
Unterschrift des Betreuers



DIPLOMARBEIT

**Valuation of fossil fueled power plants
with respect to operational constraints**

Ausgeführt am Institut für

Wirtschaftsmathematik

der Technischen Universität Wien

unter der Anleitung von

Ao.Univ.Prof. Dipl.-Ing. Mag.rer.nat. Dr.techn.
Wolfgang Scherrer

durch

Gross Peter

Vivenotgasse 37/6
1120 Wien

4. Mai 2009

ABSTRACT

This thesis consists of the implementation of an algorithm (proposed by Tseng and Barz in [1]) for determining the expected payoff of a fossil fueled power plant in MATLAB. Physical properties such as decision lead time, unit commitment constraints, output depending fuel to electricity conversion and additional cost for applying controls are included into the plant model. Hourly control of the plant with the goal of maximizing its expected payoff is assumed.

The model is solved by a combination of Dynamic Programming and Monte Carlo simulations.

In this work, the model and techniques necessary for its solution are introduced.

The structure and performance of the generated MATLAB program are presented along the results from some numerical simulations.

Contents

1	Introduction	1
1.1	Generating power for electricity markets	2
1.2	Basic models for the value of a plant	2
1.2.1	Interpreting power generation as financial options	3
1.3	The Tseng and Barz model	4
1.3.1	Constraints on control and state variable	4
1.3.2	Cost function and optimal production	6
1.3.3	Model formulation	7
2	Deterministic and Stochastic Solution	9
2.1	The deterministic case	9
2.1.1	Dynamic Programming	9
2.1.2	Non standard representation of the deterministic model	12
2.2	The stochastic model	13
2.2.1	Deriving the stochastic model	13
2.2.2	A detailed view of the stochastic model	15
2.3	Integration using Monte Carlo simulation	17
2.3.1	Monte Carlo integration	17
2.3.2	Variance reduction: antithetic variates	18
2.4	The iterative solution for the stochastic model	19
2.4.1	Describing optimal control strategies	19
2.4.2	Describing $d(\cdot)$ with indifference loci (<i>IL</i>)	20
2.4.3	Solution of the model of Tseng and Barz	20
3	Implementation in MATLAB	24
3.1	The big picture	24
3.1.1	Generating price scenarios for Monte Carlo simulation	25
3.1.2	Approximating the IL	28
3.2	The program	28
3.2.1	Data storage: The variable <i>Experiment</i>	29

3.2.2	workbench.m	30
3.2.3	J_difference.m	31
3.2.4	prices.m and pricesVred.m	33
3.2.5	Other functions	34
4	Deterministic Simulation	35
4.1	Deterministic simulation: Testing the Program	35
4.2	The standard experiment	36
4.3	Simulation of the standard experiment with deterministic prices	36
4.3.1	IL of the standard experiment with deterministic prices	37
4.3.2	Strange IL: Discontinuities and the problem of describing them as function of p_t^E	38
4.3.3	A showcase explanation for the reaction of the optimal control on a change in fuel prices	39
4.3.4	Conclusion	42
5	Empirical Results	44
5.1	The expected payoff over 168 hours (standard experiment)	45
5.1.1	Plant value in dependence of observed prices and state	45
5.1.2	Unit commitment based on different degrees of information	49
5.2	Frequency counts of scenario payoffs (standard experiment)	51
5.2.1	The distribution of $J_1(\cdot)$	52
5.2.2	The distribution of $d_1(\cdot)$	54
5.3	IL and computation time (standard experiment)	55
5.3.1	Indifference Loci	55
5.3.2	Computation time	58
5.4	Changes in the system dynamics	60
5.4.1	Short decision lead time, short unit commitment time .	61
5.4.2	The comparison of <i>gas engine</i> and <i>standard experiment</i>	63
6	Conclusion	66
6.1	Improvements of the present algorithm	66
6.1.1	Indifference Loci	66
6.1.2	Monte Carlo estimation	67
6.2	The potential of the current program	67

Chapter 1

Introduction

The main task of this thesis is the implementation of an algorithm for determining the value of a thermal power plant in MATLAB.

The thermal generation of electric power is the conversion of fossil fuels into electricity. Therefore the profit of this process is depending on the prices of two volatile goods. In addition, physical restrictions on the conversion process (arising from the properties of the plant) play a significant role as well, since they influence the possibility of an operator to react on price peaks.

In 2000, Chung-Li Tseng and Graydon Barz proposed an algorithm for determining the value of a power generating unit as described above. They employ forward moving Monte Carlo simulations to estimate the expected payoff of a plant, conditional on observed prices and for a given control strategy. Using these expected payoffs in backward moving dynamic programming yields the optimal control strategy of the plant. This approach has the benefit of allowing the implementation of physical constraints to a certain extent. On the other hand it is computationally very expensive.

The Tseng and Barz algorithm has been used in this work for plant evaluation. Note that also assumptions like the structure of the energy market¹ or the class of the price processes have been made similar as in their article [1]. One prominent reason for the similar assumptions is the possibility of comparing results. Another reason is the main focus of the work, being modeling and analyzing the effects of physical properties of a plant and not simulating energy markets.

Naturally, the article of *Operations Research* where Tseng and Barz propose this algorithm has been the main source for this text. As far as possible, their notation was used in the model formulations and in the program code.

¹hourly spot markets for electricity and natural gas

1.1 Generating power for electricity markets

The last 20 years have seen the rise of markets for electrical energy in many countries. For the producers of electricity, this has enormous consequences, since their objective switched from satisfying the demand to competing on a marketplace. One particularly interesting difference of electricity to other trading goods is its lack of storage ability, which generally results in high price volatility. However, hydro power plants appear to provide an exception from the non - storage property of electricity, since they can 'pile up' kinetic energy of water in times of low energy demand and transform it to electricity on very short notice. Therefore, the highly hydro-power supplied Nordic market for example, appears to be much less subject to short term fluctuations than markets for mainly fossil generated energy like in Germany².

This example illustrates the fundamental importance of operational properties of an electric power plant. The amount of time necessary to start power generation may be vital to sell electricity at price peaks on short notice; the effects of unit commitment constraints such as a minimum running time, lead to additional cost of production and can be difficult to quantify.

1.2 Basic models for the value of a plant

A thermal power plant can basically be characterized as a device capable of converting fuel (natural gas, coal or oil) into electricity. The pivotal element of this conversion process is given by the *heat rate* H . H measures the amount of fuel needed to generate one unit of electricity. Following the notation of Tseng and Barz, the units chosen for fuel and electricity are measured in MMBtu³ and MWh respectively. Naturally the heat rate achieved by a power plant can be dependent on the level of the output.

Taking the most simple approach, the conversion of fuel into electricity yields a payoff of $p^E - H \cdot p^F$ (p^E denoting the price of electricity in MWh, p^F the price of fuel in MMBtu) for one produced unit of power in one time period. Extending the evaluation problem to a longer time period, where the plant is assumed to be continuously producing 1 MWh electricity, leads to

$$\text{Power plant value} = \sum_{t=1}^T p_t^E - H \cdot p_t^F \quad (1.1)$$

²[2], p 15

³million British thermal units

1.2.1 Interpreting power generation as financial options

Evaluating a plant by model (1.1) is a quite strong oversimplification since it does not incorporate the stochastic nature of fuel and electricity prices. Besides, any operator would be keen on avoiding losses in case of $p^E < H \cdot p^F$ by simply turning off the plant.

In financial terms, this situation resembles the holding of European call options of spark spreads, defined as the difference of the electricity price and the generation cost⁴. This option (or swap) enables to exchange one commodity (fuel) for another (electricity) at expiry date. As a consequence, it provides the holder with the difference of the current price of electricity and the strike price of the option⁵. If the electricity price falls below the strike price, the plant is to be shut down (i.e. the option is not exercised).

The value of the plant over a whole time period resembles a sequence of spark spreads of this kind, each of them exercisable at different maturity dates t . The expected value of these options is formulated by extending 1.1 to

$$\text{Power plant value} = \sum_{t=1}^T \mathbb{E}_1 (\max(p_t^E - H \cdot p_t^F, 0)) \quad (1.2)$$

where \mathbb{E}_1 denotes the expected value conditional to the information available at time 1. As in equation (1.1), this evaluation uses a production capacity of 1 MWh.

The financial options model 1.2 has the benefits of being well researched and capturing the stochastic nature of prices as well as the decision problem. For some price processes (like geometric Brownian motion) there even exists an analytical solution.

Nevertheless, this model lacks to incorporate some fundamental issues mentioned above:

- The plant operator has the possibility to apply optimal power generation to observed prices. The plant can be started / shut down immediately if favorable / unfavorable prices appear. This neglects the possible time lag in controlling the plant (decision lead time).
- Turning the plant on or off does not imply additional expenses.

⁴[1], section 1

⁵[2] p114

- There are no constraints on unit commitment, which means that the plant can be turned on / off at every time.
- The heat rate H is independent of the level of the output.

Omitting those aspects results basically in exaggerating the operator's capabilities to control the plant. For example, the plant could be forced to produce electricity in spite of electricity prices lower than fuel cost due to previous unit commitment. This scenario is completely missing in the financial options model above.

1.3 The Tseng and Barz model

Contrary to the financial options approach above, Tseng and Barz introduce a model based on an optimal plant control problem. The operator will influence the state of the plant in order to maximize its (expected) profit over a given horizon. The feasible controls as well as the state dynamics are subject to physical restrictions.

The plant model introduced in this section is largely taken from [1].

1.3.1 Constraints on control and state variable

The plant type considered here consists of a steam turbine fed by a boiler heated by fossil fuel. Thus the main source of delay in controlling this unit can be found in the dynamics of the heat level in the boiler. On a slightly higher level of abstraction, the heat level corresponds to how long the plant has been producing electricity or turned off in one row. This amount of time is described by the state variable x_t . A running plant will be represented by a positive integer x_t , denoting the amount of time⁶ since the plant has been turned on. If the plant has been turned off, x_t will be a negative integer and its absolute value represents the time since production stopped.

The state of the plant influences its future value in two ways: First there may be restrictions in turning off (on) the plant after a too short period of production. For this purpose, t^{on} and t^{off} are chosen to represent the minimum time (in hours) the plant has to stay on or off respectively. t^{on} and t^{off} will be called minimum up- and minimum down time in the rest of the text. The control variable will be described by $u_t = 1$ (0) for turning on (off) or

⁶subsequently measured in hours

continuing running (not running) the plant at time t .

This implies the following *minimum up- and downtime constraints* on the control:

$$u_t = \begin{cases} 1, & \text{if } 1 \leq x_t < t^{\text{on}} \\ 0, & \text{if } -t^{\text{off}} < x_t \leq -1 \\ 0 \text{ or } 1 & \text{otherwise} \end{cases} \quad (1.3)$$

Also the cost of starting a cold plant is assumed to depend on the time the plant has already been shut down. It seems reasonable to assume a minimum temperature the boiler will cool towards when the plant is offline. As a result, a variable $t^{\text{cold}} \geq t^{\text{off}}$ is introduced. The detailed cost functions will be discussed later in this section.

Values of the state variable below $-t^{\text{cold}}$ and above t^{on} do not influence feasible decisions, state dynamics or the start up cost. Therefore they are of no special interest and also represented by $x_t = -t^{\text{cold}}$ and $x_t = t^{\text{on}}$ respectively.

Since it takes time to increase the heat of a cold boiler to the point where production is possible, τ is used for describing the amount of time necessary for this. Similarly ν denotes the time used for shutting down the plant. Note, that during this time no output is generated. This leads to the the formulation of *state transition constraints*:

$$x_t = \begin{cases} \min(t^{\text{on}}, x_{t-1} + 1), & \text{if } 0 < x_{t-1} \text{ and } u_{t-1} = 1, \\ -1, & \text{if } x_{t-\nu} = t^{\text{on}} \text{ and } u_{t-\nu} = 0, \\ \max(-t^{\text{cold}}, x_{t-1} - 1), & \text{if } x_{t-1} < 0 \text{ and } u_{t-1} = 0, \\ 1, & \text{if } x_{t-\tau} \leq -t^{\text{off}} \text{ and } u_{t-\tau} = 1 \end{cases} \quad (1.4)$$

For short, the states in which controls can be applied are:

$$\Phi = \{x_t | x_t = t^{\text{on}} \text{ or } -t^{\text{cold}} \leq x_t \leq -t^{\text{off}}\} \quad (1.5)$$

Figure 1.1 provides an example of how the state transition and control constraints can be applied. The blue circles represent states where the plant is running, red circles states where the plant is turned off. The arrows show the possible state transformations. The background colors represent feasible controls: Staying online or shutdown of the plant if light blue; staying offline or possible startup if light red and no applicable control if white.

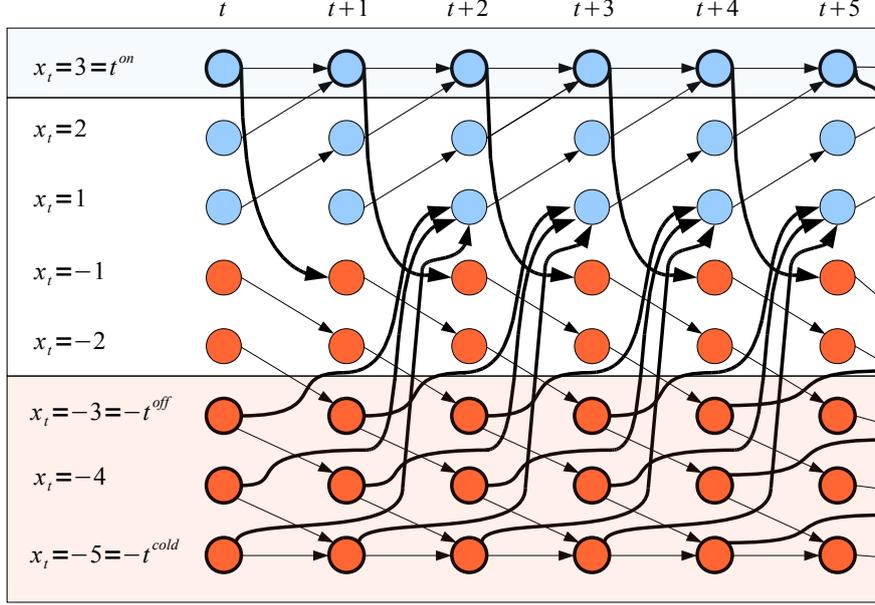


Figure 1.1: State transition ($t^{on} = t^{off} = 3$, $t^{cold} = 5$; $\tau = 2$, $\nu = 1$)

1.3.2 Cost function and optimal production

For modeling the revenue of running a power plant, the cost and the benefit of producing a unit of power has to be measured.

The benefit of a running plant is given by produced energy times price at the period or $p_t^E \cdot q$ with q being the optimal amount to be generated. The cost of production on the other hand uses the concept of the *heat rate* mentioned above. The heat rate H , the conversion rate from fuel to power, is assumed to be depending on the generated electricity:

$$H(q) = a_0 + a_1q + a_2q^2 \quad (1.6)$$

In addition to the production cost $H \cdot p_F$, the operator also has to bear the cost of any control action taken in this period. Here, the shutdown of an online plant is assumed to cost a constant amount of money C_{shut} . The effort needed for starting the plant however depends on the remaining boiler heat. Thus the startup cost, given by $C_{start}(x_t)$ for $-t^{cold} \leq x_t \leq -t^{off}$, is state depending. Similar to Tseng and Barz, the cost of control actions is modeled by

$$C = \begin{cases} C_{shut} & \text{if } x_t = t^{\text{on}} & \text{and } u_t = 0 \\ b_1(1 - \exp(x_t/\gamma)) + b_2 & \text{if } -t^{\text{cold}} \leq x_t \leq -t^{\text{off}} & \text{and } u_t = 1 \\ 0 & \text{else} \end{cases} \quad (1.7)$$

Aggregating cost and benefit of the generated electricity amount q yields a revenue of

$$f(p^E, p^F) := p^E q - (a_0 + a_1 q + a_2 q^2) p^F - C \quad (1.8)$$

Finally, from this equation the optimal power output of the plant for revealed prices can be deduced using standard calculus. At this point another physical constraint relating power production is introduced: The variable q_{min} describes the lower, q_{max} the upper bound of possible output.

$$q = \operatorname{argmax} f(q) \quad \text{subject to} \quad q_{min} \leq q \leq q_{max} \quad (1.9)$$

Assuming $a_2 \neq 0$, a unique solution satisfying the restrictions on q from (1.9) can be obtained:

$$\begin{aligned} f'(q) &= p^E - (a_1 + 2 a_2 q) p^F \stackrel{!}{=} 0 \\ q &= \left(\frac{p^E}{p^F} - a_1 \right) \frac{1}{2a_2} \\ q^{opt} &= \max \left(\min \left(q_{max}, \left(\frac{p^E}{p^F} - a_1 \right) \frac{1}{2a_2} \right), q_{min} \right) \end{aligned}$$

1.3.3 Model formulation

Expressing the optimal output as a function of prices concludes the model deduction in this section. At this point it seems convenient to retrace the most important steps of modeling the power generation procedure so far: At first, the state variable of the simulation was chosen to be the time the plant was already turned on or turned off. This choice gave way to the determination of the state dynamics (equation (1.4)) and the potential influence to be taken on them (equation (1.3)). Then, cost and benefit of power production at an arbitrary time instance were determined (equation (1.7)). Finally, the merging of these aspects with the financial options approach (model (1.2)) yields the model of Tseng and Barz ((1.10)) considering both, price uncertainty and operational properties:

$$J_1(x_1; p_1^E, p_1^F) = \max_{u_1, \dots, T} \mathbb{E}_1 \left(\sum_{t=1}^T (f(x_t; p_t^E, p_t^F) - C(x_t, u_t)) \right) \quad (1.10)$$

subject to state dynamics (equation (1.4))

$$x_t = \begin{cases} \min(t^{\text{on}}, x_{t-1} + 1), & \text{if } 0 < x_{t-1} & \text{and } u_{t-1} = 1, \\ -1, & \text{if } x_{t-\nu} = t^{\text{on}} & \text{and } u_{t-\nu} = 0, \\ \max(-t^{\text{cold}}, x_{t-1} - 1), & \text{if } x_{t-1} < 0 & \text{and } u_{t-1} = 0, \\ 1, & \text{if } x_{t-\tau} \leq -t^{\text{off}} & \text{and } u_{t-\tau} = 1 \end{cases}$$

$\mathbb{E}_t(\cdot)$ ($:= \mathbb{E}(\cdot | \mathfrak{F}_t)$) denotes the expectation conditional to the information available at t . Also, the control strategies at a given state are functions of the current price information only⁷ (see section 2.2.1 for a more detailed treatment). Their values are subject to the following control constraints (equation (1.3)):

$$u_t = \begin{cases} 1, & \text{if } 1 \leq x_t < t^{\text{on}} \\ 0, & \text{if } -t^{\text{off}} < x_t \leq -1 \\ 0 \text{ or } 1 & \text{otherwise} \end{cases}$$

⁷i.e. $u_t(x_t)$ are \mathfrak{F}_t measurable functions

Chapter 2

Deterministic and Stochastic Solution

Tseng and Barz provide an algorithm for solving their model (equation (1.10)) which uses a combination of 'forward moving Monte Carlo simulations and backward moving dynamic programming' as they describe it in their paper [1].

It is instructive to solve the problem first without stochastics. Besides, the deterministic approach was also implemented in MATLAB for debugging the more sophisticated Tseng and Barz algorithm.

This chapter provides solutions to both, a deterministic and a stochastic version of the plant evaluation problem. Also, the mathematical tools used such as Dynamic Programming and Monte Carlo simulation are introduced.

2.1 The deterministic case

Assuming perfect information about fuel and electricity prices $(p^E, p^F)_i$, $i = 1 \dots T$, an optimal solution to (1.10) can be obtained using Dynamic Programming.

