

# Reducing electricity consumption peaks with parametrised dynamic pricing strategies given maximal unit prices

Nicolas Höning<sup>1</sup> and Han La Poutré<sup>1,2</sup>

<sup>1</sup>CWI, Science Park 123, Amsterdam, The Netherlands

<sup>2</sup>Utrecht University, Princetonplein 5, Utrecht, The Netherlands

**Abstract**—Demand response is a crucial mechanism for flattening of peak loads. For its implementation, we not only require consumers who react to price changes, but also intelligent strategies to select prices. We propose a parametrised meta-strategy for dynamic pricing and identify suitable strategies for given scenarios through offline optimisation using a population model. We also model an important and novel constraint: a price cap (a maximal unit price) for consumer protection. We show in computational simulations that the maximal unit price influences the peak reduction potential of dynamic pricing. We compare our dynamic pricing approach with a constant pricing approach and show that our approach, used by a profit-optimising seller, is both peak-reducing and equally profitable.

## I. INTRODUCTION

The future technology mix in energy systems will contain more powerful electricity-consuming devices, e.g. electric vehicles and heat pumps. However, the existing infrastructure (e.g. distribution cables) has not been designed for carrying such large loads. It is crucial that our energy infrastructure can support the transition to a modern energy future, but major updates are very expensive. Dynamic pricing (DP) is a technique that can flatten demand over time and thereby enable peak shaving. Thus, otherwise necessary upgrades to network components like cables can be postponed, which can save millions in societal investments. In the so-called smart grids vision, where many more decision makers can react to prices than today, e.g. domestic consumers and small businesses, dynamic pricing is a promising route.

However, consumers are used to constant prices and are in general not welcoming prices that can vary freely and unpredictably. Therefore, promising an upper bound on prices is both of interest to retailers (to attract consumers) and government policy (to protect consumers). Then again, limiting the possible range of prices reduces the effect of dynamic pricing. Thus, we are dealing with a trade-off in this market design problem. To the best of our knowledge, we are the first to introduce a maximal price into market systems for energy in smart grids.

In our model, consumer agents have stochastically arriving jobs with deadlines for which they need to buy electricity supply (e.g. to charge a car or to run some factory equipment). A price-setting agent (the *PSA*) buys on the wholesale market and selects prices for consumers before

each time step. The *PSA* has to pay fines when peaks occur. This represents a stochastic and distributed online scheduling problem under scarcity conditions (i.e. the cable capacity) and with limitations on the range of control signals (by introducing a maximal price). Finding good dynamic pricing strategies for this problem is a complex task.

This paper contributes to the state of the art in the following ways: We offer a formalised model for the distributed online scheduling problem described above, which includes the novel aspect of a maximal price. We then propose a novel, parametrisable meta strategy to set prices dynamically, based on available knowledge about currently running jobs. Instances of this meta strategy are specified by assigning values to its parameters. Preferable parametrisation can be found if a suitable model of the population and the distribution of jobs exists. In this work, we find settings through offline optimisation, for which we use an evolutionary algorithm. We simulate computational scenarios and illustrate the peak reduction potential of this dynamic pricing (DP) strategy. We show that the maximal price has an influence on the peak reduction potential of dynamic pricing. Furthermore, we show that the *PSA* does not prefer a constant price (CP) strategy over our (DP) strategy. Finally, we show that employing a DP strategy reduces peaks, even if it is optimised for profits.

## II. BACKGROUND

Demand currently represents the least elastic part of electricity markets, on the household level even more so than on the industrial level. Several reports (e.g. [1], [2]) suggest that dynamic pricing has substantial potential to balance power in future grids, especially if more devices with significant consumption demand capacity are placed in the households (e.g. electric cars and heat pumps). Usually, this function is referred to as “Demand Response”.

Dynamic Pricing (DP) is a form of optimising revenues for a seller which involves changing the price of goods or services over time. Traditionally, it involves application areas where the capacity is fixed in the short-term and perishable [3], and thus it is a good fit for electricity retail. A dynamic pricing strategy aims at selecting the best price for a good, given a specific setting. Prices can be selected on short notice, based on conditions found in the current situation, or

a constant price is selected for a certain period. Especially when settings are complex to predict or analyse, the latter is sometimes the profit-maximising strategy [4].

