

Impact of Demand-Response on the Efficiency and Prices in Real-Time Electricity Markets*

Nicolas Gast
IC-LCA2
EPFL, Switzerland
nicolas.gast@epfl.ch

Jean-Yves Le Boudec
IC-LCA2
EPFL, Switzerland
jean-yves.leboudec@epfl.ch

Dan-Cristian Tomozei
IC-LCA2
EPFL, Switzerland
dan-cristian.tomozei@epfl.ch

ABSTRACT

We study the effect of Demand-Response (DR) in dynamic real-time electricity markets. We use a two-stage market model that takes into account the dynamical aspects of generation, demand, and DR. We study the real-time market prices in two scenarios: in the former, consumers anticipate or delay their flexible loads in reaction to market prices; in the latter, the flexible loads are controlled by an independent aggregator. For both scenarios, we show that, when users are price-takers, any competitive equilibrium is efficient: the players' selfish responses to prices coincide with a socially optimal policy. Moreover, the price process is the same in all scenarios. For the numerical evaluation of the properties of the equilibrium, we develop a solution technique based on the Alternating Direction Method of Multipliers (ADMM) and trajectorial forecasts. The forecasts are computed using wind generation data from the UK. We challenge the assumption that all players have full information. If the assumption is verified, then, as expected, the social welfare increases with the amount of DR available, since DR relaxes the ramping constraints of generation. However, DR actions alter the internal states of elastic loads; if the day-ahead market cannot observe how elastic loads are affected by DR, a large quantity of DR can be detrimental and leads to a decrease in the welfare. Furthermore, the DR operator has an incentive to under-dimension the quantity of available DR. Finally, we compare DR with an actual energy storage system. We find that storage has a faster response-time and thus performs better when only a limited amount is installed. However, storage suffers from charge-discharge inefficiency: with DR, prices do concentrate on marginal cost (for storage, they do not) and provide a better welfare.

1. INTRODUCTION

Electricity markets are developing worldwide, replacing tightly regulated systems by decentralized control mechanisms based on spot prices [22]. Demand-response (DR) mechanisms have emerged in order to take advantage of flexible demand in a system where demand is traditionally considered as inflexible [19]. DR can com-

pensate mismatches between production and consumption, thus increasing the potential deployment of renewable generation by reducing the necessity for regulation reserve. It can be implemented by having consumers react to prices [16] or to congestion signals [15].

The authors of [21] present a model of real-time electricity market that incorporates the dynamic constraints of generation and the uncertainties due to forecast errors. The system is composed of two market actors (a supplier and a consumer) that exchange energy at a spot price. They show that there exists a competitive equilibrium, *i.e.*, a price process such that market actors agree on the quantity of energy exchanged. Moreover, the market is efficient: for any competitive equilibrium, the selfish decisions of both market actors coincide with the hypothetical decisions taken by a social planner that aims to maximize the sum of every actor's utility. However, the equilibrium price exhibits large fluctuations and is never equal to the marginal production cost.

We extend [21] to model a two-stage electricity market that features generation constraints, inelastic loads, and elastic loads corresponding to a set of DR appliances. Our elastic load model captures several key-features of controllable thermostatic loads. The loads are elastic in the sense that it is possible to perform time-shifts: the consumption of each load can be delayed or anticipated, but an object's consumption cannot be arbitrarily reduced by augmenting prices. Each appliance has an internal state representing, *e.g.*, its temperature. We model individually the fatigue effect of each appliance by using *undesirable states* which have an associated penalty: if a load is delayed for too long, it reaches such an undesirable state. In practice these states are only reached in the case of a blackout. We add a second dimension to the state space: a counter indicates the waiting time before the appliance is allowed to respond to DR signals, thus modeling mini-cycle avoidance (an appliance that has just been switched on cannot be switched off immediately). In our numerical evaluations we use a mean-field model to keep track of the empirical distribution of the appliances' states.

We consider two market scenarios. In Scenario 1, the

*This work is partially supported by the EU project QUANTICOL, 600708.

appliances are controlled by the consumer who reacts to real-time prices by anticipating or delaying the consumption. In Scenario 2 a DR operator controls all the flexible appliances. She charges a fixed rate to users and makes a profit by scheduling the appliances when the market price is low. We show that in both cases, the market is efficient, *i.e.*, the selfish decisions of players coincide with the ones of a hypothetical social planner.

We develop a numerical methodology based on the Alternating Direction Method of Multipliers (ADMM) to compute the market equilibrium for a realistic forecast error model. The forecast errors are modeled using multiple trajectories [17], generated from multivariate Gaussian random variables with correlation matrices derived from traces of data from the UK.

Our numerical results show that, for low DR penetration, the gain obtained by using demand-response is almost linear in the installed DR capacity. For large DR penetration, if the state of the appliances can be observed by all players, the gain saturates to a value that does not depend on the appliances' nature. However, if the states cannot be observed by the day-ahead market, increasing the capacity of DR can be detrimental: after a certain capacity, it decreases the social welfare, as the consumption of DR appliances is harder to predict. Finally, we also show that DR operators have an incentive to under-size the DR capacity.

We also compare DR with classical energy storage [9]. We find that, at low capacity, a storage system offers better performance than DR because it reacts faster, despite losses due to charge-discharge cycles. However, for large capacities of installed DR, DR behaves similarly to a perfectly efficient storage. As such, for large capacities, DR outperforms energy storage systems, since these have a charge-discharge efficiency of 70 – 90%.

Roadmap. We discuss related work in Section 2. We present the model, the assumptions, and the main definitions in Section 3, and the social welfare theorem in Section 4. We describe our statistical model of error and our numerical methodology in Section 5. Finally, we give the numerical parameters and present the numerical evaluation of the model in Section 6. We conclude in Section 7.

2. RELATED WORK

The authors of [1] present an overview of DR programs and their integration within the electricity market. They compare price-based and incentive-based programs and present their benefits (peak reduction, bill savings, reliability) and drawbacks (deployment cost, rescheduling, metering infrastructure). New market models tailored for DR have also been investigated. For example, the authors of [20] develop a new market-clearing mechanism for load-shifting. However, it assumes perfect forecast and is limited to a small number of players.

One of the main obstacles in the development of DR

is the fact that users lack intuition about DR both before the installation – awareness of the advantages and drawbacks – but also during its use – people do not want to have to continuously monitor prices [10]. Hence, a key research topic is the means of presenting DR as a service transparent to use via smart home controllers or by having a DR operator with direct control of appliances. Multiple models are proposed, using response to prices [16] or to congestion signals [15].

DR interferes with demand and has to be taken into account when studying electricity markets. Demand response fatigue and rebound effect are known and modeled in [11] using elasticity coefficients. The response fatigue is also a large concern for the DR operator and the grid regulator. In [14], the authors use a simplified model and find that when the state of flexible appliances is not taken into account in the control, a large accumulated delayed load may manifest unexpectedly and randomly. In this paper we model the internal states of each appliance in detail and, in Section 6.3 study whether such an effect continues to occur.

Our appliance model is similar to [12, 15], who also consider thermostatic controllable loads. A difference with our model is that we add a second dimension for modeling mini-cycle avoidance. We also take into account the undesirable states that could be reached in blackout conditions. Moreover, we consider that the elastic loads can be anticipated or delayed, but that in the long run they consume a fixed amount of energy, *e.g.*, the average consumption of a heater or cooler. This is not the case in [15]. In many papers about demand-response in electricity markets, it is often considered that higher prices result in reduced consumption [1].

Even if the model has similarities, the papers [12, 15] study a different problem. They do not incorporate the market aspects and focus on the problem of the optimal control of DR from the aggregator's point of view in response to congestion signals or prices. The development of such control algorithms is a popular research subject, see for example [5] and the reference therein.

Our work builds on various papers about efficiency in electricity markets (*e.g.*, [21, 9] and the reference therein). One key contribution of the current paper is that we are able to handle a more realistic forecast model in our numerical methodology. The numerical results produced in [21, 9] were performed by assuming that the forecast error can be represented by a Brownian motion with stationary increments. We use a non-stationary model of error that is becoming a standard for wind forecast. It uses a probabilistic forecast developed by Pinson et al. [17], where the errors are represented by a finite number of possible trajectories. Using branching trajectories, we extend this model to represent the growing uncertainty of forecasts over time and the possible observations by actors.

3. SYSTEM MODEL

We consider a two-stage electricity market (day-ahead and real-time) with three actors: the consumer, the supplier, and the DR operator. The supplier has two roles: she generates the bulk of the electricity determined via the day-ahead scheduling and she provides regulation electricity to compensate for mismatches in real-time. The consumer aggregates the consumption of a population of end-users. The controllable demand of end-users is managed by the DR operator who uses the flexibility of the controllable appliances as a means to perform arbitrage on the real-time electricity market. The role of the DR operator can be assumed by the consumer, or by an independent actor. In Section 6, we will explore two type of controllable loads: fridges and boilers.

3.1 Two-stage Electricity Markets

Two-stage electricity markets are used to determine the price of electricity supplied to the consumers. In the day-ahead stage, the actors use forecasts of consumption and of renewable production to schedule the bulk production for the next day. In the real-time stage, last-minute decisions are taken to compensate mismatches.

We start from the day-ahead market of [4]. The forecast demand for the next day at time t is used to set the scheduled production $g^{da}(t)$. The scheduled production incorporates both renewable energy (volatile) and conventional energy sources. The forecast demand has two components: the non-controllable demand $d^{da}(t)$ and the controllable (flexible) demand $f^{da}(t)$. An additional fixed quantity r^{da} is produced as a precaution against forecast errors. Thus, in the day-ahead market the actors agree to trade an amount $g^{da}(t) + f^{da}(t)$ the next day at time t at a price $p^{da}(t)$, where $g^{da}(t) := d^{da}(t) + r^{da}$.

