Electricity demand response and responsiveness incentives

René Aïd* Dylan Possamaï† Nizar Touzi‡

April 26, 2019

1 Introduction

In its common form, a demand response mechanism is a contract under which the consumer benefits from cheaper electricity than the standard tariff all the year except at certain peak load periods chosen by the producer where the price is much higher. These soft mechanisms appear to cumulate the virtues of consumption reduction while providing substitutes to hardware technologies as chemical batteries or flexible gas–fired plants. There exists many forms of demand response contracts. For domestic customers, the latter form is the most common. An alternative form consists in giving an amount of money at the begining of the year and then substracting the value of the energy consume during prive events. For industrial customers, the payment during price events is indexed on the difference between the energy consumed compared to the energy he would have consumed if he had not receved any signal. This last form is refered to as Peak–Time Rebate (PTR) and the unobserved counterfactual consumption is refered to as the baseline consumption.

Moral hazard lies as a central issue in demand response contracts. Because the consumption during price event is random, it is not possible to know if the observed consumption has been reduced compared to what the consumer would have consumed if no price signal was sent, the latter being non–observable. The success of a demand response mechanism depends on the fact that consumers do react on price events. But moral hazard makes it difficult, if not possible, to quantify the responsiveness of a given consumer or group of consumers, because this quantity is non–observable.

The moral hazard problem translates in the variance of estimated consumption reduction in demand response experiments. Consumers do not react to price signal with the discipline of a gas–fired plant. When receiving a price signal, a household may give up the commitment to reduce consumption to satisfy the constraints of the every day life. The economic and power system literature reports significant variance in the responses of consumers enrolled in demand response trials. The moral hazard problem makes demand response contract less efficient and thus, less valuable than a deterministic susbtitute.

We explain here how the technics of optimal contract design in continuous–time explained in the ILB Method number 1 *Moral Hazard and Continuous–time Contract Theory* by Nizar Touzi can address the efficiency and design problems of demand response contracts. Using these technics, we show that the optimal contract has a rebate form. Furthermore, we show that it is possible and valuable for risk–averse producers to provide incentives to the consumer to increase the response to price events, i.e. to provide regular responses accross price events or to increase the predictablity of

*Université Paris–Dauphine, PSL Research University, LEDa.
†Columbia University, dp2917@columbia.edu.
‡École Polytechnique, CMAP, nizar.touzi@polytechnique.edu.

the consumption. This result is the implementation of the simple fact that uncertainty in payments deteriorates the value of a contract for a risk–averse agent. But, beyond this trivial remark, the problem is how to design a contract that achieves a reduction of the consumption with a small standard deviation when responsiveness is non–observable. Even though responsiveness is not observable, we show that the indexing of the contract to the quadratic variation of the consumption, which is an observable quantity, provides the desired result of making the consumption reduction more reliable.

This document is organised as follows. The model is explained in Section 2, the optimal contracts are provided in Section 3 and numerical illustrations in Section 4.

2 The model

We consider a single producer who has obligation–to–serve the electricity demand of a single residential consumer who enjoys a flat retail rate. The producer wonders if it is interesting or not to propose to the consumer a demand response contract to incite him to change his consumption during a given period of time $[0, T]$. It might be more efficient during peak period to pay the consumer to reduce his consumption rather than deploy costly generation power plants. Besides potential energy consumption reduction, the producer wonders if it is interesting to incite the consumer to provide a more predictable consumption during the price event. A highly random consumption induces adjustement costs that could be avoided if the consumer exhibits a more regular consumption pattern. Further, a highly random consumption reduction from one price event to another reduces the value of the contract as it increases the uncertainty of the payments to or from the consumer. Thus, the concern of the producer is to propose a contract of demand response that would allow her to avoid high generation cost and make the consumer more predictable. Nevertheless, the producer is facing a moral–hazard problem. Once the demand response contract is signed, when the price event occurs, the consumer may change his mind and prefer to consume rather than make the costly efforts of reduction or the costly effort of being regular. Because the producer only observe the consumption and not the actions of the consumer in his house, it is not possible to infer if a consumption level is the result of efforts or just chance.

