FORCES
FOUNDATIONS OF RESILIENT CYBER-PHYSICAL SYSTEMS

Effects of Risk on Privacy Contracts for Demand–Side Management

Lillian J. Ratliff

joint with Roy Dong[1], Henrik Ohlsson[1], S. Shankar Sastry[1], Carlos Barreto[2], and Alvaro A. Cárdenas[2]

UC Berkeley[1], UT Dallas[2]

Berkeley UNIVERSITY OF CALIFORNIA Massachusetts Institute of Technology NSF MICHIGAN VANDERBILT UNIVERSITY

Is my energy data private?

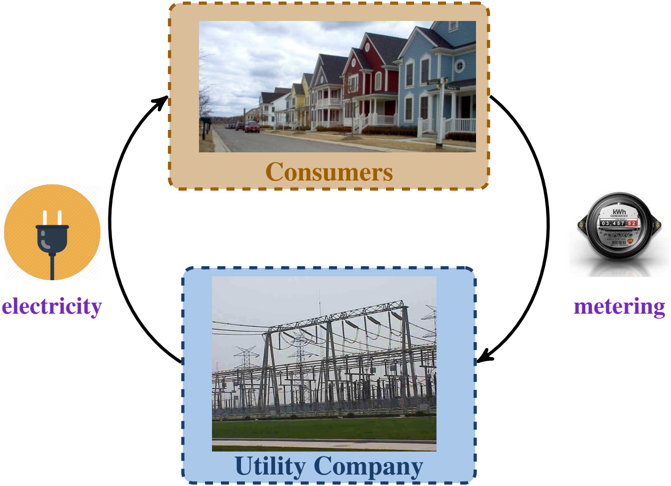

Is my energy data private?

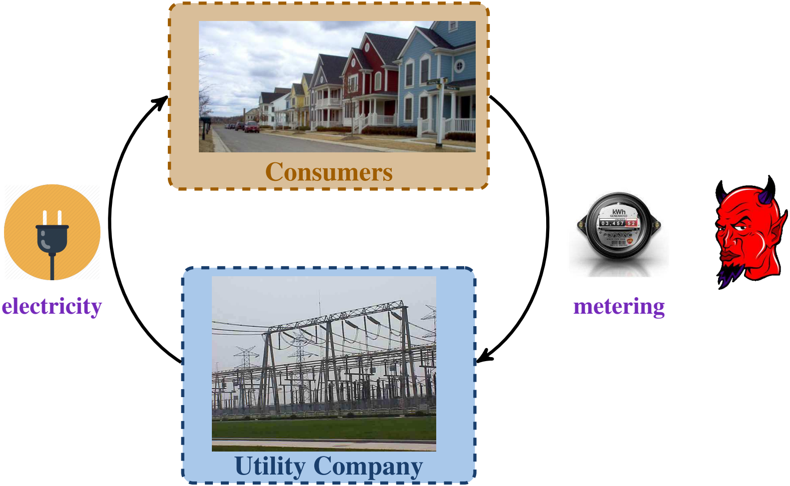

Is my energy data private?

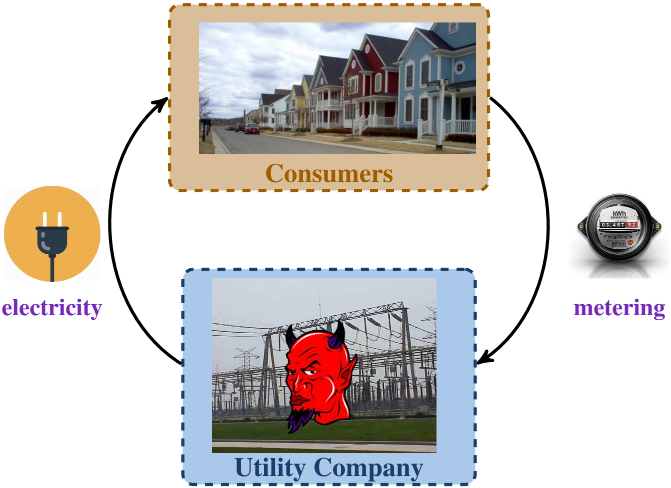

Utility–Privacy Tradeoff in Direct Load Control for Thermostatically Control Loads

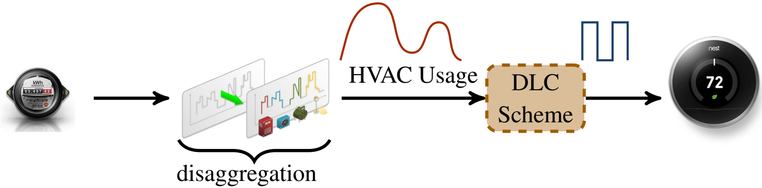

- Utility Company desires high–fidelity data for smart grid operations.
- Consumers want to protect their privacy.
- DLC performance degrades as privacy–preserving metering is increased.

