Probably Approximately Correct Nash Equilibrium Learning

Kostas Margellos

Credit

- Joint work with Dr Filiberto Fele and Dr Luca Deori

Motivation: Electric vehicle charging

- How will the mass adoption of electric vehicles change the demand for electricity?
- Can controlling charging avoid the need to increase power generation capacity?
- Is there a practical way in which this can be achieved?
- How will renewable resources may be taken into account?

Motivation

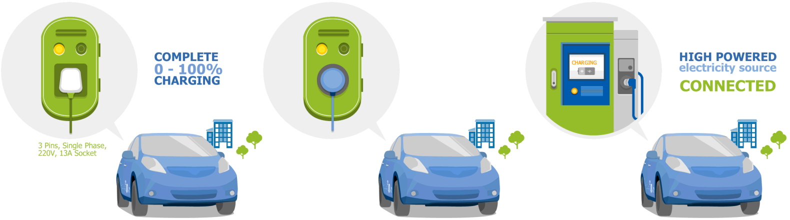

- Green car revolution – means to address the energy trilema
  - Green, cost efficient, reliable
  - Vehicles act like *virtual* storage devices
  - Not only they store but also defer their consumption in time
- In the UK: 3,500 in 2013, more than 90,000 today and the numbers keep increasing!

Motivation

- Rapid increase in the adoption of electric vehicles (EVs) in the UK
    - UK commitment to reduce emissions by 80% by 2050
    - Decreasing price of lithium ion batteries

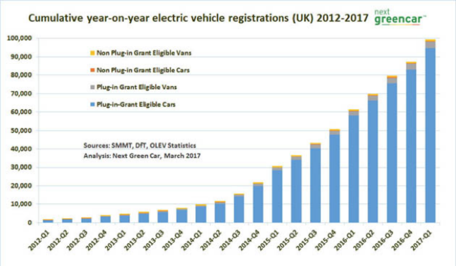

- Power consumption of various household appliances

Appliance	Power Consumption (W)
Washing Machine	700
Kettle	1800
Refrigerator	35
LCD TV	115
EV Charger	3500

Motivation

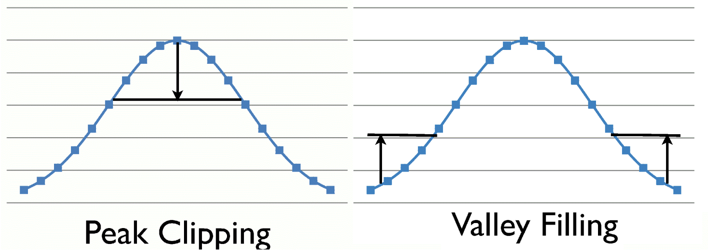

- Electric vehicles offer
    - Peak clipping, i.e. reduce peak demand by discharging at peak time instances (like storage)
    - Valley filling, i.e. charge when electricity price is lower
      $\Rightarrow$ cost savings

Electric vehicle charging control

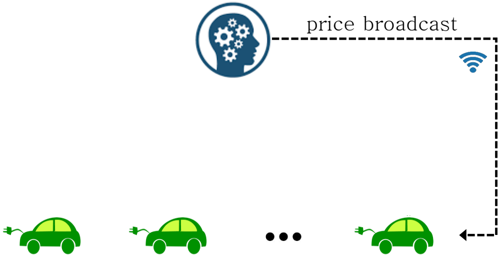

1. Aggregator / central authority sends a *price* incentive to vehicles

$$\text{price} = p\left(x_1 + \ldots + x_m\right) = p\left(\sum_i x_i\right),$$

where x_i is the consumption level of each vehicle / agent

Electric vehicle charging control

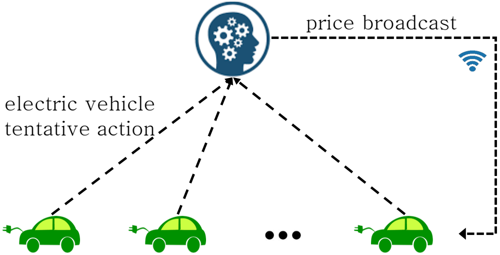

1. Aggregator / central authority sends a *price* incentive to vehicles

$$\begin{aligned} \text{price} &= p\left(x_1 + \ldots + x_m\right) \\ &= p\left(\sum_i x_i\right), \end{aligned}$$

