizing smart meters for meter-to-cash. The potential of smart meters go far beyond this basic usage and the utilities are looking for a justification for their investments. The market for energy analytics in the smart grid is estimated to be worth 9.7 billion by 2020

Transportation Systems:

- It is estimated that more than 4.2 billion hours are wasted sitting in traffic, resulting in 2.8 billion gallons of wasted fuel and costing more than 87 billion dollars annually. By utilizing tools such as intelligent transportation systems (ITS) we can actively manage our transportation network to improve safety, efficiency, and multimodal connectivity.

Other Critical Infrastructures:

- Healthcare systems, Water systems, Natural gas and oil and other energy infrastructures

Action Webs to Resilient CPS

Action Webs & Societal CPS

Incentive Design for Societal CPS

Incentive Design Framework

Learning Utilities of Players

Non-Intrusive Load Monitoring

Emerging Data Markets

Privacy Metrics

Privacy Aware Incentive Design

Conclusions and Future Research

Why Incentive Design?

- There is often a substantive gap between competitive Nash equilibria and the social planner's optimum (Hal Varian, et al).
- Due to information asymmetries and misaligned objectives, the actions taken by agents in S-CPS are not socially optimal.
- Incentives are the natural mechanism for aligning agents so that they behave in a socially optimal way.

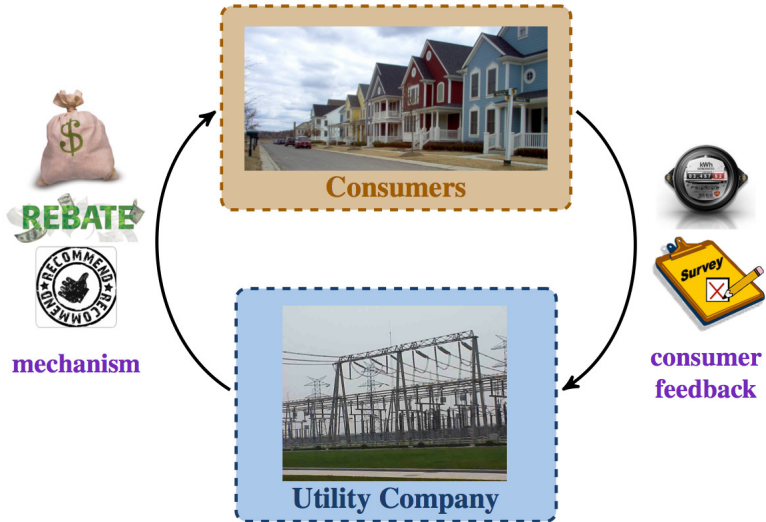
In Energy CPS:

- Consumers and Utilities are not well informed about their energy consumption patterns; incentives allow utility companies to motivate consumers to use less energy.

In Transportation CPS:


- Drivers often travel at peak hours; incentives can be used to encourage drivers to shift their departure time by only a few minutes for some reward resulting in overall reduced congestion.

Societal CPS for the Smart Grid




Incentives

- In the regulated energy market, utility companies are incentivized to reduce the overall consumption of their consumer base.
- Demand response programs incentivize customers to shift their demand thereby alleviating inaccuracies in load forecasting. Device-level incentives can be designed via non-intrusive load monitoring.
- Incentive Design needs a game theoretic model including data-driven models for agent behavior and their identification these during on-line operation.
- New Vulnerabilities: **Adversarial** agents who may **spoof** their energy signal, lie about their privacy needs or otherwise **disrupt** the energy system.