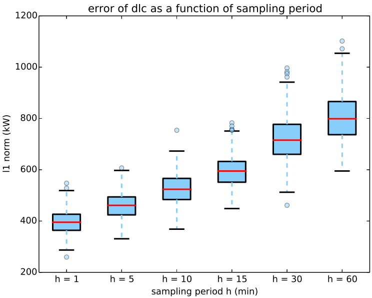

R. Dong, A. Cardenas, L. Ratliff, H. Ohlsson, S. Sastry. Quantifying the Utility–Privacy Tradeoff in the Smart Grid. Under Review in IEEE Trans. Smart Grid, 2014.

Outline

Two–Type Problem Formulation

- The utility company faces a problem of **adverse selection** since the type of the consumer is unknown.

Contract Design:
Utility company can design screening mechanisms to obtain the consumer's privacy preferences (unknown type) by offering contracts where service is differentiated according to privacy and consumers self–select based on their needs and wallet.

- Privacy settings on smart meters are viewed as a good.
- Quality of the good is either a high–privacy setting x_H or a low–privacy setting x_L where $\{x_H, x_L\} \subset \mathbb{R}$.
- The consumer's **type** is θ and it characterizes the electricity consumption privacy needs of the consumer.
- The type takes one of two values: $\theta \in \{\theta_H, \theta_L\}$ where $\theta_L < \theta_H$.
- The utility company is to design a pair of contracts: $\{(t_L, x_L), (t_H, x_H)\}$.

Ratliff, Dong, Ohlsson, Cárdenas, Sastry. Privacy and Customer Segmentation in the Smart Grid. IEEE CDC, 2014.

Individual Rationality and Incentive Compatibility

- The consumer's utility is equal to zero if he does not select a privacy setting (opt-out), and it is

$$U(x, \theta) - t \geq 0 \qquad \text{(\textbf{Individual Rationality})}$$

  if he selects the contract (t, x).
- Assumption: U is strictly increasing in (x, θ).
- All of the participants fare best when they truthfully reveal any private information asked for by the mechanism:

$$\left.\begin{array}{c} U(x_H, \theta_H) - t_H \geq U(x_L, \theta_H) - t_L \\ U(x_L, \theta_L) - t_L \geq U(x_H, \theta_L) - t_H \end{array}\right\} \qquad \text{(\textbf{Incentive-compatibility})}$$

Utility Company's Optimization Problem

- **Unit utility**: $v(x,t) = -g(x) + t$ where $g : x \mapsto g(x) \in \mathbb{R}$ is the unit cost and is assumed strictly increasing, convex, and differentiable.
- **Prior on types**: $p = P(\theta = \theta_H)$, $1 - p = P(\theta = \theta_L)$

Screening Problem:

$$\begin{cases} \max_{\{(t_L,x_L),(t_H,x_H)\}} & (1-p)v(x_L,t_L) + pv(x_H,t_H) & \\ \text{s.t.} & U(x_i,\theta_i) - t_i \geq 0,\ i = H,L & \text{(IR)} \\ & U(x_H,\theta_H) - t_H \geq U(x_L,\theta_H) - t_L & \text{(IC-1)} \\ & U(x_L,\theta_L) - t_L \geq U(x_H,\theta_L) - t_H & \text{(IC-2)} \end{cases}$$

Assume the marginal gain from increasing x is greater for type θ_H ($U(x,\theta_H) - U(x,\theta_L)$ is increasing in x). $\Rightarrow$

$$\begin{cases} t_H - t_L = U(x_H,\theta_H) - U(x_L,\theta_H) & \text{(IC-1')} \\ t_L = U(x_L,\theta_L) & \text{(IR')} \end{cases}$$

Reduced screening problem: second–best solution $\{(t_H^*, x_H^*), (t_L^*, x_L^*)\}$

$$\left. \begin{array}{l} \max_{x_H} \left\{ U(x_H,\theta_H) - g(x_H) \right\} \\ \max_{x_L} \left\{ -p(U(x_L,\theta_H) - U(x_L,\theta_L)) + (1-p)(U(x_L,\theta_L) - g(x_L)) \right\} \end{array} \right\} \quad \text{(P-1)}$$

Second–best due to information asymmetry which we will see benefits the high–type.