The real-time market deals with the inevitable mismatches that arise from forecast errors. Thus, at time t the end-users express a total aggregated non-controllable demand $D^a(t) = d^{da}(t) + D(t)$ and controllable (flexible) demand $F^a(t) = f^{da}(t) + F(t)$. The quantities $D(t)$ and $F(t)$ are the real-time components of the demand, and they can be positive or negative. While $D(t)$ is given by nature, $F(t)$ can be controlled to some extent, depending on the state of the controllable appliances and on the control signal decided by the DR operator.

The supplier deploys real-time production $G(t) > 0$ in order to compensate for mismatches. This part of the generation comes from conventional sources, and is subject to ramping constraints. The total amount of produced electricity at time t is thus $G^a(t) = f^{da}(t) + g^{da}(t) + \Gamma(t) + G(t)$, where $\Gamma(t)$ is the forecast error of renewable sources (*i.e.*, the difference between actual and forecast production). The energy produced in real-time is traded at price $P(t)$ at time t .

In this paper, we study the effect of demand-response on the real-time market. We assume that the market

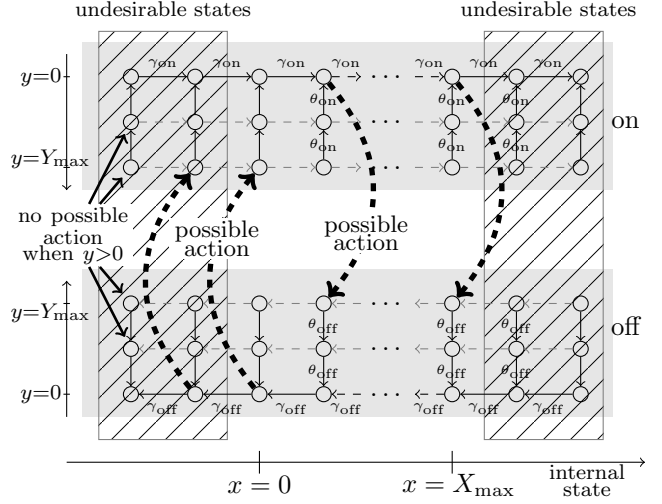


Figure 1: Markov model of an appliance. Each state is characterized by a tuple (s, x, y) , with $s \in \{on, off\}$. The x -axis represents the internal state (*e.g.*, temperature). The y -axis is the time before the device can be switched on or off.

actors are price takers: they strategically define their actions based on prices, but they cannot influence these prices. Moreover, we assume that the market actors base their decisions at time t on the knowledge of past data and actions (up to time t) and the statistics of future data. This knowledge is shared by all actors, unless otherwise specified. In our mathematical model, this means that there exists a filtration (\mathcal{F}_t) such that the decision processes are adapted to this filtration.

3.2 Controllable Appliance Model

We consider N appliances that can be controlled via demand-response. Their maximum aggregated power consumption is denoted P_{on} . At any time step, each appliance is fully characterized by its state $i = (s, x, y) \in \mathcal{M}$. The state space is defined as $\mathcal{M} = \{on, off\} \times \mathbb{Z} \times \{0, \dots, Y_{\max}\}$, where

- $s \in \{on, off\}$ indicates whether the appliance is on, in which case it consumes a power P_{on}/N , or whether it is off, in which case it consumes 0; if $s = on$ then we denote by $\bar{s} = off$ and vice-versa;
- $x \in \mathbb{Z}$ reflects the internal state of the appliance, *e.g.*, the temperature in the case of a fridge;
- $y \in \{0, \dots, Y_{\max}\}$ represents the amount of time that needs to pass before the device can react to demand-response. The device can be switched from *on* to *off* (or vice-versa) only if $y = 0$. Moreover, as soon as the state is switched as a result of a demand-response action, the value y jumps from 0 to Y_{\max} . By accounting for y , we avoid operation in mini-cycles which might damage the appliance.

In Figure 1, we represent the Markov chain that gives

the evolution of the state of an appliance. Each node on the diagram corresponds to a possible state, with

- $s = \text{on}$ for the nodes in the upper rectangle and $s = \text{off}$ for the ones in the lower rectangle,
- x is given by the position on the x -axis, and
- y is given by the corresponding y -axis.

Any state $i = (s, x, y)$ has a horizontal transition to a neighboring state ($\text{on}, x+1, y$) if $s = \text{on}$ and ($\text{off}, x-1, y$) if $s = \text{off}$. The probability of this horizontal transition is γ_s if $y > 0$, or $(1 - a_i)\gamma_s$ if $y = 0$. We consider that γ_{on} is typically larger than γ_{off} . This allows to model the different speed at which the internal state of the appliance evolves when it is *on* or *off* (e.g., a fridge spends less time *on* than *off*). When $y > 0$, we add a vertical transition to a state $(s, x, y-1)$, which decreases the y counter with a certain probability θ_s .

We partition the state space into “responsive” states $\mathcal{M}_1 \stackrel{\text{def.}}{=} \{(s, x, y) : y = 0\}$ and “non-responsive” states $\mathcal{M}_2 = \mathcal{M} \setminus \mathcal{M}_1$. A demand-response *action* is defined as a vector of probabilities $a = (a_i)_{i \in \mathcal{M}}$ indexed on the states. For a state $i = (s, x, y)$, a_i is the probability that an appliance switches its s state:

- If the appliance is in the responsive state $(s, x, 0)$, its state becomes the non-responsive state (\bar{s}, x, Y_{\max}) .
- If the state is non-responsive, the action has no effect.

We consider that if the same action a is sent to two appliances, their transitions are independent. Moreover, the effect of the actions is applied after the horizontal transitions. The actions are represented as thick dotted transitions in Figure 1.

The appliance remains in the same state with the probability that none of the transitions previously described occur: $1 - \gamma_s$ if $y = 0$ and $1 - \gamma_s - \theta_s$ if $y > 0$. These self-transitions are not represented in Figure 1 for the sake of readability. The transition probability of an appliance from state i to state j can be written:

$$\begin{aligned} \pi(i, j, \alpha) &= \mathbb{P}(m(t+1) = j | m(t) = i, a_i(t) = \alpha) \\ &= (1 - \alpha)\pi(i, j, 0) + \alpha\pi(i, j, 1). \end{aligned} \quad (1)$$

Some states of the appliance are undesirable, e.g., if the internal temperature of a fridge is too high or too low. We model this by associating a penalty with certain values of x . Specifically, as long as $x \in \{0 \dots X_{\max}\}$, this penalty is 0. When $x < 0$ or $x > X_{\max}$, the penalty is proportional to the gap between x and the desirable states. Explicitly, we define a penalty vector $\kappa = (\kappa_i)_{i \in \mathcal{M}}$ indexed by the states $i = (s, x, y)$:

$$\kappa_i = \begin{cases} 0 & \text{if } x \in \{0, \dots, X_{\max}\} \\ -\delta \cdot x / X_{\max} & \text{if } x < 0 \\ \delta \cdot (x - X_{\max}) / X_{\max} & \text{if } x > X_{\max} \end{cases}$$

for some positive constant $\delta > 0$.

When there is no DR, the appliance automatically switches *on* or *off* just before it reaches an undesirable state. In this specific scenario there is no penalty.

3.3 Mean-field Approximation

The role of each DR appliance is symmetrical. Hence, the state of the appliances is fully characterized at time t by an occupancy measure $M^N(t) = (M_i^N(t))_{i \in \mathcal{M}}$, where $M_i^N(t)$ is the fraction of appliances that are in state i at time t . We represent this occupancy measure as a row vector (of infinite dimension). Unless otherwise specified, initially, all actors know $M^N(0) = M_0$.

We consider that the number of flexible appliances N is large. It is shown in [7, Section 3.1] that, as N grows, the $M^N(t)$ goes to $M(t)$, which evolves as:

$$M_j(t+1) = \sum_{i \in \mathcal{M}} M_i(t)\pi(i, j, a_i(t)). \quad (2)$$

Moreover, as N goes to infinity, the optimal control for the mean field limit M is asymptotically optimal for the system with a finite number of appliances [7, 8].

We denote by Π^0 and Π^1 the two matrices:

$$\Pi^0(i, j) = \pi(i, j, 0) \text{ and } \Pi^1(i, j) = \pi(i, j, 1) - \pi(i, j, 0).$$

Using Equation (1), we can write the evolution of the occupancy measure as

$$M_j(t+1) = \sum_{i \in \mathcal{M}} M_i(t)\Pi^0(i, j) + \sum_{i \in \mathcal{M}} M_i(t)a_i(t)\Pi^1(i, j).$$

In what follows we assume that the occupancy measure $M(t)$ can be fully observed at t . Thus, we replace $a(t)$ by a more convenient control $U(t)$, $U_i(t) = M_i(t)a_i(t)$. Formally, we define the set X_U of processes $(U(t))$ such that $M(0) = M_0$ and for all t :

$$M(t+1) = M(t)\Pi^0 + U(t)\Pi^1, \quad (3)$$

$$U_i(t) \in [0; M_i(t)] \quad \forall i \in \mathcal{M} \quad (4)$$

The two controls a and U are equivalent in the following sense: a process $(M(t), M(0) = M_0)$ describing the evolution of the occupancy measure can be obtained via (2) using a control process $(a(t))$ if and only if there exists $U \in X_U$ such that $U_i(t) = M_i(t)a_i(t)$ for all t, i . These equations remain valid if the same actions $a(t)$ are sent to all appliances, or if each individual received an action. In that case, $a(t)$ is the average over all actions. See the discussion in [7, Section 5.1].

Finally, the power consumption and penalty, defined in §3.2, can be expressed as scalar products¹ with $M(t)$:

- The power consumption of flexible appliances is

$$F^a(t) = f^{da}(t) + F(t) = P_{\text{on}}M(t) \cdot \varsigma^\top, \quad (5)$$

where the row vector $\varsigma = (\varsigma_i)_{i \in \mathcal{M}}$ is such that $\varsigma_i = 1$ if i is an *on* state and $\varsigma_i = 0$ otherwise.