2.1.1 Dynamic Programming

Finding a solution for model (1.10) requires the knowledge of an optimal control strategy $u_{1, \dots, T}$. Basically this yields a T dimensional optimization problem with additional constraints on control and state transition.

For this class of time depending optimal control problems, the Dynamic Programming algorithm provides a solution by transformation into a recurrence equation:

Assume the problem of maximizing a 'payoff' function $F(\cdot)$ over a time horizon T (inter temporary utility maximization problem). The actual contribution to the payoff at time t is depending on the state x_t and, in an interdependent way, the control variable u_{t-1} :

$$F_1(x_1) = \max_{u_1, \dots, T} \sum_{t=1}^T f^t(x_t, u_t) \quad (2.1)$$

subject to

$$x_{i+1} = h^i(x_i, u_i) \quad (\text{state dynamics}) \quad (2.2)$$

$$x_i \in X^i \quad (\text{state space}) \quad (2.3)$$

$$u_i \in U^i(x_i) \quad (\text{feasible decisions}) \quad (2.4)$$

In other words, the task is to find an optimal path x_1, \dots, x_T which maximizes the 'payoff'. At each time step k , the direction to be taken for the next step is determined by the control variable u_k .

The key to the transformation of this problem into a recurrence equation is the so called optimality principle¹:

Each part of an optimal path has to be optimal.

Intuitively this follows from the possibility to increase the payoff by replacing any non optimal part of a path.

Now consider problem (2.1) at an arbitrary time step $s, 1 \leq s < T$. Assuming an optimal payoff at $s+1, F_{s+1}(x_{s+1})$ (and thus the optimal path, depending on the state at $s+1, \tilde{x}(x_{s+1})_{s+2}, \dots, \tilde{x}(x_{s+1})_T$) is already known, the payoff maximization over s, \dots, T can be written as

$$F_s(x_s) = \max_{u_s, \dots, T} \sum_{t=s}^T f^t(x_t, u_t) = \max_{u_s} \left(f^s(x_s, u_s) + F_{s+1}(\underbrace{h^s(x_s, u_s)}_{x_{s+1}}) \right) \quad (2.5)$$

$$F_{T+1}(x_{T+1}) = 0$$

The right equality results from the optimality principle:

It guarantees that the optimal path from $s+1$ to T , starting at x_{s+1} has to contain the states $\tilde{x}(x_{s+1})_{s+2}, \dots, \tilde{x}(x_{s+1})_T$. Equivalently, this means that

¹see [5], p24

the additional benefit at stages $s + 1$ to T necessarily yields $F_{s+1}(h^s(x_s, u_s))$. The contribution at step s is given by $f^s(x_s, u_s)$. Obtaining the optimal value (depending on the initial value at s , x_s) of this problem is now just a one stage optimization problem of one parameter (u_s). This can be solved with standard analysis or, for discrete control as in the power plant operating problem, by the mere comparison of the possible outcomes for the different u_s .

Figure 2.1 illustrates a problem like this for $f^t(x_t, u_t) = \max(0, x_t)$. The green arrows represent optimal paths, the red ones illustrate an example of a suboptimal decision.

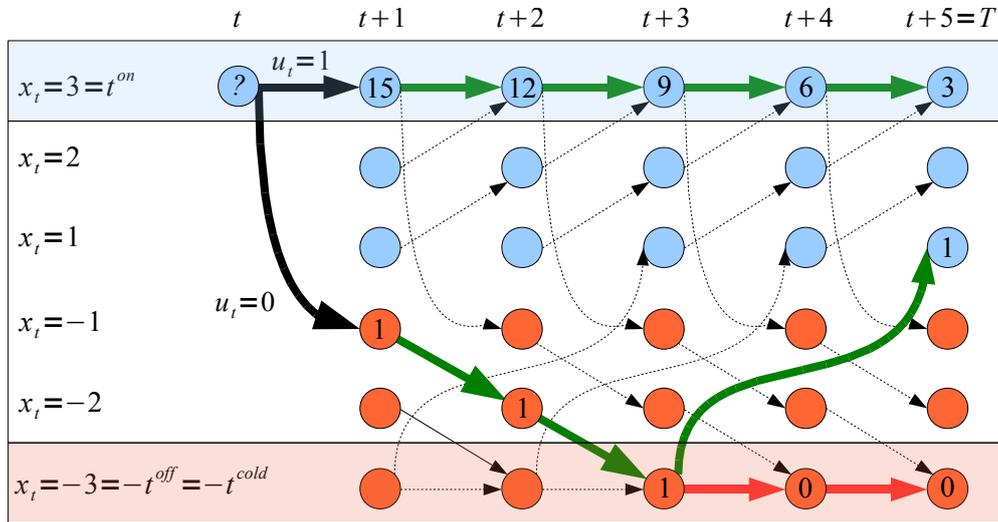


Figure 2.1: The optimization problem at step t

Using this stage depending approach, the whole recursion (2.5) can be solved by starting from the end of the horizon T . Obviously $F_T(x_T) = f(x_T)$ since no control will be applied at the end of the horizon and the boundary condition. At any further time steps t ($t < T$), the optimal payoff $F_t(x_t)$ can be computed as a function of the current state x_t by using the previously computed payoff $F_{t+1}(x_{t+1})$. The actual value of this inter temporary utility maximization problem will be found by using the initial state x_1 at $t = 1$.

Obviously, this presentation of Dynamic Programming has carefully avoided any proofs or precise arguments. For a more detailed introduction to this topic, see the seminal work of Richard Bellman [6] or the shorter book from

Gessner and Wacker [5].

With given prices, equation (1.10) turns into a Dynamic Programming problem as discussed above:

$$\begin{aligned} F_t(x_t) &= f(p_t^E, p_t^F, x_t) + \max_{u_t} [F_{t+1}(x_{t+1}) - C(x_t, u_t)] \\ F_{T+1}(x_{T+1}) &= 0 \end{aligned} \quad (2.6)$$

subject to equation (1.3), equation (1.4) and cost (1.7). The actual payoff from the power plant over the horizon T is given by $F_1(x_1)$, using the initial state of the plant as input.

With the Dynamic Programming algorithm sketched above, the deterministic plant evaluation (equation (2.6)) can be solved without difficulty.

2.1.2 Non standard representation of the deterministic model

Because of equation (1.3) it is only necessary to find an optimal control at states $x_t \in \Phi$, states $x_t \notin \Phi$ are just influenced by the state dynamics. This permits a different form of the recursions, which arranges the influence of operational constraints more clearly.

Including the restrictions on u_t and the state dynamics into the recurrence equation leads to three different cases:

- $x_t = t^{\text{on}}$ and $t \leq T - \nu$ (online plant and control feasible)

$$\begin{aligned} F_t(x_t) &= f(p_t^E, p_t^F, x_t) \\ &+ \max_{u_t} [u_t F_{t+1}(t^{\text{on}}) + (1 - u_t) (F_{t+\nu}(-1) - C_{\text{shut}})] \end{aligned} \quad (2.7)$$

- $-t^{\text{cold}} \leq x_t \leq -t^{\text{off}}$ and $t \leq T - \tau$ (offline plant and control feasible)

$$\begin{aligned} F_t(x_t) &= f(p_t^E, p_t^F, x_t) \\ &+ \max_{u_t} [u_t (F_{t+\tau}(1) - C_{\text{start}}(x_t)) \\ &+ (1 - u_t) (F_{t+1}(\max(x_t - 1, -t^{\text{cold}}))] \end{aligned} \quad (2.8)$$

- $x_t \notin \Phi$ or $t > T - \nu$ or $t > T - \tau$ (no control feasible)

$$F_t(x_t) = f(p_t^E, p_t^F, x_t) + F_{t+1}(x_t + 1) \quad \text{if } x_t > 0 \quad (2.9)$$

$$F_t(x_t) = f(p_t^E, p_t^F, x_t) + F_{t+1}(x_t - 1) \quad \text{if } x_t < 0$$

$$F_{T+1}(x_{T+1}) = 0 \quad (2.10)$$

The influence of the decision lead time on the control and state dynamics can be seen in the particular model formulation above. As well in (2.7) as in (2.8), the term $F_{t+\tau}$ or $F_{t+\nu}$ appears. Obviously, when determining the optimal unit control, one has not only to compare the plant values at different states but also at different time instances.

2.2 The stochastic model

In the first chapter, the stochastic problem was introduced as model of Tseng and Barz (model (1.10)). It was derived by the incorporation of physical restraints into an evaluation approach based on financial options. However, for the solution of the Tseng and Barz model it is instructive to interpret it as the generalization of the deterministic problem (section 2.1) to stochastic prices.

In the following section, price pairs (p_i^E, p_i^F) will be denoted as p_i and sequences of variables like $(u_i)_{i=s, \dots, T}$ as u_s^T for a more compact notation.

2.2.1 Deriving the stochastic model

Assuming a given control sequence u_s^T , the mere replacement of deterministic prices by random variables in the deterministic model (2.6) turns the future payoff into a random variable:

$$P_s(x_s, u_s^T; p_s^T) = \sum_{t=s}^T [f(p_t, x_t) - C(x_t, u_t)] \quad (2.11)$$

s.t. $x_t = h(x_{t-1}, u_{t-1})$

Obviously the future payoff is dependent on the current state, the sequence of future prices and the control strategy, since the future states x_s^T are determined by x_s and the controls u_s^T .

The value of the plant will be defined as the maximal expected payoff of the plant (over the horizon $1, \dots, T$) conditional to the price information currently available:

$$J_s(x_s; p_s^s) = \max_{u_s^T} \mathbb{E} \left(\sum_{t=s}^T (f(x_t; p_s) - C(x_t, u_t)) \middle| \mathfrak{F}_s \right) \quad (2.12)$$

$$x_{i+1} = h^i(x_i, u_i) \quad (\text{state dynamics}) \quad (2.13)$$

The usage of conditional expectations has the benefit of utilizing the memory of the electricity and fuel price processes for the evaluation. The expectation operator conditional to the information at time s is represented by $\mathbb{E}(\cdot|\mathfrak{F}_s)$, where $\mathfrak{F}_s := \sigma(p_1^s)$ denotes the σ algebra generated by the observed sequence of prices up to the current time step. The decision sequence maximizing the payoff at s is determined depending on observed prices and the current state but not on future prices $s + 1, \dots, T$. Thus, for describing the optimal control, it is necessary to define a control as \mathfrak{F}_s measurable mapping from the space of observed prices into the feasible states, assigning a value of the control variable to each price observation:

$$\begin{aligned} u_s(x_s) : \mathbb{R}^{s \times 2} &\longrightarrow U^s(x_s) \\ p_1^s &\longmapsto 0 \text{ or } 1 \end{aligned} \quad (2.14)$$

As in section 2.1, the key to the solution of the stochastic model is to rewrite equation (2.12), to exchange the $T - s + 1$ dimensional optimization problem for a recurrence equation in $T - s + 1$ steps. By virtue of the linearity of \mathbb{E} , equation (2.12) can be extended to

$$\begin{aligned} J_s(x_s; p_1^s) = & \quad (2.15) \\ \max_{u_s^T} & \left[\mathbb{E} \left(f(x_s; p_s) - C(x_s, u_s) \mid \mathfrak{F}_s \right) + \underbrace{\mathbb{E} \left(\sum_{t=s+1}^T (f(x_t; p_t) - C(x_t, u_t)) \mid \mathfrak{F}_s \right)}_{=P_{s+1}(x_{s+1}, u_{s+1}^T; p_{s+1}^T)} \right] \end{aligned}$$

Since prices at time s are known and due to the tower property of the conditional expectation, equation (2.15) can be rewritten as

$$\begin{aligned} J_s(x_s; p_1^s) = & \quad (2.16) \\ \max_{u_s} & \left[f(x_s; p_s) - C(x_s, u_s) + \max_{u_{s+1}^T} \mathbb{E} \left(\mathbb{E} (P(x_{s+1}, u_{s+1}^T; p_{s+1}^T) \mid \mathfrak{F}_{s+1}) \mid \mathfrak{F}_s \right) \right] \end{aligned}$$

Switching $\mathbb{E}(\cdot|\mathfrak{F}_s)$ with $\max_{u_{s+1}^T}$ in equation (2.16) leads to the recursive form to be introduced:

$$\begin{aligned} J_s(x_s; p_1^s) = & \quad (2.17) \\ \max_{u_s} & \left[f(x_s; p_s) - C(x_s, u_s) + \underbrace{\mathbb{E} \left(\max_{u_{s+1}^T} \mathbb{E} \left(P(x_{s+1}, u_{s+1}^T; p_{s+1}^T) \mid \mathfrak{F}_{s+1} \right) \mid \mathfrak{F}_s \right)}_{J_{s+1}(x_{s+1})} \right] \end{aligned}$$

2.2.2 A detailed view of the stochastic model

Equation (2.17) gives the compact recursive form similar to the deterministic (2.6):

$$J_s(x_s; p_1^s) = \max_{u_s} \left[f(x_s; p_s) - C(x_s, u_s) + \mathbb{E} \left(J_{s+1}(x_{s+1}) \middle| \mathfrak{F}_s \right) \right] \quad (2.18)$$

$$x_{i+1} = h^i(x_i, u_i) \quad (2.19)$$

$$u_s(x_s) : \mathbb{R}^{s \times 2} \longrightarrow U^s(x_s); u_t(x_t) \dots \mathfrak{F}_s \text{ measurable} \quad (2.20)$$

$$J_{T+1} = 0 \quad (2.21)$$

Again, the recursion in some sense mirrors the optimality principle from the dynamic programming approach: Assuming the optimal control sequence for $s + 1, \dots, T$ was known, the payoff over this period would per definition² yield $J_{s+1}(x_{s+1})$. Using this, the optimal one-step control can be determined.

The recursive form shows the appearance of two separate optimization problems³:

Implicitly included in $f(x_s; p_s)$, the optimal power output at the current state of the plant is computed for revealed prices (equation 1.9). This represents the automatic generator dispatch of an online plant, which reacts to the current demand in real time.

The second optimization problem concerns long term unit commitment, determining if it pays to keep the unit online at future states. Because of the operational constraints of the plant, this decision process is governed by volatile future prices.

Note that, since the optimal control strategy is defined as the control sequence maximizing the expected future payoff conditional to current information, the operator is assumed to be risk neutral. In other words, the operator is only interested in maximizing the plant's expected profit.

Introducing the Markov property of prices

Although fuel and electricity prices have been used since the very beginning of this work, they have not been discussed yet at all. As a matter of fact, the detailed properties of the price processes used in the actual model will be described in the following section. The reason for this is the applicability of model (2.18) to a variety of price models.

Previously in this section $\mathfrak{F}_t = \sigma(p_1^t)$ was used to denote the information

²equation (2.12)

³see also section 1.3.2

available at time t . As a result, $\mathbb{E}(\cdot|\mathfrak{F}_t)$ as well as $J(x_t)$ and $u(x_t)$ were functions of all prices p_1^t .

However, the solution to be presented will assume both fuel and electricity price processes to have the Markov property⁴.

This means that the distribution of future prices p_{t+1}^T conditional to the current price observation p_t equals their distribution conditional to the observations from p_1^t . In particular, for any function $g(p_{t+1}^T)$, the expectation conditional on \mathfrak{F}_s is equal to the expectation conditional on $\sigma(p_t)$, i.e.

$$\mathbb{E}(g(p_{t+1}^T)|\mathfrak{F}_t) = \mathbb{E}(g(p_{t+1}^T)|\sigma(p_t)) =: \mathbb{E}_t(g(p_{t+1}^T)) \quad (2.22)$$

This implies that $J_t(x_t, p_1^t)$ is a function of p_t only, which therefore will be written as $J_t(x_t, p_t)$ subsequently. As another consequence, the optimal control u_t only depends on current state x_t and current prices p_t .

Non standard representation of the stochastic model

Additional to these fundamental assumptions, section 1.3 provided a detailed description of the physical properties modeled.

As in the deterministic model, it is possible to incorporate these physical restrictions on u_t and the state dynamics into the recurrence equation:

- $x_t = t^{\text{on}}$ and $t \leq T - \nu$ (running plant and control feasible)

$$\begin{aligned} J_t(x_t; p_t) &= f(x_t; p_t) \\ &+ \max_{u_t} \mathbb{E}_t[u_t J_{t+1}(t^{\text{on}}; p_{t+1}) \\ &+ (1 - u_t) (J_{t+\nu}(-1; p_{t+\nu}) - C_{shut})] \end{aligned} \quad (2.23)$$

- $-t^{\text{cold}} \leq x_t \leq -t^{\text{off}}$ and $t \leq T - \tau$ (shut down plant and control feasible)

$$\begin{aligned} J_t(x_t; p_t) &= f(x_t; p_t) \\ &+ \max_{u_t} \mathbb{E}_t[u_t (J_{t+\tau}(1; p_{t+\tau}) - C_{start}(x_t)) \\ &+ (1 - u_t) (J_{t+1}(\max(x_t - 1, -t^{\text{cold}}); p_{t+1}))] \end{aligned} \quad (2.24)$$

- $x_t \notin \Phi$ or $t > T - \nu$ or $t > T - \tau$ (no control feasible)

$$\begin{aligned} J_t(x_t; p_t) &= f(x_t; p_t) + \mathbb{E}_t J_{t+1}(x_t + 1; p_{t+1}) \text{ if } x_t > 0 \\ J_t(x_t; p_t) &= f(x_t; p_t) + \mathbb{E}_t J_{t+1}(x_t - 1; p_{t+1}) \text{ if } x_t < 0 \\ J_{T+1}(x_{T+1}) &= 0 \end{aligned} \quad (2.25)$$

⁴again resembling the approach of Tseng and Barz

Following Tseng and Barz, $J_t(x_t; p_t)$ denotes the value of the the plant over the time period t to T , given observed prices $p_t = (p_t^E, p_t^F)$ at time t . Note, that this description of the model is consistent to the more general model (2.18) when (2.19) and (2.20) are describing the physical plant properties introduced in section 1.3.

2.3 Integration using Monte Carlo simulation

In the deterministic case, Dynamic Programming provided an excellent algorithm for the optimal plant control. In the stochastic case, the usage of conditional expectations in equations (2.23) to (2.25) urges to look for ways to compute $\mathbb{E}_t J_{t+1}(x_{t+1}; p_{t+1}^E, p_{t+1}^F)$. Generally $J(\cdot)$ is not a 'nice' function since it is not linear in prices and even not necessarily continuous. Especially those aspects suggest the usage of non - analytic methods such as Monte Carlo integration. The subsequent section is based on chapter 5 of [4] and section 2.4 from [3].

2.3.1 Monte Carlo integration

Primarily, Monte Carlo simulation is a method for evaluating integrals using (pseudo-) random numbers. The simplest problem of this kind is given by determining the value of an integral $\int f(x)dx$ over the domain $[0, 1]$. In terms of Monte Carlo simulation, this is done by identifying this integral with the expected value of the transformation $f(X)$ of a uniformly distributed random variable, $X \sim U([0, 1])$. From the properties of the uniform distribution follows

$$\theta = \int_0^1 f(t)dt = \int f(x) \underbrace{\mathbb{I}_{[0,1]}(x)}_{\text{density of X}} dx = \mathbb{E}(f(X)) \quad (2.26)$$

Therefore, evaluating the integrand f at N random samples x_i from $U([0, 1])$ and taking the mean of these values gives an estimate of θ . This estimate converges towards the exact value of the integral, following the law of large numbers. Moreover, the Monte Carlo estimator $\hat{\theta}$ is unbiased:

$$\hat{\theta}_N := \frac{1}{N} \sum_{i=1}^N f(x_i) \quad (2.27)$$

$$\mathbb{E}(\hat{\theta}_N) = \frac{\sum_{i=1}^N \mathbb{E}(f(X_i))}{N} = \mathbb{E}(f(X)) = \int_0^1 f(x)dx$$

The variance of the Monte Carlo estimator θ_N is given by

$$\mathbb{V}(\hat{\theta}_N) = \frac{\sum_{i=1}^N \mathbb{V}(f(X))}{N^2} = \frac{1}{N} \cdot \mathbb{V}(f(X)) \quad (2.28)$$

From equation (2.28) some convergence properties of the Monte Carlo estimate can be derived:

If $\mathbb{V}(f(X))$ is bounded, the speed of convergence of $\mathbb{V}(\hat{\theta}_N)$ is $O(\frac{1}{n})$. Moreover, equation (2.29) shows that the precise speed is depending on the roughness⁵ of the integrand, to be seen at the very far right.

$$\mathbb{V}(f(X)) = \mathbb{E}(f(X) - \mathbb{E}(f(X)))^2 = \int_0^1 \left(f(x) - \int_0^1 f(x)dx \right)^2 dx \quad (2.29)$$

2.3.2 Variance reduction: antithetic variates

Due to the nature of the method, the Monte Carlo estimate is blurred by an inherent error. Especially for computationally intensive problems it pays to reduce this variance by smart ways of using pseudo random numbers. Those more or less subtle methods alter the variance of the Monte Carlo estimator without changing its expected value.

