Managing electricity consumption by providing dynamic tariff incentives

To make a better use of renewable energy

Gilles Stoltz

ICML 2019 paper with: Margaux Brégère (PhD student),
Pierre Gaillard (Inria Paris), Yannig Goude (EDF R&D)

Context / Motivation

Industrial partner: EDF

Aim: maintain balance between production and consumption

Current solution: forecast consumption and adapt production

Prospective solution: encourage/discourage consumption by dynamically setting prices

Advantage: better use of renewable energy

Data set

"SmartMeter Energy Consumption Data in London Households"

https://data.london.gov.uk/dataset/smartmeter-energy-use-data-in-london-households

Public dataset – by UK Power Networks (subbranch of EDF)

Individual consumptions at half-hourly frequency in year 2013

About 1,000 customers with tariff incentives

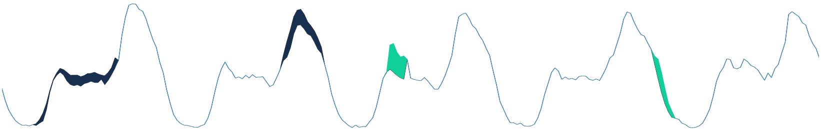

K=3 tariffs: Low (L), Normal (N), High (H)

Tariff incentives indeed have an impact on consumption!

Quantitative approach $\longrightarrow$ Design a specific pricing policy

Exploration–exploitation dilemma: Need to simultaneously

- Discover the behaviors of customers (= exploration)
- Optimize incentives sent (= exploitation)

Of course behaviors may change over time!

Bandit monitoring

- We only observe the outcome of the tariff(s) picked
- Not of what would have happened with different choices

Aim: Design pricing strategies (with theoretical guarantees)
Test them on data

Issue: Historical data obtained for a given sequence of choices

Solution: Construct first a realistic data generator

Methodology

1. Estimate a model / Build a data generator
   based on 2013 data (consumption + context)

2. Get historical contexts for 2014 + January 2015

   Generate realistic consumptions

   2.1. Use normal tariff only in 2014

   2.2. Then use a machine learning algorithm for January 2015
   and pick among all K tariff levels

Modeling of the consumption → known and effective methodology designed by EDF

Population assumed to be homogeneous (as a first approach)

(Mean) consumption Y depends on context $x_t \in \mathbb{R}^d$

Context = temperature, season, day of the week, hour of the day, etc.

Also depends on tariff $k \in \{1, \ldots, K\}$

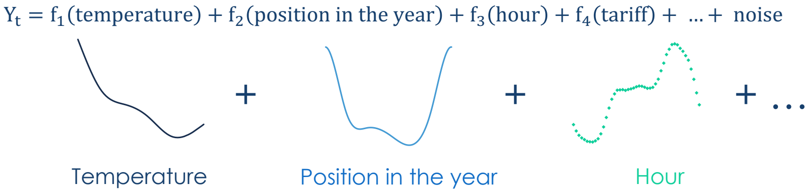

If single tariff k picked: $$Y_{t,k} = \gamma_k + \sum_{i=1}^{d} f_i(x_{t,i}) + \text{noise}$$

First model/generator: With a single tariff at any given time

Parametric model for $Y_{t,k} = \gamma_k + \sum_{i=1}^{d} f_i(x_{t,i}) + \text{noise}$

given by a

Generalized additive model (Wood, 2006) based on so-called cubic splines

$$Y_{t,k} = \gamma_k + \beta^{\mathsf{T}} \varphi(x_t) + \varepsilon_{t,k}$$

where β and γ_k are unknown, but $\varphi(x_t)$ is known

$\longrightarrow$ Need to extend this modeling to K tariffs

Final model/generator: With various tariffs at the same time

If tariffs $\{1, \dots, K\}$ are distributed in shares $p = (p_1, \dots, p_K)$

Then (cf. homogeneous population), mean consumption:

$$Y_{t,p} = \sum_{k=1}^{K} p_k Y_{t,k} = \sum_{k=1}^{K} p_k \big(\beta^\mathsf{T} \varphi(x_t) + \gamma_k + \varepsilon_{t,k}\big)$$
$$= \theta^\mathsf{T} \phi(x_t, p) + p^\mathsf{T} \varepsilon_t$$

with θ unknown, but $\phi(x_t, p)$ is known (and linear in p)

Noise: ε_t iid vectors, $\mathbb{E}[\varepsilon_t] = 0$, sub-Gaussian

$\Gamma = \mathrm{Var}(\varepsilon_t)$ estimated on data

In-sample performance: good, $r^2 = 92\%$ and MAPE $= 8.82\%$

Methodology (reminder)

