Developing and Coordinating Autonomous Agents for Efficient Electricity Markets

Andrew Perrault

USC CAIS, March 7, 2018

About me: Cornell

- Computational sustainability: spatially-balanced Latin squares
  - Improved from exponential to $O(n^2)$
  - Leads to computer-aided discovery of constructive procedures for combinatorial objects (LeBras, Gomes and Selman, 2012)
- Machine learning in robotics
  - Using supervised learning to “train” robots to manipulate objects (and generalize to unseen objects)

The School Fund

- Co-lead developer on crowdfunding platform for secondary school scholarships for students in the developing world
- 1183 years of education funded
- 12 partners in 9 countries

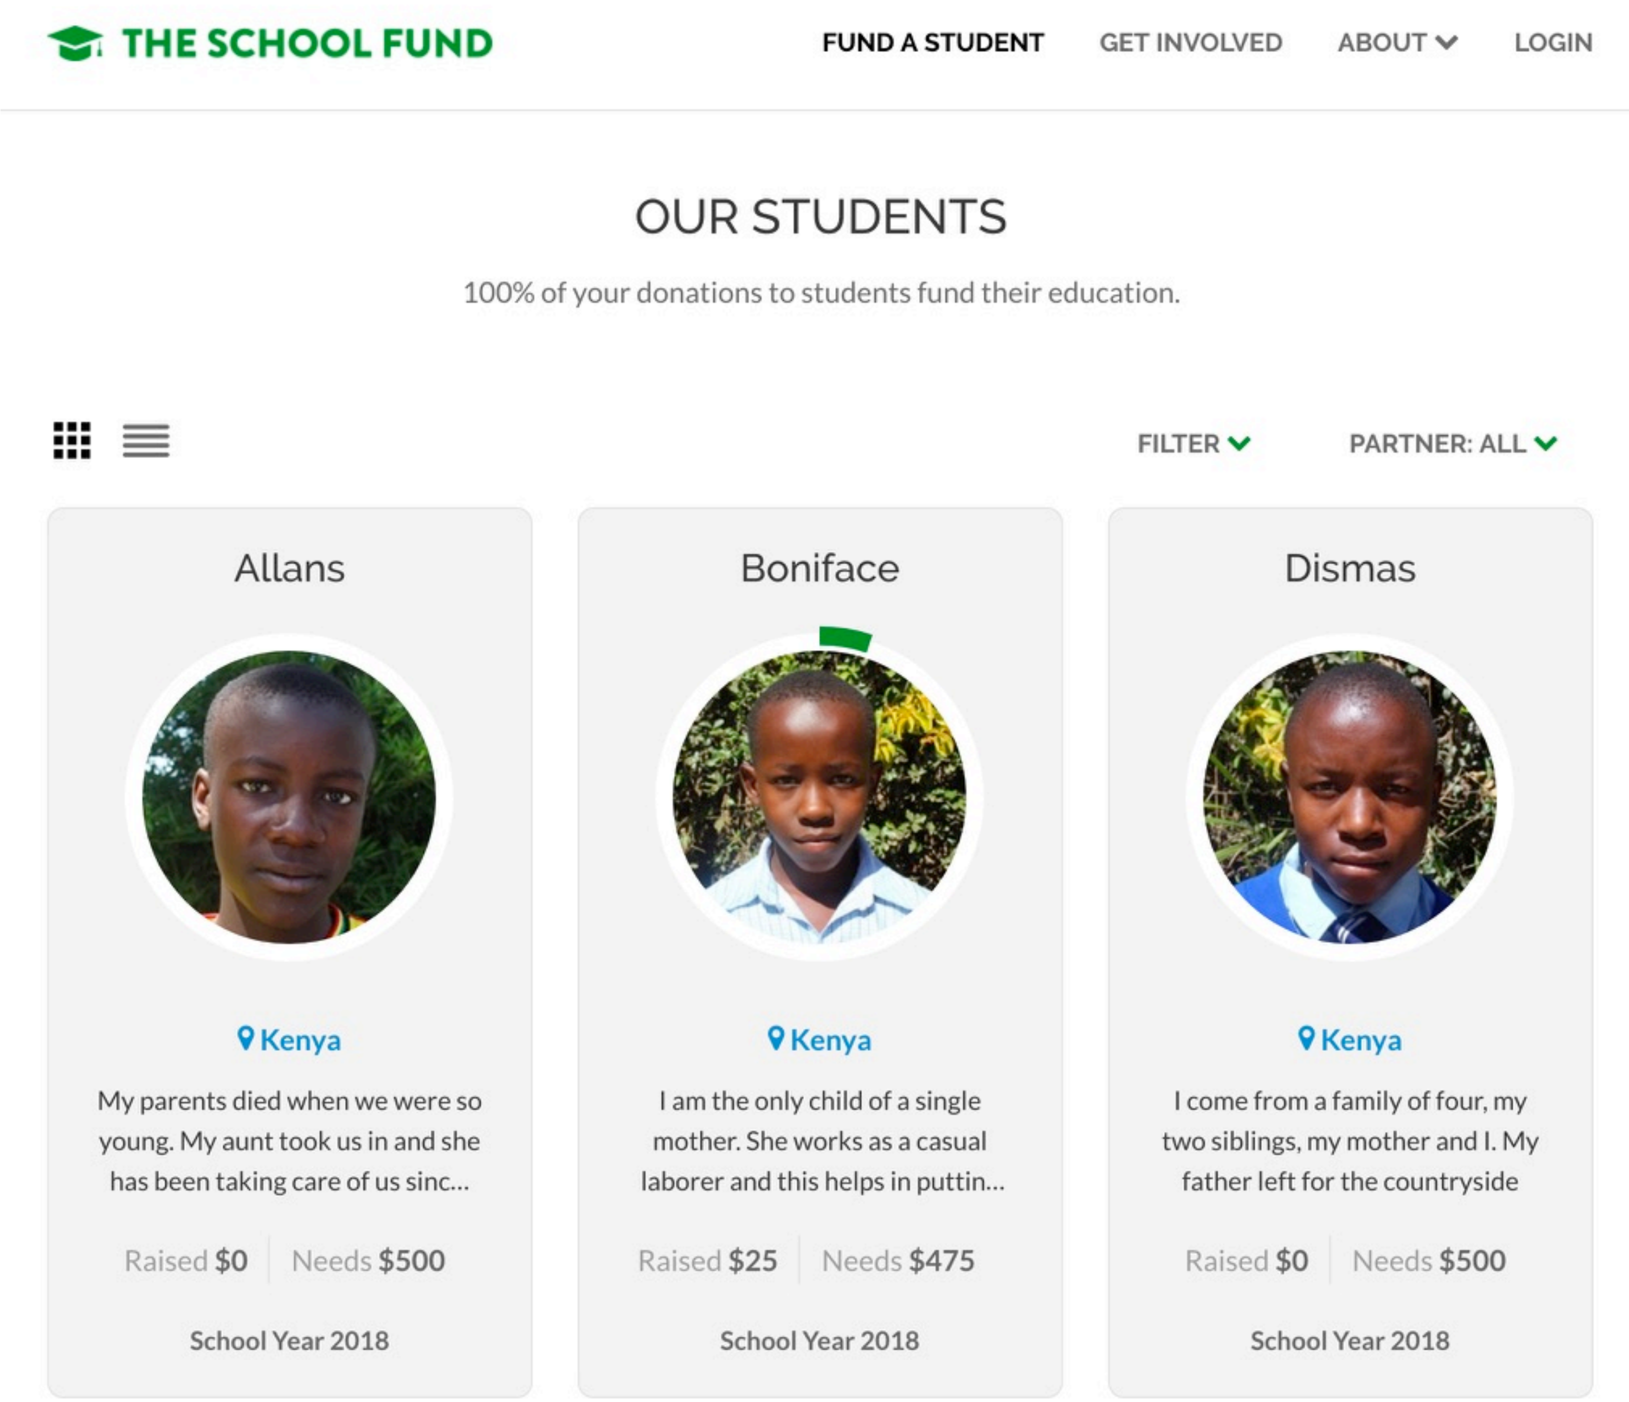

About me: University of Toronto

- Using satisfiability (SAT) to solve stable-matching problems with complementarities
  - Matching of medical residents to hospital internships uses incomplete alg.
  - SAT is complete, good performance
  - Using SAT has other advantages
    - Flexibility in adding constraints
    - Can search over the set of matchings

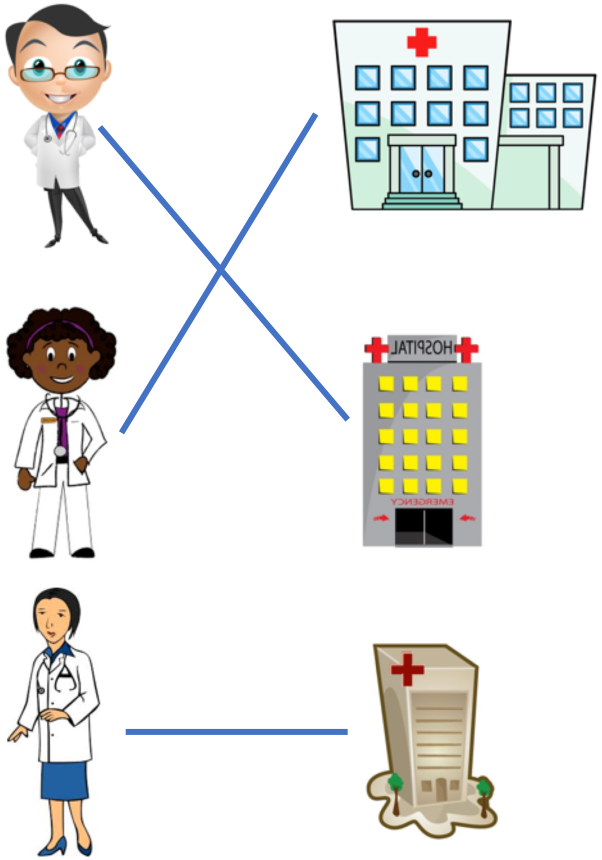

Drummond, P. and Bacchus (IJCAI-15)
P., Drummond and Bacchus (AAMAS-16)

This talk: autonomous agents for efficient electricity markets

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
  - The big picture: market design, machine learning
  - P. and Boutilier (IJCAI-17)
- Experiential elicitation for electricity management agents
  - The big picture: preference elicitation, sequential decision-making, machine learning
  - P. and Boutilier (under review)

Other work in electricity: P. and Boutilier (AAMAS-14) and P. and Boutilier (IJCAI-15)

Why study electricity in AI?