Utility Company's Optimization Problem

- **Unit utility**: $v(x,t) = -g(x) + t$ where $g : x \mapsto g(x) \in \mathbb{R}$ is the unit cost and is assumed strictly increasing, convex, and differentiable.
- **Prior on types**: $p = P(\theta = \theta_H)$, $1 - p = P(\theta = \theta_L)$

Screening Problem:

$$\begin{cases} \max\limits_{\{(t_L,x_L),(t_H,x_H)\}} & (1-p)v(x_L,t_L) + pv(x_H,t_H) \\ \text{s.t.} & U(x_i,\theta_i) - t_i \geq 0,\ i = H,L \qquad \text{(IR)} \\ & U(x_H,\theta_H) - t_H \geq U(x_L,\theta_H) - t_L \qquad \text{(IC-1)} \\ & U(x_L,\theta_L) - t_L \geq U(x_H,\theta_L) - t_H \qquad \text{(IC-2)} \end{cases}$$

Assume the marginal gain from increasing x is greater for type θ_H ($U(x,\theta_H) - U(x,\theta_L)$ is increasing in x). $\Rightarrow$

$$\begin{cases} t_H - t_L = U(x_H,\theta_H) - U(x_L,\theta_H) & \text{(IC-1')} \\ t_L = U(x_L,\theta_L) & \text{(IR')} \end{cases}$$

Reduced screening problem: second–best solution $\{(t_H^*, x_H^*), (t_L^*, x_L^*)\}$

$$\left.\begin{array}{l} \max_{x_H} \left\{U(x_H,\theta_H) - g(x_H)\right\} \\ \max_{x_L} \left\{-p(U(x_L,\theta_H) - U(x_L,\theta_L)) + (1-p)(U(x_L,\theta_L) - g(x_L))\right\} \end{array}\right\} \qquad \text{(P-1)}$$

Second–best due to information asymmetry which we will see benefits the high–type.

Characterization of Contract

First–best solution $\{(t_H^{fb}, x_H^{fb}), (t_L^{fb}, x_L^{fb})\}$: Utility company has **full information**, i.e. knows the type of the agent he is facing.

$$\max_{x,t} \left\{ -g(x) + t \;\middle|\; U(x,\theta) - t \geq 0 \right\} \Longrightarrow \max_{x,t} \left\{ -g(x) + U(x,\theta) \right\}$$

First–best (t_i^{fb}, x_i^{fb}) (full information) vs. second–best (t_i^*, x_i^*) (asymmetric information):

- The high–type always gets an **efficient allocation**: $x_H^* = x_H^{fb}$
- The high–type gets **positive information rent** (Utility company pays rent to θ_H)

$$t_H^* = t_H^{fb} - \underbrace{\left(U(x_L^*, \theta_H) - U(x_L^*, \theta_L)\right)}_{\text{information rent}}$$

- The low-type gets **zero surplus** since $t_L^* = U(x_L^*, \theta_L)$ and an **inefficient allocation** $x_L^* \leq x_L^{fb}$

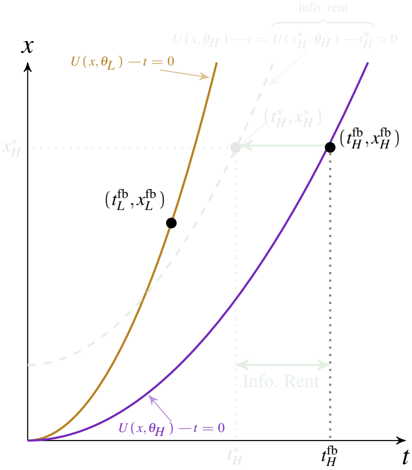

Characterization of Contract

First–best solution $\{(t_H^{fb}, x_H^{fb}), (t_L^{fb}, x_L^{fb})\}$: Utility company has **full information**, i.e. knows the type of the agent he is facing.

$$\max_{x,t}\left\{-g(x)+t \,\middle|\, U(x,\theta)-t\geq 0\right\} \Longrightarrow \max_{x,t}\left\{-g(x)+U(x,\theta)\right\}$$