We focus on one single demand response event with fixed duration, making the hypothesis that all successive demand response events exhibit the same characteristics. At time zero, the producer proposes a paying rule for demand response; at time one, the consumer accepts or rejects the contract; if the consumer accepts the contracts, then the demand response event occurs, the producer measures in continuous–time the consumption and provides the payment (positive or negative) at the end of the event.

The consumption of the consumer at time t denoted by X_t is the sum of the consumption of N different usages X_t^i, $i = 1 \cdots N$. The consumer values the electricity he can consumed and we assume he has a function that converts the level of consumption in monetary value $f(X_t)$. We consider a constant marginal value of energy κ, i.e. $f(x) = \kappa x$. The consumer can achieve a reduction of his consumption by taking actions $a_i(t)$ differentiated per usage i. These actions are costly in the sense that they require to use substitutes to achieve the same level of comfort. In times of demand response event during winter, households may use gas stoves or other heating devices. The consumer can also achieve regularity in his responses to the producer sollicitations by disciplining his consumption usages. During the price event, random events for each usage occur that may drive away the consumer from his planned consumption pattern if no actions are taken to dampen their effects (kids coming back from school with friends, sudden drop of outside brightness...). We denote by $\sigma_i W_t^i$ the random events that strike usage i where σ_i is its nominal volatility and W_t^i is a standard Brownian motion and by $b_i(t) \in (0, 1]$ the action taken by the consumer to reduce its effect on his consumption. Increasing effort corresponds to a lower value of b_i. The mitigation of the random events that occurs during the price event results in a shock $\sigma_i b_i(t) W_t^i$ which exhibit a reduced volatility as soon as b_i is lower than one.

This results in the following controled dynamics of the consumer's electricity consumption:

$$X_t = \sum_{i=1}^{N} X_t^i = x_0 + \sum_{i=1}^{N} -a_i(t)dt + \sigma_i \sqrt{b_i(t)} dW_t^i. \tag{2.1}$$

In our setting, the responsiveness of the consumer is the unobservable actions $b_i(t)$ taken to reduce the impact of his daily life events on his consumption. Of course, the higher the discipline of the consumer, the higher the cost. We suppose that the costs induced by the vector of actions a and b follows the parametric form given by

$$c(a,b) = \sum_{i=1}^{N} c_i(a_i, b_i) := \sum_{i=1}^{N} \frac{1}{2}\frac{a_i^2}{\mu_i} + \frac{1}{2}\frac{\sigma_i(b_i^{-1}-1)}{\lambda_i},$$

where μ_i and λ_i are strictly positive parameters that caracterise the cost of reduction of each usage. The higher the value of the parameter, the lower it is costly to reduce its consumption or its volatility.

Given the compensation ξ for the realised consumption and regularisation, the consumer's criterion is given by

$$J_{\mathrm{A}}(\xi; a, b) := \mathbb{E}\Big[U_{\mathrm{A}}\Big(\xi + \int_0^T \big(f(X_s) - c(\nu_s)\big)\mathrm{d}s\Big)\Big], \text{ where } U_{\mathrm{A}}(x) := -\mathrm{e}^{-rx}, \tag{2.2}$$

for some constant risk aversion parameter $r > 0$. The problem of the consumer is thus

$$V_{\mathrm{A}}(\xi) := \sup_{\substack{a\geq 0,\\ 0<b\leq 1}} J_{\mathrm{A}}(\xi, a, b) \tag{2.3}$$

i.e. maximising utility from consumption subject to the cost of effort.

We finally assume that the consumer has a reservation utility R_0. We denote by $L_0 := -\frac{1}{r}\log(-R_0)$, the certainty equivalent of the reservation utility of the consumer for the consumer.

