UNIVERSITÄT DUISBURG ESSEN

Open-Minded

Calculation of Monte-Carlo Sensitivities for a portfolio of time coupled options and application to conventional power plants

A thesis submitted for the degree of

Dr. rer. pol. (rerum politicarum)

to the Faculty of Business Administration and Economics
at University of Duisburg-Essen

Wolfgang Raabe
Münster

Chair for Energy Trading and Finance
Essen, 2015

Amtierender Dekan:	**Prof. Dr. Volker Clausen**
Tag der mündlichen Prüfung:	**16.09.2015**
Erstgutachter:	**Prof. Dr. Rüdiger Kiesel**
Zweitgutachter:	**Prof. Dr. Christoph Hanck**

Ehrenwörtliche Erklärung:

Hiermit erkläre ich, dass ich die beigefügte Dissertation selbstständig verfasst und keine anderen als die angegebenen Hilfsmittel genutzt habe. Alle wörtlich oder inhaltlich übernommenen Stellen habe ich als solche gekennzeichnet.

Ich versichere außerdem, dass ich die beigefügte Dissertation nur in diesem und keinem anderen Promotionsverfahren eingereicht habe und, dass diesem Promotionsverfahren keine endgültig gescheiterten Promotionsverfahren vorausgegangen sind.

Contents

1 Introduction

Abstract

Current European energy markets are significantly influenced by a strongly growing share of highly volatile renewable electricity generation, not only changing the absolute level and structure of electricity prices but also leading to an increased demand for conventional generation flexibility to compensate supply and demand variations. Standardised reserve energy products provide an established way to trade such flexibility, however imposing additional operational constraints on the involved generation portfolios. This enforces the ability of utilities to derive values and sensitivities of their asset portfolios subject to external reserve requirements in order to manage market price risks and perform state of the art portfolio optimisation.

In this thesis we adopt the Proxy Simulation Scheme (PSS) method of Fries and Kampen (2007) *– originally developed with a focus on fixed income markets – for the rolling intrinsic valuation of stylised power plants subject to complex technical constraints. Thereby we succeed to overcome well known numerical performance issues of standard Monte-Carlo approaches and are able to derive robust Monte-Carlo portfolio sensitivities with respect to the underlying price of electricity of both first and second order (Δ and Γ). We employ electricity prices that are affected by a strong photovoltaic production to take into account the current reality of energy markets in Europe. To our knowledge this application of the PSS methodology to energy related real option valuation has not been presented in academic literature before.*

Based on this approach we are able to analyse the impact of technical constraints including minimum up- and down-time and externally imposed reserve requirements on the risk profile of stand-alone power plant options in detail. We confirm the quality of our results via backtesting with a Delta-Gamma hedging framework and a Taylor series approach to replicate single step probability densities of option values via numerically derived sensitivities.

Furthermore we evaluate and discuss a variety of power plant option portfolios including technically more flexible and inflexible portfolios as well as larger and smaller portfolios. Thereby we are able to analyse the impact of reserve requirements on different portfolios allowing us to provide a complete value and risk assessment of varying levels of reserve requirements in each portfolio context.

Finally we compare portfolio results of simplified and numerically more cost efficient option dispatching rules with the full rolling intrinsic approach as applied otherwise throughout this thesis.

1.1 Modern electricity markets require flexible generation assets

Commencing in the late 1980s in Europe the electricity market liberalisation enforced utilities to optimise their generation assets in a new market environment which is characterised by increasing competitiveness and highly volatile market prices for electricity and fuel commodities including hard coal, gas and EU emission allowances. In addition a second fundamental market change has been initiated in the late years of the last century. Driven by favorable subsidy regimes large amounts of renewable electricity generation capacity flooded the European energy markets that traditionally relied on centralised fossil fuel fired power plants. Figure 1.1 shows this development exemplary for the German market where already today the aggregated renewable capacity is in a range comparable to the system peak demand.

In contrast to controllable conventional plants the electricity production of renewable generation assets depends on exogeneous factors on different time scales. An unexpected fog or snow cover on photovoltaic modules typically leads to significant electricity supply variations within hours or even quarter hours while a complex weather pattern may lead to unpredicted wind fluctuations within days. An unexpectedly dry winter can impact the amount of hydro electric generation during an entire season and global meteorological phenomenons like El Niño (Glantz (2001)) or North Atlantic oscillation (Vicente-Serrano and Trigo (2011)) even affect renewable production rates of several years. The comparably volatile energy output from renewable generation assets impacts both system stability and the characteristics of electricity prices in central Europe to an increasing extend as the growth rate of renewable capacity is continuously high.

In combination with historically low coal prices putting pressure on wholesale electricity prices this leads especially in Germany to a situation, where a large share of the existing highly flexible conventional generation capacity (in particular gas fired power plants) is no longer economically viable – while the system is affected by large amounts of highly volatile electricity supply at the same time. The need of a stabilizing factor to ensure security of supply enforces the availability of dedicated controllable capacity that is able to quickly ramp up and down electricity generation in order to react to short term supply and demand fluctuations.

1.2 Reserve products provide an established way to trade flexibility

Similar to the evolution of flexible generation capacity financial markets have developed products that reflect this flexibility. An established measure to incentivise the provision of flexible assets is the auctioning of reserve capacity which is typically steered by the transmission system operators (e.g. regelleistung.net (2014)). In these markets multiple reserve products are frequently traded that supplement each other and typically act on different time scales (e.g. Swider (2006)). Primary and secondary

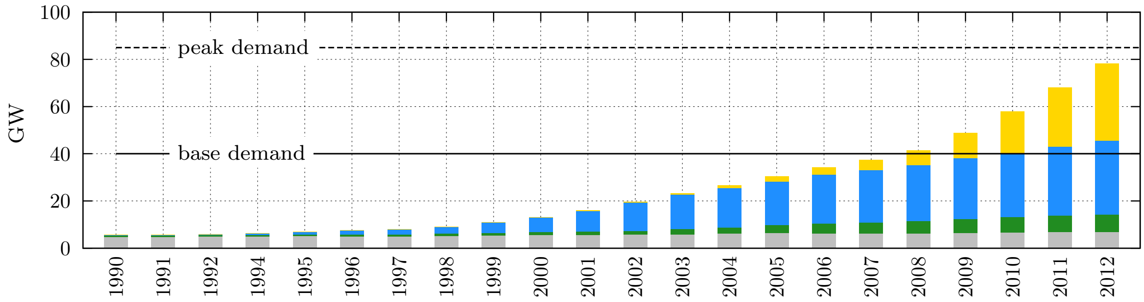

Fig. 1.1: Historical development of renewable generation capacity in Germany (gray: hydro, green: biomass & pre-mature technologies, blue: wind, yellow: photovoltaics) according to Arbeitsgruppe Erneuerbare Energien - Statistik (AGEE-Stat) (2013) in comparison to average base demand (Deutscher Bundestag (2012), Wirth (2014)) and peak demand (Winkler et al. (2013), Bundesministerium für Wirtschaft und Energie (BMWi) (2014)).

reserve capacity is exercised within seconds and in combination with tertiary reserve rebalances the system for a time period of up to one hour. Longer balancing requirements are then served via intraday trading or long term reserve capacity contracts for unplanned outages. Reserve products can usually be divided in two classes: positive reserve requires an additional production of electricity while negative reserve requires a reduction of electricity production.

A utility can offer pre-qualified controllable capacity into reserve auctions, typically at a pre-defined price for capacity and another price for produced electricity in case the asset is actually exercised. This imposes new constraints on the holder of the respective power plant portfolio: at each point in time the portfolio must be able to ramp up (or down) its electricity production according to the associated reserve contracts. As this usually has to be carried out in a short time frame, an asset providing reserve typically cannot be shut down and must be running at least at minimum stable load as up ramping would take too long. This obviously reduces the flexibility of the entire portfolio and therewith affects its value and its risk profile. An additional complexity arises as the impact of reserve requirements is not identical for all asset portfolios. Large portfolios are impacted not in the same way as smaller portfolios, rather inflexible assets show a different behavior than more flexible assets and the electricity production costs of individual assets also have a relevant impact.

1.3 Sensitivities are needed for generation portfolios subject to reserve requirements

In order to manage market price risks, be able to price reserve contracts adequately and to perform state of the art portfolio optimisation it is thus crucial for utilities to understand both value and risk profile of their asset portfolios including reserve requirements in detail.

The price risk profile of a portfolio of conventional power plants is most commonly expressed and managed via the portfolio sensitivities Delta (Δ) and Gamma (Γ). These sensitivities are defined as the first and second order derivative of the portfolio value with respect to an underlying price, which can be the forward price of electricity for a respective time period or any required tradable fuel type (e.g. hard coal, gas, oil or emission certificates). A profound knowledge of these sensitivities enables utilities to hedge their operating result against price fluctuations.

Technical constraints of physical assets require numerical approaches for the calculation of sensitivities

Usually the calculation of derivatives of the portfolio value with respect to an underlying requires knowledge of the portfolio value itself. In case of complex real options like conventional power plants simple discounted cash-flow (DCF) models cannot be used to derive a deterministic portfolio value as extrinsic value contributions (time value) need to be taken into account. This is confirmed by Brajkovic (2010) who compare DCF approaches with real option valuation explicitly for dark spread options (i.e. hard coal fired power plants) or by Borchert and Hasenbeck (2009) who compare DCF with real option valuation of a spark spread option (i.e. a gas fired power plant).

If closed-form expressions for the real option value of power plants are available they have the great advantage that sensitivities can either also be expressed by closed-form formulas or the numerical determination of derivatives is very cost effective as the underlying value is explicitly known. The academic literature provides many elegant approximations of the value of financial spread options, e.g. the famous work by Carmona and Durrleman (2003), closed-form values for spark spread options by Benth and Saltyte-Benth (2006) and closed-form solutions based on Fast Fourier Transformation by Carr et al. (1999) and Hurd and Zhou (2010). An entire risk assessment including sensitivities is published by Deng et al. (2008). Deng et al. (2001) work both on analytical value approximations of spark spread options and gas-fired power plants and Börger (2014) derive an analytical expression for the lower bound of multi-asset option prices including sensitivities. However in all cases the analysed options are rather stylised and do not include the full amount of technical requirements of real power plants.

In contrast the technically possible dispatch of physical power plants must fulfill a variety of technical constraints within the most prominent ones are minimum stable load, maximum load, minimum amount of time the plant needs to run after it has been started (minimum up-time), minimum time the plant needs to stay offline after it has been shut down (minimum down-time), start-up costs, ramp rates and must-run conditions. Figure 1.2 provides an exemplary dispatch of a constrained power plant options for a deterministic power price curve in order to demonstrate the impact of minimum up-time and load constraints. These plant specific constraints lead to temporal interdependences between dispatch decisions of the plant at different points in time (time coupling effects) which typically prohibit to specify closed form expression for the associated value or sensitivities. In addition a portfolio can be subject to external constraints affecting several power plants at the same time. This is especially true for external reserve requirements which must be fulfilled by the entire portfolio.

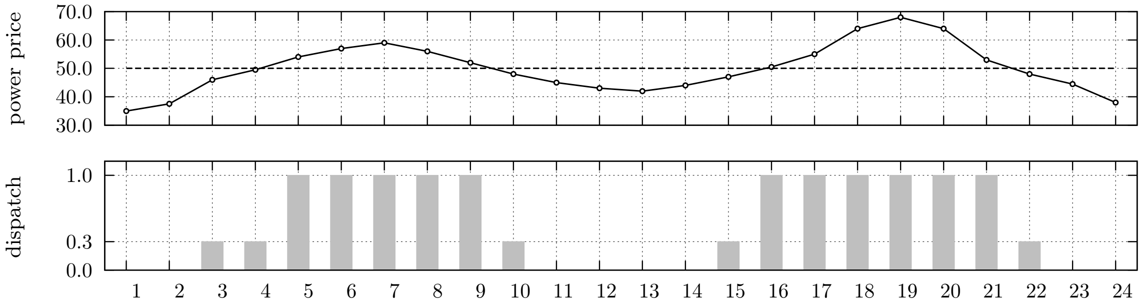

Fig. 1.2: Optimal dispatch of a technically constrained power plant options (minimum up-time of 8 hours, minimum load of 0.3, maximum load of 1.0) for a deterministic hourly power price curve quoted in €/MWh.

Such global portfolio constraints lead to interdependencies between dispatch decisions of different plants in combination with time coupling effects and further complicate approximations of values and sensitivities.

As closed-form solutions are usually not an option numerical approaches are required to derive the optimal dispatch of power plant portfolios. In the literature this task is well known as a version of the unit commitment problem (UCP), typically implying the solution of a mixed integer problem sometimes in combination with a stochastic dynamic programming framework. A good overview about the numerical implementation of different technical constraints is given in the following exemplary selection of academic work. An early paper has been presented by Yan et al. (1994) who optimise hydroelectrical generation assets. Rebennack et al. (2009) formulate short term UCPs for a public power service in the liberalised market as well as for the entire German electricity system (Rebennack et al. (2010)), however without explicit treatment of reserve requirements. Palmintier and Webster (2011) and De Vos and Driesen (2012) take reserve requirements into account and Vázquez (2006) even present an approach to derive the optimal amount of reserve capacity for a given system. Faria and Fleten (2011) use stochastic dynamic programming to determine an optimal bidding strategy of highly flexible assets into the day ahead Norwegian power market.

However in order to derive the expected value of a portfolio of power plant options it is mandatory to not only derive an optimal dispatch and payoff for one underlying price realisation but instead for multiple paths covering the entire space of possible price realisations. The classical Monte-Carlo (MC) approach in this context is to define an appropriate price process first, generate random price realisations for which dispatch and payoff are derived and calculate the expected value as the average of all payoff realisations. Hereby different ways are known to derive the dispatches, namely least squares Monte-Carlo, stochastic dynamic programming or rolling intrinsic valuation. All methods are regularly applied in current literature. Longstaff and Schwartz (2001) first proposed the concept of least squares Monte-Carlo for generic options before it was extended towards real option valuations (Boogert and de Jong (2007), Deng and Xia (2005)). Tseng and Barz (2002) use stochastic dynamic programming for technical short-term generation asset valuation while Eichhorn et al. (2010) combine the framework with so-called risk functions. Bjerksund et al. (2011) use a Heath-

Jarrow-Morton (HJM) price dynamic in combination with rolling intrinsic optimisation in order to derive the expected value of gas storage facilities.

Especially the rolling intrinsic approach is highly flexible, able to cover various price processes of different complexity as well as technical constraints while remaining intuitive at the same time. It essentially replicates the behavior of a trader who bases each next dispatch decision only on the optimal intrinsic dispatch of his assets in the context of currently available market information. Hence the rolling intrinsic valuation approach is not biased by the well-known perfect foresight error and has a high acceptance within both academic community and practitioners. For these reasons we use this approach throughout this thesis.

Standard finite difference Monte-Carlo sensitivities are restricted by insufficient numerical performance

The standard Monte-Carlo approach to calculate sensitivities uses the rolling intrinsic framework in a first step to derive the option portfolio value for a given set of forward prices. In a second step the same framework is applied to calculate portfolio values for a series of perturbed initial forward curves where one individual forward price has been shifted upwards or downwards. In a third step finite difference (FD) approximation is used to derive estimations of the first and second derivative of the portfolio value with respect to each shifted forward price. These are approximations of Δ and Γ. The approach is sometimes called bumping the model.

Assuming the usage of central differences this method requires the calculation of $2m + 1$ Monte-Carlo portfolio values for m sensitivities. In the context of portfolio management for power plants the number of requested sensitivities can easily be in the order of days per month, weeks per quarter or even more, as usually sensitivities are required for all underlying prices including various fuels. In addition the price of electricity may consider base load and peak load products which further increases the complexity. Each of the $2m + 1$ portfolio values requires the calculation of n Monte-Carlo dispatch and payoff realisations which themselves include the solution of multiple mixed-integer optimisation problems in case a rolling intrinsic approach is applied. Taking into account the computational time needed for solving only one of these complex optimisation problems, the standard Monte-Carlo approach is typically not applicable to derive portfolio sensitivities up to second order (e.g. Γ) as it lacks the required numerical performance.

Also pathwise and likelihood ratio methods show significant limitations

The standard finite difference Monte-Carlo approach for the calculation of sensitivities is sometimes referred to as an indirect method as sensitivities are approximated via differences of Monte-Carlo values which have to be derived in a first step. This can lead to numerical instabilities in case of discontinuous payoff functions and the results are generally biased due to the FD approximation.

Two further classes of Monte-Carlo methods have been discussed extensively during the last years both trying to overcome some of the restrictions of the standard method: the pathwise method and the likelihood ratio method. Both methods can be characterised as direct methods as they aim on using direct estimators of sensitivities instead of ex-post approximations. Thereby re-simulations of shifted initial forward curves are not required and both methods show a numerical cost advantage with respect to standard FD Monte-Carlo. In addition both methods are able to provide unbiased estimators. However they differ significantly within their numerical approach.

Pathwise method The pathwise method starts from a known path-wise relation between the payoff of a financial product and the parameter for which its sensitivity shall be calculated. In case of Δ and Γ this parameter would be an initial forward price. From this relation it is generally possible to derive unbiased path-wise estimators of the requested sensitivities. These path-wise estimators allow then to derive MC estimators for both value and sensitivities of the financial product by using the same MC price and payoff realisations. A comprehensive overview about the pathwise method is provided by Glasserman (2003) and Giles (2007). Papatheodorou (2005) applies it for the valuation of exotic options while Giles and Glasserman (2006), Capriotti and Giles (2010) and Capriotti and Giles (2012) combine the method with adjoint derivative computations. Lyuu and Teng (2010) work on approximations for discontinuous payoff functions.

Even though this method has advantages there are also some significant limitations that hinder its usage for the calculation of sensitivities of portfolios of real options like power plants. To our knowledge there is usually no relation available between initial forward prices and path wise payoffs of these complex options, which the method requires for each individual option. This limits the method significantly as it would need to be combined with a FD approach on a path wise level leading effectively to the same amount of MC simulations as required for standard Monte-Carlo. Finally the pathwise method is typically not applicable for the estimation of second order derivatives like Γ (Boyle et al. (1997)).

Likelihood ratio method In contrast to the pathwise method the likelihood ratio method starts with known relations between probability densities of price realisations and the parameters for which sensitivities shall be derived (Broadie and Glasserman (1996)). Based on these relations weight functions are defined for all required sensitivities which – in combination with available payoff realisations – serve as direct Monte-Carlo estimators for the value of the sensitivities. Thereby the likelihood ratio method is effectively an application of Malliavin calculus (Fournié et al. (1999), Montero and Kohatsu-Higa (2003) and Benth et al. (2003)) which is characterised by optimal weighting functions (Benhamou (2003), Chen and Glasserman (2007)).

Similar to the pathwise method there is no need for re-simulations as both value and sensitivities can be derived by using the same Monte-Carlo realisations. The likelihood ratio method is able to deal with both second order derivatives and discontinuous payoff functions as it only relies on the associated probability densities which are typically smooth enough (i.e. Lipschitz continuous, Glasserman (2003) and Shapiro et al. (2009)). However this feature can also be a disadvantage whenever in depth analytical knowledge about the probability distribution is not available. Unfortunately this is commonly the case for complex energy related price processes. In addition the likelihood ratio method provides typically lower convergence rates than the pathwise method if both are applicable. For specific problems it is possible to combine both approaches (Giles (2008)) but to our knowledge there is no consistent solution available to deal with the above mentioned restrictions for a general application in the energy related context.

In practice solution for robust Monte-Carlo sensitivities: Proxy Simulation Scheme method

In this thesis we overcome these problems by using the Proxy Simulation Scheme (PSS) method as first presented by Fries (2005), Fries (2006) and Fries and Kampen (2007) for complex interest rate options in the context of the Libor Market Model. This method is effectively a likelihood ratio method applied on the level of the numerical implementation, i.e. after the system describing stochastic differential equation (SDE) has already been discretised and is thereby well suited for computational mathematics. At this stage all discrete step-wise probability density functions (PDF) are explicitly known per definition which allows to derive weight functions for all required sensitivities based on simple finite differences directly applied on these densities.

The Proxy Simulation Scheme method combines all advantages of the likelihood ratio method with a general applicability whenever a SDE has already been numerically implemented. The sensitivity weight functions can be calculated on the fly and similar to likelihood ratio and pathwise methods no additional Monte-Carlo realisations are needed for the calculation of an arbitrary number of sensitivities. Therefore the PSS method can be added to each existing Monte-Carlo valuation framework as an additional module which derives the associated weight functions for all requested sensitivities. Apart from including this module it is not necessary to change anything within the valuation algorithm itself which can be a great advantage for practitioners.

Due to the implied FD approach all PSS Monte-Carlo sensitivity estimators are generally biased. However as the finite differences are applied on step wise probability density functions which are usually sufficiently smooth, very small shift sizes can be chosen to minimise the bias while the Monte-Carlo convergence is fairly robust against changes of the shift size.

Since its development the PSS method has been extended for specific problem sets. Fries (2007b) combine the general method with a pathwise approach and Fries and Joshi (2008) include importance sampling for advanced handling of financial products with discontinuous payoff functions. Chan and Joshi (2011) focus on minimizing the variance of applied Monte-Carlo weights (also Chan and Joshi (2012a) and Chan and Joshi (2012b)).

1.4 Main contribution and results of this thesis

In the following paragraphs we give a focussed overview about the main results of this thesis. Section 1.5 complements this summary by a structural outline of all chapters.