- The total penalty at time t is

$$M(t) \cdot \kappa^\top. \quad (6)$$

¹For two row vectors x and y , y^\top denotes the transpose of y and $x \cdot y^\top$ is the scalar product between x and y .

3.4 Market Actors

We now describe the market actors and their objectives in two different scenarios. In the first scenario, the consumer also manages the flexible demand. Based on the current or anticipated market prices and on the blackout cost she decides whether to anticipate or delay a flexible load. The amount of flexible load close to undesirable states also influences this decision. In the second scenario, the DR operator is a stand-alone market actor that directly buys energy from the market.

3.4.1 Scenario 1: Consumer Controls DR Actions

The consumer controls the quantity of energy $E_D(t)$ purchased on the real-time market at time t . In this scenario, she also manages the flexible demand $F(t)$ by influencing the state of the flexible appliances. The payoff of the consumer at time t is a sum of three terms:

- We consider linear utilities of satisfied non-controlled consumption $v \min\{D^a(t), E_D(t) + g^{da}(t) - F(t)\}$ and of satisfied controlled consumption $v' F^a(t)$.
- The disutility of a blackout is expressed as a sum of two terms: a linear term in the amount of unsatisfied non-controllable demand $-c^{bo}(D^a(t) + F(t) - E_D(t) - g^{da}(t))^+$ and a penalty of pushing controllable demand in undesirable states $-M(t) \cdot \kappa^\top$ defined in (6). As the consumer aggregates many individuals, she may curtail only part of their demand. For example, this corresponds to interrupting the service for a neighborhood in case of insufficient resources, rather than interrupting an entire city. We treat controllable demand differently, since it is actively interrupted or encouraged to consume.
- The amount of money spent on buying energy on the two markets at prices $P(t)$ and $p^{da}(t)$ is captured as $-P(t)E_D(t) - p^{da}(t)(g^{da}(t) + f^{da}(t))$.

To summarize, the total payoff of the consumer is

$$W_D^{tot}(t) := v \min\{D^a(t), E_D(t) + g^{da}(t) - F(t)\} + v' F^a(t) - c^{bo}(D^a(t) + F(t) - E_D(t) - g^{da}(t))^+ - M(t) \cdot \kappa^\top - P(t)E_D(t) - p^{da}(t)(g^{da}(t) + f^{da}(t)).$$

We strip away all the terms that are not controllable in real-time. Taking into account (5) and (6), the real-time component of the payoff is:

$$W_D(t) = M(t) \cdot (v' P_{on\zeta} - \kappa)^\top - P(t)E_D(t) - (v + c^{bo})(E_D(t) - D(t) - P_{on}M(t) \cdot \zeta^\top + f^{da}(t) + r^{da})^- \quad (7)$$

where $x^- = \max(-x, 0)$, for any real number x .

In the real-time market the consumer maximizes her expected payoff over the duration T of a day, given the scheduled day-ahead decisions. Her objective is

$$\arg \max_{\substack{E_D, U \in X_U \\ (\mathcal{F}_t)\text{-adapted}}} \mathcal{W}_D, \text{ where } \mathcal{W}_D := \mathbb{E} \left[\sum_{t=0}^T W_D(t) \right].$$

The supplier sells a quantity $E_S(t)$ on the real-time market, and produces $G(t)$ subject to the following constraints at each time t :

1. the sold quantity cannot exceed the produced quantity, *i.e.*, $E_S(t) \leq G(t) + \Gamma(t)$,
2. the real-time production has ramping limitations $\zeta^- < 0$ and $\zeta^+ > 0$: $\zeta^- \leq \frac{G(t') - G(t)}{t' - t} \leq \zeta^+, \forall t' > t$.

We denote the set of processes $(E_S(t), G(t))$ that satisfy these constraints by X_S . Thus, taking into account the marginal generation cost of the day-ahead production c^{da} and that of the real-time production c , the payoff of the supplier is:

$$W_S^{tot}(t) = (p^{da}(t) - c^{da})(f^{da}(t) + g^{da}(t)) + P(t)E_S(t) - cG(t).$$

Removing the terms that are not controllable in real-time, this payoff becomes

$$W_S(t) = P(t)E_S(t) - cG(t). \quad (8)$$

Finally, the supplier maximizes her expected payoff

$$\arg \max_{\substack{(E_S, G) \in X_S \\ (\mathcal{F}_t)\text{-adapted}}} \mathcal{W}_S, \text{ where } \mathcal{W}_S := \mathbb{E} \left[\sum_{t=0}^T W_S(t) \right].$$

Dynamic Competitive Equilibrium. We recall the notion of dynamic competitive equilibrium [4, 21].

DEFINITION 1. A dynamic competitive equilibrium $(P^e, E_D^e, E_S^e, G^e, U^e)$ is a set of (\mathcal{F}_t) -adapted price and control processes that satisfy:

$$(E_D^e, U^e) \in \arg \max_{E_D, U \in X_U} \mathcal{W}_D, \quad (9)$$

$$(E_S^e, G^e) \in \arg \max_{(E_S, G) \in X_S} \mathcal{W}_S, \quad (10)$$

$$E_D^e = E_S^e. \quad (11)$$

In the above definition, (9) means that (E_D^e, U^e) constitutes an optimal control from the selfish consumer's perspective. Similarly, (10) states that (E_S^e, G^e) is optimal from the supplier's perspective. Finally, (11) is the market constraint. Note that in (9), the consumer is not subject to the supplier's constraints and vice-versa for (10). See [4] for a discussion.

3.4.2 Scenario 2: Standalone DR operator

In this scenario, we consider that end-users are charged at two different retail prices per energy unit: one for the non-controllable appliances, and another one for the controllable appliances. The bill for the controllable appliances is collected by the DR operator. The rationale of the DR operator is that controllable appliances (e.g., fridges, boilers) are always operational and that they consume roughly a constant (predictable) quantity of energy per day. Hence, from the DR operator's viewpoint, it is as if consumers pay a constant subscription

fee per month for a provided service (e.g., to keep the fridge cool). The DR operator has the possibility of performing arbitrage on the real-time market thanks to the flexibility of the appliances.

The consumer controls only the quantity of purchased real-time energy $E_D(t)$ for serving the non-controllable demand. In this scenario, her payoff at each time t is:

$$W_D^{tot}(t) := v \min(D^a(t), E_D(t) + g^{da}(t)) + \\ - c^{bo}(D^a(t) - E_D(t) - g^{da}(t))^+ \\ - [P(t)E_D(t) + p^{da}(t)g^{da}(t)].$$

The real-time component of the consumer's payoff is

$$W_D(t) = -(v + c^{bo})(E_D(t) - D(t) + r^{da})^- - P(t)E_D(t). \quad (12)$$

The DR operator is an independent actor that controls the real-time consumption of controllable demands $F(t)$. On the real-time market, when $F(t) < 0$, she acts as a virtual energy supplier and gets paid for generating "nega-Watts". She generates revenue via arbitrage on the real-time market. Her payoff function is

$$W_F^{tot}(t) = -p^{da}(t)f^{da}(t) - P(t)F(t) + v'F^a(t) - M(t) \cdot \kappa^\top.$$

By Equation (5), $F(t) = P_{on}M(t)\kappa^T - f^{da}(t)$. Hence, her real-time controllable payoff is

$$W_F(t) = -P(t)P_{on}M(t) \cdot \zeta^\top + M(t) \cdot (v'P_{on}\zeta - \kappa)^\top \quad (13)$$

The DR operator maximizes her expected payoff subject to (\mathcal{F}_t) -adapted feasible controls $U \in X_U$, specifically:

$$\arg \max_{\substack{(U) \in X_U \\ (\mathcal{F}_t)\text{-adapted}}} \mathcal{W}_F, \text{ where } \mathcal{W}_F := \mathbb{E} \left[\sum_{t=0}^T W_F(t) \right].$$

The supplier behaves the same way as in the previous scenario and has the same payoff (8).

A similar definition of a dynamic competitive equilibrium can be written in this scenario. A competitive equilibrium is now the tuple $(P^e, E_D^e, F^e, E_S^e, G^e, U^e)$. The difference with respect to the previous case is that the consumer who was adjusting the amount of bought energy via the real-time flexible demand F^e now needs to pay for $F^e(t) = P_{on}M^e(t) \cdot \zeta^\top - f^{da}(t)$ at the price of the real-time market as a form of virtual energy generation. The market constraint (11) becomes

$$E_S^e(t) = E_D^e(t) + F^e(t) \text{ for all } t. \quad (14)$$

4. SOCIAL WELFARE THEOREM

In both presented scenarios we compute the social welfare as the sum of the actors' payoffs when the market clearing constraint is satisfied. In the case of a standalone DR operator (Scenario 2), the total payoff is the sum of Equations (8), (12) and (13). It is equal to

$$W_{tot}(t) := -(v + c^{bo})(E_D(t) - D(t) + r^{da})^- - cG(t) \\ + M(t) \cdot (v'P_{on}\zeta - \kappa)^\top + P(t)(E_S(t) - E_D(t) - F(t)).$$

When the market constraints (14) are satisfied, the bought energy is equal to the sold energy at each time step. Hence, the price disappears from the total welfare.