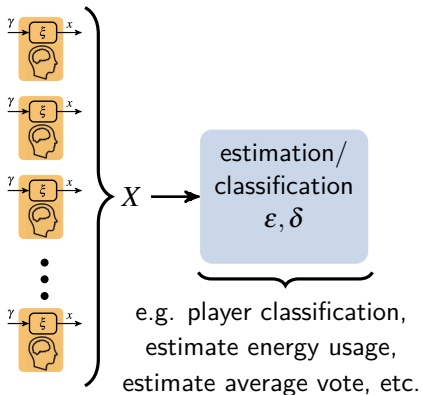


Learning with Multiple Players

Individual Agent
Decision-Making Model



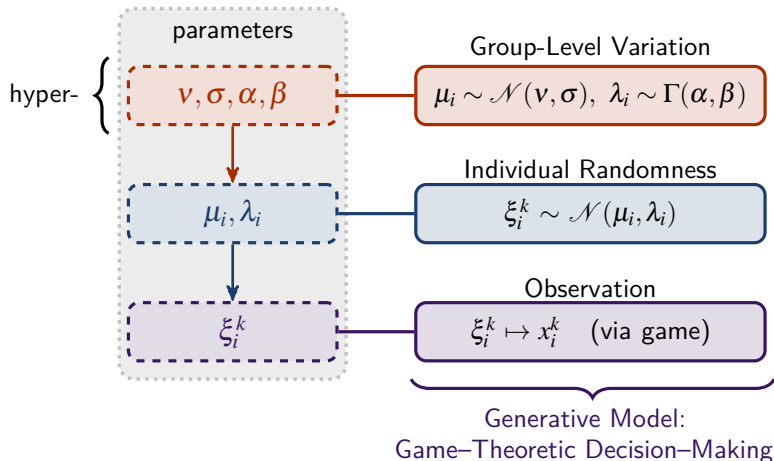
Many Decision Makers



Can we scale from the individual-level game-theoretic model of decision making to aggregate estimation / classification? task?

The Learning to Learn Framework

Learning to Learn: can improve estimation by simultaneously learning multiple similar tasks



Basic Statistical Bounds on Parameter Inference

Learning to Learn: we can provide bounds on parameter inference error.

Cramér–Rao bound for hyper-parameters $\theta = (\alpha, \beta, \nu, \sigma)$: for estimator $\hat{\theta}(X)$,

$$\underbrace{\mathbb{E}_X [(\theta_i - \hat{\theta}_i)^2]}_{\text{MSE of } \theta_i} \geq \frac{1}{n\zeta_i}, \text{ where } \zeta_i = \underbrace{-\mathbb{E}_X \left[\frac{\partial^2 \ln p(X|\theta)}{\partial \theta_i^2} \right]}_{\text{curvature}}$$

Remark: lower bound decreases by order $1/n$ (number of users)

Bayesian Cramér–Rao Bound for $\theta_r = (\mu_i, \lambda_i)$, for any estimator $\hat{\theta}_r$,

$$\mathbb{E}_{\xi^t, \theta} \left[(\hat{\theta}(\xi^t) - \theta) (\hat{\theta}(\xi^t) - \theta)^T \right] \geq \begin{bmatrix} \frac{(\alpha-1)\beta\sigma}{T\sigma - (\alpha-1)\beta} & 0 \\ 0 & \frac{2(\alpha-1)(\alpha-2)\beta^2}{T+2\alpha-2} \end{bmatrix}$$

Remark: decreases by order $1/T$ where T is number of samples

Hybrid Cramér–Rao bound applicable to the joint estimation of random and non-random parameters.

Reliable Estimation of Stopping Time Algorithm

REST: data-driven method based on concentration inequalities.

Consider for example a scenario with $n \gg 1$ occupants with lighting objectives and a planner with an energy saving objective $f : (\mathbf{x}^1, \dots, \mathbf{x}^t) \mapsto f(\mathbf{x}^1, \dots, \mathbf{x}^t) \in \mathbb{R}$ e.g.

- Estimate average lighting: $\frac{1}{t} \sum_{j=1}^t \left(\frac{1}{n} \sum_{i=1}^n [x^j]_i \right)$ Avg. vote at time j .
- Lighting energy: $\frac{1}{t} \sum_{j=1}^t g \left(\frac{1}{|S \setminus S_{\text{absent}}^j|} \sum_{i \notin S_{\text{absent}}^j} [x^j]_i \right)$ Avg. energy at time j
- Occupant classification error: $\frac{1}{n} \sum_{i=1}^n 1(h(\mathbf{x}_i^1, \dots, \mathbf{x}_i^t) \neq \mathbf{y}_i)$

The stopping time is the number of samples T needed before a specified error tolerance is reached. Standard statistical methods such as McDiarmid's Inequality and the Delta Method provide an algorithm to approximate the optimal stopping time: The Delta Method gives tighter asymptotic estimation, but the McDiarmid Inequality can be applied to more general problems, such as the classification of players

Classifying Multiple Players

Task: classify agents into different categories based on behavior.

- Incentive design: data-driven method to understand preferences
- Customer segmentation for energy management

User profiles: 132 players (33 replicates of each of the types below) in social game based on usage of shared resources in co-laboratory space

	Lighting Comfort	Incentive Award	Game Participation
Comforter	★ ★ ★	★	★ ★ ★ ★
Gamer	★	★ ★ ★	★ ★ ★ ★
Balancer	★ ★	★ ★	★ ★ ★ ★
Nonchalancer	★	★	★ ★

** each ★ indicates abstractly the amount the user types care about each of the categories

Classification of Player Category — Results

Random function of deviation from the best performance:


$$f(h; \mathbf{x}^t, \mathbf{y}) = L(h; \mathbf{x}^t, \mathbf{y}) - \inf_{h \in \mathcal{H}} \mathbb{E} [L(h; \mathbf{x}, \mathbf{y})];$$

where the loss function L is the proportion of misclassified users.