Application of the Proxy Simulation Scheme method in an energy related framework

After an introduction into the key characteristics of the Proxy Simulation Scheme method and a comparison to standard Monte-Carlo approaches including finite differences, pathwise and likelihood ratio method (chapter 2), we explain in detail how the Proxy Simulation Scheme method – originally developed with a focus on fixed income markets – is implemented in order to derive approximative values and sensitivities (Δ and Γ) of stylised power plants in a portfolio context (chapter 3). Thereby we enable the reader to apply this method with minimal additional effort and close the current gap between more theoretically focused literature and the need of industry practitioners.

Our numerical implementation combines the PSS approach with a rolling intrinsic mixed integer optimisation valuation framework to derive dispatches and expected payoffs of portfolios of technically constrained power plant options subject to external reserve requirements (section 3.2.1). We apply an hourly structure of the price of electricity that is influenced by strong photovoltaic power generation which significantly reduces prices of midday hours (section 3.1.4).

We present a thorough validation of our numerical setup (chapter 4) which assesses the technical functionality of the mixed integer solver as well as expected features of the Proxy Simulation Scheme method including Monte-Carlo convergence rates and the dependency of results on shift sizes and the alignment of forward price paths.

To our knowledge this is the first work that applies the Proxy Simulation Scheme method in an energy related context. We show in detail how to derive robust sensitivities of power plant portfolios with respect to the underlying price of electricity and assess the value of flexibility of these portfolios. Our methodical approach is well documented and can be implemented as a robust add-on to existing Monte-Carlo pricing algorithm being used in the energy industry.

Impact of technical parameters and reserve requirements on stand-alone power plant options

As a preparation for more complex portfolio analyses we provide a comprehensive assessment of the impact of technical constraints and reserve requirements on the value and sensitivities Δ and Γ of stand-alone power plant options (chapter 5). Hereby we focus on variations of minimum up- and down-times, option strike prices and the amount of applied positive and negative reserve requirement. We link our observations to developments of average hourly option dispatches in order to create a thorough understanding of associated operating modes for different combinations of technical constraints and reserve requirements.

Even the single option setups can sometimes yield rather counter intuitive results. Higher minimum up- and down-times lead to lower weekday values and higher weekday Deltas of the power plant options due to increased technical constraints affecting the option dispatch. However the same increased constraints also lead to higher Gammas on weekdays (sections 5.1.1 and 5.1.2) – which is surprising as Gamma is often referred to as a measure for the embedded optionality (or flexibility) of a portfolio. The effect can be explained by a dispatch analysis revealing that blocks of consecutive hours where the respective power plant option is dispatched are closer at-the-money for higher minimum up- and down-times.

In contrast the impact of positive and negative reserve imposes a strict must-run condition on stand-alone options which intuitively leads to a general reduction of both value and Gamma (section 5.2). Delta is decreased for increased positive reserve requirement and increased for increased negative reserve requirement as the latter leads to a must-run condition at comparably higher absolute levels.

In order to assess the quality of the Proxy Simulation Scheme sensitivities we apply a full Delta-Gamma backtesting hedging framework (section 5.4) as well as a Taylor series based replication of probability densities of single step option value changes (section 5.5). Both approaches confirm the high quality of the PSS sensitivities and clearly demonstrate their superiority over a simplistic hedging approach based on expected dispatches instead of numerical Deltas.

Impact of reserve requirements on value and sensitivities of power plant option portfolios

We start our portfolio analysis with a transition from a stand-alone power plant option subject to a minor positive reserve requirement towards a portfolio of multiple identical options subject to the same reserve requirement by dividing the original option successively in two, three and five pieces (section 6.1).

On the portfolio level all results react very intuitively: the growing fragmentation leads to increasing portfolio values, decreasing portfolio Deltas and increasing portfolio Gammas due to the growing ability of the portfolio to adapt to the imposed must-run condition of the reserve requirement. In this case Gamma can clearly be interpreted as an indicator of the portfolio flexibility.

However the PSS method also provides values and sensitivities of all individual power plant options and thereby gives more insights into complex intra-option dependencies compared to aggregated portfolio results. Thereby we are able to observe negative Gamma results for certain options on weekend days (e.g. section 6.1.2) which is a surprising finding as long option positions are typically characterised by positive Gammas. This effect can be explained via distinct dispatch regimes that occur on days with comparably low prices and which merely depend on the realised price structure in combination with historical dispatch decisions. It is a first indication for a reserve induced ranking between different options within the same portfolio leading to a clear allocation of long and short option characteristics to individual options.

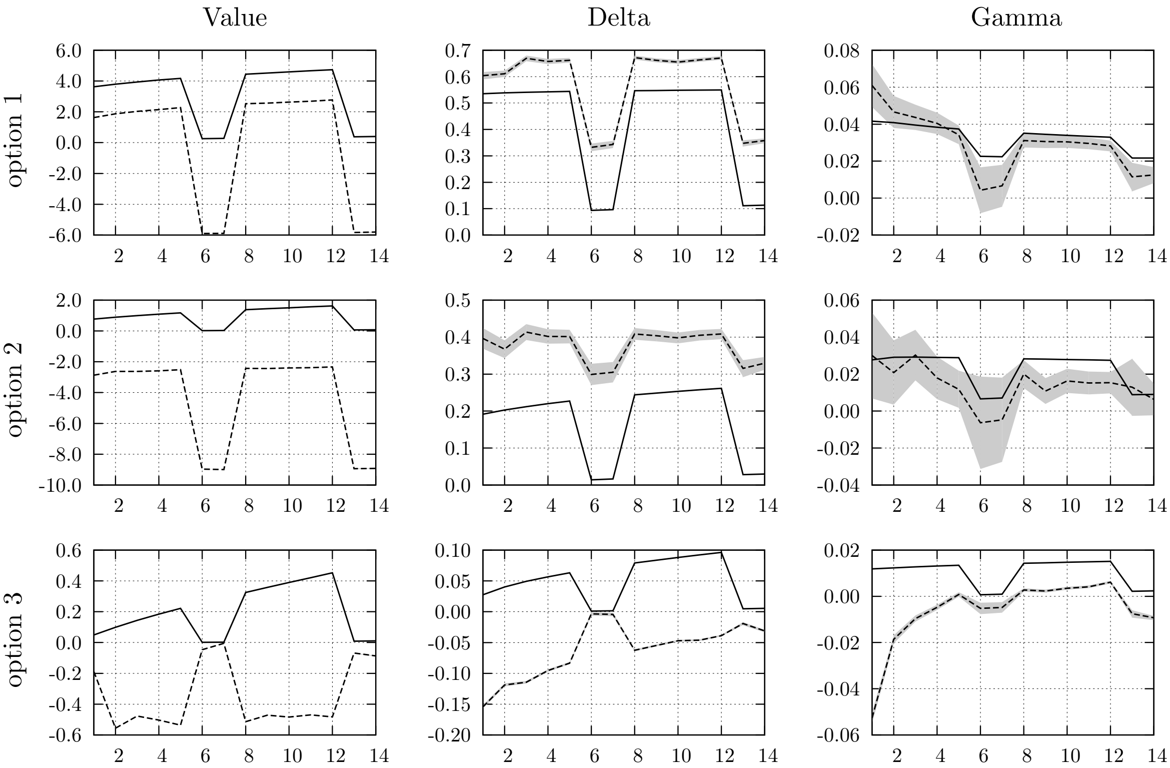

Fig. 1.3: Daily results for a time period of two weeks of all three options of the *unflexible* option portfolio subject to heavy reserve requirement $Res_{pos} = 1.0$ (dashed lines, gray areas indicate 95% confidence intervals) in comparison to analytical results of fully flexible options (solid lines). Option 1 has the lowest strike price and option 3 the highest strike price.

After this introduction to portfolio interdependencies we continue with an in depth analysis of three option portfolios that similarly include at-the-money and out-of-the money options but differ in the level of technical constraints of these options (section 6.2): the *flexible* portfolio includes options with comparably low minimum up- and down-times, the *split flexible* portfolio includes the same options but each split in two identical smaller parts and the *unflexible* portfolio contains options with significantly higher minimum up- and down-times. We compare results of all portfolios subject to an increasing level of positive reserve requirement in order to identify structural deviations between the different setups.

As long as the reserve requirements are small in comparison to the overall load flexibility the cheapest option of the flexible portfolio, which is at-the-money on weekdays, serves the major part of the reserve requirement (e.g. section 6.3.2). The other options are rarely affected only in hours with comparably high prices. This changes for a heavy reserve requirement where a clear pecking order between options with different strike prices can be observed (e.g. section 6.3.3). More expensive options are dispatched in the portfolio context only in case it is more economically attractive to re-allocate reserve serving from cheaper options to these more expensive options in order to run the cheaper options at full load. Thereby more expensive options effectively act as short options which are dispatched by cheaper options whenever this is beneficially in the portfolio context. This interdependency is confirmed by negative values, Deltas and Gammas of these options classifying them unambiguously as short option positions. The short option characteristic of the most expensive options is even more pronounced in the unflexible portfolio than in the flexible portfolio as stricter technical constraints lead to higher effective strike prices of these options (section 6.5). Figure 1.3 shows the general effect using the example of the unflexible portfolio subject to a heavy positive reserve requirement.

On a portfolio level the unflexible portfolio shows lower average weekday values in the absence of reserve requirements than the flexible portfolio due to higher technical constraints of all involved options. Remarkably even though starting from a lower value, the unflexible portfolio loses more value than the flexible portfolio when the reserve requirement is increased to a medium level (section 6.6). This changes for an additional increase to heavy reserve requirements where the unflexible portfolio looses significantly less value than the flexible portfolio until both portfolio values finally converge at a level of reserve requirement, that imposes a must-run condition on all individual options of both portfolios. This nonlinear development is explained by using the concept of online ratios, which denote the proportion of hours with comparably low prices when a certain option is running in the total amount of hours with comparably low prices.

On weekend days all options of the flexible and the unflexible portfolio are similarly deep out of the money. Hence any applied reserve requirement is served by the most cost effective options with the lowest strike prices, which are identical in both portfolios. Therefore the portfolios show similar results independently from different technical parameters of the individual options on weekend days.

In comparison to the flexible portfolio the split flexible portfolio usually yields higher values, lower Deltas and higher Gammas in case of applied reserve requirements as this more fragmented portfolio can better adapt to the imposed must-run conditions. However both portfolios show identical results for those levels of reserve that require the same amount of minimum load to stay online in both portfolios. This is an effect of our simplistic modeling setup but can nonetheless also be relevant for real generation portfolios and should be assessed in the course of investment or divestment decisions.

Dependency of value and sensitivities of reserve requirements in the portfolio context

By comparing portfolio results of all portfolios without reserve and subject to different reserve requirements we indirectly derive daily values and sensitivities (Δ and Γ) of different amounts of reserve requirement in all portfolio contexts (section 6.7).

Especially on weekdays the structural impact of reserve, and therewith the value and hedge parameters of the reserve itself, is similar in all portfolios (even though reduced in case of the split flexible portfolio): daily values are decreased, daily Deltas are increased while daily Gammas are also decreased. Remarkably this effectively means that a reserve requirement has negative daily values, positive daily Deltas and negative daily Gammas in our modeling setup. Therefore it cannot be replicated by adding a short power plant option position to the original portfolio in order "to reduce optionality", which could be an intuitive approach to approximate the impact of reserve requirements on a given physical generation portfolio. The value reduction can further be normalised by the amount of reserve requirements which directly yields the specific cost of reserve in different portfolios. Hereby it can be observed that the amount of reserve requirements that leads to the lowest specific cost of reserve is not identical for all portfolios as it heavily depends on the level of portfolio fragmentation.

Obviously the impact of reserve is heterogeneous and sometimes counter intuitive for different portfolios and depends non-linearly on the absolute amount of reserve requirement. Hence – even on the basis of our stylised modeling setup – it is highly recommended for practitioners to analyse the specific impact of reserve in detail before portfolio decisions are taken.

Assessment of simplistic dispatching rules instead of full rolling intrinsic valuation

Even though the Proxy Simulation Scheme method offers an advantage over standard Monte-Carlo approaches in terms of computational cost savings, long valuation tenors in combination with complex portfolio setups can still lead to numerical performance issues. Therefore we discuss the applicability of three alternatives for the derivation of portfolio sensitivities which focus on an additional reduction of computation time.

At first we use a truncated rolling intrinsic valuation tenor instead of the full valuation tenor within all time steps of each Monte-Carlo realisation (section 7.1). Secondly we apply a heuristic approach for the determination of the portfolio dispatches which is significantly faster than even truncated rolling intrinsic valuation (section 7.2). The third method is based on polynomial approximations of the results of a small number of Monte-Carlo realisations with negligible computational costs (section 7.3).

In case the truncated rolling intrinsic valuation tenor covers the typical length of time coupling effects in an option portfolio it leads to similar results as generated by the full rolling intrinsic tenor. Hereby we use a decay rate analysis of artificial perturbations of the dispatch of different power plant options to determine the typical lengths of time coupling. As an adequately truncated rolling intrinsic tenor can lead to significant computational cost savings it should always be considered for industry applications.

While simple dispatch heuristics hardly capture inter-option dependencies the polynomial sensitivities replicate daily result structures surprisingly well, even in case of extreme results for individual power plant options. However the sensitivities are slightly downwards biased which can be explained and consequently be corrected. Therefore this cost effective approach to derive sensitivities can be an interesting alternative for practitioners provided they are able to compensate the bias for their specific valuation problem.

1.5 Structural outline of this thesis

Chapter 2 In this chapter we present theoretical differences between the Proxy Simulation Scheme method and more standard Monte-Carlo methods including finite differences, pathwise and likelihood ratio method. Hereby we explicitly focus on a detailed comparison of advantages and disadvantages of all methods.

Chapter 3 Here we give a comprehensive overview about the numerical implementation of the Proxy Simulation Scheme method for the calculation of sensitivities of portfolios of stylised power plants. We cover the time discretisation of the underlying stochastic differential equations as well as the setup of the rolling intrinsic mixed integer valuation and the PSS extension of the standard valuation algorithm.

Chapter 4 This chapter provides a thorough validation of both technical aspects of our numerical implementation and theoretically expected characteristics of the Proxy Simulation Scheme sensitivities.

Chapter 5 This chapter contains a systematic analysis of the impact of technical parameters and reserve requirements on both value and sensitivities (Δ and Γ) of stand-alone power plant options. In addition we present backtesting results for our results which confirm the high quality and hedge effectiveness of the PSS sensitivities.

Chapter 6 Here we extend the analysis towards portfolios of multiple identical and different power plant options and explain the complex interdependencies of individual options in case of applied reserve requirements. Finally we derive daily values and sensitivities of various reserve requirements in the context of different portfolios.

Chapter 7 In this chapter we assess the applicability of alternative ways to derive portfolio sensitivities including simplified option dispatching rules via truncated rolling intrinsic and heuristic approaches and a semi-analytical Taylor series based method to derive sensitivities directly from payoff and price realisations.

2 Theoretical background on Monte-Carlo valuation methods

This chapter provides an overview about different Monte-Carlo techniques to calculate sensitivities. Starting from general Monte-Carlo theory it explicitly summarises advantages and disadvantages of Finite Difference (FD) Monte-Carlo and both pathwise (PW) and likelihood ratio (LR) method before comparing these approaches to the Proxy Simulation Scheme (PSS) method of Fries and Kampen (2007). *Hereby we closely follow* Glasserman (2003) *and refer to* Jäckel (2003) *and* Gerstner and Kloeden (2013) *for more detailed information.*

2.1 Monte-Carlo valuation

Assume an m-dimensional stochastic price process, which shall be described by the following generic stochastic differential equation (SDE):

$$\mathbf{dK}(t) = \boldsymbol{\mu}(\mathbf{K}, t)dt + \boldsymbol{\Sigma}(\mathbf{K}, t)\mathbf{dW}. \tag{2.1}$$

Here $\mathbf{dK} = (dK_1, dK_2, \ldots, dK_m)^T$ denotes the temporal evolution of the price vector $\mathbf{K} = (K_1, K_2, \ldots, K_m)^T$ subject to a drift rate $\boldsymbol{\mu} = (\mu_1, \mu_2, \ldots, \mu_m)^T$ and a set of correlated standard Brownian motions $\mathbf{dW} = (dW_1, dW_2, \ldots, dW_m)^T$ with correlation matrix Σ_{ij} for $i, j = 1, 2, \ldots, m$ (compare e.g. Deutsch (2004)).

Further assume a complex financial product whose payoff $f(Y)$ depends on

$$Y \coloneqq (\mathbf{K}(t)) \quad \text{with} \quad t \in [T_1; T_2], \tag{2.2}$$

where $[T_1; T_2]$ denotes a time interval which includes all potential trigger points for the specific financial product.

If $\mathbb{E}[\alpha|\mathcal{F}_{t_0}]$ denotes a risk-neutral expectation of a stochastic entity α at time $t = t_0$, the value $V(t_0) \coloneqq V$ of the financial product at $t = t_0$ is given by

$$V = \mathbb{E}[f(Y)|\mathcal{F}_{t_0}] = \int_\Omega f(y)\Phi(y)\mathrm{d}y, \tag{2.3}$$

where Ω denotes a measurable probability space equipped with filtration $\mathcal{F}_{t_0}$.

In case of complex options there may be no analytical solution available of the integral on the right hand side of equation (2.3) and a Monte-Carlo approach can be used to find a numerical value estimator. Hereby the integral is approximated by a weighted finite sum of discrete payoffs f for random realisations Y_i of Y with associated probability density function (PDF) Φ:

$$\hat{V}_n = \frac{1}{n}\sum_{i=1}^{n} f(Y_i). \tag{2.4}$$

According to the law of large numbers (Borel (1909)) $\hat{V}_n \to V$ with probability 1 as $n \to \infty$. In case f is square integrable the error $\hat{V}_n - V$ of the Monte-Carlo estimator is approximately normal distributed with mean 0 and standard deviation $\sigma_f/\sqrt{n}$, where σ_f is defined by

$$\sigma_f^2 = \int_\Omega \left(f(y) - V\right)^2 \Phi(y)\mathrm{d}y. \tag{2.5}$$

As V is typically unknown during the valuation process the same is also true for expression (2.5) and σ_f is usually approximated by the sample standard deviation

$$s_f = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(f(Y_i) - \hat{V}_n\right)^2}. \tag{2.6}$$

2.2 Monte-Carlo sensitivities

2.2.1 Finite difference (FD) approach

Assume that the value V depends on some parameter Θ (without loss of generality one initial forward price as elaborated in section 3.3), i.e.

$$V(\Theta) = \mathbb{E}[f(Y(\Theta))], \tag{2.7}$$

where we dropped the filtration $\mathcal{F}_{t_0}$ for better readability.

The first order sensitivity of V with respect to Θ is defined as the first derivative $\partial V(\Theta)/\partial\Theta$ (similar for higher order sensitivities). In finite difference Monte-Carlo this derivative is indirectly approximated by first calculating value estimators $\hat{V}_n(\Theta \pm h)$ for upwards and downwards shifted parameter Θ

$$\hat{V}_n(\Theta \pm h) = \frac{1}{n}\sum_{i=1}^{n} f(Y_i(\Theta \pm h)), \tag{2.8}$$

and then applying finite differences to approximate the derivative via the associated difference ratio

$$\frac{\partial V(\Theta)}{\partial\Theta} \approx \frac{\hat{V}_n(\Theta + h) - \hat{V}_n(\Theta - h)}{2h}. \tag{2.9}$$

Here we apply central differences and imply the usage of common random numbers for the calculation of $\hat{V}_n(\Theta \pm h)$.

By using Taylor expansion it can be shown that the approximation (2.9) has a general bias of order $O(h^2)$, heavily depending on the chosen shift size h. The associated optimal convergence rate is of order $O(n^{-2/5})$.

Using the same approach for the second derivative yields

$$\frac{\partial^2 V(\Theta)}{\partial \Theta^2} \approx \frac{\hat{V}_n(\Theta + h) + \hat{V}_n(\Theta - h) - 2\hat{V}_n(\Theta)}{h^2}, \tag{2.10}$$

having a similar bias of order $O(h^2)$ but a generally lower optimal convergence rate of order $O(n^{-2/7})$.

This very intuitive Monte-Carlo method has a couple of significant advantages:

- The method entirely relies on shifting initial parameters and re-evaluating the financial product (sometimes called *bumping the model*). Therefore it can be used with any existing Monte-Carlo valuation framework without changing anything within the valuation module itself.

- The method does not need any further details about the price process and the payoff distribution apart from what is required anyway for pricing the product. This makes the approach theoretically applicable without limitations for all potential problems.

But there are also disadvantages of the approach which can reduce its applicability for complex problems:

- FD sensitivities are subject to a general bias which is dependent on the size of the applied shift size h.

- In case of a discontinuous payoff function f the finite difference approximations (2.9) and (2.10) tend to numerical instabilities whenever a realisation Y_i is sufficiently close to the trigger point that causes the payoff function to jump.

- Applying central differences, the method requires at least two additional complete Monte-Carlo valuations for the calculation of each sensitivity, i.e. a calculation of n sensitivities requires $2n + 1$ full valuations. However in case of complex real options already a single valuation can lead to numerical performance problems, e.g. because the payoff realisations can only be determined by solving a time consuming mixed-integer optimisation problem. This leads to potential limitations of finite difference Monte-Carlo whenever multiple sensitivities are required for a complex financial product.

2.2.2 Pathwise sensitivities (PW)

Similar to the FD method also the pathwise method is based on an assumed relation between the payoff function f and a parameter of interest Θ. However, in contrast to indirect finite differences the pathwise method uses direct Monte-Carlo estimators for each simulated path to derive all required sensitivities in parallel to the valuation process. These estimators are calculated individually for each path, which effectively gives the method its name.