In our complex setting, we need to identify suitable solutions (that can be expected to perform well) from a parameter space with several dimensions. To this end, we have chosen to use evolutionary algorithms (EAs). We use a particular type of EA, namely Estimation-of-Distribution Algorithms (EDAs). EDAs differ from EAs in that to generate new solutions, a probability distribution in the space of solutions is first estimated from the selected solutions, from which then new solutions are drawn via random sampling. In particular, we use the iMAMaLGaM-X+, a version of EDAs for multi-objective problems [5], for which EAs have shown to be highly effective [6]. It uses mixture distributions by clustering the selected solutions, spreading the search intensity along the Pareto front in an effective manner.

### III. MODEL

Several agents  $a \in A$  are consumers of electricity who have jobs to run. Each job has a distinct arrival time and a deadline. We assume that all consumers are connected to the electricity grid via one shared cable. We measure performance over one day at a time, and thus consider a set of days  $D \in T$ . Each day is partitioned in time steps  $t \in \mathbb{N}$  of 15 minutes, so one day has  $24 * 4 = 96$  time steps.

The price-setting agent  $PSA$  buys electricity at current wholesale market prices  $\rho_w^t$  and has to pay penalties if the cable is overloaded. He sets local unit prices  $\rho^t$  for each time step, and the consumer agents decide before each time step whether to purchase electricity for their job. If  $a$  contracted power for a time step,  $a$  consumes 1 kW (0.25 kWh) and pays  $\rho^t$  to the  $PSA$ . Otherwise, the job is delayed, which is unwanted, as consumers would prefer to finish their jobs as fast as possible. We assume that all jobs have to be supplied until they are done and that supply is limitless.

#### A. Price-setting dispatch agent ( $PSA$ )

The  $PSA$  sells electricity to consumers, which he buys on the wholesale market. For each time step  $t$ , the  $PSA$  announces a unit price  $\rho^t$  for the supply of electricity. If a consumer  $a$  buys in time step  $t$ , we denote this by  $Q_a^t = 1$ ; if he does not buy, we denote this by  $Q_a^t = 0$ . We assume that unit prices are bound by the maximal unit price  $\rho_{max}$ .

We assume for simplicity of our mechanism, that  $\rho_w^t$ , the unit price on the wholesale market for time step  $t$ , is constant, so we can from now on refer to  $\rho_w$ . Furthermore, let  $Q^t$  denote the supply that the  $PSA$  has to buy on the wholesale market for time step  $t$  and let  $Q_{max}$  be the

maximal capacity, above which the cable is overheated<sup>1</sup>. The  $PSA$  is aware of a cost function  $heat(t) \rightarrow \mathbb{R}$ , which calculates the costs for overheating the cable. This cost has to be paid by the  $PSA$ . In particular, the damage to a cable through overheating at time  $t$  is simply  $heat(t) = max(0, Q^t - Q_{max})$ .

#### B. Consumers

We model a population of consumer agents  $a \in A$  (we denote the size of the population with  $|A|$ ), who all aim to minimise the costs of running their job. They differ in two aspects: the arrival times of their jobs and their reluctance to react to a price difference, i.e. to shift job execution away from their preferred slots. This reluctance is modelled as a delay cost  $\delta_a \in [0, \delta_{max}]$  per consumer. The delay costs directly influence the magnitude of price differences that a successful strategy needs to employ to be effective.

1) *Jobs*: Because each consumer agent  $a \in A$  has exactly one job, we will use the subscript  $a$  when we denote properties of the consumer's job. Each job can start after a unique arrival time  $t_a^s \in [0, 96 - L]$ , where  $L$  is the maximal length of time (measured in time steps) for which a job can remain unfinished after it started (thus, each job has a mandatory deadline  $t_a^e = t_a^s + L$ ). Start times  $t_a^s$  are drawn anew every day from a Gaussian probability function. This probability function is used for all jobs and remains the same over all days  $D \in T$ . In contrast to the  $PSA$ , consumers know in the beginning of the day when their job will start. The number of time steps for which a job needs to be supplied in order to be finished is denoted as  $W \leq L$ .

For simplicity of our mechanism, we make the assumption that if a job is supplied during time step  $t$ , it requires and receives the exact, constant input of 1 kW for the duration of  $t$ . We thus assume that physically, both consumption and delivery can, within one time step, be constant over time.