Bound on loss: $\Pr(L_t \geq \varepsilon + \kappa_t + \rho) \leq \Pr(f_t - \mathbb{E}[f_t] \geq \varepsilon) \leq 1 - \delta$


Simulation: 33 replicates of each type using occupant models generated out of data from real experiments.

Estimated stopping time



mean = 89 (blue), 78 (yellow), 94 (red)

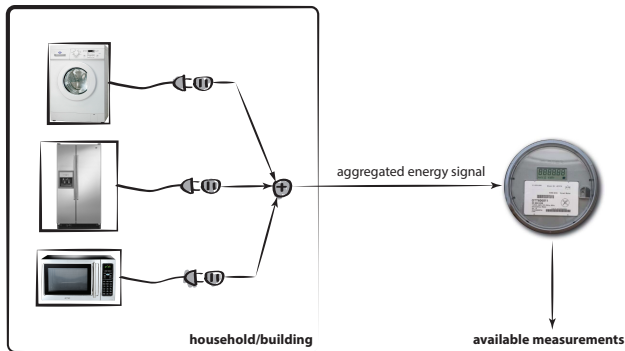
Misclassification error plateaus \approx day 80



REST captures the complexity of the problem






Energy Disaggregation

Energy disaggregation, or non-intrusive load monitoring (NILM):




The desired signal is the energy signals for each individual device.

Energy Disaggregation




Approaches to Disaggregation

Hidden Markov Models




- Unsupervised
- Requires tuning of parameters.
- The states are constant wattage levels; usage patterns and device signatures are encoded in transition probabilities.

Sparse Coding



- Supervised
- Assume inputs are sparse.
- Reconstruct the aggregate signal by selecting as few signatures as possible from a library.

Disaggregation Framework: Systems Approach

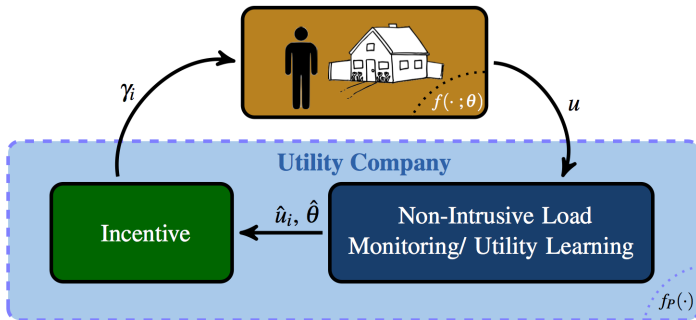


We learn *dynamical models* for the devices.

In our proposed framework:

- We have theoretical results guaranteeing recovery of the most likely device consumption signals.
- In our framework, we also learn dynamics of devices, which is useful for other Smart Grid operations.

Incentive Design Using Energy Disaggregation



- Given an upper bound on the probability of distinguishing devices, the utility company can design incentives that induce the consumer to use the desired amount of energy for a device with an error proportional to the probability of distinguishing devices.

Benefits of Energy Disaggregation

Beyond improved incentives, other benefits of disaggregated energy consumption data:

- Simply providing appliance-level power consumption information to energy consumers can achieve 20% energy savings in residential buildings, and sustain savings over the long-term.
- Disaggregation has the promise of additional reliability: fault detection, maintenance, repair, resilience, etc.


Gardner and Stern, The Short List: The Most Effective Actions U.S. Households Can Take to Curb Climate Change, 2008.

Laitner et. al, Examining the scale of the Behaviour Energy Efficiency Continuum, 2009.

Armel et. al, Is disaggregation the holy grail of energy efficiency? The case of electricity, 2013.

Hart, Nonintrusive appliance load monitoring, 1992.

Privacy Concerns with Disaggregation




Energy disaggregation^{1,2}: decomposing the whole building energy signal into device-level consumption.

Privacy Invasion?

- Inference of information about consumers: when you eat, watch TV, take a shower³.
- Such information is highly valuable and will be sought by many players: advertising⁴, law enforcement⁵, criminals⁶, power company (for designing incentives).