Starting from expression (2.9)

$$\frac{\partial V(\Theta)}{\partial \Theta} \approx \frac{\hat{V}_n(\Theta + h) - \hat{V}_n(\Theta - h)}{2h} = \frac{1}{n}\sum_{i=1}^{n} \frac{f(Y_i(\Theta + h)) - f(Y_i(\Theta - h))}{2h} \tag{2.11}$$

which depends on re-simulation, let

$$f'_{\Theta,i} := \lim_{h \to 0} \frac{f(Y_i(\Theta + h)) - f(Y_i(\Theta - h))}{2h}. \tag{2.12}$$

If this limit exists equations (2.11) and (2.12) can be combined to yield

$$\frac{\partial V(\Theta)}{\partial \Theta} = \frac{\partial}{\partial \Theta}\mathbb{E}[f(Y(\Theta))] \approx \frac{1}{n}\sum_{i=1}^{n} f'_{\Theta,i} = \mathbb{E}[f'_{\Theta}], \tag{2.13}$$

where $f'_{\Theta,i}$ are derivatives of the path wise payoffs f_i with respect to Θ.

If all required preconditions are met (f_i must be Lipschitz continuous) the interchange of differentiation and integration in equation (2.13) is justified and $\mathbb{E}[f'_{\Theta}]$ becomes an unbiased expected value of the derivative:

$$\frac{\partial V(\Theta)}{\partial \Theta} = \frac{\partial}{\partial \Theta}\mathbb{E}[f(Y(\Theta))] = \mathbb{E}\left[\frac{\partial f(Y(\Theta))}{\partial \Theta}\right] = \mathbb{E}[f'_{\Theta}]. \tag{2.14}$$

Obviously the pathwise method relies on a known relation between f and Θ in order to derive an analytical expression for f'_{Θ}. In case f explicitly depends on Y rather than on Θ (e.g. a plain vanilla Black-Scholes option) the chain rule can be applied to derive the path wise estimator via

$$f'_{\Theta} = \frac{\partial f}{\partial \Theta} = \frac{\partial Y}{\partial \Theta} \cdot \frac{\partial f}{\partial Y}. \tag{2.15}$$

When applicable the pathwise method has important advantages over the more simplistic finite difference Monte-Carlo schemes:

– In general the method allows the derivation of unbiased Monte-Carlo estimators.

– The availability of individual estimators for each required sensitivity allows to use only one set of Monte-Carlo simulations to derive both a price for a financial product and all its sensitivities at the same time without the need of re-simulations. Therefore whenever multiple sensitivities are needed for a complex product whose valuation is already a numerically challenge, the pathwise method provides a potentially huge numerical cost advantage over the FD method.

However the stricter preconditions of the pathwise method also lead to disadvantages:

– The implementation of a pathwise approach in an existing Monte-Carlo valuation framework is more complicated than in case of finite differences as typically available pricing modules need to be modified.

– The pathwise method does usually not allow to derive sensitivities for digital options as the path wise estimators are zero almost everywhere apart from a small interval around the trigger point of the associated payoff function, and therefore not meaningful. For the same reason the method is generally not applicable for the calculation of second order sensitivities which limits its applicability for the derivation of Γ.

– The method explicitly requires an expression of the pathwise derivative f'_Θ. In case of complex products this expression may include indicator functions and does not provide Lipschitz continuity, which conflicts with the preconditions of equation (2.13).

– To our knowledge there is no explicit expression f'_Θ available for complex time coupled real options like power plant options subject to technical constraints. Therefore re-simulations would be needed to calculate the pathwise estimators which would effectively lead to FD Monte-Carlo without any numerical cost advantage.

2.2.3 Likelihood ratio method (LR)

Suppose that Θ is a parameter of the underlying probability density Φ, i.e. $\Phi(y) \to \Phi_\Theta(y)$. Then equation (2.3) can be written as

$$\mathbb{E}[f(Y(\Theta))] = \int_\Omega f(y)\Phi_\Theta(y)\mathrm{d}y. \tag{2.16}$$

Suppose further that the order of differentiation and integration can be interchanged when the differential operator $\partial/\partial\Theta$ is applied on equation (2.16), which is usually possible as probability densities are sufficiently smooth functions (absolute continuity, e.g. Bingham and Kiesel (2004)):

$$\frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] = \int_\Omega f(y)\frac{\partial\Phi_\Theta(y)}{\partial\Theta}\mathrm{d}y. \tag{2.17}$$

An expansion of the integrand yields

$$\frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] = \int_\Omega f(y)\frac{\partial\Phi_\Theta(y)/\partial\Theta}{\Phi_\Theta(y)}\Phi_\Theta(y)\mathrm{d}y = \int_\Omega f(y)w_\Theta(y)\Phi_\Theta(y)\mathrm{d}y, \tag{2.18}$$

introducing the weight function

$$w_\Theta(y) = \frac{\partial\Phi_\Theta(y)/\partial\Theta}{\Phi_\Theta(y)} = \frac{\partial}{\partial\Theta}\ln\Phi_\Theta(y). \tag{2.19}$$

In equation (2.18) the average of $f(y)w_\Theta(y)$ serves as an unbiased estimator of $\frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))]$ that does only depend on unperturbed payoffs $f(y)$ which would also be used for the valuation process itself.

For the calculation of additional sensitivities only the weight function needs to be modified while the payoff realisations remain completely unchanged. Once all required weight functions are defined, they can be multiplied with payoffs of different financial products to derive all associated sensitivities at once.

In comparison to the pathwise method the likelihood ratio method provides a couple of additional benefits:

- Similarly to the pathwise method it allows the derivation of unbiased estimators.
- No additional Monte-Carlo simulations are needed to derive an arbitrary number of sensitivities for an arbitrary number of financial products in parallel to the valuation process.
- As the likelihood ratio method entirely relies on probability densities and their derivatives it generally allows to derive sensitivities of second and higher order and sensitivities for products with discontinuous payoff functions.
- The method can be added to any existing Monte-Carlo valuation algorithm in form of a module that derives weight functions and multiplies them to payoff realisations as already generated by the valuation part.

But its rather analytical approach also leads to some shortfalls:

- The likelihood ratio method requires explicitly known relations between the underlying probability densities and all parameters of interest. In case of complex price processes and products this analytical knowledge is often not available.
- In comparison to the pathwise method the likelihood ratio method shows typically lower convergence rates (e.g. Broadie and Glasserman (1996)).

2.2.4 Proxy Simulation Scheme method (PSS)

For the deduction of both pathwise and likelihood ratio method it was not relevant how the underlying price path realisations were generated. However in Monte-Carlo methods this is typically done via a numerical time discretisation scheme enabling the evolution of discrete prices or price forward curves from one point in time to the next. The Proxy Simulation Scheme method of Fries and Kampen (2007) makes explicitly use of this framework of perfect transparency by starting from an already implemented and therefore completely known time discretisation of the system describing stochastic differential equation.

Consider a time discretisation scheme which generates price path realisations Y^*. Combining this scheme with equation (2.3) yields

$$\mathbb{E}[f(Y(\Theta))] \approx \mathbb{E}[f(Y^*(\Theta))] = \int_{\Omega^*} f(y)\Phi^*_\Theta(y)\mathrm{d}y, \tag{2.20}$$

where "$\approx$" denotes potential deviations between analytical solutions of the model SDE and the numerical implementation due to the fact, that time steps cannot become infinitesimal but must stay finite.

Now consider a proxy scheme using a probability density function Φ^0 that does not depend on the parameter Θ. Following the likelihood ratio approach (section 2.2.3) assume absolute continuity between Φ^*_Θ and Φ^0, i.e.

$$\Phi^*_\Theta(y) > 0 \Longrightarrow \Phi^0(y) > 0 \tag{2.21}$$

for all $y \in \Omega^*$.

An expansion of equation (2.20) leads to

$$\begin{aligned}\mathbb{E}[f(Y^*(\Theta))] = \int_{\Omega^*} f(y)\Phi^*_\Theta(y)\mathrm{d}y &= \int_{\Omega^*} f(y)\frac{\Phi^*_\Theta(y)}{\Phi^0(y)}\Phi^0(y)\mathrm{d}y \\ &= \int_{\Omega^*} f(y)w^*_\Theta(y)\Phi^0(y)\mathrm{d}y \end{aligned}\tag{2.22}$$

with $w^*_\Theta(y) = \Phi^*_\Theta(y)/\Phi^0(y)$.

Applying the differential operator in combination with standard finite difference approximation on equations (2.20) and (2.22) yields

$$\begin{aligned}\frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] &\approx \frac{\partial}{\partial\Theta}\mathbb{E}[f(Y^*(\Theta))] \\ &\approx \frac{1}{2h}\left(\mathbb{E}[f(Y^*(\Theta+h))] - \mathbb{E}[f(Y^*(\Theta-h))]\right) \\ &= \int_{\Omega^*} f(y)\frac{1}{2h}\frac{\Phi^*_{\Theta+h}(y) - \Phi^*_{\Theta-h}(y)}{\Phi^0(y)}\Phi^0(y)\mathrm{d}y \\ &= \int_{\Omega^*} f(y)\frac{1}{2h}\left(w^*_{\Theta+h}(y) - w^*_{\Theta-h}(y)\right)\Phi^0(y)\mathrm{d}y \end{aligned}\tag{2.23}$$

with discretised weight functions

$$w^*_{\Theta\pm h}(y) = \frac{\Phi^*_{\Theta\pm h}(y)}{\Phi^0(y)}. \tag{2.24}$$

Equation (2.23) shows the Proxy Simulation Scheme method for a generic derivative with respect to an arbitrary parameter Θ. When explicit sensitivities Δ and Γ are derived, Θ is an initial forward price and the perturbations by $\pm h$ effectively change the initial conditions of the price evolution. In this case Φ^0 can be replaced by the probability density function of the original and unperturbed simulation scheme Φ^*_Θ which makes $\Phi^*_{\Theta\pm h}$ probability densities of initially perturbed schemes with identical realisations as generated by simulation scheme Φ^*_Θ:

$$\begin{aligned}\Delta = \frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] &\approx \int_{\Omega^*} f(y)\frac{1}{2h}\frac{\Phi^*_{\Theta+h}(y) - \Phi^*_{\Theta-h}(y)}{\Phi^*_\Theta(y)}\Phi^*_\Theta(y)\mathrm{d}y \\ &= \int_{\Omega^*} f(y)\frac{1}{2h}\left(w^*_{\Theta+h}(y) - w^*_{\Theta-h}(y)\right)\Phi^*_\Theta(y)\mathrm{d}y. \end{aligned}\tag{2.25}$$

Please note that in equation (2.25) $\Phi^*_\Theta(y)$ depends on Θ only in the way, that Θ serves as a fixed initial condition which is not perturbed at all. A more detailed explanation of the implementation of the Proxy Simulation Scheme method is provided in chapter 3. Similar to equation (2.10) the expression for the second order derivative Gamma is given by

$$\begin{aligned}\Gamma = \frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] &\approx \int_{\Omega^*} f(y)\frac{1}{h^2}\frac{\Phi^*_{\Theta+h}(y)+\Phi^*_{\Theta-h}(y)-2\Phi^*_{\Theta}(y)}{\Phi^*_{\Theta}(y)}\Phi^*_{\Theta}(y)\mathrm{d}y \\ &= \int_{\Omega^*} f(y)\frac{1}{h^2}\left(w^*_{\Theta+h}(y)+w^*_{\Theta-h}(y)-2\right)\Phi^*_{\Theta}(y)\mathrm{d}y. \qquad (2.26)\end{aligned}$$

The Proxy Simulation Scheme method combines all advantages of the likelihood method (section 2.2.3) with the possibility to calculate the weight functions on the fly as all step wise PDFs are known per construction. Therewith it allows to derive an arbitrary number of sensitivities for multiple products in parallel to the valuation itself without the need for re-simulation or additional information about the price process like closed form probability density functions.

Due to the FD approximation the results are generally biased, but as finite differences are applied to probability densities both shift size h and discretisation error can be kept very small.

3 Numerical implementation of the Proxy Simulation Scheme method

The purpose of this chapter is to provide both a comprehensive overview about the applied valuation model and a detailed explanation of the numerical implementation of the Proxy Simulation Scheme method as introduced in section 2.2.4. The underlying price process is chosen in a way that it allows a comparison of numerical results to analytical solutions for a limited number of artificial model setups. This is highly relevant for backtesting purposes and a thorough testing of the implementation.

The implementation entirely bases on Fortran 90 and available language bindings to the C based GNU Linear Programming Kit (GLPK (2014)).

3.1 Model SDE for daily electricity base load forward prices

The evolution of a forward curve $X_j(t) = X(t, T_j)$ with $T_j \in [t_0; T]$ and initial condition $X_j(t_0)$ shall be given by the m-dimensional system of SDEs

$$dX_j(t) = \mu_j(t)X_j(t)dt + \sigma_j(t)X_j(t)dW_j(t) \quad \text{with} \quad j = 1, 2, \ldots, m, \tag{3.1}$$

where $\mu_j(t)$ and $\sigma_j(t)$ denote the instantaneous drift and volatility of forward j and $\mathbf{W}(t) = (W_1(t), W_2(t), \ldots, W_m(t))$ is a m-dimensional set of Brownian motions with instantaneous correlation matrix $\mathbf{R}(t) = (\rho_{jk}(t))_{j,k=1,\ldots,m}$ and $dW_j(t) \cdot dW_k(t) = \rho_{jk}(t)dt$.

Under the risk neutral measure forwards can generally be considered to be martingales (i.e. driftless, compare Deutsch (2004) for a generic finance mathematical introduction and Benth and Koekebakker (2008) for an energy related discussion) and equation (3.1) reduces to

$$dX_j(t) = \sigma_j(t)X_j(t)dW_j(t) \quad \text{with} \quad j = 1, 2, \ldots, m. \tag{3.2}$$

With the help of Ito's lemma (e.g. Øksendal (2007)) and using log-coordinates, this can be written as

$$d\log(X_j(t)) \coloneqq dK_j(t) = -\frac{1}{2}\sigma_j^2(t)dt + \sigma_j(t)dW_j(t) \quad \text{with} \quad j = 1, 2, \ldots, m. \tag{3.3}$$

Furthermore, applying Cholesky decomposition on the correlation matrix $\mathbf{R}(t)$

$$\mathbf{L}(t) \cdot \mathbf{L}^T(t) = \mathbf{R}(t) \tag{3.4}$$

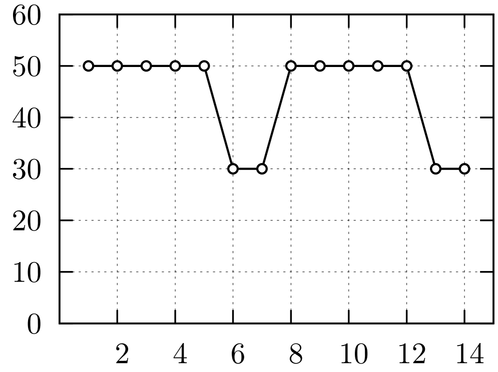

Fig. 3.1: Initial daily forward price curve for a valuation tenor of 14 days.

allows to decouple the system of SDEs (3.1), which can now be expressed in vector notation similar to equation (2.1) as:

$$d\mathbf{K}(t) = -\frac{1}{2}\mathbf{\Sigma}^2(t)\mathbf{1}dt + \mathbf{\Sigma}(t)\mathbf{L}(t)\mathbf{dU}(t)\sqrt{dt}. \tag{3.5}$$

Here, $\mathbf{dU}(t) = (dU_1(t), dU_2(t), \ldots, dU_m(t))^T$ denotes a vector of m independent standard normal distributed random numbers, $\mathbf{\Sigma}(t) = diag(\sigma_1(t), \sigma_2(t), \ldots, \sigma_m(t))$, $\mathbf{\Sigma}^2(t) = diag(\sigma_1^2(t), \sigma_2^2(t), \ldots, \sigma_m^2(t))$ and $\mathbf{1} = (1, 1, \ldots, 1)^T$.

In this thesis equation 3.5 shall describe the temporal evolution of a daily base load electricity forward price curve with a maximum tenor of 14 days. Valuation time for all simulations is $t_0 = 0$ and the associated initial condition $X_j(t_0)$ is given by a simplified generic forward curve with constant prices for weekdays and weekends (figure 3.1), where a typical historic price level is applied (EPEX SPOT SE (2014)).

3.1.1 Time discretisation

In order to solve the model SDE (3.5) numerically, the set of equations has to be discretised via an appropriate time discretisation scheme. Using a simple explicit Euler scheme (Grasselli and Pelinovsky (2008)) yields

$$\mathbf{K}^*(t_{q+1}) \approx \mathbf{K}^*(t_q) - \frac{1}{2}\mathbf{\Sigma}^{*2}(t_q)\mathbf{1}dt + \mathbf{\Sigma}^*(t_q)\mathbf{L}^*(t_q)\mathbf{\Delta U}(t_q)\sqrt{\Delta t}, \tag{3.6}$$

where discrete points of time t_q are defined via $t_{q+1} = t_q + \Delta t$ and $t_0 = 0$.

Throughout this thesis, an asterisk * denotes the numerical discretisation of an originally time continuous entity. In the applied framework of daily forward prices (section 3.1) Δt is always identical to a time step of one day, i.e. the forward curve K^* is evolved from day to day.

3.1.2 Volatility term structure

This thesis applies a term structure of volatility, which bases on a 2-factor volatility model as proposed by Börger (2007) which is structurally close to Schwartz (1997),

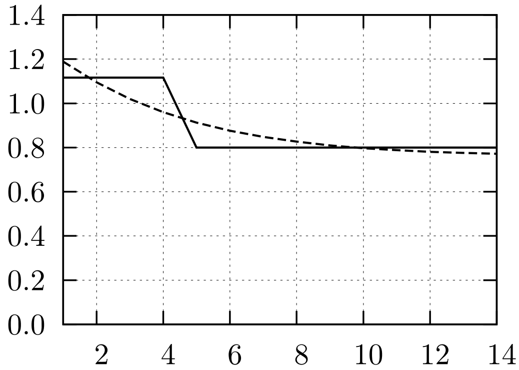

Fig. 3.2: Market volatility as taken from Benth and Koekebakker (2008) (solid line) and re-calibrated two factor volatility model of Börger (2007) as used in this thesis (dashed line).

Clewlow and Strickland (1999), Clewlow et al. (1999) and Schwartz and Smith (2000):

$$\boldsymbol{\sigma}(t,T) = \left(e^{-\kappa(T-t)}\hat{\sigma}_1, \hat{\sigma}_2\right). \tag{3.7}$$

However Börger (2007) use medium to long term forward contract prices to calibrate their model. As we focus on short term price evolutions this calibration cannot be suggested and we need to re-calibrate equation 3.7 with volatility data of traded short run forward contracts which are offered only by a small number of power exchanges (e.g. NASDAQ OMX (2014)) with recently increasing liquidity. A broad range of analyses with different targets and methodical approaches are provided by Pardo (2005), Vázquez et al. (2006), Vázquez et al. (2008), Härdle and Trück (2010), Mauritzen (2010), Bauwens et al. (2012) and Benth and Koekebakker (2008) where the latter explicitly use daily electricity forwards as a basis for the volatility term structure which is very close to the setup of this thesis.

Therefore we apply the data of Benth and Koekebakker (2008) to re-calibrate the model of Börger (2007) via least squares fitting and thereby to construct a most basic model for the short term volatility term structure (figure 3.2). The calibration yields $\kappa = 54.0$, $\hat{\sigma}_1 = 1.06$ and $\hat{\sigma}_2 = 0.76$. The model implies the following discrete term structure of volatility:

$$\sigma_j^*(t_q) = \sqrt{\hat{\sigma}_1^2 e^{-2\kappa(j-q)\Delta t} + \hat{\sigma}_2^2}. \tag{3.8}$$

3.1.3 Correlation

It is important to notice that the Proxy Simulation Scheme method does only allow to calculate sensitivities for those forwards $X_j(t_q)$ that represent individual risk factors. Generally this should not lead to any limitations as in a comprehensive model setup explicit sensitivities are only meaningful for the given set of individual risk factors.

However if sensitivities shall be derived for all m available forwards described by equation (3.6), this can only be achieved in case the correlation matrix $\mathbf{R}(t_q)$ is regular. This requirement can most easily be fulfilled by using a parameterisation method

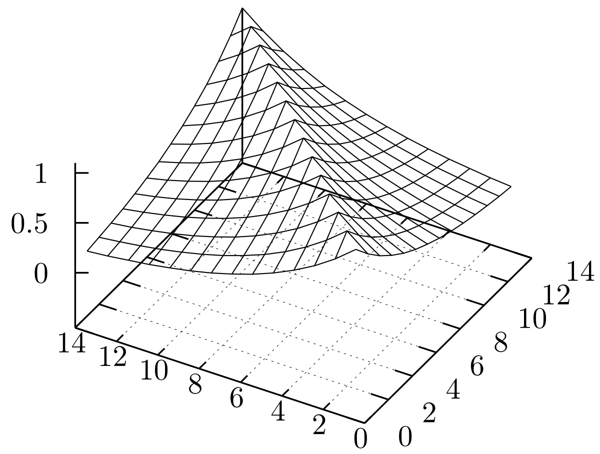

Fig. 3.3: 14×14 days correlation matrix as applied in this thesis.

adapted for the generation of regular matrices for building the correlation matrix. A good overview about available methods can be found in Brigo (2002), Börger and van Heys (2010) and Lutz (2010).