Here, the method of antithetic variates will be explained shortly since it is referred to in subsequent parts.

Consider again the problem of finding the value of (2.26). The Monte Carlo estimator for this is given by (2.27), its estimate is based on x_i from a uniform distribution on $[0, 1]$. An intuitive way of avoiding outliers among those random values is to add for each drawn x_i its mirrored value $x'_i := (1 - x_i)$ to the sample. If $f(x)$ is monotonous⁶, this reduces the variance of the estimate because of the influence of the covariance term:

$$\begin{aligned} \mathbb{V} \left(\frac{\sum_{i=1}^{\frac{N}{2}} (f(x_i) + f(x'_i))}{N} \right) &= \\ &= \frac{1}{N^2} \sum_{i=1}^{\frac{N}{2}} [\mathbb{V}(f(x_i)) + \mathbb{V}(f(x'_i)) + 2\text{Cov}(f(x_i), f(x'_i))] \\ &< \frac{\sum_{i=1}^N \mathbb{V}(f(x_i))}{N^2} \end{aligned}$$

⁵defined as the difference of a function from its integrated value, [4], p132

⁶ f monotonous $\Rightarrow \text{Cov}(f(x_i), f(x'_i)) < 0$

As intended, equation (2.30) shows that the Monte Carlo estimator stays unbiased.

$$\mathbb{E} \left(\frac{\sum_{i=1}^{\frac{N}{2}} (f(x_i) + f(x'_i))}{N} \right) = \mathbb{E}(f(x)) \quad (2.30)$$

Thus, this easily implemented method can result in a reduction of computing time. Unfortunately, the degree of this depends very much on the properties of the integral to be computed.

2.4 The iterative solution for the stochastic model

2.4.1 Describing optimal control strategies

The fundamental difficulty at assessing the plant's value is to determine the optimal control decision of the operator. Using the recursive formulation of the model, the optimal (one step) control at fixed state x_t is the mapping from \mathbb{R}_+^2 into $\{0, 1\}$, assigning to any observed price pair the feasible control which maximizes the one-step optimization in (2.23) or (2.24). Clearly, the optimal control value is $u_t = 1$ (turning on) when $J_t(x_t; p_t^E, p_t^F | u_t = 1)$ exceeds $J_t(x_t; p_t^E, p_t^F | u_t = 0)$ and $u_t = 0$ for the opposite case. This relation can be utilized to describe the optimal control:

The difference of expected payoffs is introduced as function of current state and prices:

$$d_t(x_t; p_t^E, p_t^F) := J_t(x_t; p_t^E, p_t^F | u_t = 1) - J_t(x_t; p_t^E, p_t^F | u_t = 0) \quad (2.31)$$

Then, for the optimal control strategy u^* , equation (2.32) holds:

$$u_t^* = \begin{cases} 1, & \text{if } d_t(x_t; p_t^E, p_t^F) > 0, \\ 0, & \text{if } d_t(x_t; p_t^E, p_t^F) < 0 \end{cases} \quad (2.32)$$

In other words, for each state x_t the $p_t^E \times p_t^F$ plane of observed prices at time t is divided in two parts determined by the sign of $d_t(\cdot)$. For any observed price pair, the optimal unit control variable u_t can be found by determining which of the sets it belongs to.

2.4.2 Describing $d(\cdot)$ with indifference loci (*IL*)

Needing a smart (computationally cheap) way to determine the separation of the $p_t^E \times p_t^F$ plane described by $d_t(\cdot)$, Tseng and Barz used the following approach: They looked for price pairs (p_t^E, p_t^F) satisfying $d_t(x_t; p_t^E, p_t^F) = 0$. In terms of the operational value, these are the solutions of

$$J_t(x_t; p_t^E, p_t^F | u_t = 1) = J_t(x_t; p_t^E, p_t^F | u_t = 0) \quad (2.33)$$

which represent observed prices, at which the operator is indifferent between turning on and shutting down the plant in terms of profit. Those price pairs are referred to as *indifference loci (IL)* in the rest of the text. The IL are used as a barrier between prices implicating the turn on ($u = 1$) and the turn off ($u = 0$) of the plant in the following way:

Let \bar{x}_t be a fixed state, \bar{p}_t^E an arbitrary electricity price and p_t^{F*} the fuel price satisfying $d_t(\bar{x}_t; \bar{p}_t^E, p_t^{F*}) = 0$. Then it is assumed that $d(\bar{x}_t; \bar{p}_t^E, p^F) > 0$ if $p^F < p^{F*}$ and $d(\bar{x}_t; \bar{p}_t^E, p^F) < 0$ if $p^F > p^{F*}$ respectively. In other words, in the $p_t^E \times p_t^F$ plane, the observed prices implying $u = 1$ are located in the lower right of the IL; the observations implying $u = 0$ are found in the upper left.

This method is based on the fact that $d(\cdot) > 0$ for small fuel prices (and vice versa) and the assumption that $d(\cdot)$ is monotonous in p^F ⁷.

Figure 2.2 illustrates the concept of indifference loci: The value of a plant at a given time step t is computed at various price pairs for both possible controls; $J_t(x_t; p_t^E, p_t^F | u_t = 1)$ is represented by the red plane, $J_t(x_t; p_t^E, p_t^F | u_t = 0)$ by the blue one. Naturally the optimal control will be 1 for price pairs which promise a greater profit than $u = 0$. This corresponds to the light red area in the $p_t^E \times p_t^F$ plane, the other case to the light blue area. The border between those areas is the indifference locus; for those price pairs, both controls are expected to provide the same payoff.

2.4.3 Solution of the model of Tseng and Barz

Value at time s for known $d_t(\cdot), t > s$

The value of the plant at given state x_t is determined by the recursions (2.23) to (2.25). As in the deterministic case, the expected future payoff is computed starting from the end of the horizon backwards in time.

⁷this does not apply in general, see section 4.3.2 for a counterexample

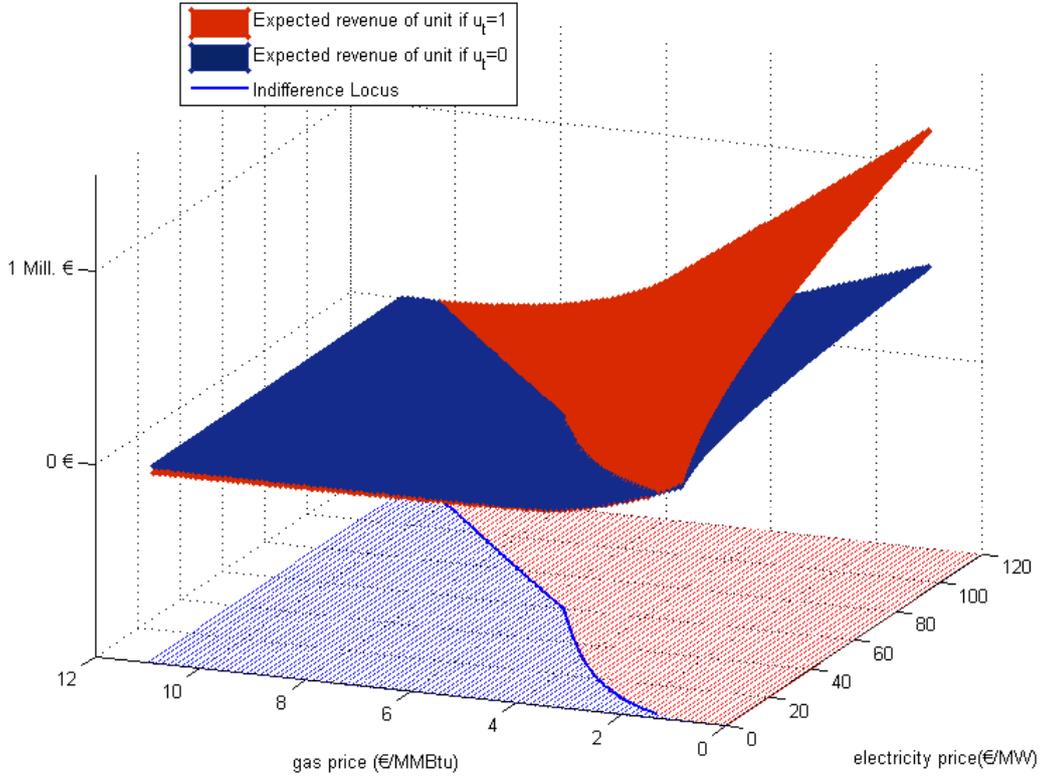


Figure 2.2: IL as intersection between $J_t(t^{\text{on}}; p_t^E, p_t^F | u_t = 1)$ and $J_t(t^{\text{on}}; p_t^E, p_t^F | u_t = 0)$

Assume the functions $(d_t(\cdot))_{t=s+1, \dots, T}$ to be known at an arbitrary time step s ($1 \leq s < T$). (This information corresponds to the optimal path in the deterministic case.)

At any given state x_s and current prices (p_s^E, p_s^F) , the expected future payoff for fixed control u_t , $\mathbb{E}_s(J_{s+1}(x_{s+1}, u_{s+1}; p_{s+1}^E, p_{s+1}^F))$ can be obtained using Monte Carlo simulations. This is done by generating a number of price scenarios starting with (p_s^E, p_s^F) according to the underlying price model. Since the sequence of controls is determined by the $(d_t(\cdot))_{t>s}$, the future payoff of the plant can be computed for each price scenario. The mean of these values is a Monte Carlo estimate of the expected payoff over $s+1, \dots, T$ conditional on observed prices (p_s^E, p_s^F) .

Estimating $\mathbb{E}_s(J_{s+1}(x_{s+1}, u_{s+1}; p_{s+1}^E, p_{s+1}^F))$ for known $d_t(\cdot), t > s$

A more instructive description of the Monte Carlo evaluation is the algorithm form⁸:

- 0 : *initial conditions*: $s, (p_s^E, p_s^F), (d_k(\cdot))_{k>s}$ and x_{s+1} are given;
 N is determined, $i = 0$
- 1 : *initialize*: set $t = s + 1, i = i + 1$ and $J^{(i)} = 0$
- 2 : *generate scenario*: if $i \leq N$ simulate price sequence $(p_k^{E(i)}, p_k^{F(i)})$,
 $k = s, \dots, T$, starting with initial (p_s^E, p_s^F) ;
 else return $(J^{(i)})_{i=1, \dots, N}$.
- 3 : *actuate control*: determine u_t by (2.32)
- 4 : *actuate payoff*: $J^{(i)} = J^{(i)} + f(x_t; p_t^{E(i)}, p_t^{F(i)}) - C(x_t, u_t)$
- 5 : *update system*: determine time step depending on u_t (either $t = t + 1$,
 $t = t + \tau$ or $t = t + \nu$), then determine x_t
- 6 : *iterate*: if $t \leq T$ go to 3, else go to 1

As indicated before, an estimate for the expected value of the plant is given by the mean of simulated plant values

$$\mathbb{E}_s(J_{s+1}(x_{s+1}, u_{s+1}; p_{s+1}^E, p_{s+1}^F)) \approx \frac{1}{N} \sum_{i=1}^N J^{(i)} \quad (2.34)$$

For the actual value of the plant at time s , equation (2.23) or (2.24) can be used, depending on the current state.

Determining IL

Having solved the evaluation problem for known future optimal decisions, the idea of the further proceeding is similar to the deterministic case. Combining equations (2.31) and (2.24) or (2.23) respectively, the IL at s can be expressed as the price pairs solving (2.35) or (2.36):

- if $x_s = t^{\text{on}}$ and $s \leq T - \nu$

$$\mathbb{E}_s J_{s+1}(t^{\text{on}}; p_{s+1}^E, p_{s+1}^F) - \mathbb{E}_s [J_{s+\nu}(-1; p_{s+\nu}^E, p_{s+\nu}^F) - C_{shut}] = 0 \quad (2.35)$$

⁸see also [1], Simulation Algorithm for J_{t_0}

- $-t^{\text{cold}} \leq x_s \leq -t^{\text{off}}$ and $s \leq T - \tau$

$$\begin{aligned} & \mathbb{E}_s [J_{s+\tau}(1; p_{s+\tau}^E, p_{s+\tau}^F) - C_{start}(x_s)] \\ & - \mathbb{E}_s J_{s+1}(\max(x_s - 1, -t^{\text{cold}}); p_{s+1}^E, p_{s+1}^F) = 0 \end{aligned} \quad (2.36)$$

Therefore the IL at s can be estimated by the algorithm above using only IL from $s + 1, \dots, T - \tau$ (or $T - \nu$). Starting at the end of the time horizon, optimal decision rules expressed by IL can be computed iteratively.

An IL does not exist at each time step, since there is no control action when the remaining evaluation time is shorter than decision lead time; that is at $t > T - \tau$ for a cold plant or at $t > T - \nu$ for an online plant respectively.

Note that an indifference locus consists of an infinite number of price pairs; therefore the computed IL will only be an approximation. Besides, the finding of the price pairs will be very expensive in terms of computing time, since at each evaluation of (2.31) the simulation of $2N$ scenarios is needed.

Chapter 3

Implementation in MATLAB

While the previous two chapters introduced the model and provided the necessary mathematical tools for solution of the plant evaluation problem, the remaining chapters will document the functionality of the created computer programs and discuss the obtained empirical results.

This chapter is entirely dedicated to the explanation of basic properties of the developed MATLAB program code.

The first section will explain the concept behind the procedure and introduce technical necessities such as simulation of price processes and the approach to the solution of (2.33).

The second section will contain brief explanations of the procedures used.

3.1 The big picture

The idea behind the Tseng and Barz algorithm was already discussed in section 2.4.3. Neglecting the details, its core task was the iterative computation of the IL, starting from the end of the evaluation period.

Looking deeper into the details, the proceeding gets more complicated: Obtaining the IL consists of finding the zeros of a function, which is evaluated by Monte Carlo simulation, which use the indifference loci of previous stages. This point of view suggests the division of the evaluation problem into following subproblems:

- (P1) Administrating the search for, and the storage of IL
- (P2) Evaluating the determining equation of IL (equation (2.33))
- (P3) Approximating the IL

(P4) Generating price scenarios for Monte Carlo simulation

(P5) Providing optimal control based on available IL

However, growing experience showed that the model is complex enough to turn

(P6) Initializing simulation parameters and storing the obtained data

into an additional problem.

While (P1), (P2), (P5) and (P6) have already been explained or are of no mathematical difficulty, methods for (P3) and (P4) have not been discussed yet.

3.1.1 Generating price scenarios for Monte Carlo simulation

Already in the previous chapters, fuel and electricity prices from hourly spot markets have been included in the models and used at Monte Carlo evaluation. However, beside the Markov property (and the existence of the population mean at Monte Carlo simulation) no underlying properties of the prices have been specified.

Indeed the solution method presented in section 2.4 can be used for many price models.

The underlying processes for both fuel and electricity prices are modeled as 'geometric' processes:

Let p_t denote the price and m_t the hourly pattern of logarithmized prices (see equation (3.2) for details). Then for $z_t := \ln(p_t) - m_t$ the price model is given by the AR(1) process

$$z_t = \alpha z_{t-1} + \sigma B_t \quad (3.1)$$

where $B_t \sim N(0, 1)$ are independent and normally distributed.

This price model is equivalent to the one used by Tseng and Barz, who generated the prices by a discrete version of a mean reverting geometric process¹.

¹[1], Appendix A; [2], p42; [3], p36

Estimating the model parameters for electricity

For fitting the parameters in (3.1) for the electricity model from N historic prices, following steps have been taken:

- Estimate the time depending pattern m_t of the price logarithms $\ln(p_t)$. Patterns for different time periods are derived by simply taking separate means of the concerned periods.

For an hourly pattern for example, this is done by

$$m_h = 24/N \sum_{i=0}^{\lfloor N/24 \rfloor} \ln(p_{h+i \cdot 24}) \quad (3.2)$$

$m_t = m_k$ for $k = t$ modulo 24.

- Subtract the pattern from the logarithms of prices $z_t := \ln(p_t) - m_t$.

- Maximum likelihood estimates for α_E and σ_E are given by

$$\hat{\alpha}_E = \frac{\sum_{i=2}^N z_i z_{i-1}}{\sum_{i=1}^N z_i^2} = 0.9086 \quad (3.3)$$

$$\hat{\sigma}_E^2 = \frac{1}{N} \sum_{i=2}^N (z_i - \hat{\alpha}_E z_{i-1})^2 = 0.0190 \quad (3.4)$$

Finally, random realizations of the estimated process can be simulated by reverting the procedure:

- Generate z_i following the AR(1) process in (3.1), with $B_i \sim N(0, 1)$
- Add the seasonal pattern m_t and derive prices by $p_t = \exp(z_t + m_t)$

The parameter values for the electricity price process used in the simulations are derived from the hourly electricity spot prices from January to September 2008. The data was taken from the webpage² of the Austrian Energy Exchange (EXAA).

Figure 3.1 compares some simulated price trajectories and the hourly price pattern to the real prices from January 1st to January 7th 2008. Of course it is inappropriate to judge the performance of the price simulation by the comparison of three random realizations with real data at an arbitrary time period. However, it is interesting to see the influence of holidays on prices: hours 1 to 24 (January 1st) and hours 120 to 144 (Sunday January 6st) show particularly low prices.

²http://www.exaa.at/market/historical/austria_germany/index.html

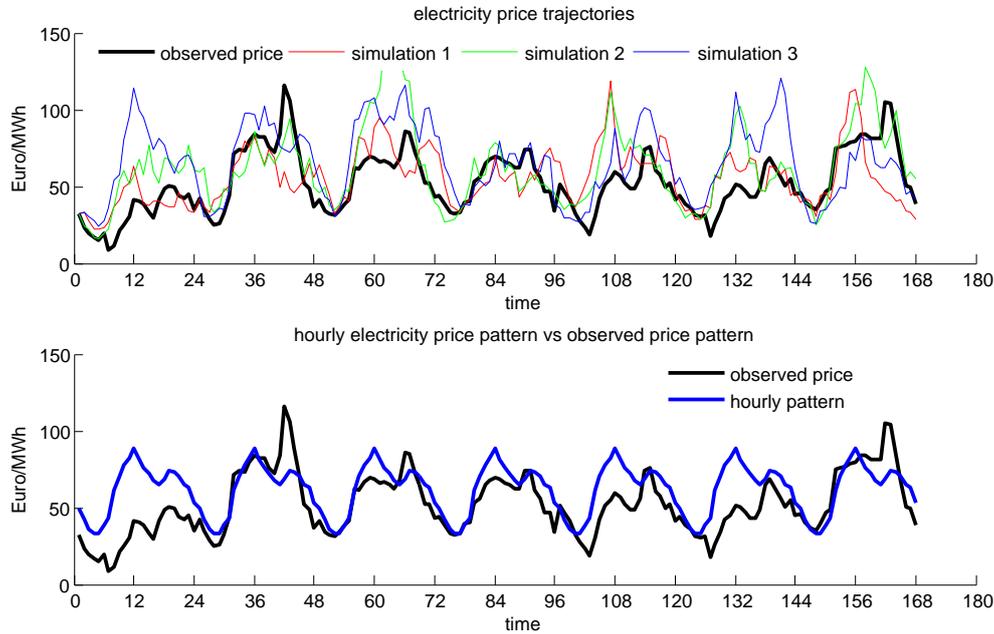


Figure 3.1: Simulations of electricity price using price model (3.1)

Estimating the model parameters for natural gas

The price process for natural gas (the fuel assumed to be used) will also be described by a geometric mean reverting process (3.1).

