Chapter 1

A Survey of Electricity and Related Markets

In the beginning of the 1990s a liberalization of the electricity and gas markets started, resulting in the emergence of markets for spot and derivative products in numerous countries and regions spread over the world. The securitization of, for instance, weather, freight and greenhouse gas emission rights, contribute to a greater flexibility for risk control for both producers and consumers operating in the energy markets. In the present book, the aim is to develop tailor-made stochastic models for the various assets traded in electricity and related markets. These assets are in many ways distinct in nature and definition compared to what we find in the more "classical" commodity markets as oil, coal, metals and agriculture. Hence, new and challenging modelling problems appear.

Our main focus will be on stochastic modelling of the electricity market. In an arbitrage-free modelling framework, the spot price does not play the same important role as it does in other commodity markets. Electricity has very limited storage possibilities. Producers may store indirectly in water reservoirs (for hydro-based electricity production) and via gas, oil or coal (for thermal electricity production). However, the consumer of electricity cannot buy for storage. This has the implication that the cost-of-carry relationship between spot and forward prices breaks down. Further consequences of the lack of storeability are strong seasonality and possible spikes in prices. The spiky behaviour of spot electricity prices is a stylised feature of these markets, and appears when, for instance, a nuclear power plant must unexpectedly be closed down, or temperature drops significantly. Power prices may soar during short periods of time, and then fall back to more normal levels shortly after (giving a "spike" in the price path). Typically, the spot price volatilities may exceed the levels usually observed in stock markets by several orders. The limited storage possibility also means that electricity markets are regional. For instance, a difference in the price of electricity between the Nordic power exchange, Nord Pool, and the German-based European Power Exchange (EEX), does not necessarily imply an arbitrage opportunity. An arbitrageur cannot buy for storage and transportation, and therefore the spot asset cannot be used to set up dynamic hedging strategies exploiting price differentials.[1] The tradable assets in such markets are typically average-based forward contracts, that deliver electricity over a specified time period.

In the title of the book we refer to *related markets* of electricity. A technical approach to define a market as being related to electricity is to look for the same modelling characteristics, such as limited storeability of the spot, seasonally dependent prices with spikes, and where the tradeable assets are average based forward contracts. An economist, on the other hand, would typically have a different focus, defining related markets as those markets interacting more or less directly with the demand or supply side of the electricity market. In the application part of this book we study, in addition to the Nord Pool electricity market, UK natural gas and the temperature markets. Both temperature and natural gas share similarities with the electricity market from a modelling point of view. Temperature is obviously not possible to store. Natural gas on the other hand, can be stored, but most often it is quite costly.[2] Limited storage capacity makes the natural gas markets the less extreme sibling to electricity from a modelling point of view. We find average based forward type contracts in all these markets. From the economic point of view the temperature market is linked to the demand side of the electricity market. Electricity demand varies with temperature when power is needed for cooling in areas with warm summer temperatures, or heating in areas with cold winters. Temperature can also affect the electricity supply side, but this effect is typically weaker. In warm summer periods nuclear power plants in continental Europe have been forced to cut on production due to lack of cold water for cooling. High temperatures will also reduce hydro production because of excessive evaporation from water reservoirs. The gas market on the other hand, is mainly linked to the electricty market through the supply side of gas fired power plants.

Before starting our analysis, we provide a survey of the three markets we

[1] Many regional markets are interconnected through cables, however, these have limited capacity preventing a full exploitation of the potential arbitrage.

[2] There is limited storage capacity in the gas pipeline system, and big flexible underground storage facilities are typically major investments.

are mainly concerned about, namely electricity, gas and temperature. Our emphasis will be on how these markets function, with particular attention to the obstacles we face when trying to model the different products offered for trade. We also give an informal discussion on the models and techniques we are going to apply in this book, together with their relevance for the markets in question.

1.1 The electricity markets

Electricity is usually labeled a "commodity", although its non-storeability has a profound effect on the infrastructure and the organization of the electricity market compared with other commodity markets.[3] Electrical power is only useful for practical purposes if it can be delivered during a period of time. This is why electricity has been called a *flow* commodity. Deregulated power markets have market mechanisms to balance supply and demand, where electricity is traded in an auction system for standardized contracts. All contracts guarantee the delivery of a given amount of power for a specified future time period. Some contracts prescribe physical delivery, while others are financially settled.

Financial power contracts are linked to some reference electricity spot price, and they are settled in cash. The market for such contracts is open to speculators, since it is not required to have consumption or production of electricity to participate in the market. We will focus our modelling efforts on the Nordic power market Nord Pool, but our results can be applied to contracts traded at other power exchanges. For instance, the base load financial contracts traded on the EEX and the French exchange Powernext are more or less identical to the Nord Pool contracts. In the following subsections we will describe both the physical and the financial electricity contracts traded at Nord Pool, along with a brief review of some of the relevant literature connected to the modelling of electricity prices.

1.1.1 *Electricity contracts with physical delivery*

By physical electricity contracts we mean contracts with actual consumption or production as part of contract fulfillment. Since capacity is restricted, and the supply and demand must balance, these markets must be

[3] [Stoft (2002)] provides a unique treatment of the interplay between economics and engineering in deregulated electricity markets. [Wolak (1997)] gives a description of worldwide electricity market organization after deregulation.

supervised by a transmission system operator (TSO). Moreover, typically, the players in these markets are restricted to those with proper facilities for production or consumption. The contracts for physical delivery are usually organized in two different markets, the real time and day ahead market.[4] This is known as the two-settlement system.

The *real-time* market (henceforth RT market) is organized by a system operator for short-term upward or downward regulation. The auction specifies both load and time period for generation or consumption. Bids in the RT market are submitted to the TSO. Bids may be posted or changed close to the operational time, in accordance with agreed rules. RT market bids are for upward regulation (increased generation or reduced consumption) and downward regulation (decreased generation or increased consumption). Both demand and supply side bids are posted, stating prices and volumes. Market participants must be able to commit significant power volumes on short notice. In the Nordic market the TSOs are Statnett (Norway), Svenska Kraftnät (Sweden), Fingrid (Suomen Kataverkko Oyj) (Finland), Elkraft System AS (Zealand - Eastern Denmark) and Eltra (Jutland/Funen - Western Denmark). TSOs list bids for each hour in priority order, according to price (merit order), and the merit order for each hour is used to balance the power system. Upward regulation is applied to resolve a grid power deficit. Then the RT market price is set at the highest price of the units called upon from the merit order. In the case of grid power surplus, downward regulation is applied, and the lowest price of the units called upon from the participation list sets the RT market price. The auctions in each country are in effect Walrasian auctions, but the specific rules for determining the hourly price of power imbalances, based on the RT market price, differ among the Nordic TSOs. In addition to the RT auction market the national TSOs have established markets for various necessary ancillary services providing balance power to market actors and securing the operational reserves needed for the system balance.

There also exists a *day-ahead* market (henceforth DA) in most deregulated electricity markets. In the Nordic area, the DA market is a non-mandatory market called Elspot and it is organized by Nord Pool. The UK Power Exchange (UKPX), Powernext and EEX are also examples of non-mandatory DA markets, contrary to, for instance, Omel in Spain. On Elspot, hourly power contracts are traded daily for physical delivery in the next day's 24-hour period (midnight to midnight). On Nord Pool's spot

[4]See Part 3 in the book by [Stoft (2002)] for a detailed description of alternative market architectures for both day-ahead and real time-markets.

market, Norwegian, Swedish, Finnish and Danish players trade in hourly contracts for each of the 24 hours of the coming day. Each morning, the players submit their bids for purchasing or selling a certain volume of electricity for the different hours of the following day. Once the spot market is closed for bids, at noon each day, the DA price is derived for each hour next day. The DA price is called the *system price*, and is common to all Nordic countries. In case of congestion due to capacity constraints, the Nordic market is divided into different bidding areas, resulting in area (or zonal) prices. Each contract is assigned a specific load for a given future delivery period. This means that, strictly speaking, the DA market is trading in electricity forward contracts with delivery over a specified hour the next day. Figure 1.1 shows a time series of weekly averages of the system price in the Nordic market.

There also exists a market that somewhat closes the gap between the DA and RT markets. This is called the Elbas market. The time span between the Elspot price fixing round and the actual delivery hour of the concluded contracts is quite long (36 hours at the most). The Elbas market supplements the Elspot and the national Nordic RT markets, as it allows a market player to adjust the market exposure between the DA auction on Elspot at noon and the actual hour of delivery the following day. Elbas is run by Nord Pool Finland Oy (formerly EL-EX Electricity Exchange Ltd.). It has been in operation since 1999. This market provides continuous power trading 24 hours a day covering individual hours in the same way as the Elspot market (one hour delivery period with 1 MW load). However, the contracts are open for trade only after the Elspot auction, so 24 new contracts are introduced daily. Just like Elspot, Elbas is a physical market for power trading in hourly contracts. The products can be traded up to one hour prior to delivery (two hours prior to delivery in Eastern Denmark). It only covers the trading areas of Finland, Sweden and Eastern Denmark, and the amount of power traded is limited by the free cross border transmission capacity. Unlike the Elspot market, Elbas allows participants to buy and sell the same physical contract several times before delivery, and the position can be closed prior to delivery, so that no actual physical delivery is necessary.