A socially optimal allocation maximizes the total expected payoff. As $W_{tot}(t)$ is increasing in $E_S(t)$, we can further simplify (15). Taking into account the constraints on $E_S(t)$, we get that $E_S(t) = G(t) + \Gamma(t)$ maximizes $W_{tot}(t)$. Using (14) to express E_D and denoting the random (uncontrolled) component of the social welfare by $Z(t) := \Gamma(t) - D(t) + f^{da}(t) + r^{da}$, a socially optimal allocation maximizes the total expected payoff:

$$W'_{tot}(t) = -(v + c^{bo})(G(t) + Z(t) - P_{on}M(t) \cdot \zeta^\top)^- \\ + M(t) \cdot (v'P_{on}\zeta - \kappa)^\top - cG(t). \quad (16)$$

We call (G^*, U^*) a *socially optimal* allocation if it is a solution to the following problem:

$$\arg \max_{G \in X_S, U \in X_U \text{ are } (\mathcal{F}_t)\text{-adapted}} \mathbb{E} \left[\sum_{t=0}^T W_{tot}(t) \right] \quad (17)$$

A direct computation shows that the total welfare is the same in Scenario 2. This implies that the socially optimal allocations are the same in both scenarios.

The following important result holds. It is proved in Appendix A.

THEOREM 1. *If $(P^e, E_D^e, E_S^e, G^e, U^e)$ is a dynamic competitive equilibrium in Scenario 1, then the allocation (G^e, U^e) is socially optimal.*

Similarly, if $(P^e, E_D^e, F^e, E_S^e, G^e, U^e)$ is a dynamic competitive equilibrium in Scenario 2, then the allocation (G^e, U^e) is socially optimal.

5. NUMERICAL SOLUTION OF THE PROBLEM

In the rest of this paper we compute numerically the equilibrium prices and welfare given by Theorem 1. To this end we use a specific model for the random component $Z(t)$ that is both tractable and realistic.

5.1 Branching Trajectory Forecasts

We use discrete time and evaluate the model of Section 3 over a period of 24h. Each time slot represents a five minute interval [18]. Hence, the real-time market lasts for $T = 288$ time slots.

Inspired by the methodology of Pinson *et al.* [17], we assume that the forecast error process Z can be represented by a finite number of 2^τ trajectories. For a trajectory $\omega \in \{1 \dots 2^\tau\}$, the value of the forecast error at time t is denoted $Z^\omega(t)$. To model the fact that the forecast for the near future (e.g., the next hour) is more accurate than the forecast for a distant future (e.g., 12 hours from the present), we consider that the possible trajectories for the forecast error coincide initially before separating in distinct branches at certain moments

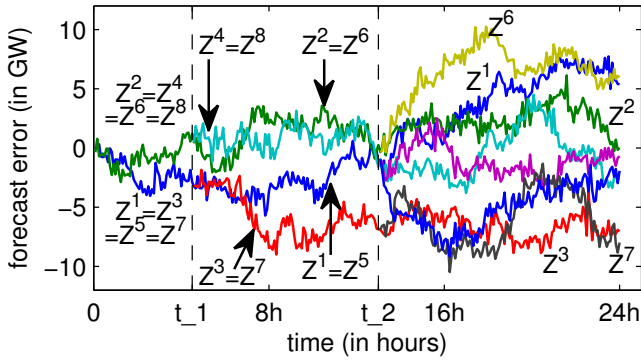


Figure 2: Example of possible trajectories of the forecast errors. To ease the presentation, this example has only eight trajectories. The numerical results of Section 6 use 512 trajectories. These trajectories are generated using data from the UK (see §6.1).

$\{t_1 \dots t_\tau\}$. Thus, over time the number of branches increases exponentially: between times t_{k-1} and time t_k , there are 2^k possible such branches of the forecast errors. We consider that initially, between times $t_0 = 0$ and t_1 , there are two equally likely branches: the odd trajectories and the even trajectories.

At each time t_k , each group of trajectories separates into two equiprobable groups of trajectories. Formally, between t_{k-1} and t_k , all trajectories ω that have the same remainder modulo 2^k share the same values:

$$\forall \omega < 2^\tau - 2^k, \forall t \in [t_{k-1}; t_k] : Z^\omega(t) = Z^{\omega+2^k}(t). \quad (18)$$

An example of trajectory is represented in Figure 2. For more clarity, we only represent eight trajectories. In the rest of the paper, the numerical evaluations use 512 trajectories. On this figure, we see that before $t_1 := 4.5h$, there are two possible forecast errors (the “odd” and “even” trajectories). Between t_1 and $t_2 = 12h$, there are four branches. After t_2 , all trajectories are distinct.

These trajectories are generated using the covariance of real forecast errors in the UK. The methodology and the algorithm used will be described in §6.1.

5.1.1 Observability assumptions

We assume that at each of the branching instants t_k , players can observe the forecast error and can know on which branch the system is evolving at t_k . However, they cannot predict which branch the forecast errors will follow after t_k . For example, a time $t_1 = 4.5h$, players can observe if the forecast error is one of the four odd (Z^1, Z^3, Z^5, Z^7) or one of the four even trajectories. The only information about the future that they have is that each trajectory is *equiprobable*.

The set of control variables (G, U, \dots) is restricted to the set of admissible controls, which are the sequence of causal decisions: a decision taken at time t can use

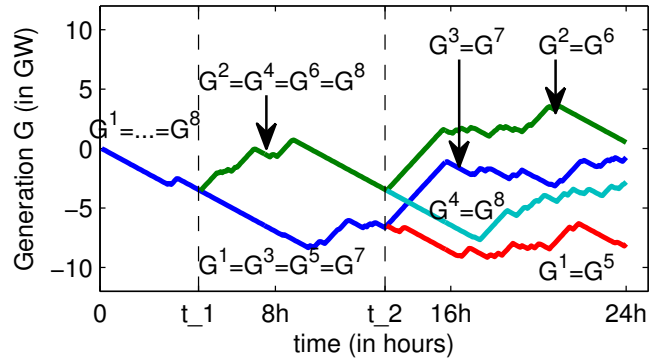


Figure 3: Example of an admissible control for G corresponding to the forecast errors of Figure 2.

the information available at time t as well as the statistics about future trajectories but not the exact future realization of the forecast errors. Using our observation model, this means an admissible control for G and U is a sequence of values $G^\omega(t)$ and $U^\omega(t)$ that satisfies the generation and DR constraints (3) such that for all $\omega < 2^\tau - 2^k$ and all $t \in [t_k; t_{k+1}]$:

$$G^\omega(t) = G^{\omega+2^k}(t); \quad U^\omega(t) = U^{\omega+2^k}(t). \quad (19)$$

For example, if the forecast error is represented by the trajectories of Figure 2, an example of admissible control for G is depicted in Figure 3. This control is causal. At time $t_0 = 0$ it knows that the forecast error may follow one of the two branches. The chosen control for the interval $[0; t_1]$ is a compromise that yields the best average performance. At the observation point t_1 , players observe whether the branch that was followed was that of an odd or of an even trajectory. Since each branch separates again in two branches, there are two causal sequences of controls for the interval $[t_1; t_2]$ that can be applied, depending on the observations made at time t_1 . Similarly, at time t_2 the four possible observations lead to four possible decisions.

5.2 Casting the Stochastic Problem into Coupled Deterministic Problems using ADMM

To compute the equilibrium, we use an optimization algorithm called the *alternating direction method of multipliers* (ADMM) [2]. ADMM is an iterative process to solve deterministic constrained maximization problems with concave² objective and linear constraints:

$$\max_{x \in \mathbb{R}^d, z \in \mathbb{R}^e} f(x) + g(z) \text{ subject to } Ax + Bz = c, \quad (20)$$

where f and g are two concave functions that take real values or $-\infty$; A and B are two matrices of size $d \times d_x$ and $d \times d_z$; and c is a vector of size d .

²ADMM is usually presented as a solution to solve convex minimization problems. In this paragraph, we keep the concave maximization formulation of Section 3.

Let $\rho > 0$. For each $(x, z, \lambda) \in \mathbb{R}^{d_x+d_z+d}$, the augmented Lagrangian $L_\rho(x, z, \lambda)$ is defined as:

$$L_\rho(x, z, \lambda) = f(x) + g(z) - \frac{\rho}{2} \|Ax + Bz + \lambda - c\|^2. \quad (21)$$

ADMM performs the following iterations at step $k + 1$:

$$x^{k+1} := \arg \max_x L_\rho(x, z^k, \lambda^k) \quad (22)$$

$$z^{k+1} := \arg \max_z L_\rho(x^{k+1}, z, \lambda^k) \quad (23)$$

$$\lambda^{k+1} := \lambda_k + Ax^{k+1} + Bz^{k+1} - c, \quad (24)$$

where x^1, z^1 and λ^1 have been initialized to any vector of the proper size.

It is shown in [2] that under very general assumptions, the ADMM iterations converge to a solution of the optimization problem (20), regardless of the initial conditions. ADMM also converges when f or g are not strictly concave and when they take infinite values.

The social maximization problem (17), is a stochastic optimization problem where the expectation is taken over all possible forecast processes. As we assume that the forecast errors can be represented by a finite number of trajectories, this problem can be written

$$\arg \max_{G, U \in \text{admissible}} \sum_{\omega} W_{tot}(G^\omega, U^\omega) \cdot \text{proba}(\omega), \quad (25)$$

where the set admissible controls G, M denotes controls that satisfy generation constraints X_S , DR constraints X_U , as well as causality constraints (19).

In particular, Equation (19) implies that the decisions taken at a time t depend only on the information up to time t . As such, they are a compromise between the possible optimal decisions taken by a hypothetical omniscient controller that can observe the future.