On her side, the producer bears two kind of costs: generation cost to meet consumption of the consumer on real–time basis $g(X_t)$ and direct consumption volatility cost h. The marginal generation cost is assumed constant, i.e. $g(x) = \theta x$. We denote by $\delta := \kappa - \theta$, the *energy value discrepancy.* The case $\delta \geq 0$ corresponds to off–peak hours (the energy is cheaper to produce than it has value for the consumer) while negative δ corresponds to peak–load hours (the energy is more costly to produce than it has value for the consumer). The direct consumption volatility cost h represents the costs induced by the non-predictable part of the consumption. Her performance criterion is defined by

$$J_{\mathrm{P}}(\xi; a, b) := \mathbb{E}\Big[U\Big(-\xi - \int_0^T g(X_s)\mathrm{d}s - \frac{h}{2}\langle X\rangle_T\Big)\Big], \text{ with } U(x) := -\mathrm{e}^{-px}. \tag{2.4}$$

The two following contracting problems are considered here:

- *First best contracting* corresponds to the benchmark situation where the producer has full power to impose a contract to the consumer and to dictate the Agent's effort

$$V^{\mathrm{FB}} := \sup_{\substack{\xi; a\geq 0,\\ b\in(0,1]}} \big\{J_{\mathrm{P}}(\xi, a, b) : \; J_{\mathrm{A}}(\xi, a, b) \geq R_0\big\}, \tag{2.5}$$

- *Second best contracting* allows the consumer to respond optimally to the producer offer.

$$V^{\mathrm{SB}} := \sup_{\xi} J_{\mathrm{P}}(\xi; a, b) \quad \text{s.t.} \quad V_A(\xi) \geq R_0. \tag{2.6}$$

3 The optimal contracts

We now provide the optimal contracts for the cases cited above. The first–best case is a standard optimisation problem under constraints and the solution is obtained using KKT conditions. The second–best situation is a more complex problem. Its solution is obtained using the technics developed in Sannikov (2008) [7] and Cvitanic et. al. (2018) [3] for which a resume is given in the ILB Method #1 [9]. We recall here the general form of the optimal contract in the second–best case. For any payment rates Z_t and Γ_t, the contract

$$\xi = Y_0 + \int_0^T Z_t dX_t + \frac{1}{2}\int_0^T (\Gamma_t + rZ_t^2) d\langle X\rangle_t - \int_0^T (H(Z_t, \Gamma_t) + f(X_t))dt, \tag{3.1}$$

where H is the Hamiltonian of the consumer's problem, ensures that $V_A(\xi) = R_0$.

First–best contract

The optimal first-best contract is given by $\xi_{\text{FB}} = \xi_{\text{FB}}^{\text{F}} + \xi_{\text{FB}}^{\text{V}}$ where

$$\xi_{\text{FB}}^{\text{F}} := L_0 - \kappa X_0 T, \quad \xi_{\text{FB}}^{\text{V}} := \int_0^T c(a_{\text{FB}}(t), b_{\text{FB}}(t))dt + \int_0^T \pi_{\text{FB}}^{\text{E}}\big(x_0 - X_t\big)dt - \frac{1}{2}\int_0^T \pi_{\text{FB}}^{\text{V}} d\langle X\rangle_t.$$

and

$$\pi_{\text{FB}}^{\text{E}} := \frac{r}{r+p}\kappa + \frac{p}{r+p}\theta, \qquad \pi_{\text{FB}}^{\text{V}} := \frac{p}{r+p}h.$$

Under optimal contract, the consumer's efforts are

$$a_{\text{FB}}(t) := \mu\delta^{-}(T-t), \quad \text{and} \quad b_{\text{FB}}(t) := 1 \wedge \Big(\lambda_j(h + \rho\,\delta^2(T-t)^2)\Big)^{-\frac{1}{2}}.$$