To describe the state of  $a$ 's job in a time step  $t$ , we denote with  $rmng : a, t \rightarrow \mathbb{N}$  the amount of electricity which  $a$  still needs to buy (see also the previous section) *before*  $t$ . Initially,  $rmng(a, t_a^s) = W$ . If  $Q_a^t = 1$ , then  $rmng(a, t+1)$  is reduced by one. If  $t_a^e > t + rmng(a, t)$ , then  $a$ 's job has flexibility of when to buy electricity:  $a$  can choose not to supply the job for up to  $t_a^e - t - rmng(a, t)$  time steps by delaying it. Note that all jobs have (the same) flexibility on arrival, as for all  $a$ , it holds that  $t_a^e - t_a^s = L$  and  $rmng(a, t_a^s) = W$ . So,

$$rmng(a, t) = 4 - \sum_{i=t_a^s}^{t-1} Q_a^i \quad (1)$$

<sup>1</sup>In today's reality, there exists a second threshold  $Q_{cut} > Q_{max}$ , above which supply has to be cut off. Penalties have to be paid only for cutoffs. We assume that  $Q_{cut}$  is not breached in our scenarios and that there will be some negative response for the  $PSA$  to not preventing a peak (i.e. for overheating and thus limiting the lifetime of the cable or because spot market prices increase during a peak).

Furthermore,  $dlyd : a, t \rightarrow \{0, 1\}$  (short for *delayed*) is 1 if the job is not finished but also does not run *during*  $t$ . So,

$$dlyd(a, t) = \begin{cases} 1 & \text{if } rmng(a, t) > 0 \wedge Q_a^t = 0 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

2) *The consumer's cost minimisation problem:* Each consumer  $a$  needs to make an informed decision whether to buy or not in each time step  $t \in [t_a^s, t_a^e]$  if  $rmng(a, t) > 0$ . For this, he needs some kind of expectation over future costs. Consumers learn which prices are to be expected in which time step based on experience in previous days.

In this section, we provide a simple algorithm for  $a$  to make this decision in a rational way. This algorithm should be performed in polynomial time, to keep computation time within acceptable bounds. We therefore provide a greedy algorithm to solve the minimisation problem given in 3.

$$\arg \min_{Q_a^t} \left[ Q_a^t \rho^t + dlyd(a, t) \delta_a + EC(a, t, Q_a^t) \right] \quad (3)$$

s.t.  $Q_a^t \in \{0, 1\}; \quad Q_a^t = 1$  if  $t + rmng(t, a) \geq t_a^e$

where  $\rho^t \in [0, \rho_{max}]$  is the announced unit price and  $EC(a, t, Q_a^t)$  is a function to estimate costs of future time steps  $i \in [t + 1, t_a^e]$ , given  $Q_a^t$ .  $EC$  is given by Algorithm 1. The algorithm estimates purchasing costs based on  $\bar{\rho}^E$ , a vector of expected prices.  $\bar{\rho}^E$  contains, for each time slot of the current day, a weighted running average of the prices that were announced in previous days for this time slot. We denote with  $E[\rho^i] \in \bar{\rho}^E$  the local price that  $a$  expects in time step  $i$ . The outcomes of the state functions  $rmng$  and  $dlyd$  for any future time steps  $i > t$  are based on purchasing decisions  $Q_a^j$  which the algorithm already assumed for time steps  $j \in [t + 1, i - 1]$ .

---

**Algorithm 1** For a given  $t$ ,  $a$  estimates  $EC_a^t$ , the costs of completing his job in future time steps.

---

```

 $EC_a^t = 0$ 
for all  $i \in [t + 1, t_a^e]$  do
  // If job needs supply:
  if  $rmng(a, i) > 0$  and  $t_a^s \leq i \leq t_a^e$  then
    // If necessary or cheaper than waiting, buy now.
    if  $i + rmng(a, i) \geq t_a^e$ 
      or  $E[\rho^i] < dlyd(a, i) \delta_a + E[\rho^{i+1}]$  then
         $EC_a^t += E[\rho^i]$ 
    // Otherwise do not buy now.
  else
     $EC_a^t += dlyd(a, i) \delta_a$ 
  end if
end if
end for

```

---

## IV. PRICING STRATEGIES

In this section, we propose a meta strategy for dynamic pricing (DP). This work deals with uniform prices, i.e. a unit price is set for the following time step which is valid for all consumers. Furthermore, we make two design choices:

- 1) *Dynamic pricing:* The price in each time step is a function of the state of active jobs (only jobs with  $t_a^s \leq t \leq t_a^e$  are known to the *PSA*). We suppose the *PSA* can assess how much work each active jobs still has to do and when its deadline is.
- 2) *Offline optimisation:* Given that designing successful strategies is hard in a complex setting like this, we optimise their parametrisation offline, given some knowledge about distribution of job arrivals. A pricing strategy, once chosen, is employed by the *PSA* unchanged for a number of days.