1.1.2 *Financial electricity contracts*

Specifications and rules of trading for financial electricity contracts vary among the different power exchanges. The fact that these contracts are

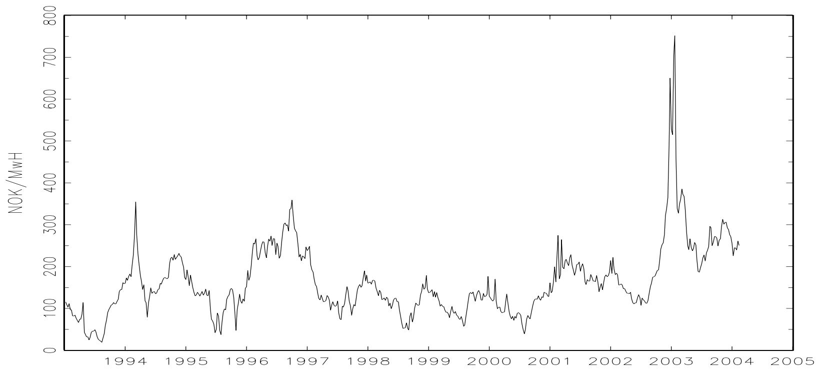

Fig. 1.1 Time series of spot prices from Nord Pool in the period 1993-2004 (weekly average of DA prices).

settled financially against a reference price, implies that the market place for financial electricity contracts does not require central coordination. They can be considered as side bets on the physical system. We will describe contracts traded on Nord Pool only.

Although contracts for future delivery of power are called futures or forwards, this denomination may be misleading. The basic exchange traded contracts at Nord Pool are written on the (weighted) average of the (hourly) system price over a specified *delivery period*. During the delivery period the contract is settled in cash against the system price, hence, financial electricity contracts are in fact swap contracts, exchanging a floating spot price against a fixed price. However, to be in line with the market jargong, we call these contracts *electricity futures* in this book.[5] When we refer to the spot price in our mathematical modelling, this should be interpreted as whatever reference price which a given exchange has tied its financial contracts to. The specified reference price is typically the DA price described in the previous section. In this way the financial electricity contracts are not the relevant risk management vehicles for hedging RT electricity price risk. Contracts on Nord Pool are not traded during the delivery period, and market participants typically close their position prior to the delivery period. We shall consider only electricity futures dynamics in the trading period in our mathematical models. The trading period is the time period the contract is available for trading. The term "time to maturity" used for

[5]We will use the term swap, but then as a common reference to electricity futures and gas futures.

fixed maturity forward contracts is replaced by *time to delivery*.

Nord Pool has facilitated trading in financial electricity contracts since 1995. Since the contracts are settled against hourly DA prices (the Nord Pool system price), the underlying amount of electrical energy is determined by

$$\mathrm{DP} \times 24\,\mathrm{MWh}\,,$$

with DP being the "delivery period" measured in days. These are base load contracts. To be able to compare contracts with different delivery periods, prices are listed in Euros (EUR) for 1 MWh of power delivered as a constant flow during the delivery period.

Since the start in 1995, contract specifications have changed several times. Peak load contracts were available the first couple of years, but were taken from the market due to low liquidity. However, in the summer of 2007 they were reintroduced. There has also been a change of delivery periods for monthly and seasonal contracts. Block contracts with delivery periods of exactly four weeks, have been replaced by monthly contracts with delivery period equal to the respective calendar month. Quarterly contracts have replaced the former three-season regime. The quarterly contracts were first introduced for the year 2005.

In the first trading day in January each year, four new quarterly contracts (Q1, Q2, Q3 and Q4) are listed. The new Q1 contract trades for two years, the new Q2 contract trades for two years and three months, etc. A new yearly contract that trades for three years is also introduced. Thus, it is possible to hedge the exposure to electricty prices in the Nordic market three to four years into the future at all times. Every month a monthly contract is unlisted, and a new one is introduced that trades for six months. Each week one weekly contract is unlisted, and a new one is introduced that trades for eight weeks. New daily contracts are introduced every Thursday. The Friday contract only trades for one day. All contracts trade until the last trading day prior to the delivery period.

The contracts differ when it comes to how settlement is carried out during the trading period. Daily and weekly contracts are futures contracts. The value of such a contract is calculated daily, reflecting changes in the contract's market price. These changes are settled on a margin account for each participant. The electricity futures with monthly, quarterly and yearly delivery are forward-style contracts.

Nord Pool's financial market also includes option contracts and Contracts for Differences (CfD). Call and put options are written on the electricity futures contracts, and they are of European type. Exercise day is set as the third Thursday in the month before the delivery period of the underlying contract starts. The options are traded on quarterly and yearly contracts. The activity on the option market on Nord Pool is, at the time of writing this book, rather low. Hence, it is not easy to derive implied volatilities from this market. The EEX is also offering options written on electricity futures. Asian options written directly on the system price are frequently traded in the bilateral over-the-counter (OTC) market. As mentioned above, area prices may differ from the system price in case of congestion. CfDs are defined as the area price minus the system price. The different tradeable area prices are Oslo, Stockholm, Copenhagen, Aarhus, Helsinki and northern Germany. The CfDs are defined similar to the electricity futures contracts with identical delivery periods. However, delivery periods shorter than one month do not trade. Using CfDs in combination with electricity futures allows a market participant to effectively remove the price risk associated with congestion.[6]

1.2 The gas market

Natural gas is an important fuel for heating and when generating electricity. For instance, in 2002 one-third of the electricity production in the UK came from gas fired power plants, with a prospect of 60% by 2020 (see [Geman (2005)]). The figure for US is that 14% of gas demand comes from electricity generation. The gas markets, foremost in UK/Europe and the US, have been liberalized over the years, with some structural differences and similarities with the electricity markets. In this section we give a brief overview of the specifics of the gas markets actively traded in the US and UK.

The gas markets are located around different hubs, which are connection and arrival points for gas transportation systems and where there are infrastructure capabilities like, for instance, storage and a concentration of buyers and sellers. Two important hubs are Henry Hub located in Louisiana (US) at the Mexico Gulf and the National Balancing Point (NBP) in the UK. The latter is a *notional* hub without any physical location, where all UK gas flows through. The market for short-term delivery of gas is usually

[6]See http://www.nordpool.no/nordpool/financial/index.html for details on the different financial contracts traded on Nord Pool. The interested reader is recommended to read [Kristiansen (2004)] for more on the pricing of CfDs.

referred to as the *spot market*, and the trading is mostly OTC. Futures contracts ensure the delivery of gas over longer time periods like weeks, months, quarters, or even years, and the settlement of these resembles closely electricity futures. Although the largest portion of the trade in futures takes place in the OTC market, some exchanges also offer futures with physical delivery of gas through a hub. We shall refer to futures contracts in the gas market as *gas futures*, following the terminology used in the industry.

Gas prices, very much like electricity prices, exhibit sudden spikes during periods of high demand or shortage of production (or low storage), as can be seen in Fig. 1.2 presenting gas spot prices at the NBP. This gives

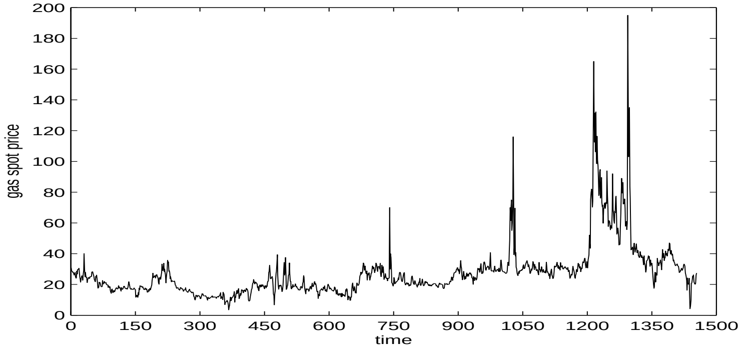

Fig. 1.2 Gas spot prices (Heren DA index) at the NBP for the period 6 February 2001 till 24 October 2006.

rise to a price dynamics having a higher volatility than what is normally observed in other commodity markets (like oil, say). Furthermore, especially in the UK market, the gas prices are seasonal since demand is very much dependent on temperature. Due to capacity constraints, one may even observe negative prices of gas from time to time (however, naturally rather infrequent). On the other hand, gas can be stored, which makes it possible to use for hedging. This links the analysis of the gas markets closer to more "classical" commodities like oil.

