Analysis of Hedging Strategies for Hydro Power on the Nordic Power Market

Patrik Gunnvald
Viktor Joelsson

April 2, 2015

Abstract

Hydro power is the largest source for generation of electricity in the Nordic region today. This production is heavily dependent on the weather since it dictates the terms for the availability and the amount of power to be produced. Vattenfall as a company has an incentive to avoid volatile revenue streams as it facilitates economic planning and induces a positive effect on its credit rating, thus also on its bottom line. Vattenfall is a large producer of hydro power with a possibility to move the power market which adds further complexity to the problem. In this thesis the authors develop new hedging strategies which will hedge more efficiently. With efficiency is meant the same risk, or standard deviation, at a lower cost or alternatively formulated lower risk for the same cost. In order to enable comparison and make claims about efficiency, a reference solution is developed that should reflect their current hedging strategy. To achieve higher efficiency we focus on finding dynamic hedging strategies. First a prototype model is suggested to facilitate the construction of the solution methods and if it is worthwhile to pursue a further investigation. As this initial prototype model results showed that there were substantial room for efficiency improvement, a larger main model with parameters estimated from data is constructed which encapsulate the real world scenario much better. Four different solutions methods are developed and applied to this main model setup. The results are then compared to reference strategy. We find that even though the efficiency was less then first expected from the prototype model results, using these new hedging strategies could reduce costs by 1.5 % - 5%. Although the final choice of the hedging strategy might be down to the end user we suggest the strategy called BW to reduce costs and improve efficiency. The paper also discusses among other things; the solution methods and hedging strategies, the term optimality and the impact of parameters in the model.

Acknowledgement

We would like to thank our supervisor at Vattenfall, Olof Nilsson, for the continuous support, encouragement and feedback. We would also like to thank our advisor at KTH, Prof. Boualem Djehiche for the support and valuable input.

Finally we would like to thank Alexander Aurell, Daniel Boros, Magnus Bergroth and Rickard Gunnvald for the long but great years we have been studying together at KTH.

Stockholm, March 31, 2015
Patrik Gunnvald and Viktor Joelsson

Contents

1 Introduction

Hydro power as a source of electricity generation is considerable in the Nordic countries. In Sweden alone hydro power constitutes roughly 45 % of the total power generation. It is not only the sheer size that makes it an interesting area to study but also because of its character. When large amounts of hydro power is accessible and produced due to e.g. rainfall combined with full reservoirs, the price of electricity will usually decrease as a result of the increased supply but relatively constant demand (this will be further elaborated in Section 2). More precisely this results in a negative correlation between the hydro balance and the spot price of electricity. Furthermore, the annual generation from hydro power in Sweden is around 65.5 TWh but the variation in downfall could lead to as much as 90 TWh or as little as 40 TWh produced annually. For companies who have large amounts of hydro power, this is troublesome.

Vattenfall is one of those companies with a large power generation from hydro power and thus has to deal with this problem. The company normally produces in excess of 30 TWh of hydro power each year, representing almost half the hydro power generation in Sweden or about 15 % of the total annual generation. It is in Vattenfall's interest to have a non volatile revenue stream. The revenue of interest in this paper stems from the electricity that is generated from hydro power and then sold at the power market. Having stable cash flows into the company is beneficial as it among other things facilitates economic planning and has a positive effect on the credit rating. The better credit rating the company has, the lower the cost will be for borrowing money as it is considered a more stable and creditworthy company. This is one of the main reasons why there is a need for hedging. In this thesis hedging will be to sell forward contracts of electricity on the Nordic power exchange Nasdaq OMX, i.e. an agreement to deliver electricity during a given period in the future at a predetermined price.

This thesis will analyze ways of hedging electricity produced by hydro power. Today, Vattenfall applies a static hedging strategy which results in a certain amount of power to be hedged each year. This strategy does not take into account new information that becomes available over time. Examples of information that reasonably could affect the strategy decisions are e.g. changes in the electricity price, new weather forecasts or repairs on hydro facilities leading to a change in expected volume of power generation. This means that regardless if the price of electricity is higher or lower in some months time, the strategy will still be the same. Similarly if Vattenfall after some time expects to generate more or less power than they did at the time of the strategy simulation - the strategy will still be the same.

1.1 Problem formulation

Can Vattenfall hedge its production of electricity from hydro power more efficiently, i.e. at a lower cost and/or lower risk, than what is done today using a more dynamic approach to model hedging which takes new information into account?

1.2 Methodology outline

One way of achieving this is to construct an objective function f of the revenue which penalize volatile incomes or penalize 'risk', i.e. a function of the form

$$f = E[R] - \lambda Var[R]$$

where R is the revenue as a function of the volume to be hedged H and λ is a constant determining the fictional penalty for taking risk, or having volatile revenue streams. In this case the level of risk means the level of standard deviation of the revenue. The constant λ can also be interpreted as a level of risk aversion – the higher λ the greater risk aversion. The idea is that for every λ find an optimal strategy of hedge decisions from today until delivery of the electricity. With these optimal values of f it is possible to construct an efficient frontier for varying λ which will be able to say what hedge strategy to use at any given level of risk. This will be more thoroughly explained in Section 5.

To accomplish this some of the main tasks involve:

- Define the objective function f
- Define the revenue as a function of the volume to be hedged, H
- Model the dynamics of electricity price and expected produced volume from hydro power, since if these changes over time, so will the revenue and likely also the hedge strategy.
- Transfer the price and volume movements into a binomial tree which will be used for the optimization of f. Each node in the tree will contain information about the price and volume. Generating a binomial will require estimating parameters such as volatility of price, volume and the probabilities of reaching certain outcomes taking into account the correlation between them.
- Develop a method that optimizes f with hedge decisions in line with what Vattenfall does today. Alter λ in order to construct a sort of efficient frontier for f (the efficient frontier is constructed by optimizing f for different λ and extracting the σ associated with that solution. We will loosely speak about this as the efficient frontier). This will serve as a reference solution and will help to answer the problem formulation; namely if hedging of hydro power can be done more efficiently.
- Develop new dynamic methods of finding an optimal hedge strategy to compare with the reference solution.

1.3 Our contribution and main findings

Our contribution is to develop new hedging strategies that takes information into account and are more effective than the strategy currently applied today. The dynamic hedging strategies will change the hedging decision depending on new available information e.g. changes in price or expected hydro power generation. We develop a reference case or reference solution (called *QP*), that reflects the current strategy of Vattenfall applied today. The other developed strategies are *BR*, *BW*, *Full Opt* and *Precom*. These new strategies are analyzed and compared to the reference case QP in order to be able to determine their effectiveness. Further details about these strategies and how they work can be found in Section 5, however briefly they can be described as:

QP: Static strategy. Optimizes f as seen from $t = 0$. Can make one hedge decision at each time step. Reference solution, reflects how Vattenfall currently hedges.

Precom: Static strategy. Optimizes f as seen from $t = 0$. Can make one hedge decision at each node.

Full Opt: Dynamic strategy. An iterative version of Precom. Starts at the initial node at $t = 0$ and then works its way forward, optimizing every subtree in the whole problem using the Precom approach at different times and nodes.

BR: Dynamic strategy. Game theoretic approach. Starts at $t = 0$ and works its way forward through the tree. Remembers previous hedging but does only use information in the following 4 nodes in the next time step.

BW: Dynamic strategy. Game theoretic approach. Uses a dynamic programming procedure but due to time inconsistency it does not give an optimal solution, rather a "best response" solution. Starts at $t = T$ and iterates backwards to $t = 0$.

We find that hedging indeed can be done more efficiently. As shown in Table 4 BW, Full Opt and Precom all have about 1.5 %-5 % lower cost, or better efficiency, for the same standard deviation in revenue as the current strategy QP. The cost reductions depend on the risk level σ. Risk reduction is of course increasing with cost, this means that an alternative formulation of lower cost for the same risk level could be lower risk for the same cost. However, only the cost difference numbers are tabulated. After discussing the results we conclude that BW is the best strategy.

1.4 Limitations

Vattenfall does have other assets to hedge as well and not only hydro power. This has to be taken into consideration when implementing the results. We and Vattenfall were aware of this at the outset of this thesis. However the hydro power constitutes a large portion of the Vattenfall portfolio. It will be even more so if Vattenfall sells its German lignite power (which is currently being proposed). Furthermore it is notably difficult to hedge away a lot of variance in hydro power due to the uncertainty in production, thus a major part of the unhedged variance in Vattenfall's total portfolio of all assets stems from hydro power. Sometimes that is used as an argument to not drill down in complex hedging questions for other types energy. This implies that progress made in the hydro area can have spillover effects in other areas and it could also serve as a starting point for other hedging strategies in those areas.

1.5 Structure of the thesis

In Section 2 we give an introduction to the characteristics of electricity as an asset, the Nordic power market, the geography and it's relevance for the transmission of power and availability of hydro power. Section 2 also covers Vattenfall's role in this market. Section 3 contains the relevant theory for the methods used in this paper, especially multi-dimensional binomial trees, quadratic programming, stochastic dynamic programming and game theory. In Section 4 there is a description of the data. Section 5 covers the methodologies used in this paper for estimating parameters, deriving functions, modeling strategies and explains more in detail how the strategies/algorithms work. In Section 6 the results are presented, along with a discussion of the results in Section 7. Finally in Section 8 we present our conclusions and in Section 9 there are some suggestions for further studies.

2 About the Nordic power market

2.1 Energy

The society is, and has been for the last century, heavily dependent on energy. The energy consumption has increased drastically especially the last 50 years due to among others: population increase, more energy demanding processes but also as a consequence of a more globalized and connected society. From transport to industry to households, all need energy every moment of the year. To be able to meet the demand for energy from the society, different sources of energy are used such as fossil fuels (oil, gas, coal, etc.), nuclear and renewables (wind, sun, water, etc.). This dependence leads to a constant debate about energy; what sources to use, how much to use and where to get it from. The control of energy has become increasingly important for states as it powers everything from heating to communication and military forces. There is a strong correlation between economic growth and energy consumption. Even though the causality is still debated, recent reports (Aliakbari 2014) show that they either jointly influence each other or that abundant energy leads to economic growth.

2.1.1 Electricity as a commodity

A large part of the energy consumed today is produced in power plants that are not situated where the vast majority of the energy is consumed. They generate electricity which then can be transferred via the power grid to other locations where it is consumed. Electricity needs to be consumed instantaneous and can not be stored, this is a special property for this commodity. There are exceptions e.g. if one has a pump power plant one can pump water to a higher level above ground and then use the water in a conventional hydro power plant. There is an ongoing discussion whether electric cars could store some of the electricity in their batteries in the future. As of today, however, it is a vanishing small amount so in general it is considered non storable. The property of not being able to store electricity is important since it means that everything produced needs to be consumed and this will affect the price.

2.2 Power Markets

2.2.1 Nordic power

In the Nordic region the power is mainly generated from hydro power plants and this differs from rest of Europe due to the different geographical characteristics. In Figure 1 one can see that more than 50 precent is generated by hydro power and only 12 percent is generated by fossil fuel.

The countries that the Nordic power market consists of are Norway, Sweden, Denmark, Finland, Estonia, Latvia and Lithuania. Some of these countries are divided into smaller price areas. Norway is divided into five different areas, Sweden four and Denmark two.

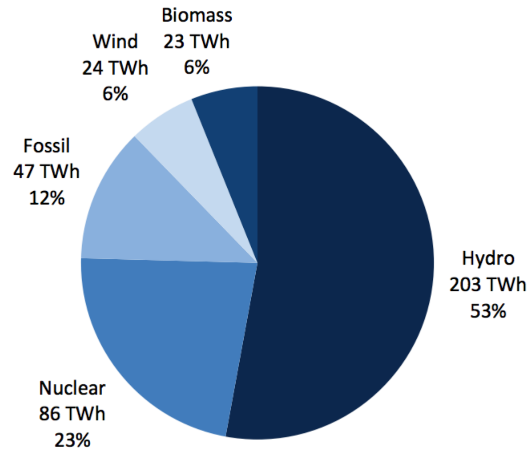

Figure 1: Generation of power in the Nordic region by source during 2013. *(Source: Nordic Energy Regulators)*

In total there are 15 different price areas which are shown geographically in Figure 2. The reason for the existence of different price areas is that the transmission system have bottlenecks. This means that it is not always possible to transfer power from the locations where it is produced to where the demand is sufficiently. The continental Europe power region is connected to the Nordic power region through transmission system and therefore it is possible to transfer power from continental Europe to the nordic power grid or the other way around. The Nordic power region has an electricity/power market called Nord Pool Spot which is owned by the transmission operators in all countries in the region where the trading of physical power is done.

2.2.2 The merit order curve

The trading of physical power is done on Nord Pool Spot on the day-ahead market. This market is a market where one can buy or sell power for physical delivery for all hours under the following day. All bids from buyers and seller should be sent in before 12.00 and then an algorithm calculates the price for each hours during the following day. The price for each hour is set where the demand curve meets the supply curve. The bids sent in by the producers contains how much power they want to sell and at what price. The buyers send in how much they want to buy and at what price. A large actor with many different sources of generation sends in different bids depending on which generation source is used to generate the power. These bids are then stacked into a merit order curve, where the bids are ordered from the lowest to highest price to build up the supply curve. An example of a merit order curve can be seen in the top of Figure 3. The spot price will be set where the supply and demand curves meets. In the bottom of the same figure one can see what happens with the price if an increased production of power

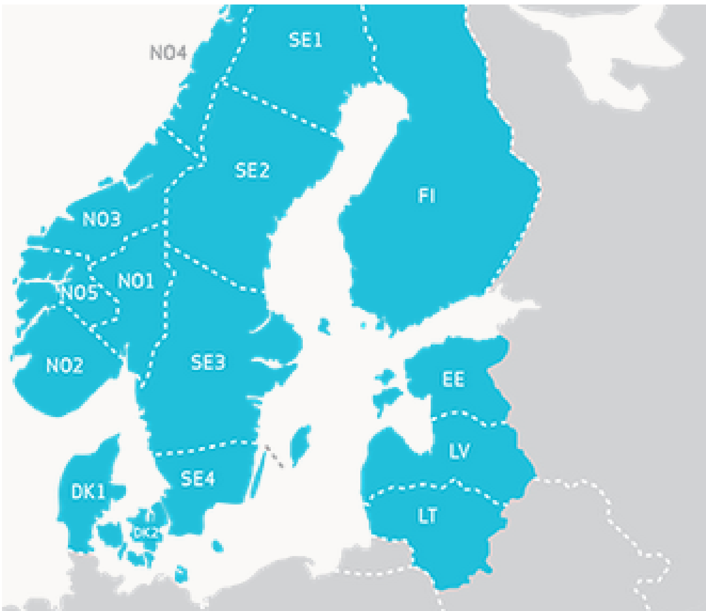

Figure 2: Map over Nordic price areas.
(Source: NordPoolSpot)

generated by a source with low cost happens. The price will decrease cause the elasticity of the demand is much lower then the elasticity of the supply. Therefore the marginal cost of the production where the demand curve meets the supply curve will drive the price.

The price which is set according to this procedure each day is called the system price, abbreviated SYS. This price is a theoretical price that would be the price on the whole market if there where no constraints in the transmission system. Since there are many bottlenecks the price will differ for the different price areas and the SYS price will be theoretical. The reason for the SYS being theoretical is if the demand in one region is higher than its production power needs to be transferred to that region through the transmission system. If it is not possible to transfer the demanded quantity, the price will increase. However, it can be the other way around and then the price will be lower than the SYS price. In Sweden it is rather usual that the price areas in the northern parts have a price that is below the price for the southern parts. This is due to the fact that a large proportion of the Swedish power production is generated in the hydro power plants which are mainly located in the north, while the consumption is in the southern parts of Sweden.

2.2.3 Nordpool & Nasdaq OMX

Electricity is, as pointed out, a non-storable commodity therefore one can not buy it at specific point in time for use of it later on. However, there exist a market on the exchange Nasdaq OMX where one can trade forwards and futures which makes it possible to settle

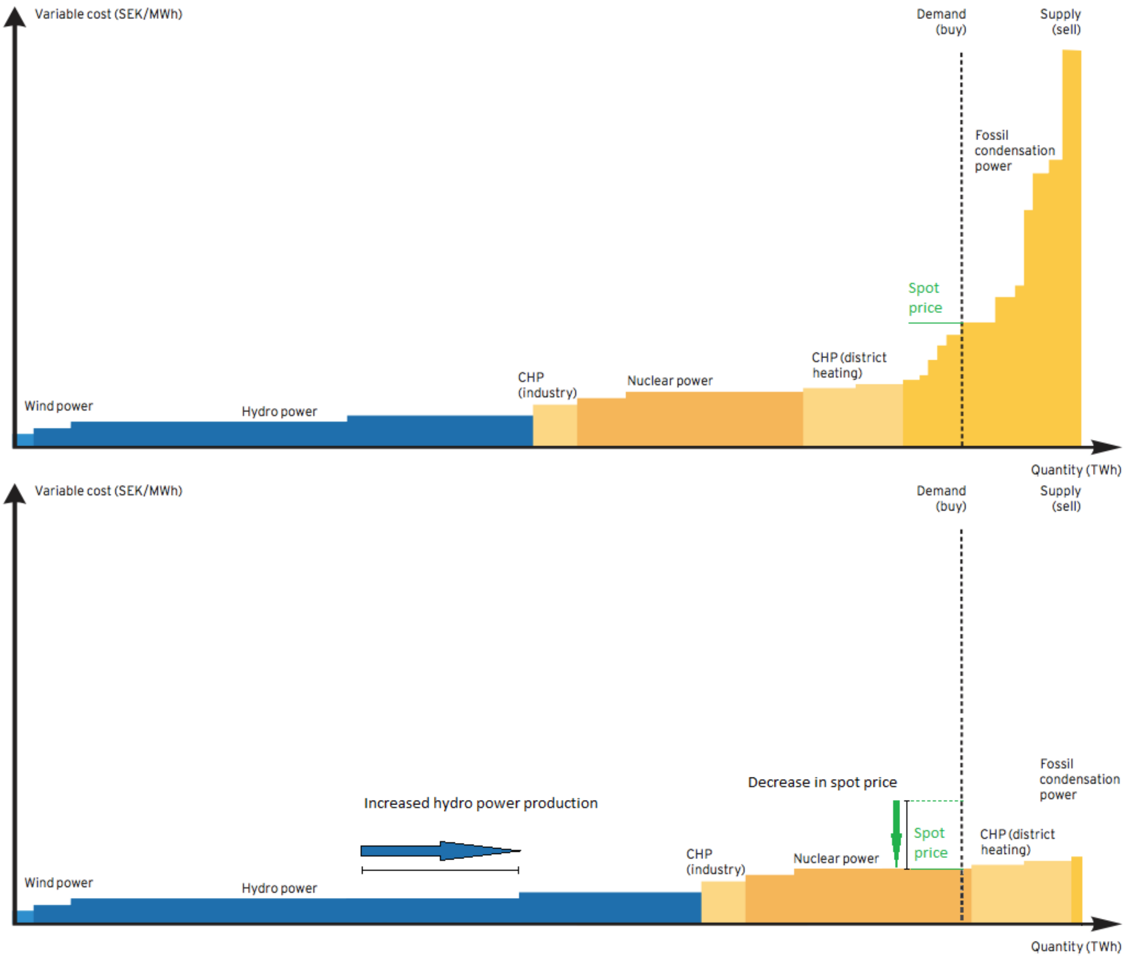

Figure 3: *Top:* Example of a merit order curve.
Bottom: Example of how the merit order curve shifts to the right when additional hydro power is produced and thereby decreasing the price.
(Source original picture: Vattenfall)

a price for power during a given time period. The contracts are traded in MW which will be delivered during every hour of the decided time period, this is called base load. All contracts are financially settled, so the price difference between the decided price and the spot price will be settled between the buyer and the seller of the contract. If the holder wants to buy the actual physical power, then the holder needs to take the money and buy it on Nord Pool Spot instead. The contracts can be bought with time durations from days to years. As an example a buyer can buy 1 MW for year 2017 and will then receive the difference between the forward price and the spot price each day during 2017, which can be both positive and negative.