We are interested in formulating pricing strategies that compute the price  $\rho^t$  for the current time step  $t$  and achieve good values (over a range of days  $D \in T$ ) on an objective function  $O$  given by

$$\arg \min_{\vec{x}} \sum_{D \in T} O(D) \quad (4)$$

where  $\vec{x}$  is a vector of parameters which is optimised offline. We discuss possible implementations of  $O$  below.

Prices are computed online with a stylised mathematical function. Its general form allows the offline optimisation algorithm to explore a large variety of possible functions through parametrisation. However, its online application requires no extensive computation and our strategy is thus able to be executed fast. We try out two meta strategies, the first being quadratic functions (denoted by  $\rho_{quad}^t$ ) and the second being exponential functions (denoted by  $\rho_{exp}^t$ ):

$$\begin{aligned} \rho_{quad}^t(\omega^t) &= x_0 + x_1 \omega^t + (x_2 \omega^t)^2 \\ \rho_{exp}^t(\omega^t) &= x_0 + (x_1 \omega^t)^{x_2} \end{aligned} \quad (5)$$

where the parametrisation vector is  $\vec{x} = \{x_0, x_1, x_2\}$  and  $\omega^t$  describes the likelihood of consumers with active jobs to buy in  $t$ , given the state of their job:

$$\omega^t = \sum_{a \in A} \frac{rmng(t, a)}{t_a^e - t} \quad (6)$$

The term  $\frac{rmng(t, a)}{t_a^e - t}$  is initially  $\frac{W}{L}$  for all jobs and can increase up to 1 if  $a$  has no more time to delay his job. Thus,  $\omega^t \in [0, |A|]$ . Note that  $\omega^t$  does not include values of  $\delta_a$ , which is private information about delay costs of each consumer. The parametrisation vector  $\vec{x}$  is optimised offline, using an accurate model of the population (including the local scheduling procedure  $EC$  from Algorithm 1 and the statistical distribution of  $t_a^s$  and  $\delta_a$  values). In the online setting, the *PSA* uses only the optimised pricing functions

Table I: Experimental settings

Parameter	Description	Setting
$ A $	number of agents	20
$T$	number of days	20
$L$	length of time window for a job	8 (2 hours)
$W$	workload of a job	4 (1 hour)
$\rho_w$	wholesale market price	€ .1 / kWh
$Q_{max}$	cable threshold	{8, 4, 2}
$\delta_{max}$	max. delay cost factor	€ 0.05/15 min
$J_{mean}, J_{std}$	mean and st. dev. of $t_a^s$ (job arrival)	20, 2

(which require only the current state  $\omega^t$ ) and does not need to know  $\delta_a$ . With  $\rho_{max}$  given,  $\rho^t$  is given by

$$\rho^t = \min(\rho_{max}, \rho_{quad}^t(\omega^t)) \text{ or } \rho^t = \min(\rho_{max}, \rho_{exp}^t(\omega^t)) \quad (7)$$

We are interested in **lowering the maximal price**  $\rho_{max}$ , in order to offer a consumer-friendly operation. However,  $\rho_{max}$  should remain higher than  $\rho_w$ , in order to give room for price differences such that the pricing can lower peaks and at the same time avoid losses for the PSA.

There are two possible implementations for the objective function  $O$ . Both form a trade-off with  $\rho_{max}$ , as the DP strategy depends on a range in which to vary prices. First, we are interested in **reducing peaks** by dynamic pricing:

$$O(D) = \sum_{t \in D} heat(t) \quad (8)$$

Alternatively, we can **maximise the profit of the PSA**:

$$O(D) = \sum_{t \in D} \sum_{a \in A} Q_a^t(\rho^t - \rho_w) - heat(t) \quad (9)$$

## V. SIMULATIONS

This section describes computational experiments. We model three scenarios which differ in the scarcity of cable capacity. In each scenario, we test two strategies - a constant pricing (CP) strategy and our dynamic pricing (DP) strategy described in Section IV. We report on differences between quadratic and exponential DP meta-strategies in the end. The detailed parameter settings can be read from Table I.