Unfortunately, contrary to the model assumption of an hourly market, natural gas is traded on a daily basis. This problem can be avoided by fitting a price process with daily fluctuations. Hourly gas prices can be obtained either by interpolation between the resulting values, or by keeping one price for every 24 hours.

However, Tseng and Barz state to use a model based on hourly time steps³. The parameters of this process are deduced from the parameters of the daily process in a not specified way. One advantage of simulated hourly gas prices is the possibility to incorporate a correlation between gas and electricity prices as observed in reality.

Sticking to the strategy of following closely the approach of Tseng and Barz, the values for $\hat{\alpha}_F = 0.9993$ and $\hat{\sigma}_F = 0.019$, as well as a correlation of 0.4 between electricity and natural gas prices will be taken from [1].

The mean of the process m , will be estimated from real data provided by the

³[1], section 4.1

European Energy Exchange⁴.

Using these parameters, the generation of fuel price trajectories happens in the same way as the simulation of electricity prices above.

3.1.2 Approximating the IL

As previously discussed, an IL consists of an infinite number of price pairs, satisfying equation (2.33). The practical determination consists of obtaining a sufficient number of points on the IL, to determine a reasonable approximation. Tseng and Barz propose to fix p_t^E at level $\hat{p}_t^{E(i)}$ and then solve the equation

$$h(y^{(i)}) = J_t(\hat{x}_t; \hat{p}_t^{E(i)}, y^{(i)} | u_t = 1) - J_t(\hat{x}_t; \hat{p}_t^{E(i)}, y^{(i)} | u_t = 0) = 0 \quad (3.5)$$

as function of one variable, $y^{(i)}$ (replacing p_t^E). The resulting fuel electricity pairs $(y^{(i)}, \hat{p}_t^{E(i)})$ provide estimates of the IL.

Since evaluating $h(y)$ is a separate problem (P4), the tasks left are to find the zeros of this function and to approximate the IL from them.

The function *fzero.m*, already contained in MATLAB provides the method of choice to find the zeros of $h(y)$.

The actual IL will be approximated by linear or cubic spline interpolation between the zeros found by *fzero.m*.

The cubic spline interpolation is done by the MATLAB routine *interp1.m* where the property 'method' is specified as 'cubic'. This procedure approximates an one dimensional function by interpolating with a cubic polynome between the function evaluations (Piecewise Cubic Hermite Interpolating Polynomial)⁵.

Especially for a large number of Monte Carlo simulations it saves time to fit splines compared to computing a high amount of points on the IL.

3.2 The program

According to the problem division at the beginning of this chapter, the subproblems (P1) to (P6) have been assigned to different functions. The interdependences of those are depicted in figure 3.2. The blue boxes represent

⁴<http://www.eex.com/de/Marktdaten/Handelsdaten/Erdgas/Natural%20Gas%20Day-Ahead%20Chart%20—%20Spotmarkt/spot-gas-chart/B/2008-09-01>

⁵see *pchip* in MATLAB help for details

data, the red ones functions. Note the dark blue and the dark red box: They show the role of the high level data structure (*Experiment* dealing with (P6)) and the high level main function (*workbench.m* dealing with (P1)).

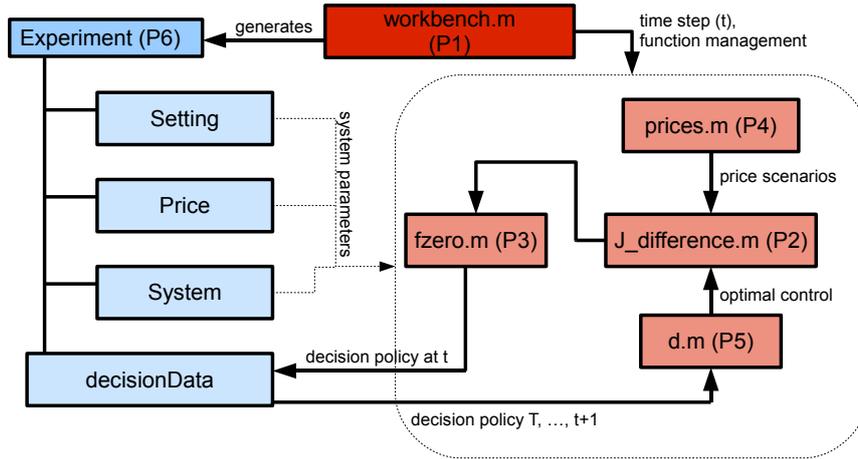


Figure 3.2: Structural properties of the program

3.2.1 Data storage: The variable *Experiment*

The task of the program is to generate the variable *Experiment* of the type *struct*. It contains all information about the simulation conducted; as well parameters for modeling the plant as the data determining the IL. This variable is usually organized in the following way:

- General simulation parameters are contained in the subfield *Setting*:
 - N ... integer number of Monte Carlo simulations
 - T ... length of evaluation period (in hours)
 - method for approximation of IL (linear or cubic spline)
 - search_range ... the interval in which a zero of (3.5) is looked for
 - il_partition ... the sequence of $\hat{p}_t^{E(i)}$ (see equation (3.5))
 - model ... the price model used ('geometric' for the model introduced above, 'prognosis' for the expected values of the 'geometric' model; 'pattern' and 'constant' for deterministic price trajectories with/without hourly pattern)

- `variance_reduction ...` variance reduction technique applied ('anthetic' or 'none')
- Parameters describing the physical plant properties like unit decision lead time, heat rate or startup cost. Values of this kind are filed under *System*
- Parameters for price process like the seasonal pattern or error variances can be found in the subfield *Price*
- The main objects of interest, the solutions of equation (2.33) are stored in *decisionData*. If spline interpolation is applied, an additional subfield *decisionSpline* is generated which contains the already estimated IL (i.e. structures to be evaluated by *ppval.m*).

The data type of *Experiment* makes it possible to include additional information in a very easy way. Therefore more subfields, for example concerning data about program performance can be included.

The advantages of collecting all data concerning a simulation in one variable are obvious: it is vital to store all simulation parameters since they are needed for the interpretation of the results. Besides, all participating functions rely on some of the model parameters. This need can be satisfied in the least confusing way by arranging all parameters in one single data structure and letting the functions get their information from there.

Note that parameters from the subfields *Setting*, *System* and *Price* describe the underlying model assumptions and therefore should not be changed during simulation. The functions *ModPar*, *sysinit* and *priceinit* are built solely for the purpose of generating those fields at the beginning of a simulation.

3.2.2 workbench.m

As shown in figure 3.2, *workbench.m* plays the role of the main function, orchestrating the specialized subfunctions and returning the *Experiment* variable. Technically, *workbench.m* is not a function but a MATLAB worksheet. The idea behind using a worksheet for this task is to supply the user with a basic framework for the calculation of the IL - and thus the power plant value - but leaving room for experiments with the procedure.

As already indicated, *workbench.m* first generates the universal data storage variable *Experiment* and then stores the (approximation of) IL in it. This

Setting up the simulation:

- (1) Generate the *Experiment* variable, using *ModPar*, *sysinit* and *priceinit*
-

Execute the backwards moving computation of IL:

- (2) Determine time step t
- (2.1) Determine states where control can be applied: $x_t \in \Phi$
- (2.1.1) Determine $\hat{p}_t^{E(i)}$ in (3.5)
- (2.1.2) make initial guess for zero in (3.5)
(usually extrapolation of previous values or the IL from $t+24$)
- (2.1.3) solve equation (3.5) using *fzero.m* and *J_difference.m*
- (2.2) store IL - data in *Experiment*
- (2.3) autosave *Experiment*
-

Table 3.1: The tasks of *workbench.m*

is done starting from the end of the time horizon. The main tasks of *workbench.m* are sketched in table 3.1.

3.2.3 J_difference.m

```
il = J_difference(x,t,pEhat,pFhat,modus,Experiment)
```

The function *J_difference* is the main workhorse of the procedure. Its primary task is to evaluate

$$J_t(x_t; p_t^E, p_t^F | u_t = 1) - J_t(x_t; p_t^E, p_t^F | u_t = 0) \quad (3.6)$$

for given x_t , t and observed prices p_t^E, p_t^F (**pEhat, pFhat**). Implicitly, also the optimal production for observed prices (see equation (1.9)) has to be determined. Parameters of the cost and revenue function are taken from *Experiment*.

As in all functions used by *workbench.m*, the current time step is specified by an integer between 1(= start of unit control) , . . . , T (= end of horizon).

The state variable x_t is described by an integer between 1 representing $x_t = t^{on}$ and $t^{on} + t^{cold}$ representing $x_t = -t^{cold}$. This is equivalent to assigning increasing integer numbers to the states, starting from $x_t = t^{on}$. The reason for using this notation in the programs is the simple storage and access of data when the state can be used as index as well. Again, this convention is used by all programs called by `workbench`.

The input parameter `modus` determines the output: the default `modus` 'difference' returns the value of (3.6); 'value' returns

$$\max_{u_t} (J_t(x_t; p_t^E, p_t^F | u_t = 1), J_t(x_t; p_t^E, p_t^F | u_t = 0)) \quad (3.7)$$

which gives an estimate of the actual future payoff at time t and state x_t . For other specifications of `modus`, the outcomes of the different Monte Carlo simulations are returned as a vector together with mean and variance of the distribution of the value of (3.6).

The algorithm used in *J_difference* corresponds to the algorithm for determining $\mathbb{E}_s J_{s+1}$, introduced in section 2.4.3. However, since the purpose of *J_difference* is to compute the difference of two evaluations of $\mathbb{E}_s J_{s+1}$, this algorithm is extended to use each generated price scenario $p^{(i)}$ twice: once for $J_t(x_t; p_t^{E(i)}, p_t^{F(i)} | u_t = 1)$ and once for $J_t(x_t; p_t^{E(i)}, p_t^{F(i)} | u_t = 0)$. If different scenarios were generated for each value of the control variable, this could lead to outliers in future prices unequally distributed between the different u_t . Thus, using the same price realizations for both possible states of control apparently results in better comparability of the Monte Carlo estimates for different controls.

3.2.4 prices.m and pricesVred.m

```
[pE,pF] = prices(pEhat,pFhat,t,model,Experiment)
```

As the name suggests, *prices.m* generates price scenarios for both fuel and electricity prices for Monte Carlo simulation in *J_difference.m*. Necessary inputs are initial prices, the current time step and a variable specifying the underlying model. Beside the geometric price model from equation (3.1), some deterministic models have been included mainly for testing purposes (see also *Experiment.Setting*). Again, constant model parameters such as the price pattern or parameters of the price process are provided by *Experiment*. Note that the returned price scenarios are of length $T - t + 1$: they are price sequences starting with the observed values p_t^E, p_t^F and realizations of the specified price process at $t + 1, \dots, T$.

```
[pE,pF]=pricesVred(pEhat,pFhat,t,model,iteration,Experiment)
```

The function *pricesVred.m* is designed for generating the price scenarios for *J_difference.m* when the variance reduction technique of antithetic variates⁶ is used (to be enabled in *Experiment.Setting*).

The generation of price trajectories is based on the samples of stochastic influences, which are drawn from a $N(0, 1)$ distribution⁷ by the random number generator.

The parameter **iteration** (which is the only difference to *prices.m*) determines the number of the current Monte Carlo simulation for which the price scenario generated will be used. When **iteration** = $N/2$ (i.e. after half of the simulations), the seed of the random number generator will be set to the same seed as in the first Monte Carlo simulation. Additionally, for the subsequent iterations the samples will be multiplied by -1 . This ensures that for each random influence, its 'mirrored value' will be in the sample.

⁶for details, see section 2.3.2

⁷see equation (3.1)

3.2.5 Other functions

d.m

(usage: `u = d(x,t,pE,pF,Experiment)`)

This function returns the optimal control for given state, time and current price observations. This is done by translating⁸ the already stored information about IL into the optimal control strategies needed for *J_difference*.

JE0difference.m

(usage: `je0 = JE0difference(x0,t0, pEhat, pFhat,modus,Experiment)`)

JE0difference.m is the deterministic equivalent of *J_difference.m*. This means that it computes the value of equation (3.6) assuming perfect information. The algorithm uses standard Dynamic Programming as presented in equation (2.5).

The main purpose of this function was to debug the Tseng and Barz algorithm, since it yields the same values as *J_difference.m* for a price model with zero variance. Besides it was also used for computing the optimal control strategy based on a price prognosis from $t = 1$. For this purpose, when `modus` is specified as 'decision', the function returns a states \times time steps matrix, with entries being zero or one. Each value corresponds to the optimal decision at the time and state specified by its position. Plant valuation based on this strategy will be discussed in section 5.1.2.

ModPar.m, sysinit.m and priceinit.m

As mentioned above, those three functions are meant to create the sub-fields of *Experiment* containing the underlying model parameters. As default they contain the same parameters as used by Tseng and Barz to evaluate the plant in section 4.1. For changing settings like the number of Monte Carlo simulations, the approximation method of IL or simulation with different parameter values, it is recommended to edit these functions. Their content is included into *Experiment* automatically at its generation by *workbench.m*.

⁸see section 2.4.2 for the connection of $d()$ and IL

Chapter 4

Deterministic Simulation

4.1 Deterministic simulation: Testing the Program

Including the possibility of validating the results obtained with the Tseng and Barz algorithm has been one of the objectives at the development of the program from the very beginning.

A quite convenient way to judge the program performance is to examine the IL for deterministic prices. Because the IL determine the optimal unit commitment strategy, they represent the pivotal result of the procedure.

Since the evaluation of the determining equation (2.33) has been implemented with two different algorithms (`JE0difference.m` using dynamic programming and `J_difference.m` using Monte Carlo simulations), the results of both functions can be compared and verified when the same price trajectories are used. Besides, for deterministic prices the Monte Carlo estimator is exact after only one simulation, which leads to a short computation time.

Deterministic prices were simulated as trajectories of a geometric price model (3.1) with zero variance. As easily derived from the autoregressive nature of the price logarithms, the resulting trajectories converge against the price pattern. The speed of convergence depends on the correlation α of the logarithmized observations. In the actual simulations, the gas price turns out to stay nearly constant due to α_F being very close to one (see figure 5.5).

4.2 The standard experiment

The physical properties of the simulated plant are similar to the ones used in the empirical study of Tseng and Barz¹. In the subsequent parts, a simulation with these settings will be referred to as *the standard experiment*:

Parameters of the standard experiment

τ :	2	startup time
ν :	2	shutdown time
t^{on} :	10	minimum online time
t^{off} :	10	minimum offline time
t^{cold} :	10	time to cool entirely
q^{min} :	250	minimum generation capacity
q^{max} :	750	maximum generation capacity
b1:	2300	cold start fuel cost for startup
b2:	950	fixed and labor cost for startup
γ :	5	proportion of startup cost to cooling
shut:	1000	unit shutdown cost
a0:	600	constant coefficient of heat rate
a1:	9.1210	linear coefficient of heat rate
a2:	0.0013	quadratic coefficient of heat rate

For the price models, the parameters estimated in 3.1.1 are used if not explicitly stated otherwise.

4.3 Simulation of the standard experiment with deterministic prices

Additional to the physical properties, parameters concerning the simulation procedure have to be specified:

The time period for which the payoff was computed was assumed to be $T = 24$ hours. The approximation of the IL was performed by linear interpolation between solutions of the determining equation (2.33) at $p_t^E = 0, 1, \dots, 170$.

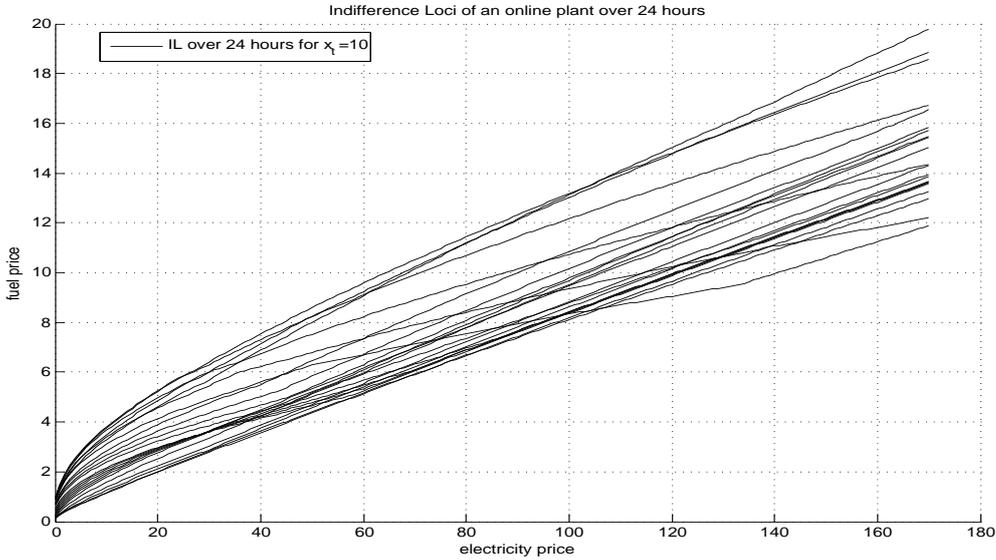


Figure 4.1: Indifference Loci of the standard experiment for deterministic prices: $x_t = 10$

4.3.1 IL of the standard experiment with deterministic prices

Figure 4.1 shows the IL for $x_t = 10$, $t = 1, \dots, 22$, figure 4.2 the IL for $x_t = -10$, $t = 1, \dots, 22$.

The division of the fuel \times electricity price plane corresponds to the zeros of $d_t(x_t; p_t^E, p_t^F)$ at certain time and state. For an observed price pair right below the corresponding IL, the optimal control action is $u = 1$; a price pair at the upper left of the IL implies $u = 0$.

At the IL from $x_t = 10$, a positive y intercept can be observed which may be explained by the nonzero shutdown cost: it can be cheaper to sell electricity with a small financial loss than to pay for the shutdown of the plant. Observation also shows, that for each period the IL of the online unit is above the IL of the offline unit. Intuitively this makes sense since the startup of a plant is connected with additional cost and thus it is of course cheaper to keep the plant online than to start it from an offline state.

¹[1], section 4

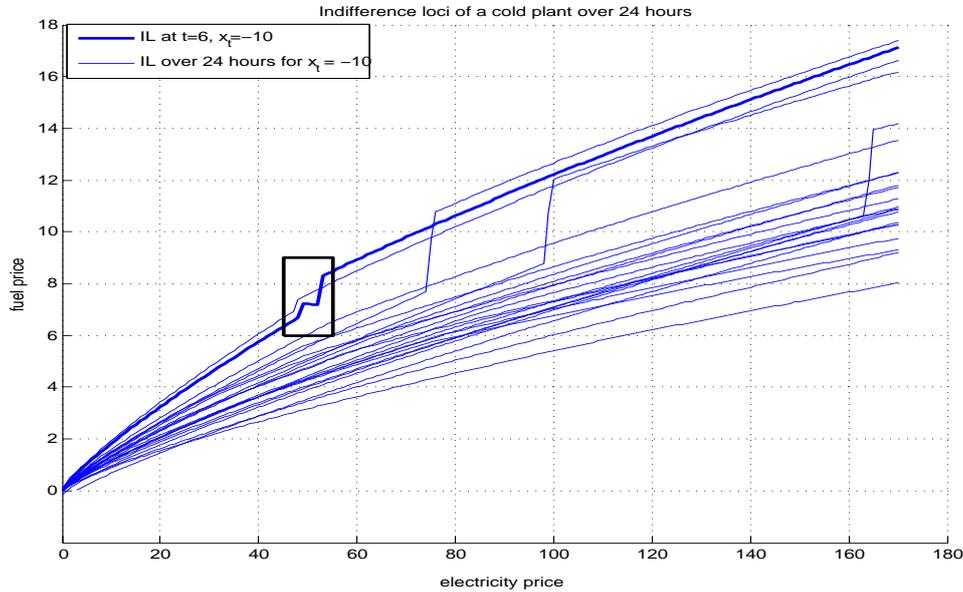


Figure 4.2: Indifference Loci of the standard experiment for deterministic prices: $x_t = -10$

4.3.2 Strange IL: Discontinuities and the problem of describing them as function of p_t^E

While the observations above correspond to the results of Tseng and Barz others seem to contradict them:

The shape of the IL differs significantly from the ones presented in [1], section 4.2. Of course it is not possible to compare results of the deterministic simulations with the stochastic case directly. However it is surprising that some IL are not continuous in electricity prices, since one would expect a marginal increase in fuel prices to have the same effect as a marginal reduction of electricity prices.