2.2.4 Hedging

A hedge is when making an investment to decrease the risk of price moments in an asset which may incur potential gains or losses for the holder of the asset. In Example 2.1 it is described why this can be of interest for a power producer:

Example 2.1. A market agent is a producer of asset A, as will be the case of an energy company with asset A being electricity produced. The agent will have a natural long position in asset A and will want to sell in advance to reduce risk and have a less volatile result. What complicates this is that if asset A is electricity, then it is non storable and has to be sold as it is produced. This means that the agent has to sell everything it produces at every point in time, either by selling at the current spot price, selling at a predetermined price from earlier agreements (i.e. positions in earlier forward contracts) or a mixture of both.

As the Nordic market consists of forward and futures agents that are interested in hedging power usage or generation can use these contracts traded on Nasdaq OMX for this purpose.

Note! To hedge in the context of this thesis is for the energy company to sell a forward contract so that the electricity can be delivered at a future date at a predetermined price.

2.3 Hydro power and hydro balance

The Nordic region is heavily dependent on hydro power (roughly 50% of the generation) and the majority is produced in Norway and Sweden. A normal year Norway has a production of 130 TWh electricity from hydro power and Sweden has 65,5 TWh. Finland also has hydro power, but only generates 12 TWh during a normal year.

A year when the catchment of water is over the average is defined as a "wet year" and when it i under the average it is defined as a "dry year". Hydro balance is defined as the difference from the normal value of snow on ground plus the difference from the normal value for the level of water in the reservoirs. The hydro balance can vary a lot and therefore the total amount of electricity produced by hydro power in the Nordic can vary with as much as 100 TWh and for Sweden it can vary in the same relative proportion, around 30 TWh. When the hydro balance is positive it decreases the price and conversely the price goes up when it is negative. In Figure 4 the hydro balance is plotted against the average spot price for the same week and one can clearly see the relationship between spot price and hydro balance – a negative correlation.

Figure 4: Spot price and hydro balance for the period 2013-2015.

2.4 Vattenfall

In Sweden there are a few big actors and one of them is Vattenfall. The company has a hydro power production of around 30 TWh which is a little less then half of the total amount of hydro power generated each year. Having such an exposure to weather, or the precipitation each year, they are interested in the best way possible to hedge their generation. A problem of being a large actor is that one can not put large volumes in the market without moving it against yourself. Therefore one has to take into account how liquid the market is and try to hedge the generation without large movements.

3 Background and application

3.1 Two-dimensional binomial tree

Consider an asset A that follows a Geometric Brownian motion (GBM), then the asset's dynamics follow the SDE:

$$dA_t = \mu_A A_t dt + \sigma_A A_t dW_t, \tag{1}$$

where W_t is a Wiener process and μ_A, σ_A are constants. The asset follows a log-normal distribution according to

$$ln(A) \sim N(\mu_A t, \sigma_A^2 t). \tag{2}$$

The movements of a GBM can be approximated using binomial trees. This is well established in the literature and stems from the fact that the binomial distribution can be approximated as normal distribution if repeated enough times using the central limit theorem. Since the GBM has independent increments, each step in the tree can be considered a repetition and thus the GBM movements can be approximated using the binomial tree approach.

The most basic example of a (one dimensional) binomial tree is where you have one asset A and two outcomes from each state, or node, in the tree. The asset can thus reach one up state u with probability p and one down state d with probability $1 - p$ according to the Figure 5. This report will consider a model with two underlying assets; for the mo-

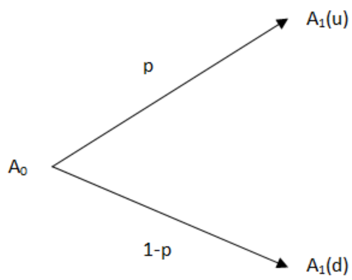

Figure 5: First step in a binomial tree

ment denoted as asset A and asset B. This results in a two-dimensional binomial model, where each asset can reach an up or a down state. Combining these outcomes in one model means that there can be four separately different outcomes in this two-dimensional binomial model. The two-dimensional binomial model can also be thought of as a lattice model which visually represented as a four sided pyramid of outcomes from each node.

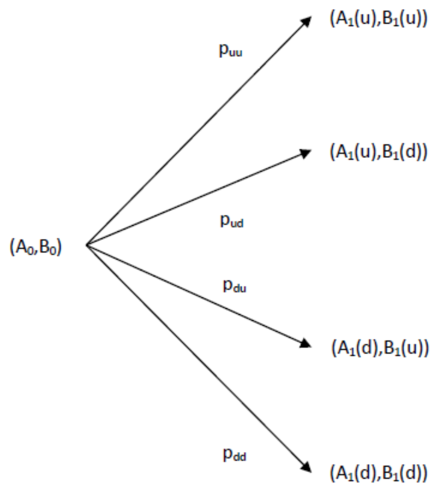

Figure 6: First step in a two-dimensional binomial tree

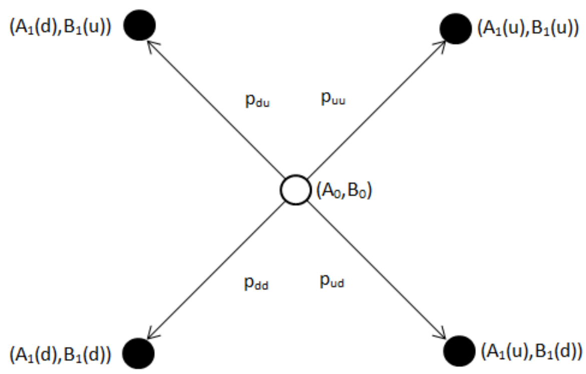

Figure 7: First step in a two-dimensional binomial tree, lattice visualization.

3.1.1 Determining probabilities

In this paper the two assets will be a volume V of electricity produced by hydro power and a forward price S per unit volume which will result in a revenue for the company as the electricity is sold (for more exact definitions please see Section 5.1). Since the electricity contracts are traded at a marketplace it is imperative to have arbitrage free prices. In order to create an arbitrage free tree, the probabilities of the up move and down move have to be correctly assigned. In the two asset case there are four probabilities, $p_{uu}, p_{ud}, p_{du}, p_{dd}$ to be determined. These probabilities are determined as in the paper by Escobar et al. (2009). As shown in the paper, making the assumption that our assets follows geometric Brownian motion dynamics (and so having log-normal probability distributions) this enables us to determine these unknown arbitrage free probabilities as the following:

$$
\begin{aligned}
p_{uu} =& p_S p_V - \frac{\sigma_S \sigma_V \Delta t p_S p_V \rho}{(e^{r\Delta t} - d_S)(d_V - e^{r\Delta t})} \\
p_{ud} =& p_S(1 - p_V) + \frac{\sigma_S \sigma_V \Delta t p_S p_V \rho}{(e^{r\Delta t} - d_S)(d_V - e^{r\Delta t})} \\
p_{du} =& p_V(1 - p_S) + \frac{\sigma_S \sigma_V \Delta t p_S p_V \rho}{(e^{r\Delta t} - d_S)(d_V - e^{r\Delta t})} \\
p_{dd} =& (1 - p_S)(1 - p_V) - \frac{\sigma_S \sigma_V \Delta t p_S p_V \rho}{(e^{r\Delta t} - d_S)(d_V - e^{r\Delta t})}
\end{aligned}
$$

where, for our assets $i \in \{S, V\}$, p_i is the probability of an up move, σ_i the volatility, d_i the size of the down move, r the risk-free rate, Δt the length of the time step and ρ the correlation between the two assets. Notice that if the volatility of one asset is zero, the two dimensional tree collapses into a one dimensional tree and the standard probabilities for a binomial tree has to be used.

The correlation and volatility can also be time dependent (Armerin, 2004), implying that the geometric Brownian motion of our assets satisfy the SDEs:

$$dS_t = \mu_S S_t dt + \sigma_S(t) S_t dW_t \tag{3}$$

$$dV_t = \mu_V V_t dt + \sigma_V(t) V_t dW_t \tag{4}$$

where W_t is a Wiener process and μ_S, μ_V are constants and σ_S, σ_V are time dependent. However, when we generate the binomial trees the volatility will be constant over each separate time step, but the volatility might be different for other time steps.

As the two-dimensional binomial tree will to a large extent be handled with programming, the following simplifying notation is introduced (for more notation see Section 5.1):

$$p_{uu} = p_1, \quad p_{ud} = p_2, \quad p_{du} = p_3, \quad p_{dd} = p_4.$$

In Section 6.1 it is shown that the assumption of log-normally distributed assets is valid.

3.2 Short on T-forward measure

There is a special risk neutral measure when using a bond maturing at time T as numeraire. Under this forward measure, forward prices are martingales. Using the definition in Björk (2009):

Definition 3.1. For a fixed T, the T-forward measure Q^T is defined as the martingale measure for the numeraire process $p(t, T)$.

where $p(t, T)$ is the price at time t of a (zero coupon) bond maturing at time T.

Example 3.2. It is known that the discounted stock price is a martingale under the risk neutral measure. Let Q be the risk neutral measure, $X(t)$ be the stock price at time t and $d(t)$ the discount factor at time t, i.e.

$$d(t) = e^{-\int_0^t r(s)ds},$$

where $r(s)$ is the interest rate. We have that

$$X(t)d(t) = E^Q[X(T)d(T) \mid \mathcal{F}_t],$$

where $\mathcal{F}_t$ is the natural filtration.
Let $G(t,T)$ be the forward price of X at time t, maturing at time T. Then

$$\begin{aligned} G(t,T) &= \frac{X(t)}{p(t,T)}, \\ G(T,T) &= \frac{X(T)}{p(T,T)} = X(T), \end{aligned}$$

and

$$\begin{aligned} G(t,T) &= \frac{E^Q[d(T)X(T) \mid \mathcal{F}_t]}{d(t)p(t,T)} = E^{Q_T}[G(T,T) \mid \mathcal{F}_t]\frac{E^Q[d(T) \mid \mathcal{F}_t]}{d(t)p(t,T)} \\ &= E^{Q_T}[G(T,T) \mid \mathcal{F}_t] \end{aligned}$$

thus the forward price $G(t,T)$ is a martingale under the forward measure Q^T.

Further, futures are martingales under the risk neutral measure Q and if interest rates are deterministic, forward and futures prices coincide. Lemma 26.9 in Björk (2009, p. 404):

Lemma 26.9 *The relation $Q = Q^T$ holds if and only if r is deterministic.*

Thus we conclude that with deterministic rates, forward prices are martingales under the risk neutral measure Q.

3.3 Quadratic optimization

Definition 3.3. Quadratic function. A function $f : \mathbb{R}^n \to \mathbb{R}$ is called a *quadratic function* if

$$f(x) = \frac{1}{2}x^T Qx + c^T x + a,$$

where $Q \in \mathbb{R}^{n\times n}$ is a symmetric matrix, $c \in \mathbb{R}^n$, $a \in \mathbb{R}$.

Lemma 3.4. If f is a quadratic function, then

1. f is convex iff Q is positive semi-definite.

2. f is strictly convex iff Q is positive definite.

The convex optimization problem for the objective function f with no constraints can be written as

$$\begin{aligned} \min_{x} \quad & \frac{1}{2}x^T Q x + c^T x + a \\ \text{subject to} \quad & x \in \mathbb{R}^n. \end{aligned} \tag{5}$$

Theorem 3.5. *Let Q be positive semi-definite. The point $\hat{x} \in \mathbb{R}^n$ is an optimal solution to (5) if and only if $Q\hat{x} = -c$.*

Theorem 3.6. *Let Q be positive definite. Then the vector $\hat{x} \in \mathbb{R}^n$ is the unique optimal solution to (5) given by $\hat{x} = -Q^{-1}c$.*

3.4 Stochastic Dynamic Programming (SDP)

Dynamic Programming (DP) is a systematic mathematical technique to solve problems with interrelated decisions and thus determining the optimal decision policy, or optimal control, for the problem. To be able to apply DP the problem needs to have certain characteristics since there is no standard way of formulating a dynamic programming problem. Two general types of DP are *deterministic dynamic programming*, where the outcome in the next state is entirely determined by the current state and the decision made, and *probabilistic dynamic programming* where it is not entirely dependent on the current state and decision, but also a probability distribution on what the next stage will be. Notice that the probability distribution itself is known at each state. Stochastic Dynamic Programming (SDP) is a form of probabilistic dynamic programming. Hillier and Libermann (2010) presents the basic characteristics of a dynamic programming problem:

1. "The problem can be divided into stages, with a policy decision required at each stage".

2. "Each stage has a number of states associated with the beginning of that stage".

3. "The effect of the policy decision at each stage is to transform the current state to a state associated with the beginning of the next stage".

4. "The solution procedure is designed to find an optimal policy for the overall problem, i.e., a prescription of the optimal policy decision at each stage for each of the possible states".

5. "Given the current state, an optimal policy for the remaining stages is independent of the policy decisions adopted in the previous stage. ... This is the principle of optimality for dynamic programming".

6. "The solution procedure begins by finding the optimal policy for the last stage".

7. "A recursive relationship that identifies the optimal policy for stage n, given the optimal policy for stage $n+1$, is available".

8. "When we use this recursive relationship, the solution procedure starts at the end and moves backward stage by stage - each time finding the optimal policy for that stage - until it finds the optimal policy starting at the initial stage. This optimal policy immediately yields an optimal solution for the entire problem, ...".

   These citations above are all from Hillier and Libermann (2010, p. 429-431).

Characteristic 5 above is sometimes called *Bellman's Principle of Optimality* and it allows the problem to be broken down into smaller subproblems and enables us to find the optimal solution for the entire problem by solving the subproblems as described by Characteristic 8. Another formulation of Characteristic 5 is that each state contains all information necessary to make an optimal decision at that time, or similarly, if a decision policy is optimal on a time interval $\{t_1, ..., T\}$ then it is also optimal for any subinterval $\{t_2, ..., T\}, \quad t_1 \leq t_2$.

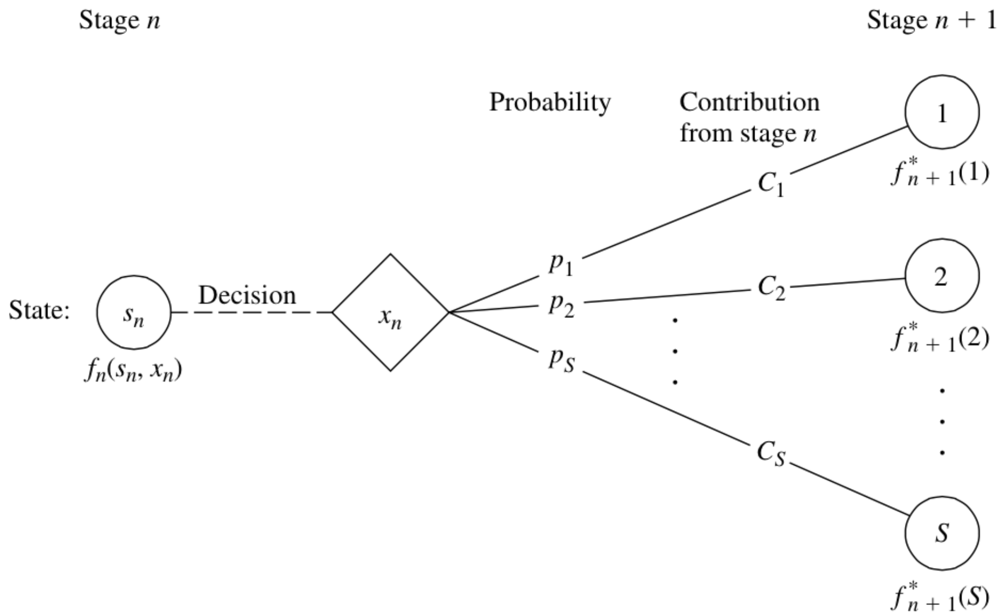

Figure 8: The basic structure for probabilistic dynamic programming. Notice that the objective function f is dependent on the state s and decision x at stage n and the probabilities $p_1, ..., p_S$ are known.
(Source: Hillier and Libermann (2010, p 452.))

3.4.1 Time consistency

Time consistency refers to something, a principle, decision or similar holding over time. In terms of optimization, the optimal control strategy is preferably constant over time meaning that as one makes an optimal strategy at time t, it is still optimal at time $t-1, t-2, ..., 0$. This is of course important if one is to make a series of interrelated decisions over time and find the optimal strategy. If a problem is time *inconsistent*, obviously problems would occur determining an optimal strategy over a time interval. More specifically Characteristic 5 or Bellman's principle is violated. In terms of backward induction this means that you would optimize your objective function for any time t, then taking one step backward to $t-1$ the strategy applied at t might not longer be optimal at this point. It is not even clear what is meant by optimality in this context as the 'optimal' strategy changes over time.

Some common sources of time inconsistent problems as given by Björk and Murgoci (2014):

1. The terminal evaluation function f is a nonlinear function of the expected value. Expected value of a nonlinear function is fine.
2. The terminal evaluation function f is not allowed to depend on the initial point.

3.5 Game theory

One way to consider an optimization problem is in the setting of game theory. This has been described by Björk and Murgoci (2014). We first introduce some concepts of game theory:
A *game* has the following constituents:

- *Players* - those involved in the game or situation
- *Strategies* - the *actions* that the players can choose from
- *Outcomes* - the results of actions taken from the players
- *Utility* - the preferences of each individual player is represented by a utility function for that player. Each player tries to maximize its own (expected) utility, this dictates which strategy or action they take.