In the market place, the energy content of gas is measured in units of 'therms' or 'British thermal units' (Btu). By definition there are 100,000 Btu in 1 therm, whereas 1 therm is the equivalent of 105.5 MJ. Since there are 3.6 GJ per MWh, we have the relation

$$1\text{therm} = 105.5\text{MJ} \cdot \frac{\text{MWh}}{3.6 \cdot 1000\text{MJ}} = 0.029306\text{MWh}\,.$$

In the US, gas transactions are denominated in Dollars per MBtu, while in the UK pence per therm is used as the unit.

1.2.1 *Futures and options on gas*

Although most of the trading of futures on gas takes place bilaterally, there exist organized markets as well. For instance, New York Mercantile Exchange (NYMEX) offers trading in standardized futures contracts with physical delivery of gas at Henry Hub over a specified month. Altogether 72 contracts are offered for trade at all times, covering the nearest consecutive months. The participants of this market have thus access to futures covering six years ahead. In addition, European options are written on the gas futures contracts.

At the Intercontinental Exchange (ICE), structured UK natural gas futures are traded. The contracts deliver gas at a fixed rate through the NBP over a specified period of time. The delivery periods are the first 10-12 consecutive months, 11-12 quarters and six seasons. There are two seasons, a summer season ranging from April to September, and a winter season from October to March. It is worth noticing the similarities with the electricity futures traded at Nord Pool, where the market is also separated into contracts with longer and shorter delivery periods, with only long delivery contracts in the long end of the curve.

There is no organized market for gas spot trading, in the sense of standardized spot contracts traded on an open exchange leading to publicly available prices. This raises the question how a gas futures contract can be benchmarked against the spot market. The lack of an objective reference price for the futures is resolved by objective indexes provided for the market. For the ICE gas futures, the Heren NBP DA index is used as a measurement of the spot price of gas at the NBP. This index is provided by Heren Energy,[7] and is published daily in the European Spot Gas Markets report.[8] The index for a specific day is the volume-weighted average of transaction prices for gas to be delivered at NBP the next day. This index constitutes the counterpart to the system price at Nord Pool, except that

[7] www.heren.com
[8] www.energypublishing.com

the latter is set for each hour the next day, and is a real trading price. The index gives, of course, just the average level of all relevant trades agreed on for the next day in the short-term market. There also exist similar indices for Henry Hub, which are used to settle NYMEX gas futures.

An important class of derivatives contracts is the so-called spark spread options. These are call and put options written on the difference between electricity and gas prices. A typical example may be a European put option on the futures price spread of the two energies, where the futures have a specified monthly or quarterly delivery period. Such options may be used for risk management of a gas fired power plant, giving the plant owner a possibility to hedge undesirable moves in the electricity and gas prices. These options are rather popular, and traded to a siginificant extent in the OTC market. Spark spreads may also be used for real option valuation of gas fired power plants.

1.3 The temperature market

In recent years the trade in contracts written on weather has emerged and become a new and interesting asset class for risk management. There are obviously close connections between energy and weather, like, for instance, an increase in power consumption during cold and warm periods. This means that both consumers and producers of energy may be interested in financial contracts that can be used to manage weather risk. Markets for weather derivatives are thus of importance for the energy industry. A thorough account on weather derivatives markets and valuation is given in [Geman (1999)] and [Jewson and Brix (2005)]. In this Section we shall concentrate on the market for temperature derivatives found at the Chicago Mercantile Exchange (CME).

From a modelling point of view, there is a close resemblance between weather and energy markets. The markets are incomplete, in the sense that hedging using the underlying is impossible. Further, there are clear evidences of mean reversion. In the energy markets this is due to the balance between demand and supply, while for temperature it may be explained by conservation of energy. Further, the typical temperature futures contracts are written on some temperature index measured over a period of time, which one may think of as a “delivery period”. In this respect, the temperature futures have “delivery” of the underlying “asset” over a period rather than at a fixed point in time. To be in line with the industry terminology, we shall refer to such contracts as *temperature futures*.

At the time of writing this book, the market for weather derivatives is steadily increasing. The OTC market offers a wide range of different weather deals, while the volumes for temperature futures and options at the CME have experienced constant growth. New weather contracts like futures and options written on the amount of snowfall in New York and the frost days in Amsterdam, have emerged, and it is expected that even more weather related derivatives will be introduced at the exchange in the future.

In our discussion we shall focus on the temperature based products actively traded at the CME. CME organizes trade in futures contracts based on four different temperature indices. These indices measure the aggregation of daily mean temperature or its deviation from a fixed level, in 18 US, two Japanese and nine European cities. In addition, there is a Frost Day index based on the temperature in Amsterdam. We refer the reader to *www.cme.com* for a detailed description of all these temperature products, which we are going to discuss.[9]

For the US cities, the contracts are written on the aggregated amount of *heating-degree days* (HDD) and *cooling-degree days* (CDD). The amount of CDD on a particular day is defined as the difference between mean temperature and 65°F (18°C), whenever this is positive. In mathematical terms, the CDD on day t is defined to be

$$\mathrm{CDD}(t) = \max\left(T(t) - c, 0\right) ,$$

where $T(t)$ is the mean temperature on day t. The mean temperature is interpreted as the average of the maximum and minimum temperature over the 24 hours of the day in consideration. The contracts are written on the accumulated amount of CDD over a month or a season.[10] The constant c denotes the threshold 65°F (or 18°C). Since most air conditioners are switched on when temperatures are above c, the CDD gives a measurement of how much air conditioning it is required, and therefore is linked to the demand for power. The temperature futures contracts written on the CDD index is for the warmer half of the year, ranging from April to October. The CDD futures are settled financially in terms of $20 per unit for the US cities. If the contract is specified as the accumulated CDD over a period

[9] We are not going to consider the snowfall contracts for New York, since these are not directly temperature linked.

[10] A season being two to seven months long.

$[\tau_1, \tau_2]$, the amount of money to be paid to the buyer of the contract is

$$\sum_{t=\tau_1}^{\tau_2} \mathrm{CDD}(t) \times \$20 \,.$$

The HDD index measures the amount of degrees below the threshold c, which is an index of how much heating it is required. It is defined as

$$\mathrm{HDD}(t) = \max\left(c - T(t), 0\right) \,,$$

and the futures contracts are written on accumulated HDD over a month or a season in the colder part of the year, lasting from October to April.

For the European cities, there is trade in HDD contracts in the winter season, while in the summer season the underlying temperature index is the so-called *cumulative average temperature* (CAT). The CAT over a period $[\tau_1, \tau_2]$ is defined as

$$\sum_{t=\tau_1}^{\tau_2} T(t) \,.$$

The currency is British pounds for the European futures contracts, which are also settled in units of 20. Finally, the Japanese futures are settled against the index called *Pacific Rim* (PRIM), being the average temperature over a period

$$\frac{1}{\tau_2 - \tau_1 + 1} \sum_{t=\tau_1}^{\tau_2} T(t) \,.$$

The currency is Japanese yen and the settlement is in units of 250,000. The contracts are listed for all the 12 months of the year.

A different class of futures and options traded at the CME is based on an index measuring the days where there is a danger of icy conditions on the runway of Schiphol airport in Amsterdam. The Frost Day index is defined as the accumulated number of days in a month or season when there is frost conditions observed at Schiphol airport. Each day counting as a frost day in the measurement period gives a contribution 1 to the index. If the temperature during a day satisfies one or more of the following three conditions, it is defined as a frost day:

(1) The temperature at 7 a.m. is less than or equal to $-3.5°$C,
(2) The temperature at 10 a.m. is less than or equal to $-1.5°$C

(3) The temperatures at 7 a.m. and 10 a.m. are less than or equal to $-0.5°$C.

At the CME, futures are listed on monthly November to March Frost Day indices. The seasonal Frost Day index is listed for the whole period of November to March. The trade unit is EUR 10,000 times the index. Note that the Frost Day index over a measurement period becomes a natural number including zero. The upper limit is the number of measurement days for the index in question. A frost day on day t is mathematically defined as

$$\mathrm{FD}(t) = 1\Big(\{T(t+7/24) \leq -3.5\} \cup \{T(t+10/24) \leq -1.5\} \cup \{\{T(t+7/24) \leq -0.5\} \cap \{T(t+10/24) \leq -0.5\}\}\Big)\,.$$

Here $1(\cdot)$ is the indicator function. Moreover, we use the convention that the time t is measured in days, implying that $t+7/24$ is at 7 a.m., and $t+10/24$ is at 10 a.m. of the day in question. The Frost Day index over a measurement period $[\tau_1, \tau_2]$ is then defined as

$$\mathrm{EUR}10,000 \times \sum_{t=\tau_1}^{\tau_2} \mathrm{FD}(t)\,.$$

In the market place, only weekdays are taken into account when finding the Frost Day index. Thus, in the summation above, we should disregard weekends. To avoid the introduction of a new (and more messy) notation we consider here a summation over all the days in the measurement period.