First–best (t_i^{fb}, x_i^{fb}) (full information) vs. second–best (t_i^*, x_i^*) (asymmetric information):

- The high–type always gets an **efficient allocation**: $x_H^* = x_H^{fb}$
- The high–type gets **positive information rent** (Utility company pays rent to θ_H)

$$t_H^* = t_H^{fb} - \underbrace{\left(U(x_L^*,\theta_H) - U(x_L^*,\theta_L)\right)}_{\textit{information rent}}$$

- The low–type gets **zero surplus** since $t_L^* = U(x_L^*, \theta_L)$ and an **inefficient allocation** $x_L^* \leq x_L^{fb}$

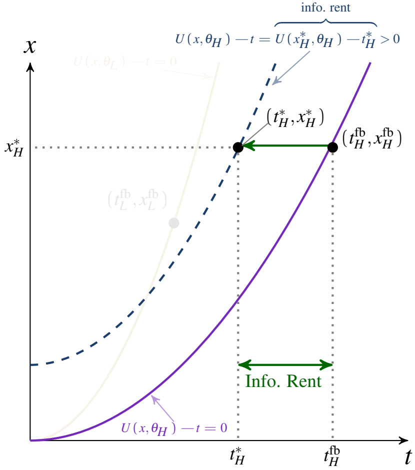

Characterization of Contract

First–best solution $\{(t_H^{fb}, x_H^{fb}), (t_L^{fb}, x_L^{fb})\}$: Utility company has **full information**, i.e. knows the type of the agent he is facing.

$$\max_{x,t}\left\{-g(x)+t \;\middle|\; U(x,\theta)-t \geq 0\right\} \Longrightarrow \max_{x,t}\left\{-g(x)+U(x,\theta)\right\}$$

First–best (t_i^{fb}, x_i^{fb}) (full information) vs. second–best (t_i^*, x_i^*) (asymmetric information):

- The high–type always gets an **efficient allocation**: $x_H^* = x_H^{fb}$
- The high–type gets **positive information rent** (Utility company pays rent to θ_H)

$$t_H^* = t_H^{fb} - \underbrace{\left(U(x_L^*, \theta_H) - U(x_L^*, \theta_L)\right)}_{\text{information rent}}$$

- The low-type gets **zero surplus** since $t_L^* = U(x_L^*, \theta_L)$ and an **inefficient allocation** $x_L^* \leq x_L^{fb}$

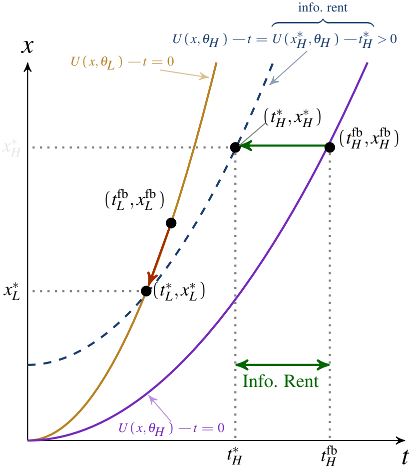

Who bears the risk of privacy loss?

- Given the probability of privacy breach as a function of privacy setting and the associated value of the loss of privacy as a function of type, how does the optimal contract change?
- What does this mean for **security** and **insurance** investment?

Who s the risk of privacy loss?

Consumer's Utility : $\overbrace{U(x,\theta)}^{\text{no risk of privacy loss}}$ $\overbrace{\mathbf{U}(x,\theta) = U(x,\theta) - \underbrace{(1-\eta(x))}_{\textbf{privacy breach probability}}\underbrace{\ell(\theta)}_{\textbf{loss}}}^{\text{with risk of privacy loss}}$

- **Individual Rationality: $\mathbf{U}(x,\theta) - t \geq 0$**

$$\mathbf{U}(x_L^*,\theta_L) - t_L^* = U(x_L^*,\theta_L) - t_L^* - (1-\eta(x_L^*))\ell(\theta_L) = -(1-\eta(x_L^*))\ell(\theta_L) \leq 0$$

- **low–type might opt–out**
- **Incentive Compatibility: $\mathbf{U}(x_L,\theta_L) - t_L - \mathbf{U}(x_H,\theta_L) + t_H \geq 0$**

$$\underbrace{U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^*}_{\geq 0} - \underbrace{(\eta(x_H^*) - \eta(x_L^*))}_{\geq 0}\underbrace{\ell(\theta_L)}_{\geq 0}$$

- $(\eta(x_H^*) - \eta(x_L^*))\ell(\theta_L) \geq U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^* \implies$ **the low-type might choose (t_H^*, x_H^*) and thus does not report truthfully**

Who s the risk of privacy loss?