¹Dong, Ratliff, et. al., Allerton 2013 ²J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. Reynolds, S. Patel, IEEE CS, 2011 ³M. Livovich, D. Mulligan, S. Wicker, IEEE Security Privacy, 2010. ⁴R. Anderson, S. Fluoria, 9th Workshop on Economics of Information, 2010. ⁵G. Smith, *Marijuana bust shines light on utilities*, Post and Courier, 2012. ⁶Government Accountability Office, 2011.

Privacy Concerns with Disaggregation



Energy disaggregation^{1,2}: decomposing the whole building energy signal into device-level consumption.

Privacy Invasion?

- Inference of information about consumers: when you eat, watch TV, take a shower³.
- Such information is highly valuable and will be sought by many players: advertising⁴, law enforcement⁵, criminals⁶, power company (for designing incentives).

¹Dong, Ratliff, et. al., Allerton 2013 ²J. Froehlich, E. Larson, S. Gupta, G. Cohn, M. Reynolds, S. Patel, IEEE CS, 2011 ³M. Livovich, D. Mulligan, S. Wicker, IEEE Security Privacy, 2010. ⁴R. Anderson, S. Fluoria, 9th Workshop on Economics of Information, 2010. ⁵G. Smith, *Marijuana bust shines light on utilities*, Post and Courier, 2012. ⁶Government Accountability Office, 2011.

Action Webs to Resilient CPS

Action Webs & Societal CPS

Incentive Design for Societal CPS

Incentive Design Framework

Learning Utilities of Players

Non-Intrusive Load Monitoring


Emerging Data Markets

Privacy Metrics


Privacy Aware Incentive Design

Conclusions and Future Research


The Emerging Data Market



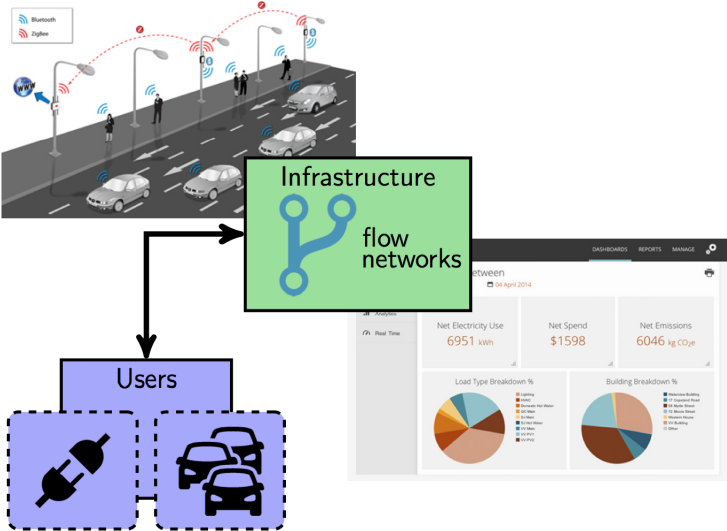
Smart, Connected Infrastructure



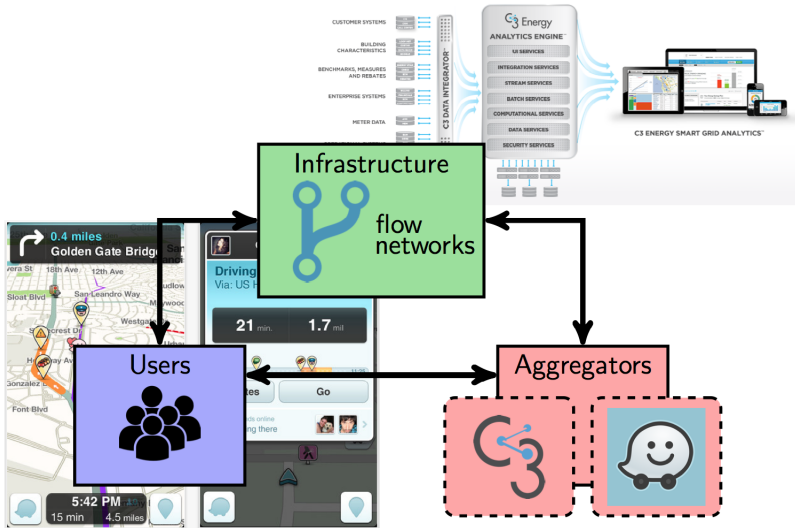
Operational Efficiency Informed by Usage Patterns