The temperature market at CME also includes options written on the different futures. The options are plain vanilla European call and put options, with a temperature futures as the underlying asset.

1.4 Other related energy markets

Natural gas and temperature are not the only two markets related to electricity, there are others not discussed in detail in this book. Different oil products are certainly important fuels for electricity generation. However, since storage is easier, and since the forward market for oil has been studied quite extensively over the past decades (see, for instance, [Schwartz (1997)]), we do not discuss the oil market in this book. We only briefly comment on the market for coal.

The economic link between coal and electricity is strong, as coal fired plants represent an important part of total power supply in many electricity markets. The market for coal has historically been a physical market, with big players on both the supply and the demand side. The contractual agreements have typically been long-term and bilateral. In recent years there has been increasing London-based OTC trading activity in forward type contracts. The forward contracts bear close resemblence with the average-based contracts which we describe in this book. Heren Energy collects fair prices from different market players for coal delivered at certain landing points around the world. This is done each week. The OTC traded forward contracts are settled financially on average price fixings during the settlement month. Exchange traded coal contracts have been around for a while, but so far this market has yet to see a big trading activity. The link to the electricity market has made the EEX to launch exchange traded coal futures.

A market with a more indirect link to the electricity market is the growing financial shipping market. The development of the freight derivatives market has spanned more than two decades, starting with the Baltic International Freight Futures Exchange (BIFFEX) market in 1985 and, since 1992, gradually developing into an active OTC Forward Freight Agreement (FFA) market. Towards the end of 2001, the development came full circle, with the emergence of electronic trading of route-specific cleared tanker derivatives on the Oslo-based IMAREX platform. Since 2005, large global clearing houses such as London Clearing House and NYMEX have also offered clearing of traditional voice-brokered FFA contracts. However, both contracts are cash settled against the average spot freight rate for a particular route, as published daily by the Baltic Exchange, over a specified future period of time and for a specified quantity of cargo and vessel type.

From a modelling point of view, the contracts are identical to the electricity and gas futures contracts studied in detail in this book. Research in the area of freight derivatives has been primarily concerned with various aspects of market efficiency, first in the BIFFEX futures market and later in the OTC FFA market. The research topics include the applicability of the unbiasedness hypothesis, hedging effectiveness for ship operators, and the interaction between the spot and forward markets. [Aadland and Koekebakker (2007)] and [Koekebakker, Aadland and Sødal (2007)] are early attempts to study freight rate dynamics in a continuous time model similiar to the modelling framework advocated in this book. The dry bulk shipping market has been very volatile in recent years, and freight rates have peaked. In some periods freight has been a significant cost factor for coal fired power plants with short-term contracts. This linkage between shipping and electricity markets may become even more important in the future.

A market segment that has increased in size over the years is the freight market for natural gas. Gas can be liquified and compressed through a cooling process (liquified natural gas is known as LNG). LNG carriers are very expensive, but the new building activity has been high for this vessel type in recent years due to the increase in gas fired power plants around the world. As storage facilities for natural gas are both expensive and require a certain infrastructure, storage capacity is scarce. Recently it has been speculated that LNG ships are hired also as floating storage devices to exploit 'LNG arbitrage' in natural gas pipeline systems. If these trends escalate in the future, it will make sense to include certain shipping markets as integral parts of the electricity market.

Finally, we comment on the emerging CO2 emission market. In 1997, many governments adopted the Kyoto protocol accepting mandatory constraints on reduction of greenhouse gases emission. The Kyoto protocol contains different mechanisms to reduce emissions: International Emission Trading, Joint Implementation, and Clean Development Mechanism. Carbon dioxide is by far the most important greenhouse gas, the runner up is methane. The protocol promotes permit trading as the champion mechanism to reduce CO2 emissions.

There is a close economic link between prices on CO2 emissions and electricity. Increased cost of pollution increases costs for many power producers. Producers can either buy enough emissions and keep on polluting, or they can invest in cleaning technology (or both). From a modelling point of view, the issue of non-storability is a shared characteristic with the electricity market. The basic 'commodity' is overall CO2 emissions. Since it is not possible to trade physical emissions, spot certificates have been introduced as tradeable assets. According to [Daskalakis *et al.* (2006)] there are four active emission allowance markets: the European Union Emission Trading Scheme (EU ETS), the UK Emission Trading System (UK ETS), the New South Wales GHG Abatement Scheme and the voluntary Chicago Climate Exchange (CCX). The EU ETS is dominant with a share of approximately 97% of the total transactions during the first three quarters of 2006 ([Daskalakis *et al.* (2006)]).

Financial research in this market is still scarce, and the research conducted has to our knowledge been focusing exclusively on EU ETS. An early discussion of emission trading and reports from an expert survey is given in [Uhrig-Homburg and Wagner (2006)]. The EU ETS is a joint effort by EU member states to efficiently reach their Kyoto goals. The EU ETS breaks down the emissions trading to the company level. Companies in industries with big emission needs, are allocated a certain amount of EU Allowances (EUAs). One EUA gives the holder the right to emit one tonne of CO2. If a company does not use all its allocated permits, due, for instance, to new environmental friendly technology, the surplus EUAs can be sold. Other companies, where new technology is more expensive or perhaps does not exist, can buy additional permits if needed. Spot EUAs are sold at Powernext, EEX, Nordpool, Energy Exchange Austria (EXAA), European Climate and Exchange (ECX) and Climex (see [Uhrig-Homburg and Wagner (2007)]). Some research has examined spot price dynamics of EUAs, see [Benz and Trück (2006)] and [Paolella and Taschini (2006)] for econometric studies. In the paper [Fehr and Hinz (2006)], the authors build a microeconomic equilibrium model for price formation of carbon emission rights.

Forward contracts on EUAs are also traded. The EUAs can only be used within a particular trading period. The first trading period is 2005 – 2007, the second trading period coincides with the Kyoto commitment period spanning from 2008 to 2012. From a modelling point of view, we observe the interesting feature that the traditional cost-of-carry relationship should hold for forwards that mature within a trading period. But since there exists no EUAs for the second trading period, there can be no spot-forward relationship (see [Uhrig-Homburg and Wagner (2007)], [Borak *et al.* (2006)] and [Daskalakis *et al.* (2006)] for discussions and empirical results). The trading period system suggests that the price dynamics of EUAs changes over time depending in particular on the total emissions. If towards the end of a trading period, cumulative emissions in the period are high, then EUAs would be in high demand, supplies of EUAs are scarce, and we would expect very high prices. In the case of low cumulative emissions, we would expect the opposite, and the prices of EUAs would plummet. [Seifert, Uhrig-Homburg and Wagner (2006)] propose an equilibrium model consistent with such predictions.

1.5 Stochastic modelling of energy markets

The energy related markets consist in general of three different segments, a market for physical spot trading, futures contracts on the spot with either physical or financial settlement over a period, and an option market with the futures contracts as underlying. The exception is the market for temperature, where there is obviously no trading in the "spot". Thus, modelling of the energy markets can be separated into three tasks: spot price modelling, derivation or modelling of futures, and pricing of options. In this section we discuss the different modelling issues, to establish a common foundation for the theoretical and empirical analyses which will be our focus in the subsequent chapters. The discussion here will be kept at an informal level, to leave space for fixing the ideas and highlight the approaches we are going to use.

We emphasise that in this book we refer to swaps being futures contracts with delivery over a period. This will be used as a general reference including electricity and gas futures. Wherever it is natural, we use the terminology "electricity futures" and "gas futures" instead of the general notion "swaps". Temperature futures will be discussed separately. By forwards we understand exclusively contracts with a fixed delivery time. We will be consistent in this separation throughout the book.

A dynamics for the spot price evolution is desirable for several reasons. Models describing the uncertainty in the spot price is of interest for traders operating in these markets. However, they are also used as the reference point for settlement of forward and futures contracts, and thus is a basic input in understanding the dynamics of these derivatives. The spot dynamics will be based on Ornstein-Uhlenbeck (OU) processes, which model mean reversion in a natural way. The stochastic driver may allow for jumps, where we can explain spikes in electricity prices, for example. Finally, it is paramount to allow for seasonal variations, since the demand after electricity and gas vary with temperatures, which are highly dependent on season. In the literature, one usually connects stationarity properties to OU processes. When we include seasonlity (in, for instance, the jump occurrence and size), the traditional notion of stationarity breaks down. From this point of view, the terminology "OU process" may not be natural. However, we keep the name "OU process" in this book, which seems to be the standard use. We interpret it as a dynamic model with certain mean reversion properties.