The strategies and outcomes are as seen above very general concepts. The games themselves can be *single shot* i.e. only occur one once, or repeated a finite or an infinite amount of times. Players can take actions simultaneously or after one another in sequential games, constrains can be added or removed, the values and beliefs leading to the utility of a certain player can be altered which results in different actions. Altogether this results in that a very large number of (social) situations can be viewed as games.

Throughout this paper we consider only *rational* players.

Definition 3.7. Game. More formally a *game* can be defined as:
$G = \langle J, S, u \rangle$ is a game where
J is the set of players,
$S = \times_{j \in J} S_j$, where $\times$ is the Cartesian product.
S_j is player j's strategy set,
s_j is player j's strategy, $s_j \in S_j$
$s = (s_1, s_2, \ldots, s_m)$ is a strategy profile, where m is the number of elements in J.
$u : S \to \mathbb{R}^{|J|}$ is the combined utility function.

Example 3.8. The prisoner's dilemma
In this setup there are two *players*, A and B who are arrested for a series of crimes and put in prison. They are assumed to be completely rational, they are separated and not able to communicate. There is not enough evidence available to give them both the maximum penalty of 5 years in prison, but there is enough to convict them for a lesser crime resulting in 1 year in prison. During interrogation they are each separately given the possibility to cooperate with the law enforcement and snitch on the other prisoner or to remain silent. Cooperation will be rewarded with lesser time in prison. In this game, each player then has the *strategies* C cooperate and R remain silent. Using Definition 3.7 we have:

$$
\begin{aligned}
J &= \{A, B\} \\
S_1 &= \{R, C\} \\
S_2 &= \{R, C\}
\end{aligned}
$$

The *outcomes* are the following:
- Both player A and player B chooses strategy R, giving them both 1 year in prison.
- Players A chooses C while player B chooses R and remains silent, A is rewarded and set free while B is convicted and receives the maximum penalty of 5 years (and vice versa).
- Both players cooperate and snitches on each other getting them both convicted for the maximum penalty, but due to cooperation they only receive 4 years in prison.

For simplicity define the *utility* functions as:

$$
\begin{aligned}
u_1(s_1) &= -y_1 \\
u_2(s_2) &= -y_2
\end{aligned}
$$

where y_1 and y_2 are player 1 and 2 years in prison respectively. Meaning that the players' utility functions are totally selfish - they care only about their own time in prison and the utility or payoff is linear to the time in prison. Let *(a,b)* be the utilities for player A and B respectively, the game can then be visualized as:

		Player B	
		R	C
Player A	R	(-1,-1)	(-5,0)
	C	(0,-5)	(-4,-4)

Notice that

- If player A believes player B will choose strategy R, player A will prefer to play C since 0 > -1.
- If A believes B will choose C, then player A still prefers to play C since -4 > -5.

The conclusion is that a purely selfish and rational player A will play C no matter what, which will leave B no option but to play C since -4 > -5. Since the game is symmetric the same reasoning can be done from player B's standpoint. The end result is that both end up in a worse situation than if both stayed silent. What also is interesting is that the utility of the individual players is very important. If both players had an egalitarian utility function, e.g. maximizing their total utility, both would play R. Knowing the utility of the other player, one can determine the best response to maximize the own utility.

3.5.1 Subgame perfect Nash equilibrium

Extensive-form games can be used to describe sequencing moves among players. A usual visual interpretation of this is in the form of a *game tree*. Sequential games implies that there is a flow in the tree in only one direction, e.g. left to right or top to bottom. One player makes the decision at each stage of the game and so the tree is built according to the sequence in which the players make their decisions. We assume complete information, i.e. all information about e.g. players' strategies and payoffs are available to all players. A game tree consists of one distinct root node or initial node as well as other nodes which can be of three types:

i) *Chance nodes*, where the transition to the next state is determined by some probability distribution.
ii) *Decision nodes*, where the transition to the next state is determined by a player.
iii) *End nodes*, where all decisions are made and the payoff for player i is decided by u_i.

The extensive-form game is called *finite* if there exists a finite number of nodes. The game is played when the sequence of decisions from root node to end node is made. Notice that these sequences themselves can be viewed as games, or *subgames*.

Definition 3.9. Subgame. Let G be a game and let x be any node in the game tree that is not an end node, then a *subgame* $G(x)$ is a game that has initial node x and consists of x and all the following nodes that can be reached from x.

The payoffs and information from the original game G is inherited in the subgame.

Nash equilibrium is a concept of game theory where in a game of two or more players no one has an incentive to deviate from the strategy on its own, i.e. if any player got the opportunity to change its strategy, given that nobody else did, no one would change it since no one is better off choosing any other strategy. Using the definition of a game from Section 3.5 the Nash equilibrium can be defined as follows:

Definition 3.10. Nash Equilibrium. A strategy profile $s^* \in S$ is a *Nash Equilibrium* of the game $G = \langle J, S, u \rangle$ if

$$\forall j \in J, s_j \in Sj : \quad u_j(s_j^*, s_{-j}^*) \geq u_j(s_j, s_{-j}^*),$$

where s_{-j}^* is a strategy profile of all other players than player j.

In other words, each player's strategy is a *best response* to the opponents' strategies. Notice that in Example 3.8, $(-4, -4)$ is a Nash equilibrium.

Definition 3.11. Subgame Perfect Nash Equilibrium. Let $B_k = \{d_0, d_1, \ldots, d_{k-1}\}$ be the history of the game until stage k consisting of the decisions up to stage k.
Let $G(B_k)$ be a subgame that starts after the history B_k. A Nash equilibrium of $G(B_1)$ is a *subgame perfect (Nash) equilibrium* if the strategy profile forms a Nash equilibrium in all subgames $G(B_k), k = 2, \ldots, T$.

Usually the subgame perfect equilibria in finite games are found using backwards induction. In fact as mentioned by Voorneveld (2015), the methodology of finding subgame perfect (Nash) equilibrium is analogous to that of dynamic programming. Although in a game theory context, the optimal decision is conditioned on the decision maker's opponents decision, which in turn is optimal from their point of view (with their preferences and utilities).

3.5.2 Time inconsistency in game theory

In the setting of game theory time inconsistency can be thought of as changing preferences over time. A player seeks to maximize its utility at any time t by deciding a strategy, but as time flows the player's preferences, and so its utility, changes over time. This results in the previous decided optimal strategy not longer being optimal. As a simple example consider a player choosing between strategy 1: a direct flight to place A taking 10h for \$1000, or strategy 2: a flight to A via B taking 20h and costing \$750. The players preferences today when booking the flight might be "I will still be on a plane for a long time regardless, and I might well just spend some shopping time at place B while changing flights" and so chooses strategy 2. However as the day of travel comes and the player has already been traveling for 9h and still has another 11h to go, the player might gladly pay \$250 or more to reach destination A in an hour. i.e the preferences has changed, and the optimal decision strategy at time of booking is inconsistent with the optimal strategy at time of travel.

4 The data set

In order to build a two dimensional binomial tree that better represents the reality we try to capture in our problem, we have to estimate real world parameters. For this purpose a data sample from Vattenfall is used. The data in this sample contains 1000 simulations of the [forward] price for delivery each month over the period 2015 to 2018 from Vattenfall's own price simulation model. In the data set there are also 46 [expected produced] volume simulations. From this data the volatility parameters can be estimated, please see Section 5.2 for more details. This information enables the construction of a binomial tree with time-varying standard deviations, according to the simulated price and volume paths. As the tree will be used for yearly products, we have chosen to calculate yearly standard deviation for price and volume respectively.

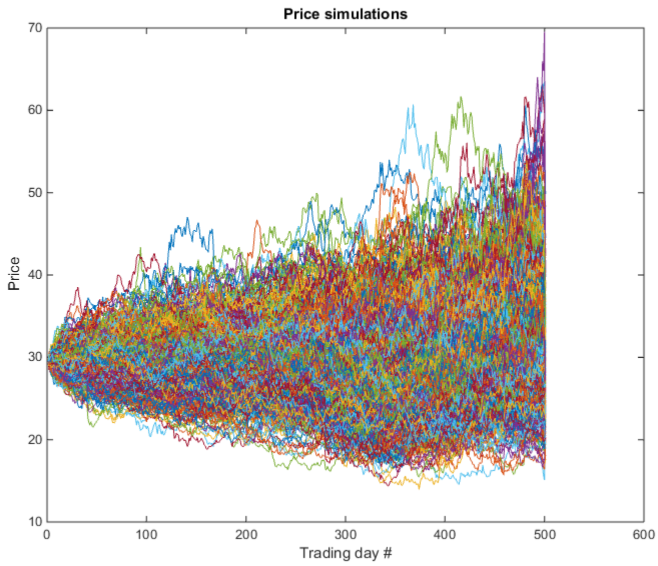

Figure 9: An example of the data obtained from the price simulations. Each simulation consists of 1000 paths and is specific for a certain delivery month a certain year.

The price simulations are specific for a certain delivery month a certain year. The binomial tree structure requires a standard deviation for the corresponding yearly price. To obtain this price the following is calculated:

$$price_y = \frac{1}{12} \sum_{m=1}^{12} w_{m,y} \cdot price_{m,y} \tag{6}$$

where $price_{m,y}$ and $w_{m,y}$ is the price and weight for month m in year y, respectively. The weight factor is due to the different length of the months and thus the monthly prices have a varying impact on the yearly price, as they should. The end result is similar to the example simulation shown in Figure 9, but with yearly price data. This procedure is done for all years in the data sample.

The volume data obtained from simulations are given as volume production per week. To get a yearly outcome, all weeks for that year are added together. This is done for all four years where we have data in the sample. This is because we are looking at forward contracts with yearly duration. Summing the weeks removes the seasonal effects.

For the correlation estimation, historical production plans and historical forward prices are used. A production plan is a plan over how much hydro power that will be produced over the year given all information available.

5 Modelling

5.1 Notations and preliminaries

Please also see Figure 10 which visualizes some of the definitions below.

- V_t: Random process representing volume to be produced under period that starts at T, at time t. Also called *volume* (at time t). V_t is assumed to be a GBM and due to its nature the produced volume is assumed to have zero drift. V_t is a martingale under the risk neutral measure Q.

- S_t: Random process representing forward price at time t of electricity for delivering at time T. Also called *price* (at time t). S_t is assumed to be a GBM with zero drift since it is a forward price. S_t is a martingale under the risk neutral measure Q.

- n: The number of nodes (or outcomes) in the tree. As seen from the combinations of outcomes of our two assets $n = 4$.

- $\sigma_S(t_1, t_2)$: The standard deviation of S in the time period $[t_1, t_2)$.

- $\sigma_V(t_1, t_2)$: The standard deviation of V in the time period $[t_1, t_2)$.

- *up state*, $u_S(t_1, t_2)$: The fractional increase of the price from time t_1 to the next time step t_2,
$$u_S(t_1, t_2) = e^{\sigma_S(t_1,t_2)\sqrt{t_2-t_1}}$$

- *down state*, $d_S(t_1, t_2)$: The fractional decrease of the price from time t_1 to the next time step t_2,
$$d_S(t_1, t_2) = e^{-\sigma_S(t_1,t_2)\sqrt{t_2-t_1}}$$

- *up state*, $u_V(t_1, t_2)$: The fractional increase of the volume from time t_1 to the next time step t_2,
$$u_V(t_1, t_2) = e^{\sigma_V(t_1,t_2)\sqrt{t_2-t_1}}$$

- *down state*, $d_V(t_1, t_2)$: The fractional decrease of the volume from time t_1 to the next time step t_2,
$$d_V(t_1, t_2) = e^{-\sigma_V(t_1,t_2)\sqrt{t_2-t_1}}$$

- *State* i: The state that the assets can reach in the next time step from their current state. From every state, or node, they can transition to 4 different nodes. That is, $i = 1$ corresponds to price and volume moving up from previous state, $i = 2$ to price up and volume down, $i = 3$ to price down volume up and finally $i = 4$ to both price and volume moving down.

- $p_{1,t}$: The probability that the next state after time t is where both S and V are in the up state, $p_{1,t} = P(S_t \cdot u_S, V_t \cdot u_V \mid S_t, V_t)$

- $p_{2,t}$: The probability that the next state after time t is where S is in the up state and V in the down state, $p_{2,t} = P(S_t \cdot u_S, V_t \cdot d_V \mid S_t, V_t)$
- $p_{3,t}$: The probability that the next state after time t is where S is in the down state and V in the up state, $p_{3,t} = P(S_t \cdot d_S, V_t \cdot u_V \mid S_t, V_t)$
- $p_{4,t}$: The probability that the next state after time t is where both S and V are in the down state, $p_{4,t} = P(S_t \cdot d_S, V_t \cdot d_V \mid S_t, V_t)$
- $V_{t+1|i}$: Volume at time $t+1$ (for delivery in period T), given ending up in state i when transitioning from t to $t+1$,

$$V_{t+1|1} = V_{t+1|3} = V_t \cdot u_V$$

$$V_{t+1|2} = V_{t+1|4} = V_t \cdot d_V$$

- $S_{t+1|i}$: Forward price at time $t+1$ (for delivery in period T), given ending up in state i when transitioning from t to $t+1$,

$$S_{t+1|1} = S_{t+1|2} = S_t \cdot u_S$$

$$S_{t+1|3} = S_{t+1|4} = S_t \cdot d_S$$

- $\mathcal{F}_t$ is the natural filtration at time t.
- $\mathcal{G}^i_{t+1}$ is the filtration generated by the σ-algebra $\sigma(S_{t+1|i}, V_{t+1|i})$.
  Note that $\mathcal{G}^i_{t+1} \subset \mathcal{F}_t$ since $S_{t+1|i}, V_{t+1|i}, \quad i = 1, 2, 3, 4$ can be determined by S_t, V_t.
- H_t: Volume to be hedged under period $[t : t+1)$.
- X_t: All hedges done up to t, i.e.

$$X_t = \sum_{\tau=0}^{t} H_\tau$$

- c_t: Parameter of the quadratic hedge cost function at time t.
- R_t: Revenue during period 0 to t.
- $\tilde{R}_t$: Revenue during period t to T.
- $\tilde{R}^*_t(X_{t-1})$: $\tilde{R}_t$ as a function of X_{t-1} and best response hedging during period t to T.
- $E^Q[\bullet]$ will be written in the short hand notation $E[\bullet]$ in this paper.
- $Y^*_{t+1}(X_{t-1} + H_t) = E[(\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}]$ is used as a simplifying notation. Note that Y_t and R_t are functions of the total amount hedged at that time.

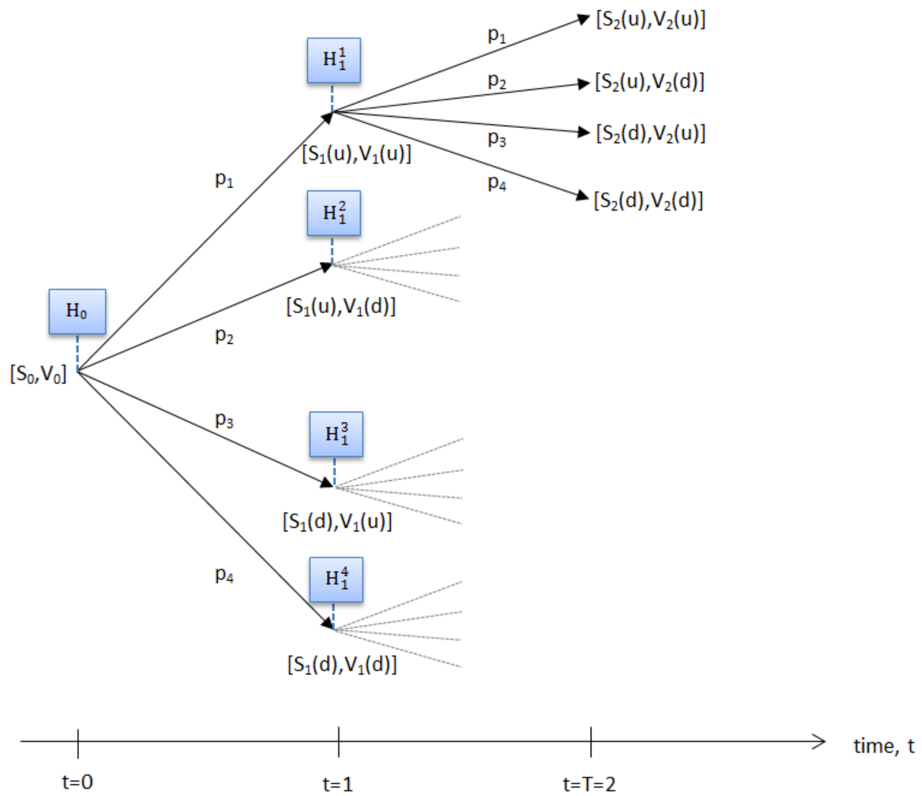

Figure 10: The binomial tree model visualizing some of the variables and definitions. In each node a revenue can be calculated which in turn enables the calculation of the objective function f.

5.2 Estimating parameters

Price volatility: To estimate the volatility of the price, $\hat{\sigma}_s$, we use the data from the 1000 simulations of the forward price in the data set. The estimate of the constant volatility from the initial time 0 to time t can be calculated according to

$$\hat{\sigma}_s(0,t) = \sqrt{\frac{\frac{1}{n}\sum_{i=1}^{n}(s_i(t) - \mu(t))^2}{t}}, \quad \text{where} \quad \mu(t) = \frac{1}{n}\sum_{i=1}^{n} s_i(t) \tag{7}$$

and where $s_i(t)$ is the price of simulation i at time t and n is the number of simulations, in this case $n = 1000$. The calculation of the estimations of the intermediate volatilities between two time steps was then done using the following

$$\hat{\sigma}_s(t,T) = \sqrt{\frac{T\hat{\sigma}_s^2(0,T) - t\hat{\sigma}_s^2(0,t)}{T-t}}. \tag{8}$$

These estimates of the price volatilities are used for the construction of the binomial tree.

Volume volatility: The expected volume to be produced is based on the outlook of the hydro balance. Since it is impossible to predict how much precipitation that will fall far away in the future the assumption that nothing is known about outcome the hydro balance until you are less then 2 years before delivery of the contract is used. Therefore the volatility of the volume will be zero for all time steps where there are more than 2 years to delivery. This will affect tree in the way that the two dimensional tree will only have price movements during the years of zero volume volatility.
As described in Section 4 there are no path simulations for the volume, instead there are 46 observations for each year, the next five years. To estimate the yearly volume volatility, $\hat{\sigma}_v$, the same procedure as we used for the price volatility estimations, but with $n = 46$ and $s_i(t)$ replaced by $v_i(t)$.

Correlation: The correlation $\hat{\rho}$ between price and volume is estimated using Vattenfall's production plans for their volumes and forward prices from Nasdaq OMX. The data is from the two year period 2010-2012. Production plans are plans on how much hydro power Vattenfall will produce each year given all information available at the time. The estimated correlation is calculated between the production plan for each of the four coming years and the price of the forward contracts those years.

