Multiple-Profile Prediction-of-Use Games

Andrew Perrault and Craig Boutilier

{perrault, cebly}@cs.toronto.edu

Computer Science
UNIVERSITY OF TORONTO

Electricity Market

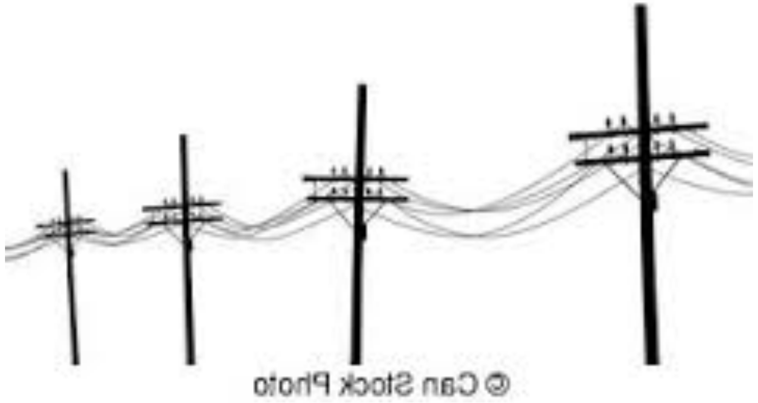

Consumer pays per kWh used, a *fixed-rate tariff*

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 (POU) Tariffs

- Each consumer makes a prediction ahead of time
  - They are charged based on:
    - How much they consume
    - How accurate their prediction was
- Consumers can form groups and be treated as one large consumer
  - But they can only do this if they can agree on how to split the costs

Contributions

- Extend POU games to support multiple profiles
  - Extension remains convex
  - Creates new enforcement problems addressed by separating functions
- Experimentally validate our approach using learned utility models

Intro to Cooperative Games

Cooperative Games

- Set of agents N
- Can form *coalitions*
  - Characteristic value function $v\colon 2^N \to \mathbb{R}$ represents value each coalition can achieve
- Agents can defect to other coalitions, but are forced to cooperate within coalition
  - Coalition can enforce contracts
- Definition (*superadditivity*): $v(S \cup T) \geq v(S) + v(T)$
  - *Grand coalition* of all agents maximizes utility

Benefit Sharing

- Challenge: how to share benefits among its members?
- Def. (*stability*): no agent has incentive to defect to another coalition
- Two major approaches:
  - Core allocation: strong stability guarantees, but hard computation
  - Shapley value: fairness, “easy” to approximate, no stability guarantee
- If game is *convex* (has a supermodular characteristic function):
  - Shapley value (and some approximations) is a core allocation (Shapley, 1971)
  - Can cheaply get fairness of Shapley value and stability of core simultaneously

Prediction-of-Use Games

Robu et al. (2017) POU Model

- Each household has a distribution over consumption in *next time period*—a *profile*
- Households can form coalitions
  - Coalition's profile is sum of members' profiles
- Each coalition predicts a *baseline* $b \in \mathbb{R}$

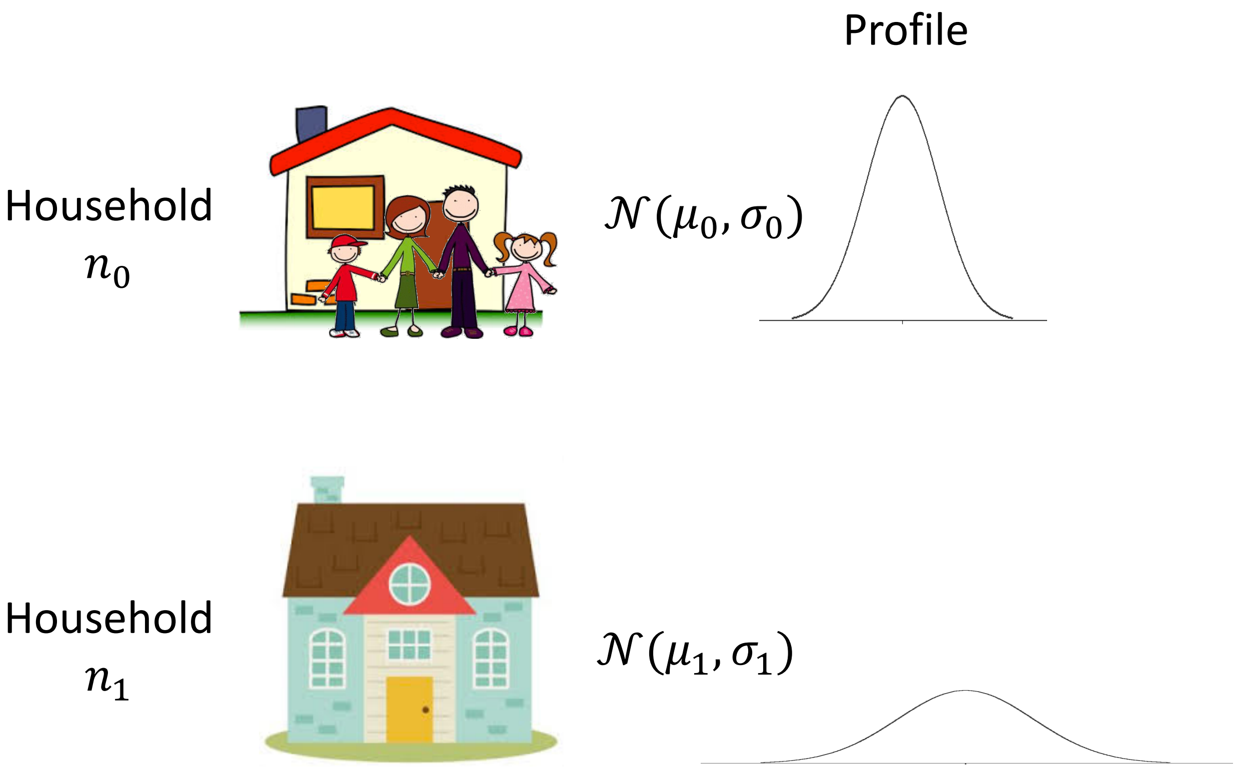

Valentin Robu, Meritxell Vinyals, Alex Rogers, and Nicholas Jennings. Efficient Buyer Groups with Prediction-of-Use Electricity Tariffs. *IEEE Transactions on Smart Grid* (2017).

Robu et al. (2017) POU Model

- Three-parameter *POU tariff*:
  - Charge p for *realized consumption*
  - Charge $\overline{p}$ for each unit over baseline b
  - Charge $\underline{p}$ for each unit under baseline b
- Closed-form for optimal b
- Characteristic function is total cost in expectation
- **Characteristic function is convex**

.