Figure 4.3 shows a more detailed view of the difference between the future payoffs at $t = 6$. For this picture, $d_6(x_6 = -10; p_6^E, p_6^F)$ was computed on a grid defined by 200×200 price pairs. The contours resulting from this method (red) give a more accurate description of the underlying IL. Figure 4.3 shows that it is apparently not possible to determine one unique fuel price for every electricity price, at which the operator is indifferent between the control actions. This observation contradicts a fundamental assumption used for the computation of IL:

According to equation (3.5) the IL can be constructed by assigning to every

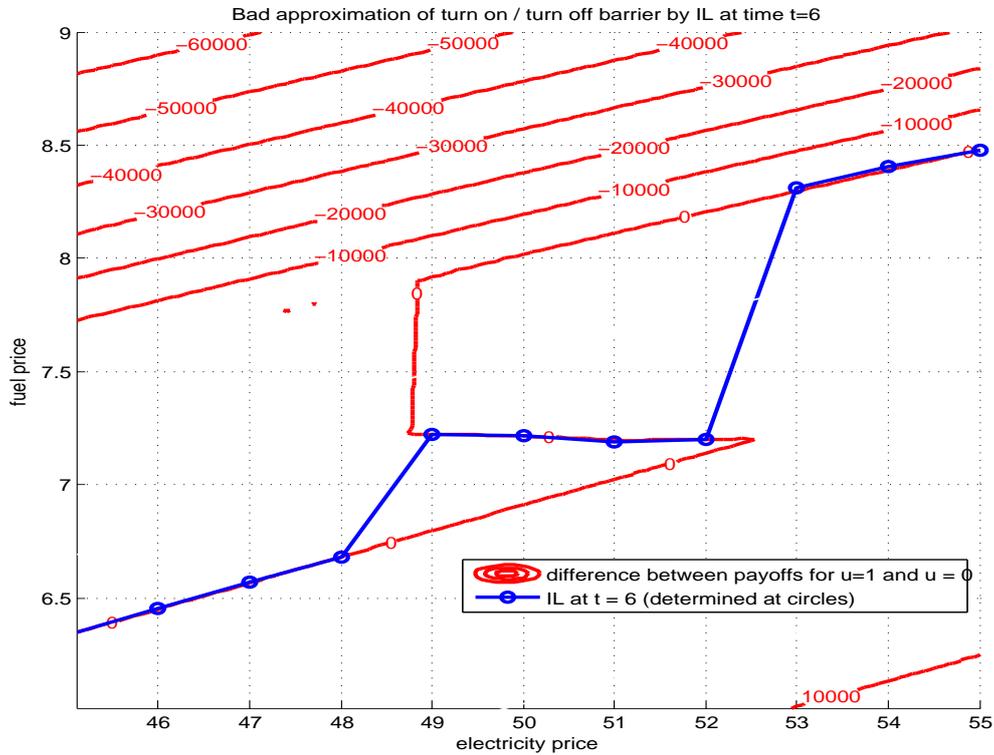


Figure 4.3: Differences between the future payoffs for $u = 1$ and $u = 0$ at $t = 6$ (the price range corresponds to the black frame in Figure 4.2)

electricity price $\hat{p}_t^{E(i)}$ the unique fuel price y^i satisfying $d_t(\hat{x}_t; \hat{p}_t^{E(i)}, y^{(i)}) = 0$.

4.3.3 A showcase explanation for the reaction of the optimal control on a change in fuel prices

The shape of the IL seems to be influenced mainly by the hourly pattern in electricity prices and the minimal unit commitment time. As already discussed in section 1, the minimal unit commitment time determines how long the plant has to stay running after going online (or vice versa for going offline). After this period, there is no restriction on the control anymore.

The influence of the price pattern alone results in a shift of IL upwards or downwards, corresponding to the behavior of the pattern. In other words, the size of the area right below the IL is depending very much on the period

of the day at which the IL applies. It pays for example for many price combinations to turn on the plant in the morning to sell electricity during the price peaks at noon. The opposite effect appears for night hours. The minimal unit commitment time plays an important role in connection with price patterns as illustrated for the optimal decision based on selected price observations below:

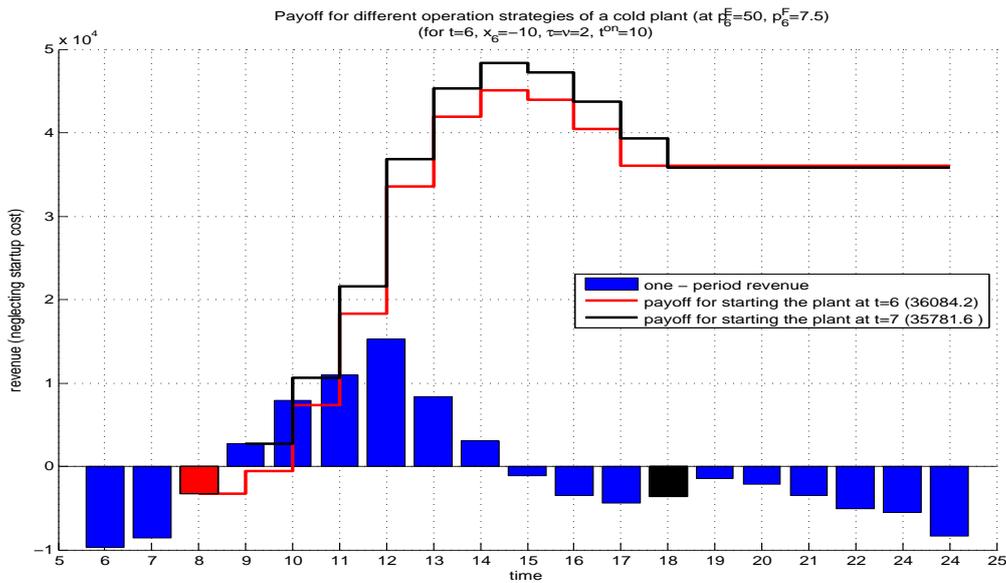


Figure 4.4: Revenue of going online at $t = 6$ and $t = 7$ for $(p_6^E, p_6^F) = (50, 7.5)$

Figures 4.4 to 4.6 explain graphically the generation of the payoff of the plant at $x_6 = -10$ over $t = 6, \dots, 24$. The blue bars represent the payoff² of production at the respective hour. The red step function represents the development of the maximal payoff if the plant is turned on at $t = 6$; the black step function shows the maximal payoff if the plant is turned on at $t = 7$.

The initial prices are $(p_6^E, p_6^F) = (50, 7.5)$ in figure 4.4, $(p_6^E, p_6^F) = (50, 7)$ in figure 4.5 and $(p_6^E, p_6^F) = (50, 6.5)$ in figure 4.6.

In figure 4.3 the optimal decision for all three price pairs can be seen: $(50, 7.5)$ lies in the area where $J_6(\cdot|u = 1) - J_6(\cdot|u = 0) > 0$ and thus it is optimal to start the plant at $t = 6$. For $(50, 7)$ on the other hand, $J_6(\cdot|u = 1) - J_6(\cdot|u = 0) < 0$ which paradoxically suggests not to start the

²payoff = $f_t(x_t; p_t^E, p_t^F)$ for $x_t > 0$

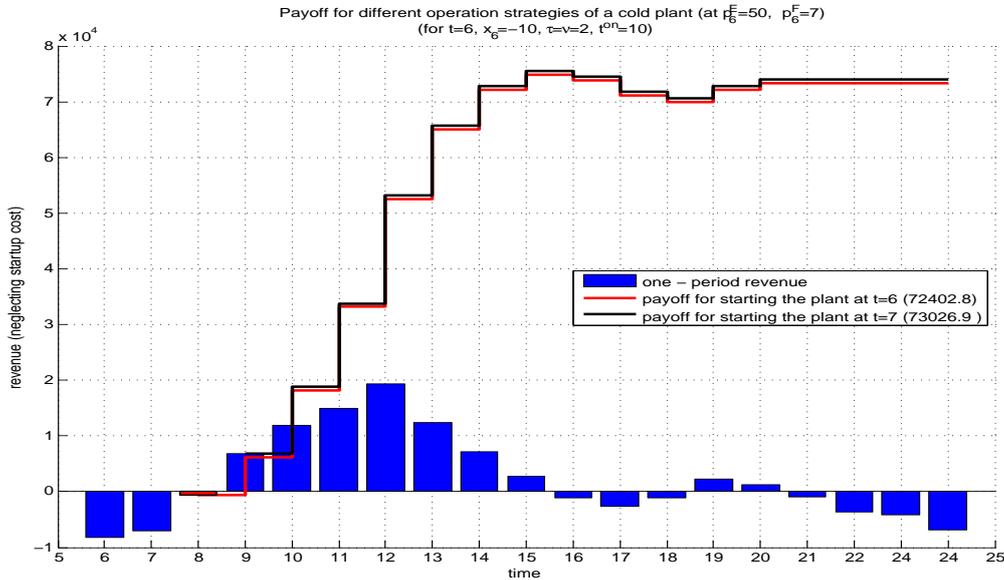


Figure 4.5: Revenue of going online at $t = 6$ and $t = 7$ for $(p_6^E, p_6^F) = (50, 7.0)$

plant at $t = 6$. Finally the optimal decision for $(50, 6.5)$ is again to start the plant.

From graphic treatment of the three cases, the reason for the different optimal strategies can be retraced:

For prices $(50, 7.5)$, the operator will try to profit from high revenues from hour 9 to 14. The minimum unit commitment constraint however demands production over 10 hours.

Thus, starting the plant at $t = 6$ results in production from hour 8 (decision lead time = 2) to hour 17, at which the plant is turned off since there is no more profit to be made. Although the power generation at four hours has generated a negative payoff, this was compensated by the payoff at the other six hours.

Starting the plant at $t = 7$ on the other hand results in avoiding the loss at hour 8 (red bar in figure (4.4)) but producing at loss at hour 18 (black bar in figure (4.4)). Since the black bar exceeds the red bar, it does not pay to start at $t = 7$ compared to $t = 6$.

For prices $(50, 7)$ (figure (4.5)), the situation has changed. Although the production at $t = 8$ still results in a loss smaller than production at $t = 18$, now hours 19 and 20 promise positive payoff for production. Since the profit from hours 19 and 20 is greater than the loss at 18, a plant started at $t = 6$

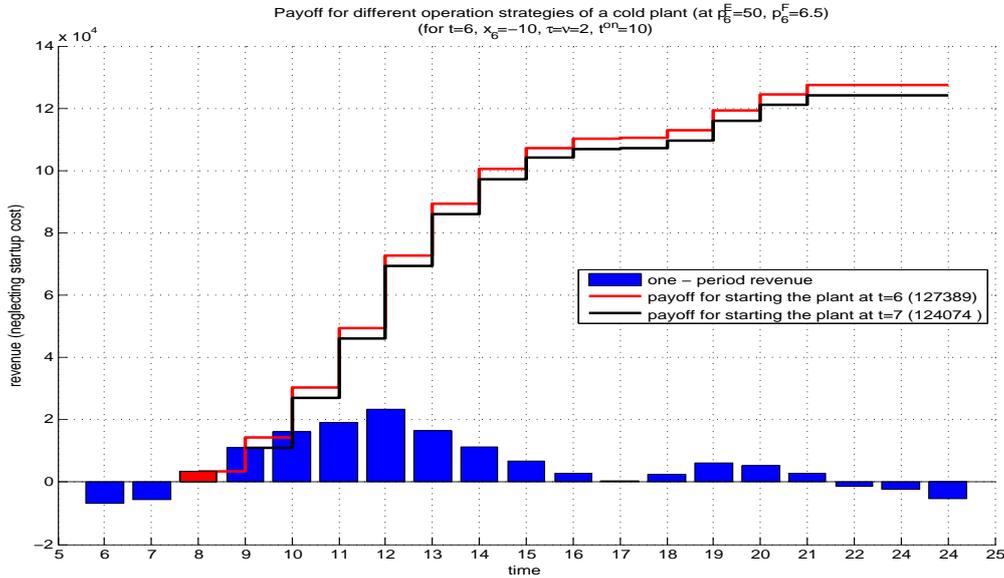


Figure 4.6: Revenue of going online at $t = 6$ and $t = 7$ for $(p_6^E, p_6^F) = (50, 6.5)$

will be shut off at $t = 20$. Therefore, now the optimal strategy is not to start the plant at $t = 6$ (but at $t = 7$).

An additional small change in fuel prices results in the situation depicted in 4.6. At prices $(50, 6.5)$, production at every hour from 8 to 21 results in a positive payoff. The obvious optimal production strategy is to use the opportunity of starting the plant at $t = 6$ and producing until $t = 21$.

This example illustrates some aspects of the effects arising from the implementation of price patterns and unit commitment constraints. Simulations show that this particular behavior vanishes when either price pattern or commitment constraints are excluded.

4.3.4 Conclusion

The results presented in this chapter show that defining the IL as a function of electricity prices leads to a bad approximation of the actual IL by the algorithm used.

This approximation is especially inexact, when the IL is estimated by spline interpolation between few points found to be on the indifference locus. As a function of electricity prices only, the IL are not necessarily continuous; us-

ing spline interpolation however, they are described by a continuous smooth function of p^E .

Note that all results in this chapter were obtained for deterministic price sequences. This makes the direct comparison with results based on stochastic price models inappropriate.

Still, the deterministic results provide sufficient evidence to doubt the general applicability of the IL approximation introduced by Tseng and Barz.

To obtain a generally applicable approximation, the IL would have to be expressed as a function of both fuel and electricity prices. For computing the actual IL at each time step, the difference between the expected payoff for both controls would have to be determined on a grid of price pairs $p_t^E \times p_t^F$. The actual IL would be represented as the contour line of zero difference (as done in figure 4.3). Although this method is simple and provides a better approximation, the extreme computational effort needed creates a new problem. The high computational requirements arise since the grid has to be chosen very narrow to catch phenomena like the ones discussed above.

In the remainder of this work, the approximation of IL will stay unchanged in spite of the observations of this chapter. Also the IL will be approximated by cubic spline interpolation.

This is done simply to avoid an increase of the already uncomfortably high computation time. As a result, suboptimal approximations of the IL are obtained. However, as discussed in section 5.1.2, for stochastic prices still a reasonable estimation of the expected future payoff can be obtained using this suboptimal control.

Chapter 5

Empirical Results

The subsequent part is dedicated to the discussion of empirical results obtained by the implemented algorithm.

In the first three sections of this chapter, the plant dynamics are described by the parameters from the *standard experiment* introduced in chapter 4. Those sections discuss the outcome of the numerical simulation of this model. In the course of this, the results of two different stochastic electricity price models will be compared: The deterministic pattern of one price process is based on the mean price at the specific hour of the day (hourly pattern, determined as in equation (3.2)) as already discussed in section 3.1.1. To determine the impact of this pattern on the payoff, the same model is also simulated for electricity prices with time independent pattern. In this case, the time dependent pattern is replaced by the overall mean of observed prices.

For both electricity price models, the same fuel price process as described in section 3.1.1 is used.

In section 5.4, a plant with different physical characteristics as the *standard experiment* will be simulated. The performance of this unit, designed for production at price peaks (short decision lead time, short unit commitment time and low production efficiency) will be compared to the standard unit.

The optimization horizon of all simulations in this chapter is one week ($T=168$ hours).

5.1 The expected payoff over 168 hours (standard experiment)

5.1.1 Plant value in dependence of observed prices and state

Following the method introduced before, the expected future payoff is a function of the current price observations and the state of the plant. Figures 5.1 and 5.2 depict the main result of the simulation: estimates of the expected payoff of the power plant over 168 hours.

The evaluation period is assumed to start at midnight where the electricity price is usually at a low level (~ 50 €/ MWh). As mentioned before, the fuel prices are independent of the hour of the day; they move around a mean of 7.4 €/ MMBtu.

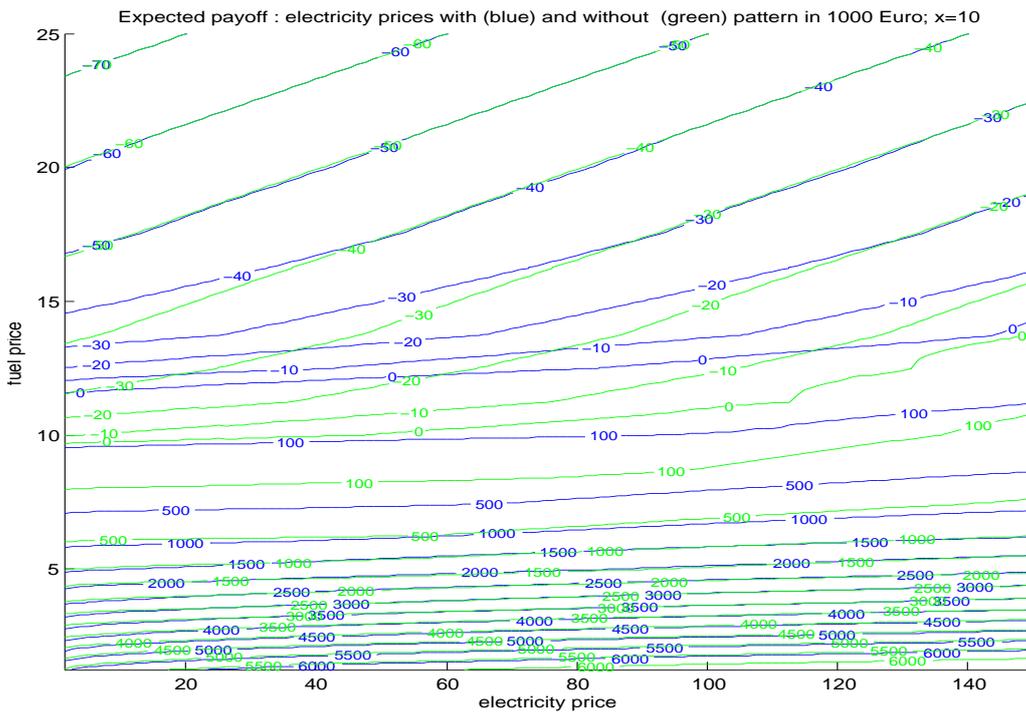


Figure 5.1: Estimates of the expected payoffs for online plant: electricity price with (blue) and without hourly pattern (green)

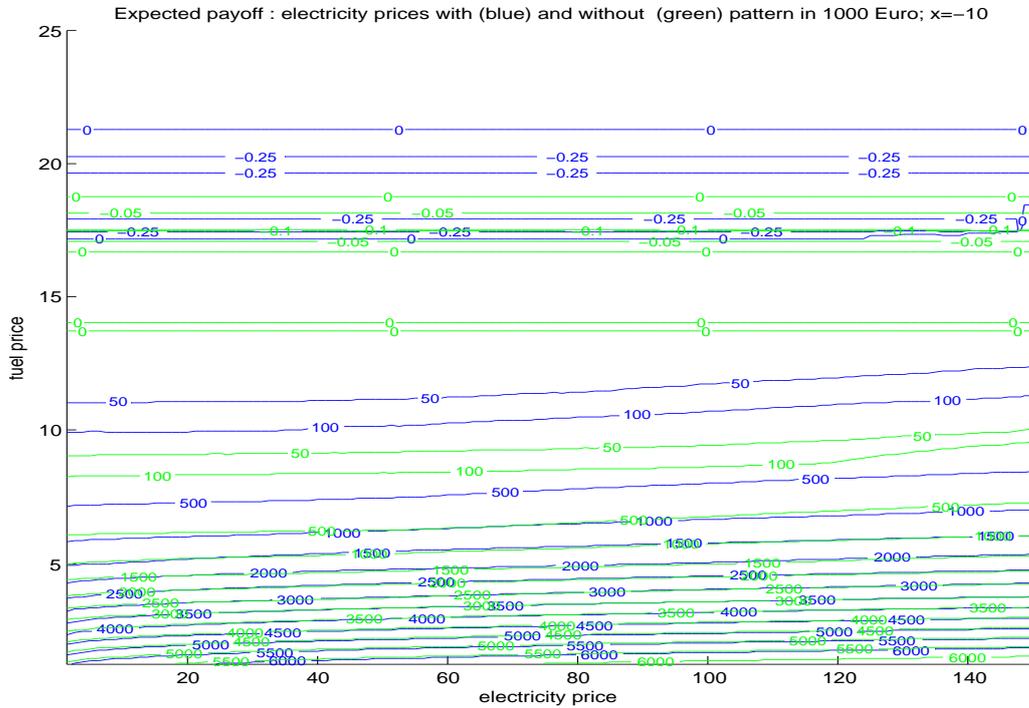


Figure 5.2: Estimates of the expected payoffs for offline plant: electricity price with (blue) and without hourly pattern (green)

The estimated expected payoff for a running plant is shown in figure 5.1; figure 5.2 depicts estimated expected payoffs of an offline plant. In both figures, the payoffs based on simulated electricity prices with hourly patterns are compared with the payoffs resulting from simulations without time dependent pattern in electricity prices. The number of scenarios used for the Monte Carlo estimate was $N=250$. For reasons of comparability, the same random numbers were used for computing the expected payoff at each initial price constellation. The following results can be observed for both price models:

- *Negative expected payoffs for online plant:* If the plant is running at the beginning of the evaluation period, it is quite obvious that negative payoff can be expected for some price constellations. High fuel prices for example can lead to a situation where production does not pay at any instance of the evaluation period. Nevertheless an online plant has to produce at $t=1$ at negative payoff and subsequently to be shut down with additional cost arising.