5.3 Equilibrium Computation using ADMM

The formulation of ADMM, given by Equation (20), requires the optimization space to be partitioned into two groups of variables with no objective functions that depends on both groups, like $h(x, z)$. To perform this separation, we enlarge our optimization space. In addition to the original variables $G^\omega(t)$, $M^\omega(t)$ and $U^\omega(t)$, we add, for each $t < T$ and ω , the variables $\overline{M}^\omega(t)$, $\overline{G}^\omega(t)$, $\underline{G}^\omega(t)$, $E_D^\omega(t)$, $E_F^\omega(t)$, $E_S^\omega(t)$, $\underline{U}^\omega(t)$, $\overline{U}^\omega(t)$ and the linear constraints given by Equation (19) and

$$M(t+1) = \overline{M}(t)\Pi^0 + U(t)\Pi^1; \quad M(t) = \overline{M}(t); \quad (26)$$

$$\underline{U}(t) = U(t); \quad \overline{U}(t) = \overline{M}(t) - U(t); \quad E_F(t) = P_{on} M(t) \cdot \varsigma^\top \quad (27)$$

$$\overline{G}(t) = G(t) = \underline{G}(t); \quad G(t) = E_S(t). \quad (28)$$

The first group of variables (the ‘‘X variables’’) contains the variables $M^\omega(t)$, $\underline{U}^\omega(t)$ and $\overline{U}^\omega(t)$, $G^\omega(t)$. The corresponding objective function $f(M, \underline{U}, \overline{U}, G)$ is:

$$\sum_{t, \omega} M^\omega(t) \cdot (v' P_{on} \varsigma - \kappa)^\top - c G^\omega(t) + \text{Pos}(\underline{U}^\omega, \overline{U}^\omega), \quad (29)$$

where $\text{Pos}(x, x') = 0$ if all coordinates of the vectors x and x' are non-negative and $-\infty$ otherwise.

The second group of variables contains the variables $\overline{M}, U, \overline{G}, \underline{G}, E_D, E_F, E_S$. The corresponding objective function $g(U, \overline{M}, \overline{G}, \underline{G}, E_D, E_F, E_S)$ is equal to:

$$\begin{aligned} & \sum_{t, \omega} -(v + c^{bo}) E_D(t)^- + \text{Pos}(\zeta^+ - \underline{G}^\omega(t+1) + \overline{G}^\omega(t)) \\ & + \text{Pos}(\underline{G}^\omega(t+1) - \overline{G}^\omega(t) - \zeta^-) \\ & + \text{Equal0}(E_S^\omega(t) + Z^\omega(t) - E_D^\omega(t) - E_F^\omega(t)) \end{aligned} \quad (30)$$

where $\text{Equal0}(x) = 0$ if all coordinates of x equal 0 and $-\infty$ otherwise.

5.4 Analysis of the ADMM Algorithm

When the constraints given by Equations (26-27) and the two terms $\text{Pos}(\cdot)$ of Equation (29) are satisfied, the evolution of $M(t)$, given by Equation (3) is respected. Similarly, the constraints given by (28) and the two term $\text{Pos}(\cdot)$ of Equation (30) ensure that the ramping constraints of G are satisfied.

This implies that our algorithm converges to the set of socially optimal allocations. Moreover, it also computes competitive equilibria. The following theorem summarizes these results. It is proved in Appendix B.

THEOREM 2. *Let E_D^k, E_S^k, G^k, U^k be the values after the k th iteration and let $(\lambda_G^k)^\omega(t)$ be the multipliers associated with the constraints $E_S^\omega(t) = G^\omega(t)$. Then:*

- (i) *There exists P such that $\lim_{k \rightarrow \infty} \rho(\lambda_G^k)^\omega(t) = P^\omega(t)$.*
- (ii) *Any subsequence of E_D^k, E_S^k, G^k, U^k has a subsequence that converges. Let $E_D^\infty, E_S^\infty, G^\infty, U^\infty$ be its limit. Then*
 - *$(P, E_D^\infty + E_F^\infty, E_S^\infty, G^\infty, U^\infty)$ is a competitive equilibrium for Scenario 1.*
 - *$(P, E_D^\infty, E_F^\infty, E_S^\infty, G^\infty, U^\infty)$ is a competitive equilibrium for Scenario 2.*
- (iii) *If the maximization algorithm has a unique solution, the algorithm converges to this solution. If there are multiple optimal solutions, they form a convex set and the algorithm converges to this set (for the Euclidean distance).*

Remark. Even if we consider piecewise-linear objective functions, ADMM only requires the objective to be concave. Hence, the same algorithm can be applied to more general cost functions and Theorem 2 also holds.

6. NUMERICAL EVALUATIONS

6.1 Parameters and Trajectory Generation

As in [9], we use $\zeta^+ = 1\text{GW}/\text{h}$, $\zeta^- = 3\zeta^+$. We set the ratio between the blackout cost and the generation cost to $(v + c^{bo})/c = 10$. The values of v , c^{bo} and c are normalized so the flexibility provided by fridges L parametrized below leads to a welfare gain of 100 when P_{on} is large.

We compare three cases of appliances [13]:

- **Fridges L** (large inertia fridges) – the *on* period is $X_{\max}/\gamma_{\text{on}} = 60\text{min}$ and the off period is 120min, with a mini-cycle prevention period $Y_{\max}/\theta_{\text{on}} = 0$ or 20min. It corresponds to $X_{\max} = 6$, $\gamma_{\text{on}} = \theta_{\text{on}} = \theta_{\text{off}} = 0.5$, $\gamma_{\text{off}} = 0.25$.
- **Fridges S** (small inertia fridges) – same as *Fridges L* but with $X_{\max} = 3$. The only difference is that the *on* period lasts 30min and the *off* period lasts 60min.
- **Boilers** – the *on* period is 4h and the off period is 20h. It corresponds to $X_{\max} = 12$, $\gamma_{\text{on}} = 0.25$, $\gamma_{\text{off}} = 0.05$ and $\theta_{\text{on}} = \theta_{\text{off}} = 0.5$.

The random transitions of our model of appliance account for various random evolutions of the internal temperature, for example, due to fridge doors opening. The second fridge model corresponds to a fridge with a smaller inertia. It consumes the same average power but has less flexibility, typically because the temperature dead-band is narrower. The ratio between the duration of *off* and *on* periods is 2 for fridges and 5 for boilers. Our numerical evaluation shows that, qualitatively, the behaviors of the three models are the same.

6.1.1 Trajectory Generation using UK Data

To construct wind forecast errors, we use wind production and day-ahead wind production forecast in the time interval from June 2009 to April 2012, obtained from the BMRA data archive (elxonportal.co.uk). We normalize the values of production and forecast to maintain a constant wind capacity³ of 26GW all throughout the 1300 days. We obtain 1300 samples of forecast error trajectories $\{\varepsilon_d(t) : t \in [0; 24h]; d \in \{1 \dots 1300\}\}$. $\varepsilon_d(t)$ is the error for the day d at time $t \in [0; 24h]$.

A generic method to generate forecast error trajectories has been introduced in [17]. To generate one trajectory, the authors of [17] compute the covariance Σ of the process ε_d (a 288×288 matrix in our case) and then generate a multivariate normal vector of covariance Σ . We adapt their method to our case of branching forecast. Algorithm 1 generates a set of trajectories such that each one has covariance Σ , and such that two trajectories that share a branching point depend on each other only via their values before the branching point. These properties are summarized in Theorem 3.

THEOREM 3. *Assume that Σ has full rank. Then, Algorithm 1 generates a set of trajectories such that:*

- individually, each trajectory Z^ω is a multivariate normal vector of covariance Σ ;*
- If a subset of trajectories $Z^{i_1} \dots Z^{i_k}$ is fixed, then the distribution of another trajectory Z_i is the distribution of a multivariate normal vector of covariance Σ subject to respecting Equation (18).*

This theorem is proved in Appendix C. The proof can be adapted if Σ is not full rank but it is more technical.

³This scenario is envisioned for the UK in 2020, where 20% of the total electricity consumption is covered by wind.

Input: Covariance matrix Σ , sequence $t_1, \dots, t_{\tau-1}$

Output: Trajectories Z^1, \dots, Z^{2^τ}

```

1  $A \leftarrow$  lower Cholesky decomposition of  $\Sigma$ ;  $t_0 \leftarrow 0$ ;
2  $N_1 \leftarrow$  vector of  $T$  i.i.d. normal random variables;
3 for  $k \leftarrow 0$  to  $\tau - 1$  do
4   for  $\omega \leftarrow 1$  to  $2^k$  do
5      $N^{\omega+2^k} \leftarrow [N^\omega(1 \dots t_k); \hat{N}^{\omega+2^k}]$ , where
6      $\hat{N}^{\omega+2^k}$  is a vector of  $T - t_k$  i.i.d. normally
7     distributed random variables;
8   end
9 end
10 for  $\omega \leftarrow 1$  to  $2^\tau$  do
11    $Z^\omega \leftarrow A \times N^\omega$ ;
12 end

```

Algorithm 1: Generation of 2^τ branching trajectories satisfying Theorem 3. The notation $N^\omega(1 \dots t_k)$ denotes the first t_k elements of the vector N^ω . All vectors are column vectors. $[A; B]$ is the concatenation of two column vectors to form one column vector.

As a consequence, all trajectories $Z^1 \dots Z^{2^\tau}$ are equally likely. Moreover, as τ grows, the process obtained by following the branching trajectories “approaches” a normally distributed random process with covariance Σ .

6.2 Impact of DR Penetration and of the Nature of Elastic Loads

For a given power flexibility P_{on} , we denote $\mathcal{W}^*(P_{\text{on}})$ the optimal value of social welfare in Equation (17). $\mathcal{W}^*(0)$ corresponds to the social welfare without DR. For $P_{\text{on}} > 0$, the optimal welfare $\mathcal{W}^*(P_{\text{on}})$ depends on the type of appliance, while $\mathcal{W}^*(0)$ does not.