Consumer's Utility : $\overbrace{U(x,\theta)}^{\text{no risk of privacy loss}}$ $\overbrace{\mathbf{U}(x,\theta) = U(x,\theta) - \underbrace{(1-\eta(x))}_{\text{privacy breach probability}}\underbrace{\ell(\theta)}_{\text{loss}}}^{\text{with risk of privacy loss}}$

- **Individual Rationality:** $\mathbf{U}(x,\theta) - t \geq 0$

$$\mathbf{U}(x_L^*,\theta_L) - t_L^* = U(x_L^*,\theta_L) - t_L^* - (1-\eta(x_L^*))\ell(\theta_L) = -(1-\eta(x_L^*))\ell(\theta_L) \leq 0$$

- **low–type might opt–out**
- **Incentive Compatibility:** $\mathbf{U}(x_L,\theta_L) - t_L - \mathbf{U}(x_H,\theta_L) + t_H \geq 0$

$$\underbrace{U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^*}_{\geq 0} - \underbrace{(\eta(x_H^*) - \eta(x_L^*))}_{\geq 0}\underbrace{\ell(\theta_L)}_{\geq 0}$$

- $(\eta(x_H^*) - \eta(x_L^*))\ell(\theta_L) \geq U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^* \implies$ the low–type might choose (t_H^*, x_H^*) and thus does not report truthfully

Who s the risk of privacy loss?

Consumer's Utility : $\underbrace{U(x,\theta)}_{\text{no risk of privacy loss}}$ $\qquad$ $\overbrace{\mathbf{U}(x,\theta) = U(x,\theta) - \underbrace{(1-\eta(x))}_{\text{privacy breach probability}}\underbrace{\ell(\theta)}_{\text{loss}}}^{\text{with risk of privacy loss}}$

- **Individual Rationality**: $\mathbf{U}(x,\theta) - t \geq 0$

$$\mathbf{U}(x_L^*,\theta_L) - t_L^* = U(x_L^*,\theta_L) - t_L^* - (1-\eta(x_L^*))\ell(\theta_L) = -(1-\eta(x_L^*))\ell(\theta_L) \leq 0$$

- **low–type might opt–out**
- **Incentive Compatibility**: $\mathbf{U}(x_L,\theta_L) - t_L - \mathbf{U}(x_H,\theta_L) + t_H \geq 0$

$$\underbrace{U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^*}_{\geq 0} - \underbrace{(\eta(x_H^*) - \eta(x_L^*))}_{\geq 0}\underbrace{\ell(\theta_L)}_{\geq 0}$$

- $(\eta(x_H^*) - \eta(x_L^*))\ell(\theta_L) \geq U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^* \implies$ the low–type might choose (t_H^*, x_H^*) and thus **does not report truthfully**

Who s the risk of privacy loss?

Consumer's Utility : $\overbrace{U(x,\theta)}^{\text{no risk of privacy loss}}$ $\qquad \overbrace{\mathbf{U}(x,\theta) = U(x,\theta) - \underbrace{(1-\eta(x))}_{\textbf{privacy breach probability}}\underbrace{\ell(\theta)}_{\textbf{loss}}}^{\text{with risk of privacy loss}}$

- **Individual Rationality**: $\mathbf{U}(x,\theta) - t \geq 0$

$$\mathbf{U}(x_L^*,\theta_L) - t_L^* = U(x_L^*,\theta_L) - t_L^* - (1-\eta(x_L^*))\ell(\theta_L) = -(1-\eta(x_L^*))\ell(\theta_L) \leq 0$$

- **low–type might opt–out**
- **Incentive Compatibility**: $\mathbf{U}(x_L,\theta_L) - t_L - \mathbf{U}(x_H,\theta_L) + t_H \geq 0$

$$\underbrace{U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^*}_{\geq 0} - \underbrace{(\eta(x_H^*) - \eta(x_L^*))}_{\geq 0}\underbrace{\ell(\theta_L)}_{\geq 0}$$

- $(\eta(x_H^*) - \eta(x_L^*))\ell(\theta_L) \geq U(x_L^*,\theta_L) - t_L^* - U(x_H^*,\theta_L) + t_H^* \implies$ the low–type might choose (t_H^*, x_H^*) and thus **does not report truthfully**