Limitations of POU Games

- The only decision agents have in POU games is what profile to declare
- The choice of profile is made before the game starts
- Agent have utility functions—choosing the best profile is an optimization
- Optimal choice depends on what other agents choose

Multiple-Profile Prediction-of-Use Games

Multiple-Profile POU (MPOU) Games

- Each profile has a value
- Each household is assigned a profile by the coalition
- Characteristic function (value of a coalition):
  sum of assigned profile values minus expected costs, under best possible assignment

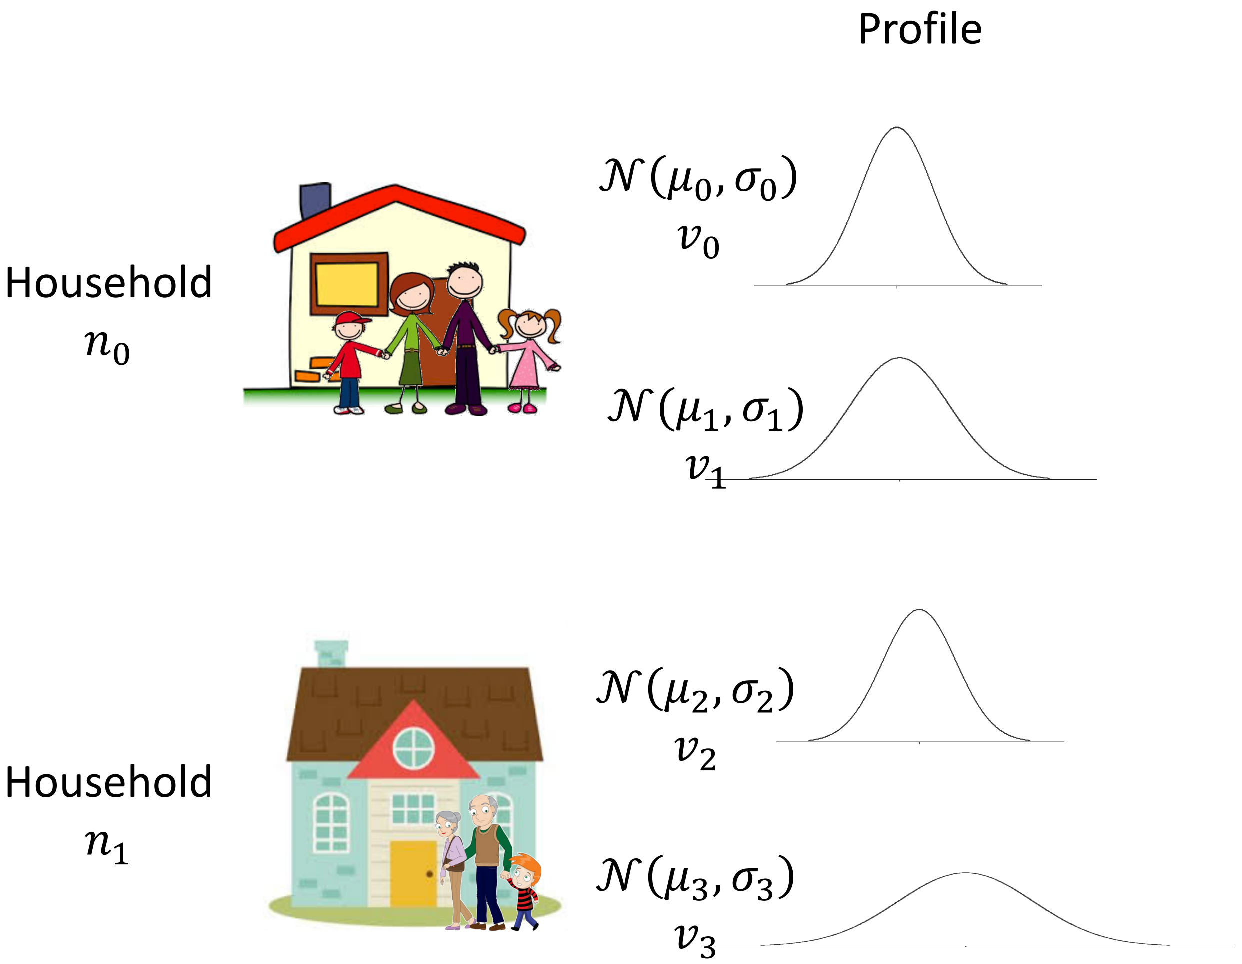

Cost Sharing in MPOU Games

- Theorem: MPOU games are convex
- Additional complexity does not interfere with convexity 😀
- However, having multiple profiles creates a new issue

Enforceability

- Coalition assigns a profile to each agent
- Actions are only partially observable in MPOU
  - Coalition knows each agent's profiles
  - *Selected* profile 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
- Definition: $D(x)$ is a separating function under assignment *A* of agents to profiles if:
  - $\mathbb{E}_{A(i)}\big(D(x)\big) + v\big(A(i)\big) > \mathbb{E}_{\overline{A}(i)}\big(D(x)\big) + v\left(\overline{A}(i)\right)$ (incentive)
  - $\mathbb{E}_{A(i)}\big(D(x)\big) = 0$ (zero-expectation)

Empirical Results

Goal of Experiments

- Measure social welfare difference between POU, MPOU and fixed-rate tariff
  - Use agent utility functions learned from pecanstreet.org data