The next modelling point is to establish the connection between the spot and futures/forward price dynamics. In markets like oil, say, there exists an extensive theory for the spot-forward relation, including storage costs and convenience yields. For electricity and temperature, and perhaps gas, the connection is not at all clear since the underlying spot is not storable. One may explain the relation through a market price of risk, which essentially is the specification of a risk-neutral probability. The existence of a delivery period for the electricity and gas futures (swaps) puts restrictions on the class of spot models feasible for analytical pricing.

Alternatively to explaining the forward and swap prices by the underlying spot, one may adopt the Heath-Jarrow-Morton (HJM) approach from interest rate theory (see [Heath, Jarrow and Morton (1992)]). Rather than trying to establish a spot-futures/forward relation via the specification of a risk-neutral probability, the HJM approach suggests to directly assume a dynamics for the forward and swap price evolution. This can be done in terms of market dynamics, or under the risk-neutral measure. Modelling the swap price dynamics, where the energy delivers over a period, creates challenges that are not present in the fixed income markets theory (see, for example, [Musiela and Rutkowski (1998)] for a discussion of the HJM approach in fixed income markets).

Having a forward and swap price dynamics, our final task is to look at how to price options. With a risk-neutral dynamics available for the forward and swap price, this entails in an exercise of calculating a conditional expectation of the pay-off from the option, which for many of our models can be done more or less explicitly by using Fourier techniques. The question of hedging will also be analysed, however, leaving out a significant part related to incomplete markets.

We discuss now these modelling aspects in more detail, trying to explain our choice of models and approaches that we are going to consider in the following chapters.

1.5.1 *Spot price modelling*

In mathematical finance, the traditional models are based on stochastic processes driven by a Brownian motion $B(t)$, also called a Wiener process. The most frequently encountered model for the price dynamics $S(t)$ of a financial asset is the geometric Brownian motion (see [Samuelson (1965a)]), being the exponential of a drifted Brownian motion. It is defined as

$$S(t) = S(0) \exp\left(\mu t + \sigma B(t)\right) ,$$

with μ and $\sigma > 0$ being constants. Brownian motion is a process with independent and stationary increments, where the increments are normally distributed. This implies that the logarithmic returns (or logreturns, for short), defined as logarithmic price changes over a time interval Δt

$$\ln S(t + \Delta t) - \ln S(t),$$

will become independent and stationary, a reasonable property in view of the market efficiency hypothesis (see, for example, [Fama (1970)]). In addition, logreturns are normally distributed.

A natural and frequently used generalization of the geometric Brownian motion is the exponential of a Lévy process (see, for example, [Barndorff-Nielsen (1998)] and [Eberlein and Keller (1995)]),

$$S(t) = S(0) \exp(L(t)).$$

Lévy processes $L(t)$ open for the possibility to model price jumps and leptokurtic behaviour of asset prices on small time scales. These processes have independent and stationary increments, with Brownian motion being a special case. With these models at hand, we may incorporate the possibility of large price variations, and even skewness in the price fluctuations. However, due to stationarity, the variation in prices is homogeneous over the year, and we cannot allow for more variable prices during winter than summer.

Energy markets, and in particular electricity markets, are seasonally varying markets. By appropriate modelling of the mean level of energy prices, one may remove much of the seasonal features observed in prices, however, there are still distinctive characteristics which call for models that may vary with time. In the electricity market we observe seasonality in the jump size and frequency. For instance, in the Nord Pool market spikes are most frequent in the winter period. Further, the temperature dynamics underlying weather derivatives turns out to have seasonal features like a time-dependent volatility. In other markets, like gas, we see similar seasonal variations in the dynamics, explained by demand being weather dependent.

The classical model for commodity markets is the Schwartz model (see [Schwartz (1997)]), which is an extension of the geometric Brownian motion allowing for mean reversion. In the simplest case, it may be defined as

$$S(t) = S(0) \exp(X(t)), \tag{1.1}$$

where

$$dX(t) = \alpha(\mu - X(t))\,dt + \sigma\,dB(t)\,. \tag{1.2}$$

Extending to Lévy process innovations, we still preserve the homogeneity in jump size and frequency, and we will not be able to explain the observed seasonal features. A reasonably flexible class of models are provided by the *independent increment* (II) processes, which generalize Lévy processes in a way that the increments are independent, but not necessarily stationary. This opens up for multi-factor models of the Schwartz type which may have one or more factors with seasonally dependent jump frequencies and sizes in addition to mean reversion. In other words, the logarithmic spot price is represented as a sum of OU processes driven by II processes. In this way we may model the typical spikes observed for electricity spot prices by having an OU process with big, but rare, jumps coupled with a strong mean reversion. The II processes allow for a stochastic analysis which can be utilised for calculating derivatives prices. This class of models is a reasonable compromise between modelling flexibility and analytical tractability, and will be our stochastic driver in the spot price dynamics. Further, by substituting $S(0)$ in (1.1) by a deterministic funtion $\Lambda(t)$, we can model explicitly a seasonally varying mean level.

We may argue in favour of arithmetic models rather than geometric ones for the spot price evolution, that is, we may assume that the spot price dynamics is represented as a sum of OU processes directly. This makes analytical pricing of swap contracts feasible for a large class of models, a possibility not shared with most geometric models. In this book we shall introduce a class of arithmetic models where we ensure positivity of spot prices, using the specific choice of increasing II processes as stochastic drivers.

The question of estimating such models on data is not an easy one. For some simple one-factor models, this may be a straightforward task, as we shall demonstrate in many examples. However, if the jumps are seasonal we immediately face problems when trying to fit the stochastic model to spot data. For multi-factor models this may be an even more challenging problem, involving highly sophisticated estimation techniques. It is outside the scope of this book to give an exhaustive presentation and application of the different estimation approaches. We shall present many examples, where we can use simple estimation techniques. Emphasis is put on transparency and on showing the connection between data and model without having to implement advanced estimation procedures. Admittedly, to apply our models at full strength, this is not satisfactory. We will indicate possible estimation approaches along the way for the convenience of the reader.

The traditional models in mathematical finance belong to the class of semimartingale processes. The reason for this is the existence of so-called equivalent (local) martingale measures, being probability measures equivalent to the objective (or market) P probability, and such that the discounted price dynamics is a (local) martingale. Existence of such probabilities, which are often coined *risk-neutral probability measures*, leads to markets where there are no arbitrage possibilities, since the martingale property of the discounted prices makes it impossible to create portfolios with a sure win (we refer the reader to [Bjørk (1998)] and [Bingham and Kiesel (1998)] for excellent accounts on this theory). In markets like temperature or electricity the underlying spot (being either temperature itself, or the spot electricity), is not tradeable in the sense of being an asset that can be liquidly bought or sold, and kept in a portfolio over time. Hence, the spot is not accounted for as being a tradeable asset, and will not be a part of the definition when fixing a martingale measure. Therefore, any probability measure Q being equivalent to the objective probability P is also an equivalent martingale measure. This has the implication that we do not need to restrict the class of spot price models to be semimartingales. However, all our models will be semimartingales, since this is a convenient class of processes from an analytical point of view.

On the other hand, the swap and temperature futures markets are liquid, and these contracts have to be priced so that arbitrage opportunities do not exist. Thus, it is required that the dynamics of the forward and swap price possesses the semimartingale property in order to ensure the existence of risk-neutral probabilities. Connecting the spot dynamics with the forward and swap price leads to formulas representable in terms of conditional expectations of the spot dynamics. We will discuss this in more detail in the following subsection.

Let us elaborate on the spot price dynamics for the electricity market. As we discussed in Sect. 1.1, the spot markets of electricity quote prices on an hourly basis (or half-hourly in some markets). This means that strictly speaking, the spot price should be modelled as a time series. Thus, it will not make sense to talk about the spot price of electricity at any time t. On the other hand, we know that electricity futures are settled against the hourly spot prices, but are traded in a continuous market in the sense that the actors can buy or sell at any time as long as they find a counterpart in the market. Hence, contrary to most other commodity markets where there is liquid trading in both spot and futures/forwards, we face the situation of a discrete-time spot and a continuous-time futures market.

Let us introduce a continuous-time stochastic process $\widetilde{S}(t)$ being the *unobserved* instantaneous spot price of electricity, that is, the price of electricity at time t with delivery in the interval $[t, t+dt)$. Associated to the process is a filtration $\mathcal{F}_t$ modelling the stream of information. We can think of the process $\widetilde{S}(t)$ as the price market participants know they would have to pay *if* they could buy electricity at time t with infinitesimal delivery time (that is, like a shock of electricity).