- Reduce greenhouse gas (*GHG*) emissions
  - California target: 80% below 1990 levels by 2050
- Electrify and decarbonize electricity generation

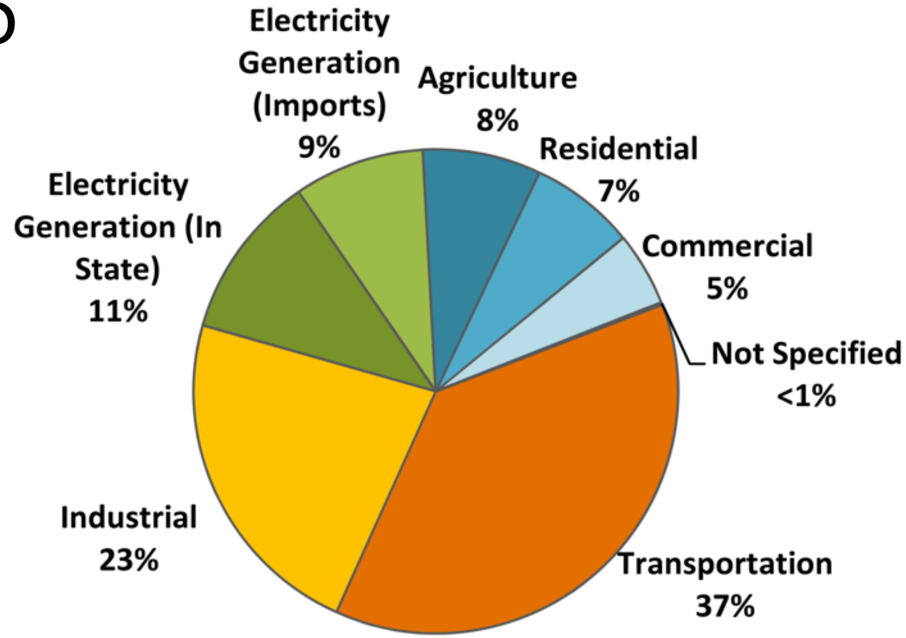

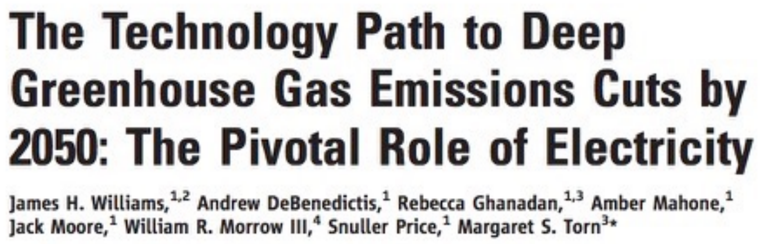

The Technology Path to Deep Greenhouse Gas Emissions Cuts by 2050: The Pivotal Role of Electricity

James H. Williams,[1,2] Andrew DeBenedictis,[1] Rebecca Ghanadan,[1,3] Amber Mahone,[1] Jack Moore,[1] William R. Morrow III,[4] Snuller Price,[1] Margaret S. Torn[3]*

Williams et al. (2012)

Achieving high renewable penetration

- Overbuild

- Transmission

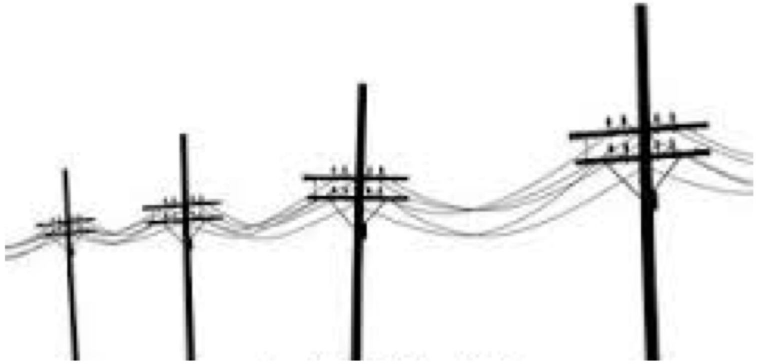

- Storage

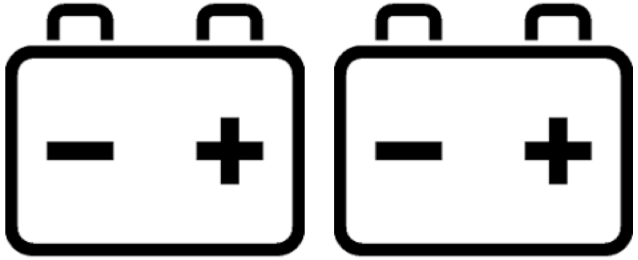

Change consumer demand

- Changeable consumer demand, currently:
  - Air conditioning (around 30% of *all* electricity at peaks in summer)
  - Various other appliances (dishwashers, washing machines, dryers)
- After electrification:
  - Personal vehicles (about 60% of all energy used in transportation)
  - Space heating and water heating

The problem with demand response

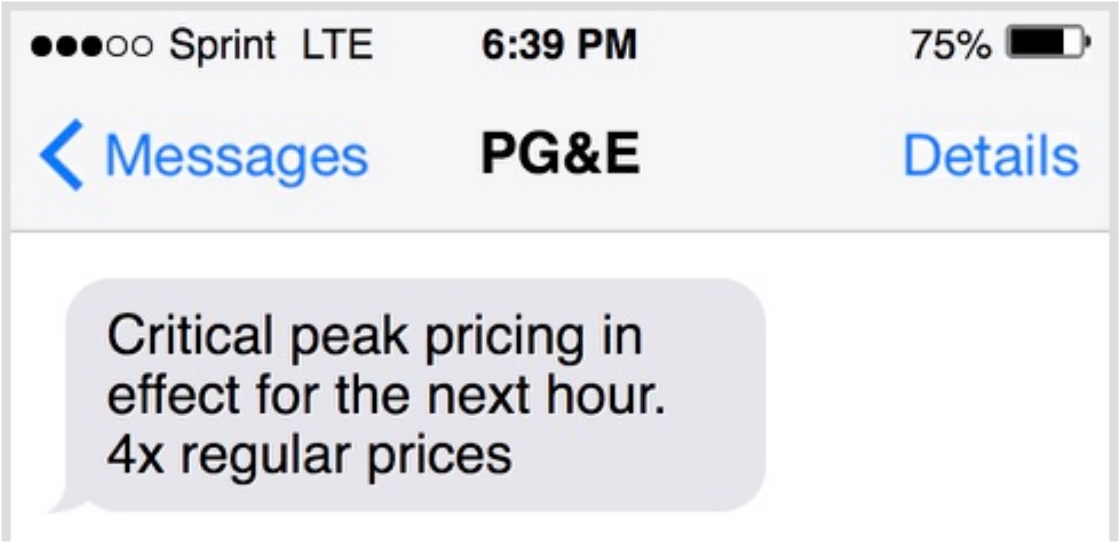

- Do you notice?
- What do you do?
- Need a response policy

The value of autonomous agents (Part I)

- Redesign electricity markets around autonomous agents to increase market efficiency while paying attention to strategic aspects