where x_i is the consumption level of each vehicle / agent

2. Agents solve some local problem and broadcast an update x_i

Electric vehicle charging control

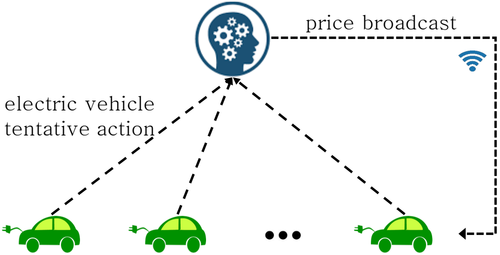

1. Aggregator / central authority sends a *price* incentive to vehicles

$$\text{price} = p\Big(x_1 + \ldots + x_m\Big) = p\Big(\sum_i x_i\Big),$$

where x_i is the consumption level of each vehicle / agent

2. Agents solve some local problem and broadcast an update x_i
3. A new price is calculated and process is repeated

Challenges

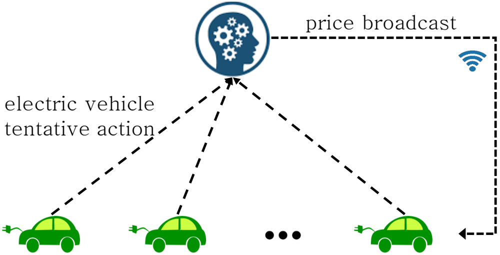

- Vehicles act as non-cooperative entities ⇒ multi-agent game
- What happens if price is uncertain? ⇒ price volatility

1. Data driven Nash equilibrium computation
   ⇒ Equilibria become random variables
2. Nash equilibrium efficiency
   ⇒ How far from the social welfare optimum?

Data driven Nash equilibrium computation

Non-cooperative game

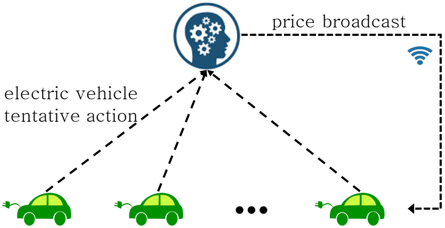

Agents' description

Cost function: $\sum_t x_{it} \times p_t\left(\sum_j x_{jt}\right)$ [quantity × price]

Constraints: $\sum_t x_{it} = E_i,$ [prescribed charging level]

$x_{it} \in [\underline{x}_{it}, \overline{x}_{it}]$, for all t [consumption limits]

Non-cooperative game

- Agents are selfish, non-cooperative entities
- Interested in minimizing some cost when other agents' strategies are fixed

$$J_i(\textcolor{blue}{x_i}, x_{-i}) = f_i(\textcolor{blue}{x_i}, x_{-i}) + \textcolor{red}{\max_{k=1,\ldots,M}} g(\textcolor{blue}{x_i}, x_{-i}, \textcolor{red}{\theta_k})$$

- Price is subject to volatility $\Rightarrow$ Represent uncertainty by means of scenarios!

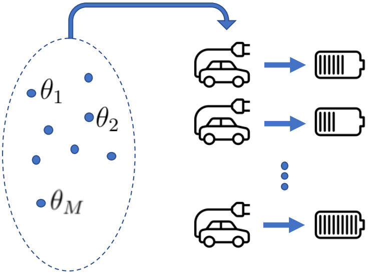

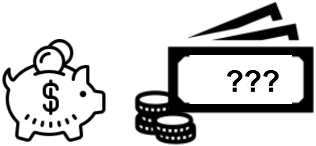

- Uncertainty on electricity price
- Historical data available

Non-cooperative game

- Agents are selfish, non-cooperative entities
- Interested in minimizing some cost when other agents' strategies are fixed

$$J_i(x_i, x_{-i}) = f_i(x_i, x_{-i}) + \max_{k=1,\dots,M} g(x_i, x_{-i}, \theta_k)$$

- Price is subject to volatility $\Rightarrow$ Represent uncertainty by means of scenarios!