What we do *observe* in the market is the price of electricity with delivery over a specified hour. Let us say that the hour is the time interval $[t_i^d, t_{i+1}^d)$, with $i = 0, 1, \ldots, 23$ denoting the hour and d the day. Thus, t_i^d is the start of hour i on day d. Entering a spot contract will then give us the following expenses

$$\int_{t_i^d}^{t_{i+1}^d} \widetilde{S}(u)\, du\,, \tag{1.3}$$

if we would know the instantaneous spot price. The hourly spot price in the market is set before the delivery takes place. Thus, a natural assumption is that the hourly spot price is the best prediciton of (1.3), *given* the information up to start of delivery. Hence, the hourly spot price should be

$$S_i^d = \mathbb{E}\left[\int_{t_i^d}^{t_{i+1}^d} \widetilde{S}(u)\, du \,|\, \mathcal{F}_{t_i^d}\right]. \tag{1.4}$$

This definition makes the time series S_i^d measurable with respect to $\mathcal{F}_{t_i^d}$, meaning that the hourly spot price contains all market information up to the start of delivery, but not into the delivery period.

Let us approximate the integral inside the conditional expectation in (1.4) with

$$\int_{t_i^d}^{t_{i+1}^d} \widetilde{S}(u)\, du \approx \widetilde{S}(t_i^d)\,,$$

using the convention that time is measured in hours and thus $t_{i+1}^d - t_i^d = 1$. From the measurability of $\widetilde{S}(t_i^d)$ it follows that

$$S_i^d \approx \mathbb{E}\left[\widetilde{S}(t_i^d) \,|\, \mathcal{F}_{t_i^d}\right] = \widetilde{S}(t_i^d)\,.$$

This argues in favour of defining a spot price process in the market as

$$S(t) = \widetilde{S}(t)\,, \tag{1.5}$$

where we observe the spot price at time moments t_i^d, that is, $S(t_i^d)$ are the observations of an underlying continuous-time spot price process of electricity. This is the assumption usually made (implicitly) in the literature when modelling a spot price of electricity. Note that this connection makes it possible to estimate the parameters of the unobserved process $\widetilde{S}$ directly.

The above arguments for linking the unobserved process $\widetilde{S}$ to the electricity spot price may be questioned from several different angles. The spot price is determined in the market from bids in an auction, which results in prices for all hours the next day. It is not clear how the available information is taken into account in this price determination. It is therefore not simple to understand the connections linking the assumed continuous-time process, the filtration $\mathcal{F}_t$ and the electricity spot price, if such a modelling approach is at all valid. To keep matters simple, we assume a continuous-time stochastic process for the electricity spot price, and relate it to the observed spot via (1.5). As we discuss in the next subsection, this setup will also lead to natural connections between the spot and electricity futures price. Note that gas can in principle be purchased OTC at any time (continuously), thus we may assume that the gas spot price is a continuous-time process. Temperature may obviously be viewed as a continuous-time process as well.

1.5.2 *Forward and swap pricing in electricity and related markets*

The key driving factor for the swap price is the underlying spot. The relation between spot and swap prices are of crucial importance to the players in the energy market, and one of the central topics in this book. In the current subsection we aim at illustrating some of the ideas and problems encountered when deriving swap prices for the energy markets.

Suppose $S(t)$ is a stochastic process[11] defining the price dynamics of the spot, and $r > 0$ is the constant risk-free interest rate. For simplicity, let us discuss forward contracts first. Assume that we have entered a forward contract delivering the spot at time τ. Denoting $f(t, \tau)$ the forward price

[11]Since we are not going to give a rigorous treatment of the forward-spot relation, we do not go into details on the exact dynamics of the spot and the hypotheses required. We leave the mathematical details for later chapters.

at the time $t \leq \tau$ of entry of the contract, the payoff from the position is

$$S(\tau) - f(t, \tau)$$

at delivery time τ. From the theory of mathematical finance (see, for example, [Duffie (1992)]), we know that the value of any derivative is given as the present expected value of its payoff, where the expectation is taken with respect to a risk-neutral probability Q. Hence, since a forward contract is entered at no cost,

$$\mathrm{e}^{-r(\tau-t)}\mathbb{E}_Q\left[S(\tau) - f(t,\tau)\,|\,\mathcal{F}_t\right] = 0\,.$$

Here, $\mathcal{F}_t$ is the filtration containing all market information up to time t, and $\mathbb{E}_Q$ is the expectation operator with respect to the risk-neutral measure.

The forward price is set at time t, and therefore cannot include any more information about the market than given by $\mathcal{F}_t$, which therefore implies that it must be adapted to this filtration. Thus, we obtain the following formula for the spot-forward relationship

$$f(t,\tau) = \mathbb{E}_Q\left[S(\tau)\,|\,\mathcal{F}_t\right]\,. \tag{1.6}$$

This definition yields an arbitrage-free dynamics of the forward price process $t \mapsto f(t,\tau)$, since this process is a martingale under Q. In effect, the relation (1.6) implies that the forward price is the best risk-neutral prediction at time t of the spot price $S(\tau)$ at delivery. In this book we exclusively consider the situation where the interest rate r is constant. We know then that forward and futures prices coincide. We will not make any distinction between the two, and reserve the terminology "forwards" for these contracts.

Suppose now that the spot can be liquidly traded in a market (like a stock, say). Then we can perfectly hedge a short position in the forward contract by a long position in the spot, financed by borrowing at the risk-free rate r. This hedging strategy is known as the *buy-and-hold strategy*, and uniquely defines the forward price. Since Q is a risk-neutral probability, it follows by definition that the discounted spot price $S(t)\mathrm{e}^{-rt}$ is a martingale under Q, and thus we get

$$f(t,\tau) = S(t)\mathrm{e}^{r(\tau-t)}\,. \tag{1.7}$$

This is the well-known connection between a forward contract and the underlying spot in a market where the two assets can be traded frictionless (a complete market). From (1.7) we easily see that the forward price converges to the underlying spot price when time to delivery $\tau - t$ approaches zero.

When running a buy-and-hold strategy in a commodity market, the commodity must be stored. Thus, the hedger will be incurred additional costs reflected in the forward price (1.7) as an increased interest rate to be paid. On the other hand, holding the commodity has a certain advantage over being long a forward contract due to the greater flexibility. For instance, the access to a gas storage facility means that one can sell gas when prices are high, and store if prices are low. Furthermore, if you run a gas fired power plant, you ensure production with such a storage facility. These opportunities are lost when holding a forward contract instead. The notion of *convenience yield* is introduced to explain this additional benefit accrued to the owner of the physical commodity. If it is assumed that the convenience yield comes at a constant rate, the hedging argument leading to (1.7) is modified exactly as if the spot would be a dividend paying stock. Hence, letting the convenience yield rate be c, and the storage costs be measured at a rate s, we get

$$f(t,\tau) = S(t)\mathrm{e}^{(r+s-c)(\tau-t)}\,. \tag{1.8}$$

Obviously, to measure the convenience yield is a more delicate task than the dividends paid from a stock. Note that the relation (1.8) can be derived from (1.6) via an appropriate choice of a risk-neutral measure for reasonable spot price models. Therefore, it may be more convenient and give more flexibility to start out with (1.6) as the definition of the forward price. Choosing Q will correspond, loosely speaking, to specifying the convenience yield. We will choose this approach for gas. We refer to [Geman (2005)] and [Eydeland and Wolyniec (2003)], and the references therein, for more details on the convenience yield and storage in commodity markets.

In the electricity market, these considerations break down since electricity is a non-storable commodity. For temperature, it does not make sense to talk about any trading in the underlying, which also makes the hedging arguments senseless. However, since the forward contracts need to have a price dynamics being arbitrage-free, we use (1.6) as a definition of the forward price, but now based on *any* equivalent probability measure Q. Recall that a risk-neutral probability turns all *tradeable* assets into martingales after discounting. Since both electricity spot and temperature are not tradeable in the usual sense, we are left with the bank account, which trivially becomes a martingale under any equivalent measure Q after discounting. In conclusion, we cannot pin down a unique forward price dynamics based on arbitrage arguments.