In this paper we use a modification of a standard parameterisation method which has originally been published by Schoenmakers and Coffey (2003):

$$\rho^*_{jk}(t_q) = \exp\left\{ \frac{|(j-q)-(k-q)|}{m-1} (-\log \hat{\rho}) \right\}, \tag{3.9}$$

with $1 < j \leqq k \leqq m$ and $\hat{\rho} = 10\%$ which leads to a basic correlation matrix that still allows for a huge variety of evolutions of the daily forward curve structure (figure 3.3).

3.1.4 Hourly price adjustment factors

Equations (3.6), (3.8) and (3.9) describe the daily evolution of an electricity forward price curve with daily granularity. In addition all European power exchanges offer day-ahead auctions of electricity delivery for all individual hours of the next day. The resulting hourly price structure typically serves as an important guideline for scheduling and final marketing decisions for power plants, even though physical plant dispatches may still be affected by short-term intraday supply and demand variations. In this thesis we apply hourly price adjustment factors to derive hourly prices via direct multiplication from previously defined daily base load forward prices.

Especially in Germany the shape of hourly prices is currently subject to a significant and sustainable change. Due to the increasing impact of photovoltaic electricity production the historically peak price structure during midday hours is almost completely eroded for days with specifically high solar radiation and replaced by comparably low midday prices in combination with increased prices in shoulder hours. This effect can already be observed in Hauser (2012). In the meantime the amount of installed photovoltaics capacity has been doubled in Germany (e.g. BSW (2013)) and this trend seems to be robust in the context of rather unchanged subsidy schemes for private

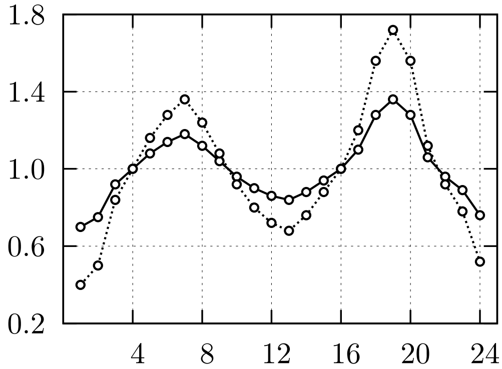

Fig. 3.4: Hourly curve adjustment factors (hca factors) for weekdays (solid line) and weekends (dotted line) as applied in this thesis.

customers (BMWI (2012)) and dramatically falling prices for photovoltaic modules (pvXchange (2014)).

In this thesis we assume an even more heavily photovoltaics driven and highly stylised hourly price structure, which could be observed in a future electricity system with high photovoltaic generation capacity backed by flexible but expensive conventional assets (figure 3.4). It is characterised by low peak load prices (i.e. prices at midday hours) and relatively increased prices in shoulder hours, where the latter is driven by a systematically tight market in these hours due to comparably low solar radiation and relevant demand. We assume this effect to be even more distinct at weekends as the generally lower peak demand may shorten the number of hours where costly conventional backup capacity is required, which further increases the specific cost of electricity production. Therefore we use an exaggerated weekday price structure for weekends where the hca factor α_{we}^l valid for hour l of a weekend day is derived from the respective hca factor α_{wd}^l via

$$\alpha_{we}^l = 2 \cdot \alpha_{wd}^l - 1. \tag{3.10}$$

Hourly prices will be derived by multiplication of daily base load forward prices $X_j^*(t_q)$ and the respective, pre-defined hourly curve adjustment factor $\alpha_{\iota(j)}^l$ of hour l where $\iota(j) = wd$ in case j denotes a weekday and we (for weekend) otherwise.

3.2 Monte-Carlo pricing

As stated in section 2.1 the structural approach of Monte-Carlo pricing is as follows:

1. generation of a sufficiently large number of price evolutions Y_i^*,

2. calculation of associated economically optimal payoffs $f(Y_i^*)$ of a financial product and

3. according to the law of large numbers the average payoff $\hat{V}_n$ converges against the true value of the product V a.s.:

$$\hat{V}_n = \frac{1}{n} \sum_{i=1}^{n} f(Y_i^*) \to V. \tag{3.11}$$

However in case of complex options like stylised power plants already the calculation of payoffs can be a numerical challenge due to a potentially large number of technical constraints (section 3.2.1) requiring a mixed integer optimisation to derive both optimal dispatch and associated payoff (section 3.2.2). Furthermore a naive calculation of price realisations and ex-post derivation of optimal dispatches and payoffs would systematically lead to upwards biased valuation results due to application of a super optimal exercise strategy (sometimes called the "perfect foresight error", e.g. Fries (2007a)). To avoid this bias we will use a rolling intrinsic valuation approach throughout this thesis (section 3.2.3).

3.2.1 Option portfolio setup

We evaluate a number of stylised power plants $p = 1, \ldots, N$ which are set up as financial options with minimum up- and downtime (t^p_{up}, t^p_{down}), minimum and maximum load (P^p_{min}, P^p_{max}) and individual strike prices κ^p. We treat all options as single commodity options being in the money whenever an hourly electricity price $\alpha^l_{\iota(j)} \cdot X_j$ exceeds the associated strike prices κ^p.

Thereby we focus on the calculation of sensitivities with respect to the electricity price and not on the derivation of more elaborated sensitivities like spread Deltas or Gammas (e.g. Burger et al. (2007)). However this is not critical in our opinion as the market standard in commodity risk assessment are still ceteris paribus sensitivities, where fuel prices are held constant anyway during the calculation of electricity hedge parameters Delta and Gamma.

Technical constraints of all options are close to values of real assets, i.e. up- and downtimes between 1 and 20 hours and strike prices between 50 and 70 €/MWh. The maximum load is scaled to unity apart from some examples where a reduced maximum load is explicitly mentioned. Minimum load is 30% of maximum load when not explicitly stated differently.

In addition to technical constraints the option portfolio can also be subject to an external reserve requirement. This can either be a positive reserve requirement Res_{pos} or a negative Res_{neg}. A positive requirement of x means that in each hour of the simulated time tenor the set of running power plant options needs to be able to produce an additional load of x, while a negative requirement of y means that in each hour of the simulated time tenor the running power plant options need to be able to reduce load by a total amount of y while staying online. This is both compliant with the standard approach in literature (e.g. Palmintier and Webster (2011)) and sensible, as only a limited number of very flexible physical power plants is able to serve reserve requirements from an idle state. In this thesis we focus on Res_{pos} and apply absolute levels of reserve requirements between 0.2 and 1.5, i.e. a share of 1/5 and 3/2 of the usual maximum load of one option.

3.2.2 Mixed integer optimisation for payoff calculation

As mentioned in the introduction to section 3.2 the calculation of dispatches and payoffs of power plant options is usually no trivial problem due to existing technical constraints and reserve requirements. Hereby technical constraints lead to "time coupling" effects, i.e. interdependencies between dispatch decisions of an option at different points in time, while reserve constraints have an impact on the portfolio level leading to additional coupling also between different options.

This setup can be formulated as a partially relaxed mixed integer optimisation problem (Burger et al. (2007)) thereby making the economically optimal hourly dispatch equivalent to the associated numerical solution of this problem. The optimal dispatch maximises the portfolio value function

$$V_{\text{Port}}(t_q, \boldsymbol{\beta}, \boldsymbol{s}, \boldsymbol{\gamma}) = \sum_{p=1}^{N} \sum_{j=q}^{m} \sum_{l=1}^{24} \left(\gamma_{j,l}^p P_{max}^p + (\beta_{j,l}^p - \gamma_{j,l}^p) P_{min}^p\right) \cdot \left(\alpha_{\iota(j)}^l X_j^*(t_q) - \kappa^p\right), \tag{3.12}$$

where $\beta_{j,l}^p \in \mathbb{N}$ and $s_{j,l}^p, \gamma_{j,l}^p \in \mathbb{R}$ and

$$\beta_{j,l}^p \in \{0, 1\}, \tag{3.13}$$

$$0 \leqq s_{j,l}^p \leqq 1, \tag{3.14}$$

$$0 \leqq \gamma_{j,l}^p \leqq 1. \tag{3.15}$$

$\beta_{j,l}^p = 1$ denotes that power plant p is running in hour l of day j (otherwise 0) and $s_{j,l}^p = 1$ indicates whether power plant p is started in hour l of day j (otherwise 0). $\gamma_{j,l}^p$ varies between 0 and 1 and defines the load of power plant p in hour l of day j.

In order to fulfill all technical requirements as explained in section 3.2.1, the maximisation of $V_{\text{Port}}(t_q, \boldsymbol{\beta}, \boldsymbol{s}, \boldsymbol{\gamma})$ is subject to the following set of constraints:

$$\beta_l^p - \beta_{l-1}^p \leqq s_l^p \quad \text{and} \quad \beta_l^p \geqq \gamma_l^p \tag{3.16}$$

guarantee that the decision variable s_l^p is set to 1 whenever a power plant option is started and that its associated load is either 0 or located in the interval $[P_{min}^p; P_{max}^p]$. Here and in the following we explicitly ignore the dependency on j to indicate that equation (3.16) enforces only a relation between consecutive hours l independent of the respective day j.

$$\beta_l^p \geqq \sum_{t=l-t_{up}^p+1}^{l} s_t^p \quad \text{and} \quad \beta_l^p \leqq 1 - \sum_{t=l+1}^{l+t_{down}^p} s_t^p \tag{3.17}$$

ensure that a power plant option is online for at least t_{up}^p hours after each start-up and offline for at least t_{down}^p hours after each shut-down.

In addition the portfolio may need to serve positive (or negative) reserve requirements Res_{pos} (or Res_{neg}), which is equivalent to the following two additional conditions being valid for each hour l of the valuation tenor:

$$\sum_{p=1}^{N} \beta_l^p P_{max}^p - (\gamma_l^p P_{max}^p + (\beta_l^p - \gamma_l^p) P_{min}^p) \geqq \text{Res}_{pos}, \tag{3.18}$$

$$\sum_{p=1}^{N} (\gamma_l^p P_{max}^p + (\beta_l^p - \gamma_l^p) P_{min}^p) - \beta_l^p P_{min}^p \geqq \text{Res}_{neg}. \tag{3.19}$$

Again, please note that in this setup only running power plants are able to serve reserve requirements.

The implementation bases on Fortran 90 and associated language bindings to the C based GNU Linear Programming Kit (GLPK (2014)) as provided by Kelly (2014).

3.2.3 Rolling Intrinsic valuation to reduce perfect foresight error

While calculating the Monte-Carlo approximation of an option value it is generally not desirable within each Monte-Carlo simulation step to first calculate price realisations for the entire valuation tenor and then to derive the associated optimal dispatch. A concrete example within this thesis would be to first deriving electricity price realisations for all $14 \cdot 24$ hours of the valuation tenor and then calculating the associated optimal dispatch of the option portfolio for these price realisations.

The reason against this approach is obvious: the method implicitly assumes that all future price realisations are already known when the optimal dispatch is calculated and therefore effectively assumes a perfect foresight of the future. In reality dispatch decisions at a certain point in time are always subject to rather limited information including historical price realisations, the current status of the option portfolio and forward prices which may serve as estimators of future price realisations. Therefore the perfect foresight method will systematically overestimate option values by using more information for the dispatch calculation as available in reality (perfect foresight bias due to a violation of measurability requirements, e.g. Fries (2007a)).

In order to reduce this bias we follow a popular Monte-Carlo approach, the so called rolling intrinsic valuation (e.g. Breslin et al. (2009)). This approach replicates a trader who defines the portfolio dispatch at each point in time only for the next point, based on an intrinsic optimisation of the entire remaining valuation tenor considering only the current state of information.

In this thesis the applied price process is a daily price forward curve which is evolved also on a daily basis. Consequently the rolling intrinsic valuation is applied as follows: On day t_q all daily forward prices $\boldsymbol{X}^*(t_q)$ are known for the remaining valuation tenor $[t_{q+1}; t_m]$ and hourly forward prices are given by multiplication of $\boldsymbol{X}^*(t_q)$ with the associated hourly curve adjustment factors α^l. Solving the hourly mixed integer optimisation problem (equation 3.12) for the remaining tenor provides the optimal intrinsic dispatch only based on currently available information. The resulting hourly dispatch for day t_{q+1} serves as final dispatch decision for all hours of this day. Now the valuation time is rolled from t_q to t_{q+1} including the temporal evolution of all daily forward prices within the remaining valuation tenor $[t_{q+2}; t_m]$. By solving the optimisation problem for this reduced tenor the final dispatch decision for all hours of day t_{q+2} is subsequently defined. The method is proceeded until the dispatch of the last day of the valuation tenor t_m is finally derived.

Please note that even though this method is very popular due to its intuitive and very flexible approach it is numerically rather costly as it requires the solution of m mixed integer optimisation problems for each Monte-Carlo simulation step. Depending on the complexity of the mixed integer problem even pure valuation can therefore lead to performance issues.

3.3 Proxy Simulation Scheme sensitivities (Delta and Gamma)

As explained in section 2.2.4 the Proxy Simulation Scheme framework allows an expression of the first order derivative (Delta) via ratios of the probability density of the implemented price evolution scheme Φ^*_Θ and the probability densities of the same scheme with upwards and downwards shifted initial conditions $\Phi^*_{\Theta\pm h}$. These ratios are called weight functions $w^*_{\Theta\pm h}$ as they are used in the Monte-Carlo expression of equation (2.25) to weight the portfolio payoff, therewith providing a direct estimator of the derivative Delta:

$$\begin{aligned}\Delta_\Theta = \frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] &\approx \int_{\Omega^*} f(y)\frac{1}{2h}\frac{\Phi^*_{\Theta+h}(y)-\Phi^*_{\Theta-h}(y)}{\Phi^*_\Theta(y)}\Phi^*_\Theta(y)\mathrm{d}y \\ &= \int_{\Omega^*} f(y)\frac{1}{2h}\left(w^*_{\Theta+h}(y)-w^*_{\Theta-h}(y)\right)\Phi^*_\Theta(y)\mathrm{d}y \\ &\approx \frac{1}{n}\sum_{i=1}^{n} f(Y^*_i)\cdot\frac{1}{2h}\left(w^*_{\Theta+h,i}(Y^*_i)-w^*_{\Theta-h,i}(Y^*_i)\right). \end{aligned} \qquad (3.20)$$

At least from a theoretical point of view this expression is very intuitive. However, a proper implementation of the concept requires a detailed understanding of how the path wise weights $w^*_{\Theta+h,i}(Y^*_i)$ and $w^*_{\Theta-h,i}(Y^*_i)$ are defined.

3.3.1 Numerical calculation of the PSS weights

In case of calculating portfolio sensitivities Δ and Γ the parameter Θ is identical to one component of the initial daily forward price curve $\mathbf{X}^*(t_0)$. Assume

$$\Theta := X^*_k(t_0 = 0). \qquad (3.21)$$

In the following we use the upwards shifted scenario as an example to explain the approach in detail.

In order to derive $w^*_{\Theta+h,i}(Y^*_i)$ consider an upwards shifted simulation scheme $\mathbf{X}^{up}(t_q)$, whose only difference to $\mathbf{X}^*(t_q)$ is given by modified initial conditions

$$\mathbf{X}^{up}(t_0) = \mathbf{X}^*(t_0) + h \cdot \mathbf{e}_k. \tag{3.22}$$

Here $\mathbf{e}_k$ denotes a unit vector in direction $k \leq m$. Equation (3.22) essentially defines a shift-scenario of the original simulation scheme. The simple central finite difference approximation of a sensitivity with regard to the kth component of $\mathbf{X}^*(t_0)$ would be to price the financial with an upwards and a downwards shifted simulation scheme, subtract the results and divide by two times the shift size h. In contrast, the Proxy Simulation Scheme framework enables to calculate both prices as expected values of differently weighted realisations of the basis scheme $\mathbf{X}^*(t_q)$ (compare section 2.2.4). The shifted schemes can be seen as *virtual* schemes, as they will never physically be simulated.

Equation (3.22) can be written in log-coordinates:

$$\mathbf{K}^{up}(t_0) = \mathbf{K}^*(t_0) + \delta^{up}_k \cdot \mathbf{e}_k, \tag{3.23}$$

with

$$\delta^{up}_k = \log(X_k(t_0) + h) - \log(X_k(t_0)). \tag{3.24}$$

According to section 2.2.4 the associated weight is defined by

$$w^*_{\Theta+h,i}(Y^*_i) = \frac{\Phi^*_{\Theta+h,i}(Y^*_i)}{\Phi^*_{\Theta,i}(Y^*_i)}. \tag{3.25}$$

This weight can be interpreted as a pathwise relation between the probability densities of one realisation Y^*_i of $\mathbf{K}^*(t_q)$ in the upwards shifted simulation scheme $\mathbf{K}^{up}(t_q)$ (numerator) and in the realisation-generating basis simulation scheme $\mathbf{K}^*(t_q)$ (denominator). Consequently, the weight needs to be recalculated for each simulation path i that leads to an individual realisation of Y^*_i.

The evaluation of the denominator $\Phi^*_{\Theta,i}(Y^*_i)$ is straight forward as it simply describes the probability density of one possible realisation of the implemented simulation scheme $\mathbf{K}^*(t_q)$ in exactly the same scheme. Therefore, it is identical to the product of the probabilities of $\mathbf{K}^*_i(t_q)$ $(q = 1, 2, \ldots, m)$ leading to $Y^*_i = (\mathbf{K}^*_i(t_1), \mathbf{K}^*_i(t_2), \ldots, \mathbf{K}^*_i(t_m))$. It is obvious from equation (3.6), that these probabilities are identical to the probabilities of the standard normal distributed multivariate vector $\mathbf{\Delta U}_i(t_q)$ $(q = 1, 2, \ldots, m)$ and therefore

$$\Phi^*_{\Theta,i}(Y^*_i) = \prod_{q=1}^{m} \left((2\pi)^{-(m-(q-1))/2} \exp\left(-\frac{1}{2} \mathbf{\Delta U}^T_i(t_q) \cdot \mathbf{\Delta U}_i(t_q) \right) \right). \tag{3.26}$$

Without loss of generality, assume as a starting point for the derivation of the numerator $\Phi^*_{\Theta+h,i}(Y^*_i)$ a shift of the first element of $\mathbf{K}^*(t_0)$, i.e.

$$\mathbf{K}^{up}(t_0) = \mathbf{K}^*(t_0) + \delta^{up}_1 \cdot \mathbf{e}_1. \tag{3.27}$$

Than $\Phi^*_{\Theta+h,i}(Y^*_i)$ can be interpreted as the transition probability density of a movement between t_0 and t_1, which makes the perturbed scheme $\mathbf{K}^{up}(t_q)$ identical to the base scheme $\mathbf{K}^*(t_q)$ after the first time step:

$$\begin{aligned}\mathbf{K}^{up}(t_0) &= \mathbf{K}^*(t_0) + \delta_1^{up} \cdot \mathbf{e}_1 \\ \mathbf{K}^{up}(t_q) &\stackrel{!}{=} \mathbf{K}^*(t_q) \quad \text{for} \quad q = 1, 2, \ldots, m.\end{aligned} \tag{3.28}$$

Bear in mind, that $\mathbf{K}^{up}(t_q)$ can be seen as a virtual numerical simulation scheme. Asking for the transition probability of the above mentioned movement is therefore effectively asking for the probability density of a vector of standard normal distributed random variables $\mathbf{\Delta U}^{up}(t_0)$ leading to the requested movement:

$$\begin{aligned}\mathbf{K}^{up}(t_1) &= \mathbf{K}^{up}(t_0) - \frac{1}{2}\mathbf{\Sigma}^{*2}(t_0)\mathbf{1}\Delta t + \mathbf{\Sigma}^*(t_0)\mathbf{L}^*(t_0)\mathbf{\Delta U}^{up}(t_0)\sqrt{\Delta t} \\ &= \mathbf{K}^*(t_0) + \delta_1^{up} \cdot \mathbf{e}_1 - \frac{1}{2}\mathbf{\Sigma}^{*2}(t_0)\mathbf{1}\Delta t + \mathbf{\Sigma}^*(t_0)\mathbf{L}^*(t_0)\mathbf{\Delta U}^{up}(t_0)\sqrt{\Delta t} \\ \stackrel{!}{=} \mathbf{K}^*(t_1) &= \mathbf{K}^*(t_0) - \frac{1}{2}\mathbf{\Sigma}^{*2}(t_0)\mathbf{1}\Delta t + \mathbf{\Sigma}^*(t_0)\mathbf{L}^*(t_0)\mathbf{\Delta U}(t_0)\sqrt{\Delta t} \\ \Longleftrightarrow \mathbf{\Delta U}^{up}(t_0) &\stackrel{!}{=} \mathbf{\Delta U}(t_0) - \frac{1}{\sqrt{\Delta t}}\mathbf{L}^{*-1}(t_0)\mathbf{\Sigma}^{*-1}(t_0)\delta_1^{up} \cdot \mathbf{e}_1.\end{aligned} \tag{3.29}$$

In case of a perturbation of the kth element of $\mathbf{K}^*(t_0)$, i.e. $\mathbf{K}^{up}(t_0) = \mathbf{K}^*(t_0) + \delta_k^{up} \cdot \mathbf{e}_k$, the schemes will not be equalised within the first time step, but linearly until the kth time step via a stepwise amendment of amount δ_k^{up}/k:

$$\mathbf{\Delta U}^{up}(t_q) \stackrel{!}{=} \mathbf{\Delta U}(t_q) - \frac{1}{\sqrt{\Delta t}}\mathbf{L}^{*-1}(t_q)\mathbf{\Sigma}^{*-1}(t_q)\frac{\delta_k^{up}}{k} \cdot \mathbf{e}_k \quad \text{for} \quad q < k. \tag{3.30}$$

The requested transition probability density $\Phi^*_{\Theta+h,i}(Y^*_i)$ can then be derived similarly to (3.26) as the product of the associated probability densities for $\mathbf{\Delta U}_i^{up}(t_q)$ $(q < k)$:

$$\Phi^*_{\Theta+h,i}(Y^*_i) = \prod_{q=1}^{k}\left((2\pi)^{-(k-(q-1))/2}\exp(-\frac{1}{2}\left(\mathbf{\Delta U}_i^{up}\right)^T(t_q)\cdot\mathbf{\Delta U}_i^{up}(t_q))\right). \tag{3.31}$$

The transition probability density in a downwards shifted scheme $\mathbf{K}^{down}(t_q)$ can be derived analogously to (3.31) by using $\mathbf{\Delta U}_i^{down}(t_q)$ $(q < k)$ instead of $\mathbf{\Delta U}_i^{up}(t_q)$ with

$$\mathbf{\Delta U}^{down}(t_q) \stackrel{!}{=} \mathbf{\Delta U}(t_q) + \frac{1}{\sqrt{\Delta t}}\mathbf{L}^{*-1}(t_q)\mathbf{\Sigma}^{*-1}(t_q)\frac{\delta_k^{down}}{k} \cdot \mathbf{e}_k \quad \text{for} \quad q < k. \tag{3.32}$$

Here, $\delta_k^{down} = \log(X_k(t_0) - h) - \log(X_k(t_0)$.

