Multiple-Profile Prediction-of-Use Games

Andrew Perrault perrault@cs.toronto.edu
Craig Boutilier cebly@cs.toronto.edu

Contributions

- Extend prediction-of-use (POU) games to support multiple profiles
  - Extension remains convex
  - Creates new enforcement problems addressed by separating functions
- Use learned utility models to experimentally validate our approach

Motivation

Consumer pays per kWh used, a *fixed-rate tariff*

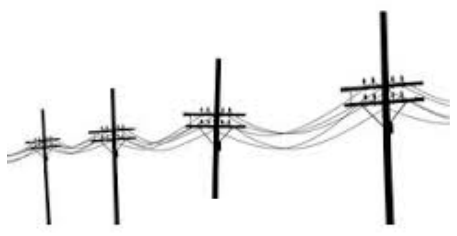

Supplier buys electricity in advance, but can also buy at the last minute for a higher price

Generator

Misalignment of incentives: Consumer's cost does not depend on predictability, but supplier's cost does

Prediction-of-Use Tariffs

- Each consumer makes a prediction ahead of time
  - They are charged based on both consumption amount and prediction accuracy
- Consumers can form groups and be treated as one large agent
  - But they can only do this if they can agree on how to split the costs

Intro to Cooperative Games

- Set of agents N which can form *coalitions*
- Characteristic value function $v: 2^N \to \mathbb{R}$ represents value that coalition can achieve
- Agents can defect away from coalition, but not from their action (contracts)
- **Definition (superadditivity).** $v(S+T) \geq v(S) + v(T)$ for all disjoint S and T
- **Theorem.** In a superadditive game, the *grand coalition* of all agents has highest social welfare.
- Most cooperative games are superadditive

Cost Sharing in Cooperative Games

- Challenge: how to divide benefits of cooperation
- $t(i)$ denotes the payment to agent i
- **Definition (stability).** No agent should have an incentive to defect to another coalition. A strong statement of stability: $\sum_{i \in S} t(i) \geq v(S), \forall S \subset N$
- **Definition (efficiency/budget-balance).** The entire value should be distributed:
$$\sum_{i \in S} t(i) = v(N)$$
- **Definition (core allocation).** Satisfies stability and efficiency
  - Core allocations are satisfying, but may not exist and are hard to compute
- *Shapley value*: reasonable or "fair" distribution, ignore competition
- **Definition (Shapley value of agent i).** Average contribution to coalition value over all *join orders*:
$$t_{Shapley}(i) = \sum_{S \subseteq C \setminus \{i\}} \frac{|S|!\,(|C| - |S| - 1)!}{|N|!} (v(S \cup \{i\}) - v(S))$$
- Shapley value always exists and is easier to compute
- **Definition (convexity cooperative game).** Characteristic function is supermodular: $v(S \cup T) + v(S \cap T) \geq v(S) + v(T), \forall S, T \subseteq N$
  - Alternatively, the value added by joining a coalition grows as the coalition grows: $v(S \cup \{i\}) - v(S) \leq v(T \cup \{i\}) - v(T), \forall S \subseteq T \subseteq N \setminus \{i\}, \forall i \in N$
- In a convex game, the Shapley value is a core allocation 😀

Prediction-of-Use (POU) Games

- Cooperative game where each agent is a household
- Each household has a distribution over future consumption—a *profile*
  - In Robu et al. model, distributions are assumed to be independent normal random variables
- Coalition's profile is sum of its members' profiles—also a normal random variable