How does the first–best contract reads? The contract is the sum of a fixed payment $\xi_{\text{FB}}^{\text{F}}$ and a payment that varies according to the realised trajectory of the consumption $\xi_{\text{FB}}^{\text{V}}$. The fixed part is the sum of the certainty equivalent of the reservation utility of the consumer minus the value of the energy for the consumer if he makes no effort. Indeed, if the consumer makes no effort, the consumption is a martingale and the expected value of the consumption is thus $\kappa T x_0$. The variable part has three components:

- the observable and enforcable cost of effort $\int_0^T c(a_{\text{FB}}(t), b_{\text{FB}}(t))dt$
- the payment or the charge for the deviation from the initial consumption level x_0 : $\int_0^T \pi_{\text{FB}}^{\text{E}}\big(x_0 - X_t\big)dt$. If the consumption decreases, the consumer receives a payment at a constant price $\pi_{\text{FB}}^{\text{E}}$. In the other case, the consumer is charged at the same price.
- the payment or the charge for the variations of the volatility $\frac{1}{2}\int_0^T \pi_{\text{FB}}^{\text{V}} d\langle X\rangle_t$. If the volatility decreases, the consumer is paid at a constant price $\pi_{\text{FB}}^{\text{V}}$. He is charged in the other case at the same price.

We find that the optimal first best contract has a form which is already implemented in the market. The general form of the contract is of a rebate contract where x_0 serves as the baseline.

More comments can be made. If the consumer is risk–neutral ($r = 0$), the price of energy reduces to its marginal cost of production: the first–best contract transfers all the uncertainty of the generation cost to the consumer as it is standard in the moral hazard optimal contract framework. Apart from this case, the first–best price for energy $\pi^{\mathrm{E}}_{\mathrm{FB}}$ is a convex combination of the marginal cost θ and value κ weighted sum by their risk–aversion ratios. It does not depends on any other parameters. In particular, it does not depend on the marginal costs of efforts because they are observed and paid separately. The first–best price for the responsiveness $\pi^{\mathrm{V}}_{\mathrm{FB}}$ is a constant fraction of the direct cost of volatility. It is zero only if there is no direct cost of volatility or if the producer is risk–neutral. As a consequence, we see that a contract that would be indexed only on the information of the cost function of the producer is optimal only in the case when the consumer is risk–neutral. The economic intuition that the marginal cost of generation triggers a socially optimal response is correct only if the consumer is risk–neutral. If not, the consumers needs a compensation payment for the risk he takes in the contract.

Regarding the induced behaviour of the consumer, we see that the consumer is required to make an effort on his average consumption only during peak hours, i.e. when $\delta \leq 0$. Indeed, during peak hours, the energy is more costly to produce than it has a value for the consumer. Thus, it makes sense for the producer to pay the consumer a price between θ and κ to get the consumer to reduce his consumption. Regarding the volatilities, reduction is performed only on those usage for which the marginal cost of effort measured by $1/\lambda_j$ is lower than the marginal cost of volatility for the producer measured by $h + \rho\delta^2(T-t)^2$. It is only in the cases where the producer or the consumer is risk–neutral or if they agree on the energy value ($\delta = 0$) that responsiveness is triggered by the mere comparison of h and the λ_j.

Second–best contract

The second–best optimal contract is given by $\xi_{\mathrm{SB}} = \xi^{\mathrm{F}}_{\mathrm{SB}} + \xi^{\mathrm{V}}_{\mathrm{SB}}$ where

$$\xi^{\mathrm{F}}_{\mathrm{SB}} := L_0 - \kappa T x_0 - \int_0^T H(z_{\mathrm{SB}}, \gamma_{\mathrm{SB}})(t)dt, \quad \xi^{\mathrm{V}}_{\mathrm{SB}} := \int_0^T \pi^{\mathrm{E}}_{\mathrm{SB}}(t)\big(x_0 - X_t\big)dt - \frac{1}{2}\int_0^T \pi^{\mathrm{V}}_{\mathrm{SB}}(t)d\langle X\rangle_t,$$