- *Negative expected payoffs for offline plant:* For an offline plant it is clearly not realistic to expect negative payoffs, since a better value can be achieved simply by not turning on the plant over the whole evaluation period. However, at some price pairs the estimate of the expected payoff turns out to be negative (for both electricity price models).

The reason for this phenomenon is the lack of robustness of the Monte Carlo estimator especially at small payoffs. As one would expect, it can be observed that the vast majority¹ of scenarios generated conditional on fuel prices of about $p_1^F = 20$ leads to zero payoff - corresponding to an offline plant for the whole evaluation period. Naturally the emerging of any outlier influences this value significantly.

- *The payoff increases (piecewise) linearly with growing electricity / nearly exponentially with falling fuel prices (figures 5.3 and 5.4):*

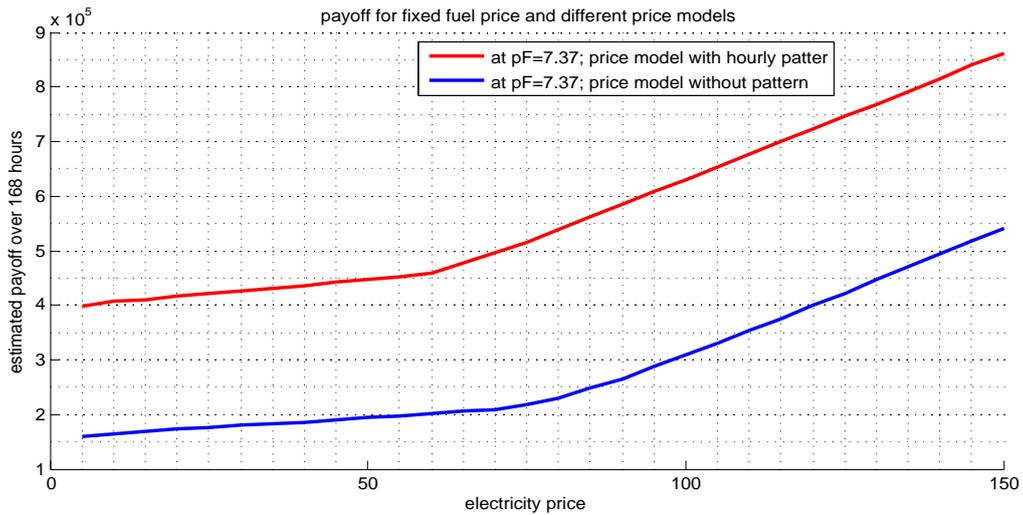


Figure 5.3: Estimated payoffs as function of electricity prices for online plant ($x_t = 10$)

This observation is a result of the fundamental difference between the behavior of electricity and fuel prices: The impact of an electricity price peak fades away much faster than a peak in fuel prices, which is modeled by $\hat{\alpha}_E = 0.9086 < 0.9993 = \hat{\alpha}_F$ ². Figure 5.5 illustrates this effect for

¹simulations at $p_1^E = 60, p_1^F = 20$ resulted in about 0.3% of the scenarios yielding a payoff different from 0

²see equation (3.1)

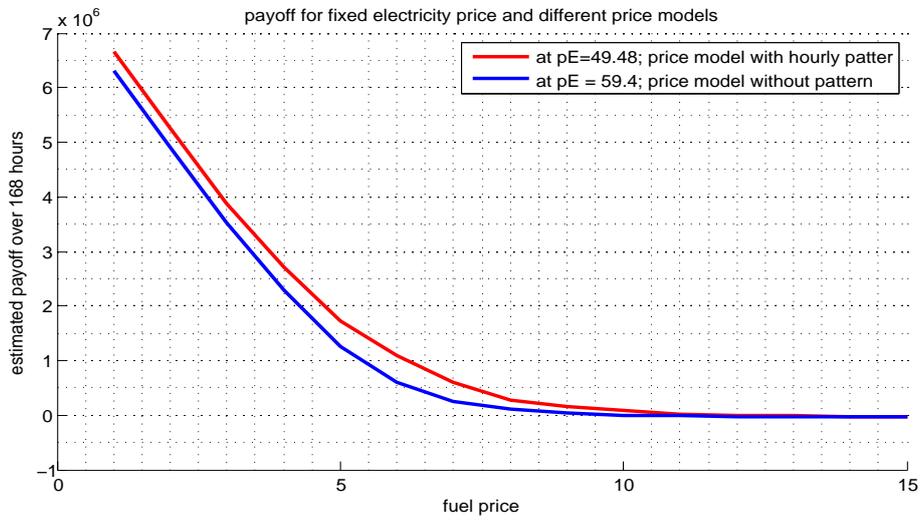


Figure 5.4: Estimated payoffs as function of fuel prices for online plant ($x_t = 10$)

the deterministic part of the price (the hourly pattern model without stochastic influence which equals the 'deterministic' price model from section 4).

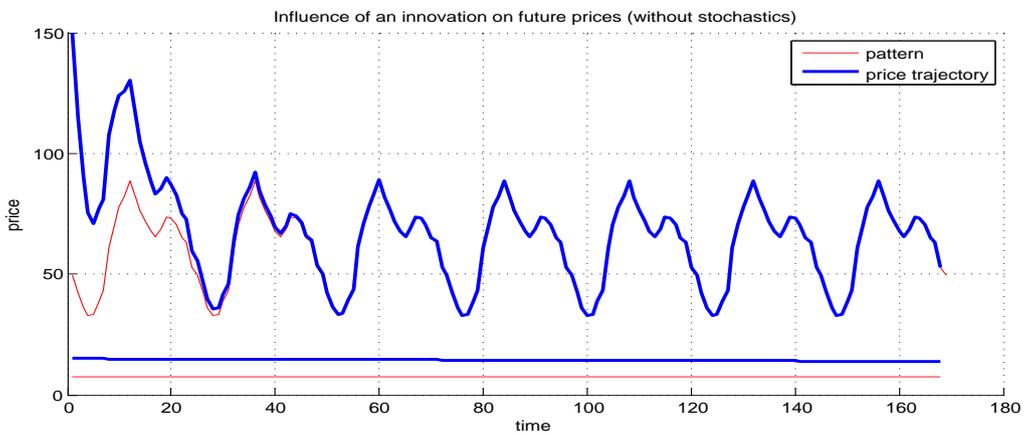


Figure 5.5: Effects of innovations on the deterministic part of price dynamics

As a result, outliers among fuel prices can affect future fuel prices for the whole evaluation period; the influence of electricity price peaks is limited to about a day. Therefore the observation of low fuel prices leads to a general tendency towards production over the whole time

horizon which results in higher additional payoff from a 'beneficial' fuel price than from a favorable electricity price.

Note that a similar effect can be observed for the simulation without hourly patterns for electricity prices.

The comparison of expected revenues in figures 5.1 and 5.2 shows a difference in expected profits for different price models: At the same price constellations the estimated payoff is higher when prices have an hourly pattern.

Note that it may be misleading to compare expected revenues for both price models directly, since initial (observed) electricity prices have different meanings according to each model. At the beginning of the evaluation period, the mean electricity price level for an hourly pattern would be at $p_1^E = 49.48$ compared to $p_1^E = 59.4$ for a model without time depending pattern. Therefore, the same observation can mean a positive innovation and thus a tendency to higher prices in terms of the one model and the opposite effect in terms of the other.

Figure 5.4 shows the estimated payoffs for both price models as function of fuel prices. The electricity price level is fixed to the mean price from the corresponding price model. For each price model these estimates describe the payoff at a situation where the current electricity price observation implies no innovation for the price processes. At least in this special case, the direct comparison of the payoffs from both models is possible. Obviously a positive expected payoff is larger for the price model with hourly pattern. However the scale of figure 5.4 hides the reverse effect ³ for the potential negative payoffs of an online plant: The losses using prices with pattern exceed the losses for prices without pattern.

5.1.2 Unit commitment based on different degrees of information

Describing the actions of the operator plays a decisive role for the estimated value of the plant. A key problem is to model the imperfect information of the operator: Assuming unit commitment based on either perfect price information or only on the expected price development at the beginning of the evaluation period ($t=1$) would be oversimplifications. However, those simpler models of the operator's decisions should give an upper (perfect information) or a lower (forecast prices at $t=1$) boundary of the future payoff.

³this effect can be detected easier in figure 5.1

Thus, to get an idea about the performance of the decision policy used, the empirical results from the previous section are put in context with the results of those simplifications.

For each of the control strategies, the value of the plant is given by the mean future payoff over a number of $N = 250$ of price scenarios using the respective control.

The results of the three methods for an online plant⁴ are depicted in figures 5.6 and 5.7. Figure 5.6 shows the direct comparison of evaluation based on the price prognosis and the Tseng and Barz algorithm; figure 5.7 depicts the comparison of the values using optimal control and the Tseng and Barz algorithm.

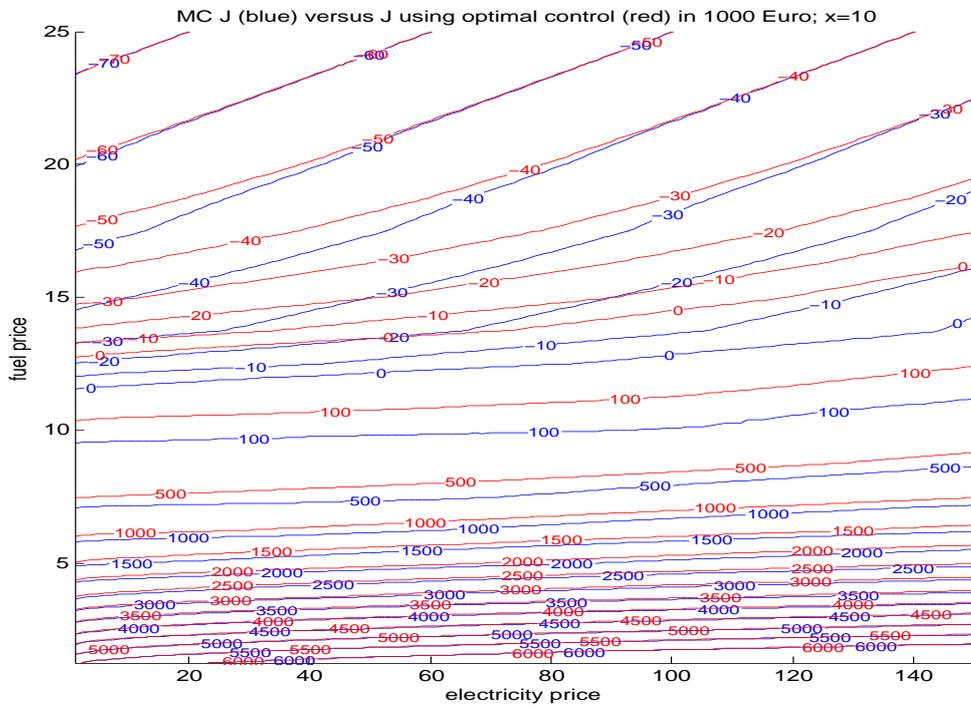


Figure 5.6: Expected payoffs: Tseng and Barz algorithm (blue) versus optimal control based on perfect information (red)

Concerning the performance of the plant, both graphics provide a satisfying result: the Tseng and Barz algorithm never predicts a greater plant value than the optimal control; the Tseng and Barz algorithm always estimates a higher future payoff than the control based on forecast prices. Also,

⁴The performance turns out to be similar for an offline plant.

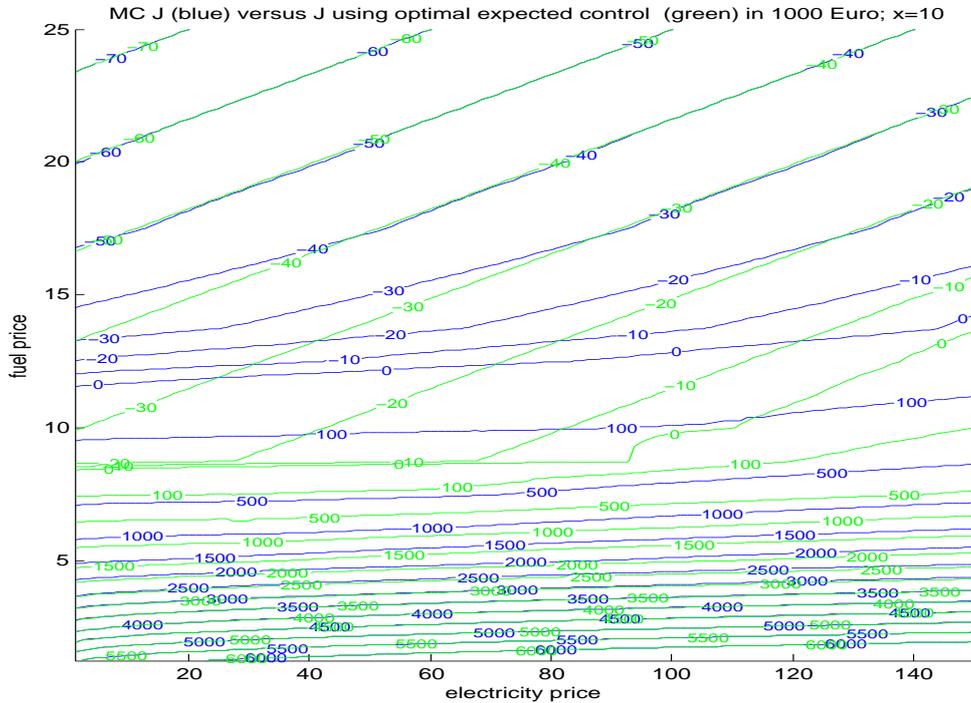


Figure 5.7: Expected payoffs: Tseng and Barz algorithm (blue) versus optimal control based on price forecasts (green)

the comparison between the strategies shows that the choice of the control strategy does not really make a difference at extreme fuel price levels, where the payoff-contours converge. Differences of significance emerge at fuel prices from about 5 to 15. Note that those prices are of especial importance since 'reasonable' observations are to be expected in this area.

5.2 Frequency counts of scenario payoffs (standard experiment)

Beside the expected payoff, the procedure of Monte Carlo estimation also provides the payoffs of single price scenarios. In this section, the relative frequency of the revenues and their differences will be discussed at the exemplary case of $(p_1^E, p_1^F) = (49.48, 7.37)$, which is the mean price observation at $t=1$ ⁵.

⁵for electricity prices with hourly pattern

5.2.1 The distribution of $J_1(\cdot)$

The following summary statistics describe the distribution of payoffs for $N = 100000$ price scenarios:

	$J(x_1, p_1^E, p_1^F u_1 = 1)$			$J(x_1, p_1^E, p_1^F u_1 = 0)$		
	Mean	Std. Dev	Skew.	Mean	Std. Dev	Skew.
$x_1 = 10$	455700	457990	1.2658	467340	454960	1.2815
$x_1 = -10$	453500	467820	1.1880	518700	466070	1.2003
$x_1^* = 10$	455110	458070	1.2385	466860	455030	1.2518
$x_1^* = -10$	454010	470410	1.2099	519350	468800	1.2221

The columns marked with x_1^* concern the distribution of payoffs simulated using antithetic variates. The positive skewness of all distributions seems reasonable, since extreme financial losses can be avoided by the operator.

It turns out, that the method of antithetic variates affects the distribution of the payoffs only insignificantly. The reason is probably the asymmetric effect of outliers on payoffs: the payoff is increased by high electricity and/or low fuel prices arbitrarily, whereas low electricity and/or high fuel prices lead to a turnoff of the plant and thus to a bounded negative payoff. As a result, generating symmetric price observations by extending the set of samples by the mirrored⁶ values does not decrease the variance of payoffs: The symmetry in the distribution of prices is not passed on to the distribution of payoffs.

Figure 5.8 depicts the histograms of payoffs for different control actions and initial states.

The shape of all histograms reflects the impact of the turnoff of the plant when future payoffs are expected to be negative. As already observed at the summary statistics, the distribution of expected payoffs is right skewed for all initial states and controls u_1 .

A more surprising observation are peaks in expected future payoffs, especially in the top three histograms in figure 5.8.