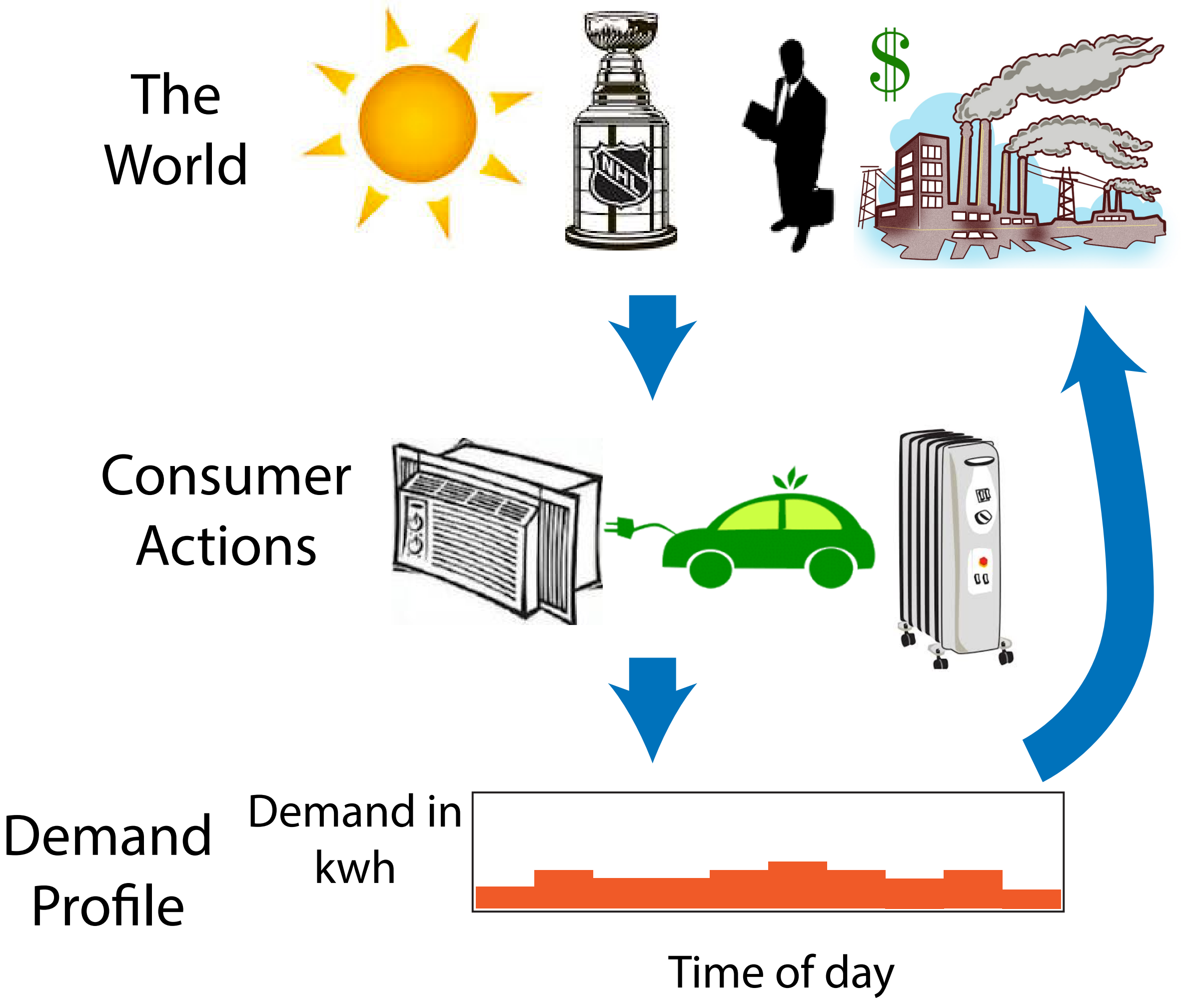

The value of autonomous agents (Part II)

- Learn a household's preferences and take actions on their behalf

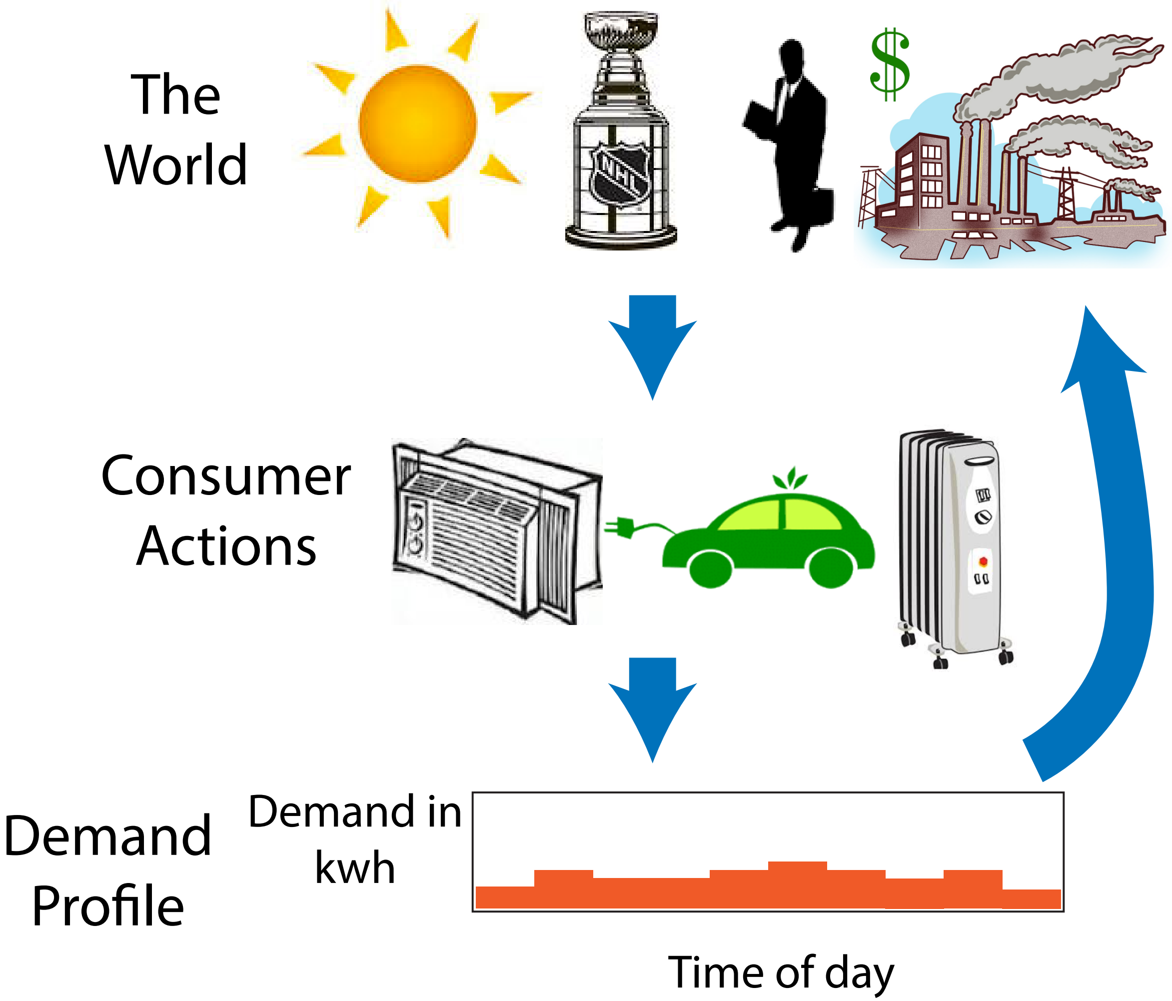

Aligning Predictability Incentives Between Consumers and Suppliers

P. and Boutilier, IJCAI-17

Electricity market

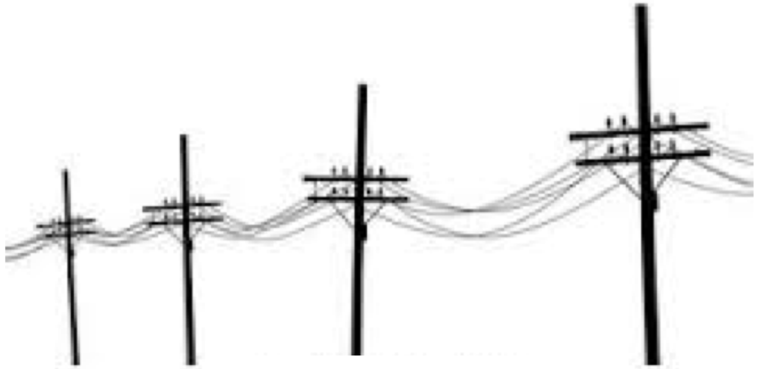

Consumer pays per kWh used, a *fixed-rate tariff*

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

Robu, Vinyals, Rogers, and Jennings. Efficient Buyer Groups with Prediction-of-Use Electricity Tariffs (2017).