1. Estimate a model / Build a data generator based on 2013 data (consumption + context)

2. Get historical contexts for 2014 + January 2015

   Generate realistic consumptions

   2.1. Use normal tariff only in 2014

   2.2. Then use a machine learning algorithm for January 2015 and pick among all K tariff levels

New machine-learning problem defined:
Target tracking for contextual bandits

Known parameters

- K tariffs
- Context set $\mathcal{X}$
- Transfer function $\phi : \mathcal{X} \times \mathcal{P} \to \mathbb{R}^m$
- Bound C on consumptions Y

Unknown parameters
(They model the behaviors)

- Coefficients $\theta \in \mathbb{R}^m$
- Covariance matrix $\Gamma = \mathrm{Var}(\varepsilon_t)$

For each round $t = 1, 2, \ldots$

1. Observe a context $x_t \in \mathcal{X}$ and a target $c_t \in [0, C]$
2. Choose an allocation of tariffs $p_t = (p_{t,1}, \ldots, p_{t,K})$
3. Observe a mean consumption $Y_{t,p_t} = \theta^{\mathsf{T}} \phi(x_t, p_t) + p_t^{\mathsf{T}} \varepsilon_t$
4. Encounter an error $(Y_{t,p_t} - c_t)^2$

$\longrightarrow$ Algorithm constructed: based on a LinUCB-approach

Realistic simulations

$K = 3$ tariff levels

Consider attainable targets: $\theta^{\mathsf{T}}\phi(x_t, 1) \leqslant c_t \leqslant \theta^{\mathsf{T}}\phi(x_t, 3)$

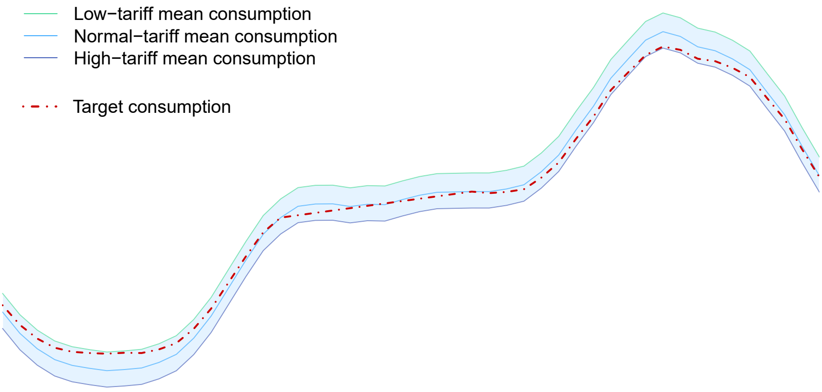

Aim: smooth out consumption

Reminder of the experiment design → provider changing its policy

- Pick the "normal" tariff for 1 year, i.e., $p_t = (0, 1, 0)$
- Then start picking different allocations with at most 2 tariffs (either 1+2 or 2+3)