Social Welfare: POU vs. Fixed-Rate

- POU suffers a large SW loss vs. fixed due to lack of coordination

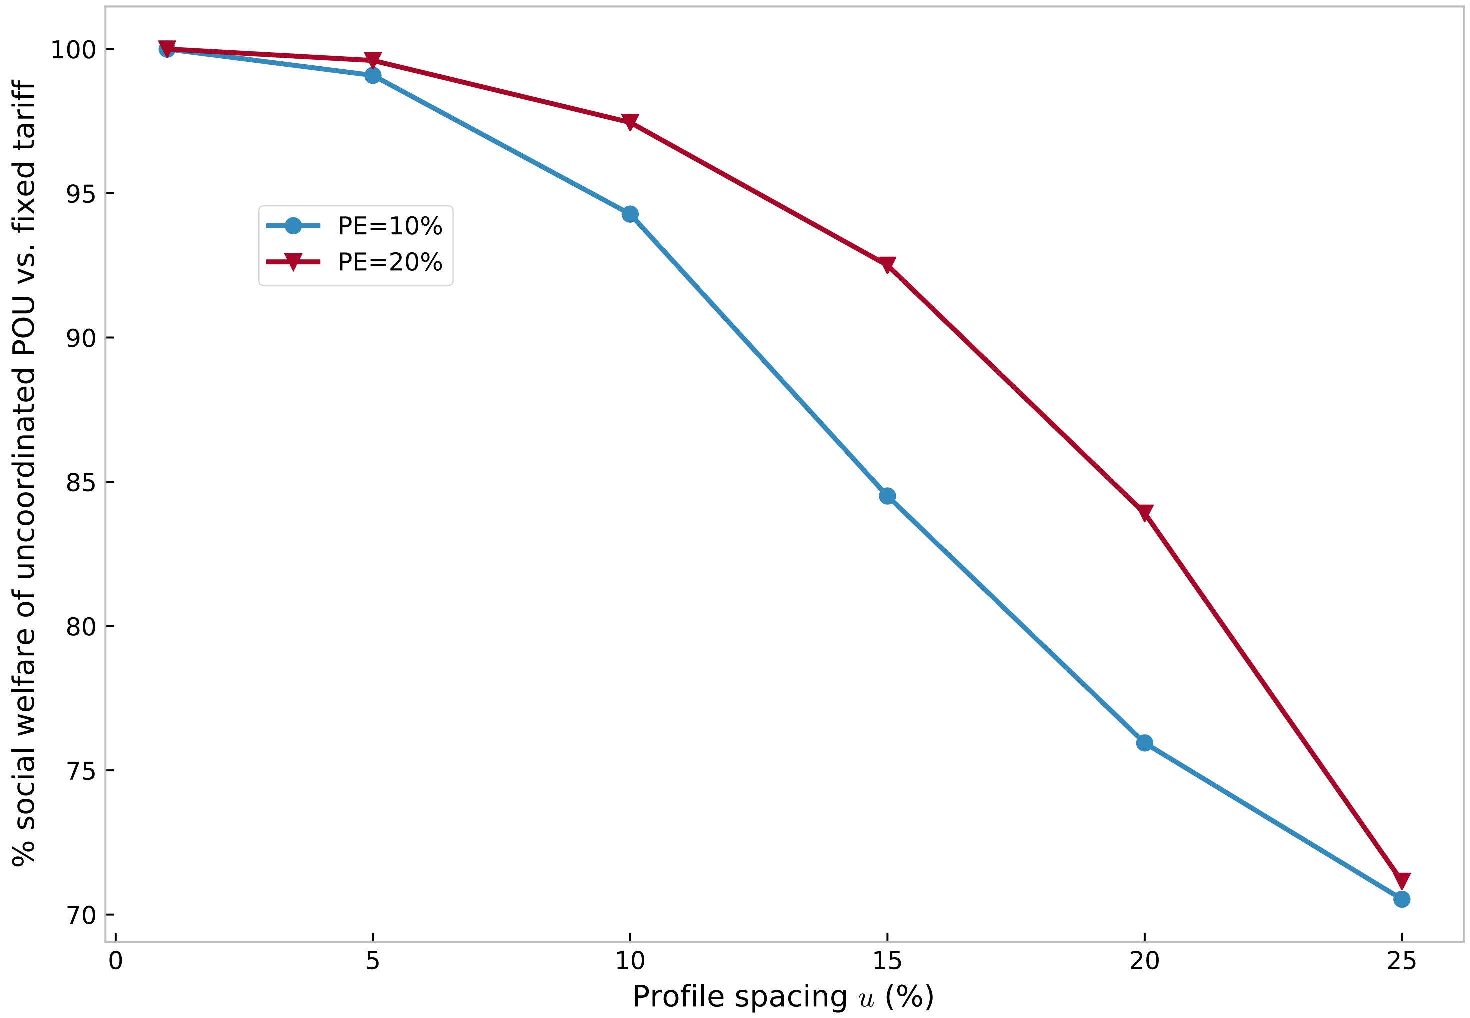

MPOU vs. Fixed-Rate

- MPOU shows a modest SW gain over fixed

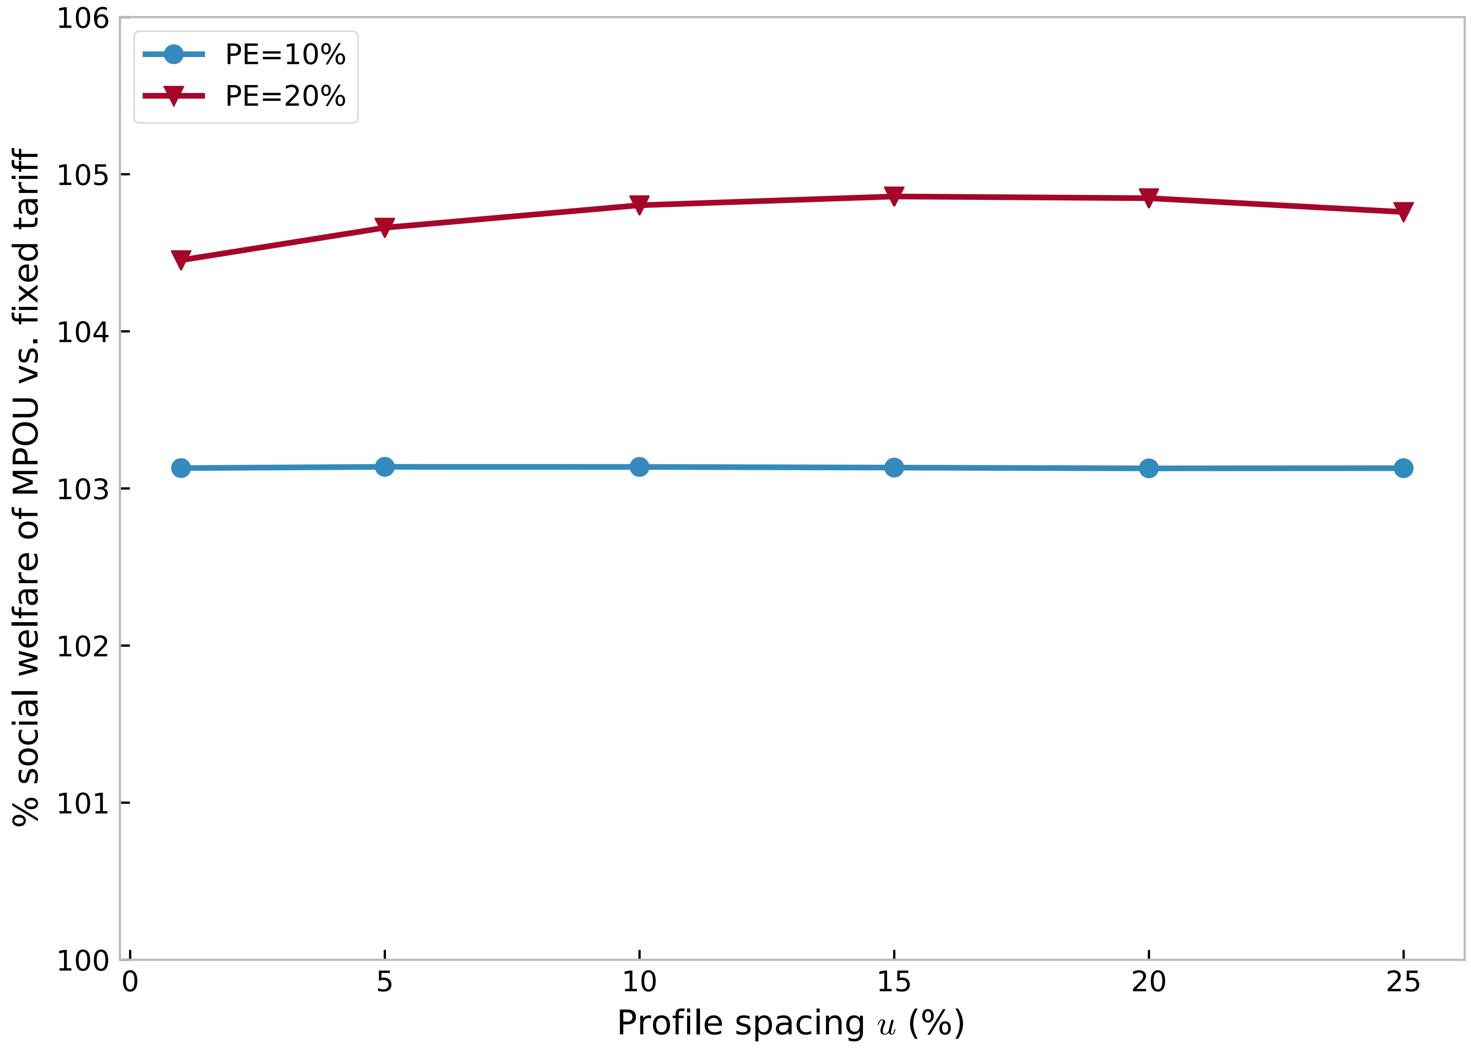

Contributions

- Extend POU games to support multiple profiles
  - Extension remains convex
  - Creates new enforcement problems addressed by separating functions
- Experimentally validate our approach using learned utility models
  - Social welfare: POU < fixed-rate < MPOU

Future Work

- POU games:
  - Manipulation
  - Correlated prediction errors
- Separating functions:
  - General applicability to principal-agent problems

Thank You

- Poster #1938

Electricity Markets

- Electricity consumption is hard to predict for suppliers
- Predictable consumers are cheaper to serve
- Residential consumers face a fixed-rate tariff
  - They are not incentivized to be predictable