### A. Scenarios

We now propose a simple indicator  $\beta$  for the scenario-specific scarcity of the cable capacity. We describe  $\beta$  as the ratio between needed time steps and available time steps. The overall number of time steps needed to run all jobs is  $|A| * W$ . Furthermore, we estimate the number of available time steps by computing the maximal number of time steps in which consumers could be supplied without overheating the cable, which is  $(6J_{std} + L)Q_{max}$ . Here,  $6J_{std}$  covers 99.7% of the distribution of job starting times and we also add one job length  $L$  to cover jobs that possibly start on the

last time step. Let us now insert values to formulate  $\beta$  in dependence of  $Q_{max}$ :

$$\beta = \frac{\text{needed}}{\text{available}} = \frac{|A| * W}{(6J_{std} + L)Q_{max}} = \frac{4}{Q_{max}} \quad (10)$$

If  $\beta = 1$  (and thus  $Q_{max} = 4$ ), the available time steps could theoretically suffice to solve the allocation problem without causing overheating, given that the needed redistribution of consumption were achievable. We note that the normal distribution of job starting times and delay costs limit the solvability of this problem. However, the given formulation of  $\beta$  allows us to anchor a scenario at the theoretically solvable setting. If  $\beta$  is increased (and thus  $Q_{max}$  is lowered), the scenario becomes more scarce. If  $\beta$  is decreased (and thus  $Q_{max}$  is increased), the scenario becomes less scarce. We run three scenarios, with  $\beta \in \{.5, 1, 1.5\}$ , thus  $Q_{max} \in \{8, 4, 2\}$ .

### B. Evaluation

In each scenario, we optimise a set of strategy instances of the DP meta strategy using the evolutionary algorithm (EA), with the minimisation of  $\rho_{max}$  as first objective. We optimise each scenario twice - once with the second objective being peak reduction (see Equation 8), and once with the second objective being profit maximisation for the PSA (see Equation 9).

The EA is parametrised according to guidelines from literature, among others populations size and maximum number of subsequent generations without an improvement. We run each market in each scenario for 1000 generations, using 5 clusters of solution populations (which the algorithm uses to model the pareto front). We evaluated that the performance of the multi-objective EA as a means to find good solutions is suitable by comparing it to a brute-force approach which evaluated 300,000 data points for one scenario.

Strategy instances are judged on outcomes averaged over ten runs per instance. For every instance, these ten runs are performed with the same set of ten seeds, used for the randomisation of  $\delta_a$  and  $t_a^s$  values. An outcome is computed as the average over the last  $\frac{T}{2}$  days of a simulation.

As a benchmark for the DP meta strategy, we implement a simple constant price strategy (CP). If the PSA charges a constant price, the amount of peaks will be the same, no matter what that price is. This is because there are no price differences, and consumers will all supply their jobs directly upon arrival, in order to avoid delay costs. If the CP strategy is optimised for profit of the PSA, the PSA always chooses  $\rho_{max}$ . Note that the DP strategy can also model the CP strategy, as the mathematical formalisation allows to charge any constant price, as well.

As benchmark for the peak reduction potential, an optimal lower bound of peaks would be useful, but the computational effort to compute it is too high. For every time step, each

consumer has  $\binom{L}{W}$  possible schedules, thus the solutions space in our setting has the size  $96 \cdot T \binom{L}{W}^{|A|} = 96 \cdot 20 \binom{8}{4}^{20}$ .

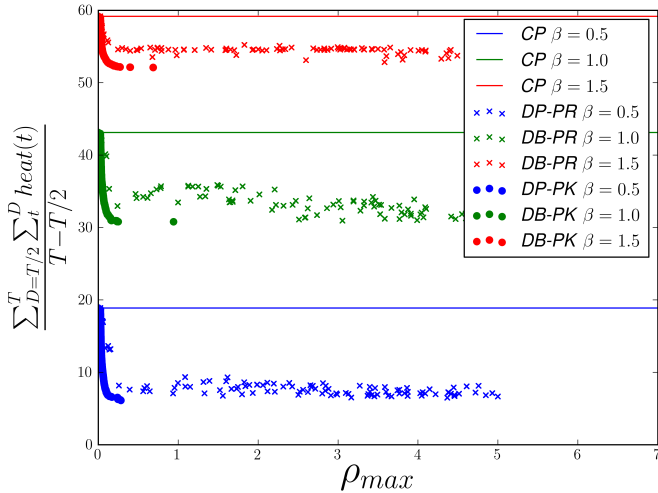


Figure 1: Average daily overheating costs against maximal prices  $\rho_{max}$  when employing the CP strategy or the DP strategy with either profit-maximising (DP-PR) or peak-reduction (DP-PK) objectives.