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

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
  - **Intro to cooperative games**
  - Prediction-of-use games
  - Multiple-profile prediction-of-use games
  - Empirical results
- Experiential elicitation for electricity management agents

Cooperative games

- Set of agents N
- Can form coalitions
  - Characteristic value function $v: 2^N \to \mathbb{R}$ represents value each coalition can achieve
- Agents can defect to other coalitions, but are forced to cooperate within coalition
  - Coalition can enforce contracts

Sharing benefits

- *Definition (stability)*: no set of agents 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)

Cooperative games and markets

- Market/cooperative game duality
  - Shapley and Shubik (1969): *exchange economies* can be formulated as cooperative games
- Useful because cooperative games are more flexible than exchange economies

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
  - Intro to cooperative games ✓
  - **Prediction-of-use games**
  - Multiple-profile prediction-of-use games
  - Empirical results
- Experiential elicitation for electricity management agents

Robu et al. (2017) POU games 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 $b \in \mathbb{R}$

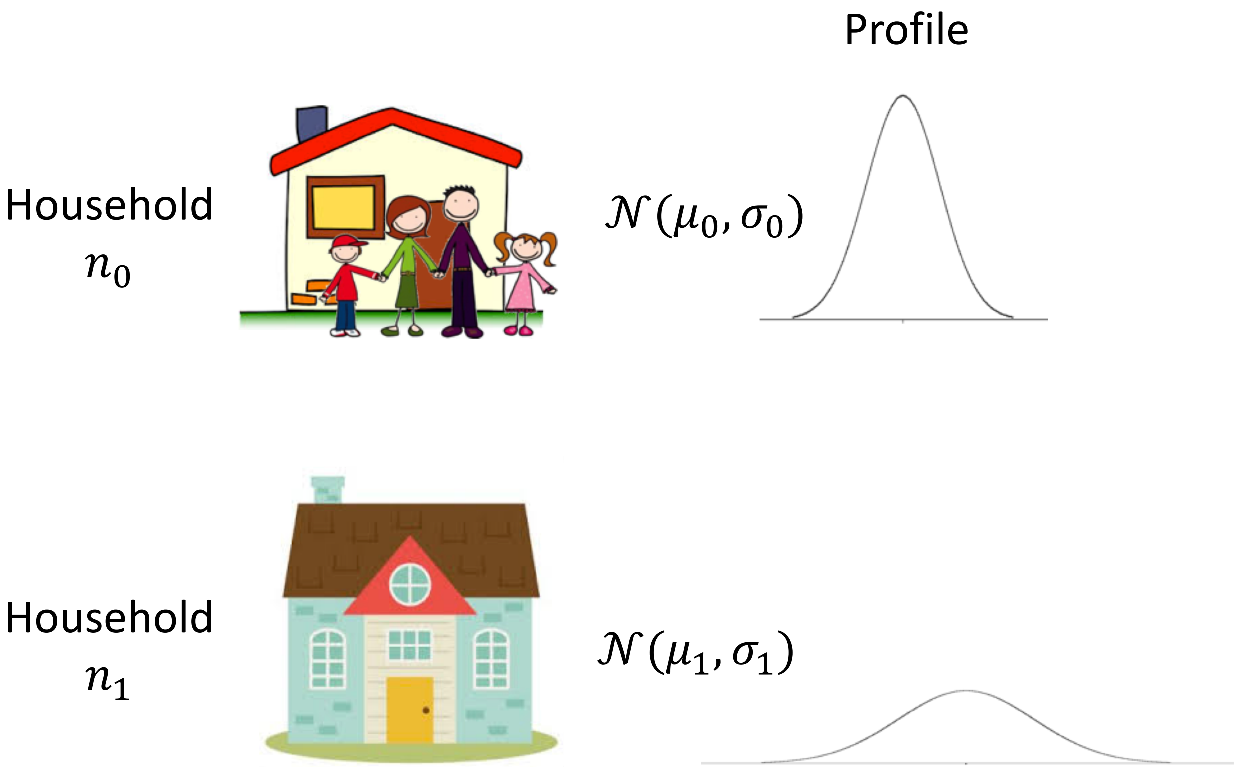

Robu et al. (2017) POU model

- 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**

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
  - Intro to cooperative games ✓
  - Prediction-of-use games ✓
  - **Multiple-profile prediction-of-use games**
  - Empirical results
- Experiential elicitation for electricity management agents

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

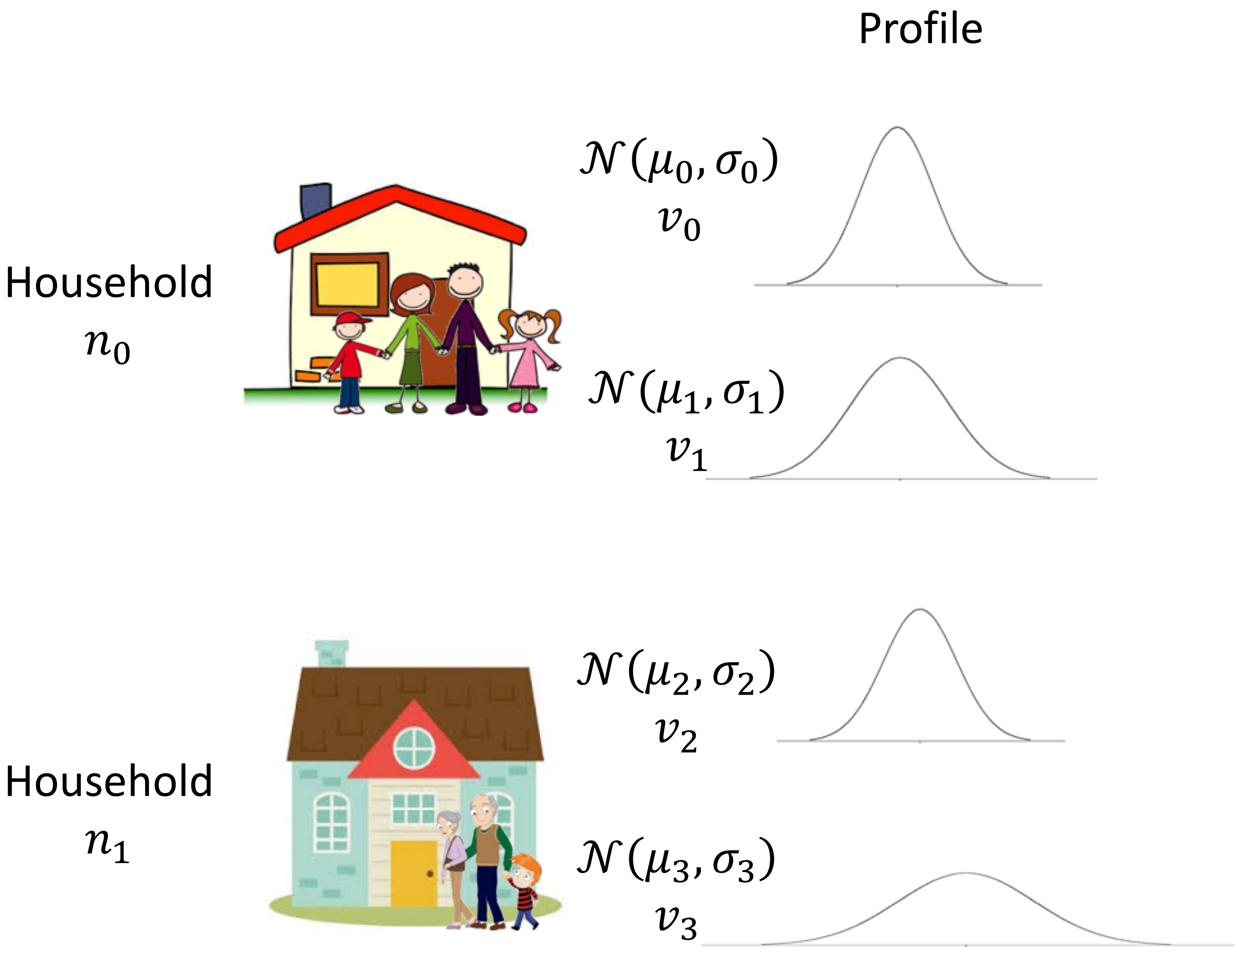

Cost sharing in MPOU games

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

Enforcing assigned actions

- 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

Example of defection

- Assigned profile with lower value (v_2), pays ex-ante according to assignment
- Uses profile with higher value (v_3)
- **Defection not fully observable**

$\mathcal{N}(\mu_2, \sigma_2)$
v_2

$\mathcal{N}(\mu_3, \sigma_3)$
v_3

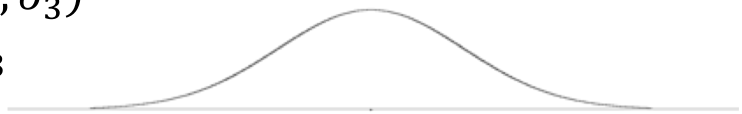

Incentivizing use of the assigned profile

- A separating function (SF) maps realized consumption to a payment from coalition to agent
- *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)
- *Theorem:* sufficient conditions for SF existence, and poly-size LP formulation

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
  - Intro to cooperative games ✓
  - Prediction-of-use games ✓
  - Multiple-profile prediction-of-use games ✓
  - **Empirical results**
- Experiential elicitation for electricity management agents

Results

- Agent utility functions learned from pecanstreet.org data
- POU: **large social welfare loss** due to lack of coordination
- MPOU: **modest SW gain**