where H is the Hamiltonian of the consumer's problem and

$$\pi^{\mathrm{E}}_{\mathrm{SB}}(t) := \kappa + z'_{\mathrm{SB}}(t), \qquad \pi^{\mathrm{V}}_{\mathrm{SB}}(t) := h + p\big(z_{\mathrm{SB}}(t) - \delta(T-t)\big)^2,$$

where z_{SB} is a deterministic function of time. Under the optimal contract, the optimal efforts of the consumer are:

$$a_{\mathrm{SB}}(t) := \mu z_{\mathrm{SB}}(t)^-, \quad \text{and} \quad b_{\mathrm{SB}}(t) := 1 \wedge (\lambda_j \gamma_{\mathrm{SB}}(t)^-)^{-\frac{1}{2}},$$

where $\gamma_{\mathrm{SB}}(t) = -h - r z_{\mathrm{SB}}(t)^2 - p(z_{\mathrm{SB}}(t) - \delta(T-t))^2$.

The second–best optimal contract shares important similarities with the first–best contract. It has also a rebate form composed of a fixed part payment $\xi^{\mathrm{F}}_{\mathrm{SB}}$ and a payment that depends on the realised consumption trajectory $\xi^{\mathrm{V}}_{\mathrm{SB}}$. We stress here the main differences.

First, the fixed part payment contains a term that is absent from the first–best, namely the quantity $\int_0^T H(z_{\mathrm{SB}}(t), \gamma_{\mathrm{SB}}(t))dt$. This quantity is a value known at the beginning of the price event. It corresponds to the maximum net benefit (payments minus costs of efforts) the consumer can get for his optimal efforts under the payment rate z_{SB} and γ_{SB}. Contrary to the first–best case where the realised costs of efforts were observable and thus paid to the consumer, in the second–best case, it is no longer possible. Thus, the producer starts by substracting the maximum net benefit from the payment to the consumer, providing incentive for the consumer to do his best effort if he wants to recover his maximum expected value.

Second, the prices for energy $\pi^{\mathrm{E}}_{\mathrm{SB}}$ and for volatility $\pi^{\mathrm{V}}_{\mathrm{SB}}$ are no longer constant in time. They are non–increasing and non–linear. In the case where the consumer is risk–neutral, the first and second–best energy prices are equal but not the volatility prices.

For comparisons purposes, it is interesting to consider the case where the consumer does not receive incentive on his responsiveness. In this case, the consumer receives payment rates only on his average consumption and not on the volatility, $\Gamma \equiv 0$. The optimal contract is explicit and given below.

Second best without responsiveness incentives

The second–best optimal contract without responsiveness incentive is given by $\xi_{\mathrm{SB}_m} = \xi^{\mathrm{F}}_{\mathrm{SB}_m} + \xi^{\mathrm{V}}_{\mathrm{SB}_m}$ where

$$\xi^{\mathrm{F}}_{\mathrm{SB}_m} = L_0 - \kappa T x_0 + \frac{1}{2}\int_0^T r z^2_{\mathrm{SB}_m}(t)|\sigma|^2 dt - \int_0^T H(z_{\mathrm{SB}_m}(t), 0)dt, \quad \xi^{\mathrm{V}}_{\mathrm{SB}_m} = \int_0^T \pi^{\mathrm{E}}_{\mathrm{SB}_m}\big(x_0 - X_t\big)dt,$$

and the price for energy $\pi^{\mathrm{E}}_{\mathrm{SB}_m}$ and the price for volatility $\pi^{\mathrm{V}}_{\mathrm{SB}_m}$ are given by:

$$\pi^{\mathrm{E}}_{\mathrm{SB}_m} := (1-\Lambda)\kappa + \Lambda\theta, \quad \pi^{\mathrm{V}}_{\mathrm{SB}_m} := 0, \quad \text{with} \quad \Lambda := \frac{p|\sigma|^2 + \bar{\mu}\mathbf{1}_{\{\delta<0\}}}{(p+r)|\sigma|^2 + \bar{\mu}\mathbf{1}_{\{\delta<0\}}},$$

where $\bar{\mu} := \sum_{i=1}^N \mu_i$.