In Figure 4 we plot the “relative” social welfare, *i.e.*, the difference $\mathcal{W}^*(P_{\text{on}}) - \mathcal{W}^*(0)$ as a function of P_{on} . We compare the three types of appliances. For *fridges L*, we plot both the case without mini-cycle avoidance ($Y_{\max} = 0$) and with mini-cycle avoidance ($Y_{\max} > 0$). As expected, the welfare is then increasing and concave in P_{on} . We observe that, as P_{on} grows, the social wel-




Figure 4: Relative social welfare $\mathcal{W}^*(P_{\text{on}}) - \mathcal{W}^*(0)$ as a function of P_{on} – the available power capacity of DR – for the three types of appliances.

fare increases and tends to stabilize at the value 100, regardless of the appliance type or the value of Y_{\max} , although for a fixed capacity of DR, increasing the mini-cycle avoidance parameter Y_{\max} diminishes its benefits. Moreover, for a given appliance type, the gain of using demand-response is almost linear in P_{on} for small values of the installed capacity P_{on} . In view of the similarity of the appliance responses, for the rest of the numerical evaluation we only show the results for the *fridges L*.

6.3 Non-Observability of DR

The previous figure assumes that the day-ahead market is able to predict without error the consumption of demand-response appliances. In practice, this can be done because each appliance evolves mainly independently of the others. As a result, the relative error of current load-prediction techniques is less than 1% [6]. However, one of the concerns of demand-response is that the presence of a large player that controls many appliances will introduce synchronization between devices, which will make the behavior of the system less predictable. For example, a small congestion can cause lots of DR loads to be delayed, which can cause a bigger problem in the future because these delayed loads were not forecast at this date [14].

In this section, we show that when the day-ahead market cannot fully observe the demand-response state, adding too much demand-response can be harmful, even if the real-time market has full information about the states of appliances. More precisely, we assume that:

- the day-ahead market cannot observe $M(0)$, the state of the demand-response at time 0. It assumes that $f^{da}(t)$ is equal to the overall average consumption of all demand-response appliances and plans the day-ahead generation $g^{da}(t)$ accordingly.
- the actor that controls flexible loads can observe the states of all DR appliances in real time.

We plot the relative social welfare $\mathcal{W}^*(P_{\text{on}}) - \mathcal{W}^*(0)$




Figure 5: Relative social welfare $\mathcal{W}^*(P_{\text{on}}) - \mathcal{W}^*(0)$ as a function of available power flexibility P_{on} when the day-ahead market cannot observe DR’s initial state, compared to the case when the day-ahead can observe this initial state. The x-axis is between 0 and 100GW.

for this scenario in Figure 5. When the flexible power is small ($P_{\text{on}} < 10\text{GW}$), the benefit of demand-response is close to the one with full information. However, when the installed demand-response capacity is large, adding more demand response decreases the total welfare. The non-observability problem is probably hidden today as the installed capacity of DR is small but this is a threat for the future. The results for *fridges S* and *boilers* are very similar and are not shown here.

6.4 Comparison with Energy Storage Systems

Demand-response is often viewed as a virtual energy storage system, which is potentially cheaper than real storage systems such as batteries. Instead of charging and discharging a battery, DR allows to anticipate or delay the consumption of appliances. In this section, we compare energy storage and flexible loads. We highlight two differences: for low penetration, storage provides more flexibility because it reacts faster. However, at high penetration, storage is less efficient because of the energy losses at each charging/discharging cycle.

We consider a storage model like in [9]: the storage has an energy capacity of B_{\max} and a maximum charging and discharging power capacity C_{\max} and D_{\max} . The cycle efficiency is η : only a fraction η of the stored energy can be retrieved. To perform a fair comparison of DR and storage, we consider a storage system that has the same flexible power: we set $C_{\max} = D_{\max} = P_{\text{on}}/2$. We chose B_{\max} to reflect the quantity of the energy stored in DR appliances. Recall that when an appliance is *on*, its internal state x increases with probability γ_{on} . When *off*, its internal state decreases with probability γ_{off} . Thus, the average consumption of the DR appliances is $P_{\text{on}}\gamma_{\text{off}}/(\gamma_{\text{on}} + \gamma_{\text{off}})$. The internal state of an appliance that avoids the undesirable states oscillates between 0 and X_{\max} . We set B_{\max} equal to the difference in consumption between the case where all DR appliances have an internal state of X_{\max} and the case where they are all in state 0. Hence, we take:

$$B_{\max} = \frac{P_{\text{on}}X_{\max}}{\gamma_{\text{on}} + \gamma_{\text{off}}} \times 5\text{min} \quad (31)$$

With the *fridges L* model, this implies $B_{\max} = \frac{P_{\text{on}}}{2} \text{GWh}$.

We adapt our ADMM algorithm of §5.3 for the case where a storage system replaces DR. The results are reported in Figure 6. We plot the welfare as a function of the flexible power P_{on} in four scenarios: presence of DR but no storage (two dashed curves); or presence of storage but no DR (two solid curves).

We observe that, at low power capacities, storage provide higher gain than DR. Two reasons explain this:

- the fatigue effect: because of internal state constraints, some appliances cannot be switched anymore.
- the mini-cycle effect $Y_{\max} > 0$: appliances that have just been switched cannot be switched instantaneously.




Figure 6: Social welfare as a function of available capacity of demand-response P_{on} or of storage. The storage flexibility is similar to the one of DR: $C_{\text{max}} = D_{\text{max}} = P_{\text{on}}/2$ and $B_{\text{max}} = (P_{\text{on}}/2)\text{GWh}$.

For large capacities, the situation is reversed. The welfare of DR and that of idealized storage saturate at the same value. This value is strictly larger than the one of 70%-efficient storage. For large power capacity, DR outperforms realistic storage (typical efficiency 70 – 90%).

6.5 Price Equilibrium

When there is no demand-response or storage, it has been shown in [21] that the price process oscillates between 0 and the choke price $v + c^{bo}$. In this case it is never equal to the marginal production cost. The situation is different in the presence of storage. In [9], we exhibit two different situations, depending on the charge-discharge efficiency of storage η :

- When $\eta = 100\%$ (perfect storage), the prices do in fact concentrate on the marginal production cost c as the quantity of storage gets large.
- When $\eta < 100\%$ (realistic storage), the prices do not concentrate on c , but exhibit two modes around $c\sqrt{\eta}$ and $c/\sqrt{\eta}$.

These results were obtained assuming that the forecast error is stationary and can be represented by a Brownian motion with stationary increments. We simulate the same scenarios using our more realistic non-stationary branching forecast model and our ADMM implementation. The price distributions are reported in Figure 7. They confirm that the results of [9] are robust to the forecast model: for $\eta = 1$, the prices concentrate on $c = 1$. For $\eta < 1$, the price distribution exhibit two modes around the values $\sqrt{\eta} \approx 0.89$ and $1/\sqrt{\eta} \approx 1.12$.

This price spread is explained by the fact that only a fraction η of the stored energy can be retrieved. As a consequence, a storage owner will store energy only if it can be sold at a $1/\eta$ times larger price. In the case of DR, the energy is virtually stored or retrieved by delaying or anticipating consumption. The corresponding prices are reported in Figure 8. It confirms that the system behaves like a 100% efficient storage system (like in §6.4): even when $Y_{\text{max}} > 0$, the prices concentrate on the marginal production cost c .




Figure 7: Distribution of prices for a system equipped with storage of capacity $C_{\text{max}} = 5\text{GW}$, $B_{\text{max}} = 5\text{GWh}$ and charge-discharge efficiency $\eta = 1$ or 0.8 . For readability, we scale the values c, v and c^{bo} to $c = 1$.

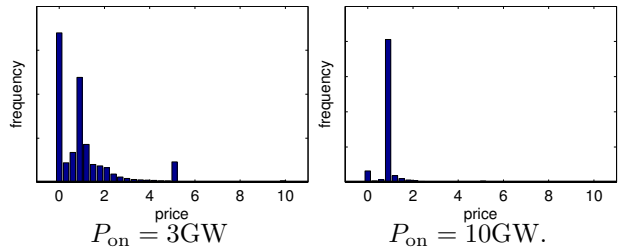


Figure 8: Distribution of prices for a system equipped with DR appliances *Fridges L* with $Y_{\text{max}} = 20\text{min}$. For readability, we scale the values c, v and c^{bo} to $c = 1$.

6.6 Incentive to Install Demand Response or Storage

The storage/demand-response operator makes money via price arbitrage. An independent storage operator buys cheap energy and resells it when prices rise. Similarly, a demand-response operator tries to shift the appliances' consumption to when the energy is inexpensive. The flexibility brought by demand-response or by storage results in more concentrated prices. In [9], we have shown that, to maximize their revenue, independent storage operators will underdimension their storage system. In this section, we show that this also holds for DR operators.

We simulate the same scenarios as in Figure 6 and we compute the expected revenue of the DR operator or of the storage operator as a function of the flexible power capacity. In all cases, this gain reaches a maximum and then decreases. This maximum is attained at $P_{\text{on}} \approx 1\text{GW}$ for batteries and *fridges L* with

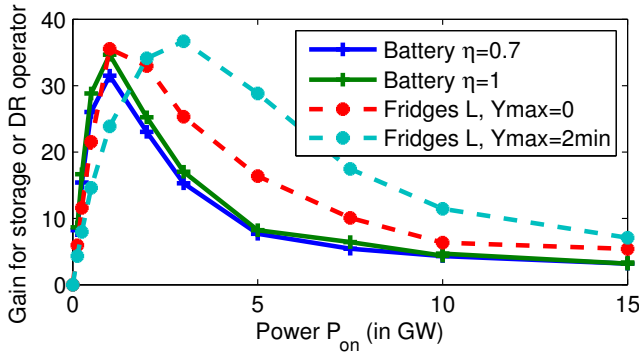


Figure 9: Revenue for the DR or storage operator as a function of the installed capacity. The studied scenarios are the same as in Figure 6.