The associated probability density of a downwards shifted scheme can finally be derived as follows:

$$\Phi^*_{\Theta-h,i}(Y^*_i) = \prod_{q=1}^{k}\left((2\pi)^{-(k-(q-1))/2}\exp(-\frac{1}{2}\left(\mathbf{\Delta U}_i^{down}\right)^T(t_q)\cdot\mathbf{\Delta U}_i^{down}(t_q))\right). \tag{3.33}$$

3.3.2 Explicit PSS weight expressions for Delta and Gamma

Knowing $\Phi^*_{\Theta+h,i}(Y_i^*)$, $\Phi^*_{\Theta-h,i}(Y_i^*)$ and $\Phi^*_{\Theta,i}(Y_i^*)$, all weights needed for the calculation of sensitivities with regard to an element of $\mathbf{X}^*(t_0)$ (i.e. Δ and Γ) are also known and the sensitivities can be derived similar to equations (2.25) and (2.26) as

$$\begin{aligned}\Delta_\Theta = \frac{\partial}{\partial\Theta}\mathbb{E}[f(Y(\Theta))] &\approx \frac{1}{n}\sum_{i=1}^{n} f(Y_i^*)\cdot\frac{1}{2h}\left(\frac{\Phi^*_{\Theta+h}(Y_i^*)}{\Phi^*_\Theta(Y_i^*)} - \frac{\Phi^*_{\Theta-h}(Y_i^*)}{\Phi^*_\Theta(Y_i^*)}\right)\\ &= \frac{1}{n}\sum_{i=1}^{n} f(Y_i^*)\cdot\frac{1}{2h}\left(w^*_{\Theta+h,i}(Y_i^*) - w^*_{\Theta-h,i}(Y_i^*)\right) \qquad (3.34)\\ \Gamma_\Theta = \frac{\partial^2}{\partial\Theta^2}\mathbb{E}[f(Y(\Theta))] &\approx \frac{1}{n}\sum_{i=1}^{n} f(Y_i^*)\cdot\frac{1}{h^2}\left(\frac{\Phi^*_{\Theta+h}(Y_i^*)}{\Phi^*_\Theta(Y_i^*)} + \frac{\Phi^*_{\Theta-h}(Y_i^*)}{\Phi^*_\Theta(Y_i^*)} - 2\right)\\ &= \frac{1}{n}\sum_{i=1}^{n} f(Y_i^*)\cdot\frac{1}{h^2}\left(w^*_{\Theta+h,i}(Y_i^*) + w^*_{\Theta-h,i}(Y_i^*) - 2\right). \qquad (3.35)\end{aligned}$$

3.3.3 Structural extension of the pricing algorithm

One of the most important advantages of the Proxy Simulation Scheme method is, that it can be used as an add-on module in combination with any already existing Monte-Carlo pricing algorithm (compare sections (2.2.3) and (2.2.4)). The associated structural approach is shown in figure 3.5.

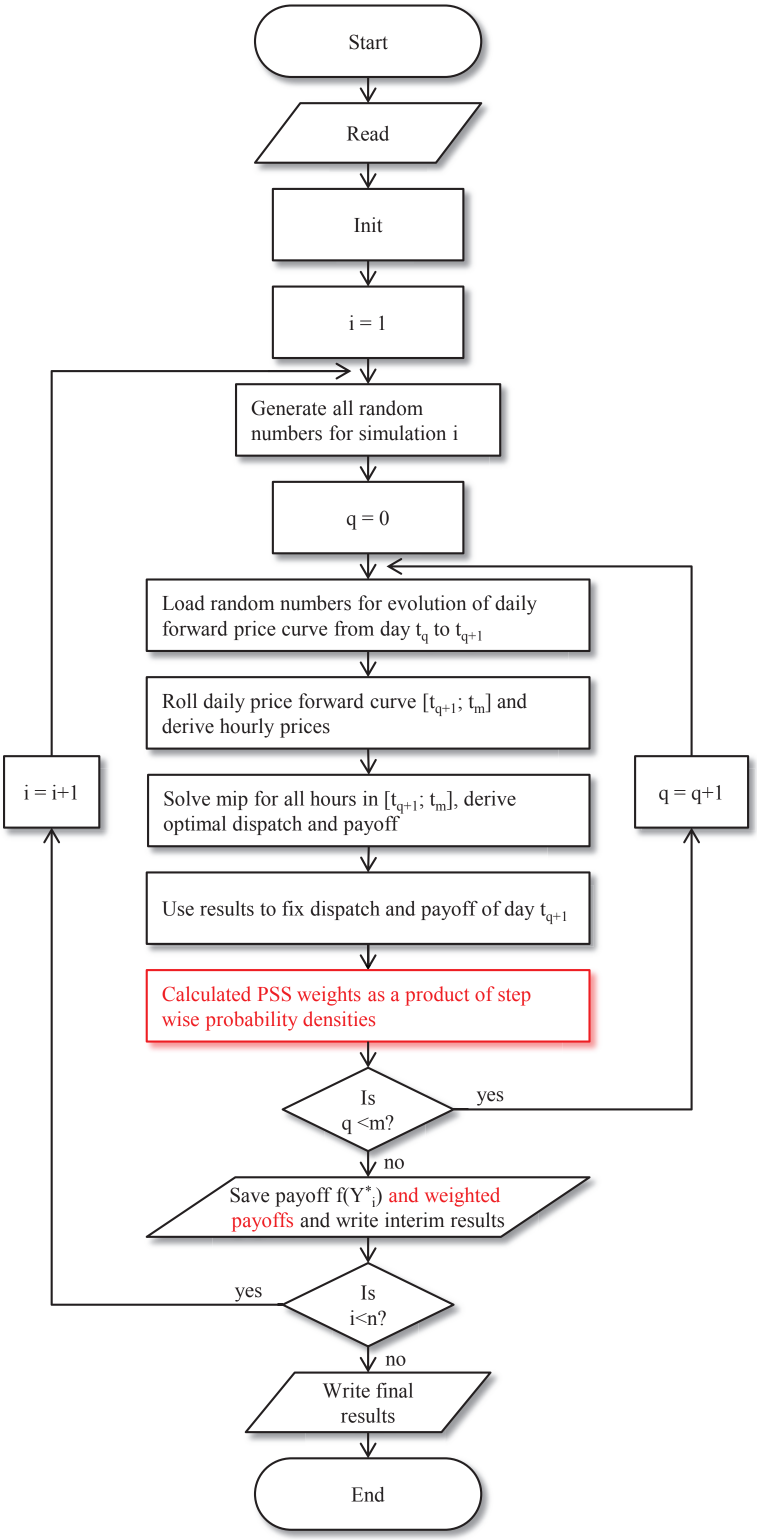

Fig. 3.5: Flowchart of a Monte-Carlo pricing algorithm including an add-on Proxy Simulation Scheme module (highlighted in red).

4 Code validation

This chapter provides a detailed validation of the numerical implementation of the Proxy Simulation Scheme method as discussed in chapter 3 beginning with an introduction to standard result types being derived and analysed throughout the upcoming chapters. In addition it presents some key characteristics of the chosen setup of the Proxy Simulation Scheme method which may be of interest especially for practitioners who intend to adapt the method.

4.1 Normalisation and units of result types

As stated in section 3.2.1 the size of evaluated power plant options is usually normalised to unity in this thesis (apart from a limited number of exceptions where the different setup is explicitly pointed out). Due to the construction of applied risk factors (compare section 3.1.3) all results including value, Delta and Gamma are generally available in daily granularity where each daily result type is effectively an aggregation of the underlying 24 hourly results.

In order to provide intuitive and easy to generalise numerical results we will divide all result types by a factor of 24 h. This approach has the great advantage that it provides results which can directly be linked to normalised daily options. This is especially true for Delta, now quoted in MW and identical to the amount of daily base load products that are required to replicate the payoff of the original option via Delta hedging. Therewith it is directly comparable to Deltas of well-known financial products like the famous Black-Scholes option (Black and Scholes (1973)). Gamma cannot be interpreted that intuitively but is fully consistent in our approach in providing the correct derivative of Delta with respect to the underlying price of electricity as quoted in €/MWh.

This approach will modify the units of our results as follows:

$$\text{Value:} \quad \left[\mathbb{E}[f(Y(\Theta))]\right] = € \implies €/\text{h} \tag{4.1}$$

$$\text{Delta:} \quad \left[\frac{\partial \mathbb{E}[f(Y(\Theta))]}{\partial \Theta}\right] = \text{MWh} \implies \text{MW} \tag{4.2}$$

$$\text{Gamma:} \quad \left[\frac{\partial^2 \mathbb{E}[f(Y(\Theta))]}{\partial \Theta^2}\right] = \frac{(\text{MWh})^2}{€} \implies \frac{(\text{MW})^2\text{h}}{€}. \tag{4.3}$$

In the following sections we will skip the explicit indication of result units for simplicity and better readability.

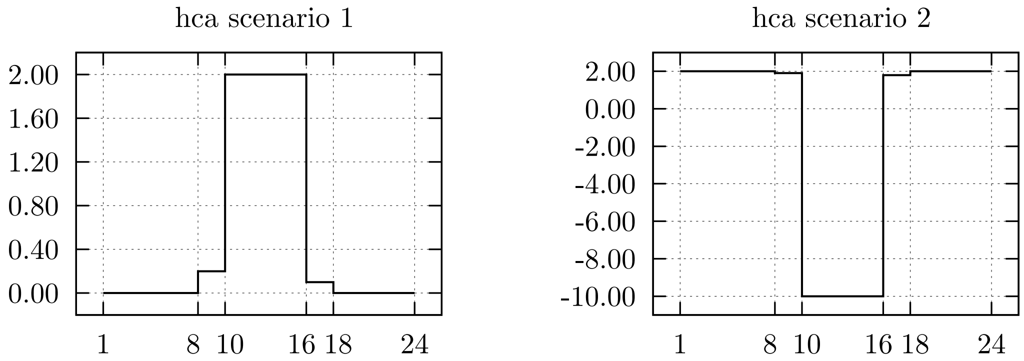

Fig. 4.1: Hourly curve adjustment (hca) factor scenarios as applied for testing purposes.

4.2 Validation of mixed integer solver implementation

In order to confirm a correct implementation of the applied mixed integer solver GLPK (2014) and associated language bindings for Fortran 90 (Kelly (2014)) in the overall valuation framework we use two extreme setups of hourly curve adjustment (hca) factors for 4 weekdays leading to power plant option values that can analytically be backtested.

Both scenarios are shown in figure 4.1. The first scenario is characterised by zero prices in all off-peak hours and a plateau of very high constant peak prices between hour 10 and hour 16 which is enclosed by two hour blocks of different, intermediate price levels. The second scenario provides the opposite price structure including very high off-peak prices and a highly negative peak price plateau which is enclosed by two hour blocks of different, positive price levels.

For power plant options whose strike prices correspond to a hca factor of 1.0 and minimum uptimes and downtimes (compare section 3.2.1) in a range between 1 and 10 hours this extreme price framework only allows a single optimal dispatch solution per option setup – even in case of reserve requirements. This dispatch leads to associated daily numerical option values. However a known dispatch solution also enables to express the option values analytically by values of strips of hourly call options in the risk neutral framework (Black and Scholes (1973)). By comparing these two sets of values we are able to test all technical components of our valuation algorithm in parallel, including price evolution scheme, mixed integer optimisation module and writeout routines.

Figure 4.2 provides average hourly dispatches for hca scenario 1 and two different option setups: 1. a technically constrained option ($t_{up} = 11$ and $t_{down} = 1$) without reserve requirement and 2. a fully flexible option ($t_{up} = t_{down} = 1$) subject to positive reserve requirement ($Res_{pos} = 0.2$). Both simulations use $n = 100\,\text{k}$ Monte-Carlo realisations.

Obviously the fully flexible option runs at minimum load (0.3) for all hours except the high priced period where it runs at 0.8, which corresponds to maximum load (1.0) minus the level of positive reserve requirement (0.2). This is fully consistent with the theoretical expectation.

In contrast to this digital dispatch pattern the technically constrained option shows

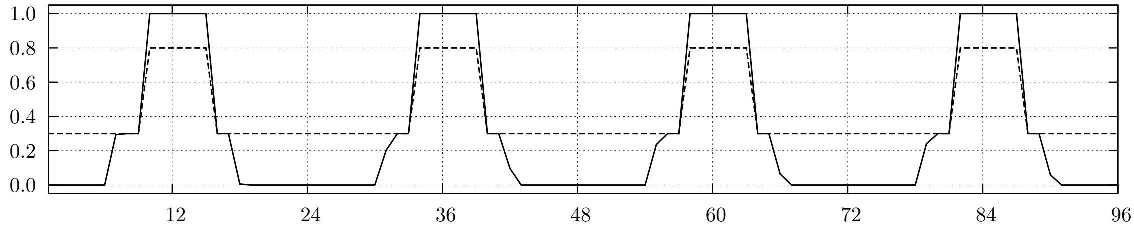

Fig. 4.2: Average hourly dispatch for all 4 days of hca scenario 1, (solid line: $t_{up} = 11$, $t_{down} = 1$, $Res_{pos} = 0.0$, dashed line: $t_{up} = t_{down} = 1$, $Res_{pos} = 0.2$). $n = 100\,\text{k}$.

hca scenario	t_{up}	t_{down}	Res_{pos}	Res_{neg}	analytical daily value	numerical daily value	relative deviation
1	6	1	0.0	0.0	12.5000	12.5000	-1.2E-06
1	8	1	0.0	0.0	11.5000	11.5001	6.1E-06
1	9	1	0.0	0.0	10.9375	10.9376	5.5E-06
1	11	1	0.0	0.0	9.7500	9.7500	-3.5E-06
1	1	1	0.2	0.0	-0.8750	-0.8749	-6.1E-0.5
1	1	1	0.0	0.2	-5.6250	-5.6252	2.8E-05
1	1	1	0.2	0.2	-8.1250	-8.1250	-4.3E-06
2	1	6	0.0	0.0	36.2500	36.2497	-7.0E-06
2	1	8	0.0	0.0	32.9167	32.9170	9.6E-06
2	1	9	0.0	0.0	31.0417	31.0417	7.5E-08
2	1	10	0.0	0.0	29.1667	29.1667	6.9E-07

Table 4.1: Theoretical and numerical daily results (normalised by factor 1/24h) for all hca scenarios and applied test combinations of technical constraints and reserve requirements.

a more complex behavior with increasingly smooth outer ramps at later days. This is due to the fact that only 10 subsequent hca factors are non zero while all other hca factors are identical to zero. The minimum up-time constraint of 11 hours requires the option to be online for 11 subsequent hours. In finding the optimal dispatch the mixed integer solver is indifferent against dispatching either hour 7 or hour 18 as both have zero hca factors and therefore zero prices. This leads to a 1 hour shifting freedom of the in the money (itm) block. However the aggregated expected dispatch of hours 7 and 18 is identical to the minimum load of the option which demonstrates that within each simulation realisation either hour 7 or hour 18 is dispatched - which is fully consistent with the applied technical constraint.

Table 4.1 provides an overview about all test scenarios including the already discussed two cases. Hereby only average daily results are shown which is a valid reduction as the extreme hca scenarios lead to identical values for all 4 days anyway. The numerical results are very close to analytical values for all scenarios with and without reserve requirements and for options with different technical constraints.

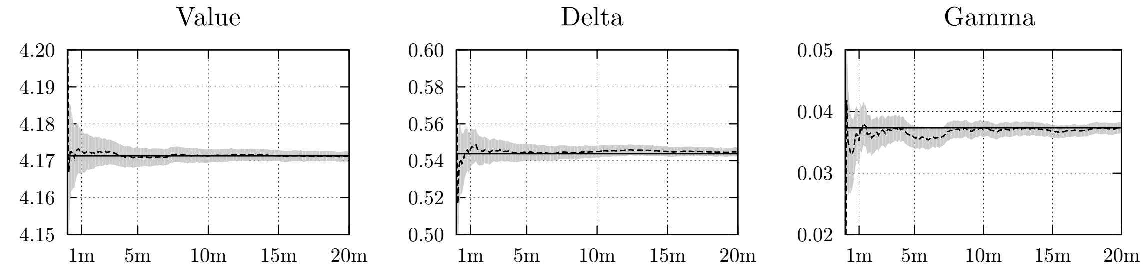

Fig. 4.3: Convergence of numerical results for day 5 (dashed lines, gray areas indicate 95% confidence intervals) of a fully flexible option in comparison to associated analytical results (solid lines).

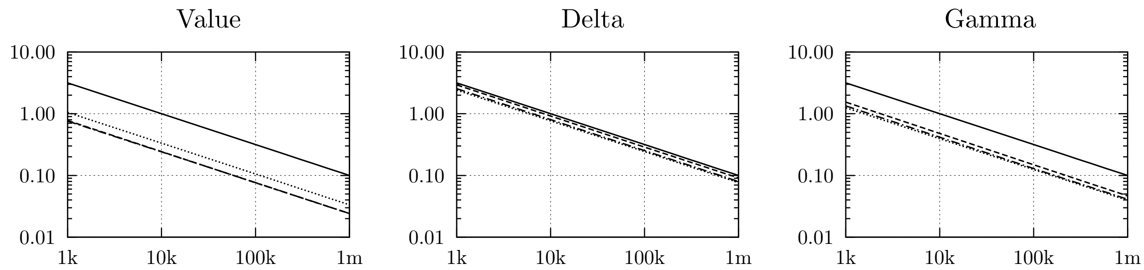

Fig. 4.4: Decay rate of 95% confidence interval of 7 day cumulated option results and comparison to theoretically expected relation $\sim 1/\sqrt{n}$ (solid lines), where n denotes the number of underlying Monte-Carlo realisations (dashed lines: $t_{up} = t_{down} = 1$, $Res_{pos} = Res_{neg} = 0.0$, dashed-dotted lines: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.0$, dotted: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.2$).

4.3 Monte-Carlo convergence

A key characteristic of each Monte-Carlo valuation algorithm is the convergence rate of numerical estimators against associated true results.

Figure 4.3 provides exemplary results and associated 95% confidence intervals for results of the fifth day of a fully flexible power plant option ($t_{up} = t_{down} = 1$) which is not subject to any reserve requirements. Hereby the confidence interval bases on the sample standard deviation (2.6) and is calculated as $1.64 \cdot s_f/\sqrt{n}$, where n denotes the number of underlying Monte-Carlo realisations. This special setup allows to compare the numerical estimators with analytical figures which are also shown in figure 4.3. Obviously the numerical results converge well against the theoretical values.

According to section 2.1 the decay rate of the confidence interval around the value estimator should be proportional to $1/\sqrt{n}$. In contrast to finite difference sensitivities this relation should also be valid for Proxy Simulation Scheme sensitivities. Figure 4.4 compares confidence interval decay rates of 7 day cumulated values, Deltas and Gammas of different power plant options with and without reserve requirements. All confidence intervals follow the theoretically expected decay rate.