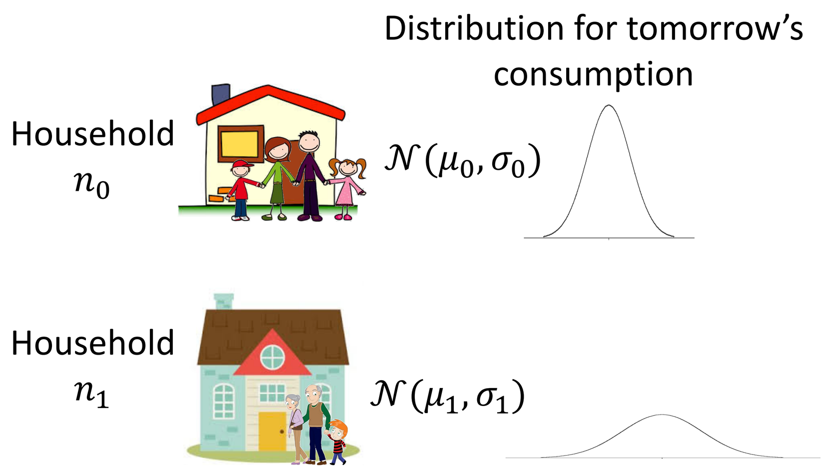

- Each coalition predicts a baseline b, and pays at realization time according to both how much it consumed and how close its prediction was
  - *POU tariff* $< p, p^+, p^- >$. Pay p for each unit consumed, p^+ for each unit above baseline and p^- for each unit less than the baseline
  - Robu et al. provide a closed form for optimal b: $b^* = \mu_i + \sigma_i \Phi^{-1}\left(\frac{\overline{p}}{\overline{p} + \underline{p}}\right)$
- Characteristic function is total cost in expectation: $v(C) = -\mu(C)p - \sigma(C)L\left(\underline{p}, \overline{p}\right)$ where $\mu(C)$ and $\sigma(C)$ are the sum of the mean and standard deviation for C
- **Theorem (Robu et al., 2017):** POU games are convex.
- Limitation: the only decision an agent makes in POU games is what profile to report
  - Even if truthful, agents have utility functions, different profiles
  - Best choice depends on the profiles other agents choose
  - This is itself a game, but it is not part of the POU model

Multiple-Profile Prediction-of-Use (MPOU) Games

- Households have multiple profiles
- Each profile has a value
- Characteristic value of a coalition: sum of profile values minus expected costs, under best possible assignment
- **Theorem:** MPOU games are convex.
- Complication: having multiple profiles interferes with contract enforcement
- Coalition assigns profile to agent
- Actions are only partially observable
  - Profile selected only known to agent
  - Coalition observes *realized* consumption

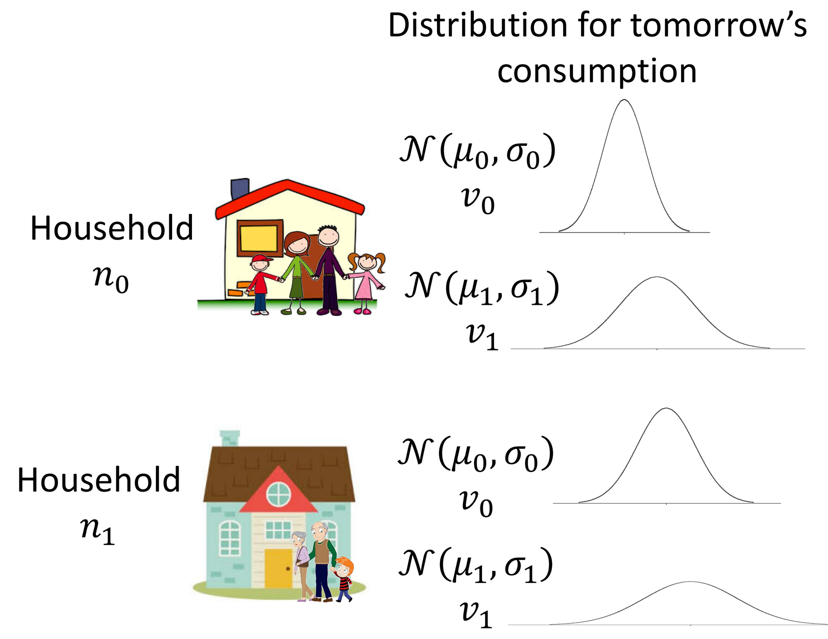

Separating Functions (SFs)