Nash equilibrium

A set of agents' strategies $(\bar{x}_i, \bar{x}_{-i})$ forms a Nash equilibrium if for all i

$$J_i(\bar{x}_i, \bar{x}_{-i}) \leq J_i(x_i, \bar{x}_{-i}), \text{ for all } x_i \in X_i,$$

where X_i is agent's i local constraint set.
No agent can improve her cost when other agents' strategies are fixed

Non-cooperative game

- Nash equilibrium $\bar{x}$ is a random variable
- How likely is it to remain unchanged when a new uncertainty realization is encountered?

$$\mathbb{P}^M\Big\{\theta_1,\ldots,\theta_M:\ \mathbb{P}\{\delta:\ \bar{x}=\bar{x}^+\} > 1-\epsilon\Big\} \geq 1-\beta$$

- Probably approximately correct Nash equilibrium learning

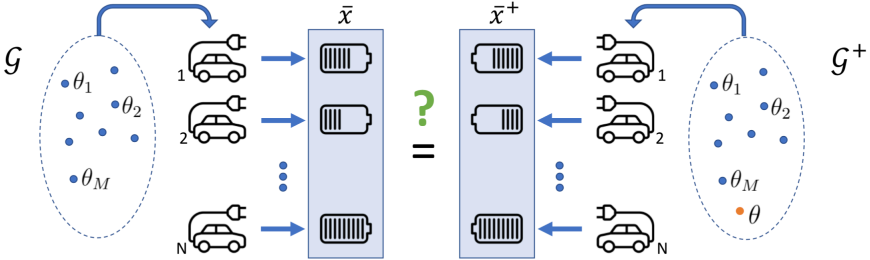

Non-cooperative game

Probably approximately correct Nash equilibrium learning

Fix $\beta \in (0,1)$, and consider the function $\epsilon(\cdot)$ such that

$$\epsilon(M) = 1 \text{ and } \sum_{k=0}^{M-1} \binom{M}{k} (1-\epsilon(k))^{M-k} = \beta$$

We then have $\mathbb{P}^M\Big\{\theta_1,\ldots,\theta_M : \; \mathbb{P}\{\theta : \; \bar{x} = \bar{x}^+\} > 1-\epsilon(d)\Big\} \geq 1-\beta$,
where d : sample compression, i.e., $\bar{x}_d = \bar{x}_M$.

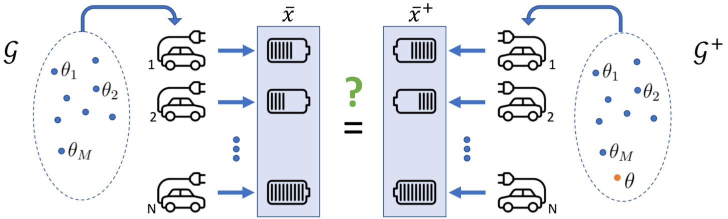

… sketch of the proof

1. Agents' problem reformulation
2. Nash equilibria as solutions of variational inequalities (VIs)
3. Using the "scenario approach"
   a la Campi, Calafiore, Garatti, Prandini, …

Sketch of the proof

STEP 1:

Epigraphic reformulation for agent i

$$\begin{aligned} \text{minimize } & f_i(x_i, \bar{x}_{-i}) + \gamma \\ \text{subject to } & x_i \in X_i \\ & g(x_i, \bar{x}_{-i}, \theta_k) \leq \gamma \text{ for all } k \end{aligned}$$

STEP 2:

- Equilibria can be characterized as solutions to VIs
- By VI sensitivity: constraint satisfaction implies equilibrium insensitivity

$$\text{if } \mathbb{P}\{\theta : \; g(\bar{x}, \theta) \leq \bar{\gamma}\} \geq 1 - \epsilon(\bar{d})$$

$$\text{... then } \mathbb{P}\{\theta : \; \bar{x} = \bar{x}^+\} \geq 1 - \epsilon(\bar{d})$$

Sketch of the proof (cont'd)

STEP 3:

Data based program

$$\begin{aligned} &\text{minimize } f_i(x_i, \bar{x}_{-i}) + \gamma \\ &\text{subject to} \\ &\qquad x_i \in X_i \\ &\qquad g(x_i, \bar{x}_{-i}, \theta_k) \leq \gamma, \forall k \end{aligned} \rightarrow (\bar{x}, \bar{\gamma})$$

Robust program

$$\begin{aligned} &\text{minimize } f_i(x_i, \bar{x}_{-i}) + \gamma \\ &\text{subject to} \\ &\qquad x_i \in X_i \\ &\qquad g(x_i, \bar{x}_{-i}, \theta) \leq \gamma, \forall \theta \in \Theta \end{aligned}$$

- What happens for a new θ $\Leftrightarrow$ Is $\bar{x}$ feasible for the robust program?
- Is this true for any $\theta_1, \ldots, \theta_M$?

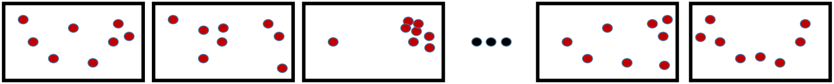

Sketch of the proof (cont'd)

STEP 3:

Data based program

$$\begin{aligned} &\text{minimize } f_i(x_i, \bar{x}_{-i}) + \gamma \\ &\text{subject to} \qquad\qquad \rightarrow (\bar{x}, \bar{\gamma}) \\ &\qquad x_i \in X_i \\ &\qquad g(x_i, \bar{x}_{-i}, \theta_k) \leq \gamma, \forall k \end{aligned}$$

Robust program

$$\begin{aligned} &\text{minimize } f_i(x_i, \bar{x}_{-i}) + \gamma \\ &\text{subject to} \\ &\qquad x_i \in X_i \\ &\qquad g(x_i, \bar{x}_{-i}, \theta) \leq \gamma, \forall \theta \in \Theta \end{aligned}$$

- What happens for a new θ $\Leftrightarrow$ Is $\bar{x}$ feasible for the robust program?
- Is this true for any $\theta_1, \ldots, \theta_M$?

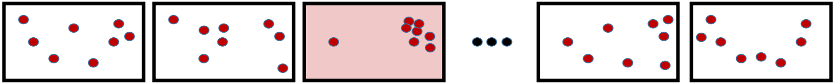

Sketch of the proof (cont'd)

Data based program

$$\begin{aligned} \text{minimize } & f_i(x_i, \bar{x}_{-i}) + \gamma \\ \text{subject to } & \qquad\qquad \rightarrow (\bar{x}, \bar{\gamma}) \\ & x_i \in X_i \\ & g(x_i, \bar{x}_{-i}, \theta_k) \leq \gamma, \forall k \end{aligned}$$

Robust program

$$\begin{aligned} \text{minimize } & f_i(x_i, \bar{x}_{-i}) + \gamma \\ \text{subject to } & \\ & x_i \in X_i \\ & g(x_i, \bar{x}_{-i}, \theta) \leq \gamma, \forall \theta \in \Theta \end{aligned}$$

Feasibility link [Campi & Garatti, 2008] and [KM, Prandini & Lygeros, 2015]

Fix $\beta \in (0, 1)$. With confidence $\geq 1 - \beta$, $\bar{x}$ is feasible for the robust program with probability $\geq 1 - \epsilon(d)$, i.e.

$$\mathbb{P}\Big(\theta : \; g(\bar{x}_i, \bar{x}_{-1}, \theta) \leq \bar{\gamma}\Big) > 1 - \epsilon(d) \text{ with prob. } \geq 1 - \beta$$