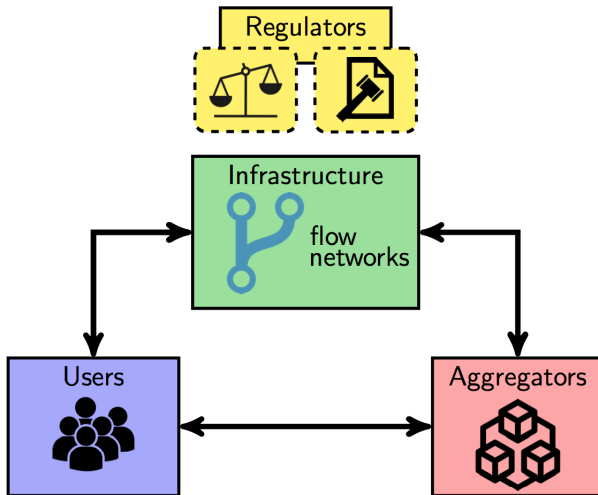
Closing the Loop — Integrating the End-User



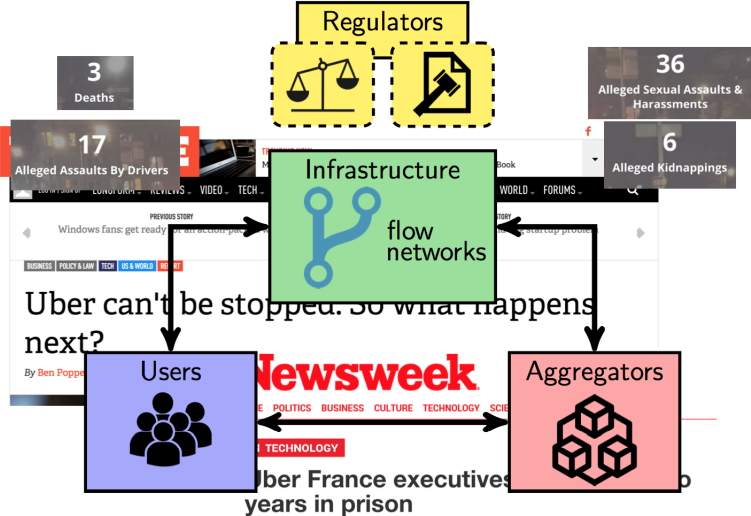
Emerging Data Market — Data as a Commodity



Old School Regulation




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



UBER ENCOURAGES DRIVERS TO OPERATE ILLEGALLY AND PAYS THEIR FINES

Technologically-Aware Regulation and Policy



Big CPS Data and Privacy

Data Minimization Principle (NIST Internal Report 7628)

Limit the collection of data to only that necessary for Smart Grid operations, including planning and management, improving energy use and efficiency, account management, and billing.

Non-actionable! Privacy is ontologically subjective

Inference Metrics:

- **Inferential Privacy**¹: Suppose an adversary uses an inference algorithm $u_a(y)$, then $p(u_a(y) = u^*) \leq p(u_{\text{MAP}}(y) = u^*)$.
- Information theoretic metrics²
- Differential Privacy³

¹R. Dong, L. Ratliff, *et al.*, HiCoNS 2014. ²Sankar, *et al.*, IEEE TSG 2012. ³Dwork, Microsoft Research Tech. Report, 2006.

Inferential Privacy Framework

- Consumers have a private parameter ξ , distributed according to some prior p_ξ . Assume ξ takes r different values, with $r < \infty$.
- This private parameter determines usage parameters: $u|\xi \sim p_{u|\xi}$.
- The usage parameters determine observable energy consumption: $y|u, \xi \sim p_{y|u}$.
- Let $p_{y|\xi}(y|\xi) = \int p_{y|u}(y|u)p_{u|\xi}(u|\xi)du$.
- The adversary knows the model p_ξ , $p_{u|\xi}$, and $p_{y|u}$ and observes y .

Inferential privacy: The model is α *inferentially private* if for any estimator $\hat{\xi} : y \mapsto \hat{\xi}(y)$, we have:

$$P(\hat{\xi}(y) \neq \xi) \geq \alpha$$

Maximum A-Posteriori Estimator

The Maximum A-Posteriori (MAP) estimator is given by:

$$\hat{\xi}_{MAP}(y) = \arg \max_{\hat{\xi}} p_{\xi}(\hat{\xi}) p_{y|\xi}(y|\hat{\xi})$$

Optimality of the MAP estimator: The model is α *inferentially private* with:

$$\alpha = P(\hat{\xi}_{MAP}(y) \neq \xi)$$

Furthermore, it is not α' inferentially private for any $\alpha' > \alpha$.

Often, it is not easy to calculate this value of α .