5.3 The revenue function, R

The revenue is generated by selling produced electricity to the market. The details of the revenue function are explained below:

1. The revenue is generated by two parts, electricity sold on the spot market and electricity sold by forward contracts.
2. To hedge, or sell electricity with forward contracts, is associated with a cost of transaction and a cost related to the liquidity of the market (i.e. low liquidity can move the market against yourself). As the cost associated with liquidity is by far the dominating one, the cost of the transaction will be neglected.
3. The cost of hedging should be more expensive further from delivery, due to the fact that those contracts are less liquid.
4. The cost of hedging should increase non-linearly with increased hedged volume to reflect liquidity, supply and demand.

To take the properties of the cost of hedging into account a quadratic hedge cost function is used

$$\sum_{t=0}^{T-1} c_t H_t^2, \tag{9}$$

where c_t and H_t is defined as in Section 5.1.

Combining all of the above and let the revenue be defined as

$$\begin{aligned} R_T =& V_T S_T - \sum_{t=0}^{T-1} H_t S_T + \sum_{t=0}^{T-1} H_t S_t - \sum_{t=0}^{T-1} c_t H_t^2 \\ =& V_T S_T - \sum_{t=0}^{T-1} H_t (S_T - S_t + c_t H_t), \end{aligned} \tag{10}$$

where the first term the right hand side is the outcome of price and volume produced, terms two and three are the contribution of the hedged volume and the fourth term is the cost associated with hedging.

5.4 The cost of hedging

How to determine the parameter c_t? Initially, one might think that the cost of hedging might be a diminishing small portion of payment to enter contracts in the market. As described earlier, being a large player one has to think about not only the actual cost of hedging, but also how much the market can move against you as a result of putting large volumes in the market. Thus the parameter c_t is related to how much is traded at the market by Vattenfall compared to the market liquidity at that time. The calculations to determine c_t is done with more detail when solving the optimization problem, but can't be disclosed here due to confidentiality reasons.

5.5 The objective function, f

In this thesis the idea is to

$$\begin{aligned} \max_{H} \quad & E[R_T] \\ \text{s.t.} \quad & \sigma \leq \sigma_{target}, \end{aligned} \tag{11}$$

where R_T is the revenue, σ is the standard deviation of the revenue and σ_{target} is the target standard deviation set by the company, which may not be exceeded.

Since the standard deviation is always positive, both sides of the constraint above can be raised to the power of two, $\sigma^2 \leq \sigma^2_{target}$, and still be valid. The problem can now be relaxed as

$$\max_{H} \quad E[R_T] - \lambda[\sigma^2 - \sigma^2_{target}], \tag{12}$$

where λ is a factor adjusting the variance impact on the expected return. The target variance, σ^2_{target}, is a constant and therefore the maximization problem (12) can be written as

$$\max_{H} \quad E[R_T] - \lambda[\sigma^2]. \tag{13}$$

As follows from above the object function f, which this thesis will strive to optimize, is defined as

$$f = E[R_T] - \lambda Var[R_T], \tag{14}$$

where $\lambda \in [0, \infty)$ is a risk aversion factor that can be adjusted to increase/decrease the penalty of variance on f and is used to construct the efficient frontier.

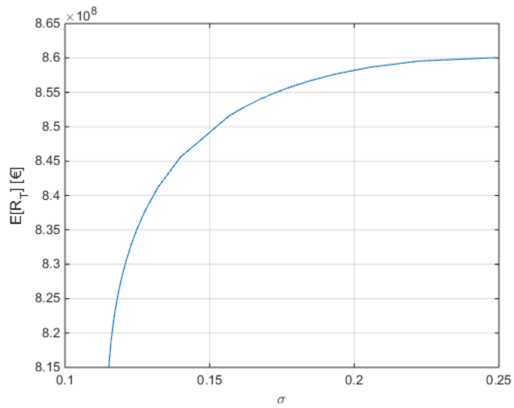

Figure 11: Expected value against the standard deviation.

5.6 Solution Methods

5.6.1 Quadratic Programming

The objective function f in (14) can be written as a quadratic problem, if one hedge decision per time step is taken (derivation see Appendix A.1)

$$\begin{aligned} f(\bar{H}) =& E[S_T V_T] - \sum_{t=0}^{T-1} c_t H_t^2 \\ & - \lambda(E[S_T^2 V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] \\ & - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] \\ & - 2E[\sum_{t=0}^{T-1} H_t S_T \sum_{t=0}^{T-1} H_t S_t] - (E[V_T S_T])^2), \end{aligned} \tag{15}$$

where $\bar{H} = (H_0, H_1, \ldots, H_{T-1})^T$ and each H_t is the volume to be hedged during time period $t, \quad t = 1, \ldots, T-1$.

The objective function, f, written as in Equation 15 can be optimized using quadratic programming. To facilitate implementation in Matlab we minimize $-f$, which is equal to maximizing f. Hence, the negative objective function, $-f$, can be rewritten on matrix form with the following structure for the quadratic programming

$$-f = \frac{1}{2}\bar{H}^T Q \bar{H} + \bar{L}^T \bar{H} + \phi, \tag{16}$$

$$Q = -2(\lambda(E[S_T^2] \begin{bmatrix} 1 \\ \vdots \\ \vdots \\ 1 \end{bmatrix} - 2\beta + \alpha) + \begin{bmatrix} c_0 & 0 & \ldots & & 0 \\ 0 & c_1 & 0 & \ldots & 0 \\ \vdots & & \ddots & & \vdots \\ 0 & \ldots & & & c_{T-1} \end{bmatrix}, \tag{17}$$

$$\bar{L} = -2\lambda(\gamma - E[V_T S_T^2] \begin{bmatrix} 1 \\ \vdots \\ \vdots \\ 1 \end{bmatrix}), \tag{18}$$

$$\phi = -(E[V_T S_T] - \lambda(E[S_T^2 V_T^2] - E[S_T V_T]^2)), \tag{19}$$

$$\alpha = \begin{bmatrix} E[S_0^2] & E[S_0 S_1] & \dots & E[S_0 S_{T-1}] \\ E[S_1 S_0] & E[S_1^2] & & \vdots \\ \vdots & & \ddots & \\ E[S_{T-1} S_0] & \dots & & E[S_{T-1}^2] \end{bmatrix}, \tag{20}$$

$$\beta = \begin{bmatrix} E[S_T S_0] & \frac{1}{2}E[S_T(S_0+S_1)] & \dots & \frac{1}{2}E[S_T(S_0+S_{T-1})] \\ \frac{1}{2}E[S_T(S_1+S_0)] & E[S_T S_1] & & \vdots \\ \vdots & & \ddots & \\ \frac{1}{2}E[S_T(S_{T-1}+S_0)] & \dots & & E[S_T S_{T-1}] \end{bmatrix}, \tag{21}$$

$$\gamma = \begin{bmatrix} E[V_T S_T S_0] \\ E[V_T S_T S_1] \\ \vdots \\ E[V_T S_T S_{T-1}] \end{bmatrix}. \tag{22}$$

In equation (16) ϕ is a constant and therefore will not be included during the optimization. The problem is now set up as a quadratic programming problem with no constraints and can be solved accordingly. We note that this solution will only give one hedge decision at each time step, applied to all nodes in that time step. The other methods will be able to provide a [different] hedge decision at every node in the tree, even if they are at the same time t.

5.6.2 Best Response Backwards Induction (BW)

A first thought would be to try and find the optimal solution to the maximization of f via Stochastic Dynamic Programming (SDP). However, considering this problem formulation, a pure SDP solution cannot be used due to the fact that knowledge of previous hedging is needed in order to make an optimal hedge decision in the current node. Dynamic programming requires the current state to contain all information of previous behaviour that is needed to make an optimal decision at this point. This is also called the *Markov property*. As the states in this problem contains information about volume and price, but not about amount previously hedged, the problem violates the principle of optimality for dynamic programming. As described by Hillier and Lieberman (2010), any problem lacking this principle cannot be formulated as a dynamic programming problem, i.e. we have violated Characterstic 5 as described in Section 3.4. This problem can at least partially be mitigated which we will discuss more soon. Furthermore, even if all information was known at each state, the objective function as described in Section 5.5 cannot be optimized with SDP due to the time inconsistency of the variance term. This is quickly realized via

$$Var[X] = E[X^2] - (E[X])^2$$

where the first term is fine since we are allowed to have the expected value of a nonlinear function. The second term however is a nonlinear function of the expected value and will cause time inconsistency.

However, one way of thinking about this optimality problem, since it cannot be optimized in the true meaning of the word, is to apply game theory as earlier described in Section 3.5. This means that the *utility function* f has to be maximized using a strategy containing different amounts of hedges at each time period and node.

This situation can be thought of as a series of subgames with one subgame at each time step and node. Further, there is one player at each time step, responsible for making an as good hedge decision as possible at that time, i.e. we have a sequential game. The players can be thought of as a reincarnation of ourselves but with different preferences. This reflects the time inconsistency of our problem where the utility function changes for each time step. Each reincarnation of ourselves tries to maximize the utility function f as seen from that point of view in time.

Now, the last player makes its best decision at time $T-1$ to find the subgame perfect Nash equilibrium where he/she will not deviate. Now, with backwards induction, knowing what the last player at $T-1$ would play, the player at $T-2$ will choose a strategy that will maximize f at this point, *already knowing* what strategy the player at $T-1$ will play once it is that player's turn, since the utility function is known. Iterating this backwards, each player at each time step have predicted what the players after will play and would so be able to decide what they should play. Once it reaches the player at $t=0$, the hedge strategy decision chain is complete and the optimal decision policy, or hedge strategy, is found. Notice the difference that the optimal strategy decision is not optimal for the entire problem, but rather an optimal strategy as a response to the other players optimal strategies from *their* point of view, i.e. the optimal strategy is conditioned on the other players decisions which are made according to their utilities.

Adding information

As mentioned earlier our specific problem does not contain all information necessary to make a best response decision at all times. The player at time t can predict all other players strategies at times $t+1, ..., T-1$, but as the objective, or utility, function is structured it also depends on the decision of previous players. If the player at time $T-1$ does not know the previous players' strategies, he/she cannot make a decision. And if the player at $T-1$ cannot make a decision, obviously the player at $T-2$ cannot either. The information that is missing is to make a best response decision is the *total amount of previous hedging*, X_{t-1} (if the decision is to be made at time t). This can be realized from Equation 24 together with the expression for expected revenue in Equation 25 and variance of revenue in Equation 26.

Consider being at time t. To find the best response decision a discrete number X_{t-1} is used as input at each node at time t. Loosely speaking, this corresponds to the question to the player at time t: "If you know that the total amount X_{t-1} is hedged previously, what hedge decision would you make?". Note that the player at time t of course requires a lot of other information, but the information of previous hedging is the only information that is missing in the node. Therefore, for every different X_{t-1} input, we get a best response hedge decision as a part of the overall hedging strategy.

Consider the following example: Let L be a discrete list of numbers, containing {-100, -95, ..., -5, 0, 5, ..., 100}. Then at a given node n at time t we try to find the best hedge decision H^* at that point. We will then get:
$H^*_{t,-100}$ for $X_{t-1} = -100$
$H^*_{t,-95}$ for $X_{t-1} = -95$
...
$H^*_{t,100}$ for $X_{t-1} = 100$
At each node we thus produce a list of best hedge decisions. Which one to actually use depends on what decisions are made previously. Once we have iterated backwards to $t = 0$, the player at $t = 0$ will know that previous hedging is 0 since we start out unhedged. The chain of events is then the following:
The player at $t = 0$ have all information to make a best response hedge decision, assume the decision is to hedge 15 units. At next time step $t = 1$, we can then enter the created list, and pick out $H^*_{1,15}$ in every node (remember that there are four nodes at $t = 1$). That is, given that player 0 hedged 15 units, assume $H^*_{1,15}$ corresponds to e.g. the player at time 1 hedging 20 units in the top node, 10 units in the upper middle node, 15 units in the lower middle node and 5 units in the bottom node. Then this process continues, for all four nodes at time $t = 2$ that can be reached from the top node at time $t = 1$, we go into the list at every node and pick $H^*_{2,35}$ etc. Note that this also constrains what hedge decisions that can be made, a player can not take e.g. a non-integer hedge decision if the discrete numbers in the list are all integers. This is what we call *resolution* for BW.

Using this list, a best response decision can be made in all nodes at time t, for all $X_{t-1} \in L$. Similarly at time $t-1$ lists are added to each node containing information about X_{t-2} etc. One way of looking at it is that information is added to the state where it was previously lacking.

Of course this information consists of a finite number of discrete elements and not a continuum as one would wish. If a list of this sort is created, then one can iterate backwards and find the best response hedging strategy via backward induction, as described above. In theory it is desirable to always have more discrete numbers in the list L to use as input since in general the step-wise solutions tend to those of a continuous solution. Although in practice this has to be balanced against computational time. The aim however is to always use an as fine resolution as possible.

The utility function

Using the notation in Section 5.1 it is clear that $\tilde{R}_t$ will be the same as R_T once $t = 0$ and so the solutions are comparable. For a time t we have that the revenue $\tilde{R}_t$ satisfies

$$\tilde{R}_t = -c_t H_t^2 - (X_{t-1} + H_t)(S_{t+1} - S_t) + \tilde{R}^*_{t+1}(X_{t-1} + H_t). \tag{23}$$

At each time t, seek to maximize the (utility/objective) function

$$f = E[\tilde{R}_t \mid \mathcal{F}_t] - \lambda Var[\tilde{R}_t \mid \mathcal{F}_t]. \tag{24}$$

As shown in Section A.2 we have that

$$E[\tilde{R}_t \mid \mathcal{F}_t] = -c_t H_t^2 + \sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}], \tag{25}$$

and

$$\begin{aligned} Var[\tilde{R}_t \mid \mathcal{F}_t] = {} & (X_{t-1} + H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] \\ & + \sum_{i=1}^{4} p_i Var[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}] + Var[Y^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{F}_t] \\ & - 2(X_{t-1} + H_t) Cov[S_{t+1}, Y^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{F}_t], \end{aligned} \tag{26}$$

which yields an expression for f. Notice that since backward induction is used, $E[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}]$ and $Var[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}]$ are known from the previous time step calculation.

At time $t = T - 1$ we have a special case since have

$$\tilde{R}^*_{t+1}(X_{t-1} + H_t) = \{t = T - 1\} = \tilde{R}^*_T = S_T V_T,$$

which yields the following (see Section A.2.1 for derivation)

$$E[\tilde{R}_{T-1} \mid \mathcal{F}_{T-1}] = -c_{T-1} H_{T-1}^2 + \sum_{i=1}^{4} p_i S_{T|i} V_{T|i} \tag{27}$$

and

$$\begin{aligned} Var[\tilde{R}_{T-1} \mid \mathcal{F}_{T-1}] = {} & (X_{T-2} + H_{T-1})^2 Var[S_T \mid \mathcal{F}_{T-1}] + Var[S_T V_T \mid \mathcal{F}_{T-1}] \\ & - 2(X_{T-2} + H_{T-1}) Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}] \end{aligned} \tag{28}$$

which yields an expression for f at time $t = T - 1$.

We again reiterate that this is not an optimization in the true meaning of the word but rather a best response solution where the individual steps are optimized with respect to the utility function at those respective times, or in game theory terms a subgame perfect Nash equilibrium is found.

5.6.3 Best Response Forward (BR)

An alternative way to consider this problem is to see it as a game but instead of using backwards induction, a best response forward iteration is used. Loosely speaking, the rationale behind it is the following:

- Start at time $t = 0$. Knowing *only* the outcomes and their respective probabilities one step ahead, what is the best strategy decision (amount to hedge) to use in order to maximize the objective/utility function f?
- At time $t = 1$, knowing what strategy decision was made in the earlier time step and the outcomes one step ahead, what is the best strategy decision to make?
- At time t, knowing what strategies that are made in previous time steps $t-1, t-2, ..., 0$ and the outcomes one step ahead, what is the best strategy decision to make? And similarly up to time $t = T-1$ where the last decision is made.

It is clear that this solution will not take the whole problem into account when making the hedge decisions. Mathematically we have that:

The revenue function at time $t+1$ is

$$R_{t+1} = R_t - c_t H_t^2 - H_t \cdot (S_{t+1} - S_t) - \sum_{k=0}^{t-1} H_k \cdot (S_{t+1} - S_t) + (S_{t+1}V_{t+1} - S_t V_t) \quad (29)$$

and again we want to maximize the function f

$$f = E[R_{t+1} \mid \mathcal{F}_t] - \lambda Var[R_{t+1} \mid \mathcal{F}_t] \quad (30)$$

where

$$E[R_{t+1} \mid \mathcal{F}_t] = R_t - c_t H_t^2 + Cov(S_{t+1}, V_{t+1} \mid \mathcal{F}_t) \quad (31)$$

and

$$\begin{aligned} Var[R_{t+1} \mid \mathcal{F}_t] = (H_t + \sum_{k=0}^{t-1} H_k)^2 Var(S_{t+1} \mid \mathcal{F}_t) - 2(H_t + \sum_{k=0}^{t-1} H_k) Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t) \\ + Var(S_{t+1}V_{t+1} \mid \mathcal{F}_t). \end{aligned} \quad (32)$$

First order condition yields the optimal H_t,

$$H_t = \frac{Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t) - \sum_{k=0}^{t-1} H_k \cdot Var[S_{t+1} \mid \mathcal{F}_t]}{(c_t/\lambda) + Var[S_{t+1} \mid \mathcal{F}_t]}. \quad (33)$$

The derivation of this can be seen in Appendix A.3.

5.6.4 Precommitment (Precom)

Alternative solutions can be used depending on how one chooses to view the optimality problem and may be a question of which strategy you want to use in the company. One way to look at it is to optimize the objective function f today, but without taking into account that the subtrees not might be optimal. This entails that as time flows, or steps are taken in the tree, the solution from that point might be suboptimal. This hedging strategy is created with one optimal hedge decision at every node to make the outcome at $t = 0$ look as promising as possible.

This is done numerically by using Matlab's function *fminunc* to find the hedge decisions which optimize the f-value today at $t = 0$. For further information about the algorithms used by this function we refer to Matlab's documentation.

5.6.5 Iterated optimization (Full Opt)

This method of finding a solution to the optimization problem uses the Precom method, but with an iterative approach. The idea is the following:

- Optimize the objective function at $t = 0$ using *fminunc*, giving each node an optimal hedge decision that should be made at that point.
- Take one step in time to $t = 1$, keep the hedge decision H_0 made at time $t = 0$, optimize all four subtrees using the same function, *fminunc*.
- Continue to all nodes in $t = 2$, keeping the decisions made at $t = 1$ (and $t = 0$ too of course), optimize these subtrees, etc.
- When all subtrees in the final step has been optimized, the hedge decisions that are at each node in the whole tree is considered to be the optimal hedge strategy.