Again, the contract has a rebate form similar to the two preceeding cases. The price for energy $\pi^{\mathrm{E}}_{\mathrm{SB}_m}$ is constant, as in the first–best case, and it is also a convex combination of the marginal cost θ and value κ of energy. During off peak period ($\delta > 0$), the weight is the risk–aversion ratio and thus, the energy price is equal to the first–best case. But, during peak–load period, the weights depend now not only on the risk–aversion ratios but also on the volatility of the consumption σ and on the marginal cost of effort of the consumer μ. The energy price is no longer a function only of the risk aversion parameters and of the marginal cost and values of the energy.

4 Numerical illustration

We illustrate some aspects of the implementation of a responsiveness incentive mechanism for electricity consumption reduction. We calibrated our model on the publicly available data set of Low Carbon London Pricing Trial using the idea that a standard demand respons trial corresponds to the implementation of an optimal contract without responsiveness incentives. The data gathered by the Low Carbon London Project of demand-side response (DSR) trial performed in 2012-2013 can be downloaded freely at London DataStore website (https://data.london.gov.uk). The detailled calibration process and parameters value can be found in Aïd et. al. (2018) [1].

Figure 1 illustrates the different energy and volatility prices for the first–best, the second–best and the second–best without responsiveness incentives. The second–best price of energy without responsiveness incentives is significantly different from the first–best and also from the marginal cost of energy. The incentive for responsiveness leads to a non–constant price of energy. It lies between the marginal cost of energy and the second–best price without responsiveness incenticve. Further, it is higher at the beginning of the price event to trigger a quick response of the consumer. The price for volatility follows the same pattern of decreasing value. Note that it significantly different from the first–best price of volatility.

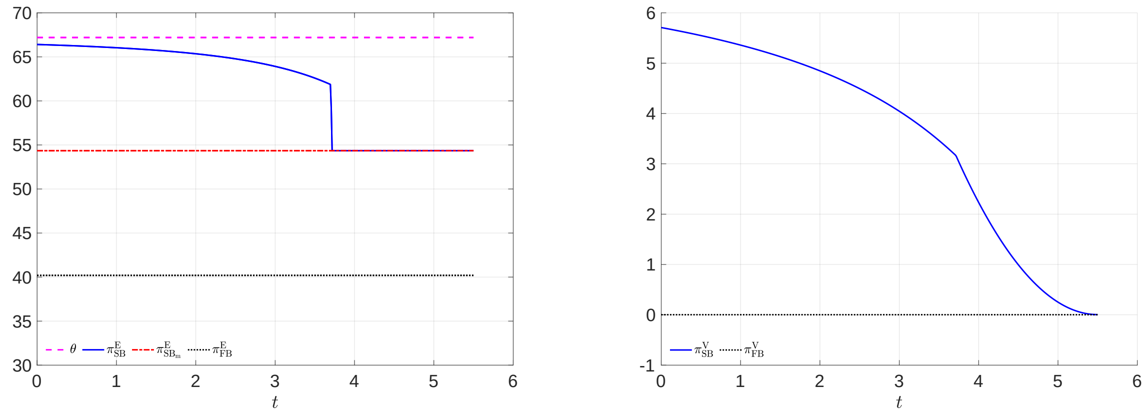

Figure 1: Prices for energy (left) and volatility (right).