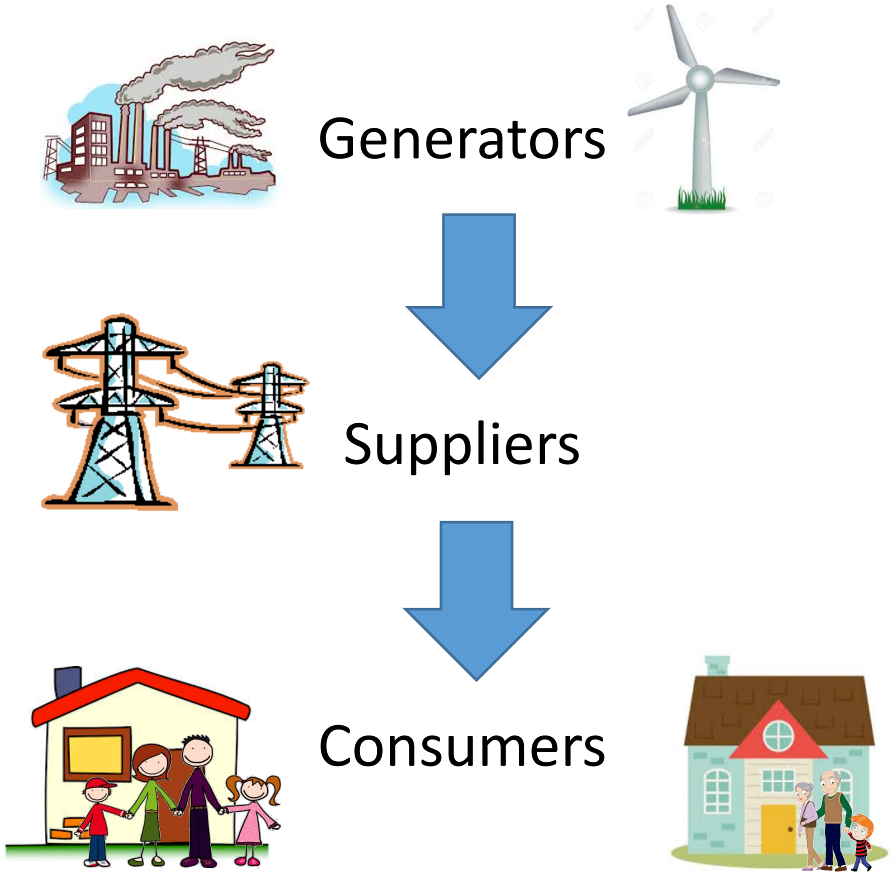

SF Calculation, Two-Profile Example

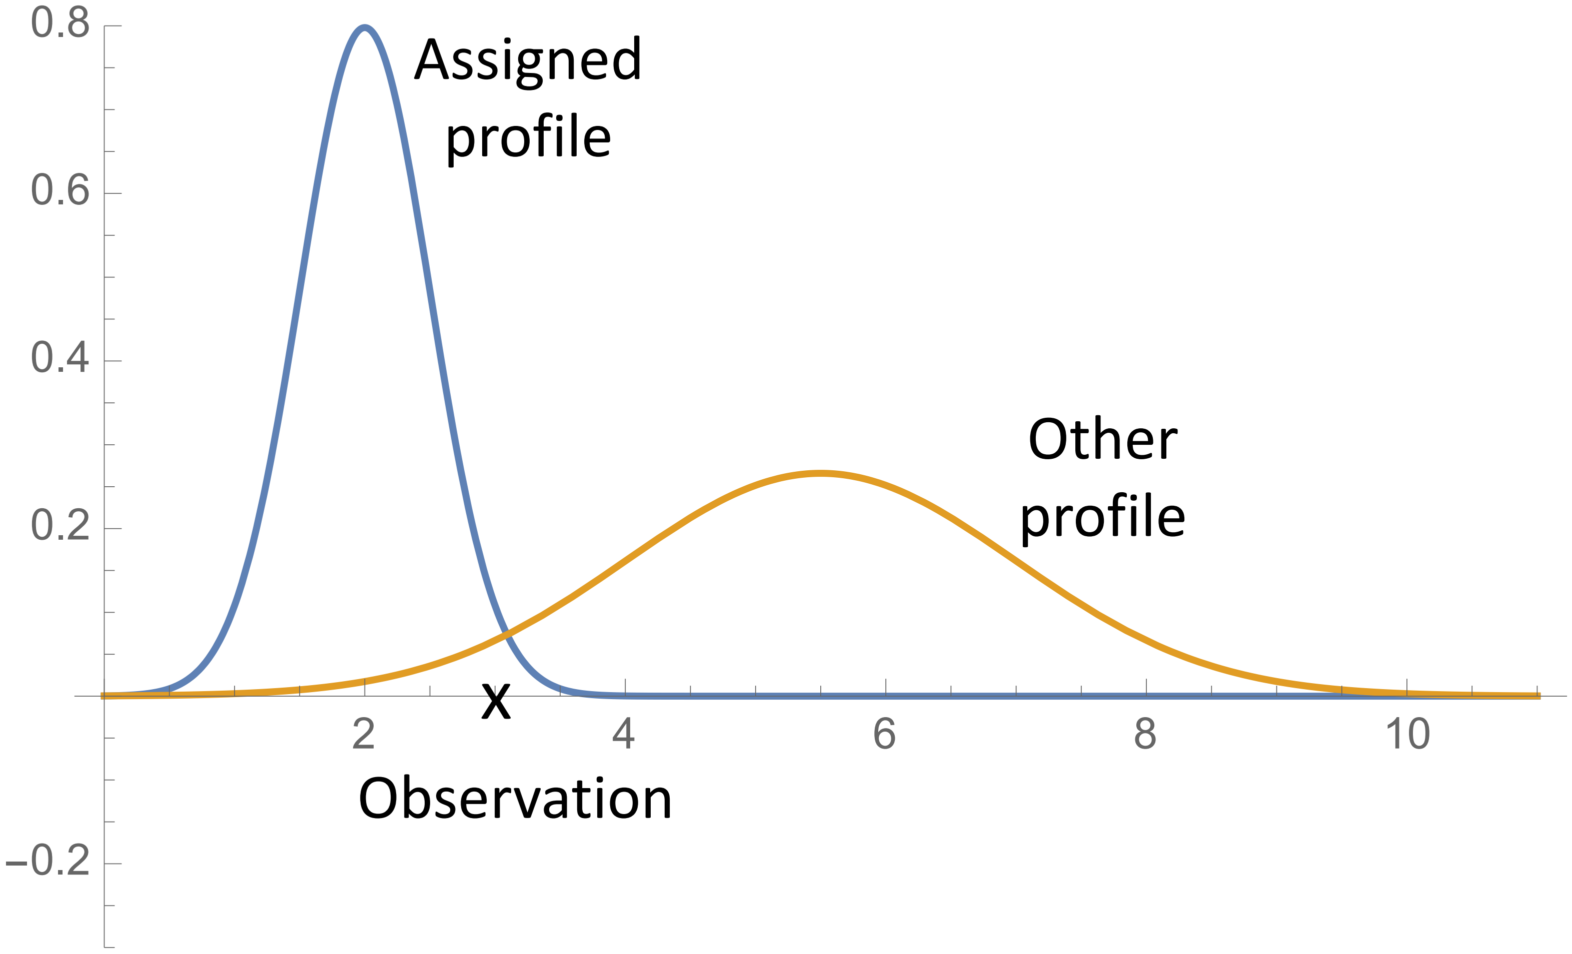

SF Calculation, Two-Profile Example

- Theorem: separating function for two-profiles:

PDF(assigned profile) – PDF(other profile)