In particular for $u_1 = 0$ (first and third from the top), the peaks occur at one single value and not at an interval. Those observations resemble the scenarios, where the power plant is turned off and stays offline for the rest of the evaluation period. As a consequence, the payoff only consists

⁶section 2.3.2

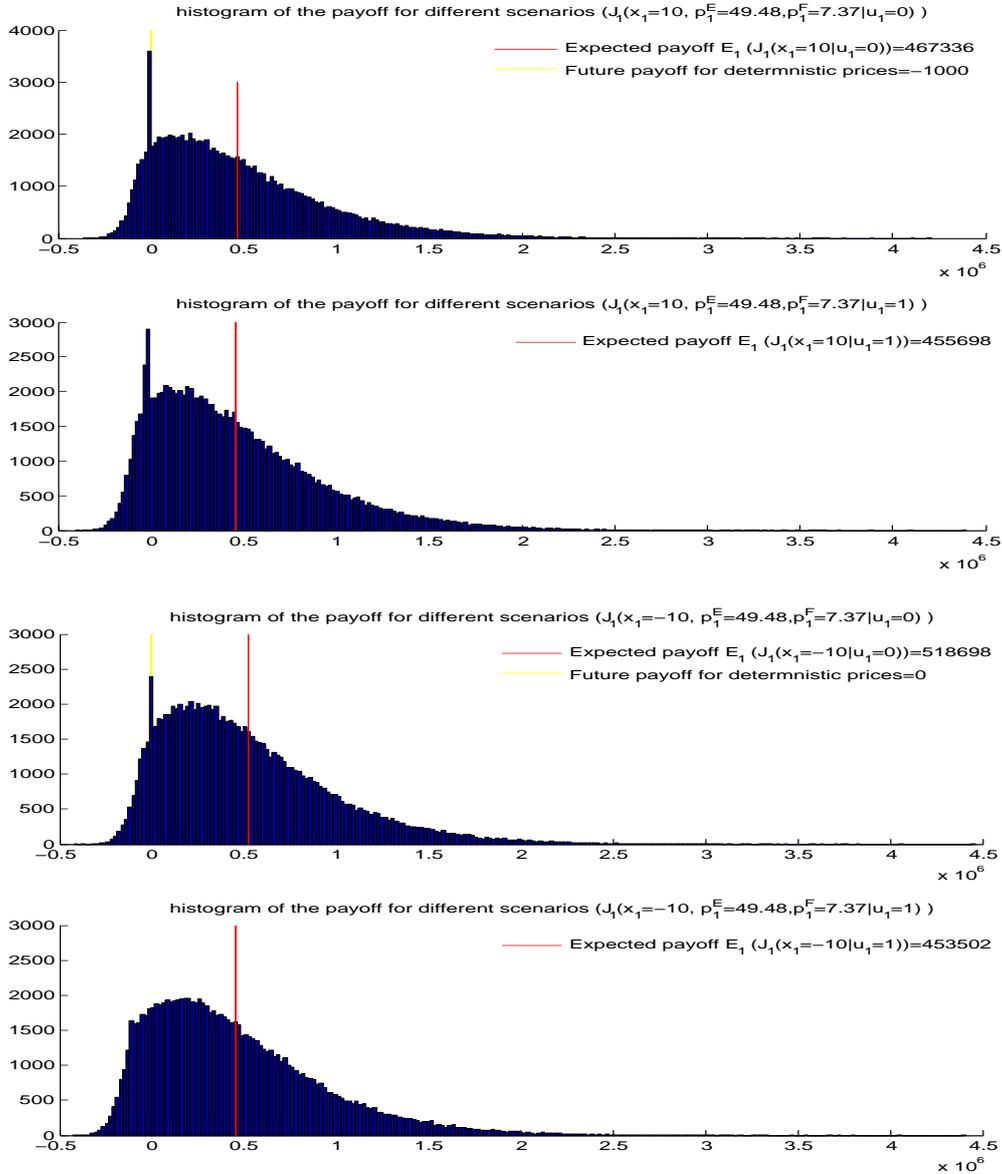


Figure 5.8: Payoffs at $x_1 = 10$ (first and second from the top) and $x_1 = -10$ (below), $(p_1^E, p_1^F) = (49.48, 7.37)$; $N=100000$

of the production at observed prices from the first period which yields a deterministic result.

The situation is slightly different for $u_1 = 1$: For an online plant, this strategy implies production in the following period ($t=2$) at unknown prices. Turning off the plant at $t=2$ and keeping it offline for the rest of the optimization

period however, yields a payoff affected by the price stochastics of only one time period. Thus in this case, the scenarios where $u_1 = 1$, $u_t = 0$, $t = 2, \dots, T$ are only subject to stochastic prices at one time period which causes a cluster of payoffs in a small interval.

The comparatively smooth histogram for payoffs at $x_1 = -10$ and $u_1 = 1$ is a consequence of the unit commitment time: If a cold unit is turned on, it has to stay online for the subsequent 10 hours. As a consequence, every payoff scenario with this initial setting is influenced by price stochastics significantly. This prevents the strong clustering of payoffs observed above.

5.2.2 The distribution of $d_1(\cdot)$

The summary statistics of $d_1(\cdot) = J_1(\cdot|u = 1) - J_1(\cdot|u = 0)$ for $(p_1^E, p_1^F) = (49.48, 7.37)$, $N=100000$ assuming electricity prices with hourly pattern (the same setting as above) are listed below:

	$d_1(x_1; 49.48, 7.37)$		
state	Mean	Std. Dev	Skew.
$x_1 = 10$	-11638	59275	1.9824
$x_1 = -10$	-65196	41457	-0.0309
$x_1^* = 10$	-11756	59641	1.9181
$x_1^* = -10$	-65348	41256	-0.1180

Naturally the standard deviations have been decreased by the generation of the differences of expected payoffs. Note that in addition to this numerical effect, the usage of the same random numbers for both, $J_1(\cdot|u = 1)$ and $J_1(\cdot|u = 0)$ also contributes to decreasing the variance of the differences⁷. Figure 5.9 consists of the histograms of the differences. The upper two graphics show $d_1(x_1 = 10; 49.48, 7.37)$ with different ranges, the lower depicts $d_1(x_1 = -10; 49.48, 7.37)$.

Those histograms are shaped by the same effects as above: The additional loss of keeping the online plant running for one more hour is located in the area of the extreme peak of the histogram⁸. This indicates, that in many cases it is optimal to shut down the plant as soon as possible; thus the difference between the expected payoffs for $u_1 = 1$ and $u_1 = 0$ consists of the loss at $t=2$ to a great extent.

The histogram at the very bottom, showing the situation for the offline plant, lacks a similar cluster of payoff differences. Again as in the previous section,

⁷see Section 3.2.3

⁸the expected payoff of operating at $t=2$ is -11197

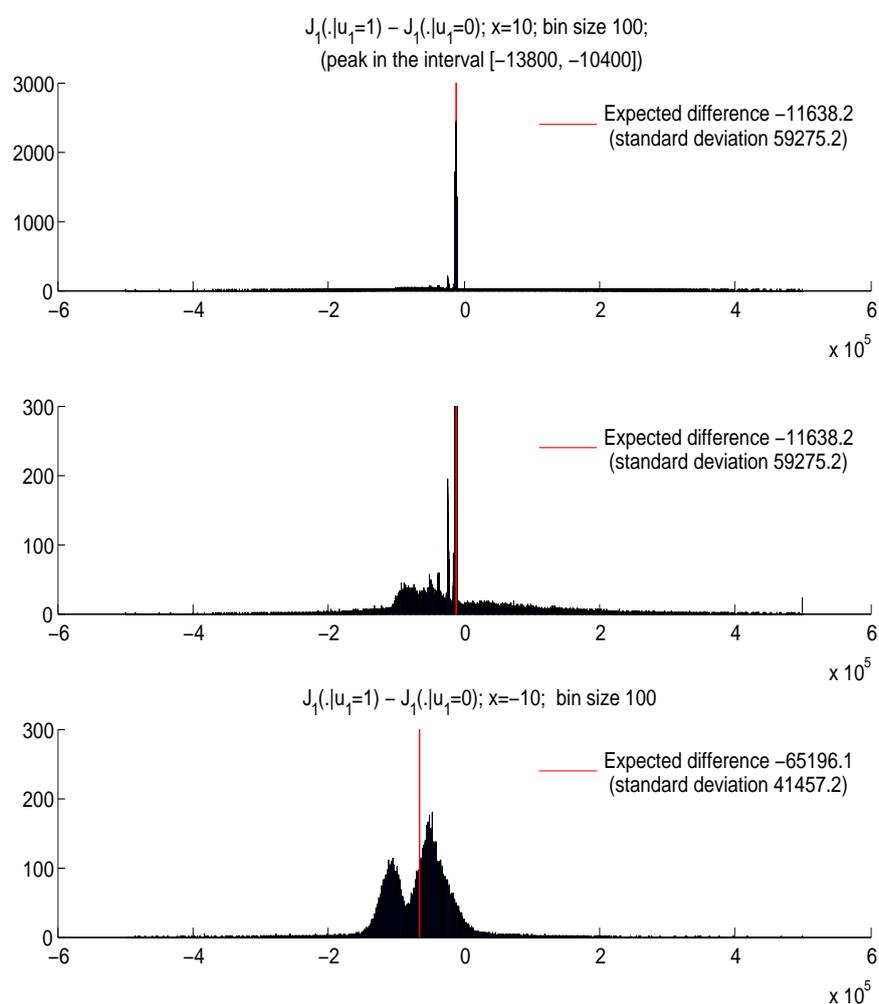


Figure 5.9: Histograms of $d_1(\cdot; 49.48, 7.37)$

this is due to the unit commitment constraint, which makes the payoff for $u_1 = 1$ subject to stochastic influences.

5.3 IL and computation time (standard experiment)

5.3.1 Indifference Loci

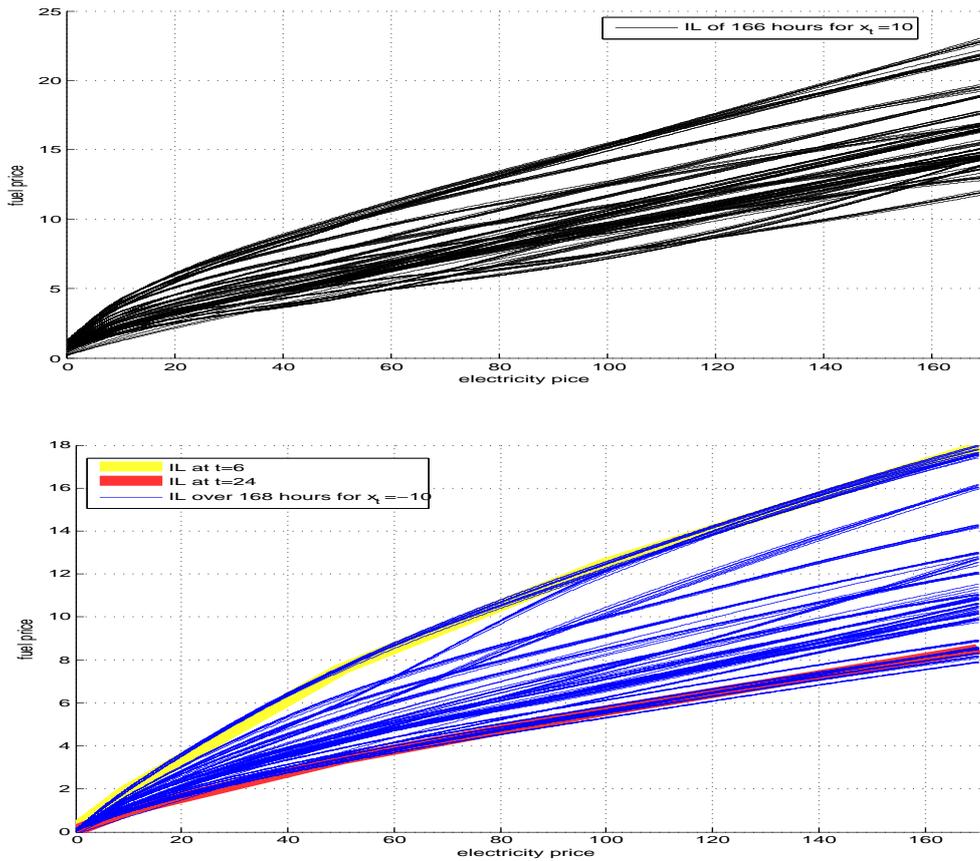


Figure 5.10: IL of 166 hours of the standard experiment for stochastic prices: $x_t = 10$ (top) and $x_t = -10$ below

Figure 5.10 shows the IL for an online and an offline plant. Due to the decision lead times ($\tau = \nu = 2$), there are 166 IL for the optimization period.

Note that the IL have been determined following section 3.1.2. For fixed electricity prices $p_t^E = 1, 10, 50, 100, 170$ ⁹, a fuel price p_t^F was com-

⁹the value for $p_t^E = 0$ is an extrapolation from the approximated IL to guarantee stability of the zero search, since sometimes there is no corresponding positive solution

puted for which (p_t^E, p_t^F) is on the IL. The actual IL were approximated by cubic spline interpolation between those price pairs. This explains the smooth shape of the IL for the stochastic case.

It is not uncommon that electricity prices above 170 appear among simulated prices¹⁰. For this kind of prices, the IL is approximated by spline extrapolation. This method provides nearly linear extensions of the IL for electricity prices up to more than 200 €; however is not appropriate to use the extrapolation of IL at electricity prices above 250 €.

As already observed in the deterministic case, the IL for an online plant at given time are always above the IL of the offline plant. Also a tendency towards positive intercepts for IL at $x_t = 10$ and the opposite for $x_t = -10$ can be found again. This phenomenon also appears for prices without hourly pattern; similar to the deterministic case its cause is likely to be found in the nonzero startup / shutdown cost.

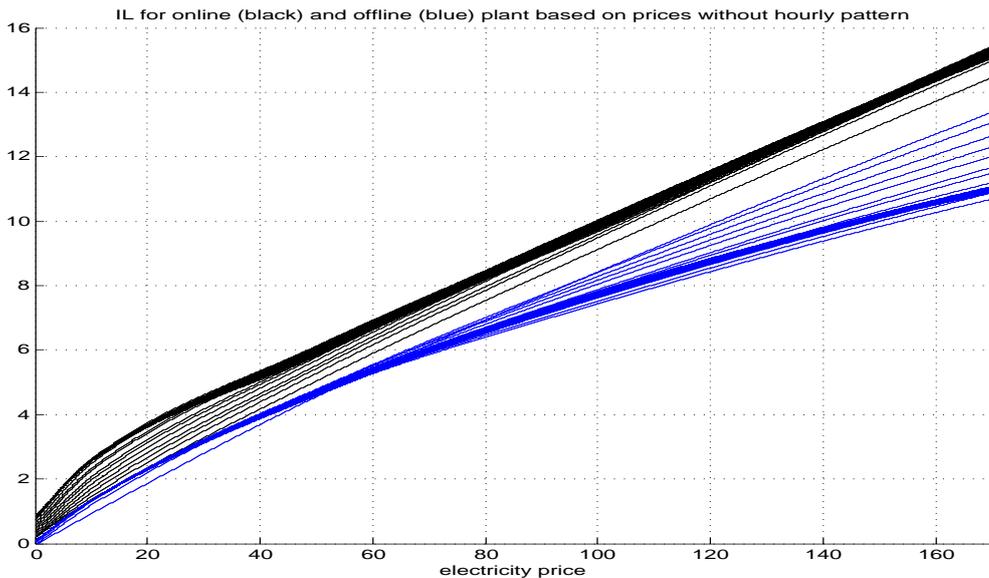


Figure 5.11: IL of 166 hours for price processes without hourly pattern

The influence of the price pattern on the optimal decision strategy can be assessed by comparison of figure 5.10 and 5.11.

For prices with pattern, the IL converge depending on their respective hour of the day. This can be seen at the clusters of IL in figures 5.12 and 5.12,

¹⁰the maximal observation amongst historic data used for estimation of the model parameters was a price peak of 248.27 €/MWh

which illustrate sequences of IL at corresponding hours of the day for an online and an offline plant.

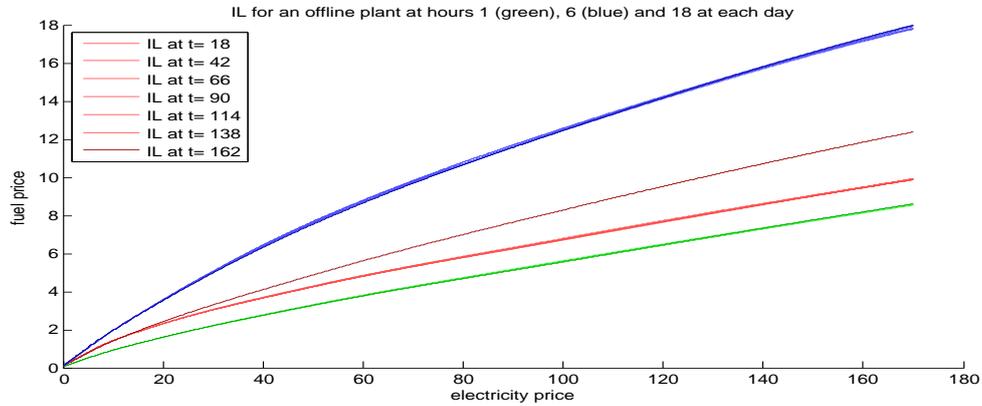


Figure 5.12: IL (offline plant) at $t = 1/6/18$ for each day

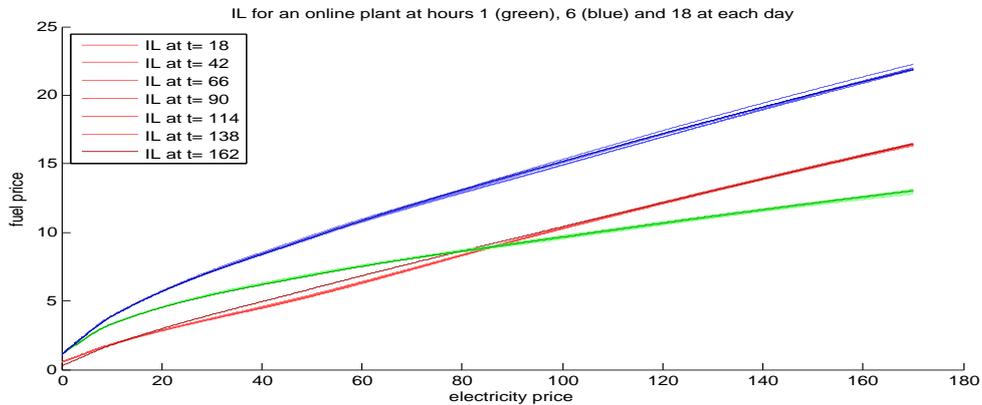


Figure 5.13: IL (online plant) at $t = 1/6/18$ for each day

In the figures above, the IL closest to the end of the optimization period T are colored darker than the IL at different hours of the day. Apparently the IL converge with increasing (time) distance from the end of the optimization horizon. Note that this phenomenon can be observed for each hour of the day and both offline and online plant, which is also reflected in the computation time (figure 5.15). The example of the IL at $t = 6$ and $t = 24$ gives an intuitive reason for this behavior: each day, the electricity price is assumed to have a peak at noon and a low during the night. As a result, the expected

future electricity price - and thus the expected revenue from power generation - is higher before peaks in the hourly pattern.

For prices without hourly patterns, converging IL can also be observed (figure 5.11). The absence of the hourly pattern however leads to the convergence against one optimal decision rule.

5.3.2 Computation time

The computation time for the simulation of the standard experiment with electricity prices using hourly patterns over a horizon of 168 hours is given in figure 5.14.

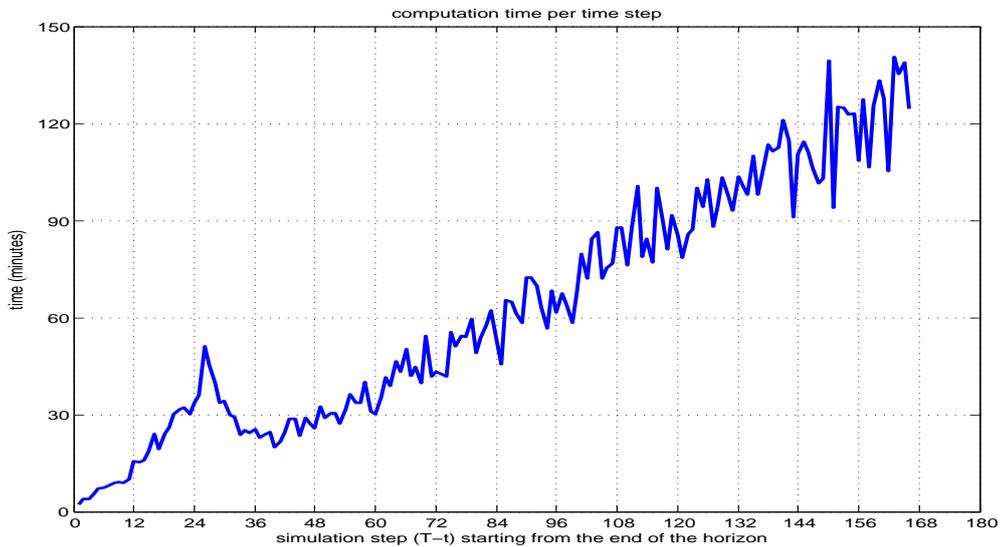


Figure 5.14: Time needed to compute the IL of one step

The most obvious observation is the mere amount of the time needed for simulation¹¹. As a rule of thumb, the evaluation over a period of 168 hours takes about one week.