- A separating function maps realized consumption to a payment
  - From coalition to agent
  - To incentivize use of the assigned profile
- $D(x)$ is a separating function under assignment A of agents to profiles if
$$\mathbb{E}_{A(i)}(D(x)) + v(A(i)) > \mathbb{E}_{\overline{A}(i)}(D(x)) + v\left(\overline{A}(i)\right) \quad (incentive)$$
$$\mathbb{E}_{A(i)}(D(x)) = 0 \quad (zero\text{-}expectation)$$
- Incentive condition makes agent use assigned profile
- Zero-expectation condition means expected payments not affected
- **Theorem:** separating function for two profiles: PDF(assigned profile) – PDF(other profile).
- No known closed-form for arbitrary # of profiles

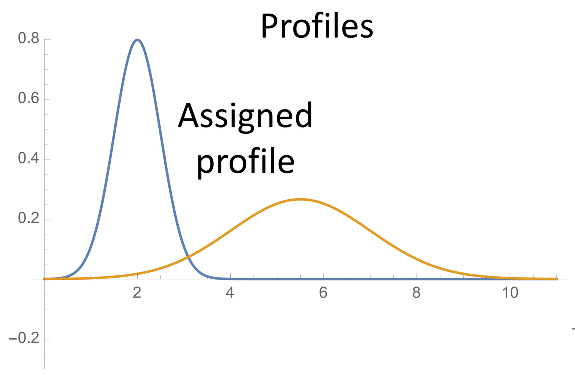

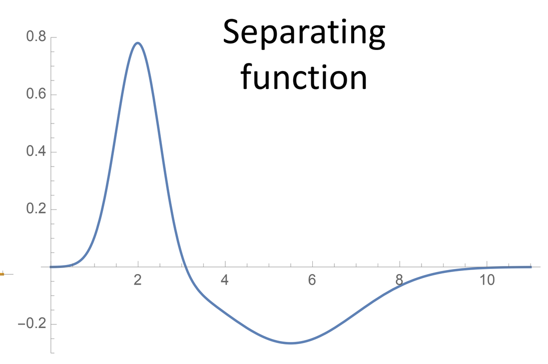

- Can search over linear combination of profile PDFs with compact linear program
- No existence proof although strong sufficient conditions for existence
- Search for SF using weaker condition $\mathbb{E}_{A(i)}(D(x)) + v(A(i)) > \mathbb{E}_{\overline{A}(i)}(D(x))$
- Convert using linearity of expectation
- However, SF introduce variance—we study empirically

Experiments

- Goals: measure social welfare between POU, MPOU and fixed-rate tariffs
- Study variance costs of introducing SFs

Instance Generation

- Learn utility models from electricity use data (pecanstreet.org)
- Generate agents by sampling from utility functions
- Create revenue-equivalent fixed-rate and POU tariffs
  - *Predictivity emphasis (PE):* how much penalty for bad predictions
- Generate profiles for each agent
  - Profile spacing: measure of how dissimilar profiles are