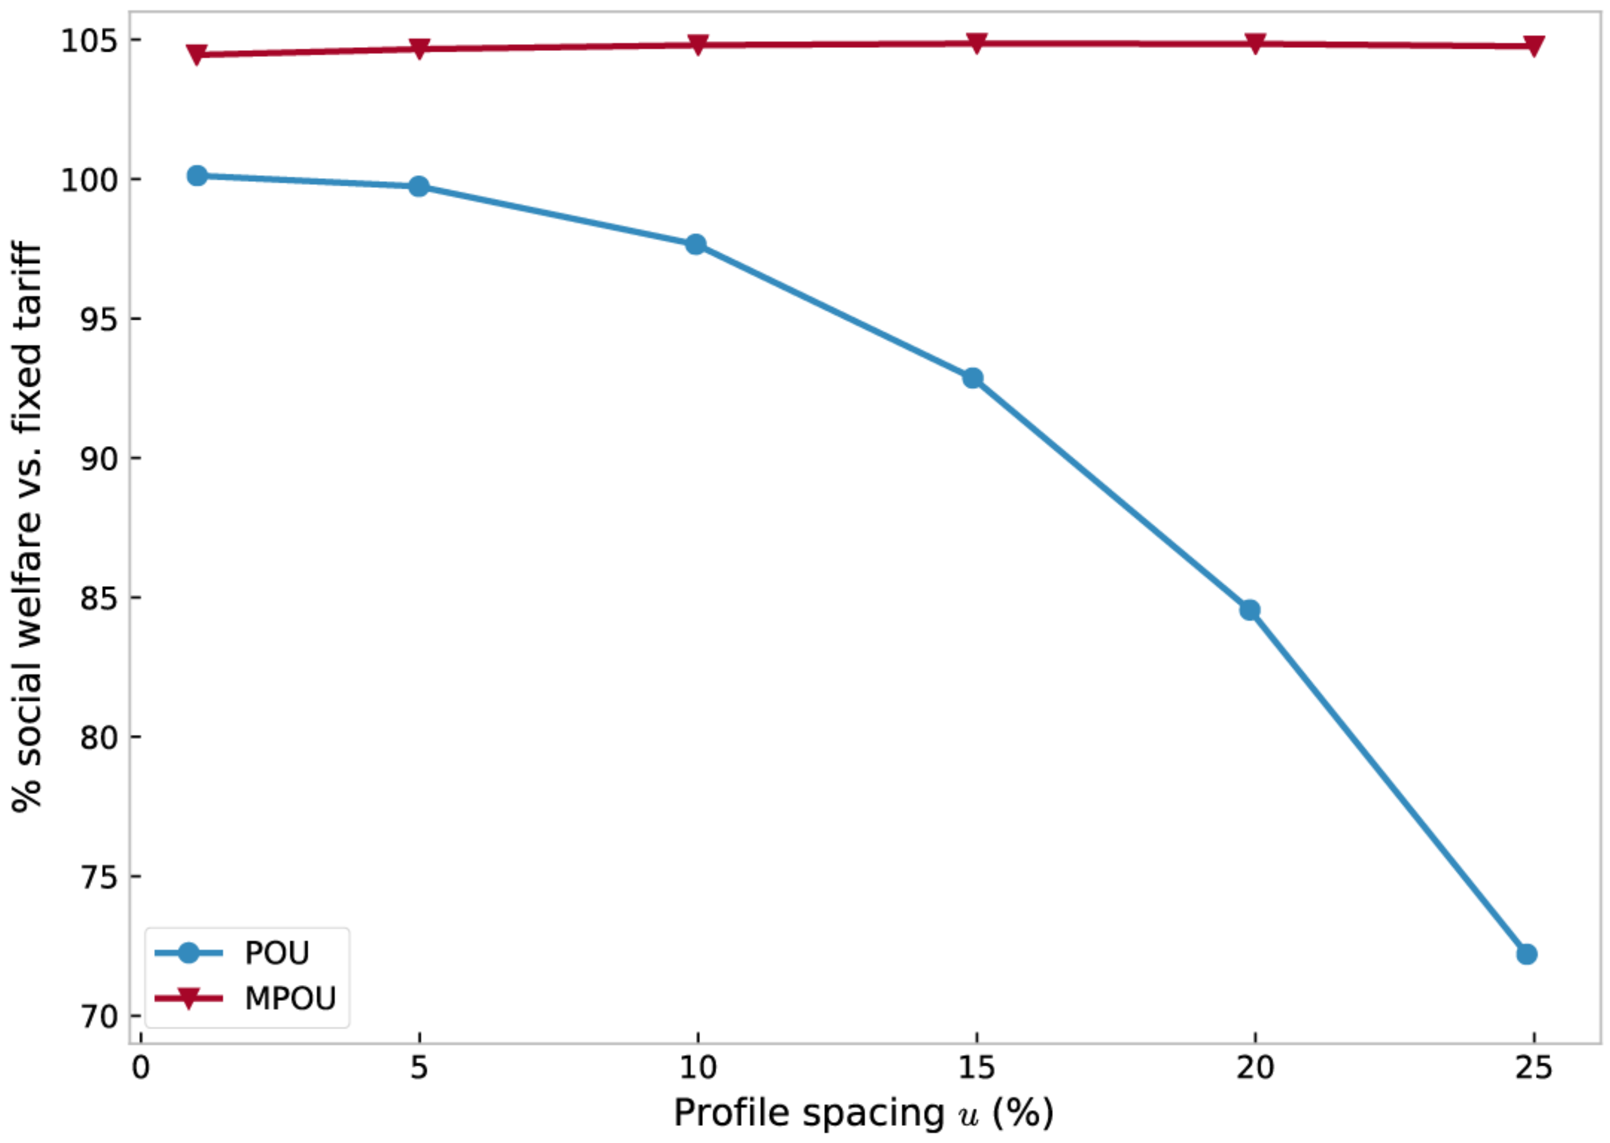

Contributions

- 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

Experiential Elicitation for Electricity Management Agents

P. and Boutilier, under review

Representing humans in electricity interactions

- Passive observation is not enough
  - Need preference elicitation (PE)

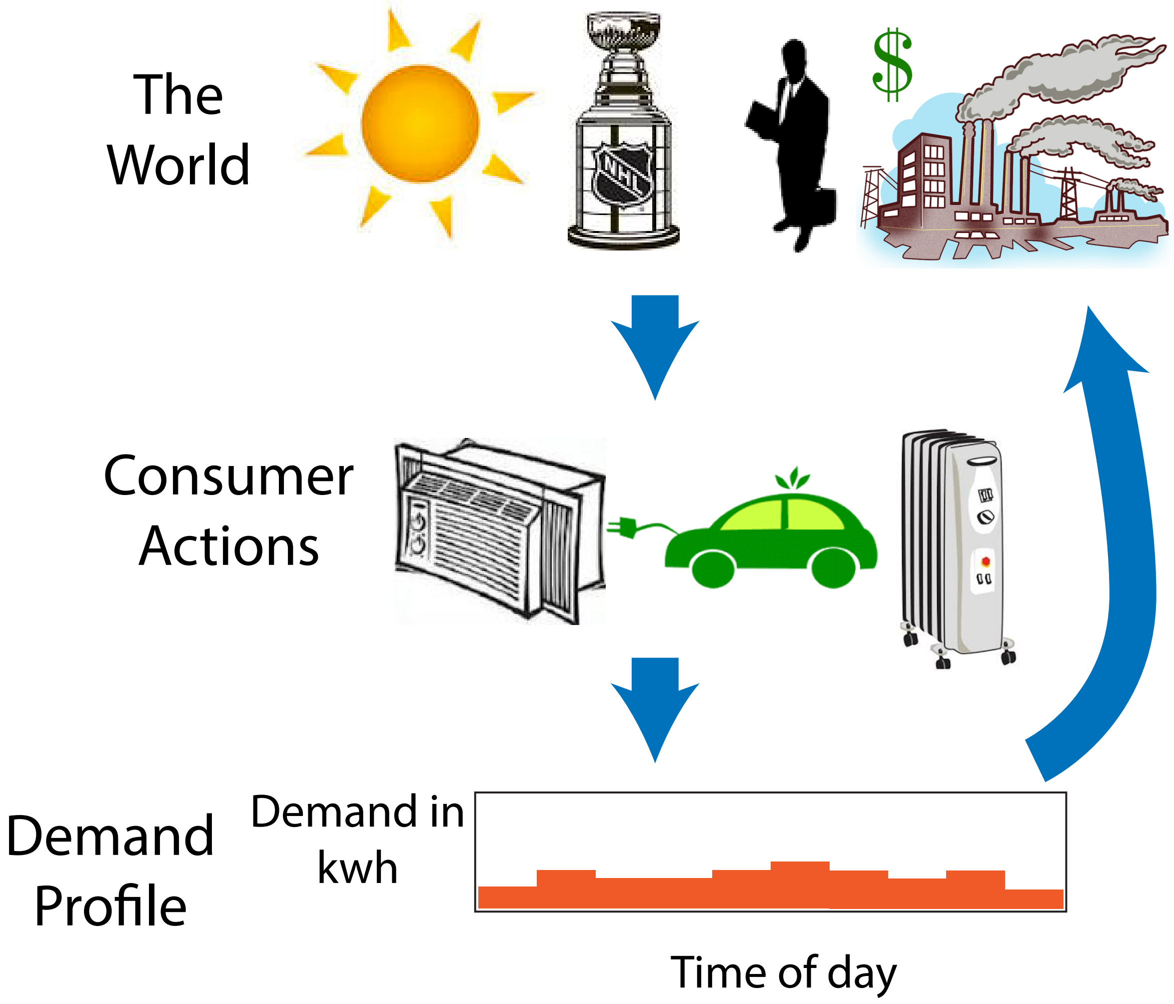

Dual process theory

- System 1 (intuition): fast, frequent, emotional
- System 2 (reasoning): slow, infrequent, logical