$Y_{\max}=0$. When $Y_{\max}=20\text{min}$, this maximum is reached at $P_{\text{on}} \approx 3\text{GW}$. More problematically, the maximum is attained at a value that is far from optimal for the social welfare: This entails that even if we disregard installation costs, the operator seeking to maximize her revenue will not deploy more than 1GW (or 3GW) of DR capacity. In Figure 4, we observe that these values of P_{on} only lead to a gain representing 40 – 60% of the potential welfare benefit of DR/storage. Hence, to maximize her revenue, a DR operator’s incentive is to deploy a suboptimal capacity from a social perspective.

7. CONCLUSION

In this paper we studied the effects of demand-response in a real-time electricity market. We modeled a population of flexible appliances that can anticipate or delay their consumption. This model accounts for fatigue via undesirable states and enforces mini-cycle avoidance via an additional dimension in the state space. We showed that when demand-response is traded by price-taking actors, the dynamic competitive equilibrium is socially efficient. We considered a realistic and tractable trajectory forecast error model where trajectories separate into branches over time. We parametrized the model using real wind forecasts from the UK. We solved numerically the model by using ADMM. We showed that if the appliance states are unknown, too much demand-response can be detrimental. We compared demand-response to an energy storage system and concluded that a sufficiently large amount of demand-response is as beneficial as a perfect storage system. Demand-response relaxes the ramping constraints of generation and smooths prices, that concentrate around the marginal generation cost.

8. REFERENCES

- [1] M. H. Albadi and E. El-Saadany. A summary of demand response in electricity markets. *Electric Power Systems Research*, 78(11):1989–1996, 2008.
- [2] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends® in Machine Learning*, 3(1):1–122, 2011.
- [3] S. P. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [4] I. Cho and S. Meyn. Efficiency and marginal cost pricing in dynamic competitive markets with friction. *Theoretical Economics*, 5(2):215–239, 2010.
- [5] A. J. Conejo, J. M. Morales, and L. Baringo. Real-time demand response model. *Smart Grid, IEEE Transactions on*, 1(3):236–242, 2010.
- [6] E. A. Feinberg and D. Genethliou. Load forecasting. In *Applied mathematics for restructured electric power systems*, pages 269–285. Springer, 2005.
- [7] N. Gast and B. Gaujal. A mean field approach for optimization in discrete time. *Discrete Event Dynamic Systems*, 21(1):63–101, 2011.
- [8] N. Gast, B. Gaujal, and J.-Y. Le Boudec. Mean field for markov decision processes: from discrete to continuous optimization. *Automatic Control, IEEE Transactions on*, 57(9):2266–2280, 2012.
- [9] N. Gast, J.-Y. Le Boudec, A. Proutière, and D.-C. Tomozei. Impact of storage on the efficiency and prices in real-time electricity markets. In *Proceedings of the fourth international conference on Future energy systems*, pages 15–26. ACM, 2013.
- [10] J.-H. Kim and A. Shcherbakova. Common failures of demand response. *Energy*, 36(2):873–880, 2011.
- [11] D. S. Kirschen, G. Strbac, P. Cumperayot, and D. de Paiva Mendes. Factoring the elasticity of demand in electricity prices. *Power Systems, IEEE Transactions on*, 15(2):612–617, 2000.
- [12] S. Koch, J. L. Mathieu, and D. S. Callaway. Modeling and control of aggregated heterogeneous thermostatically controlled loads for ancillary services. In *Proc. PSCC*, pages 1–7, 2011.
- [13] G. Koutitas and L. Tassiulas. A delay based optimization scheme for peak load reduction in the smart grid. In *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, page 7. ACM, 2012.
- [14] J.-Y. Le Boudec and D.-C. Tomozei. Satisfiability of elastic demand in the smart grid. In *ENERGY 2011, The First International Conference on Smart Grids, Green Communications and IT Energy-aware Technologies*, pages 163–168, 2011.
- [15] S. Meyn, P. Barooah, A. Busic, and J. Ehren. Ancillary service to the grid from deferrable loads: the case for intelligent pool pumps in florida. *ACE (GW)*, 100(50):50, 2013.
- [16] A.-H. Mohsenian-Rad and A. Leon-Garcia. Optimal residential load control with price prediction in real-time electricity pricing environments. *Smart Grid, IEEE Transactions on*, 1(2):120–133, 2010.
- [17] P. Pinson, H. Madsen, H. A. Nielsen, G. Papaefthymiou, and B. Klöckl. From probabilistic forecasts to statistical scenarios of short-term wind power production. *Wind energy*, 12(1):51–62, 2009.
- [18] Pjm operational data, 2013. <http://www.pjm.com/markets-and-operations/energy/real-time/lmp.aspx>.
- [19] K. Spees and L. B. Lave. Demand response and electricity market efficiency. *The Electricity Journal*, 20(3):69–85, 2007.
- [20] C.-L. Su and D. Kirschen. Quantifying the effect of demand response on electricity markets. *Power Systems, IEEE Transactions on*, 24(3):1199–1207, 2009.
- [21] G. Wang, M. Negrete-Pincetic, A. Kowli, E. Shafiepoorfar, S. Meyn, and U. Shanbhag. Dynamic competitive equilibria in electricity markets. In *Control and Optimization Theory for Electric Smart Grids*. Springer, 2011.
- [22] R. Wilson. Architecture of power markets. *Econometrica*, 70(4):1299–1340, 2003.

APPENDIX

A. PROOF OF THEOREM 1

The proof of this result is similar to the one of [9, Theorem 2] and [21]. We recall the main steps here. We give the proof for Scenario 1. The proof for Scenario 2 is similar, with three actors instead of two.

Let $\mathcal{W}(G, U)$ be the social welfare. For a given price P , we denote $\mathcal{W}_D(E_D, U, P)$ the welfare of the consumer when she takes the decisions E_D, U and by $\mathcal{W}_S(E_S, G, P)$, the welfare of the supplier for the decisions E_S, G .

We consider the social optimization problem (17) and we relax the constraint $E_S = E_D$. Denoting the corresponding Lagrange multiplier by P , the Lagrangian $\mathcal{L}(E_D, E_S, U, G, P)$ is equal to:

$$\begin{aligned} \mathcal{W}(G, U) + \mathbb{E} \left[\sum_{t=0}^t P(t)(E_D(t) - E_S(t)) \right] \\ = \mathcal{W}_D(E_D, U, P) + \mathcal{W}_S(E_S, G, P). \end{aligned}$$

When the constraint $E_D = E_S$ is respected, the social welfare $\mathcal{W}(G, U)$ equals $\mathcal{W}_D(E_D, U, P) + \mathcal{W}_S(E_S, G, P)$.

Assume that $(E_D^e, E_S^e, G^e, U^e, P^e)$ is a competitive equilibrium. It follows that

$$\begin{aligned} \sup_{G, U} \mathcal{W}(G, U) &= \sup_{\substack{G, U, E_S, E_D \\ \text{s.t. } E_S = E_D}} \mathcal{L}(E_D, E_S, G, U, P^e) \\ &\leq \sup_{G, U, E_S, E_D} \mathcal{L}(E_D, E_S, G, U, P^e) \\ &= \sup_{E_D, U} \mathcal{W}_D(E_D, U, P^e) + \sup_{E_S, G} \mathcal{W}_S(E_S, G, P^e) \\ &= \mathcal{W}(E_D^e, U^e, P^e) + \mathcal{W}_S(E_S^e, G^e, P^e) \\ &= \mathcal{W}(G^e, U^e). \end{aligned}$$

The last equality holds because a competitive equilibrium implies $E_D^e = E_S^e$. The one before last equality holds because each player maximizes his social welfare.

Hence, the allocation (G^e, U^e) is socially optimal. \square

B. PROOF OF THEOREM 2

B.1 Convergence of ADMM

The problem formulated in §5.3 is a maximization problem with a concave objective and linear constraints. Moreover, the problem is feasible, *i.e.*, there exists a allocation such that all constraints are satisfied (for example $G^\omega(t) = G(0)$, and $M^\omega(t+1) = M^\omega(t)\Pi^1$). Hence, the Slater's conditions are satisfied and the unaugmented Lagrangian has a saddle point [3, Chapter 5].

We use the generic notation of §5.2 and denote the variables after iteration k by x^k, z^k, λ^k . By [2, §3.2.1], the existence of the saddle point implies:

(I) the variables x^k, z^k approach feasibility:

$$\lim_{k \rightarrow \infty} Ax^k + Bz^k = c.$$

(II) λ^k converges to a dual optimal point λ^∞ .

(III) $f(x^k) + g(z^k)$ converges to the optimal value of the social optimal problem:

$$\lim_{k \rightarrow \infty} f(x^k) + g(z^k) = \sup_{x, z \text{ s.t. } Ax+Bz=c} f(x) + g(z)$$

The augmented Lagrangian $L_\rho(x, z, \lambda)$ goes to infinity as x or z goes to infinity. Hence, the iterations x^k, z^k are bounded and any subsequence x^k, z^k has converging subsequence. Let x^∞, z^∞ be its limit. By (I), (x^∞, z^∞) is a feasible allocation. Moreover, by (III), it is socially optimal allocation. This shows (i) and (iii). Since λ^∞ is a dual optimal point, $(x^\infty, z^\infty, \lambda^\infty)$ is a saddle point. In particular, (x^∞, z^∞) is a maximizer of $L_\rho(x, z, \lambda^\infty)$.

B.2 Definition of the Price as a Multiplier

In our ADMM formulation of §5.3, the augmented Lagrangian $L_\rho(x, z, \lambda)$, is equal to:

$$\begin{aligned} \sum_{t, \omega} -(v+c^{bo})E_D^\omega(t)^- - cG^\omega(t) + \text{Pos}(\underline{U}^\omega) + \text{Pos}(\overline{U}^\omega) \\ + \text{Pos}(\zeta^+ - \underline{G}^\omega(t+1) + \overline{G}^\omega(t)) + \text{Pos}(\underline{G}^\omega(t+1) - \overline{G}^\omega(t) - \zeta^-) \\ + \text{Equal0}(E_S^\omega(t) + Z^\omega(t) - E_F^\omega(t) - E_D^\omega(t)) \\ - \frac{\rho}{2} \|Ax + Bz - c + \lambda\|^2, \end{aligned}$$

where the linear equation $Ax + Bz = c$ represents all linear constraints (26, 27, 28) and (19).