Repeat this 200 times

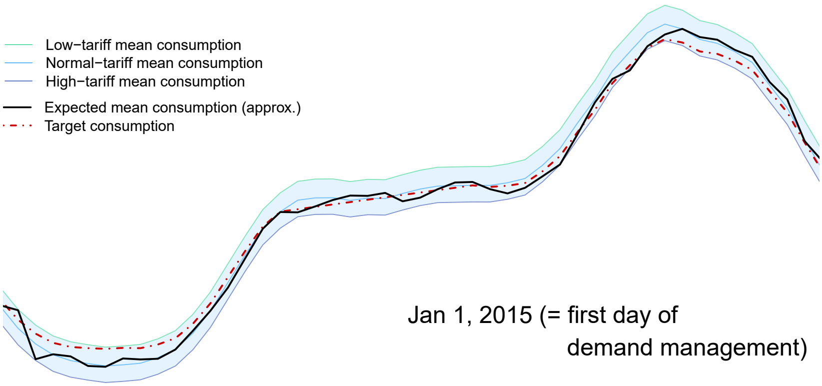

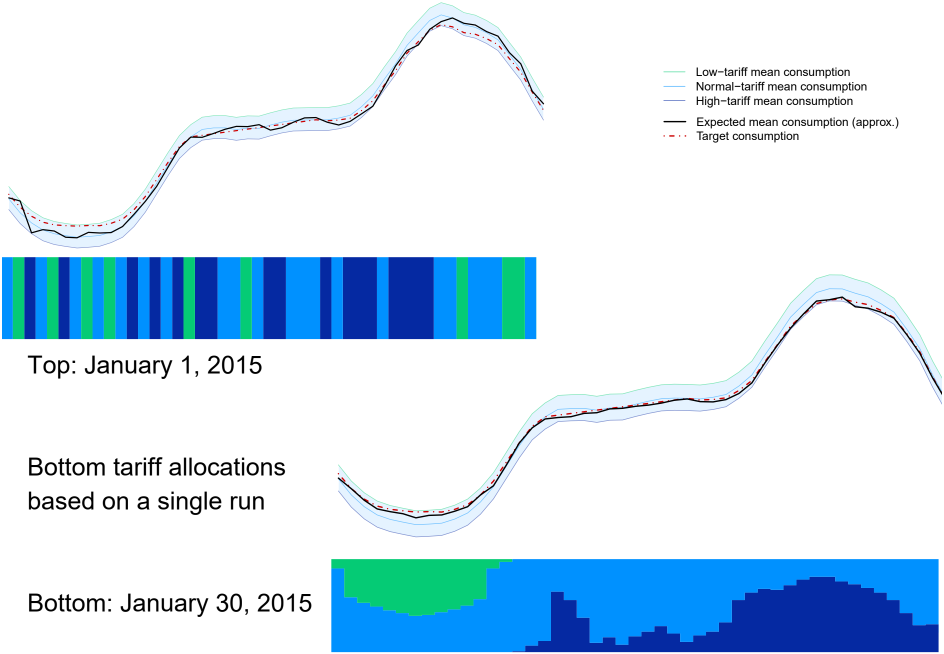

Top: January 1, 2015

Bottom tariff allocations based on a single run

Bottom: January 30, 2015

What's next? [Work in progress]

- The case of inhomogeneous consumers
  - Create clusters of clients according to their profiles
  - Tailor allocations picked to each cluster

- Rebound effect

And now, some dirty details about the algorithm...

Estimation of θ as for the LinUCB algorithm

(Li et al., 2010; Chu et al., 2011; Abbasi-Yadkori et al., 2011)

For some $\lambda > 0$: at the beginning of round $t \geqslant 2$,

$$\widehat{\theta}_{t-1} \in \underset{\widetilde{\theta} \in \mathbb{R}^m}{\arg\min} \left\{ \lambda \left\| \widetilde{\theta} \right\| + \sum_{s=1}^{t-1} \left(Y_{s,p_s} - \widetilde{\theta}^{\mathsf{T}} \phi(x_s, p_s) \right)^2 \right\}$$

$\widehat{\theta}_{t-1}$ is essentially $1/\sqrt{T}$–close to the real parameter θ

Variance Γ also needs to be estimated online

Thus, conditional error

$$\mathbb{E}\left[(Y_{t,p_t} - c_t)^2 \,\middle|\, \mathcal{F}_{t-1} \right] = \left(\theta^{\mathsf{T}} \phi(x_t, p_t) - c_t \right)^2 + p_t^{\mathsf{T}} \Gamma p_t$$

estimated by the confidence interval

$$\left(\left[\widehat{\theta}_{t-1}^{\mathsf{T}} \phi(x_t, p) \right]_C - c_t \right)^2 + p^{\mathsf{T}} \widehat{\Gamma} p \pm \alpha_{t,p}$$

Confidence intervals $\left(\left[\widehat{\theta}_{t-1}^{\mathsf{T}}\phi(x_t, p)\right]_C - c_t\right)^2 + p^{\mathsf{T}}\widehat{\Gamma}p \pm \alpha_{t,p}$

Play optimistically
(it is a trade-off between exploitation and exploration):

Pick $$\operatorname*{arg\,min}_{p\in\mathcal{P}}\left\{\left(\widehat{\theta}_{t-1}^{\mathsf{T}}\phi(x_t, p) - c_t\right)^2 + p^{\mathsf{T}}\widehat{\Gamma}p - \alpha_{t,p}\right\}$$

Theoretical guarantee:

$$\sum_{t=1}^{T}(Y_{t,p_t} - c_t)^2 \lesssim \mathcal{O}(T^{2/3}) + \sum_{t=1}^{T}\min_{p\in\mathcal{P}}\left\{\left(\theta^{\mathsf{T}}\phi(x_t, p) - c_t\right)^2 + p^{\mathsf{T}}\Gamma p\right\}$$

Cumulative error $\lesssim$ Regret + Performance of the best constant p

Technological support for startups

By Agence maths–entreprises

Helps connecting with academic researchers in machine learning and funding projects (50%–50%, even more)
to develop/improve your core technology

Up to 30 kE (e.g., for 1–year shared postdoc)

Recent example: UncharTech (Sébastien Toth, H16)

Contact: I'm a board member, come and talk to me!