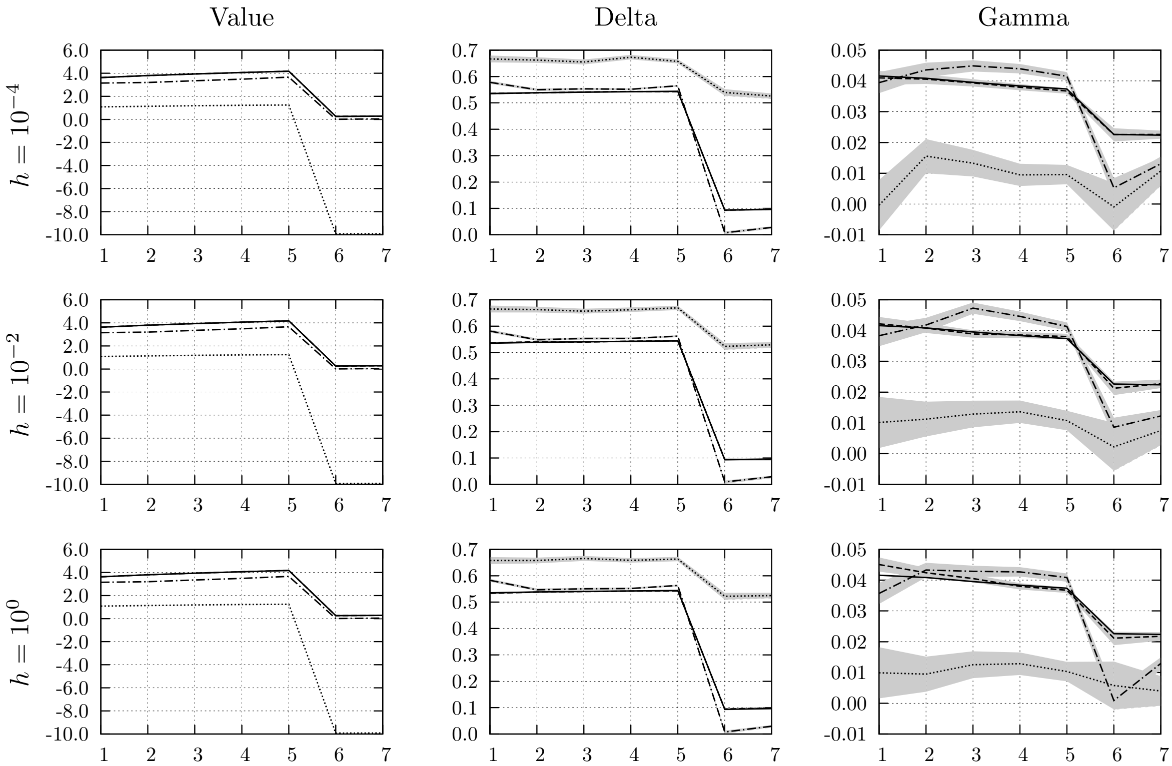

Fig. 4.5: Dependence of numerical results of a 7 day valuation tenor on the shift size h (dashed lines: $t_{up} = t_{down} = 1$, $Res_{pos} = Res_{neg} = 0.0$, $n = 20$m, dashed-dotted lines: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.0$, $n = 7$m, dotted lines: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.2$, $n = 1.1$m) in comparison to analytical results of a fully flexible option (solid lines).

4.4 Impact of shift size

As explained in section 2.2.4 the Proxy Simulation Scheme method includes a finite difference approximation of required derivatives leading to a general discretisation bias. However, as the approximation is applied to probability densities which are sufficiently smooth, both shift size h and the bias can be very small.

In order to demonstrate this property figure 4.5 shows daily numerical results of a fully flexible power plant option and a technically constrained power plant option with and without reserve requirements each derived by using three different shift sizes h in a range between 10^{-4} and 10^{0}. In addition figure 4.5 provides analytical results of the fully flexible option for comparison. The simulations do not base on identical numbers of Monte-Carlo realisations which is the reason for different confidence intervals especially visible in case of Gamma.

Comparing numerical results of the fully flexible option to theoretical values reveals a systematic bias for the largest shift size of 10^{0}. This bias decreases when h is reduced from 10^{0} to 10^{-2} and is no longer relevant for the smallest applied shift size of 10^{-4}. At the same time the confidence interval is rather independent of the shift size.

Therefore we will use a shift size of $h = 10^{-4}$ for all Monte-Carlo simulations throughout this thesis.

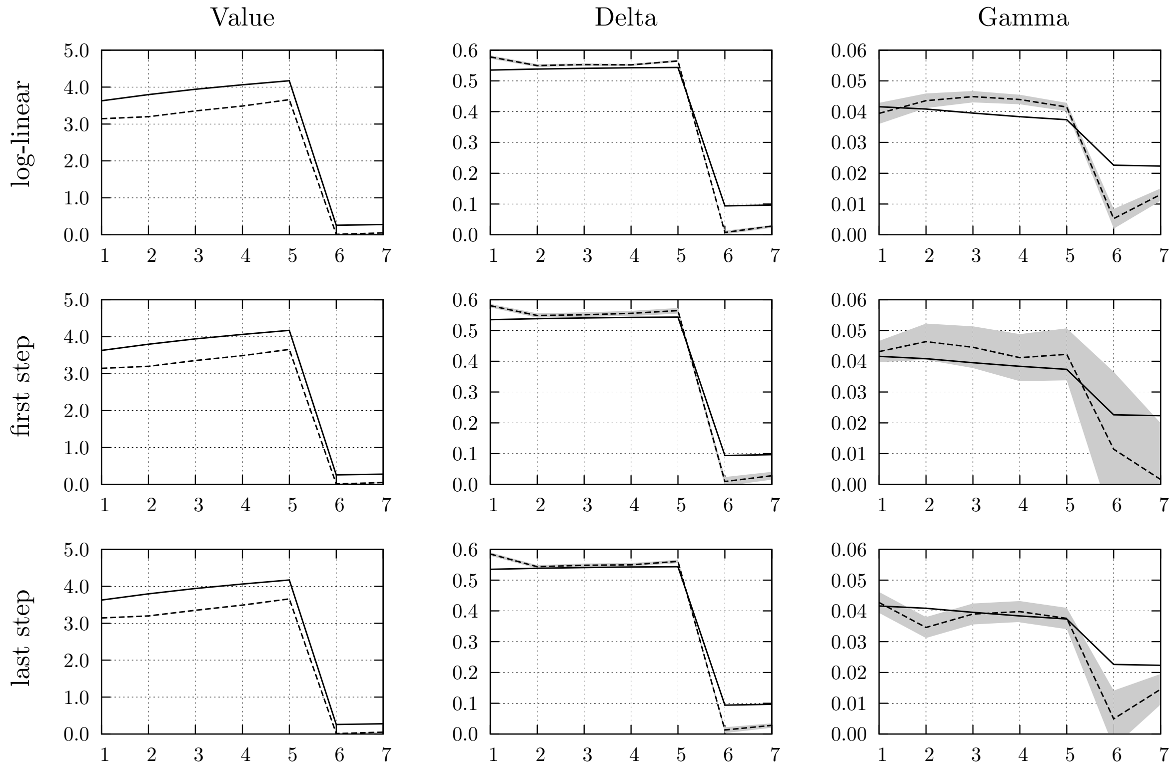

Fig. 4.6: Dependence of numerical results on the method of aligning the virtual simulation paths ($t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.0$, $n = 7$m).

4.5 Alignment of forward price paths

As elaborated in section 3.3.1 the Proxy scheme and the basis simulation scheme are not aligned in a single time step when a sensitivity with respect to the kth forward price (i.e. the kth element of $\mathbf{X}^*(t_0)$) is derived, but is rather equalised linearly in all k time steps between t_0 and t_k. This approach can be called log-linearly as the initial perturbation and its step wise equalisation are expressed in log-coordinates.

We chose this log-linear alignment of forward price paths on purpose because it has a beneficial impact on convergence rates of the Monte-Carlo results. Figure 4.6 shows exemplary results of a technically constrained power plant option without reserve requirements which are derived by aligning price paths log-linearly, entirely within the first time step and entirely within the last time step (the latter with respect to each associated day). In addition figure 4.6 also shows theoretical results of a fully flexible option for comparison.

Obviously the alignment approaches have a significant impact on associated convergence rates. The daily confidence intervals decrease in the log-linear method with increasing time to maturity while exactly the contrary effect is true for an alignment in the first time step. An alignment in the last time step produces similar confidence levels like the log-linear method for the first forward price but this level stays constant instead of being reduced for higher time to maturity as observed for the log-linear alignment. These effects are even more obvious in figure 4.7, which shows widths of

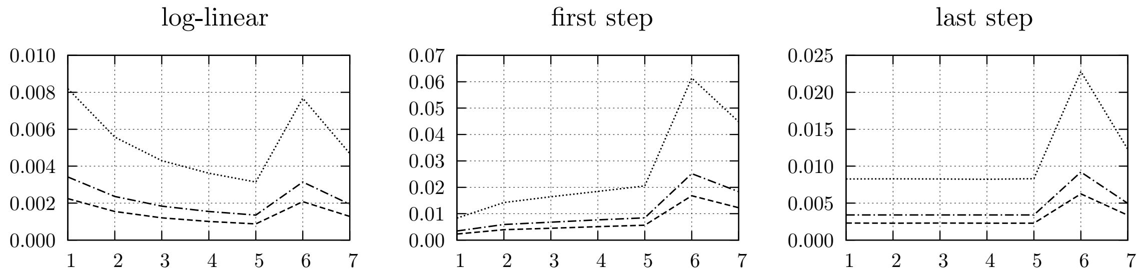

Fig. 4.7: Dependence of the width of Γ confidence intervals on the method of aligning the virtual simulation paths (dashed: $t_{up} = t_{down} = 1$, $Res_{pos} = Res_{neg} = 0.0$, $n = 20$m, dashed-dotted: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.0$, $n = 7$m, dotted: $t_{up} = 12$, $t_{down} = 8$, $Res_{pos} = Res_{neg} = 0.2$, $n = 1.1$m).

confidence intervals of daily Gamma for multiple option setups each using all three different alignment approaches. (This figure also reveals a general confidence level drop on weekend days 6 and 7 which is further discussed in section 4.6.)

The observed confidence level dependencies can be explained by the size and the underlying structure of the weight functions (3.25) being used in the Proxy Simulation Scheme expressions of the sensitivities (3.34). These weights explicitly depend on modified vectors of standard normal distributed random variables $\mathbf{\Delta U}^{up}(t_0)$ and $\mathbf{\Delta U}^{down}(t_0)$ (equations (3.30) and (3.32)) where the absolute size of the modification depends on the associated volatility term structure (3.8) which is effectively a part of the denominator. Therefore the absolute size of the weight function will generally increase with lower applied volatilities. As the weight functions are directly multiplied with option payoffs to produce Monte-Carlo samples whose average value serves as final Monte-Carlo estimator, their absolute size has a direct influence on the convergence rates of the estimators.

This effect is the reason for the observed confidence dependencies. When price paths are fully aligned in the first time step the applied volatility for a sensitivity with respect to the kth forward price is identical to $\sigma_k^*(t_0)$ (equation 3.8) which decreases with increasing k. This consequently leads to a decreased Monte-Carlo confidence level for later days. This argument is further supported by figure 4.8 providing numerical results including confidence intervals for a path alignment within the first time step subject to two different, artificial volatility term structures. The first one is a constant term structure on a comparably high level of $\sigma_1^*(t_0)$ whereas the second emulates a very sharp drop of volatility with increasing time to maturity. As expected the confidence interval stays rather constant in the first case while it strongly increases in the second case.

When price paths are fully aligned in the last time step before the kth forward price is fixated the applied volatility is $\sigma_k^*(t_{k-1})$ which is always identical to the largest element $\sigma_1^*(t_0)$ of the volatility term structure. This leads to similar confidence levels for the entire valuation tenor. Please note that this level is the same comparably low level as observed for the first forward price when price paths are aligned in the first time step.

In case of a log-linear alignment of price paths an additional effect influences the

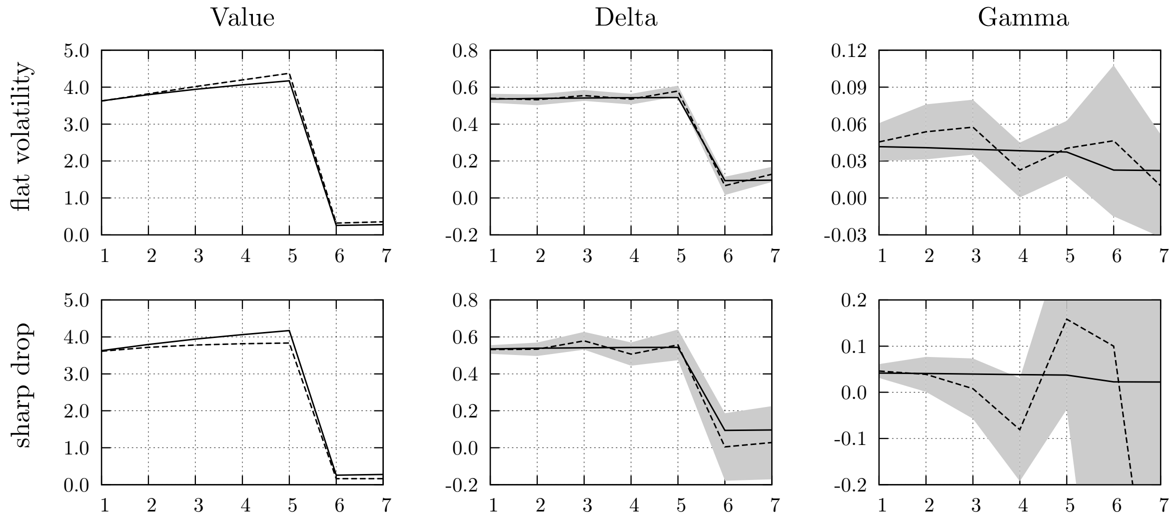

Fig. 4.8: Dependence of numerical results (dashed lines) and associated 95% confidence intervals (gray areas) on the volatility term structure when virtual paths are aligned in the first time step ($t_{up} = t_{down} = 1$, $Res_{pos} = Res_{neg} = 0.0$, $n = 0.3$m). Analytical results are provided for comparison (solid lines). Top: $\sigma^* = 1.18$, Bottom: $\sigma_1 = 1.5$, $\sigma_2 = 0.0$, $\kappa = 100$.

confidence rates at later days. The absolute confidence level of sensitivities with respect to the first forward price is similar to results of the other two alignment methods. This is again due to the fact that all approaches must align the first forward price within the same first time step using the same volatility $\sigma_1^*(t_0)$. However where the weight functions of the other two methods depend only on random numbers as used in one associated time step, the weight functions of the log-linear method include weighted random numbers of multiple, independent time steps – the more the higher the time to maturity. These random numbers are entirely uncorrelated for two consecutive time steps and therefore lead to a variance reduction of the weight functions being directly visible in figure 4.7.

4.6 Standard error on weekend and weekdays

As already mentioned in section 4.5 it can be observed that numerical confidence levels are generally lower on weekends than on weekdays (figures 4.6, 4.7 and 4.8).

Similar to the impact of volatility as discussed in section 4.5 this is a direct effect of the structure of the weight functions (3.25). The volatility is given in the denominator of the argument of the exponential functions that are part of the weights, leading to a confidence increasing effect for higher applied volatilities. In contrast the shift size h is part of the logarithmic shift sizes $\delta_k^{up/down}$ (3.24) and therewith part of the numerator of the argument of the exponential functions. Since the logarithmic shift sizes also depend on the absolute price level $\delta_k^{up/down}$ increase with decreasing prices. Due to the positioning of $\delta_k^{up/down}$ in the numerator this leads to a lower confidence level for lower weekend prices.

This effect can be eliminated by using a variable shift size h which compensates all

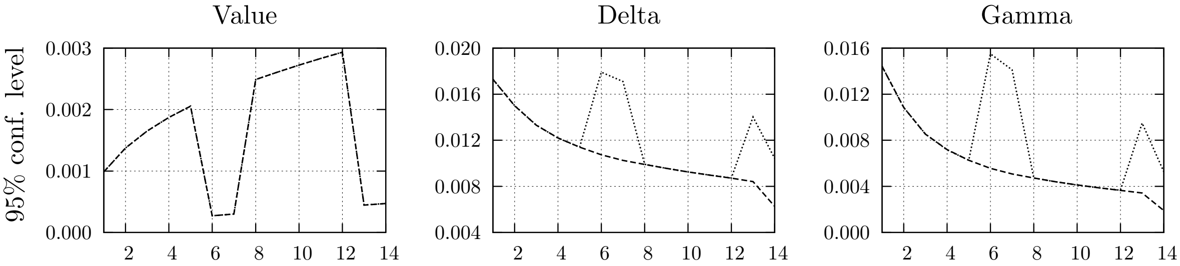

Fig. 4.9: 95% confidence levels of daily Value, Delta and Gamma results for fully flexible option based on constant shift sizes of $h = 10^{-4}$ for all days (dashed lines) and on variable shift sizes, i.e. $h_{wd} = 10^{-4}$ and $h_{we} = 3/5 \cdot 10^{-4}$ (dotted lines).

impacts on the logarithmic shift sizes due to differences of initial forward prices. This is demonstrated by figure 4.9 which shows confidence levels of numerical results of a fully flexible power plant option being derived by 1. using a constant shift size $h = 10^{-4}$ for all days and 2. using the same shift size for weekdays but an appropriately down scaled shift size for weekends. However this shift modifying approach can generally raise numerical stability issues in case initial forward prices vary significantly. Therefore we will stick to the constant absolute shift size $h = 10^{-4}$ throughout this thesis.

5 Single option sensitivities

Before numerical results of the Proxy Simulation Scheme method are analysed for portfolios of multiple power plant options in chapter 6 this chapter presents the impact of different technical constraints and external reserve requirements only on stand-alone power plant options. Even these simplified setups can lead to unintuitive results whose underlying drivers need to be identified before being able to separate them from more complex portfolio interdependencies in later analyses.

After a detailed discussion of these findings, followed by a short introduction to differences between Delta and expected dispatch, we provide a comprehensive verification of all numerical results including sensitivities via both a Delta-Gamma hedging backtesting framework and the comparison of single step probability densities of option value changes.

5.1 Impact of technical constraints

As explained in sections 3.2.1 and 3.2.2 technical constraints are typically not identical for different power plant options and need to be fulfilled separately for each option. In the absence of global portfolio constraints (like reserve requirements) these local constraints will lead to time coupling effects on the dispatch of each individual power plant but not to any kind of interdependencies between dispatch decisions of different power plant options.

In this thesis we will especially analyse the impact of minimum up- and down-times $t^p_{up/down}$ and option strikes κ^p with and without additionally applied reserve requirements $Res_{pos/neg}$. Section 5.1 focuses on technical constraints of stand-alone power plant options while combinations with reserve requirements are provided in section 5.2.

5.1.1 Minimum up time

Figure 5.1 shows daily numerical results of four power plant options with different minimum up-times in a range between 4 and 48 hours, minimum down-time of 1 hour and identical strikes $\kappa = 50$. For comparison figure 5.1 also provides daily analytical results of a fully flexible hourly option based on standard Black-Scholes option evaluation (Black and Scholes (1973)).

For a minimum up-time of $t_{up} = 4$ there are no significant deviations between analytical and numerical results for weekdays (i.e. day 1 to day 5) whereas for weekends (day 6 and 7) both value and Delta are slightly reduced. This can be explained via the hourly price shape (section 3.1.4) in combination with the different average price

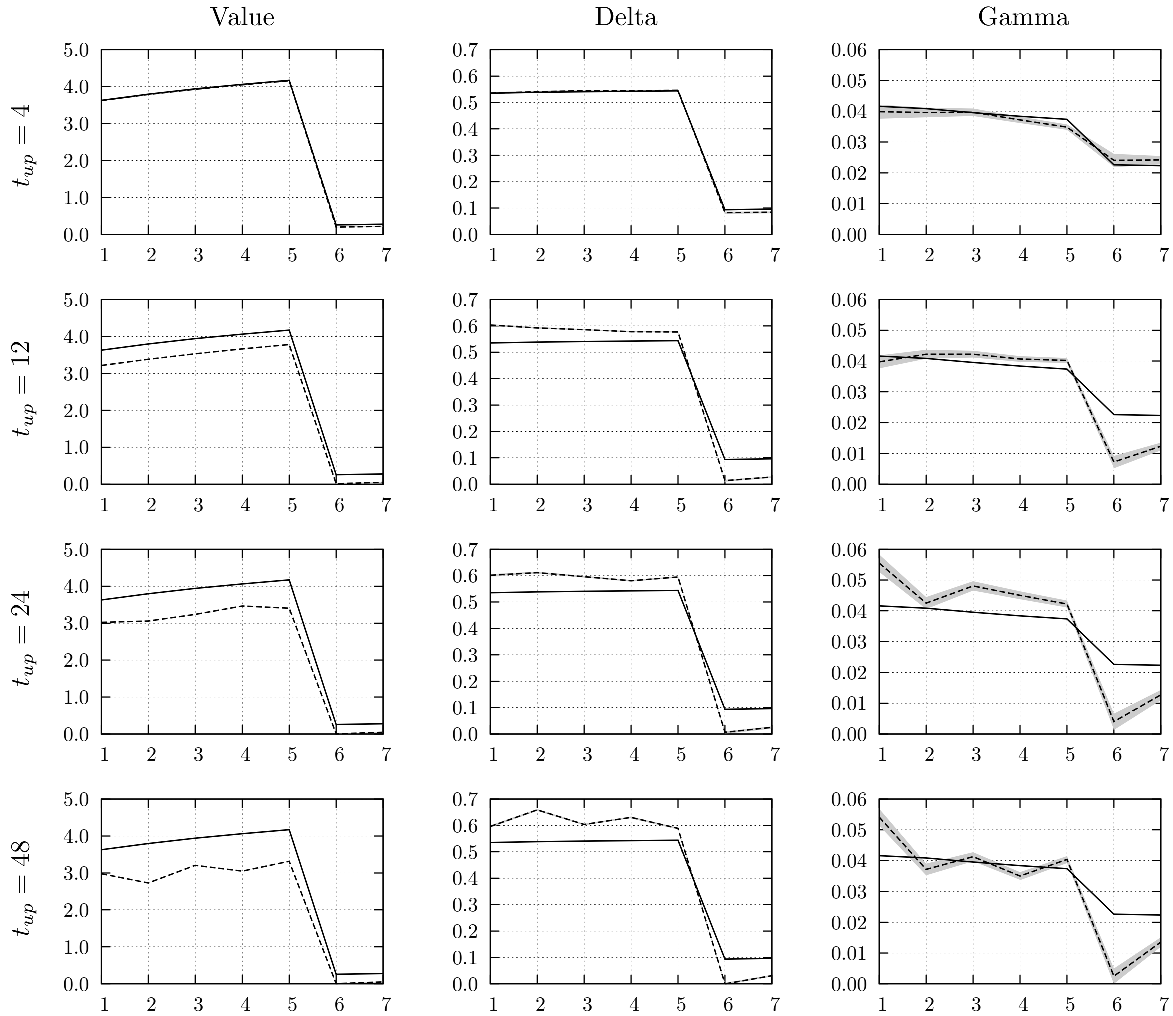

Fig. 5.1: Dependence of numerical Values, Deltas and Gammas on the minimum up-time t_{up} of the power plant option for a tenor of 7 days (dashed lines, $t_{down} = 1$, $n = 10\text{m}$, $\kappa = 50$). Analytical benchmark results of a fully flexible option are provided in comparison (solid lines).