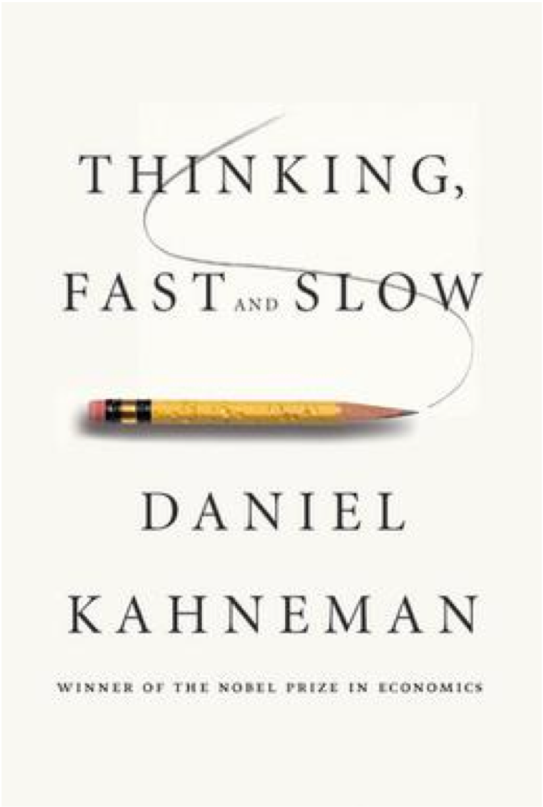

Kahneman (2011)

Dual process theory in AI

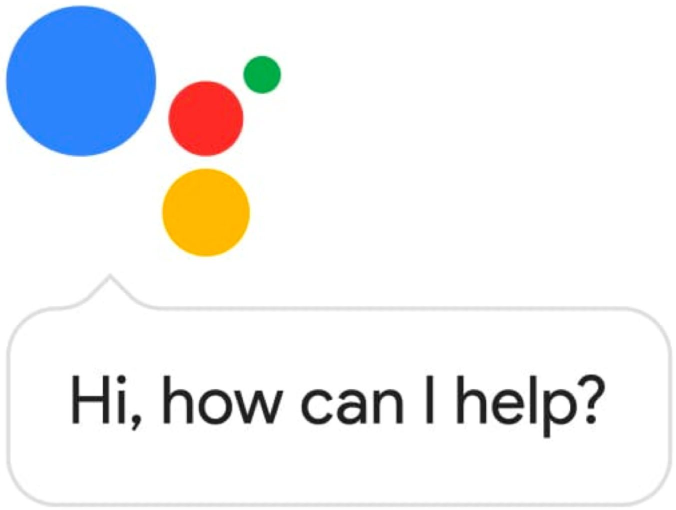

- System 1: intelligent assistants in daily tasks
  - **Energy use**
  - Personal assistants
- System 2: infrequent, impactful decisions
  - Security games
  - Apartment/house choice

- **Preference elicitation (PE) has focused on system 2 decisions**
  - Assume that a user can respond accurately regardless of query asked

Experiential elicitation

- Can better engage System 1 by asking questions about the current context
  - What are you willing to pay to decrease the temperature by 1°F for one hour?
- Beneficial side effect: delay queries until they are relevant
  - Increases discounted utility

Contributions

- Introduce a model of experiential elicitation
- Develop an instance of the model focused on electricity use
    - Introduce a relevant query type: the relative value query (RVQ)
    - Gaussian process-based models naturally accommodate RVQs
    - GP-based RVQ model performs well on synthetic data

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
- Experiential elicitation for electricity management agents
    - **Model of experiential elicitation**
    - Experiential elicitation with relative value queries
    - Empirical results

Markov decision processes

- Markov decision processes (MDPs): sequential decision-making under uncertainty
  - States S
  - Actions A
  - Transition function $\{P_{sa}\}$
  - Reward function r
  - Discount factor γ

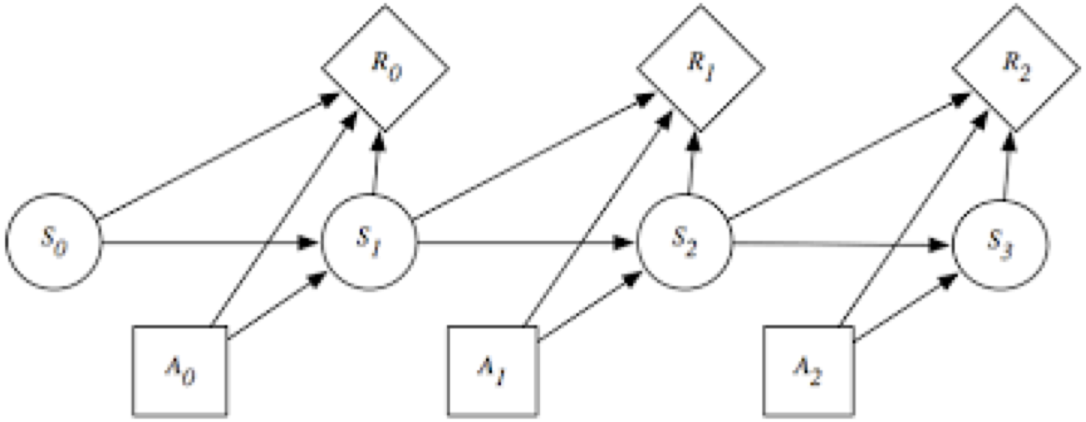

- Goal: find a policy that maximizes discounted reward over time
- Smart thermostat application: Shann and Seuken (2013, 2014)

Model of experiential elicitation

- MDP with unknown reward function
  - Agent knows R: distribution over reward functions
- Set of queries Q, each with:
  - Set of responses N_q
  - Response distribution D_q (with *known* dependence on state history and r)
- Query cost function C
  - Depends on state history and query

Model of experiential elicitation

- Repeat infinitely:
  - The agent asks the user zero or more queries
    - User responds according to response function
    - Agent incurs query cost
  - Agent chooses an action & state transitions according to transition function
- Goal: maximize reward minus query cost, w.r.t. to unknown reward function
- *Observation*: with risk neutral agent, can find optimal policy through MDP reduction

Relationship to other models

- *Theorem*: reducible to *partially-observable Markov decision process (POMDP)*
- Intuitively related to reinforcement learning

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
- Experiential elicitation for electricity management agents
  - Model of experiential elicitation ✓
  - **Experiential elicitation with relative value queries**
  - Empirical results

Relative value queries

- Relative value query (RVQ) asks the user about *difference* in utility between states
  - What are you willing to pay to decrease the temperature by 1°C for one hour?
  - Formally $(\boldsymbol{x}_0, \boldsymbol{x}_1)$, with response y: user's estimate of $r(\boldsymbol{x}_0) - r(\boldsymbol{x}_1)$
- *Observation*: relative value queries are enough to solve MDP
- Question: how do we make state value estimates given RVQs?

Gaussian processes (GPs)

- Universal function approximator
- Estimate at each point plus uncertainty info
- **Naturally accommodates RVQs**
  - Difference of Gaussian is Gaussian
- (Myopic) expected value of information (EVOI) estimation: sample query responses from GP posterior, average change in policy value across samples

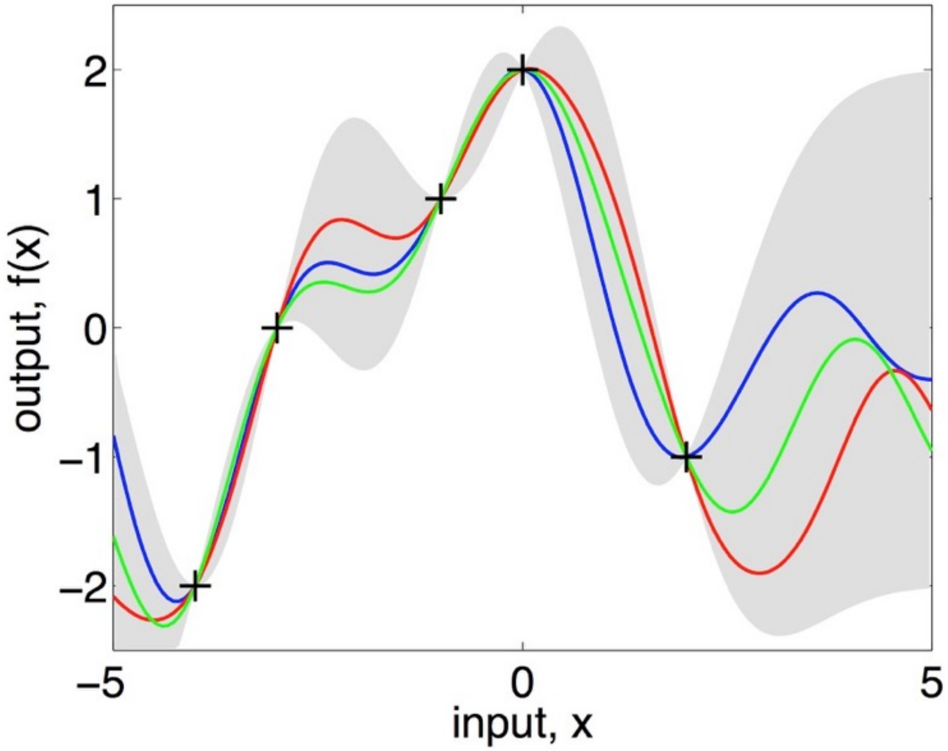

Rasmussen and Williams (2006)