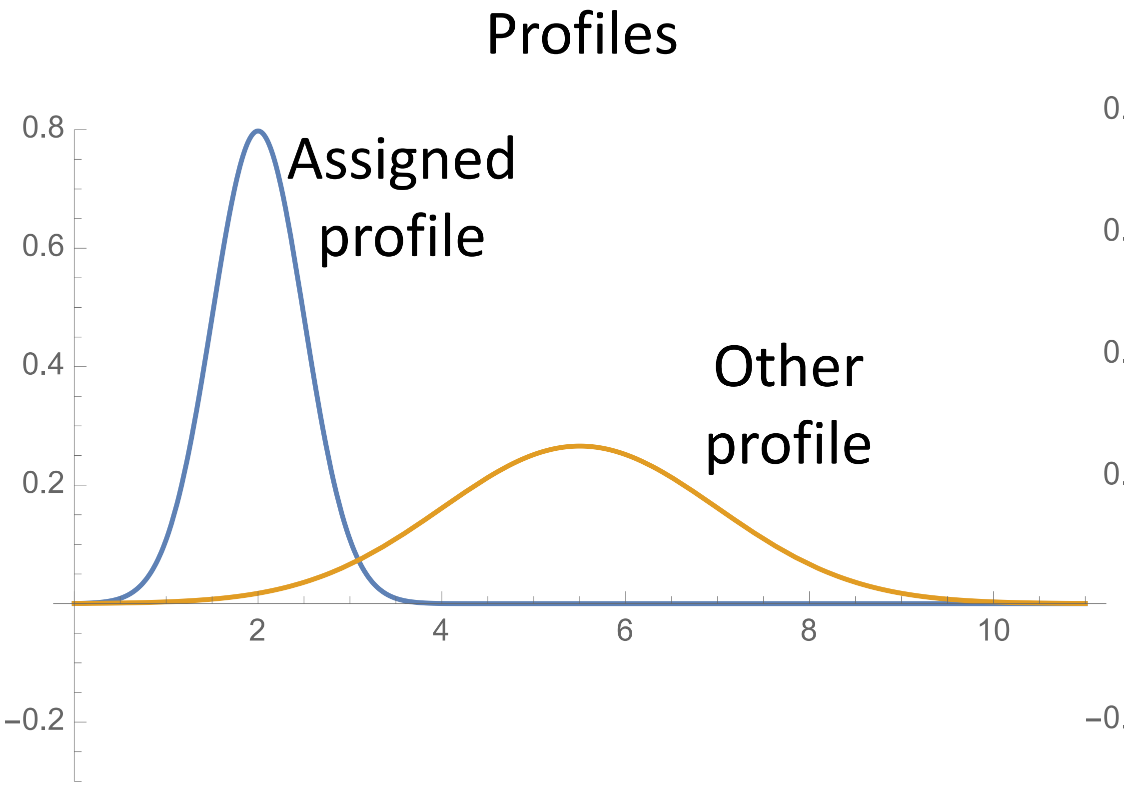

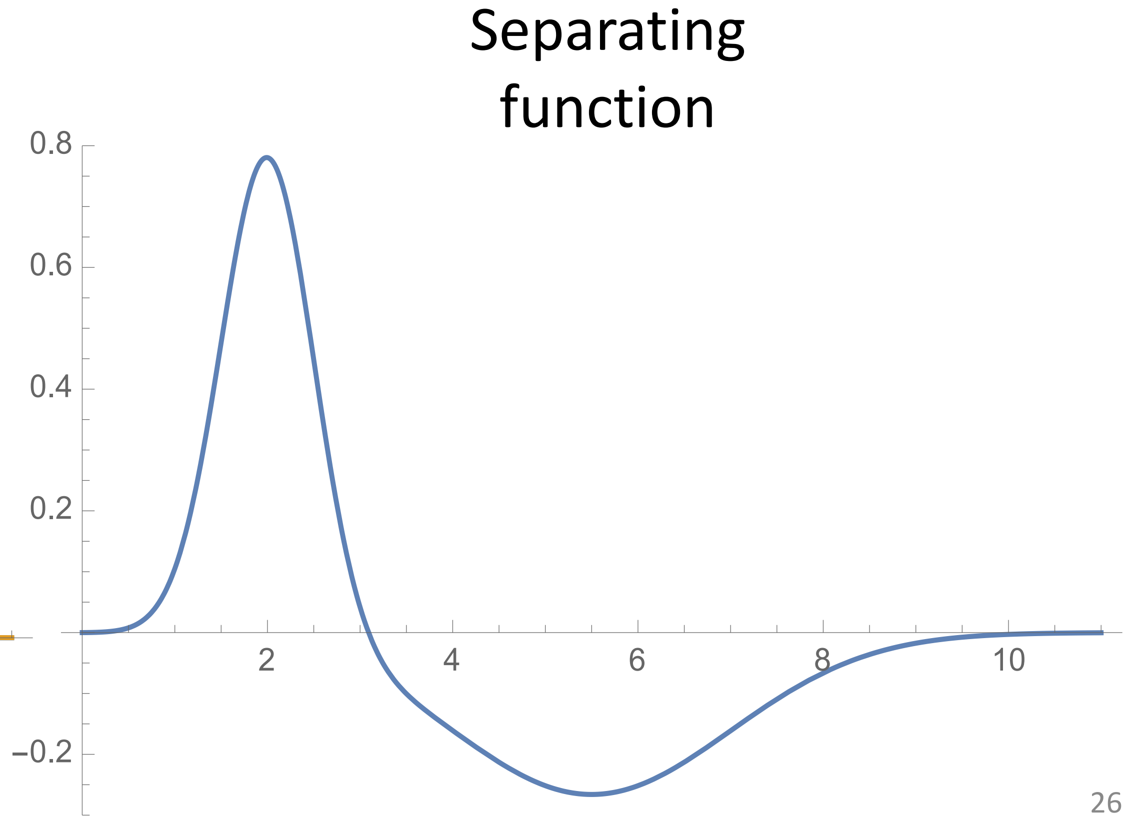

SF Calculation, Arbitrary Number of Profiles

- No closed form that we know of
- Can compute using a linear program
- We suspect they always exist in practice

Addressing the Problems of Weak SFs

- Key observation: weak separating condition maintained under affine transformations
  - A weak separating function can have arbitrary power through scaling
  - Can make $\mathbb{E}_{A(i)}\big(D(x)\big) = 0$ through translation. Thus, payments not affected