Using this approach, it is true that the solution at time $t = 0$ will look less favourable since the hedge decisions have been altered. But, this is a trade off made to make the subtrees be more optimized. We recognize that the H_0 is optimized given the previous, unchanged hedge decision but present it as an alternative way of thinking of the optimization problem.

5.7 The prototype model

In order to construct algorithms for solving this problem, starting with a simpler model is often useful to get a grasp on how they work (or not work when searching for errors). This involved constructing a two-step binomial tree, with made up data. The prototype model will make it possible to start work on the algorithms before the data is acquired and parameters estimated. It will also give a hint on whether the solutions actually become better and the project is worth pursuing in full scale. It is therefore important that this tree reflects the characteristics of the final tree, using proper simulated data. However the results from the prototype model should be considered with some caution, the main reason for constructing the prototype model is to make sure the algorithms/solution methods work correctly before introducing larger trees and estimated parameters.

In this prototype model it is assumed that the standard deviation and probabilities do not change over time, i.e. they are not different between time steps (this is implemented in the main model). Furthermore the prototype model does only contain two time steps. This gives us the following two dimensional binomial tree:

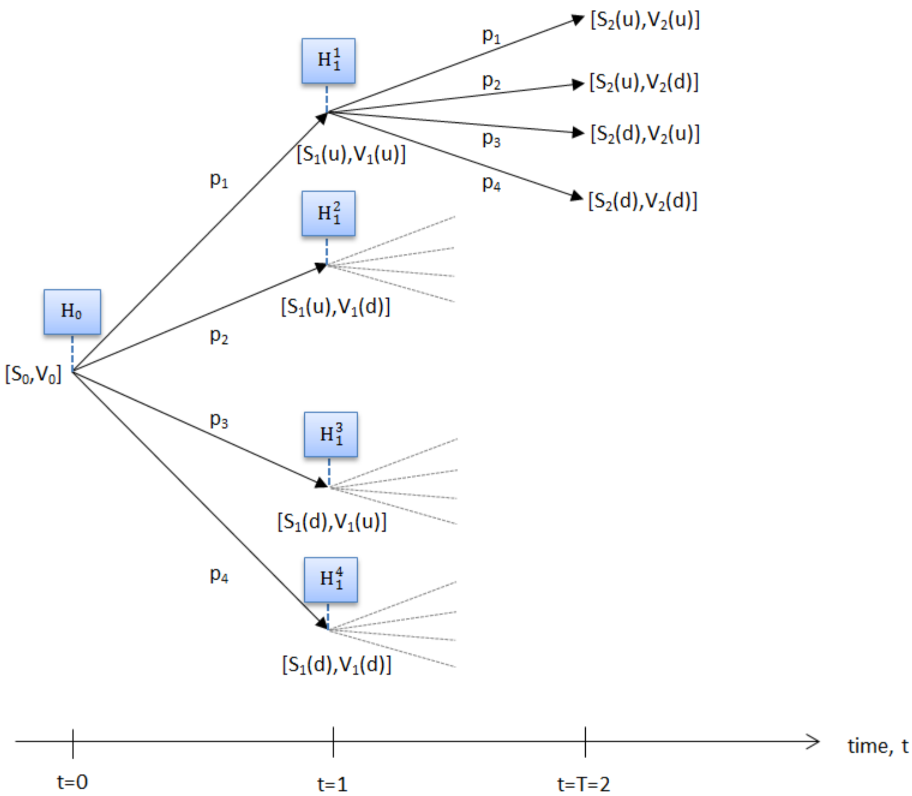

Figure 12: The prototype model. For the Quadratic Programming solution all hedge decisions at a given point in time are the same, i.e. $H_t^1 = \ldots = H_t^N$ where $N = 4^t$. Of course, hedge decisions at different times are most likely not equal.

6 Results

6.1 Log-normality of distributions

In Section 3.1.1 an assumption was made about the two assets following a log-normal distribution. In order to validate this assumption a QQ-plot for both the price data and the volume data is constructed.

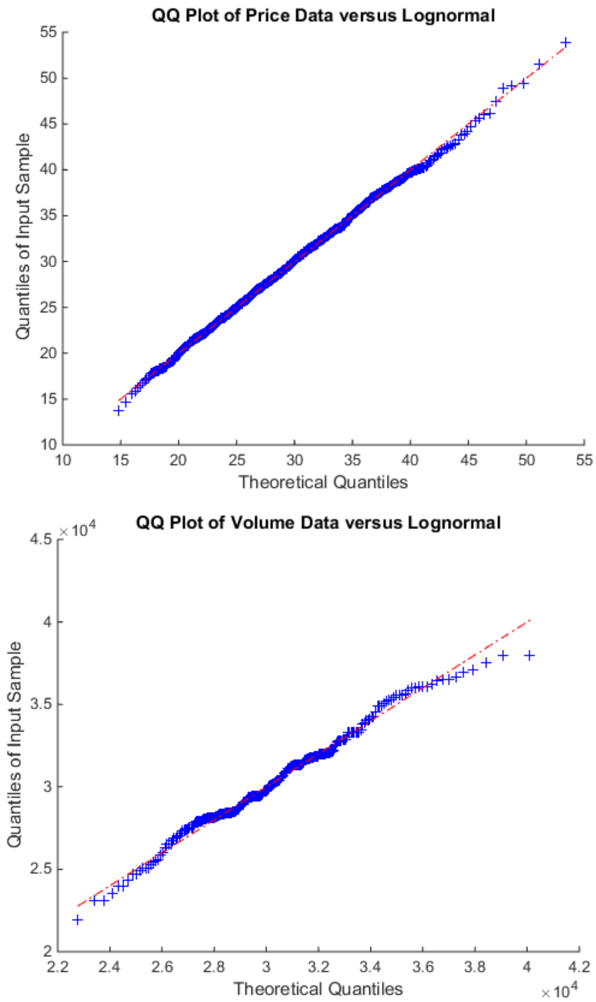

Figure 13: QQ-plot of the simulated price and volume data in our data sample versus the theoretical quantiles of a log-normal distribution.

The graphs show that both the price and volume data fits well with a log-normal distribution and that the assumptions of the assets' distributions were valid. The authors believe that the reader of this report is acquainted with the concepts of the quantile-quantile plot (QQ plot) and therefore leave out any further explanation.

6.2 The prototype model

As an illustrative example, the following parameters are used:

Table 1: Parameters in the simulation of the prototype model.

	Parameter	t_0	t_1	$t_2 = T$
Prototype Model	(S_0, V_0)	$(40, 100)$	–	–
	σ_S	0.1	0.1	–
	σ_V	0.1	0.1	–
	ρ	-0.5	-0.5	–
	Δt	1 year	1 year	–
	c_t	0.8	0.4	–

Hedging resolution used in BW: 0.5 MW. All other solutions have continuous resolution.
λ interval: [0, 0.1]
λ step size: 0.01

With the parameters as above the solutions of the QP-, the BW- and the BR-models were calculated and the results can be seen in Figure 14 and 15. In Figure 14 it is evident that both BW and BR are more effective than the QP strategy. In Figure 15 one can similarly see that the cost associated with the decrease of the standard deviation is less for the BW compared with the other two.

As a measure of effectiveness for the f-value we consider for any method m and given value of σ:

$$Effectiveness^{m}_{f,\sigma} = \frac{m_{f,\sigma}}{QP_{f,\sigma}} - 1 \quad for\ m \in \{BR, BW, Full\ Opt, Precom\}, \tag{34}$$

as a percentage number. Similarly we have that for the cost of hedging c

$$Effectiveness^{m}_{c,\sigma} = 1 - \frac{m_{c,\sigma}}{QP_{c,\sigma}} \quad for\ m \in \{BR, BW, Full\ Opt, Precom\}, \tag{35}$$

as a percentage number so that any positive percentage number always reflects a more effective method at that point. For the f-value if effectiveness is a positive (negative) number it means a certain percentage higher (lower) value than *QP*. For the cost of hedging if effectiveness is a positive (negative) number it means a certain percentage cheaper (more expensive) cost than *QP*.

There are tabulated values of the effectiveness of the methods for some σ:s both with respect to f and *cost*. However, exact percentage numbers of effectiveness for the f and *cost*-functions are here left out as it is not of major interest. More importantly the prototype model serves to develop our strategies and give the initial results to see if it can be of interest to apply this methods on a more sophisticated model setup, which it indeed does. The percentage numbers of effectiveness will instead be presented for the Main Model.

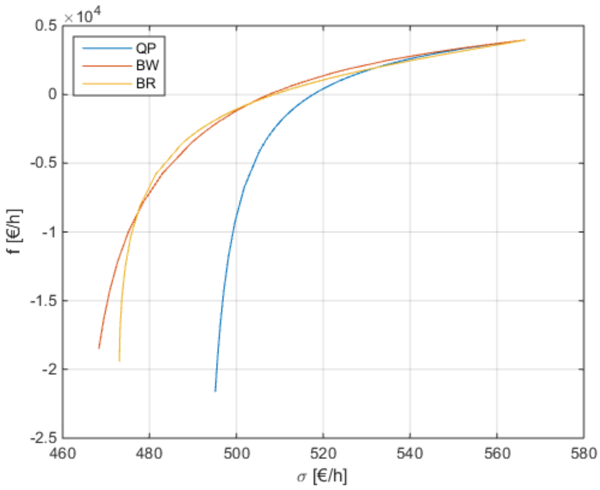

Figure 14: The object function plotted against the standard deviation as lambda varies.

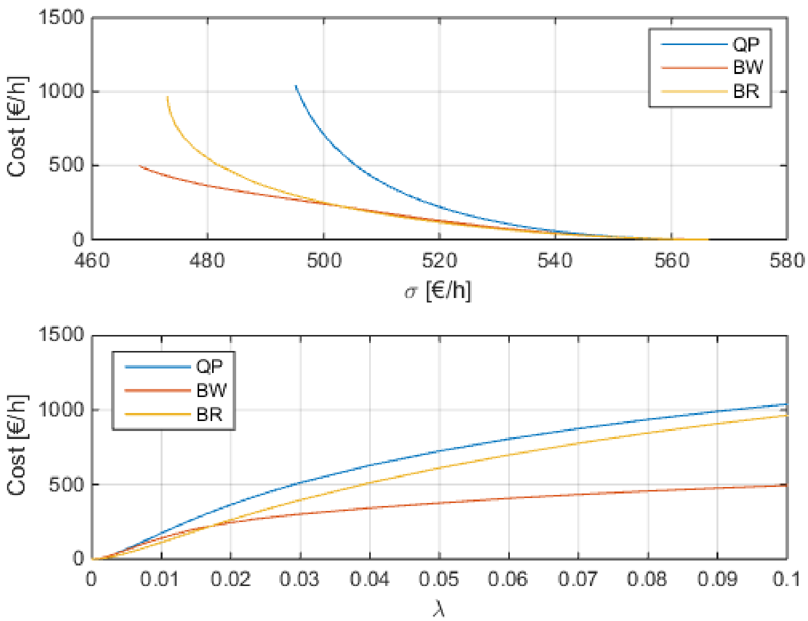

Figure 15: *Top:* Cost of hedging plots against standard deviation.
Bottom: Cost of hedging plotted against λ.

6.3 The main model

These results are obtained from the simulations. The simulations are done using the following parameters:

Table 2: Parameters in the simulation of the main model.

	Parameter	t_0	t_1	t_2	t_3	$t_4 = T$
Main Model	(S_0, V_0)	$(29, 3400)$	–	–		
	$\hat{\sigma_S}$	0.1290	0.1159	0.1382	0.0729	–
	$\hat{\sigma_V}$	0	0	0.0573	0.1076	–
	$\hat{\rho}$	0	0	-0.1	-0.445	–
	Δt	1 year	1 year	1 year	1 year	–
	cost $[10^{-3}]$	0.4833	0.4833	0.2417	0.1611	–

Hedging resolution used in BW: 8 MW. All other solutions have continuous resolution.
λ interval: $[0, 3]\cdot 10^{-5}$
λ step size: $0.01 \cdot 10^{-5}$ in $[0, 0.1]\cdot 10^{-5}$
$0.1 \cdot 10^{-5}$ in $[0.1, 3]\cdot 10^{-5}$
Notice that the correlation is becoming increasingly negative as time to delivery comes closer.
What is striking is the similarities between the reference solution QP and BW, Full Opt and to some extent BR in Figure 16. For all λ in our interval, the QP, BW and Full Opt show very similar results. The BR solution is worse from an optimization point of view since the higher f for the same σ the better. For higher λ the BR align with the other solutions except Precom. The f values also start to drop quite rapidly for larger λ, suggesting that the hedge is done almost as good as possible and no lower σ is reached once the trajectory is vertical.

One should bear in mind that the objective function f in Figure 16 does not show the actual cash flow of the company, but rather a tool for finding a good hedge strategy. The λ in front of the variance-term is a fictional penalty for taking risk. From a company's point of view, once this strategy is used, the actual revenue or the actual costs for the hedge strategy might be more practical.

The cost versus λ is created mainly of interest to see how much each solution spends on hedging for a given λ. Notice that the amount spent is not a direct indication on how much is hedged as the cost of hedging varies with time.

As f is constructed, the cost of hedging has to outweigh the penalty on risk from λ. Figure 20 shows the impact of higher risk aversion. The trees squeezes to reduce variance and lowers a bit in the beginning due to high cost from large amounts of hedging. For high λ at the end of the simulation most strategies start out with an initial hedge of

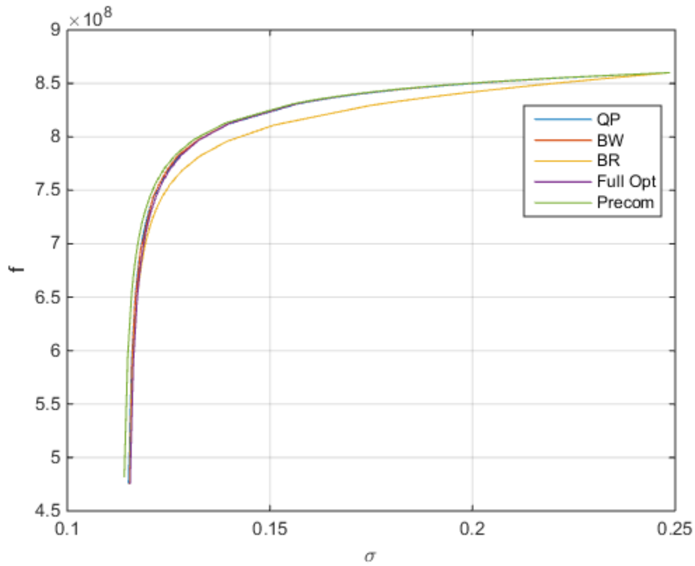

Figure 16: The object function plotted against the standard deviation as lambda varies.

Table 3: Effectiveness in terms of f of the main model.

Main Model

σ	QP f-value	BR diff [%]	BW diff [%]	Full Opt diff [%]	Precom diff [%]
0.22	$8.55 \cdot 10^8$	-0.56	0.02	0.02	0.02
0.20	$8.50 \cdot 10^8$	-0.94	0.03	0.03	1.10
0.17	$8.41 \cdot 10^8$	-1.42	0.03	0.03	0.06
0.15	$8.23 \cdot 10^8$	-1.66	0.07	0.03	0.14
0.13	$7.90 \cdot 10^8$	-2.01	0.12	-0.08	0.51
0.12	$7.28 \cdot 10^8$	-2.41	0.16	-1.06	1.58

almost the full amount of production.

The volume only has non-zero standard deviation the last two years before delivery. This is as earlier mentioned because when there is a long time to delivery, the hydro balance at delivery is expected to be normal, implying no change in the expectation of produced volume. Of course zero standard deviation in volume implies zero correlation between price and volume during that period. The revenues in Figure 20 are taken from the BW solution.

Figure 21 shows what hedge decisions are made over time and at what node for various λ, i.e. taking one decision at one node, that node will have four following nodes. In each of those nodes another decision will be made. The QP solution only gets to choose

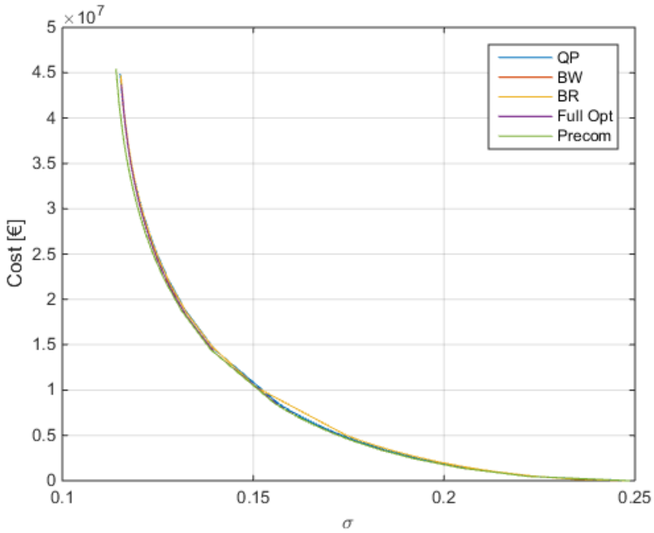

Figure 17: Cost of hedging plots against standard deviation.

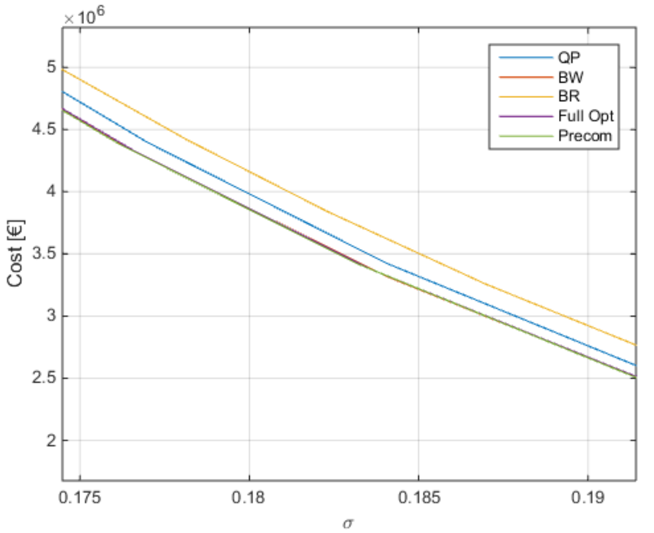

Figure 18: Zoomed, cost of hedging plots against standard deviation.

one hedge decision for all nodes at a certain time, why there is only one line in contrast to the other solutions. Also recall that if $\sigma_V = 0$ which it is for all years except the two closest to delivery, the four outcomes will reduce to two (unique) outcomes. This explains why one line at times only splits to two new lines instead of four. Note that for a tree spanning over four years, the decisions are made at times $t = \{0, 1, 2, 3\}$.