Query response and cost models

- Use distance between query points and closest (discounted) historical state: $||\boldsymbol{x}_{q0} - \boldsymbol{x}_{q1}|| + \min_i (1+\delta)^{t-i}(||\boldsymbol{x}_{q0} - \boldsymbol{x}_i|| + ||\boldsymbol{x}_{q1} - \boldsymbol{x}_i||)$
  - Distance reduces response quality, increases cost

- Wednesday
- March 2018, 5 pm
- 20.5°C interior
- 17°C exterior

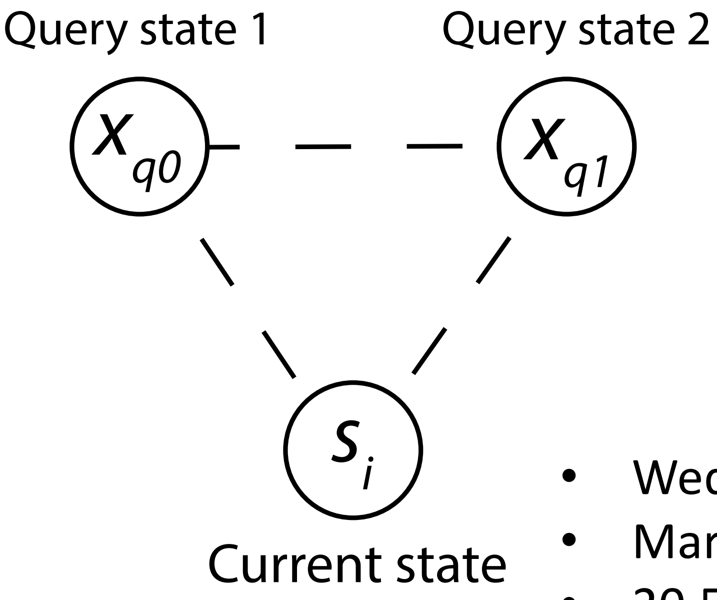

- Wednesday
- March 2018, 5 pm
- 21°C interior
- 17°C exterior

- Wednesday
- March 2018, 4 pm
- 20.5°C interior
- 17°C exterior

Outline

- Why study electricity in AI?
- Aligning predictability incentives between consumers and suppliers
- Experiential elicitation for electricity management agents
  - Model of experiential elicitation ✓
  - Experiential elicitation with relative value queries ✓
  - **Empirical results**

Experimental setup

- Synthetic data based on pecanstreet.org
- 3.5 million state MDP
- 100 households with different utility functions

Query strategies

- Meta-strategy: sample transition from believed best action and believed 2nd best action
    - Info is immediately relevant
    - Inexpensive and accurate

Results: reward accrued

- By time 1000, GP_EVOI is getting **75-80%** of OMNI
- Other strategies fail to outperform NULL
- Reason: GP_EVOI achieving same decision quality while paying less query cost

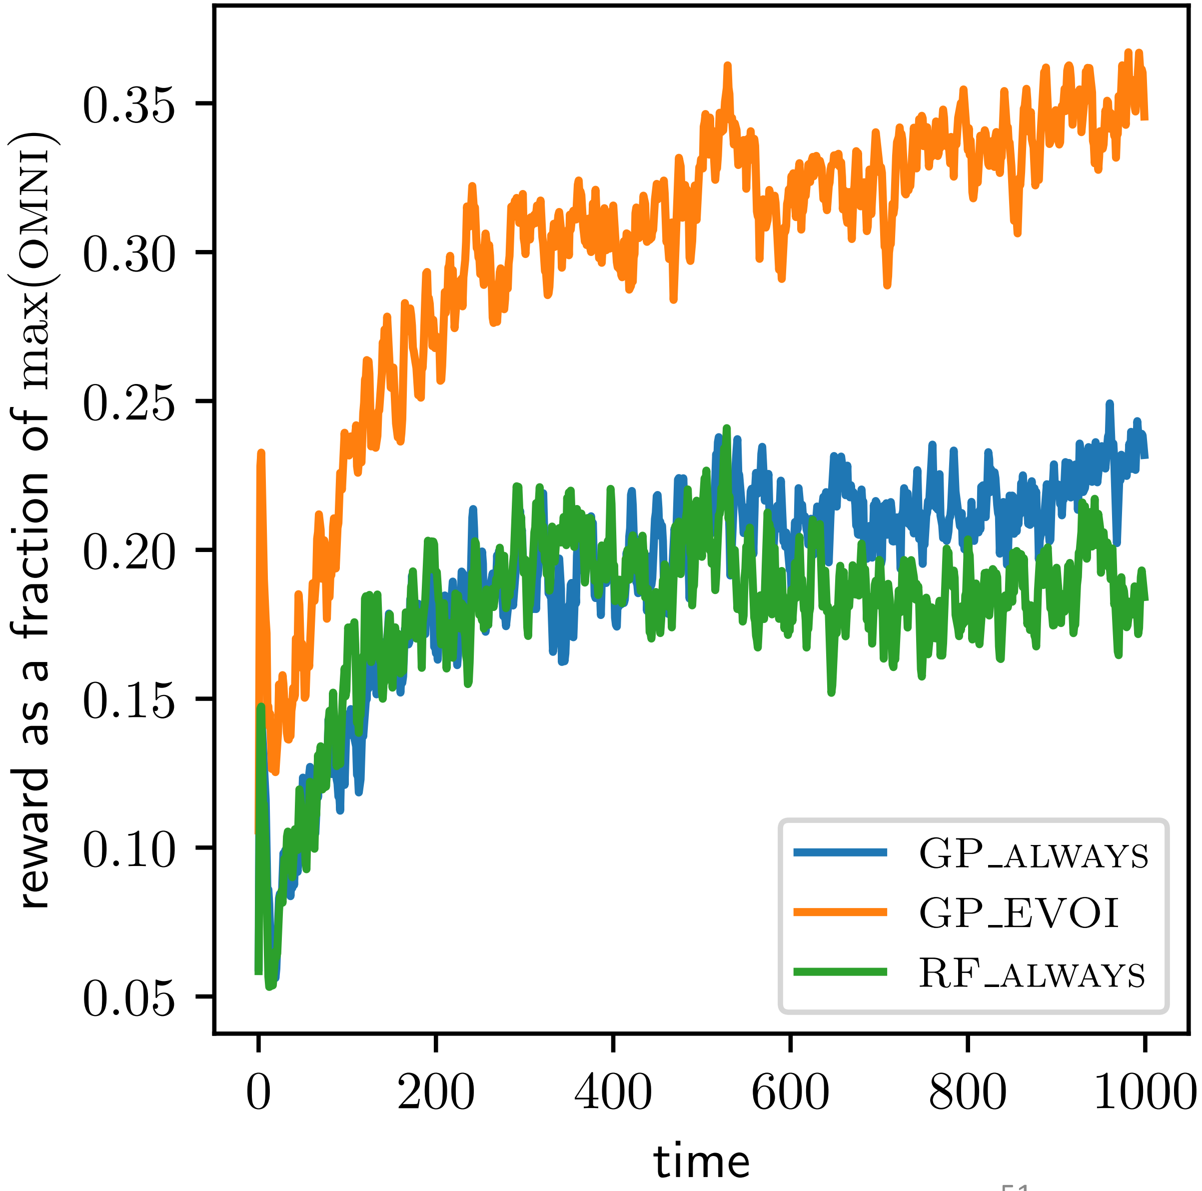

Results: decision quality

- GP_EVOI achieves **comparable decision quality**
  - **Asks many fewer queries**
    - ~44 by time 100
    - ~180 by time 1000