The second observation of interest is the decrease of computation time needed per step after 24 hours. The reason for this is to be found in a change in the initial guess of the zero of equation (3.5): In the first 24 hours, an extrapolation of already determined points on the IL was used; later, the result from the corresponding time of the day before was taken as initial guess of the IL.

¹¹The computations were performed using a Pentium 4 2.6 GHz with 512 MB RAM

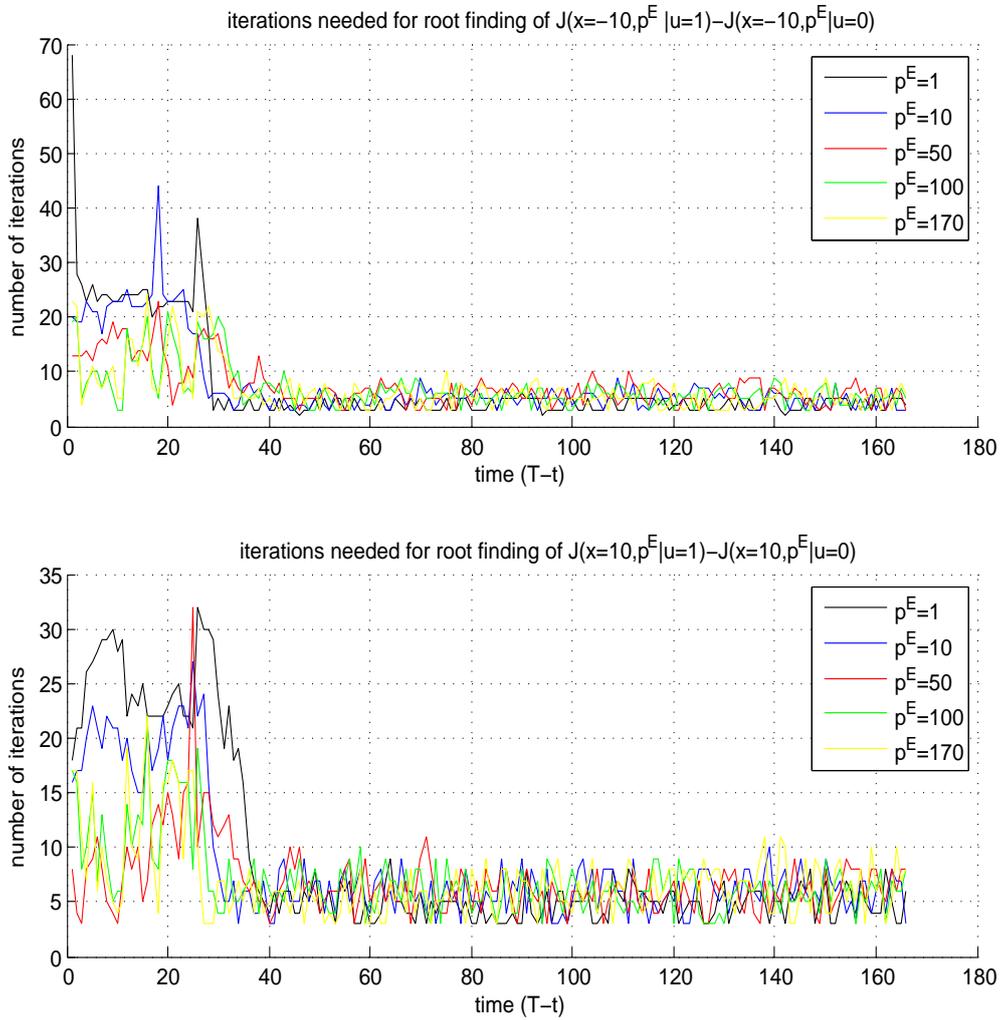


Figure 5.15: Evaluations needed for root finding of equation (3.5)

Figure 5.15 displays the number of evaluations necessary to find the zero of equation (3.5). Indirectly, this gives an impression of the speed of convergence of the IL: Since the initial guess of the zero search procedure is the corresponding zero from 24 hours earlier, the number of iterations corresponds to the distance of previous IL to the one currently determined. Figure 5.15 shows, that the number of iterations needed for root finding decreases below 10 in less than 40 time steps and stays at this level for the remaining periods. This behavior reflects the clustering of IL. The fact that a certain number of iterations needed at every time step is due to clustered but not identical IL.

5.4 Changes in the system dynamics

The simulations conducted up to this point used different price models to assess the validity of the algorithm (chapter 4) and to illustrate basic effects (section 5.1.1). This section will compare simulations for different system dynamics.

The parameters of the standard experiment (see section 4.2) correspond to a thermal power plant, where electricity generation from steam plays at least a certain role¹². This can be seen in particular at the variable startup cost, reflecting the heat in a boiler which cools down with increasing time since shutdown. Of course other parameters (heat rate and startup / shutdown time) also are supposed to differ for other plant types.

Note that the plant simulated in this section does not rely on parameter values derived from a real plant. It represents a more or less realistic fictional scenario by which the capabilities of the program should be demonstrated.

5.4.1 Short decision lead time, short unit commitment time

A prototypic example of a plant type with these characteristics is the *gas engine*. In a unit like this, a gas turbine (similar to a large jet turbine) powering a generator, produces low amounts (< 100 MWh) of electricity at a comparatively low efficiency. This kind of plant has a very short startup (a matter of minutes) and virtually no unit commitment time; thus it is used for generating power at peak loads.

Modeling the short startup and unit commitment time is achieved in two ways: The variables for unit commitment and startup time can be set to one (which is the lowest possible value for the simulation). The simulation of real time production is done by re-modeling the minimum production capacity to zero. This means that it is now possible to generate no electricity at an online plant. Note that also the 'production of no electricity' may generate expenses for an online plant, if the heat rate has a positive intercept (as in the standard experiment). This can be interpreted as the cost of keeping the generator on standby.

Figure 5.16 compares the heat rates of the standard experiment and of the gas engine. The heat rate of the gas engine is chosen to have a steeper

¹²Either steam is generated directly by burning natural gas (thermal power plant) or a gas turbine generates electricity and the turbine's emissions are used for a steam plant

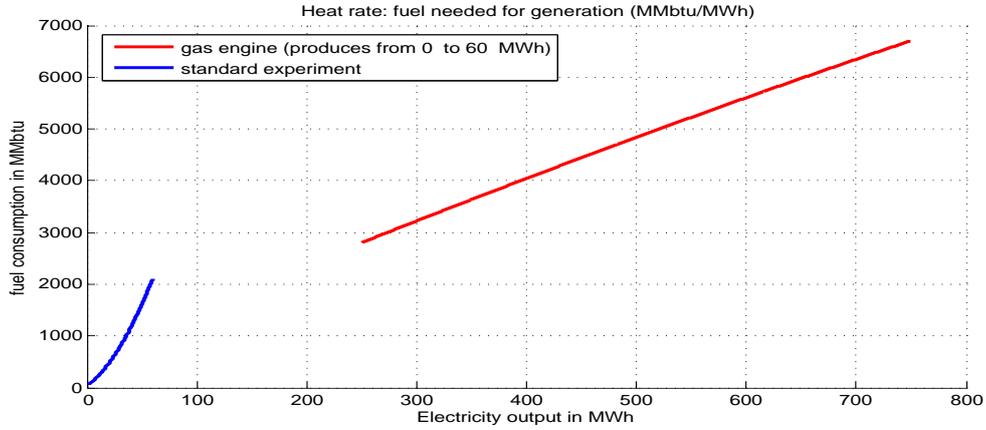


Figure 5.16: Heat Rates: Standard Experiment vs Gas Engine

ascent to model the lower efficiency of electricity generation.

The detailed parameters of the gas engine model are as follows:

Parameters for gas engine plant

τ :	1	startup time
ν :	1	shutdown time
t^{on} :	1	minimum online time
t^{off} :	1	minimum offline time
t^{cold} :	1	time to cool entirely
q^{min} :	0	minimum generation capacity
q^{max} :	60	maximum generation capacity
b1:	0	no cold start fuel cost
b2:	500	fixed and labor cost for startup
shut:	500	unit shutdown cost
a0:	30	intercept of heat rate: standby cost
a1:	10	linear coefficient of heat rate
a2:	0.31	quadratic coefficient of heat rate

Beside the model parameters, the simulation of the *gas engine* provided another novelty: The convergence of IL was utilized to save computation time.

When the computation of point p_t on an IL at time t ($t < T - 48$)¹³ took less than 10 iterations, the subsequent IL at $t - 24k$ ¹⁴ was assumed to contain p_t as well. This approach led to a total computation time of 'only' 59 hours.

¹³this ensures that the corresponding previous IL has been taken as initial guess

¹⁴note that the computation of IL proceeds backwards in time

5.4.2 The comparison of *gas engine* and *standard experiment*

Naturally the expected payoff of this simulation cannot be compared to the results of the *standard experiment* directly, because of the different capacities of the units.

One way of scaling the estimated payoff of the plants to make them comparable¹⁵ is the determination of the payoff per output (€/ MWh) as shown in figure 5.17:

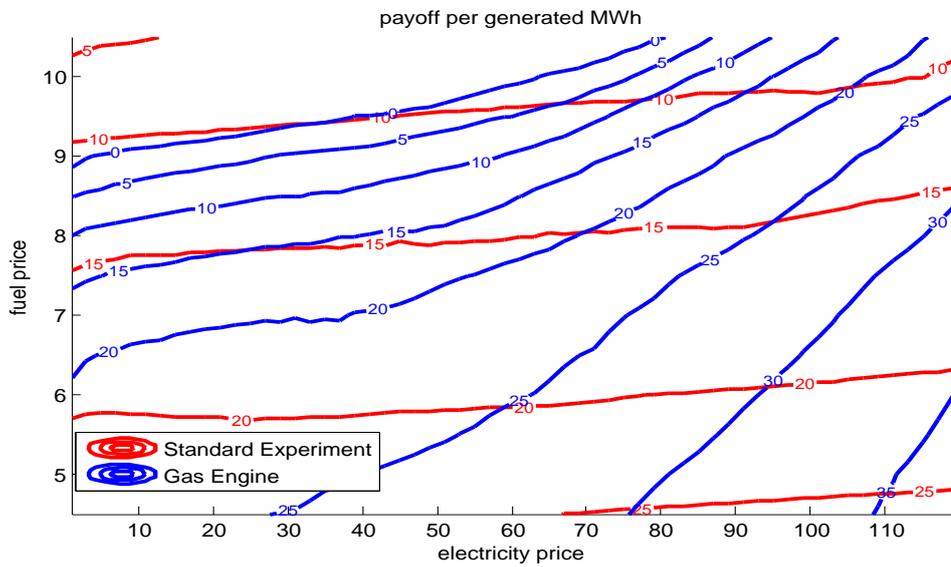


Figure 5.17: Payoff per MWh (N=250) for $x_1 = 1, t = 1$

The payoff for generated electricity at the *gas engine* shows a big dependency on the electricity price level. This reflects the concept of power generation at price peaks only: a high electricity price can be utilized immediately. Naturally the plant resembling the *standard experiment* shows a different behavior. Similar as the overall payoffs, the payoffs per output are mainly subject to the level of the fuel price (see section 5.1.1 for details).

The fundamental difference in the operation of the two plant types is illustrated in figure 5.18 which compares the amount¹⁶ of the plant's production capacity used over the evaluation period (average utilization). As expected, the *gas engine* only uses a small part of its total capacity.

¹⁵the evaluations at each price pair and each unit rely on the same price scenarios

¹⁶mean output per overall capacity(= $168 \cdot q^{\max}$)

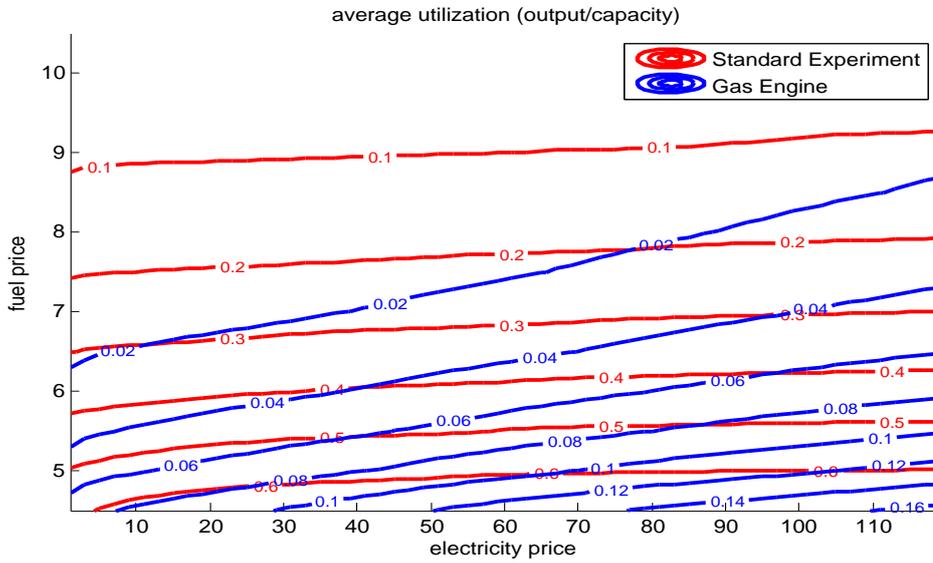


Figure 5.18: Output per capacity ($N=250$) for $x_1 = 1, t = 1$

Note that also for the *standard experiment*, the average utilization does not reach 70% for fuel prices above 4.5 €/per MMBtu. Figure 5.19 uses the

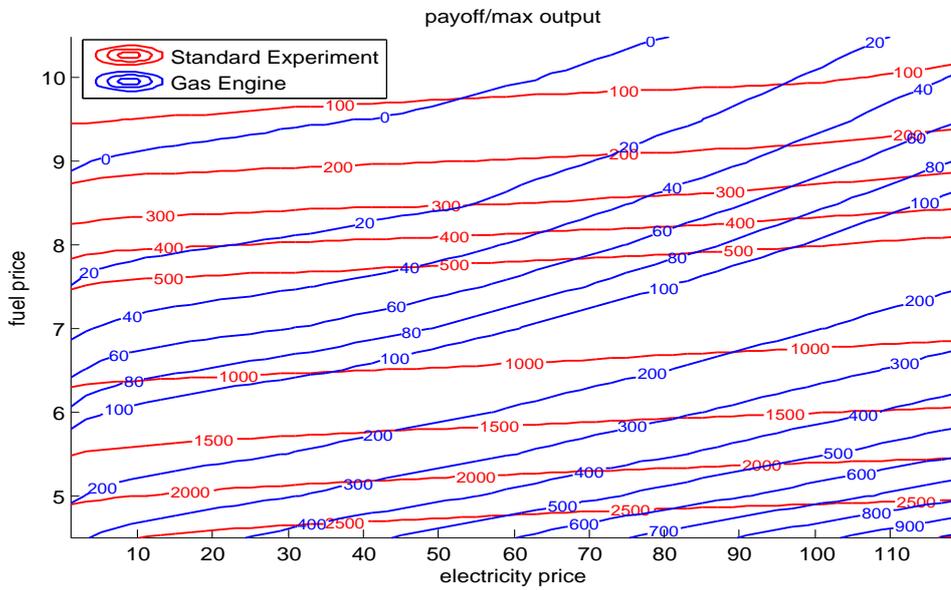


Figure 5.19: Payoff/ q^{\max} ($N=250$) for $x_1 = 1, t = 1$

maximal production capacity per hour (q^{\max}) to scale the payoffs of the two

simulations. Again the graphic illustrates the greater importance of electricity prices for the *gas engine*. Besides, it shows that price constellations exist where it is still possible to generate profit with the standard plant, when production with the gas engine powered plant does not pay any more.

Chapter 6

Conclusion

In the present work, the evaluation of a thermal power plant has been realized according to the algorithm proposed by Tseng and Barz in [1]. This final chapter provides an overview of the insights gained during the implementation and testing of the resulting program.

6.1 Improvements of the present algorithm

6.1.1 Indifference Loci

The most unpleasant observation at testing the algorithm was the bad description of IL by a function of electricity prices only (section 4). Possibly this problem can be avoided by using fuel prices instead of electricity prices to describe the IL. A definitively better description of IL could be achieved by using electricity and fuel prices to determine the IL (see section 4.3.4).

Another problem is the huge time consumption of the computations involved. Saving computing time by estimating the IL from few points or choosing a low number of Monte Carlo simulations may result in considerable inaccuracies. A promising way to reduce computation time would be the parallel computation of price pairs on one IL.

It can also be possible to benefit from the convergence behavior of the IL (figure 5.15). Over a longer time period it can be of advantage to approximate IL by their corresponding predecessor instead of computing them again. This approach was used experimentally in section 5.4.2. Note that apparently the convergence of the IL depends very much on the price processes used (section 5.3.1).

6.1.2 Monte Carlo estimation

Another possibility to lower the computational burden could be provided by variance reduction methods for the Monte Carlo estimator.

Apparently, the complexity of the distribution of payoffs requires the application of more sophisticated methods than antithetic variates. The analysis of the histograms of the payoffs for different scenarios in section 5.2.2 shows that this method is not appropriate to improve the estimation.

Due to the limited computer resources available, all simulations have been conducted with a quite small number of scenarios for the Monte Carlo simulations ($N=750$). Of course it is debatable if this number would be sufficient to generate results for a purpose beyond the mere demonstration of functionality as in this work. Beside the inaccuracy resulting from the low number of scenarios used for the Monte Carlo estimator, its lack of robustness at some price constellations can lead to inappropriate estimations of payoffs (see figure 5.2).

6.2 The potential of the current program

Probably, the most interesting application of the program is not the determination of the exact payoff of a plant but the comparison of the performance of units with different physical characteristics and the impact of different price developments on the power plant value. Since the algorithm was used for artificial examples only, the empirical results are limited to demonstrations of the performance of the evaluation method:

The comparison between the payoffs from the *standard experiment* resulting from two price models was demonstrated in section 5.1.1. For the plant type modeled, the confrontation of comparable initial price observations showed that electricity prices with hourly pattern provide a higher profit than prices without pattern.

Different types of plants were compared in section 5.4, illustrating the impact of unit commitment time, decision lead time and the efficiency of power generation on the production strategy.

The current settings for the simulations performed (especially the number of Monte Carlo simulations) in connection with the problems discussed above return results of a questionable precision. Nevertheless, the computed payoffs turned out to be reasonable estimates compared to the results based on less subtle control strategies (see section 5.1.2). The obvious difference between the results of the three evaluation methods in this section shows that it may pay indeed to use a complex evaluation approach.

Bibliography

- [1] TSENG CHUNG-LI, BARZ GRAYDON: *Short-term generation asset valuation: A real options approach*; Operations Research Vol. 50, No. 2, March-April 2002, pp 297-310
- [2] Weber, Christoph, *Uncertainty in the electric power industry, Methods and Models for Decision Support*, Springer Science+Business Media, Inc. 223 Spring Street, New York, NY 10013, USA, 2005
- [3] Seydel, Ruediger: *Tools for Computational Finance*, Springer-Verlag Berlin Heidelberg, 2002
- [4] Gentle, James E.: *Random Number Generation and Monte Carlo Methods*, Springer- Verlag New York, Inc., 1998
- [5] Gessner P., Wacker H.: *Dynamische Optimierung*, Carl Hanser Verlag Muenchen, 1972
- [6] Bellman, Richard: *Dynamic Programming*, Princeton University Press, Princeton New Jersey, 1957