SFs Introduce Variance

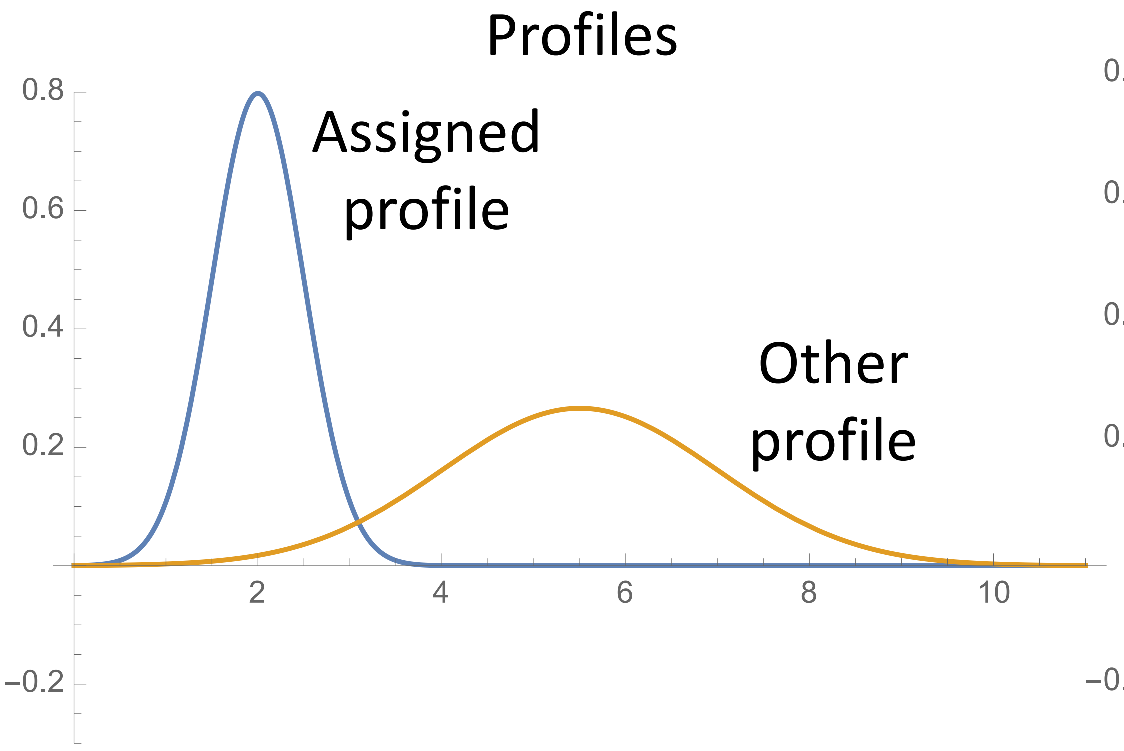

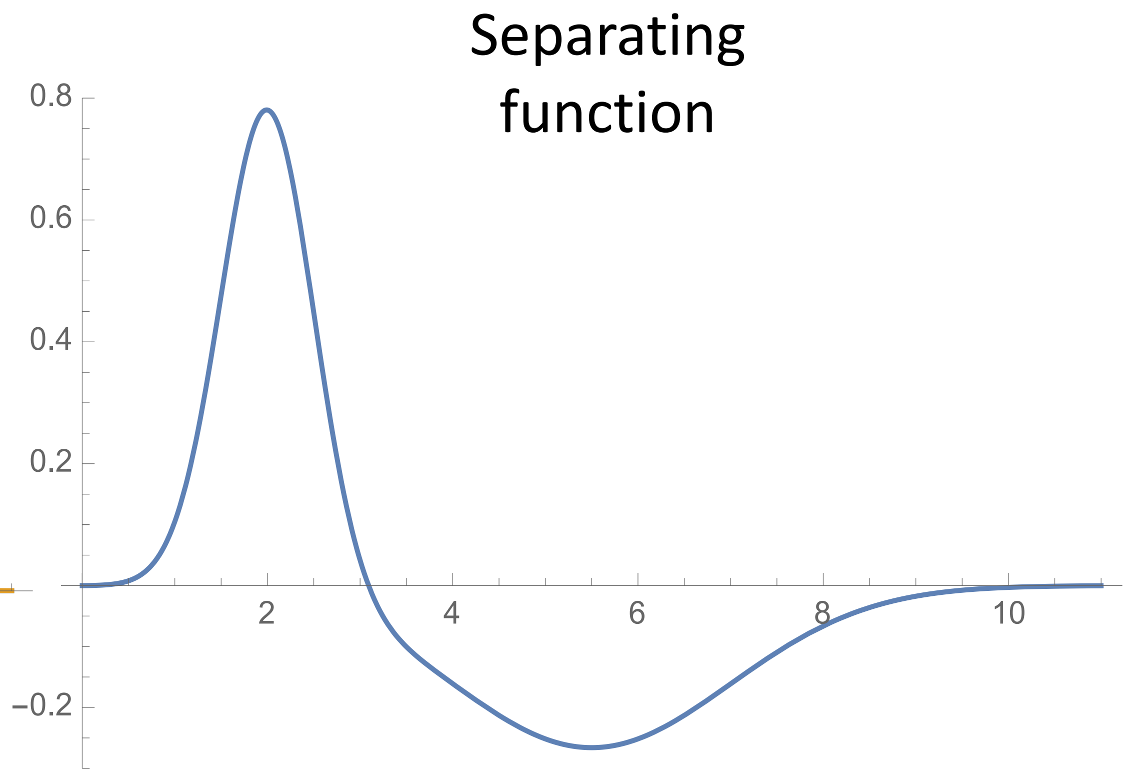

Utility Models

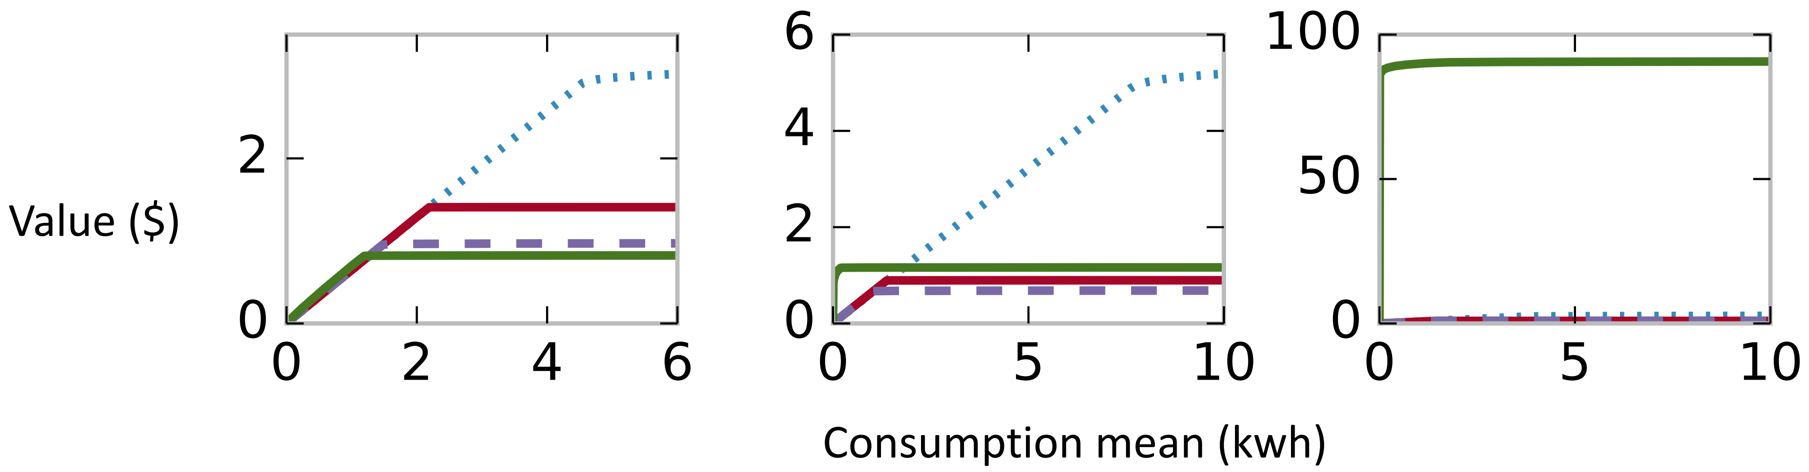

Free for academic use at pecanstreet.org

Instance Generation

- Generate agents by sampling utility functions
- Create revenue-equivalent fixed-rate and POU tariffs
  - *Predictivity emphasis (PE)*: parameter for how much agents are penalized for prediction errors relative to fixed-cost
- Generate profiles for each agent
  - *Profile spacing*: measure of how dissimilar generated profiles are

Variance Introduced by SFs

- Record average variance of SF as a fraction of Shapley payment
  - Only for agents that require SFs (1-10%) of total
- Substantial variance introduced
- If time periods are independent, variance decreases in aggregate