The *rational expectation hypothesis* in interest rate theory has also been considered in relation to forward prices in commodity markets. In this context, it says that the forward price is the best prediction of the spot price at delivery, or, in mathematical terms,

$$f(t,\tau) = \mathbb{E}\left[S(\tau) \,|\, \mathcal{F}_t\right] . \tag{1.9}$$

In view of (1.6), the rational expectation hypothesis means choosing $Q = P$ as the risk-neutral probability. In reality, it is not to be expected that the rational expectation hypothesis holds. The theory of normal backwardation argues that producers of a commodity will wish to hedge their revenues by selling forwards, and thereby willing to accept a discount on the expected spot price. Thus, in normal backwardation, we should have $f(t,\tau) < \mathbb{E}[S(\tau) \,|\, \mathcal{F}_t]$, saying that the hedgers are willing to pay a premium for getting rid of the spot price risk. The risk premium is defined as

$$\mathrm{RP}(t,\tau) \triangleq f(t,\tau) - \mathbb{E}\left[S(\tau) \,|\, \mathcal{F}_t\right] , \tag{1.10}$$

which is negative when the market is in normal backwardation. [Geman and Vasicek (2001)] find evidence of a positive risk premium in the Pennsylvania-New Jersey-Maryland (PJM) electricity market for contracts with a short time to maturity, and explain this by the market's aversion for the high volatility and thereby willingness to pay high prices to ensure delivery. For longer matured contracts, the sign of the risk premium changes in their study. [Longstaff and Wang (2004)] perform a non-parametric study of the PJM market, obtaining evidence of significant positive risk premium for the short-term contracts. Their study is extended by [Diko, Lawford and Limpens (2006)], who analyse risk premia in the three markets EEX, Powernext, and Dutch market APX. A term structure for the risk premium is found, which varies significantly from the short- to the long-term segment of the market. [Benth, Cartea and Kiesel (2006)] present a framework for explaining the sign of the risk premium in terms of the certainty equivalent principle and jumps in the spot price dynamics.

If the forward price is set under a risk-neutral probability Q as in (1.6), then the risk premium measures exactly the difference between the risk-neutral and the "market probability" predicitions. The choice of Q determines the risk premium, and opposite, having knowledge of the risk premium determines the choice of the risk-neutral probability. It is common to select a parametric class of risk-neutral probabilities, to explain the risk premium. These risk-neutral probabilities introduce a parametric change of the drift of the spot. To explain this further, suppose for simplicity that the spot is defined as a drifted Brownian motion

$$S(t) = \mu t + \sigma B(t)\,,$$

with $\sigma > 0$. Consider a change of measure given by the Girsanov transformation (see, for example, [Bjørk (1998)]). For a constant θ, there exists a probability Q equivalent to P such that

$$B^{\theta}(t) = B(t) - \theta t$$

is a Brownian motion under Q. Hence, we find

$$\begin{aligned}
\mathrm{RP}(t,\tau) &= f(t,\tau) - \mathbb{E}[S(\tau)\,|\,\mathcal{F}_t] \\
&= (\mu+\sigma\theta)\tau + \sigma\mathbb{E}_Q[B^{\theta}(\tau)\,|\,\mathcal{F}_t] - \mu\tau - \sigma\mathbb{E}[B(\tau)\,|\,\mathcal{F}_t] \\
&= \sigma\theta\tau + \sigma B^{\theta}(t) - \sigma B(t) \\
&= \sigma\theta(\tau - t)\,.
\end{aligned}$$

We see that the risk premium is positive if and only if θ is positive. It seems to be a common view that the risk premium is modelled as a change in the drift of the spot dynamics, or implictly, a Girsanov-type change of probability (see, for example, [Clewlow and Strickland (2000)]).

In this book we use the Esscher transform as the way to select risk-neutral probabilities. The Esscher transform is a parametric structure-preserving change of measure which generalizes the Girsanov transform for Brownian motion to a general II process. The drift of the spot dynamics will be changed using the Esscher transform, along with a change in the jump frequency and size, through possibly time-dependent parameters. These parameters are called the *market prices of risk*, and are closely linked to the risk premium. In normal backwardation, the risk premium is negative, equivalent to a negative market price of risk in the above context. However, for power commodities, the sign of the market price of risk may change depending on the time horizon in question. For instance, [Cartea and Williams (2006)] show that in the gas market, in the long-term the sign is positive, whereas in the short-term sign may change. [Weron (2005)] finds a changing sign of the market price of risk in the Nord Pool market, when considering Asian-style options, whereas [Cartea and Figueroa (2005)] argue for a negative market price of risk in the UK electricity market.

A complication in the electricity, gas and temperature markets is the fact that we have swap contracts traded, and not forwards that deliver the underlying energy at a fixed maturity time. Gas and electricity are flow commodities, in the sense that being long a swap ensures receiving a flow of the commodity over a specified time period. Similarly, temperature futures are contracts written on different temperature indices measured over specified periods like months or quarters of a year. To be able to calculate expressions for swap prices, we must constrain the class of models seriously if we want to avoid simulation-based pricing. As we shall see in a moment, the swap price is expressed through the average (or a weighted average) of the spot price over the delivery period. For exponential models like geometric Brownian motion or the Schwartz model, this may be difficult to calculate analytically. Arithmetic models may resolve this problem, and we shall discuss a class of such models for which the spot price is restricted to be positive. Interestingly, we cannot any longer expect to have the convergence of swap prices to the spot price when time to delivery approaches zero.

We discuss the pricing of electricity futures in more detail. As discussed in the previous subsection, the electricity spot price is strictly speaking *not* a continuous-time process. The reference price for the electricity futures contracts is given as the hourly price for electricity in the spot market, and therefore a time series model for the spot price dynamics should be used in determining the electricity futures price. Let us explore the consequences of this view by starting with the spot-forward relation in (1.6).

Consider an electricity futures contract with finanical delivery over a time interval $[\tau_1, \tau_2]$. The payoff from being long such a contract entered at time t is

$$\sum_{t_i=\tau_1}^{\tau_2-1} S(t_i) - (\tau_2 - \tau_1)F(t, \tau_1, \tau_2)\,,$$

where the electricity futures price is denoted $F(t, \tau_1, \tau_2)$ at $t \leq \tau_1$. In the electricity market, the futures price is customarily denominated in terms of currency per MWh, which means that the total amount paid is $F(t, \tau_1, \tau_2)$ times the length of the delivery period. The hourly prices between τ_1 and τ_2, with $\tau_2 - 1$ being the last hour before end of the delivery period, are denoted by $S(t_i)$. If we suppose that the electricity futures is settled at the end of the delivery period, the price will be defined through the relation

$$\mathrm{e}^{-r(\tau_2-t)}\mathbb{E}_Q\left[\sum_{t_i=\tau_1}^{\tau_2-1} S(t_i) - (\tau_2-\tau_1)F(t,\tau_1,\tau_2)\,|\,\mathcal{F}_t\right] = 0\,,$$

yielding

$$F(t,\tau_1,\tau_2) = \mathbb{E}_Q\left[\frac{1}{\tau_2-\tau_1}\sum_{t_i=\tau_1}^{\tau_2-1} S(t_i)\,|\,\mathcal{F}_t\right]. \tag{1.11}$$

If we consider an hourly spot price model as in (1.4), we need to calculate

$$F(t,\tau_1,\tau_2) = \mathbb{E}_Q\left[\frac{1}{\tau_2-\tau_1}\sum_{t_i=\tau_1}^{\tau_2-1} \mathbb{E}\left[\int_{t_i}^{t_{i+1}} \widetilde{S}(u)\,du\,|\,\mathcal{F}_{t_i}\right]\,|\,\mathcal{F}_t\right],$$

or, by appealing to $\mathcal{F}_t \subset \mathcal{F}_{t_i}$,

$$F(t,\tau_1,\tau_2) = \mathbb{E}_Q\left[\frac{1}{\tau_2-\tau_1}\int_{\tau_1}^{\tau_2} \widetilde{S}(u)\,du\,|\,\mathcal{F}_t\right]. \tag{1.12}$$

This integral formulation of the electricity futures price will be used throughout the book. Note that by interchanging the expectation and integration in (1.12), the electricity futures price can be viewed as an average of forward prices maturing over the delivery period.

Let us discuss briefly the consequences of modelling the hourly electricity spot prices directly as a time series without an underlying continuous-time process. Suppose that $S(t_i)$ is a time series defined at the hourly time moments t_i, $i = 0, 1, 2, \ldots, n$, where $t_0 = 0$ and $t_n = \tau_2$. Next, we assume that we have a discretely defined filtration $\mathcal{F}_{t_i}$ associated to the spot process. This is naturally enlarged to all times t by setting $\mathcal{F}_t = \mathcal{F}_{t_i}$ for $t \in [t_i, t_{i+1})$, which means that there is no new information coming from the spot price process before next time instance t_{i+1}. This implies that

$$\begin{aligned} F(t,\tau_1,\tau_2) &= \mathbb{E}_Q\left[\frac{1}{\tau_2-\tau_1}\sum_{t_i=\tau_1}^{\tau_2-1} S(t_i)\,|\,\mathcal{F}_t\right] \\ &= \mathbb{E}_Q\left[\frac{1}{\tau_2-\tau_1}\sum_{t_i=\tau_1}^{\tau_2-1} S(t_i)\,|\,\mathcal{F}_{t_i}\right] \\ &= F(t_i,\tau_1,\tau_2)\,. \end{aligned}$$

Hence, the electricity futures price becomes constant over each hour, that is, it becomes a time series process rather than a continuous-time stochastic process. This behaviour is not what we observe in the actual market, where electricity futures prices move according to trades taking place also within every hour. We mend this by introducing a continuous-time unobserved spot price process as above, which then introduces more information in the model.