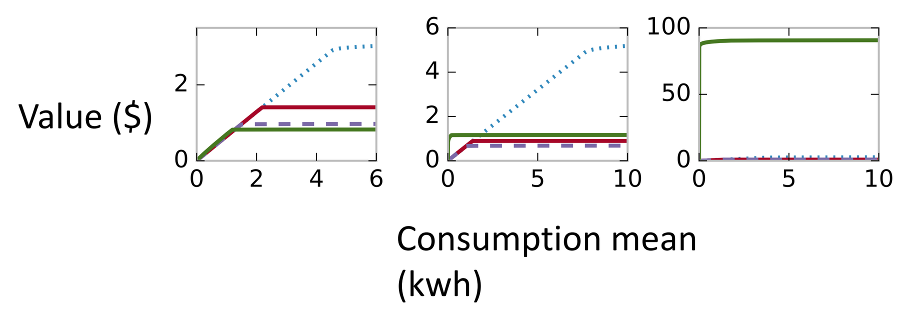

Social Welfare

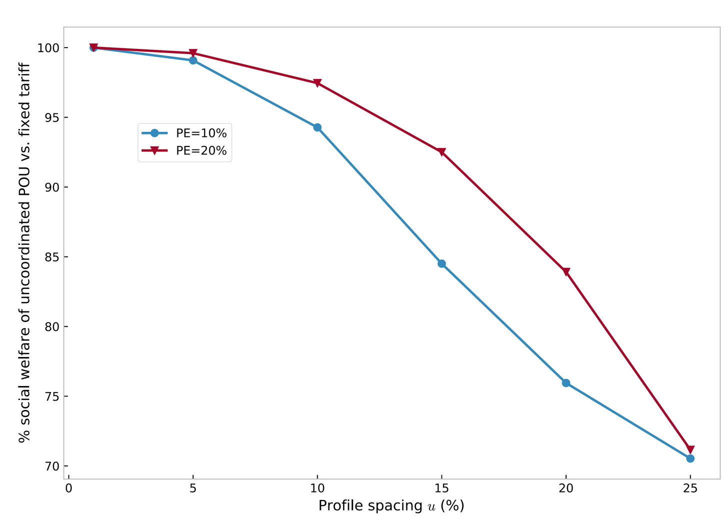

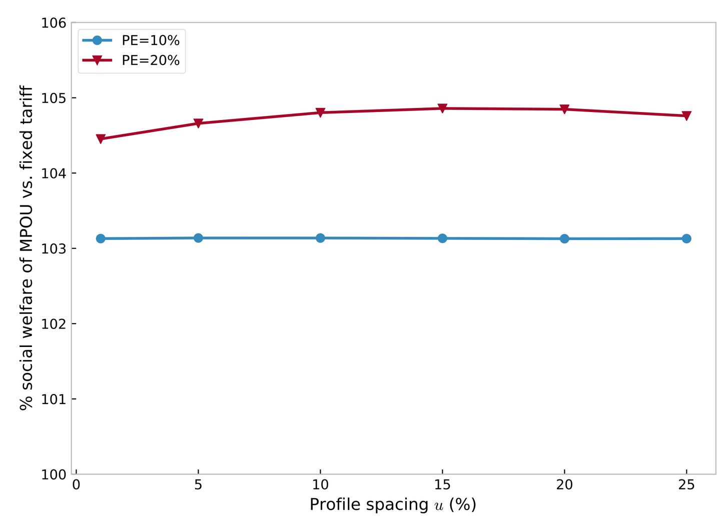

- POU setting: agents choose profile with highest net utility as if alone
- Because grand coalition forms, individual agents overestimate cost of variance
- Result is net loss of social welfare relative to fixed-rate tariff
- MPOU setting: modest gain over fixed-rate
- Numbers subject to change: limited data about PEs and agent value for variance
  - But direction of effects is clear

Variance Introduced by SFs

- Record average variance of SF payment as a fraction of Shapley payment
  - Only for the 10% of agents that need SFs
- Substantial variance introduced, but: 1) SFs payments independent and 2) our SFs do not explicitly minimize variance

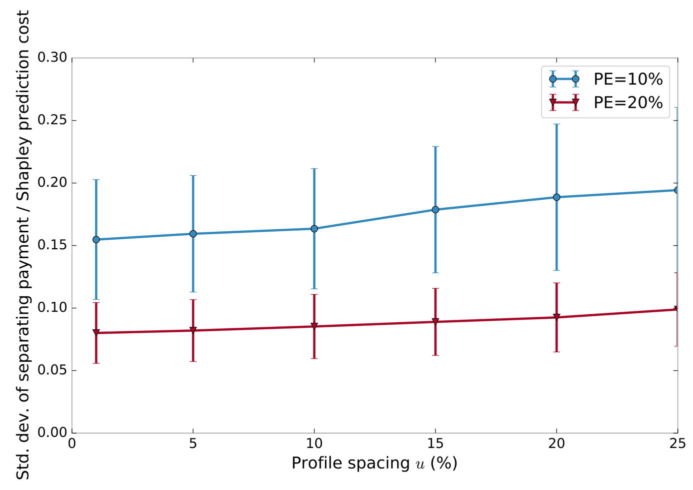

Future Work

- POU games: 1) manipulability and 2) correlated errors
- Separating functions: more applications?