Sketch of the proof (cont'd)

Feasibility link

Fix $\beta \in (0,1)$. With confidence $\geq 1-\beta$, $\bar{x}$ is feasible for the robust program with probability $\geq 1-\epsilon(d)$, i.e.

$$\mathbb{P}\Big(\theta : \; g(\bar{x}_i, \bar{x}_{-1}, \theta) \leq \bar{\gamma}\Big) > 1-\epsilon(d) \text{ with prob. } \geq 1-\beta$$

- On which parameters does ϵ depends on?

$$\epsilon(M) = 1 \text{ and } \sum_{k=0}^{M-1} \binom{M}{k} (1-\epsilon(k))^{M-k} = \beta$$

$$\text{for } k = d: \; \epsilon(d) \approx \frac{2}{M}\Big(d + \ln\frac{1}{\beta}\Big)$$

  - Logarithmic in β: $1-\beta$ can be set close to one
  - Linear in M^{-1}: The more data the better the result
  - Linear in d: cardinality of sample compression

Sketch of the proof (cont'd)

Cardinality of sample compression d

- For convex agents' objective functions and constraint sets

$$d \leq \text{\# decision variables} = \text{\# agents} \times \text{\# time-steps}$$

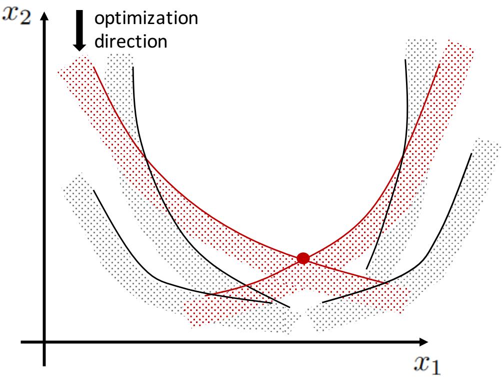

Non-cooperative game – the convex case

Probably approximately correct Nash equilibrium learning

Fix $\beta \in (0,1)$, and consider the function $\epsilon(\cdot)$ such that

$$\epsilon(M) = 1 \text{ and } \sum_{k=0}^{M-1} \binom{M}{k} (1-\epsilon(k))^{M-k} = \beta$$

We then have $\mathbb{P}^M \Big\{ \theta_1, \dots, \theta_M : \ \mathbb{P}\{\theta : \ \bar{x} = \bar{x}^+\} > 1 - \epsilon(d) \Big\} \geq 1 - \beta$,
where $d =$ # agents $\times$ # time-steps.

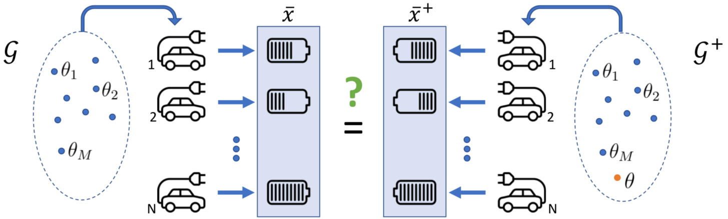

Simulation results

- Agents' cost function

$$f_i(x_i, x_{-i}) + \max_k g(x_i, x_{-i}, \theta_k)$$
$$= x_i^\top (A_0 \sigma(x) + b_0) + \max_k \sigma(x)^\top (A_k \sigma(x) + b_k)$$

where

  - A_k: diagonal with entries from log-normal distribution
  - b_k: entries from uniform distribution

- Agents' constraint set

$$X_i = \{x_i : \sum_t x_{it} = E_i,\ x_{it} \in [\underline{x}_{it}, \overline{x}_{it}],\ \forall t\}$$

i.e., energy and charge rate limits

Simulation results

- Probability of Nash equilibrium altering (constraint violation)

d	4	6	7	9
Empirical [%]	0.98	1.09	1.26	1.33
Theoretical [%]	8.06	9.76	10.55	12.06