Let $(\lambda_G^k)^\omega(t)$ and $(\lambda_F^k)^\omega(t)$ be the multipliers after the k th iteration of the ADMM, corresponding to the constraints $E_S^\omega(t) = G^\omega(t)$ and $E_F(t) = F(t)$ respectively. The values of E_S^{k+1} and E_F^{k+1} after the k th iteration maximize the quantity

$$\begin{aligned} -\frac{\rho}{2} (\|E_S - G^{k+1} + \lambda_G^k\|^2 + \|E_F - F^{k+1} + \lambda_F^k\|^2) \\ - (v+c^{bo})E_D(t)^-, \quad (32) \end{aligned}$$

subject to $E_S + Z - E_F - E_D = 0$.

A necessary condition for E_S^{k+1} and E_F^{k+1} to maximize this equation is that

$$E_S^{k+1} - G^{k+1} + \lambda_G^k = -(E_F^{k+1} - F^{k+1} + \lambda_F^k). \quad (33)$$

By Equation (24), the multipliers λ_G^{k+1} and λ_F^{k+1} are

$$\begin{aligned} \lambda_G^{k+1} &:= \lambda_G^k + E_S^{k+1} - G^{k+1} \\ \lambda_F^{k+1} &:= \lambda_F^k + E_F^{k+1} - F^{k+1} \end{aligned}$$

By Equation (33), these two quantities are opposite. This implies that after this iteration, the two multipliers λ_G^{k+1} and λ_F^{k+1} are opposite (for all t, ω).

Let $(\lambda_G^\infty)^\omega(t)$ be the limit of $(\lambda_G^k)^\omega(t)$ as k grows. We define the price process $P^\omega(t)$ as:

$$P^\omega(t) = \rho \lambda_G^\omega(t). \quad (34)$$

B.3 This Price Leads to an Equilibrium

We now show that the price given by Equation (34) leads to a competitive equilibrium for Scenario 2. The proof is similar for Scenario 1.

Let us denote by $x^\infty = (M^\infty, \underline{U}^\infty, \overline{U}^\infty, G^\infty)$, $z^\infty = (\overline{M}^\infty, U^\infty, \overline{G}^\infty, \underline{G}^\infty, E_D^\infty, E_F^\infty, E_S^\infty)$ and λ^∞ , three limiting values of x^k , z^k , and λ^k of our ADMM algorithm, as k goes to infinity.

B.3.1 G^∞ Maximizes the Supplier's Problem

We show that G^∞ is an optimal schedule from the supplier point of view. The vector (x^∞, z^∞) maximizes $L_\rho(x, z, \lambda^\infty)$. Hence, $(G^\infty, \overline{G}^\infty, \underline{G}^\infty)$ maximizes

$$\sum_{t,\omega} -cG^\omega(t) - \frac{\rho}{2}[E_S^\omega(t) - G^\omega(t) + \lambda_G^\omega(t)]^2 \quad (35)$$

$$+ \text{Pos}(\zeta^+ - \underline{G}^\omega(t+1) + \overline{G}^\omega(t)) + \text{Pos}(\underline{G}^\omega(t+1) - \overline{G}^\omega(t) - \zeta^-) - \frac{\rho}{2}[G^\omega(t) - \overline{G}^\omega(t) + \lambda_{\overline{G}}^\omega(t)]^2 - \frac{\rho}{2}[G^\omega(t) - \underline{G}^\omega(t) + \lambda_{\underline{G}}^\omega(t)]^2$$

The quantity $\frac{\rho}{2}\|E_S^\omega - G^\omega + \lambda_G^\omega\|^2$ rewrites as

$$\frac{\rho}{2}(\|\lambda_G^\omega\|^2 + \|E_S^\omega - G^\omega\|^2) + P^\omega \cdot (E_S^\omega - G^\omega)^\top.$$

Hence, since G^∞ maximizes (35) and satisfies $G^\infty = \overline{G}^\infty = \underline{G}^\infty$, it also maximizes $f_\rho(G)$ defined for $\rho > 0$ as

$$f_\rho(G) = \sum_{t,\omega} (P^\omega(t) - c)G^\omega(t) - \frac{\rho}{2}\|E_S^\omega(t) - G^\omega(t)\|^2.$$

An optimal schedule G for the supplier maximizes $f_0(G)$ subject to the generation constraints. Let G_S be an optimal schedule for the supplier and let us show that G^∞ also optimal for the supplier.

For $\theta \in [0; 1]$, we define an allocation $G_\theta := \theta G_S + (1 - \theta)G^\infty$ and $g_0(\theta) = f_0(G_\theta)$ and $g_\rho(\theta) = f_\rho(G_\theta)$. As $G^\infty = E_S$ we have

$$g_0(\theta) = \sum_{t,\omega} (P^\omega(t) - c)G_\theta^\omega(t); \quad (36)$$

$$g_\rho(\theta) = \sum_{t,\omega} (P^\omega(t) - c)G_\theta^\omega(t) - \theta^2 \frac{\rho}{2}\|G_S - G^\infty\|^2. \quad (37)$$

As G_θ is a convex combination of two feasible allocation. Hence, it also satisfies the generation constraints. As G_S is optimal for the supplier, the function $g_0(\theta)$, attains its maximum in $\theta = 1$. Moreover, this function is concave in θ and therefore, it has a right derivative in $\theta = 0$ that satisfies $g_0'(0) \geq 0$. Similarly, the function g_ρ is concave and attains its maximum in $\theta = 0$. Hence, it has a right derivative in $\theta = 0$ that satisfies $g_\rho'(0) \leq 0$.

It should be clear from the expressions of g_0 and g_ρ , given by Equations (36) and (37), that $g_0'(0) = g_\rho'(0)$. This implies that $g_0'(0) = 0$. Hence, $\theta = 0$ is also an optimum for g_0 , which implies that G^∞ is optimal for the supplier.

B.3.2 DR Operator's Problem

We showed in §B.2 that the multipliers for $E_F = F$ and for $E_S = G$ are opposite. Hence, the optimality of E_F^∞, U^∞ for the DR operator under the price P follows from the same approach as for the supplier (by comparing the situations $\rho = 0$ and $\rho > 0$).

B.3.3 Demand's Problem

Replacing E_S by $E_D + E_F - Z$, maximizing Equation (32) on E_D subject to $E_S + Z - E_D - E_F = 0$ is equivalent to maximizing

$$-\frac{\rho}{2}(\|E_D + E_F - Z - G^{k+1} + \lambda_G^k\|^2 + \|E_F - F^{k+1} + \lambda_F^k\|^2) - (v + c^{bo})E_D(t)^-,$$

without constraint.

As the second term does not depend on E_D , this problem becomes maximizing:

$$-\frac{\rho}{2}\|E_D + E_F - Z - G\|^2 - P(t)E_D(t) - (v + c^{bo})E_D(t)^-.$$

Hence, the optimality of E_D for the Demand under the price P follows from the same approach as for the supplier (by comparing $\rho = 0$ and $\rho > 0$).

C. PROOF OF THEOREM 3

The matrix Σ is the covariance of the forecast error process ε . As such, Σ is a positive-definite symmetric matrix of size $T \times T$ and has real entries. Thus, there exists a real matrix A of size $T \times T$ such that:

- A is a lower-triangular matrix, *i.e.*, for all $t \in \{1, \dots, T\}$ and $s \in \{t + 1, \dots, T\}$, $A_{t,s} = 0$.
- $\Sigma = A \times A^\top$.

A is called the lower Cholesky decomposition of Σ . It also exists when Σ is not full rank.

Let Z^ω be a trajectory. By construction, the vector N^ω is a vector of T *i.i.d.* normally distributed random variables. Moreover, we have $Z^\omega(t) = \sum_s A_{t,s} N^\omega(s)$. Hence, the covariance $\mathbb{E}[Z^\omega(t)Z^\omega(t)^\top]$ is equal to:

$$\mathbb{E}[Z^\omega(t)Z^\omega(t)^\top] = \sum_{s,s'} \mathbb{E}[A_{t,s} N_s A_{t',s'} N_{s'}] = (AA^\top)_{t,t'},$$

which implies that property (i) holds.

Let us fix a sub-set of trajectories $Z^{\omega_1} \dots Z^{\omega_k}$ and let Z^ω be another trajectory. Let τ be the largest branching time of Z^ω shared with one of the trajectories $Z^{\omega_1} \dots Z^{\omega_k}$. By construction, $Z^\omega = AN^\omega$. The matrix A and the vector N^ω can be decomposed as:

$$A = \begin{bmatrix} B & 0 \\ C & D \end{bmatrix} \text{ and } N^\omega = \begin{bmatrix} N_1 \\ N_2 \end{bmatrix},$$

where the matrices B, C and D have size $\tau \times \tau$, $(T - \tau) \times \tau$ and $(T - \tau) \times (T - \tau)$ and N_1 has size τ .

By construction, the distribution of Z^ω given trajectories $Z^{\omega_1}, \dots, Z^{\omega_k}$ is the distribution of Z^ω given $N^\omega(1), \dots, N^\omega(\tau)$. Let $z = B \cdot [N^\omega(1), \dots, N^\omega(\tau)]$. Because A has full rank, B has also full rank and is invertible. Hence, the distribution of Z^ω conditioned on $Z^\omega(t) = z_t$ for $t \leq \tau$ is the same as the distribution of Z^ω given $N^\omega(1) \dots N^\omega(\tau)$. Moreover, the former is the conditional distribution of a multivariate Gaussian random variable of covariance Σ subject to Equation (18). \square