Inferential Privacy Approximations

Le Cam's method (total variation distance): The model is α inferentially private with:

$$\alpha = \max_{\xi_1 \neq \xi_2} \left[\min(p_\xi(\xi_1), p_\xi(\xi_2)) \cdot \frac{1}{2} \int |p_{y|\xi}(y|\xi_1) - p_{y|\xi}(y|\xi_2)| dy \right]$$


Fano's inequality (Kullback-Leibler divergence): The model is α inferentially private with:

$$\alpha = \frac{1}{\ln(r-1)} \left[\ln(r) - \frac{1}{r} \sum_{\xi_i, \xi_j} \int p_{y|\xi}(y|\xi_1) \ln \left(\frac{p_{y|\xi}(y|\xi_1)}{p_{y|\xi}(y|\xi_2)} \right) dy - \ln(2) \right]$$


These only require pairwise calculations!

Inferential Privacy Examples

Inferential Privacy Guarantee: α provides a lower bound on the probability an adversary makes an incorrect inference.




Determining if users are energy savers or wasters as a function of sampling duration.





Identifying income level from aggregate power consumption, as a function of sampling rate.

Quantifying the Societal Efficiency–Privacy Tradeoff



Efficiency-Privacy Tradeoff:

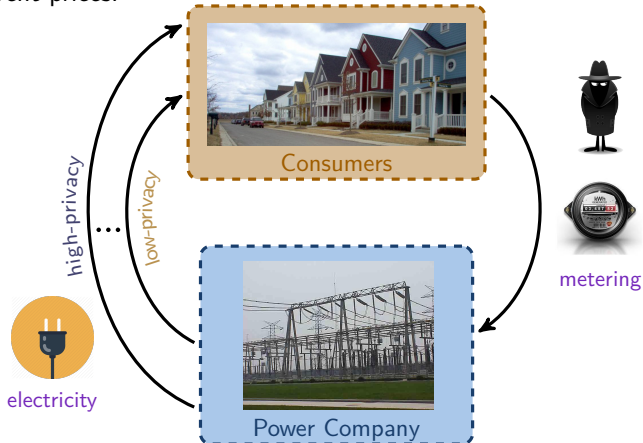
Direct Load Control (DLC) performance degrades as privacy-preserving metering is increased.



Privacy Settings on Smart Meters

- Utility desires high-fidelity data for smart grid operations.
- However, high-fidelity data increases risk of privacy breach (adversarial inference) and consumers want to protect their privacy.

Incentive Design Solution: Offer privacy-setting options on smart meter at different prices.



N -Privacy Settings Incentive Design Formulation

- Privacy settings on smart meters are viewed as a good.
- The consumer's **type** is θ and it characterizes the electricity consumption privacy needs of the consumer.
- The type $\theta \in \Theta = \{\theta_1, \dots, \theta_n\}$ where θ_1 is the lowest valuation of privacy and θ_n the highest with $\theta_i \geq \theta_j$ for $i > j$.
- The utility company faces a problem of **adverse selection** since the type of the consumer is unknown.
- $x : \Theta \rightarrow \mathbb{R}$ denotes the quality of the good. Since Θ is finite, we can define $x_i = x(\theta_i)$.
- The utility company designs a **menu of contracts**: $\{(t_i, x_i)\}_{i=1}^n$ in order to maximize its profit.

Individual Rationality and Incentive Compatibility

- Consumer's Utility: $U(x, \theta) - t$
- It is desirable that the consumer voluntarily participates and selects the contract designed for her type.
- **Individual Rationality** constraint ensures voluntary participation:

$$U(x, \theta) - t \geq U_o, \quad \text{where } U_o \text{ is the utility of the outside option}$$

- **Incentive Compatibility** constraints ensure the consumer reports her type truthfully, i.e. selects the contract designed for her:

$$U(x_i, \theta_i) - t_i \geq U(x_j, \theta_i) - t_j \quad \forall i, j$$

Utility Company's Optimization Problem

- **Prior on types:** $P(\theta = \theta_i) = p_i$, $\sum_{i=1}^n p_i = 1$, **Unit Cost to Utility:** $g(x)$

$$\text{Screening Problem: } \begin{cases} \max_{\{(t_i, x_i)\}_{i=1}^n} & \sum_{i=1}^n p_i (t_i - g(x_i)) \\ \text{s.t.} & U(x_i, \theta_i) - t_i \geq 0, \quad \forall i \quad (\text{IR}) \\ & U(x_i, \theta_i) - t_i \geq U(x_j, \theta_i) - t_j, \quad \forall i, j \quad (\text{IC}) \end{cases}$$

There are $n(n-1) + n$ constraints!