- Valley filling behaviour: Charging when price is low

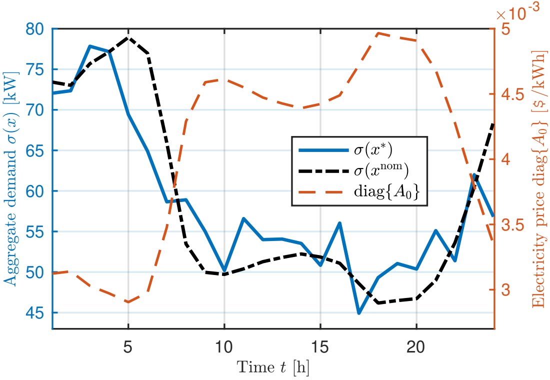

Nash equilibrium efficiency

Nash equilibrium efficiency

Recall the Nash equilibrium definition

A vector of vehicles' charging strategies is a Nash equilibrium if

$$J_i(\bar{x}_i, \bar{x}_{-i}) \leq J_i(x_i, \bar{x}_{-i}) \text{ for all } x_i \in X_i$$

Assumptions

1. For simplicity assume $J_i = g$, i.e., common across agents
2. Price is deterministic and affine in aggregate demand

How far are Nash equilibria from social welfare optima?

Nash equilibria as social optima

- Denote the objective function of vehicle i by

$$g(x_i, x_{-i}) = \sum_t x_{it}\Big(p_t \sum_{j, j \neq i} x_{jt} + p_t x_{it}\Big)$$

- For each i, consider the penalty term

$$g_a(x_i) = \sum_t p_t (x_{it})^2$$

Nash equilibrium as social optimum

The optimal solution of

$$\text{minimize} \ \sum_i g(x_i, x_{-i}) + g_a(x_i)$$

$$\text{subject to: } x_i \in X_i \text{ for all } i$$

is a Nash equilibrium of the multi-vehicle game

Nash equilibria as social optima

Nash equilibrium

The optimal solution of

$$\text{minimize} \sum_i g(x_i, x_{-i}) + g_a(x_i)$$
$$\text{subject to: } x_i \in X_i \text{ for all } i$$

is a Nash equilibrium of the multi-vehicle game.

- The penalty term makes the optimization program strictly convex
- It admits a unique minimizer $\Rightarrow$ unique Nash equilibrium.

Social optima vs Nash equilibria

Social optimum

$$\text{minimize} \quad \sum_i g(x_i, x_{-i})$$

$$\text{subject to:} \quad x_i \in X_i \text{ for all } i$$

Nash equilibrium

$$\text{minimize} \quad \sum_i g(x_i, x_{-i}) + g_a(x_i)$$

$$\text{subject to:} \quad x_i \in X_i \text{ for all } i$$

- Both of them are obtained as solutions to optimization programs
- Nash equilibrium is optimum for the a problem which trades between total cost and penalty term
- Penalty term acts like variance

Simulation results

- "Valley filling" property for both cases
- Albeit different solutions; is this always the case? (see vistas)
  - Nash equilibrium $\neq$ social optimum
  - Nash equilibrium $=$ social optimum (of another problem)

Case study for 5 vehicles

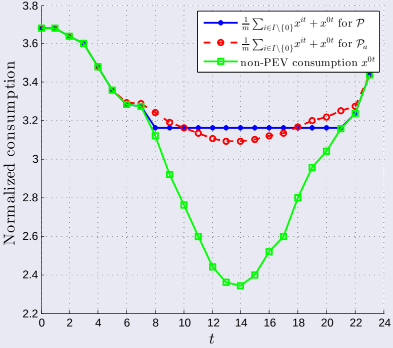

From non-cooperative to cooperative ...