The Figure 2 (Left and Middle) the gain from responsiveness control and the reduction of of the volatility as functions of the duration of price event T in hours and the energy value discrepancy $|\delta| = \theta - \kappa$. The red dot in the pictures represents the nominal situation. We observe that there is a threshold of values of energy value discrepancy and price event duration under which no benefit should be expected from the responsiveness incentives. The lower the energy value discrepancy, the longer the price event should be to ensure a significant benefit of responsiveness control. The incentive on volatility needs time or a large energy value discrepancy to show its benefits. But, on the other, the reduction of volatility is less prone to this dependence on the energy value discrepency. Even modest differences can induce subtantial reduction of volatility for a standard duration of a price event.

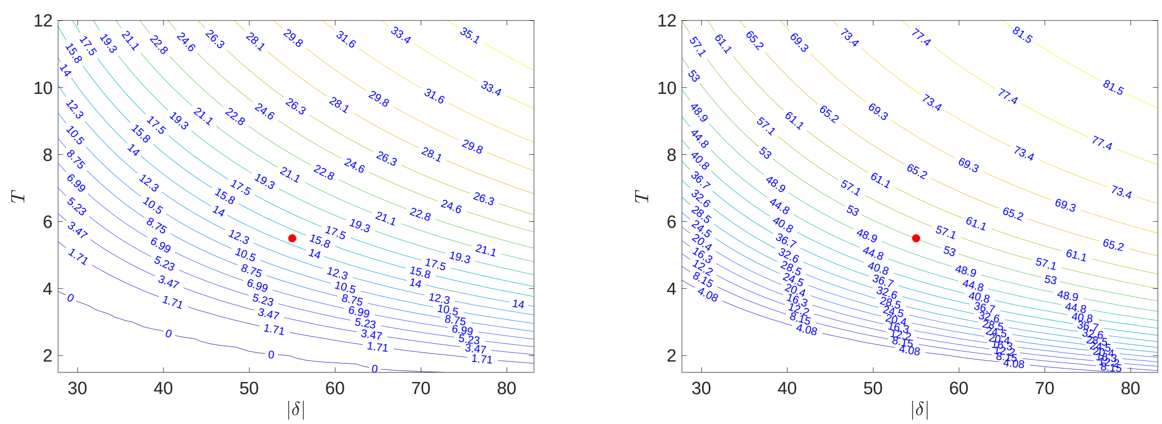

Figure 2: (Left) Gain for the producer from responsiveness control and (Right) volatility reduction as a function of the price event duration T and the absolute value of the energy value discrepancy δ; all in percentage.

Figure 3 shows the total payment and its decomposition between its fixed part and the certainty

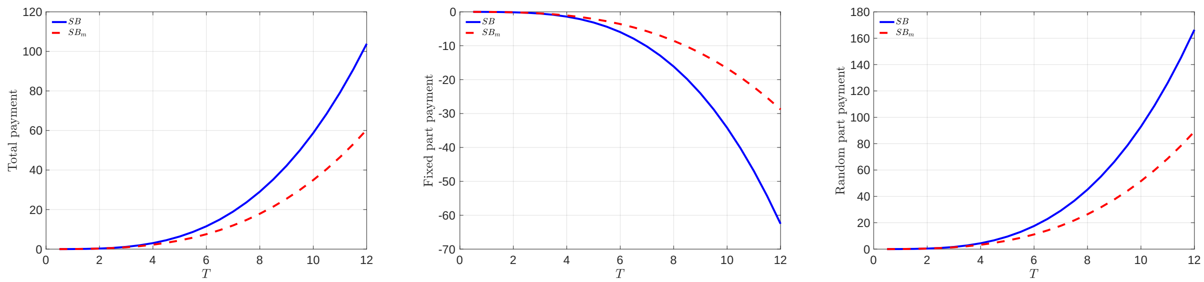

Figure 3: Total (left), fixed part (middle) and certainty equivalent of the random part (right) of the optimal payment with responsiveness control (blue) and without (red) as a function of the price event duration T in pence.

equivalent of the random part as a function of the duration of the price event T with and without responsiveness control. In both cases, the total payment is positive and increases when the duration of the effort becomes large. As no surprise, the payment with responsiveness control is larger than without it because it requires more efforts from the consumer. The noticeable result comes from the decomposition of the contract between its deterministic part and its random part. The producer charges more the consumer when implementing responsiveness incentive than without but, provides higher certainty equivalent. The implementation of regular response from the consumer starts by charging him a lot more but also by rewarding him a lot more in case of appropriate result. The longer the consumer is asked to make an effort, the higher this difference should be.