Utility Company's Optimization Problem

- **Prior on types:** $P(\theta = \theta_i) = p_i$, $\sum_{i=1}^n p_i = 1$, **Unit Cost to Utility:** $g(x)$

$$\text{Screening Problem: } \begin{cases} \max_{\{(t_i, x_i)\}_{i=1}^n} & \sum_{i=1}^n p_i (t_i - g(x_i)) \\ \text{s.t.} & U(x_i, \theta_i) - t_i \geq 0, \quad \forall i \quad (\text{IR}) \\ & U(x_i, \theta_i) - t_i \geq U(x_j, \theta_i) - t_j, \quad \forall i, j \quad (\text{IC}) \end{cases}$$

There are $n(n-1) + n$ constraints!

Utility Company's Optimization Problem

- **Prior on types:** $P(\theta = \theta_i) = p_i$, $\sum_{i=1}^n p_i = 1$, **Unit Cost to Utility:** $g(x)$

$$\text{Screening Problem: } \begin{cases} \max_{\{(t_i, x_i)\}_{i=1}^n} & \sum_{i=1}^n p_i (t_i - g(x_i)) \\ \text{s.t.} & U(x_i, \theta_i) - t_i \geq 0, \quad \forall i \quad (\text{IR}) \\ & U(x_i, \theta_i) - t_i \geq U(x_j, \theta_i) - t_j, \quad \forall i, j \quad (\text{IC}) \end{cases}$$

Assumption: (Spence-Mirrlees single-crossing condition)

Marginal gain from increasing the privacy setting x is greater for type θ_{i+1} than type θ_i , i.e. $U(x, \theta_{i+1}) - U(x, \theta_i)$ is increasing in x .

$$\text{Screening Problem: (Redux)} \quad \begin{cases} \max_{\{(t_i, x_i)\}_{i=1}^n} & \sum_{i=1}^n p_i (t_i - g(x_i)) \\ \text{s.t.} & U(x_1, \theta_1) - t_1 = 0, \quad (\text{IR-1}) \\ & U(x_i, \theta_i) - t_i \geq U(x_{i-1}, \theta_i) - t_{i-1}, \quad (\text{LDIC-i}) \\ & \forall i \in \{2, \dots, n\}, \\ & x_i \geq x_j \text{ for each } \theta_i \geq \theta_j \end{cases}$$

Solution to the above referred to as (Second-Best in the Contracts literature): $Y^{\text{sb}} = \{(t_i^{\text{sb}}, x_i^{\text{sb}})\}_{i=1}^n$

The Value of Information

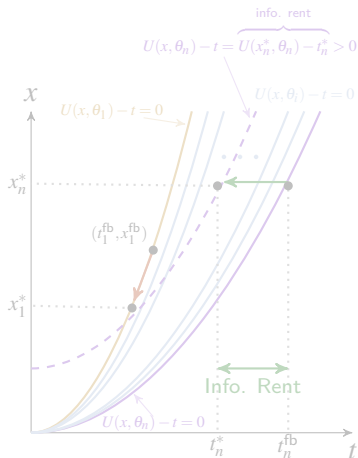
First-Best Solution:

$$\left[\max_{(x,t)} \{t - g(x) \mid U(x, \theta) - t \geq 0\} = \max \{U(x, \theta) - g(x)\} \right] \implies Y^{\text{fb}} = \{(t_i^{\text{fb}}, x_i^{\text{fb}})\}_{i=1}^n$$

Basic Insights:

- Type θ_n (highest valuation of privacy) gets the socially optimal privacy setting, $x_n^{\text{fb}} = x_n^{\text{sb}}$
- All other types get an inefficient allocation, $x_i^{\text{fb}} \geq x_i^{\text{sb}}$
- Type θ_1 (lowest valuation) gets **zero surplus**, $U(x_1^{\text{sb}}, \theta_1) = t_1^{\text{sb}}$
- All other types enjoy some **positive information rent** and hence, **free-ride** on society,

$$IR(\theta_i) = \sum_{j=1}^{i-1} (U(x_j^{\text{sb}}, \theta_{j+1}) - U(x_j^{\text{sb}}, \theta_j)) \geq 0$$



The Value of Information


First-Best Solution:

$$\left[\max_{(x,t)} \{t - g(x) \mid U(x, \theta) - t \geq 0\} = \max \{U(x, \theta) - g(x)\} \right] \implies Y^{\text{fb}} = \{(t_i^{\text{fb}}, x_i^{\text{fb}})\}_{i=1}^n$$