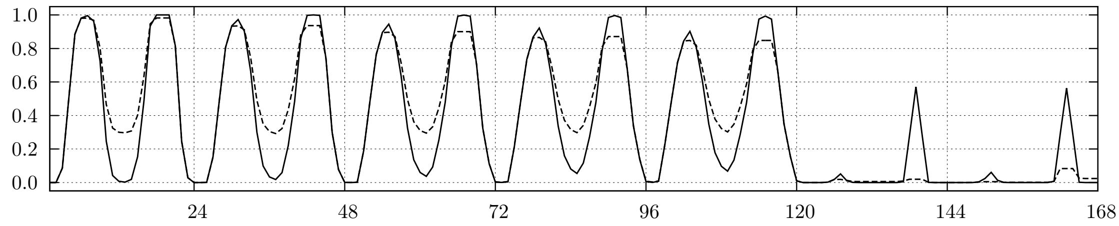

Fig. 5.2: Average hourly dispatch for all 7 days of a technically constrained option with $t_{up} = 12$ and $t_{down} = 1$ (dashed line) in comparison to the dispatch of a fully flexible option (solid line).

level on weekdays and weekends. On weekdays in the money blocks, i.e hours where $\alpha^l \cdot X_j(t_j) > \kappa^p$, are typically longer than 4 hours while on weekends this length generally decreases as the options are overall deep out of the money. Therefore the technical up-time constraint $t_{up} = 4$ has an impact on option dispatches on weekends while dispatches on weekdays are rather unaffected.

This changes when the minimum up-time is increased to $t_{up} = 12$. Now the value is reduced for all days and therewith close to zero on weekend days while Delta and Gamma are systematically increased on weekdays but decreased on weekend days. This can directly be explained by a comparison of the structure of the average dispatch in the constrained case and the average dispatch of a fully flexible hourly option (figure 5.2). The increased up-time constraint leads to more frequent option exercises during midday hours of weekdays. This results in lower daily values as these hours are typically out of the money, but also in higher Deltas as the power plant option runs additional hours which is equivalent to an increased sensitivity with respect to electricity price changes. Due to the applied hourly price shape the longer online block is closer at the money than the smaller hourly blocks in case of a fully flexible option, which leads to the increased daily Gamma values for $t_{up} = 12$. To support this explanation figure 5.3 shows Black-Scholes Gammas for an option that mirrors the typical power plant option dispatch in case of $t_{up} = 12$, i.e. the option is online from hour 3 to hour 23 while running at maximum load of 1.0 when $\alpha^l \geq 1$ and at minimum load of 0.3 when $\alpha^l < 1$. Obviously this can only serve as a rough approximation of the real effect but even this approach is able to replicate both structure and direction of the Gamma deviation. Interestingly Gamma is on average higher on weekdays for $t_{up} = 12$ than for $t_{up} = 4$ in this setup, even though Gamma is commonly interpreted as a measure for the flexibility of a financial product and following this logic could be expected to decrease for stricter technical constraints.

According to figure 5.2 a minimum up-time of 12 reduces the average dispatch on weekend days significantly. This is due to the fact that only a small amount of consecutive hours is in the money on these days which makes a 12 hour dispatch usually uneconomical. This effect is comparable to an additional out of the money shift of the option which is consistent with the observed reduction of value, Delta and Gamma on weekends.

A further increase of minimum up-time to $t_{up} = 24$ intensifies the impact on both value and Delta of the option while the daily structure of Gamma shows additional

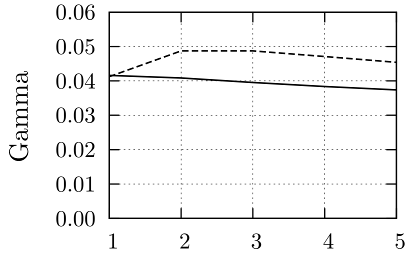

Fig. 5.3: Analytical Black-Scholes Gamma results of a constrained option with a single in the money block of 21 hours (dashed line) in comparison to the analytical Gamma of a fully flexible option (solid line).

features. Gamma is significantly larger on the first day and comparably smaller on the second day. This reflects the high flexibility of the power plant option with respect to the point in time during the first day when it is started (if at all), which subsequently limits the flexibility of the dispatch on the second day as the option needs to stay online at least for 24 hours after a start-up.

This effect becomes even more dominant when the minimum up-time is further increased to $t_{up} = 48$. Now dispatch decisions on the first day influence and dominate all dispatch decisions of later days which leads to a comparably high Gamma on the first day and smaller Gammas on consecutive days. In addition all numerical results now show a spike structure with period length of two days. This is consistent with the minimum up-time constraint of 48 hours which generally limits the dispatch flexibility on each first day after the day when a dispatch decision has been taken. Similar to the option with $t_{up} = 24$ also the more inflexible option with $t_{up} = 48$ is per construction of the mixed integer optimisation problem flexible in the decision when to start-up on the first day (if at all) and will typically be started in the evening price peak. This leads to a relatively high value, a relatively low Delta and a high Gamma on the first day. However this dispatch decision implies a must-run condition during the second day which leads to comparably low value, high Delta and low Gamma. The must-run condition expires during day three which lifts value and Gamma while leading to a lower Delta - similar to the first day.

Both options with $t_{up} = 24$ and $t_{up} = 48$ are almost never dispatched on day 6 which leads to numerical results close to zero. These values increase again for day 7. This effect is a numerical artifact which is due to the open right boundary condition of the mixed integer optimisation problem. The applied solver does only take into account technical constraints within the valuation tenor including all hours of day seven in figure 5.1. In case minimum up- or down-times exceed the number of remaining hours of the valuation tenor, these technical constraints will effectively be lowered to match the number of remaining hours. This leads to an artificially higher option flexibility during the last day in this numerical setup.

Figure 5.4 aggregates all results as provided in figure 5.1 in order to identify general trends for weekday and weekend results. We will use this kind of visualisation regularly in chapter 6 to show dependencies of results of power plant option portfolios.

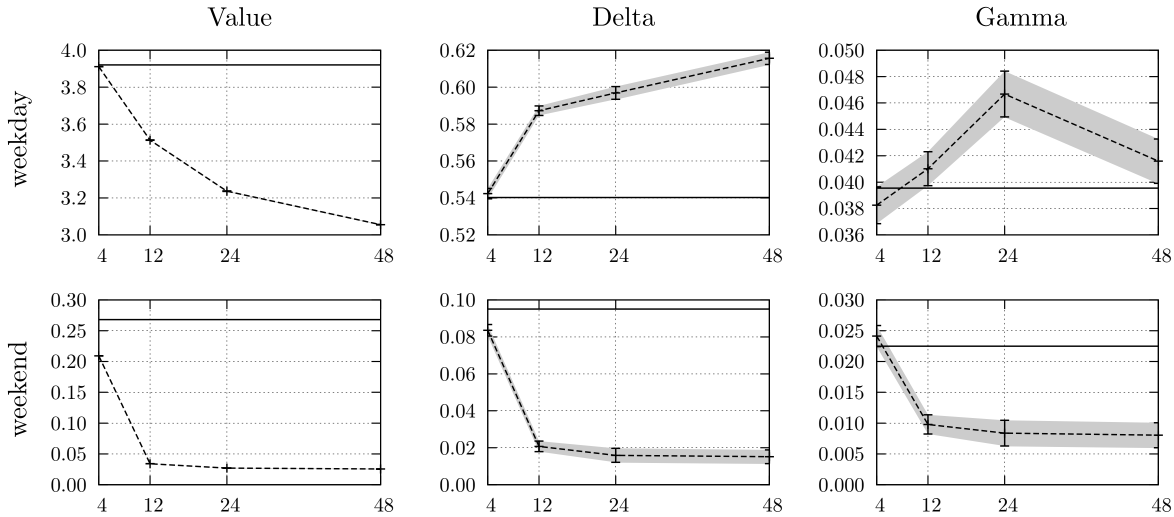

Fig. 5.4: Dependence of average daily Values, Deltas and Gammas (dashed lines) for weekdays and weekend days on the minimum up-time t_{up} in comparison to analytical results of a fully flexible option (solid lines). 95% confidence intervals denoted by error bars and gray areas. $n = 10$m.

5.1.2 Minimum down time

The impact of increasing minimum down-time is of similar size as the impact of minimum up-time (section 5.1.1), although structurally not identical in all cases. Figure 5.5 shows daily numerical results for four power plant options with different minimum down-times between 4 and 48 hours, minimum up-time of 1 hour and identical strikes $\kappa = 50$. For comparison also daily analytical results for a fully flexible hourly option are provided.

Similarly to a low minimum up-time (figure 5.1) a minimum down-time of $t_{down} = 4$ has no relevant impact on the numerical results. This changes for a minimum down-time of $t_{down} = 8$ and the impact on both value and Gamma is comparable to the impact of a minimum up-time of $t_{up} = 12$. However where this minimum up-time leads to an increase of Delta on weekdays, $t_{down} = 8$ leads to a significant Delta reduction. This can be explained by the structure of the average hourly dispatch of the constrained power plant option as provided in figure 5.6. In contrast to the generally increased average dispatch for $t_{up} = 12$ (figure 5.2) $t_{down} = 8$ results in longer shut-downs periods during midday and night hours. This subsequently leads to a decreased average daily dispatch and the observed decreased daily Delta on weekdays. The minor down-time constraint does not impact weekend results significantly because the period between in the money hours is typically longer than the applied minimum down-time, leading to almost unchanged dispatches and therefore only minor impacts on numerical results on weekends.

The results for $t_{down} = 24$ and $t_{down} = 48$ show similar impacts as already explained for higher minimum up-times in section 5.1.1. The numerical values deviate significantly from analytical results of a fully flexible option and are overlaid by a spike structure. This spike structure has a slightly different shape than the structure for

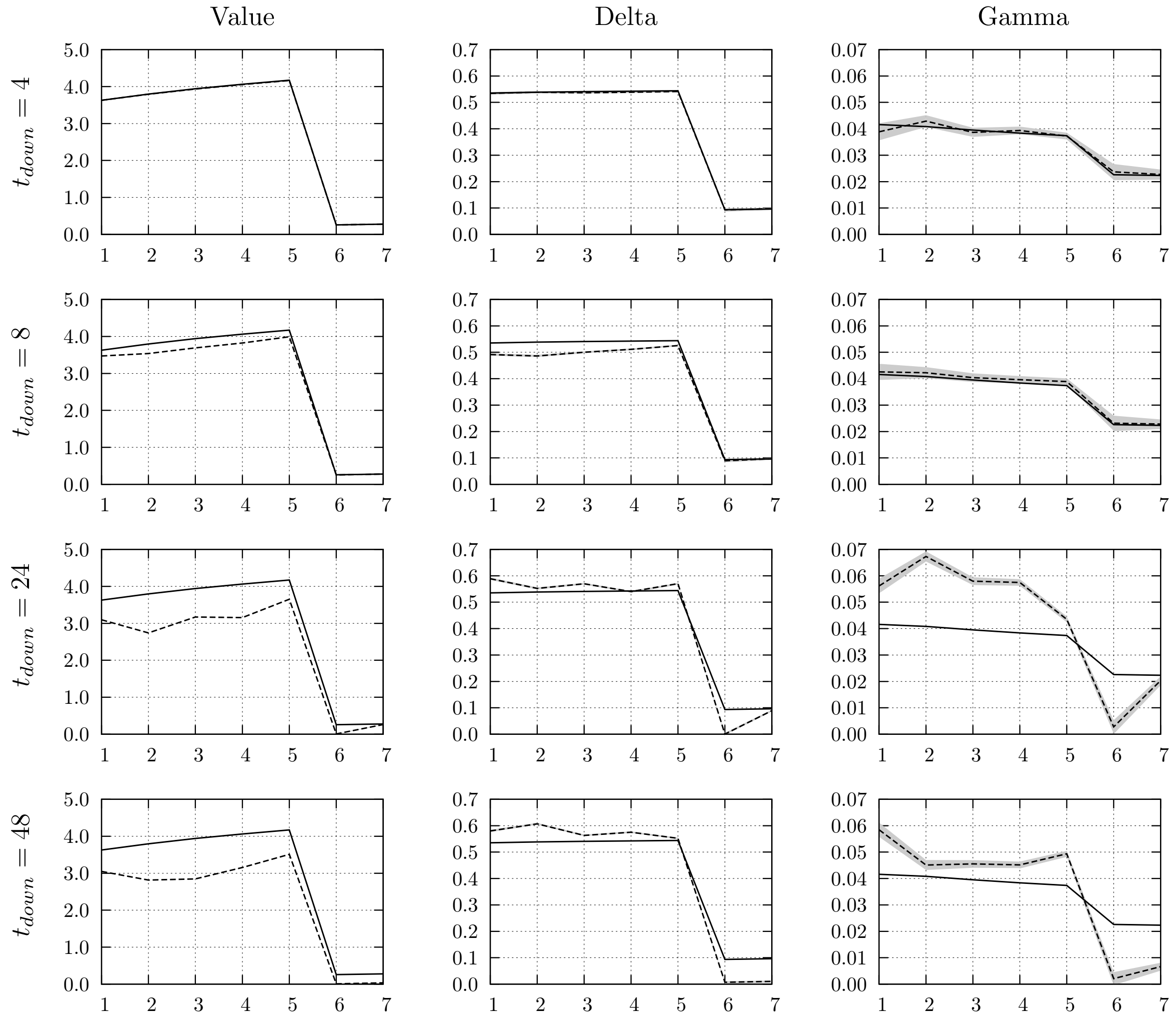

Fig. 5.5: Dependence of numerical Values, Deltas and Gammas on the minimum down-time t_{down} of the power plant option for a tenor of 7 days (dashed lines, $t_{up} = 1$, $n = 10\text{m}$, $\kappa = 50$). Analytical benchmark results of a fully flexible option are provided in comparison (solid lines).

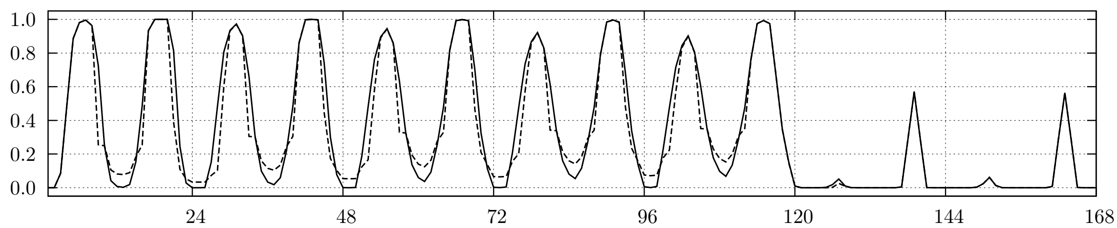

Fig. 5.6: Average hourly dispatch for all 7 days of a technically constrained option with $t_{up} = 1$ and $t_{down} = 8$ (dashed line) in comparison to the dispatch of a fully flexible option (solid line).

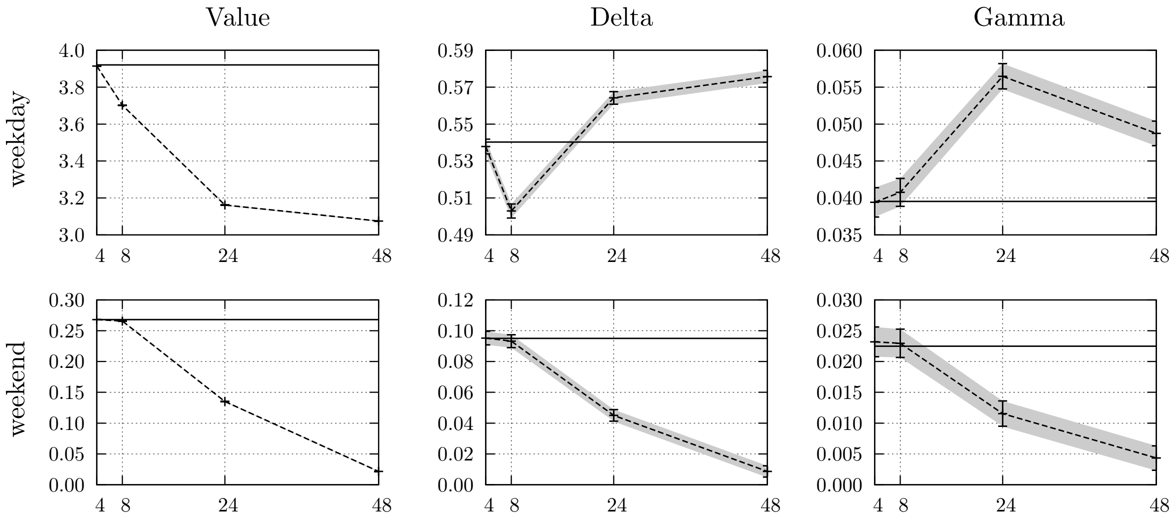

Fig. 5.7: Dependence of average daily Values, Deltas and Gammas (dashed lines) for weekdays and weekend days on the minimum down-time t_{down} in comparison to analytical results of fully flexible option (solid lines). 95% confidence intervals denoted by error bars and gray areas. $n = 10$m.

higher minimum up-times because relevant dispatch decisions (e.g. when to shut down on the first day of the valuation tenor) are usually taken at different points in time than in case a minimum up-time constraint is valid. Again weekday Gammas are positively impacted by increased technical constraints even though it could intuitively be expected that the overall flexibility of the power plant option should rather be decreased.

Similarly to figure 5.4 figure 5.7 provides an aggregated view on the individual option results which clearly shows the discussed trends.

5.2 Impact of reserve requirement

Reserve requirements need to be served on a portfolio level and therefore impose global constraints on the involved set of power plant options. In contrast to local constraints like option specific technical constraints reserve requirements typically lead to interdependencies between dispatches of different power plant options within the portfolio. In a first step we discuss the impact of reserve requirements on stand-alone options in this section before entire portfolios are analysed in chapter 6.

5.2.1 Positive reserve requirement

Figure 5.8 shows numerical results of a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$) without reserve requirements and subject to 0.3 and 0.5 positive reserve requirement in comparison to analytical results of a fully flexible option. According to the general model setup (section 3.2.1) the option's minimum and maximum load are 0.3 and 1.0. Subject to reserve requirements the option is forced

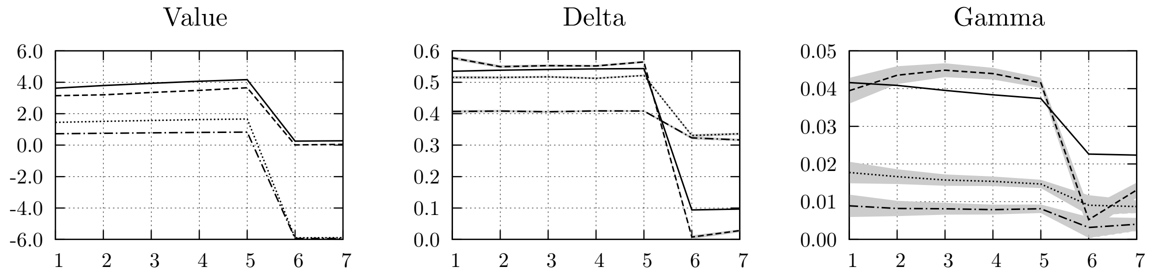

Fig. 5.8: Numerical results of a technically constrained power plant option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$ without reserve requirements (dashed lines) and subject to positive reserve requirements of 0.3 (dotted lines) and 0.5 (dashed-dotted lines) in comparison to analytical results of a fully flexible option (solid lines). $n = 3$m.

to stay online in all hours (strict must-run condition) while being able to ramp up the amount of positive reserve requirement in each hour within the valuation tenor. This limits the load flexibility of the power plant option to the interval [0.3, 0.7] in case of 0.3 reserve requirement and [0.3, 0.5] in case of 0.5 reserve requirement.

The effect of this limitation is clearly visible in figure 5.8. For $Res_{pos} = 0.3$ the option value decreases significantly on weekdays and even more on weekends which is due to a combination of an increased amount of hours with negative value contribution triggered by the overall must-run condition and the reduced effective maximum load of 0.7 limiting the positive value contribution of in the money hours. The first effect leads to a general increase of Delta as more electricity is produced in out of the money hours while the second effect leads to lower electricity production during in the money hours and therefore a negative impact on Delta. In the chosen model setup the latter effect predominates on weekdays and yields a slightly reduced Delta on weekdays while the must-run condition leads to a significantly increased Delta on weekends. Please note that the relative impact size of both effects may vary for different hourly price shapes or option strikes. Gamma is generally reduced due to the overall reduction of flexibility of the power plant option which is here consistent with the interpretation of Gamma as a measure for flexibility of a financial product or a portfolio.