4.1 Practical issue

We have seen in the preceeding section that the optimal contract to induce an increase in the responsiveness of consumers to price signal should be written on the quadratic variation of the consumption. In discrete time $0 = t_0 < t_1 \cdots t_n = T$, the quadratic variation $\langle X \rangle_T$ can be approximated by

$$\langle X \rangle_T \approx \sum_{i=0}^{n-1} \left(X_{t_{i+1}} - X_{t_i} \right)^2.$$

Regarding energy consumption, households make a clear connection between the fact that they reduce the heating of the house and the reduction of their consumption. But, it is not obvious that domesic consumers first would understand why they should be charged a price proportional to that quantity and second, how the different actions they take during the day are related to it. The genuine form of the contract might not be acceptable for consumers. It is necessary to make simplifications to increase its potential acceptability.

A way one could think of is to index on the event–per–event variations of the average consumption. On event k, the producer measure the quantity $\bar{X}_k := \int_0^T X_t dt$ and charges the consumer a cost proportional to the quantity $\mathbb{V}_k := |\bar{X}_k - X_{\mathrm{c}}|$ where X_{c} is a contractualised targeted consumption. This contract is clearly sub–optimal: we loose the fact that the incentive price for responsiveness should be higher in the beginning of the price event, we loose the fact that the contract should be written on the quadratic variation, and not on the L^1–norm and so on. But, we earn a simple way to provide incentives to the consumer to remain as close as possible to a given pattern of consumption. The quantity $\mathbb{V}_k$ is measured in kWh and can be understood to be charged at a price measured in €/kWh just as the energy.

5 Perspectives

The use of modern tools of optimal contract theory sheds a new light on the design of electricity demand response contract. It is possible to reduce the average consumption, while improving at the same time the responsiveness of the consumer. The calibration of our model to pricing trial data predicts that the cost of efforts of the consumer to reduce his average consumption will lead to significant benefits for producers and significant increase in the responsiveness of consumers and thus, on the efficiency of demand response programs. These predictions are testable. If our claim is true, the indexing of the payment to consumer on their regularity of consumption across price events should deeply enhance the efficiency of demand response programs. Proper experiments could test this prediction.

References

[1] R. Aïd, D. Possamaï, N. Touzi. Optimal electricity demand response contracting with responsiveness incentives. arxiv, https://arxiv.org/abs/1810.09063, 2018.

[2] J. Cvitanić, J. Zhang. Contract theory in continuous–time models. Springer, 2010.

[3] J. Cvitanić, D. Possamaï, N. Touzi. Dynamic programming approach to principal–agent problems. *Finance and Stochastics*, 22(1):1–37, 2018.

[4] A. Faruqui, S. Sergici. Household response to dynamic pricing of electricity: a survey of 15 experiments. *Journal of Regulatory Economics*, 38:193–225, 2010.

[5] B. Hölmstrom, P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328, 1987.

[6] J.-J. Laffont and D. Martimort. *The theory of incentives: the principal–agent model*. Princeton University Press, 2002.

[7] Y. Sannikov. A continuous–time version of the Principal–Agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.

[8] J. Schofield, R. Carmichael, S. Tindemans, M. Woolf, M. Bilton, G. Strbac. Residential consumer responsiveness to time–varying pricing. Report A3 for the Low Carbon London LCNF project: Imperial College London, 2014.

[9] N. Touzi. Moral hazard and continuous–time contract theory. Institut Louis Bachelier Methods #1.