Implementing the first optimisation goal (peak reduction) in the DP strategy, **we can show the peak reduction potential of the DP strategy, given  $\rho_{max}$** . There is clearly a trade-off between the price range in which DP can operate (limited by  $\rho_{max}$ ) and the peak reduction it can achieve. The plots labelled DP-PK in Figure 1 show that the outcomes form a Pareto front. It also becomes clear that if  $\rho_{max} \leq \in 0.5$ , the possibilities to reduce peaks decrease. If  $\rho_{max} \geq \in 0.5$ , the peak reduction potential does not increase further.

Implementing the second optimisation goal (profit maximisation) in the DP strategy, we can compare the revenue optimisation potential for the PSA of CP and DP. It turns out that **with the DP strategy, profits are not significantly different as with the CP strategy**. In both cases, profits increase linearly with  $\rho_{max}$ . From the results, we can also estimate which is the minimal  $\rho_{max}$ , below which the PSA should expect to be making losses. In addition, we can show how much peaks occur if the PSA is optimising profits (as he realistically would do). The plots labelled DP-PR in Figure 1 make clear that **if the PSA optimises the DP strategy to maximise revenues, peaks are still considerably lower than with the CP strategy** (but of course higher than when the DP strategy is optimised for peak reduction). The PSA is not charging  $\rho_{max}$  constantly, but prefers to actually charge dynamic prices in order to avoid some peak costs.

An important consideration is how the ratio between peak costs and revenues from selling electricity influences the performance of our DP strategy. A straightforward analysis

shows that this ratio affects how much the PSA considers peak reduction as a priority when choosing a profit-maximising strategy. We have so far simply accumulated the magnitude of the total peaks (see Section III-A) and taken these sums as pure peak costs. Assume there exists a factor  $\phi$ , with which peaks can be multiplied, in order to get their true costs. So far, we have used  $\phi = 1$ . It is straightforward to assume that, if  $\phi = 0$ , the PSA would assume the CP strategy and always charge  $\rho_{max}$ . If  $\phi = \infty$ , the PSA would resort to make peak reduction its sole objective, which would result in the outcomes of our peak-optimised optimisations.

Finally, we compared two mathematical classes of DP meta strategies that the EA could explore (see Equation 5). Both the quadratic and exponential classes result in outcomes that are not significantly different.

## VI. CONCLUSIONS

Dynamic pricing is an important tool for reducing peaks in future electricity grids, where scheduling decisions are made online by independent actors. Mandating an upper constraint on unit prices is an important ingredient for consumer protection. We offer a formalised model for such a distributed online scheduling problem and propose a meta strategy for dynamic pricing. We show how to find suitable strategies through offline optimisation. We show in computational simulations that the peak reduction potential of our strategy depends on the maximal price. Furthermore, we show that the PSA does not prefer a constant price (CP) strategy over our dynamic price (DP) strategy. Finally, we show that employing a DP strategy reduces peaks, even if it is optimised for profits.

## REFERENCES

- [1] A. Faruqi and S. Sregici, "Household response to dynamic pricing of electricity - a survey of the experimental evidence," The Brattle Group, San Francisco, Tech. Rep., 2009.
- [2] D. Kathan, "Demand Response And Advanced Metering," Federal Energy Regulatory Commission, Tech. Rep., 2009.
- [3] W. Elmaghraby and P. Keskinocak, "Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions," *Management Science*, vol. 49, no. 10, pp. 1287–1309, 2003.
- [4] G. Gallego and G. van Ryzin, "Optimal dynamic pricing of inventories with stochastic demand over finite horizons," *Management Science*, vol. 40, no. 8, pp. 999–1020, 1994.
- [5] P. Bosman and T. Alderliesten, "Incremental Gaussian Model-Building in Multi-Objective EDAs with an Application to Deformable Image Registration," in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO)*, 2012.
- [6] M. Farina, K. Deb, and P. Amato, "Dynamic Multiobjective Optimization Problems: Test Cases, Approximations, and Applications," *IEEE Transactions on Evolutionary Computation*, vol. 8, no. 5, pp. 425–442, Oct. 2004.