Basic Insights:

- Type θ_n (highest valuation of privacy) gets the socially optimal privacy setting, $x_n^{\text{fb}} = x_n^{\text{sb}}$
- All other types get an inefficient allocation, $x_i^{\text{fb}} \geq x_i^{\text{sb}}$
- Type θ_1 (lowest valuation) gets **zero surplus**, $U(x_1^{\text{sb}}, \theta_1) = t_1^{\text{sb}}$
- All other types enjoy some **positive information rent** and hence, **free-ride** on society,

$$IR(\theta_i) = \sum_{j=1}^{i-1} (U(x_j^{\text{sb}}, \theta_{j+1}) - U(x_j^{\text{sb}}, \theta_j)) \geq 0$$



Large Numbers of Consumers – Problem Formulation

Consider now that we have $N \gg 1$ consumers where consumer j 's types is $\theta^j \in \Theta^j = \{\theta_1^j, \dots, \theta_n^j\}$.

$$\left\{ \begin{array}{ll} \max_{\{Y^j\}_{j=1}^N} & \sum_{j=1}^N \sum_{i=1}^n p_i^j (t_i^j - g(x_i^j)) \\ \text{s.t.} & U(x_1^j, \theta_1^j) - t_1^j = 0 \quad (\text{IR-1,j}) \\ & U(x_i^j, \theta_i^j) - t_i^j = U(x_{i-1}^j, \theta_i^j) - t_{i-1}^j \quad (\text{LDIC-i,j}) \\ & x_i^j \geq x_k^j \text{ when } \theta_i^j \geq \theta_k^j \\ & \forall i \in \{2, \dots, n\}, \forall j \in \{1, \dots, N\} \end{array} \right.$$

where $Y^j = \{(t_i^j, x_i^j)\}_{i=1}^n$

Remark: Due to lack of network effects of information in this model, the problem decomposes into a single problem per consumer.

Future Work: Contracting over DLC (curtailment agreement) AND privacy setting.

Privacy Aware Contracts — Summary and Ongoing Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- Designed contracts in which electricity service is differentiated according to privacy to manage the efficiency-privacy tradeoff.
- Qualitative insights that show **high-type free-rides**, motivation to study **insurance-security investment**, and need for **regulation (subsidies)** to achieve social optimum.
- Similar analysis can be done on insurance contracts.

Open Questions:

- What is the impact of multiple consumers with varying levels of consumption and different privacy preferences on social welfare?
- Can we sequentially update the prior on preferences?
- What about the counterpart to privacy, namely, security/theft?

Privacy Aware Contracts — Summary and Ongoing Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- Designed contracts in which electricity service is differentiated according to privacy to manage the efficiency-privacy tradeoff.
- Qualitative insights that show **high-type free-rides**, motivation to study **insurance-security investment**, and need for **regulation (subsidies)** to achieve social optimum.
- Similar analysis can be done on insurance contracts.

Open Questions:

- What is the impact of multiple consumers with varying levels of consumption and different privacy preferences on social welfare?
- Can we sequentially update the prior on preferences?
- What about the counterpart to privacy, namely, security/theft?

Action Webs to Resilient CPS

Action Webs & Societal CPS

Incentive Design for Societal CPS

Incentive Design Framework

Learning Utilities of Players

Non-Intrusive Load Monitoring

Emerging Data Markets

Privacy Metrics

Privacy Aware Incentive Design

Conclusions and Future Research

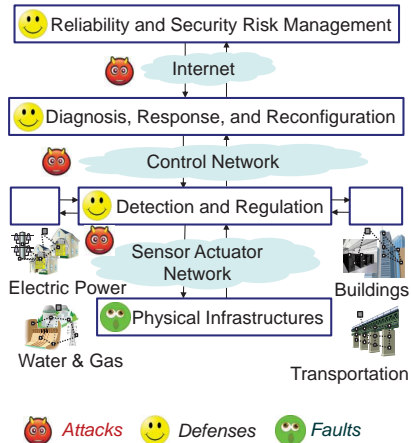
Towards a Theory of Resilient CPS

Issues Addressed

- Incentive Design
- Disaggregation and Fundamental Privacy Bounds
- Privacy Aware Contract Design: Free Riding and Adverse Selection
- Insurance against Loss of Privacy

Next Steps

- New Vulnerabilities, Attacks and Defenses
 - Financial zttacks
 - Increased reliability and preventive maintenance
 - Incentivize investments in security, privacy



Thank you for your attention. Questions?

Shankar Sastry
sastry@coe.berkeley.edu