Figure 21 show these hedge decisions for three of the solutions; the reference solution QP,

Table 4: Effectiveness in terms of *cost* of the main model.

Main Model

σ	QP *cost*-value [€]	BR diff [%]	BW diff [%]	Full Opt diff [%]	Precom diff [%]
0.22	$6.08 \cdot 10^5$	-9.29	4.36	4.69	4.85
0.20	$1.84 \cdot 10^6$	-6.54	3.21	3.43	3.49
0.17	$4.72 \cdot 10^6$	-3.84	2.91	2.87	3.18
0.15	$1.09 \cdot 10^7$	2.06	2.69	2.50	3.12
0.13	$2.05 \cdot 10^7$	0.51	2.24	2.34	3.93
0.12	$3.13 \cdot 10^7$	0.15	1.58	1.55	5.07

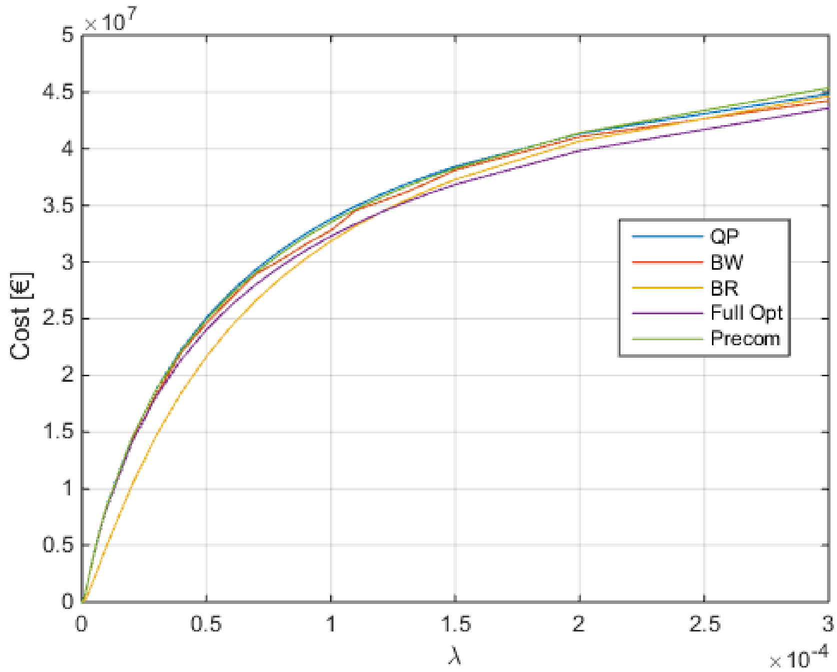

Figure 19: Cost of hedging plotted against λ.

the BW solution and the Precom solution which showed a slightly better performance in the earlier results. The figure shows a general pattern of quite heavy initial hedging from time $t = 0$ and then negative hedging at the final time step before delivery. This behaviour is especially pronounced for higher λ. The overall pattern between the solutions is similar even if λ is changed. One notable exception is for the Precom solution with high λ, where the negative hedging in the last period is greater than that in the other solutions. Recall that the initial volume $V_0 = 3400$ so a hedge of that amount would translate into 100% hedging of the total expected produced volume already in the first

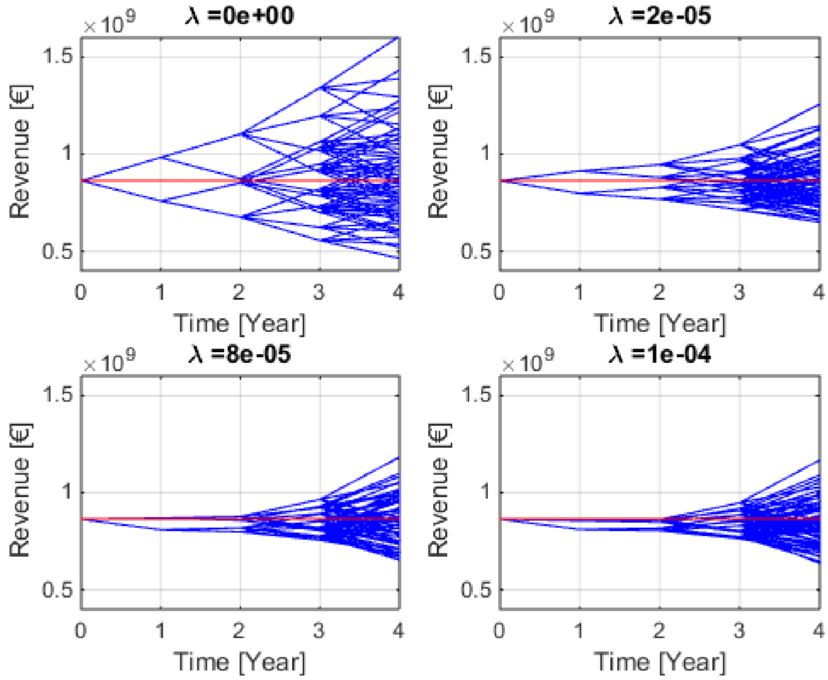

Figure 20: The evolution of all possible revenue paths as λ increase. The nodes in the tree are showing the possible revenue outcomes each year (integer t's). These are connected with lines to form paths and show every node's preceding and following nodes.

year.

Furthermore in Figure 22 the revenue is split into three separate parts - spot, hedge and cost - to see what sort of revenue streams that are expected as of $t = T$. Spot is the volume that is still free/unhedged at time of delivery T.

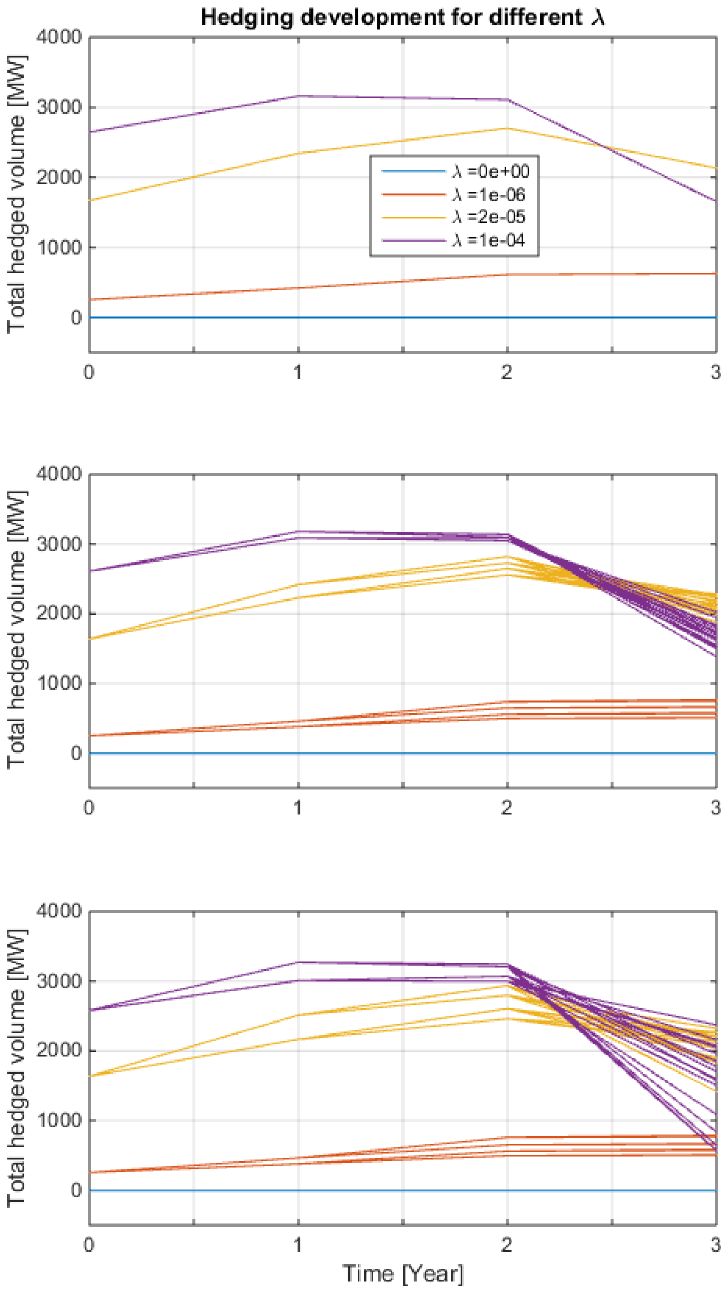

Figure 21: The cumulative amount of hedging done at time t by different solution models. The different colored lines reflect different λ.
Top graph: QP
Middle graph: BW
Bottom graph: Precom

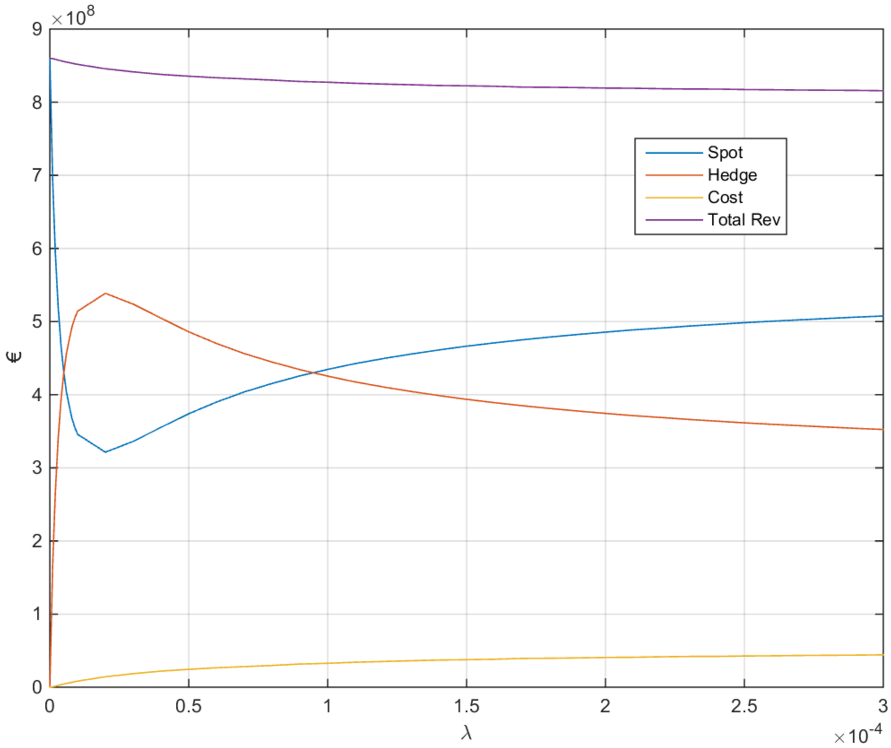

Figure 22: The (expected) revenue at $t = T$ split up by spot, hedge and cost. This graph is from the BW strategy. These are shown over different λ. As can be seen, when $\lambda = 0$ all income stems from spot and the cost is zero since no hedging has been done and consequently the income from hedging is also zero. As λ increases the amount of hedging initially rapidly increases and then slowly decreases from around $\lambda = 0.2 \cdot 10^{-4}$. This behavior can also be seen in Figure 21 where for higher λ, more negative hedging (buying back forwards) is done in especially the last time step. One interpretation of this is that the method exploits the increased negative correlation that exists closer to delivery.

6.3.1 Increasing steps

The volume volatility is zero for all but the last two years which might have an effect on the similarities of the different solution methods. Therefore an investigation of the last two years is made. Here there are six time steps over the last two years where $\sigma_V \neq 0$ and thus we have $\rho \neq 0$ as well as four separately identifiable outcomes. In this test Full Opt is left out. This is due to a number of reasons:

i) In the main results Full Opt was quite similar to QP and BW (this is also true for countless other simulations done but not presented here. If anything Full Opt is slightly below QP and BW in the f-graph.
ii) it will have worse f-outcomes evaluated at $t = 0$ than Precom since the hedges are changed from the optimal Precom ones.
iii) The main objective is to see if the differences compared to the main results increase with more time steps (i.e. a larger tree) and more steps with non-zero correlation and volume volatility.

The results from this simulation can be seen in Figure 23 and Figure 24.

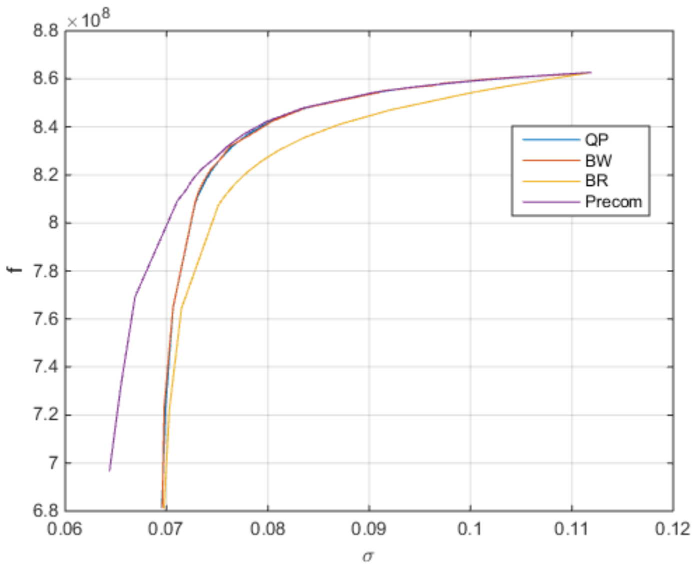

Figure 23: The object function plotted against the standard deviation as λ varies.

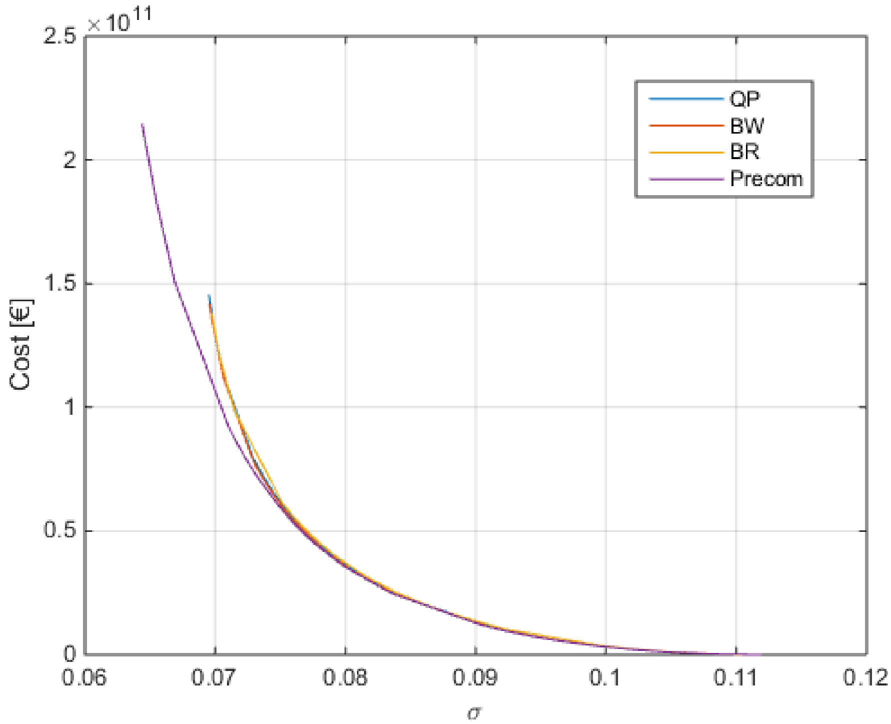

Figure 24: Cost of hedging plots against standard deviation.

6.3.2 Robustness of parameters

To know what impact a change in the parameters would have for the solutions the problem was simulated again with scaled parameters σ_s, σ_v, ρ and c. The parameters are one at the time multiplied by two or divided by two. The result of the robustness test is that the solutions change but the internal order remain intact. That is, the conclusions of the results remain even though all solutions squeeze together or get more spread out depending on which variable is changed and whether you increase or decrease it. In Figure 25 one can see that when the cost is smaller the solutions are closer together and when it is increased the solutions are getting more spread out.

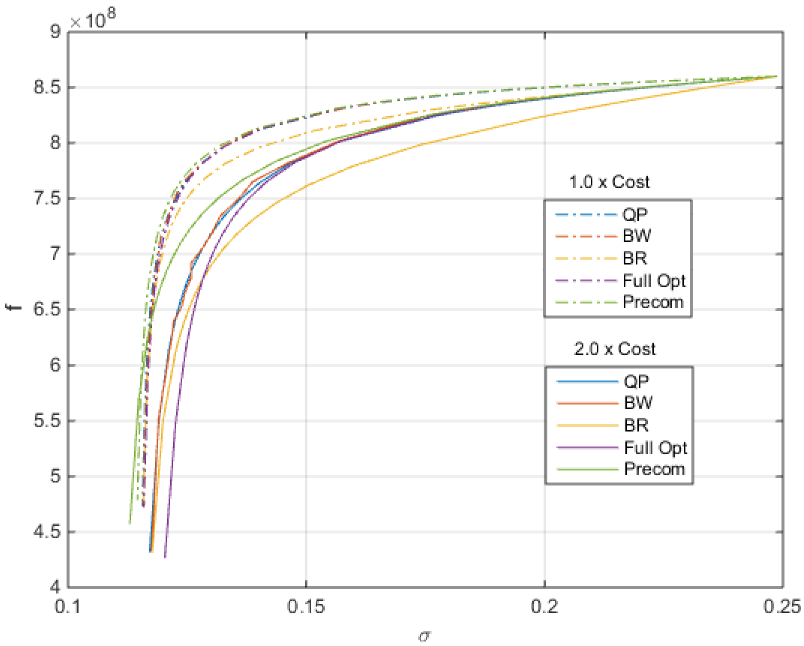

Figure 25: Robustness of the solution when multiplying the cost parameter with the factor 2. As can be seen, the internal order is to a large extent intact but the separation between them is more evident. The simulation with $0.5 \times \mathit{Cost}$ is left out due to visibility reasons.

7 Discussion

7.1 Solution methods

The simulated data obtained from Vattenfall contains price simulations for the four years 2015-2018 (see Section 4 for further details). Therefore the main focus has been to develop an as good model as possible during that time period. The main model covers all these four years with one time step each year. The different solution methods give quite similar results, as can be seen in Figure 16, which is a bit surprising but may be due to a number of reasons. First of all, even though the time span is four years, the tree in itself is relatively small with four time steps. An increase in number of steps might have enabled some of the solutions to differ more from the reference solution QP. This is something that seems logical considering the results from Section 6.3.1. In those results six time steps are taken over the last two years. Even though we would like to run simulations for larger trees, the non recombining trees made it too computationally intense to run these simulations and the methods are not fully optimized for computational performance. This will be further discussed in Section 9.