Effects of Risk on Privacy Contracts: $\underbrace{\{(t_i^*, x_i^*)\}_{i\in\{H,L\}}}_{\text{with privacy loss risk}}$ and $\underbrace{\{(t_i^*, x_i^*)\}_{i\in\{H,L\}}}_{\text{without privacy loss risk}}$

Proposition

- *Independent of* p, $x_H^* \geq x_H^*$.
- *The privacy setting* x_L^* *(resp.* x_L^**) is* ***decreasing*** *w.r.t.* p. *Thus,* t_L^* *is also decreasing.*
- *The privacy setting for type* θ_L *is further characterized by the following:*

$$\begin{cases} x_L^* \geq x_L^*, & \text{if } p \leq \frac{\ell(\theta_L)}{\ell(\theta_H)}, \\ x_L^* < x_L^*, & \text{if } p > \frac{\ell(\theta_L)}{\ell(\theta_H)}. \end{cases}$$

- *If* $p > \frac{\ell(\theta_L)}{\ell(\theta_H)}$, *then* $t_L^* < t_L^*$, $t_H^* < t_H^*$
- *If* $p > \frac{\ell(\theta_L)}{\ell(\theta_H)}$, *the information rent is higher without risk:*
  $U(x_L^*, \theta_H) - U(x_L^*, \theta_L) > \mathbf{U}(x_L^*, \theta_H) - \mathbf{U}(x_L^*, \theta_L)$

To promote participation, $\downarrow t_L$ and/or $\uparrow x_L \implies$ decrease in benefit and fees collected.

Hence, there is an incentive for the utility company to purchase insurance and/or invest in security.

DLC Example — Optimal Contracts

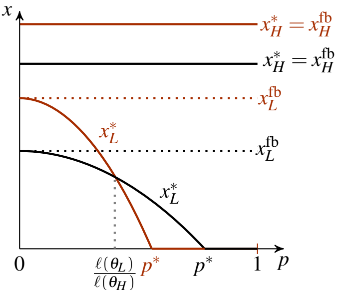

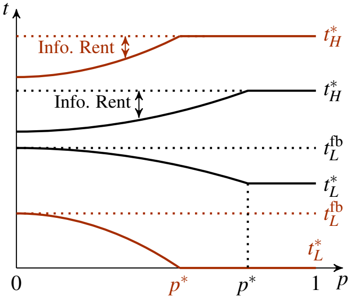

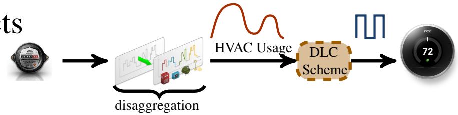

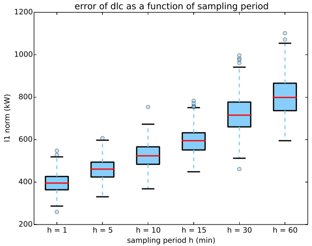

- Utility company's unit cost: $g(x) = \frac{1}{2}\zeta x^2$, $0 < \zeta < \infty$.
- Consumer's utility function: $U(x, \theta) = x\theta$, $\mathbf{U}(x, \theta) = x\theta - (1 - \eta(x))\ell(\theta)$
- Let $1 - \eta(x) = m(1 - x)$, $m > 0$ and $x \in [0, 1]$
- Critical values: $p^* = \frac{\theta_L + m\ell(\theta_L)}{\theta_H + m\ell(\theta_H)}$, $p^* = \frac{\theta_L}{\theta_H}$

DLC Example — Profit and Social Welfare

- Utility company's profit: Π^* (w/ risk), Π^*, w/o risk
  $\Pi(t_L, x_L, t_H, x_H) = (1-p)(-g(x_L)+t_L) + p(-g(x_H)+t_H)$
- Social Welfare: W^* (w/ risk), W^*, w/o risk
  $W(p, t_L, x_L, t_H, x_H) = \Pi(t_L, x_L, t_H, x_H) + p(U(x_H, \theta_H) - t_H) + (1-p)(U(x_L, \theta_L) - t_L)$