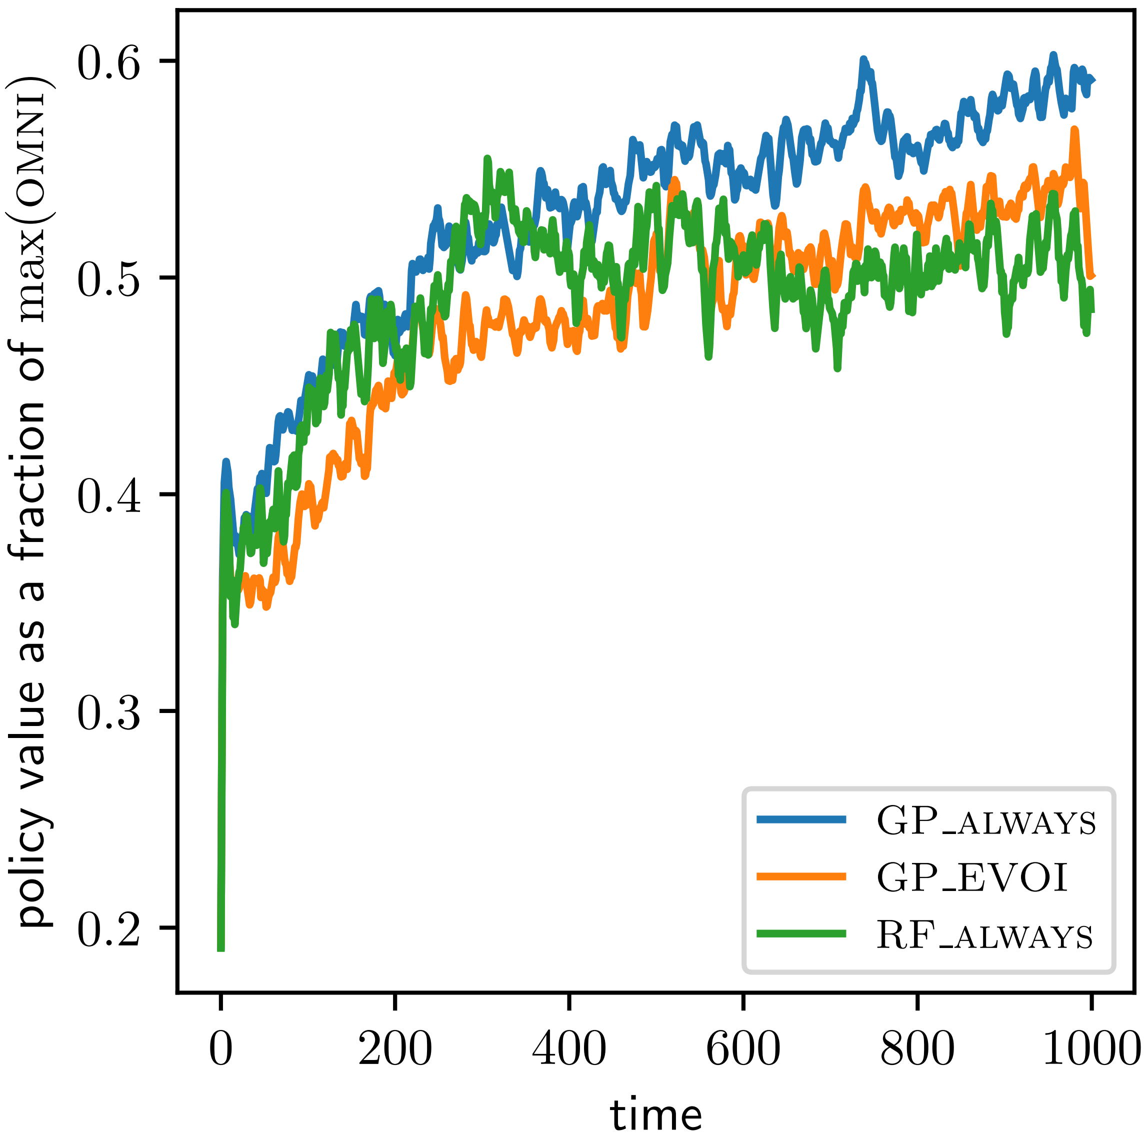

Results: query cost

- Fewer queries translates to much lower cost

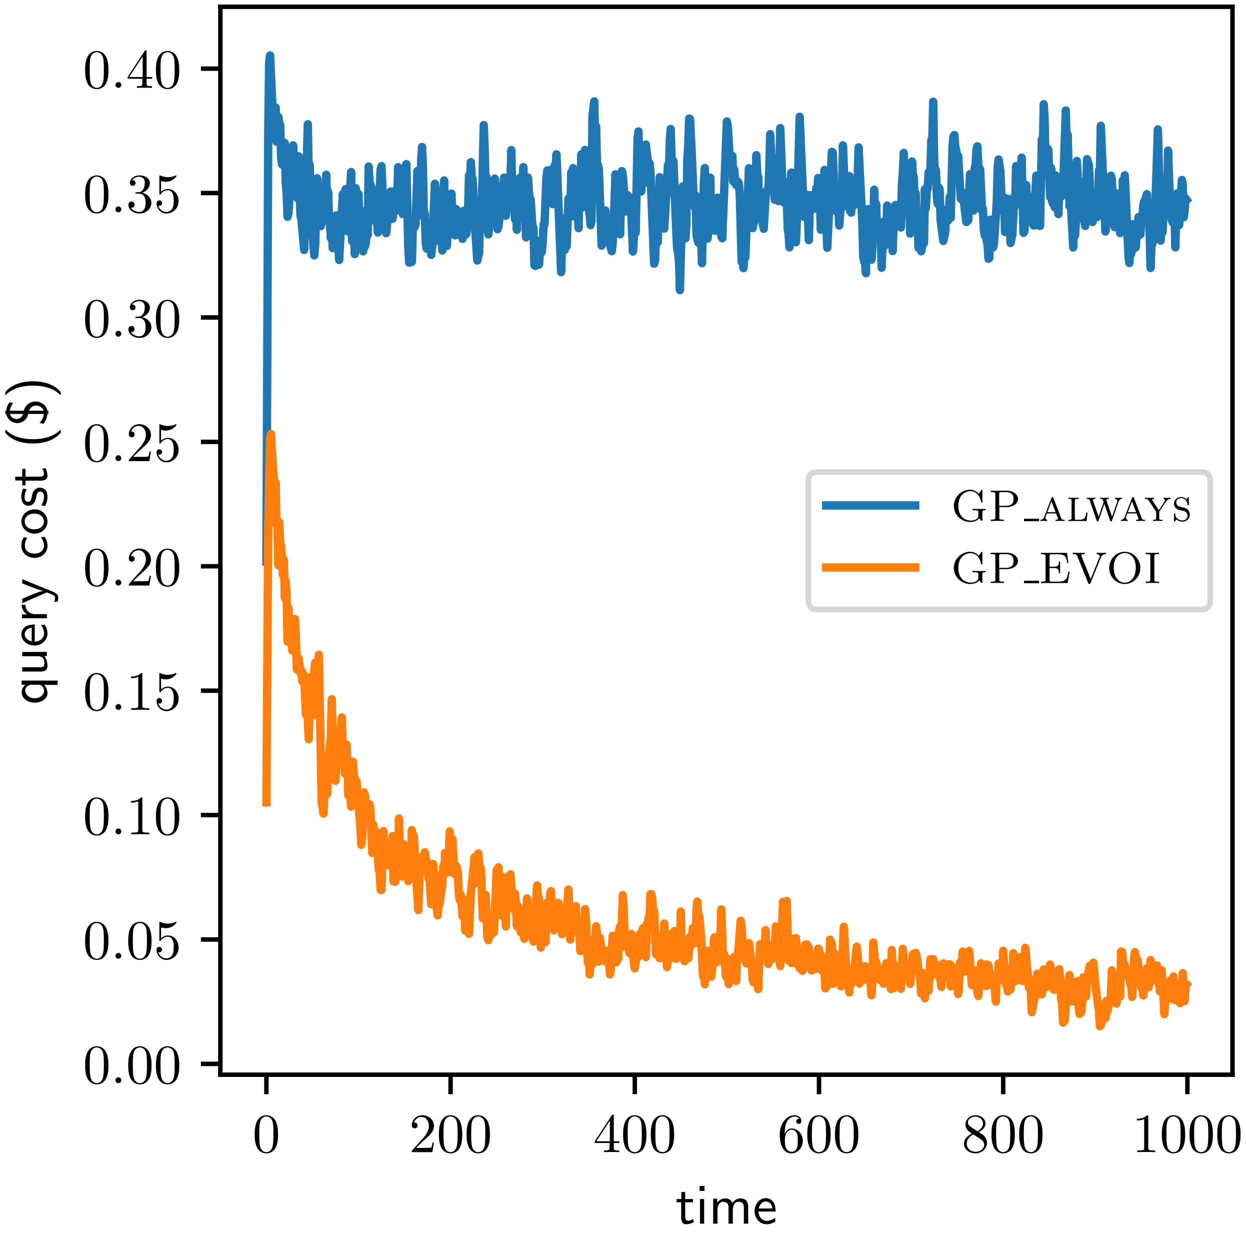

Contributions

- Introduce a model of experiential elicitation
- Develop an instance of the model focused on electricity use
  - Introduce a relevant query type: the relative value query (RVQ)
  - Gaussian process-based models naturally accommodate RVQs
  - GP-based RVQ model performs well on synthetic data

Future Work

Future work

- Strategic interactions between market design and preference elicitation
- POU games with correlated prediction errors
- Separating functions: more technical work? Wider applicability?
- Applications outside electricity: sharing of scarce resources
  - Cloud computing?
- Experiential elicitation
- More data?

Extra Slides

Making predictions from differences

- Related to rank learning
- Can convert any supervised learning alg.
  - Standard: train *f* on $(\boldsymbol{X}, \boldsymbol{Y})$ (matrices)
  - Differences: train *f* on $(\boldsymbol{X_0}, \boldsymbol{X_1}, \boldsymbol{Y})$ *and* $(\boldsymbol{X_1}, \boldsymbol{X_0}, -\boldsymbol{Y})$
- Problem: how to *estimate expected value of information (EVOI)?*

Query strategies

- Meta-strategy: sample transition from believed best action and believed 2nd best action
  - Info is immediately relevant
  - Inexpensive and accurate
- Query strategies:
  - RF_ALWAYS: Random forest-based model that always queries
  - GP_ALWAYS: GP-based model that always queries
  - GP_EVOI: GP-based model that queries *if query EVOI is higher than cost*
  - OMNI: take the best action (without querying)
  - NULL: take the null action (without querying)