- What if vehicles were "anarchists", interested in individual objectives?
- What is the "price of anarchy"[1], i.e.
  relative difference between social welfare optimum and Nash value?
- One can't do much ... but if it is many of them

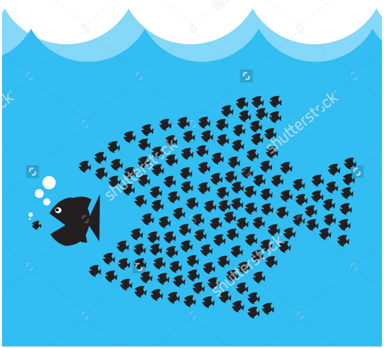

[1] Koutsoupias & Papadimitriou, 1999

From non-cooperative to cooperative ...

Social optimum $x^\star$

$$\text{minimize} \sum_i g(x_i, x_{-i})$$

subject to: $x_i \in X_i$ for all i

Nash equilibrium $x^\star$

$$\text{minimize} \sum_i g(x_i, x_{-i}) + g_a(x_i)$$

subject to: $x_i \in X_i$ for all i

- Let's compare the optimal values

$$J^m(x^\star) = \sum_i g(x_i^\star, x_{-i}^\star)$$

$$J^m(x^\star) = \sum_i g(x_i^\star, x_{-i}^\star)$$

Vehicles are heterogeneous, and heterogeneity is modelled assuming certain parameters in their constraints are randomly chosen from some distribution

Price of anarchy

- As the population size m grows, the social optimum tends to the value of the game, i.e.

$$\lim_{m \to \infty} \frac{J^m(x^\star)}{J^m(x^\star)} = 1$$

for almost all realizations of the (random) heterogeneity parameters.

- Price of anarchy tends to zero!

- Vehicles tend to cooperate even though they are selfish individuals

Price of anarchy

- "Valley filling" property for both cases
- Nash equilibrium $\neq$ social optimum. Is this always the case?
- Not as m increases ... they tend to coincide

Case study for $m = 5$ and $m = 100$ vehicles

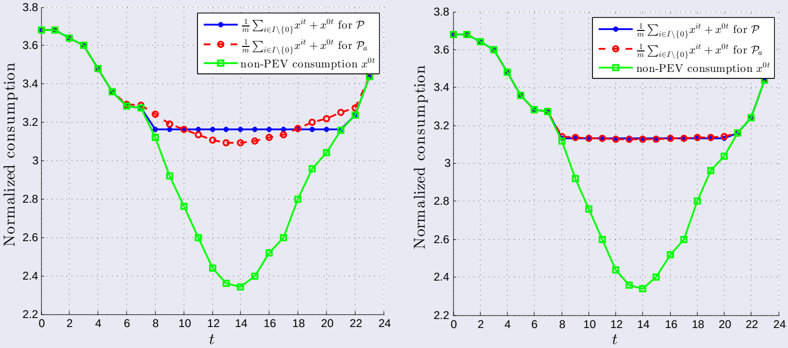

Summary

- Data driven Nash equilibrium computation
    - Probabilistic equilibrium sensitivity
    - *A priori* robustness certificates
- Nash equilibrium efficiency
    - How far are Nash equilibria from social optima?
    - Price of anarchy characterization
- Other results – future work
    - Decentralized equilibrium computation via best response algorithms
    - *A posteriori* robustness certificates

References

Margellos, Prandini & Lygeros (2016)
On the Connection Between Compression Learning and Scenario Based Single-Stage and Cascading Optimization Problems
IEEE Transactions on Automatic Control, 60(10), 2716-2721.

Deori, Margellos & Prandini (2018)
Price of anarchy in electric vehicle charging control games: When Nash equilibria achieve social welfare
Automatica, 96(10), 150-158.

Fele & Margellos (2020)
Probably approximately correct Nash equilibrium learning
IEEE Transactions on Automatic Control, to appear.

Romao, Margellos & Papachristodoulou (2021)
On the exact feasibility of convex scenario programs with discarded constraints
IEEE Transactions on Automatic Control, under review.

Thank you for your attention!
Questions?

Contact at:
kostas.margellos@eng.ox.ac.uk