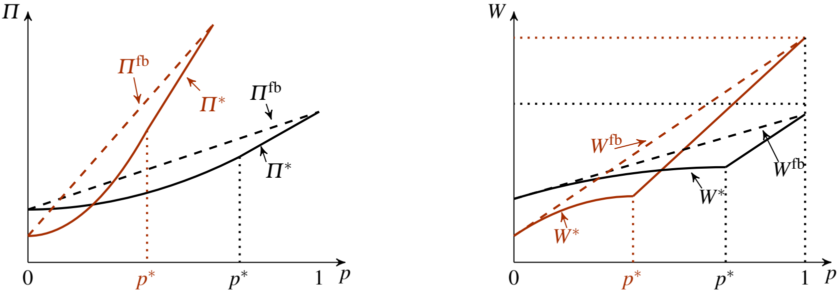

- There are values of p for which no one does well when there is risk; both the social welfare and the utility company's profit are lower

L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, S. Sastry. Effects of Risk on Privacy Contracts for Demand–Side Management. Under Review in IEEE Trans. Smart Grid, 2014.

Summary and Future Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- We modeled electricity service as a product line differentiated according to privacy and we found the following.
  - Privacy loss risks **decrease** the level service offered to each consumer type.
  - We remark that people who value high privacy more, need to be compensated more to participate in the smart grid.
  - The utility company has an **incentive** to purchase **insurance** and invest in **security** when there are loss risks.
- Using knowledge of consumer preferences, the utility company can incentivize consumers to choose a low privacy setting. We are investigating **dynamic contracts** in which the utility estimates the distribution of the population at each step.
- We are currently investigating the **security–insurance** investment tradeoff in the presence of privacy loss risks as well as the design of **insurance contracts** for utility companies given a compensation policy for consumers.

L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, S. Sastry. Effects of Risk on Privacy Contracts for Demand–Side Management. Under Review in IEEE Trans. Smart Grid, 2014.

Summary and Future Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- We modeled electricity service as a product line differentiated according to privacy and we found the following.
  - Privacy loss risks **decrease** the level service offered to each consumer type.
  - We remark that people who value high privacy more, need to be compensated more to participate in the smart grid.
  - The utility company has an **incentive** to purchase **insurance** and invest in **security** when there are loss risks.
- Using knowledge of consumer preferences, the utility company can incentivize consumers to choose a low privacy setting. We are investigating **dynamic contracts** in which the utility estimates the distribution of the population at each step.
- We are currently investigating the **security–insurance** investment tradeoff in the presence of privacy loss risks as well as the design of **insurance contracts** for utility companies given a compensation policy for consumers.

L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, S. Sastry. Effects of Risk on Privacy Contracts for Demand–Side Management. Under Review in IEEE Trans. Smart Grid, 2014.

Summary and Future Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- We modeled electricity service as a product line differentiated according to privacy and we found the following.
  - Privacy loss risks **decrease** the level service offered to each consumer type.
  - We remark that people who value high privacy more, need to be compensated more to participate in the smart grid.
  - The utility company has an **incentive** to purchase **insurance** and invest in **security** when there are loss risks.
- Using knowledge of consumer preferences, the utility company can incentivize consumers to choose a low privacy setting. We are investigating **dynamic contracts** in which the utility estimates the distribution of the population at each step.
- We are currently investigating the **security–insurance** investment tradeoff in the presence of privacy loss risks as well as the design of **insurance contracts** for utility companies given a compensation policy for consumers.

L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, S. Sastry. Effects of Risk on Privacy Contracts for Demand–Side Management. Under Review in IEEE Trans. Smart Grid, 2014.

Summary and Future Work

- Implementing **privacy-aware** data collection policies results in a **reduction in the efficiency** of grid operations.
- We modeled electricity service as a product line differentiated according to privacy and we found the following.
  - Privacy loss risks **decrease** the level service offered to each consumer type.
  - We remark that people who value high privacy more, need to be compensated more to participate in the smart grid.
  - The utility company has an **incentive** to purchase **insurance** and invest in **security** when there are loss risks.
- Using knowledge of consumer preferences, the utility company can incentivize consumers to choose a low privacy setting. We are investigating **dynamic contracts** in which the utility estimates the distribution of the population at each step.
- We are currently investigating the **security–insurance** investment tradeoff in the presence of privacy loss risks as well as the design of **insurance contracts** for utility companies given a compensation policy for consumers.

L. Ratliff, C. Barreto, R. Dong, H. Ohlsson, A. Cardenas, S. Sastry. Effects of Risk on Privacy Contracts for Demand–Side Management. Under Review in IEEE Trans. Smart Grid, 2014.