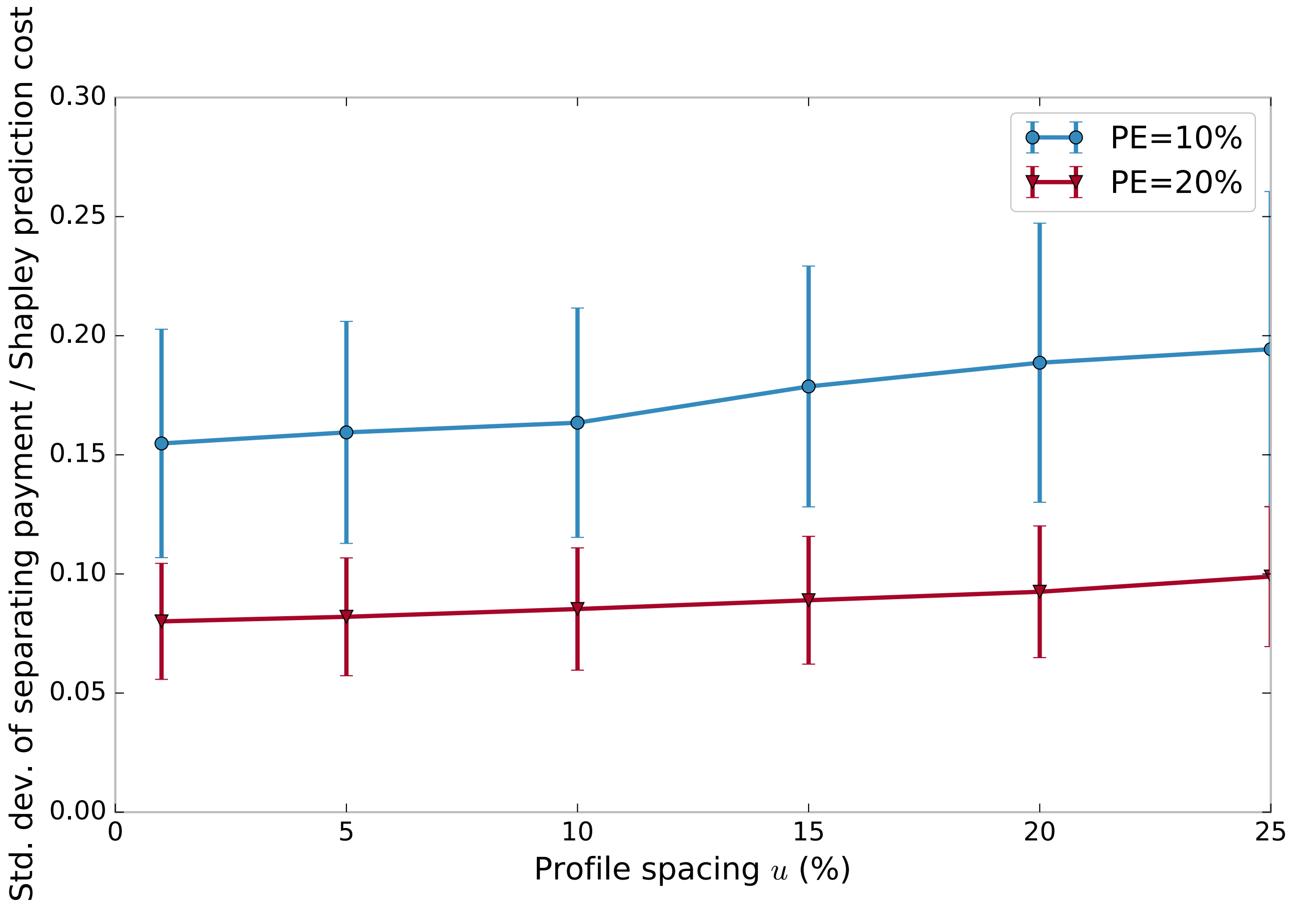

Core Allocations

- Let $t(i)$ be the payment to agent i
- *Budget balance:* distributes all benefits: $\sum_{i \in N} t(i) = v(N)$
- *Stability:* no defections possible: $\forall S \subset N, \sum_{i \in S} t(i) \geq v(S)$
- Very satisfying, but, in general:
  - May not exist
  - May be hard to compute

Shapley Value

- Intuition: agent's average contribution to coalition value
- Guarantees "fairness"
- Budget-balanced, but not guaranteed to be stable
- Easy to approximate
- If a

Scalability

- 100k agents with 4 profiles each takes 90 minutes
- Largest bottleneck is Shapley value computation: $n \log n$ linear programs (LPs) where n is number of agents
  - Each LP has nk variables where k is number of profiles
  - Need to calculate coalition values n times for each sample
  - Need $\log n$ samples
- Separating function LPs have k^2 variables each

Instance Generation

1. Generate agent population by sampling utility function and adding noise

2. Calculate optimal profile for each agent under fixed rate tariff with rate p_{fixed}

(using generated profiles)

3. Calculate revenue equivalent POU tariff $< p, p^+, p^- >$ with desired predictivity emphasis *PE* (portion of revenue that is based on predictivity)

4. Calculate optimal profile under POU tariff, ignoring variance costs (*base profile*)

5. Create additional profiles according to profile spacing *u*

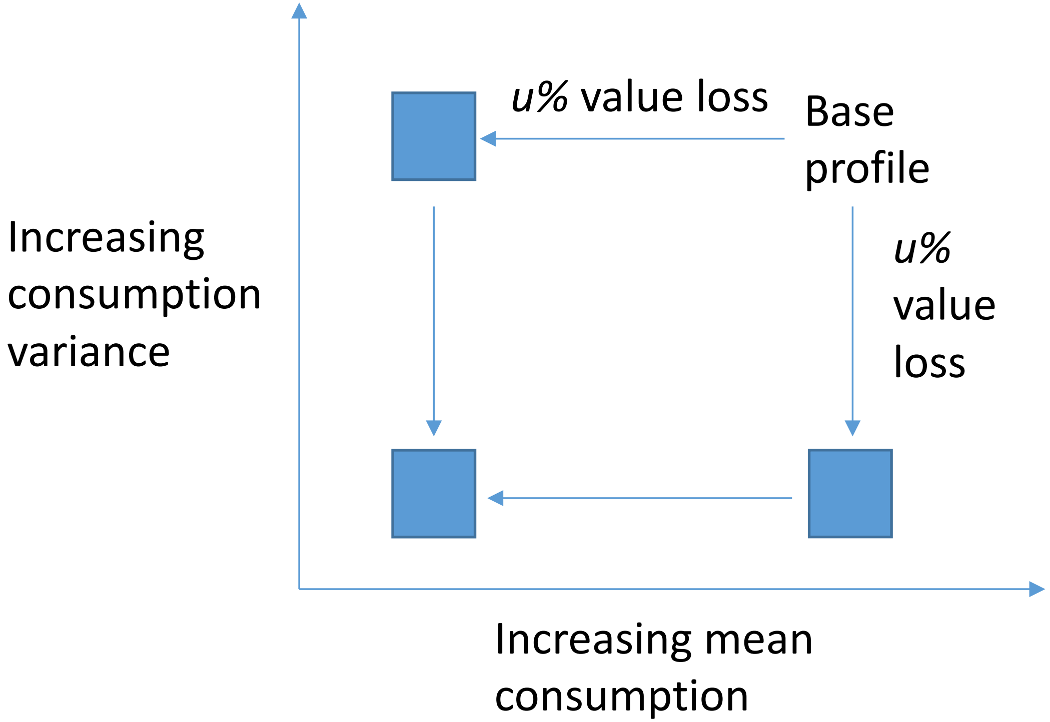

Profile Spacing

- Base profile maximizes utility ignore variance → only need to consider profiles that reduce variance
- But we don't know what spacing will maximize social welfare (SW)