Another reason for the similarities of the results might be found in the zero volume volatilites. As the time of delivery is very far away in time (over two years) the expected produced volume does not change. This results in zero volume volatility for all years but the two closest years to delivery. This leads to only two distinctly different solutions which once again might improve the QP solution relative to the others since the number of unique nodes to take into consideration only becomes two rather than four. We remind once again that the QP solution can only choose *one* hedge decision at each time step, rather than one hedge decision for each node in the given time step. It is also implied that more outcomes/number of nodes might increase the relative performance of solutions such as BW, Full Opt and Precom compared to QP. The non-zero volatilities for the two years closest to delivery was also a reason for letting the six time step model run over those exact years.

However, even though the minimum standard deviation only might be decreased with roughly 1% the cost could be reduced with approximately 1.5% - 5% for a given σ this is still a meaningful improvement. The generation of power from hydro power plants is a significant part of Vattenfall's generation and will therefore have an impact on their bottom line result. The lower cost for same σ of course also means that for the same amount of cash, a risk reduction in terms of volatility of revenue can be achieved. With a large amount of hydro power being more efficiently hedged, even a small improvement of the hedging strategy can lead to a large improvement for the company. Furthermore since the Nordic hydro constitutes such a big part of Vattenfall's portfolio and the hedging of hydro production is more difficult compared to other assets, groups within Vattenfall might reason that it is not worth the effort to decrease other assets' risks as it is so small compared to the difficult-to-hedge hydro power. This means that if the hydro power hedging can improve, it could spread ripples to other groups within Vattenfall.

The main reason for the difficulty to hedge this due to the uncertainty in production. Considering Figure 16 one can see that the f-value almost starts to drop down vertically for higher λ. The cost and λ are trying to balance each other - as the penalty of volatility increases, it is more worthwhile to spend a higher cost on hedging. However when the f-value drops down and little to no more hedging is done, it seems like the solution methods have come to a point where they have hedged as good as they can. Increasing λ only pulls down the f-value without any notable decrease in volatility, yet the volatility is not close to zero. This might be a reflection of the difficulties that appear when facing uncertain production.

7.2 Hedging

In Figure 21 for $\lambda = 0$ no hedge was done in any solution model which is totally in line with what is expected. If no penalty for taking risk/having variance exists but there is a cost to reduce risk/variance, no such measures will be taken. All of the solution models behaves like this and it is also the reason why they start our at the same initial point in e.g. Figure 16. Also note that the solution methods are allowed to find any hedge that will maximize the objective function, even if it is a negative one. A negative hedge would mean that the company essentially goes long electricity price, and for a company producing electricity it is already long by default. If this is done, or will be done, in practice is not clear.

Figure 21 also shows that initially quite a lot hedging can be done to offset the penalty from high variance even though the cost is high. This is probably a direct response to the objective function being defined as it is. The hedging done *before* $t = 0$ is assumed to be zero, which means that the f-function might be far from maximized at the initial point. Then as one time step is taken, the method only has to account for the changes. This results in relatively finer adjustments until the last decision where again a lot of adjustments are made. It is interesting that the solutions chooses to unhedge those amounts in the last step but the cost of hedging is lower closer to delivery due to higher liquidity, which might be one explanation to a larger change occurring in this step. Another explanation is probably that the methods exploit the negative correlation that is more pronounced closer to delivery.

7.3 Optimality

All solutions have had the objective function f evaluated in $t = 0$ to be comparable. However, the way the respective solution method has obtained this f-value is different. Some solutions use a recursive, iterative approach while others maximize the f-value as seen from $t = 0$. The BR, BW and Full Opt belong to the iterative methods while QP and Precom optimize as seen from $t = 0$.

Comparing the BR and BW solutions, they use different approaches to find best responses. BR sees only one step ahead, but knows exactly what decisions have been made previously in time. BW predicts later players decisions via the utility functions and chooses a best response with respect to those decisions and its own utility function. The big advantage of BW is that all players in the time steps make an optimal decision taking all later (predicted) decisions into account. In that respect BR is kept a bit in the dark which might be a reason for the slightly worse results that can be seen in Figure 16. Due to the way they find their solutions it is reasonable to believe that this relationship of BW outperforming BR is likely to hold over time. For all the BW player at time $t = 0$ knows, he/she might still be a part of a subgame but yet manages to outperform BR.

Considering the Precom method, it is allowed to individually change all hedge decisions to make it optimal from a $t = 0$ point of view. This is of course difficult to beat since it is allowed to choose a continuous (i.e. non-integer) volume to hedge in each node that requires a decision. This method can fully focus on optimizing the objective function as good as possible today. The other methods are iterative except QP, but QP is restrained in the sense of only being able to make one hedge decision that should be applied to all outcomes at a specific time. The question then arises what is optimal? Is it optimal today or over time? How should one consider optimality over time?

Comparing our solutions the one which have the lowest cost given a standard deviation is the Precom strategy, even though it's only slightly better than BW. However, this strategy requires that you stick with the hedging decisions decided today, even though more information will be available as time goes forward. What might sound more intuitive would be to have an alternative that tries to take the best decision in every part of the tree. Such a method would when time passes have a strategy that could take into account all the decisions and in the past, as well as the current information today and the knowledge of probable outcomes going forward. This would be the big advantage of the SDP method *if* it was usable. Then one will know that the optimal decision policy in each node at each time would coincide with the optimal decision policy for the entire problem at $t = 0$. The strategy closest to the SDP method is the game theoretic approach BW.

The BW solution evaluated today gives very similar f-values as the QP method. In that regard it is not worse or better evaluated at $t = 0$. But it does still have different hedging decisions at each node in a specific time step, adapted to the outcome in that node. It also has made the hedging decisions by predicting later decisions in time. It seems reasonable to believe that it would fare better over time and therefore there is no reason for not using it over QP. According to our results, the strategy can also achieve the same amount of risk reduction at a lower price. How BW fares compared to Full Opt is interesting. Both change decisions over time to get optimal subtrees (at the cost of optimality as of $t = 0$, which can e.g. be seen by comparing Precom and Full Opt) and are showing quite similar results. One difference might be that BW responds optimally to predicted later decisions, while Full Opt calibrates the decisions today optimally with regard to later decisions which will be changed as the iteration method progresses through the tree. This is again the same trade off of optimality as is seen between Precom and Full Opt at time $t = 0$. Ceteris paribus, it would be foolish to choose a sub optimal method at $t = 0$, but one should probably take into consideration what would happen if the same graph as Figure 16 was to be produced at a later time, i.e. what constitutes optimality over time. We believe that the Full Opt can be a evaluation today over how Precom actually performs over time, especially if one do not fully commit to all the years but rather is re-run the Precom strategy after a year. In that case BW would according to results and logic be the preferred strategy.

7.4 Computational performance

For practical reasons it is also an obvious drawback if the solution method is very numerically intense. How big of a drawback depends of course on how often it has to be run and what time/computer power is available at the company in question. Some solutions, namely Precom, Full Opt and BW might be hard to use for very large trees in their current form. Matlab's function *fminunc* which is used when calculating the Precom solution needed more function evaluations and iterations than default options, already for quite small trees of three or four time steps. Since Full Opt is an iterated version of Precom this of course also applies to the Full Opt solution. When it comes to BW it is more a question of resolution when adding the information (see Section 5.6.2). As the tree becomes larger and time of delivery is further away, it is preferable to increase the resolution to prevent a too jagged solution graph. It should be mentioned that there might be significant performance improvements to be made by writing more efficient code/search algorithms or similar. This might e.g. be by concentrating on fewer λ, having more efficient search algorithm's when looking for an optimal solution etc. This has been implemented to some extent by us when writing the code and choosing methods, but since it is not a main objective of this thesis there is probably room for improvement.

8 Conclusions

8.1 Effectiveness of hedging

It is possible to hedge more effectively using the Precom, BW or Full Opt method rather than QP which is done today. Even though the minimum risk only could be lowered with roughly 1% in terms of standard deviation, a cost reduction of 1.5 % - 5 % is according to the results possible (although they are of course simulations and no guarantee for real world results). The difference is somewhat lower than what was expected after seeing the initial results from the prototype model. However, for other similar problems with different parameters and inputs, the prototype model shows that an alternative to the reference solution QP might bring substantial effectiveness gains. In the case presented here there might also be a bigger difference when simulating more time steps in the tree. Recall that the Precom method is allowed to make a hedge decision in each node and so it will always outperform QP which only can make a hedge decision at each time step. This is because the QP solution is a subset of the Precom solution. The difference between these methods increases as the amount of decisions in each time step increases which can be seen in Figure 23. It is likely that both factors contribute to the difference.

How to think about optimality is important and it might be up to the preference of the end user. The iterative methods BR, BW and Full Opt take some of this into account, and based on the reasoning in Section 7.3 BW seems according to us be the most viable of those three in that regard. In the same section is it discussed that BW seems to be at least at par with QP looking at the objective function f. It is reasonable to believe that it will stay at least at par or improve over time as it is more dynamic. Furthermore BW shows cost reductions of 1.5 %- 5% depending on risk level. This of course also translates into lowering risk for the same amount of cash spent. It should therefore be possible to hedge more effectively than today using the BW method rather than QP. The core problem that brings forth the question about optimality is that none of the methods can guarantee a fully optimal solution over all nodes and times in the true meaning of the word, which a SDP solution would do. How to value a dynamic method such as BW against a static one from an $t = 0$ point of view is a more difficult question. This is especially true if the static one shows significant cost/risk improvements and this might again be up to the end user or a given company's strategic preference. However the results from Precom and BW are quite similar and that speaks in favor for the BW method. Viewing Full Opt as a dynamic case of Precom then BW would be preferred as it shows slightly better results than Full Opt. Furthermore, from our point of view the logic behind the BW strategy is more satisfying.

8.2 Parameters

The parameters going into the model are important and therefore it would be advisable to make a thorough study of the models being accurate to reality. It is our belief that Vattenfall has sophisticated models for simulating price and volume, as well as reliable and long term data. Robustness tests have been carried out and the solutions for the different models will not change order, i.e. the one which gives the best result will still be the best one. However they will squeeze together or spread out when the parameters are changed. It should also be mentioned that as parameters gradually were adjusted to real world (or simulated real world) data, the solutions became more robust to individual changes of any parameter. That is, changing volatilities seemed to have a bigger impact on the solution when the correlation and cost were set as in the prototype model, than changing the volatility for the main model's parameters. Still, if e.g. costs are increased, the difference in the solutions will be more evident and the choice of strategy more important. This is of relevance since just a slight efficiency improvement is interesting for Vattenfall.

8.3 Computational intensity

Some solution methods are more numerically intense than others. BR and QP are calculated very quickly relative to the other strategies Precom, Full Opt or BW. Therefore performance enhancing measures might be needed. The perhaps most time consuming simulation is the BW solution where information has to be added in the nodes of the tree in order to solve the problem. More effective methods of adding information and look for optimal solutions might improve the performance of this strategy.

9 Suggestions for further development

9.1 Recombining trees

As the binomial tree is built, each node splits into four new nodes. Quite quickly the number of calculations become large as the tree grows. This has implications for practical purposes and a suggested performance enhancing measure is to build recombining trees. To better see the advantage of recombining trees compared to non-recombining trees we will consider how the number of nodes grow with the tree size. Let n be the number of outcomes from one node (i.e. n=4 in this paper), N the amount of nodes at time step t, where $t \in \{0, 1, 2, ...\}$. Then the tree follows

$$N = n^t$$

and the total amount of nodes in the entire tree N_{tot} will be

$$N_{tot} = \sum_{t=0}^{timesteps} n^t$$

where *timesteps* is the number of time steps in the tree. With recombining nodes the two-dimensional binomial tree will follow

$$N_{recomb} = (t+1)^2$$

or similarly

$$N_{recomb,tot} = \sum_{t=0}^{timesteps} (t+1)^2$$

We have tabulated some values in Table 5 where the advantage of a recombining tree becomes obvious.

Table 5: Comparison between a recombining and non-recombining two dimensional binomial tree for $n = 4$.

Time, t	0	1	2	3	4	5	10	20
$\mathbf{N_{tot}}$	1	5	21	85	341	1 365	1 398 101	$1.47 \cdot 10^{12}$
$\mathbf{N_{recomb,tot}}$	1	5	14	30	55	91	506	3 311
$(\mathbf{N_{recomb,tot}/N_{tot}})$	1	1	0.67	0.35	0.16	0.07	0.0004	$2.26 \cdot 10^{-9}$

The problem that arises when confronted with real world data and estimating volatilities is that they change over time. Thus the tree will not generally be recombining as if one had constant volatility. The pattern of the two dimensional recombining tree can be seen easily when visualizing the first steps of a recombining binomial tree as a lattice in Figure 26, Figure 27 and Figure 28.

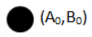

Figure 26: Initial point at $t = 0$ in a recombining binomial tree. The evolution below would be the same for an arbitrary point.

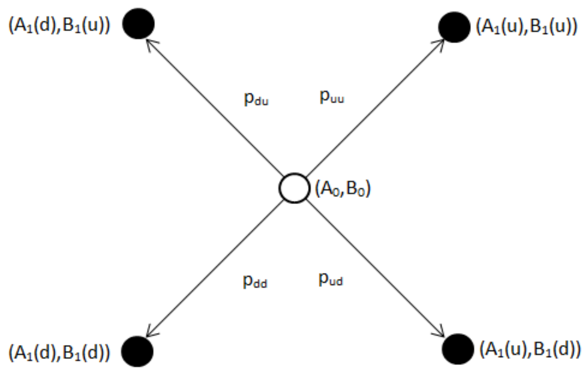

Figure 27: Nodes in a recombining tree from $t = 0$ (white) to $t = 1$ (black) visualized as a lattice.

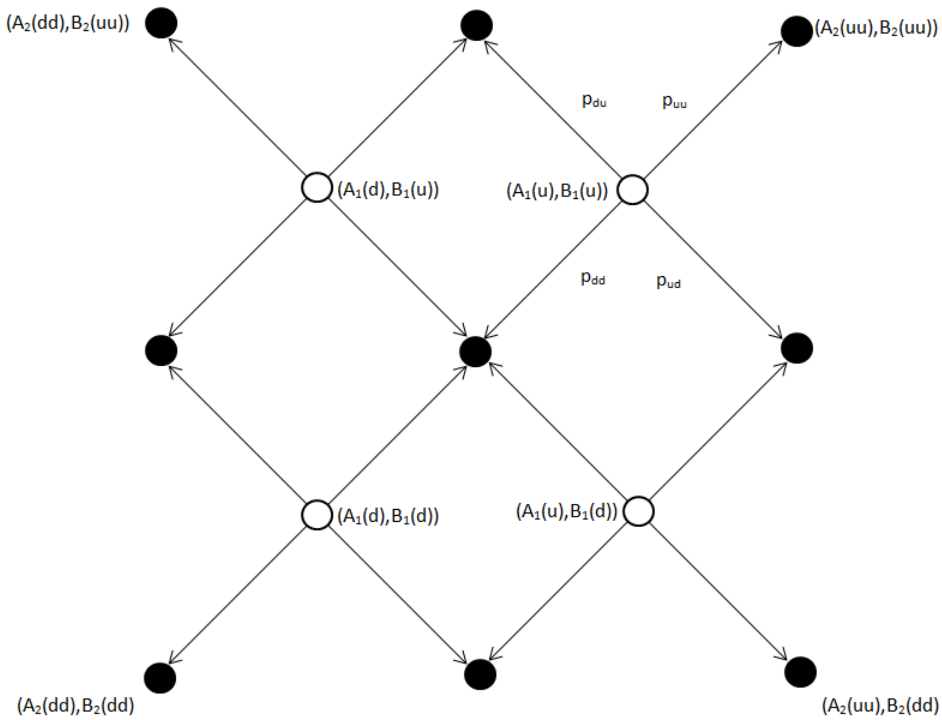

Figure 28: Nodes in a recombining tree from $t = 1$ (white) to $t = 2$ (black) visualized as a lattice.

9.2 Intra year

This thesis covers hedging strategies up to time of delivery, T. However there might of course be important to have a hedging strategy during year of delivery too. This model cannot instantly be used to model that scenario without some tweaks. One thing that needs to be taken into consideration is the fact that after e.g. six months, half the power would have been delivered and the volatility of the remaining power to be delivered therefore has to be adjusted. Furthermore there might be intra year patterns that has to be modelled in order to have an efficient hedging strategy.

9.3 Limiting downside

This approach is based on the rhetorical question "Do I have the same risk aversion regardless of how my portfolio performs?". Probably not. In other cases the company itself might have some requirements of not hitting below a certain revenue/profit level for various purposes. In both these scenarios the willingness to hedge might increase if the revenues fall. If there are some margin left, then one might want to be fully hedged even at a high cost if that would mean zero risk of getting below the profit level. To put it in the context of this thesis, if one would enter e.g. a state where both price and volume falls compared to the previous state, λ would increase. This would reflect an incentive to limit downside. Similarly if it would be an up state, one might feel that there is well enough distance from any target minimum revenue to not hedge in the same way. This would be an interesting extension to our model - alter λ in every node in the tree in order to capture varying risk aversion depending on portfolio performance.