We remark that temperature is naturally a continuous-time process, even though the indices used as underlying for the temperature futures contracts are discretely monitored. The gas spot can also be viewed as a process in continuous time, since one can buy a delivery of gas on short notice at a desired (in principle) time. This makes electricity as a rather particular case for the above discussion on discrete- vs. continuous-time models.

The HJM approach in the interest rate markets proposes to model the forward rates directly rather than the spot rates. This approach has been suggested to be used for modelling the forward price dynamics in commodity markets. In particular, [Bjerksund, Rasmussen and Stensland (2000)], [Keppo *et al.* (2004)], [Benth and Koekebakker (2005)] and [Kiesel, Schindlmayer and Börger (2006)] have done this for the contracts in the Nord Pool and EEX electricity markets, while a discussion of the approach to general energy markets can be found in [Clewlow and Strickland (2000)]. Note that both [Bjerksund, Rasmussen and Stensland (2000)] and [Clewlow and Strickland (2000)] suggest to use the HJM approach to model forward contracts, while in [Benth and Koekebakker (2005)] electricity futures, the actual contracts traded in the market, are considered. A large portion of this book is devoted to the application of the HJM approach, with a particular view towards the electricity markets. Some issues arise when trying to apply the HJM theory to electricity (and gas) futures.

First of all, what kind of contracts should the HJM approach be used on. Following the interest rate method directly, one may be tempted to model the forwards contracts, as it is done by [Bjerksund, Rasmussen and Stensland (2000)] and [Clewlow and Strickland (2000)]. However, in the electricity market we do not have data for such contracts, and the question of how to estimate the model to market observations arises. One way out is to smoothen the observed electricity futures prices, in order to transform the data to forward prices. Alternatively, one may integrate up the forward prices to get an implied dynamics for the observed futures. Instead of using the HJM technique on forwards that is not traded in the market, one may instead consider modelling the electricity futures directly.

In the Nord Pool market, electricity futures contracts with overlapping delivery periods are traded. For example, you can enter a yearly contract, or four quarterly contracts covering the entire year. Hence, certain consistency conditions need to be satisfied for the price dynamics of the contracts in order to avoid arbitrage opportunities. In the strict sense, the HJM method models the electricity futures price dynamics for arbitrary delivery periods. As we shall see, it is difficult to state models satisfying the arbitrage conditions and at the same time being analytically tractable. Furthermore, the condition rules out models like geometric Brownian motion.

To resolve this problem, we follow the path given by the LIBOR[12] models in interest rate theory (see, for example, [Brigo and Mercurio (2001)]). We model exclusively those contracts that are traded in the market, and in addition have delivery periods which cannot be decomposed into other traded contracts. With this way of modelling, we are much more free to state reasonable stochastic dynamical models which can easily be estimated on data and used for risk management analysis.

A possibly undesirable consequence of the HJM approach for electricity futures price modelling is the loss of a connection with the underlying spot price. Given an electricity futures price dynamics, one cannot trace back a spot price dynamics except in trivial and not relevant cases. This is a serious matter on one hand, since the spot is namely the reference index for the futures. On the other hand, one may view the electricity spot market as itself being a futures market, where the contracts have hourly settlement periods throughout the day.

1.6 Outline of the book

The basic modelling tools in this book are II processes and mean-reverting stochastic processes driven by these. In order to understand the models, and price products like swaps and options, we need a stochastic analysis for the II processes. The theory on stochastic integration and differentiation (for example, Itô's Formula) for such processes is surveyed in Chapter 2. The purpose of the chapter is to provide the reader an easy reference for the fundamental results and notions which will be useful in the modelling and pricing analysis of energy markets. The chapter is *not* meant to give a complete theory, for which the reader is referred to the existing literature in the area, for instance [Jacod and Shiryaev (1987)]. For us, the most useful results will be the Lévy-Kintchine representation, Itô's Formula, the

[12] London interbank offer rate.

stochastic Fubini Theorem and Bayes' Formula. However, to have a complete theoretical foundation for the analysis, we also need to understand stochastic integration with respect to II processes. The chapter includes examples of some of the most used stochastic processes in finance, and in particular energy markets. The reader being eager to process to the modelling and analysis of electricity and related markets, may skip reading this chapter and only use it for reference.

In Chapter 3 we model spot prices in energy markets based on OU processes. We analyse both geometric and arithmetic models, and present in particular an arithmetic model which preserves positivity of prices. The models are multi-factor, driven by both Brownian motion and pure jump processes, with possible seasonally dependent jump size and intensity. Stochastic simulation of these models is discussed in a case study of the arithmetic model.

Based on these spot models, we derive the forward and swap price dynamics in Chapter 4. We recall here that forwards in our use of the terminology are contracts with a fixed maturity time, whereas swaps are used as a general reference to electricity and gas futures. After some general considerations, we apply the Esscher transform to construct risk-neutral probabilities. The Esscher transform preserves the distributional properties of the jump processes, and can be thought of as a generalization of the Girsanov transform used for Brownian motions. Forward prices for the arithmetic and geometric spot models introduced in Chapter 2 are derived. For the swaps, the geometric models do not in general admit any explicit formulas for the price dynamics. Choosing an arithmetic spot model, we can derive an explicit swap dynamics. The issue of currency conversion, being relevant, for instance, on the Nord Pool market, is discussed in detail.

Our spot models and derived swap price dynamics are applied to the UK gas market in Chapter 5. A simple one-factor model with both Brownian motion and jump-driven increments are considered, a frequently used dynamics for energy spot prices. Recursive filtering is implemented for identifying the jumps in the spot price series. The heavy-tailed normal inverse Gaussian distribution for the spot price innovations is considered and estimated on data as well. The different spot price models are next used as a basis for deriving gas futures prices. We analyse the theoretical prices in view of the observed gas futures term structure in the UK market, and discuss the market price of risk, that is, the choice of the parameters in the Esscher transform. The chapter also contains discussions on how multi-factor spot models can be estimated on data, incorporating, for instance, OU processes having different speeds of mean reversion.

The HJM approach to the modelling of forward and swap prices is presented in Chapter 6. The different modelling issues regarding forward prices and swaps are investigated in detail, along with the incorporation of jump processes. As we show, the no-arbitrage condition for the term structure dynamics of the swap price rules out most of the relevant models. To resolve this issue, we introduce market models for the swaps, much in the spirit of LIBOR models for fixed income markets.

When applying the HJM approach to electricity markets, one may base the electricity futures price dynamics on a model for non-traded forwards. To estimate such models, one needs to derive forward data from the observed electricity futures prices. An algorithm for the derivation of smooth forward curves in electricity markets is presented in Chapter 7. The algorithm may be applied to gas markets as well. We demonstrate the algorithm at work on Nord Pool electricity futures data, and further apply it to study the term structure of volatility of electricity.

The smoothing algorithm is also applied in Chapter 8, where we empirically analyse the Nord Pool electricity futures market using HJM-based models. The smoothing algorithm enables us to derive a data set which is structured and more easy to use in an empirical investigation of the market. A principal component analysis reveals certain structures for the short- and long-term market, and motivate a parametric multi-factor market model, including seasonal volatility with maturity effect. The model is fitted to market data.

Following is a more theoretical chapter dealing with the pricing and hedging of options traded for energies. Chapter 9 presents pricing formulas for call and put options based on the various proposed spot, forward and swap models. The option prices become generalizations of the Black-76 formula when the underlying models are depending on Brownian motions only. For models with jumps, we use a Fourier approach to derive formulas for the prices. Issues of hedging are discussed for these options. The pricing of spread and Asian options are analysed for arithmetic multi-factor models, where reasonably explicit formulas are available based on the cumulant functions of the jump processes. A case study on the pricing of spark spread options in the UK market is presented, based on a direct modelling approach for the spread between electricity and gas.

The final Chapter 10 is devoted to the market for temperature futures. We present continuous-time mean reversion models being generalizations of autoregressive moving average time series. Applying these to temperature data, we find that the "volatility" of temperature has a clear seasonal pattern. The temperature models allow for rather explicit pricing of the typical futures traded on CME. The chapter includes a thorough empricial analysis of Stockholm temperature data in view of the proposed models.