A Appendices

A.1 Derivation of objective function for quadratic optimization

Derivation of objective function from the revenue function (10).

$$
\begin{aligned}
R_T =& V_T S_T - \sum_{t=0}^{T-1} H_t (S_T - S_t) - \sum_{t=0}^{T-1} c_t H_t^2 \\
=& V_T S_T - \sum_{t=0}^{T-1} H_t S_T + \sum_{t=0}^{T-1} H_t S_t - \sum_{t=0}^{T-1} c_t H_t^2
\end{aligned}
$$

$$
\begin{aligned}
R_T^2 =& [V_T S_T - \sum_{t=0}^{T-1} H_t S_T + \sum_{t=0}^{T-1} H_t S_t - \sum_{t=0}^{T-1} c_t H_t^2]^2 \\
=& (V_T S_T)^2 + (\sum_{t=0}^{T-1} H_t S_T)^2 + (\sum_{t=0}^{T-1} H_t S_t)^2 + (\sum_{t=0}^{T-1} c_t H_t^2)^2 - 2V_T S_T^2 \sum_{t=0}^{T-1} H_t \\
&+ 2V_T S_T \sum_{t=0}^{T-1} H_t S_t - 2V_T S_T \sum_{t=0}^{T-1} c_t H_t^2 - 2 \sum_{t=0}^{T-1} H_t S_T \sum_{t=0}^{T-1} H_t S_t \\
&+ 2 \sum_{t=0}^{T-1} H_t S_T \sum_{t=0}^{T-1} c_t H_t^2 - 2 \sum_{t=0}^{T-1} H_t S_t \sum_{t=0}^{T-1} c_t H_t^2
\end{aligned}
$$

$$
\begin{aligned}
E[R_T^2] =& E[S_T^2 V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] + (\sum_{t=0}^{T-1} c_t H_t^2)^2 \\
&- 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] - 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 \\
&- 2E[\sum_{t=0}^{T-1} H_t S_T \sum_{t=0}^{T-1} H_t S_t] + 2E[S_T] \sum_{t=0}^{T-1} H_t \sum_{t=0}^{T-1} c_t H_t^2 \\
&- 2 \sum_{t=0}^{T-1} H_t E[S_t] \sum_{t=0}^{T-1} c_t H_t^2
\end{aligned}
$$

Under the assumption that no arbitrage exists we have that $E[S_t] = S_0, \quad t \in [0, T]$

$$
\begin{aligned}
E[R_T^2] =& E[S_T^2V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] + (\sum_{t=0}^{T-1} c_t H_t^2)^2 \\
& - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] - 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 \\
& - 2E\Big[\Big(\sum_{t=0}^{T-1} H_t S_T\Big)\Big(\sum_{t=0}^{T-1} H_t S_t\Big)\Big] + 2S_0 \sum_{t=0}^{T-1} H_t \sum_{t=0}^{T-1} c_t H_t^2 \\
& - 2S_0 \Big(\sum_{t=0}^{T-1} H_t\Big)\Big(\sum_{t=0}^{T-1} c_t H_t^2\Big) \\
=& E[S_T^2V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] + (\sum_{t=0}^{T-1} c_t H_t^2)^2 \\
& - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] - 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 \\
& - 2E\Big[\Big(\sum_{t=0}^{T-1} H_t S_T\Big)\Big(\sum_{t=0}^{T-1} H_t S_t\Big)\Big]
\end{aligned}
$$

$$
\begin{aligned}
E[R_T] =& E[V_T S_T] - \sum_{t=0}^{T-1} H_t E[S_T] + \sum_{t=0}^{T-1} H_t E[S_t] - \sum_{t=0}^{T-1} c_t H_t^2 \\
=& E[V_T S_T] - \sum_{t=0}^{T-1} c_t H_t^2
\end{aligned}
$$

$$
(E[R_T])^2 = (E[V_T S_T] - \sum_{t=0}^{T-1} c_t H_t^2)^2 = (E[V_T S_T])^2 - 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 + (\sum_{t=0}^{T-1} c_t H_t^2)^2
$$

$$
Var[R_T] = E[R_T^2] - (E[R_T])^2
$$

$$
\begin{aligned}
Var[R_T] =& E[S_T^2 V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] + (\sum_{t=0}^{T-1} c_t H_t^2)^2 \\
& - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] - 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 \\
& - 2E\Big[\Big(\sum_{t=0}^{T-1} H_t S_T\Big)\Big(\sum_{t=0}^{T-1} H_t S_t\Big)\Big] - (E[V_T S_T])^2 + 2E[V_T S_T] \sum_{t=0}^{T-1} c_t H_t^2 - (\sum_{t=0}^{T-1} c_t H_t^2)^2 \\
=& E[S_T^2 V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] \\
& - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] \\
& - 2E\Big[\Big(\sum_{t=0}^{T-1} H_t S_T\Big)\Big(\sum_{t=0}^{T-1} H_t S_t\Big)\Big] - (E[V_T S_T])^2
\end{aligned}
$$

$$
\begin{aligned}
f =& E[R_T] - \lambda Var[R_T] = E[S_T V_T] - \sum_{t=0}^{T-1} c_t H_t^2 \\
& - \lambda (E[S_T^2 V_T^2] - (\sum_{t=0}^{T-1} H_t)^2 E[S_T^2] + E[(\sum_{t=0}^{T-1} H_t S_t)^2] \\
& - 2E[V_T S_T^2] \sum_{t=0}^{T-1} H_t + 2E[V_T S_T \sum_{t=0}^{T-1} H_t S_t] \\
& - 2E\Big[\Big(\sum_{t=0}^{T-1} H_t S_T\Big)\Big(\sum_{t=0}^{T-1} H_t S_t\Big)\Big] - (E[V_T S_T])^2)
\end{aligned}
$$

The objective function is and will be used for the quadratic programming.

A.2 Derivation of objective function for BW

Using the notations in Section 5.1, we have the following

Revenue:

$$\tilde{R}_t = -c_t H_t^2 - (X_{t-1} + H_t)(S_{t+1} - S_t) + \tilde{R}^*_{t+1}(X_{t-1} + H_t) \tag{36}$$

Expected value of revenue:

$$\begin{aligned} E[\tilde{R}_t \mid \mathcal{F}_t] &= -c_t H_t^2 - (X_{t-1} + H_t) E[(S_{t+1} - S_t) \mid \mathcal{F}_t] + E[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{F}_t] \\ &= -c_t H_t^2 + E[E[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}] \mid \mathcal{F}_t] \\ &= -c_t H_t^2 + \sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{G}^i_{t+1}] \end{aligned} \tag{37}$$

Expected value of revenue squared:

$$
\begin{aligned}
E[\tilde{R}_t^2 \mid \mathcal{F}_t] =& E[(-c_t H_t^2 - (X_{t-1}+H_t)(S_{t+1}-S_t) + \tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{F}_t] \\
=& E[(-c_t H_t^2)^2 - ((X_{t-1}+H_t)(S_{t+1}-S_t))^2 + (\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \\
&+ 2c_t H_t^2 (X_{t-1}+H_t)(S_{t+1}-S_t) - 2c_t H_t^2 \tilde{R}^*_{t+1}(X_{t-1}+H_t) \\
&- 2(X_{t-1}+H_t)(S_{t+1}-S_t)\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
=& c_t^2 H_t^4 + (X_{t-1}+H_t)^2 (S_t^2 - 2S_t E[S_{t+1} \mid \mathcal{F}_t] + E[S_{t+1}^2 \mid \mathcal{F}_t]) \\
&+ E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{F}_t] - 2c_t H_t^2 E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
&- 2(X_{t-1}+H_t)E[(S_{t+1}-S_t)\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
=& c_t^2 H_t^4 + (X_{t-1}+H_t)^2 (S_t^2 - 2S_t^2 + E[S_{t+1}^2 \mid \mathcal{F}_t]) \\
&+ \sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{G}^i_{t+1}] - 2c_t H_t^2 E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i E[(S_{t+1}-S_t)\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
=& c_t^2 H_t^4 + (X_{t-1}+H_t)^2 (E[S_{t+1}^2 \mid \mathcal{F}_t] - E[S_{t+1} \mid \mathcal{F}_t]^2) \\
&+ \sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{G}^i_{t+1}] - 2c_t H_t^2 E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
=& c_t^2 H_t^4 + (X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] + \sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{G}^i_{t+1}] \\
&- 2c_t H_t^2 E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]
\end{aligned}
$$

Squared expected value of the revenue:

$$
\begin{aligned}
(E[\tilde{R}_t \mid \mathcal{F}_t])^2 =& (-c_t H_t^2 + \sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2 \\
=& c_t^2 H_t^4 + (\sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2 \\
&- 2c_t H_t^2 \sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]
\end{aligned}
$$

Variance of the revenue:

$$
\begin{aligned}
Var[\tilde{R}_t \mid \mathcal{F}_t] = E[\tilde{R}_t^2 \mid \mathcal{F}_t] - (E[R_t \mid \mathcal{F}_t])^2 = (X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] \\
&+ \sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t))^2 \mid \mathcal{G}^i_{t+1}] \\
&- (\sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2 \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
=&(X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] \\
&+ \sum_{i=1}^{4} p_i \{Var[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] + (E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2\} \\
&- (\sum_{i=1}^{4} p_i E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2 \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
=&(X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] + \sum_{i=1}^{4} p_i Var[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
&+ \sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]^2 \\
&- (\sum_{i=1}^{4} p_i E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}])^2 \\
&- 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) E[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
=&\{Y^*_{t+1}(X_{t-1}+H_t) = E[(\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]\} = \\
=&(X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] + \sum_{i=1}^{4} p_i Var[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}] \\
&+ Var[Y^*_{t+1}(X_{t-1}+H_t)] - 2(X_{t-1}+H_t)\sum_{i=1}^{4} p_i (S_{t+1|i}-S_t) Y^*_{t+1|i}(X_{t-1}+H_t) =
\end{aligned}
$$

$$=(X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] + \sum_{i=1}^{4} p_i Var[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]$$

$$+Var[Y^*_{t+1}(X_{t-1}+H_t)] - 2(X_{t-1}+H_t)(\sum_{i=1}^{4} p_i(S_{t+1|i}-S_t)(Y^*_{t+1|i}(X_{t-1}+H_t)+E[Y^*_{t+1|i} \mid \mathcal{F}_t])$$

$$=(X_{t-1}+H_t)^2 Var[S_{t+1} \mid \mathcal{F}_t] + \sum_{i=1}^{4} p_i Var[\tilde{R}^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{G}^i_{t+1}]$$
$$+ Var[Y^*_{t+1}(X_{t-1}+H_t)] - 2(X_{t-1}+H_t)Cov[S_{t+1}, Y^*_{t+1}(X_{t-1}+H_t) \mid \mathcal{F}_t] \tag{38}$$

A.2.1 Special case: $\mathbf{t = T - 1}$

When $t = T-1$, then

$$\tilde{R}^*_{t+1}(X_{t-1}+H_t) = \{t = T-1\} = \tilde{R}^*_T = S_T V_T \tag{39}$$

similarly

$$E[\tilde{R}^*_T(X_{t-1}+H_t) \mid S_{T|i}, V_{T|i}] = S_{T|i}V_{T|i} \tag{40}$$

since the outcomes are known and does not change with different X or H anymore.

Important! When the following calculations are made, t will sometimes be kept as subscript instead of changing it to $T-1$ to facilitate comparisons with and replacements in terms in earlier derivations. However, going forward in this section the subscript t will be equivalent to $T-1$.

Equation (36) becomes

$$\tilde{R}_t = -c_t H_t^2 - (X_{t-1}+H_t)(S_{t+1}-S_t) + S_T V_T \tag{41}$$

and

$$E[\tilde{R}_t \mid \mathcal{F}_t] = -c_t H_t^2 + \sum_{i=1}^{4} p_i S_{T|i} V_{T|i} \tag{42}$$

with equation (40) we now have

$$E[\tilde{R}^*_T(X_{t-1}+H_t) \mid S_{T|i}, V_{T|i}] = S_{T|i}V_{T|i}$$
$$Var[\tilde{R}^*_T(X_{t-1}+H_t) \mid S_{T|i}, V_{T|i}] = 0$$

then

$$\begin{aligned} Var[Y^*_{t+1}(X_{t-1}+H_t)] &= Var[E[\tilde{R}^*_T(X_{t-1}+H_t) \mid S_{T|i}, V_{T|i}] \\ &= Var[S_{T|i}V_{T|i}] = Var[S_T V_T \mid \mathcal{F}_{T-1}] \end{aligned}$$

will be the variance over the four possible spot revenues reachable for the corresponding state i at time $t = T - 1$. Similarly the covariance becomes

$$Cov[S_{t+1}, Y^*_{t+1}(X_{t-1} + H_t) \mid \mathcal{F}_t] = Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}]$$

We thus have that for $t = T - 1$, the variance of the revenue:

$$\begin{aligned} Var[\tilde{R}_{T-1} \mid \mathcal{F}_{T-1}] = (X_{T-2} + H_{T-1})^2 Var[S_T, \mid \mathcal{F}_{T-1}] + Var[S_T V_T \mid \mathcal{F}_{T-1}] \\ - 2(X_{T-2} + H_{T-1}) Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}] \end{aligned} \tag{43}$$

and the objective function f becomes:

$$\begin{aligned} f = E[\tilde{R}_{T-1} \mid \mathcal{F}_{T-1}] - \lambda(Var[\tilde{R}_{T-1} \mid \mathcal{F}_{T-1}]) = -c_{T-1} H^2_{T-1} + \sum_{i=1}^{4} p_i S_{T|i} V_{T|i} \\ - \lambda\Big((X_{T-2} + H_{T-1})^2 Var[S_T \mid \mathcal{F}_{T-1}] + Var[S_T V_T \mid \mathcal{F}_{T-1}] \\ - 2(X_{T-2} + H_{T-1}) Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}]\Big) \end{aligned}$$

Similarity to minimum variance in last time step

We want to maximize f changing H_{T-1} so we calculate $f'(H_{T-1}) = 0$ and solve for H_{T-1}.

$$\begin{aligned} f'(H_{T-1}) = -2c_{T-1} H_{T-1} - \lambda\Big(2(X_{T-2} + H_{T-1}) Var[S_T, \mid \mathcal{F}_{T-1}] \\ - 2Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}]\Big) \end{aligned}$$

Setting $f'(H_{T-1}) = 0$ is equivalent to

$$\begin{aligned} c_{T-1} H_{T-1} + \lambda H_{T-1} Var[S_T, \mid \mathcal{F}_{T-1}] = \lambda\Big(Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}] \\ - X_{T-2} Var[S_T, \mid \mathcal{F}_{T-1}]\Big) \end{aligned}$$

this gives

$$H_{T-1} = \frac{\lambda\Big(Cov[S_T, S_T V_T \mid \mathcal{F}_{T-1}] - X_{T-2} Var[S_T, \mid \mathcal{F}_{T-1}]\Big)}{c_{T-1} + \lambda Var[S_T, \mid \mathcal{F}_{T-1}]} \tag{44}$$

A.3 Derivation of objective function for BR

Revenue:

$$R_{t+1} = R_t - c_t H_t^2 - H_t(S_{t+1} - S_t) - \sum_{k=0}^{t-1} H_k \cdot (S_{t+1} - S_t) + (S_{t+1}V_{t+1} - S_t V_t) \quad (45)$$

Now, using the simplifying notation $\sum_{k=0}^{t-1} H_k = X_{t-1}$ from the definitions we have the expected revenue:

$$\begin{aligned}
E[R_{t+1} \mid \mathcal{F}_t] =& E[R_t - c_t H_t^2 - H_t(S_{t+1} - S_t) - X_{t-1}(S_{t+1} - S_t) + (S_{t+1}V_{t+1} - S_t V_t) \mid \mathcal{F}_t] \\
=& R_t - c_t H_t^2 - H_t E[(S_{t+1} - S_t) \mid \mathcal{F}_t] - X_{t-1} E[(S_{t+1} - S_t) \mid \mathcal{F}_t] \\
&+ E[(S_{t+1}V_{t+1} - S_t V_t) \mid \mathcal{F}_t] \\
=& R_t - c_t H_t^2 + E[(S_{t+1}V_{t+1} - S_t V_t) \mid \mathcal{F}_t] \\
=& R_t - c_t H_t^2 + Cov(S_{t+1}, V_{t+1} \mid \mathcal{F}_t) \qquad (46)
\end{aligned}$$

Variance of the revenue:

$$\begin{aligned}
Var[R_{t+1} \mid \mathcal{F}_t] =& Var[R_t - c_t H_t^2 - H_t(S_{t+1} - S_t) - X_{t-1}(S_{t+1} - S_t) + S_{t+1}V_{t+1} - S_t V_t \mid \mathcal{F}_t] \\
=& Var[-H_t S_{t+1} + H_t S_t - X_{t-1} S_{t+1} + X_{t-1} S_t + S_{t+1} V_{t+1} \mid \mathcal{F}_t] \\
=& Var[-H_t S_{t+1} - X_{t-1} S_{t+1} + S_{t+1} V_{t+1} \mid \mathcal{F}_t] \\
=& (H_t + X_{t-1})^2 Var(S_{t+1} \mid \mathcal{F}_t) - 2(H_t + X_{t-1}) Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t) \\
&+ Var(S_{t+1}V_{t+1} \mid \mathcal{F}_t) \qquad (47)
\end{aligned}$$

With

$$f = E[R_{t+1} \mid \mathcal{F}_t] - \lambda Var[R_{t+1} \mid \mathcal{F}_t] \quad (48)$$

First order condition yields the optimal H_t

$$\begin{aligned}
0 =& -2c_t H_t - \lambda\Big(2H_t Var(S_{t+1} \mid \mathcal{F}_t) + 2X_{t-1} Var(S_{t+1} \mid \mathcal{F}_t) \\
&- 2Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t)\Big) \\
&\Leftrightarrow H_t\Big(2c_t + 2\lambda Var(S_{t+1} \mid \mathcal{F}_t)\Big) = 2\lambda\Big(Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t) - X_{t+1} Var[S_{t+1} \mid \mathcal{F}_t]\Big)
\end{aligned}$$

$$\Leftrightarrow H_t = \frac{Cov(S_{t+1}, S_{t+1}V_{t+1} \mid \mathcal{F}_t) - X_{t+1} Var[S_{t+1} \mid \mathcal{F}_t]}{\frac{c_t}{\lambda} + Var[S_{t+1} \mid \mathcal{F}_t]} \quad (49)$$

References

[1] Aliakbari E, McKitrick R, 2014, *Energy Abundance and Economic Growth: International and Canadian Evidence*, Fraser Institute, available at `http://www.friendsofscience.org/assets/documents/energy-abundance-and-economic-growth.pdf`, accessed at 13 March 2015

[2] Armerin F, 2004, *Stochastic Volatility, A Gentle Introduction*, Department of Mathematics, Royal Institute of Technology, Stockholm, Sweden. Available at `https://people.kth.se/~armerin/document/stochvol.pdf`, accessed at 24 February 2015

[3] Björk T, 2009, *Arbitrage Theory in Continuous Time, Third Edition*, Oxford University Press, 2009

[4] Björk T, Murgoci A, 2014, *A theory of Markovian time-inconsistent stochastic control in discrete time*, Finance and Stochastics, Vol 18 Issue 3, 545-592

[5] Escobar M, Götz B, Seco L, Zagst R, 2009, *Pricing of spread options on stochastically correlated underlyings*, The Journal of Computational Finance, Vol 12 Number 3, 31-61

[6] Hiller F, Lieberman G, 2010, *Introduction to Operations Research, Ninth Edition*, McGraw-Hill, International Edition 2010

[7] Nordic Energy Regulators: *Nordic Market Report 2014*. Available at `http://www.nordicenergyregulators.org/wp-content/uploads/2014/06/Nordic-Market-Report-2014.pdf`, accessed at 25 February 2015

[8] NordPoolSpot: *Bidding areas*. Available at `http://www.nordpoolspot.com/How-does-it-work/Bidding-areas/`, accessed at 3 March 2015

[9] Vattenfall: *Europes energy markets*. Available at `https://www.vattenfall.com/en/file/Europes-energy-markets_8459821.pdf`, accessed at 25 February 2015

[10] Voorneveld M, *SF2972: Game Theory* (Lecture Notes), Department of Mathematics, Royal Institute of Technology, Stockholm, Sweden. Available at `http://www.math.kth.se/matstat/gru/sf2972/2015/lecture6_2015.pdf`, accessed at 6 March 2015