When the amount of positive reserve requirement is increased to $Res_{pos} = 0.5$ an additional value decrease can be observed on weekdays. However this additional effect is lower than the impact of imposing a reserve requirement at all when comparing results without reserve requirements and subject to a minor reserve requirement of $Res_{pos} = 0.3$. This is due to the fact that the must-run condition is similarly valid for both amounts of reserve requirements and an additional increase from 0.3 to 0.5 only impacts the effective maximum load during in the money hours. This is also the explanation for almost no additional value drop on weekends as the absolute amount of in the money hours is significantly lower on weekends than on weekdays. Consistently with this argumentation Delta is significantly reduced on weekdays and only slightly downwards shifted on weekends. Gamma is again generally reduced with respect to results for $Res_{pos} = 0.3$ which is a results of the additionally decreased load flexibility on all days.

Figure 5.9 shows the average impact of an increased positive reserve requirement

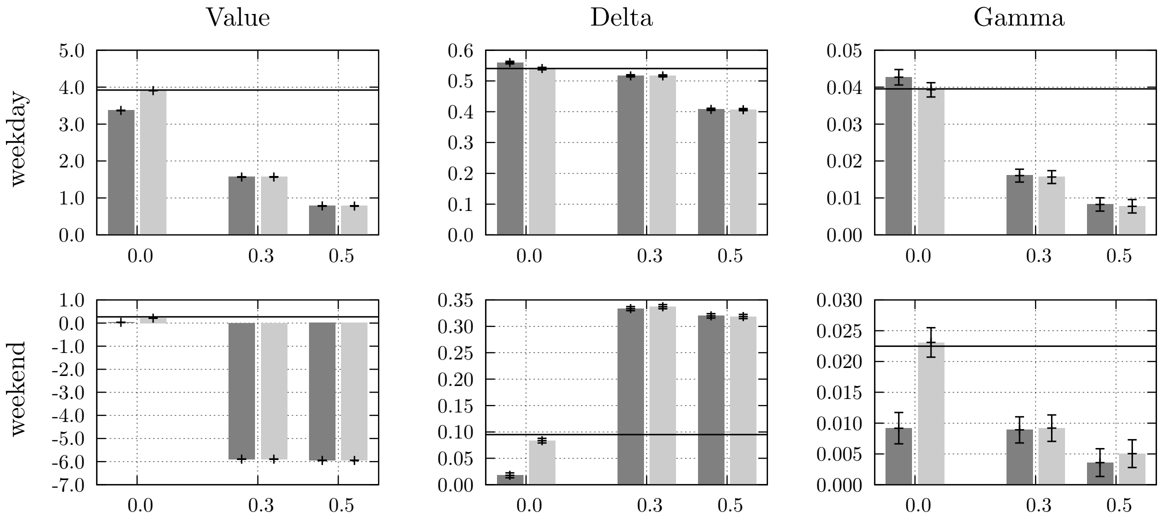

Fig. 5.9: Dependence of average daily numerical results of a technically constrained power plant option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$ (dark gray columns) and $t_{up} = t_{down} = 4$ and $\kappa = 50$ (light gray columns) on the amount of positive reserve requirement in comparison to constant analytical results of a fully flexible option (solid lines). 95% confidence intervals denoted by error bars. $n = 3$m.

on two different stand-alone power plant options separated by weekdays and weekend days. The first option has already been discussed above (i.e. $t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$), the second options is more flexible in terms of minim up- and down time ($t_{up} = t_{down} = 4$, $\kappa = 50$). Similarly to figures 5.4 and 5.7 this visualisation provides a good overview about the discussed trends for the technically more constrained option. The comparison to the more flexible option shows in addition that different up- and down-times do not have any impact on numerical results in case of existing reserve requirements (within given confidence levels). This is consistent with the chosen model setup that does not take ramp rates into account. In this setup the option is able to ramp an arbitrary percentage of the maximum load from one hour to the next hour (which is not unrealistic for flexible physical generation assets). In case the option needs to stay online in all hours due to the reserve requirement, up- and down-time constraints are no longer relevant as the option can ramp instantly between minimum load and (effective) maximum load, independently of the exact value of t_{up} and t_{down}. This leads to identical dispatches and therefore identical numerical results for both options. However this coincidence is only true for stand-alone options. In case of multiple option portfolios the impact of increasing reserve requirements clearly depends on technical parameters of the involved power plant options as discussed in chapter 6.

5.2.2 Negative reserve requirement

Similarly to figure 5.8 figure 5.10 shows numerical results of a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$) without reserve requirements and subject to 0.3 and 0.5 negative reserve requirement in comparison to analytical results of a fully flexible option.

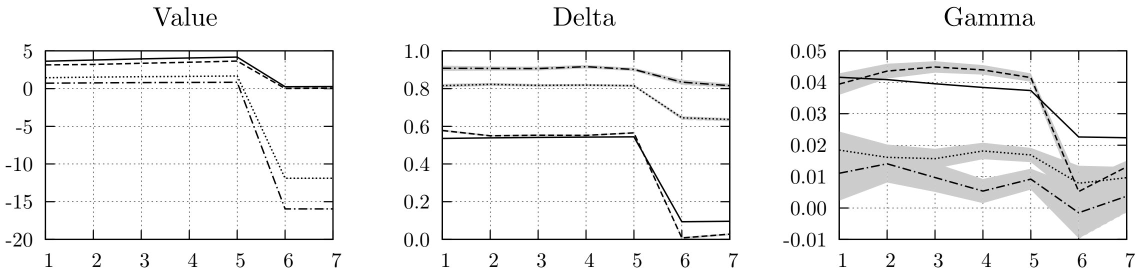

Fig. 5.10: Numerical results of a technically constrained power plant option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$ without reserve requirements (dashed lines) and subject to negative reserve requirement of 0.3 (dotted line) and 0.5 (dashed-dotted lines) in comparison to analytical results of a fully flexible option (solid lines). $n = 3$m.

Where the positive reserve requirements in section 5.2.1 limit the option's load to the intervals [0.3, 0.7] for $Res_{pos} = 0.3$ and [0.3, 0.5] for $Res_{pos} = 0.5$, the applied negative reserve requirements of 0.3 and 0.5 limit the load to the intervals [0.6, 1.0] and [0.8, 1.0] respectively. In contrast to a positive reserve requirement that reduces the effective maximal load, a negative reserve requirement increases the effective minimum load. This directly leads to a significantly stronger impact of the must-run condition resulting in significant value drops especially on the weekend and a general Delta increase for all days of the valuation tenor. The absolute size of the effective load intervals are exactly the same as in the setup with positive reserve requirement. Therefore also the flexibility of the option dispatches are identical being the reason for the same level of Gamma in figures 5.8 and 5.10.

Similar to figure 5.9 figure 5.11 shows the average impact of an increased negative reserve requirement on two different power plant options separated by weekdays and weekend days. As discussed for positive reserve the numerical results do not depend on t_{up} and t_{down} in case of existing reserve requirements - no matter of positive or negative reserve requirements - as long as the analysis is performed for a stand-alone power plant option.

As the impact of negative reserve requirements is very similar to the impact of positive reserve requirements apart form the absolute size we will usually focus on positive reserve requirements in the following analyses of this thesis.

5.2.3 Combination of positive and negative reserve requirements

Figure 5.12 provides numerical results of a technically constrained power plant option ($t_{up} = t_{down} = 4$, $\kappa = 50$) subject to both 0.3 positive and 0.3 negative reserve requirement in comparison to analytical results of a fully flexible option.

These requirements limit to options's load flexibility to the interval [0.6, 0.7] which leads to Gamma results close to zero for all days of the valuation tenor. Delta is generally increased and shows an almost flat daily structure within the allowed load interval. This is consistent with the limited flexibility of the option that makes its dispatch very similar to the dispatch of a linear product like a forward or future contract (with contracted base load capacity of ≈ 0.65 on weekdays and ≈ 0.6 on weekend days).

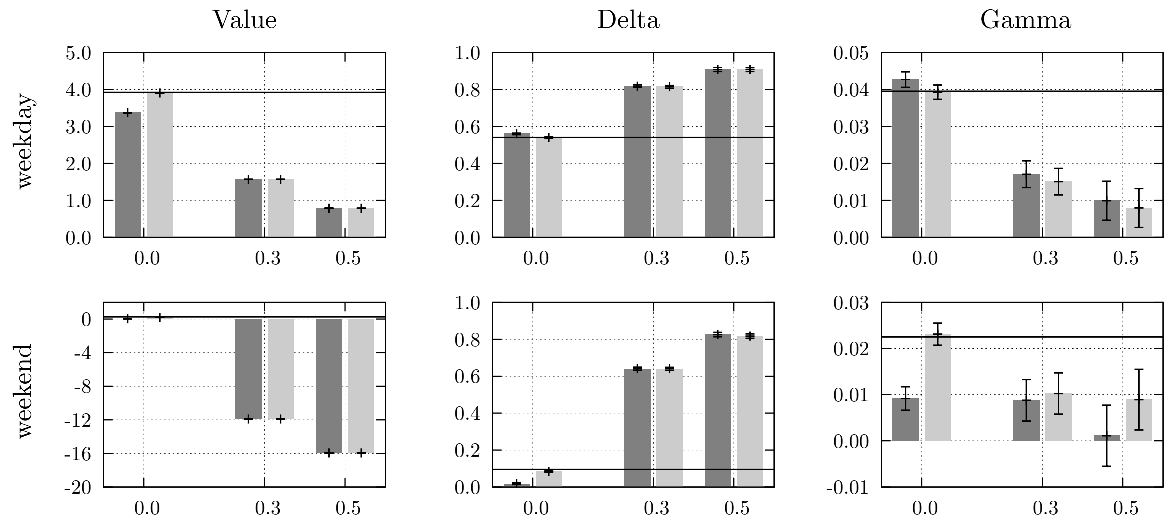

Fig. 5.11: Dependence of average daily numerical results of a technically constrained power plant option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$ (dark gray columns) and $t_{up} = t_{down} = 4$ and $\kappa = 50$ (light gray columns) on the amount of negative reserve requirement in comparison to constant analytical results of a fully flexible option (solid lines). 95% confidence intervals denoted by error bars. $n = 3$m.

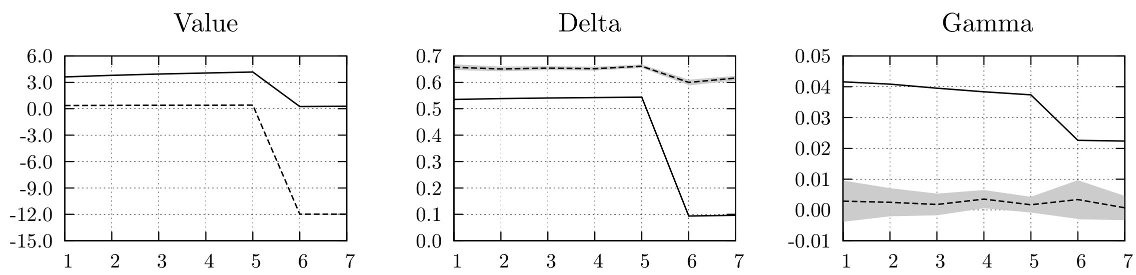

Fig. 5.12: Numerical results of a technically constrained power plant option with $t_{up} = t_{down} = 4$ which is subject to positive and negative reserve requirement of $Res_{pos} = Res_{neg} = 0.3$ (dashed line) in comparison to analytical results of a fully flexible option (solid lines). $n = 3$m.

With respect to numerical results for either positive or negative reserve requirement (figures 5.8 and 5.10) the value is further decreased on weekdays which is a result of the stricter must-run condition. The value on weekend days is similar to previous results of the option subject to only negative reserve requirement as the option is so deep out of the money on weekend days that the reduction of upside flexibility due to the additional positive reserve requirement is negligible.

5.3 Delta vs. expected Dispatch

In comparison to the real sensitivity Delta the expected dispatch of a power plant option can be derived rather easily and numerically cost efficiently and has usually a high acceptance of both practitioners and decision-makers in the industry sector. However there is a systematical bias between the two result types because the expected dispatch is not able to cover the entire non-linearity of a typical option product. This is the reason why especially in the area of risk controlling and risk management a naive utilisation of expected dispatch instead of Delta is usually not accepted.

In this section we will discuss the bias for some exemplary cases before we will demonstrate the explicit requirement of the sensitivity Delta for a risk assessment of constrained power plant options via backtesting results in section 5.4.

5.3.1 Only technical constraints

Figure 5.13 shows numerical Deltas of different technically constrained power plant options with identical strikes $\kappa = 50$ without reserve requirements in comparison to both associated expected dispatches and analytical Deltas of a fully flexible option.

The general downwards bias between Delta and expected dispatch is clearly visible for all analysed options even though the absolute level of deviation decreases when both results converge against zero as it is the case for $t_{up} = 12$ and $t_{down} = 8$ on weekend days. The overall bias might seem to be small but it has to be taken into account that hedging on the basis of expected dispatch instead of Delta would lead to a hedging error of up to 20% for the analysed options which effectively means that up to 20% of the electricity production would systematically remain unhedged and therefore subject to market price risk. This is clearly unacceptable for both portfolio management and risk controlling purposes.

Figure 5.14 provides an aggregated visualisation of the discussed results that demonstrates the systematical bias even more clearly.

5.3.2 Technical constraints and reserve requirements

Similar to figures 5.13 and 5.14 figures 5.15 and 5.16 show numerical Deltas of different technically constrained power plant options with identical strike $\kappa = 50$ subject to positive reserve requirement ($Res_{pos} = 0.3$) in comparison to both associated expected dispatches and analytical Deltas of a fully flexible option.

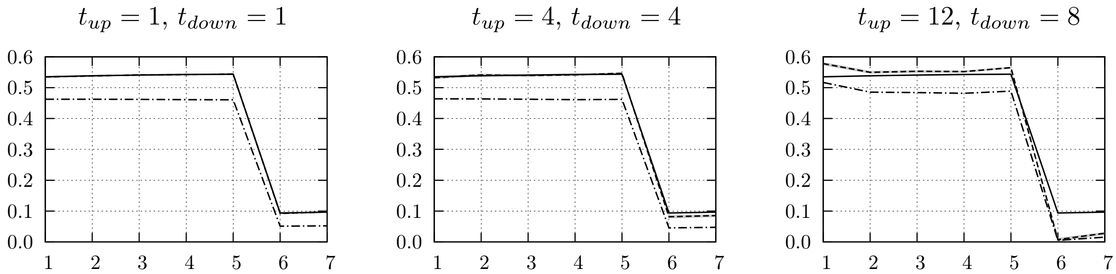

Fig. 5.13: Comparison of daily numerical Deltas (dashed lines) with daily expected dispatches (dashed-dotted lines) of options with different technical constraints, and analytical results of fully flexible options (solid lines). Dispatch calculations base on 100k realisations each while different n is used for Delta derivation (left: 20m, middle: 10m, right: 7m).

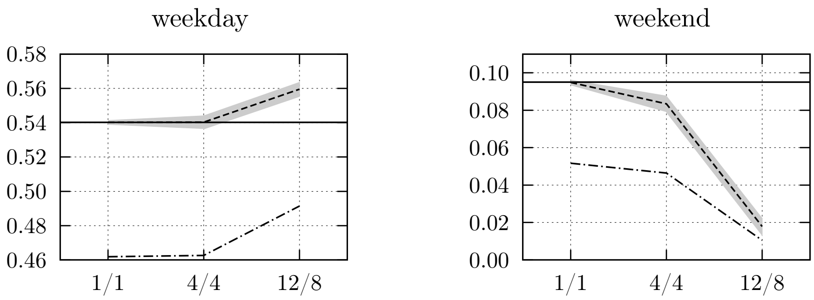

Fig. 5.14: Dependence of weekday and weekend day averaged Delta (dashed lines) and expected dispatch (dashed-dotted lines) on the size of technical constraints (indicated by t_{up}/t_{down}) in comparison to analytical results of a fully flexible option (solid lines). Simulation specifications are identical to figure 5.13.

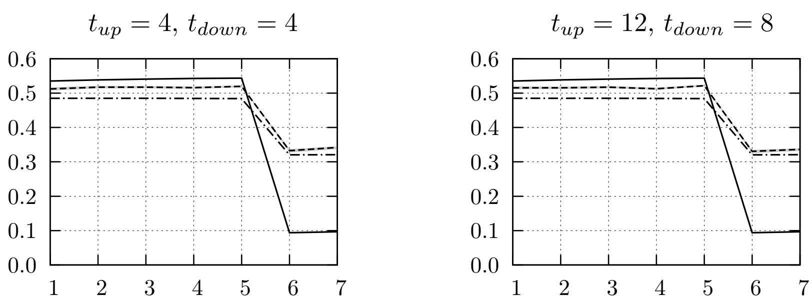

Fig. 5.15: Comparison of daily numerical Deltas (dashed lines) with daily expected dispatches (dashed-dotted lines) for options with different technical constraints and subject to positive reserve requirement of $Res_{pos} = 0.3$. Analytical results of a fully flexible option are provided for comparison (solid lines). Dispatch calculations base on 100k realisations each and Delta results on $n = 3$m.

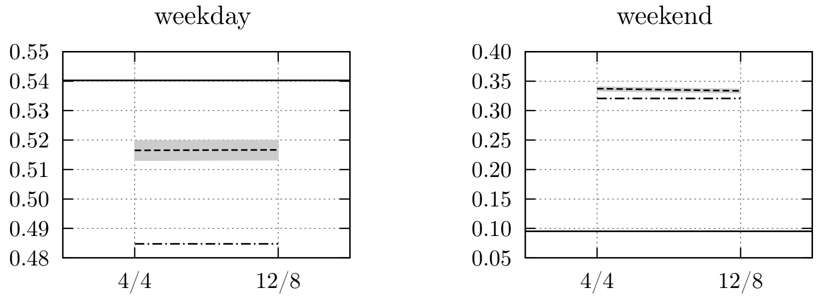

Fig. 5.16: Dependence of weekday and weekend day averaged Delta (dashed lines) and expected dispatch (dashed-dotted lines) on the size of technical constraint (indicated by t_{up}/t_{down}) when the options are subject to a positive reserve requirement of $Res_{pos} = 0.3$ in comparison to analytical results of a fully flexible option (solid lines). Simulation specifications identical to figure 5.15.

Again the systematical bias between Delta and expected dispatch is clearly visible, however its size is smaller than in case of absent reserve requirements. This is due to the fact that the bias is a direct result of the nonlinearity of the option payoff functions which is decreased by imposed reserve constraints that reduce the ability of the option to adjust its load to different hourly price levels. Anyway the hedging error is still in a range of 10% of the overall electricity production.

As expected from the discussion in section 5.2.1 different up- and down-time constraints do not have an effect on Delta or average dispatch in case of an existing positive reserve requirement.

5.4 Backtesting via Delta-Gamma hedging framework

As analytical results for sensitivities of constrained power plant options are only available for a couple of rather artificial cases it is usually not possible to backtest numerical results of this thesis against theoretical values. An alternative in such a case is to use a backtesting hedging framework that performs a step wise Delta-Gamma hedging of the evaluated option at each discrete point in time of the numerical valuation framework. A comparison of the distributions of hedged and unhedged option payoff realisations then allows to verify the hedge efficiency and thereby the quality of the applied numerical sensitivities Delta and Gamma.

This approach involves a combination of inner and outer Monte-Carlo simulations as described in section 5.4.1 which is numerically very cost demanding. Therefore we use this approach only to verify the quality of sensitivities for power plant options without additional reserve requirements and use a second backtesting method in section 5.5 to assess numerical results of options that are also subject to reserve requirements.

5.4.1 Backtesting hedging framework implementation

The implemented backtesting hedging framework consists of an outer Monte-Carlo valuation module and the standard Proxy Simulation Scheme implementation that is called by the outer module within each time step. The outer module serves three main purposes:

1. It derives outer Monte-Carlo dispatch and payoff realisations for all days of the chosen valuation tenor.

2. Within each time step of these realisations it passes the current status of the outer simulation including historical option dispatches and the current daily price forward curve over to the Proxy Simulation Scheme implementation, which derives daily Deltas and Gammas via an inner (or embedded) Monte-Carlo simulation and passes them back to the outer module.

3. Based on these sensitivities the outer module defines the amount of daily forwards and at the money Black-Scholes options that is needed to hedge the power plant option's risk position.

Thereby the backtesting implementation is able to generate outer payoff realisations of the power plant option for each day of the valuation tenor and in addition cumulated payoff realisations of the defined hedge products. This allows to directly compare payoff distributions of unhedged and hedged power plant options.

For reasons of numerical performance we use a valuation tenor of four weekdays. The number of outer Monte-Carlo realisations is always set to 100 while the amount of inner Monte-Carlo realisations is varied between 100k and 500k. Outer simulation results are translated to smooth probability distributions via standard Gaussian kernel density estimation (Kroese et al. (2011) and Wessa (2012)).

5.4.2 Results for a fully flexible option

As a reference for the backtesting hedging approach the method is at first used for backtesting numerical sensitivities of a fully flexible option. This allows to use the method also in combination with analytical sensitivities providing a natural quality check of the backtesting method itself.

Figure 5.17 provides all backtesting hedging results for a fully flexible power plant option. The graphs on the left hand side show standard deviations of unhedged daily payoffs for all four days of the valuation tenor in comparison to the standard deviations of Delta- and Delta-Gamma-hedged payoffs all based on the same 100 outer Monte-Carlo realisations. The graphs on the right hand side show associated kernel density estimators of the probability distributions of four days cumulated option payoffs.

The graphs on top in figure 5.17 base on a backtesting hedging framework that uses analytical sensitivities for hedging the fully flexible power plant option. This setup should theoretically lead to the best possible hedging efficiency within the implemented daily model as the analytical sensitivities exactly match the risk profile of the option up to second order effects. The hedging effect of Delta is clearly visible both in terms of significantly reduced daily standard deviations and a narrowed probability density of the cumulated payoffs. The hedging effect of Gamma is smaller being consistent with the model setup as the power plant option is at-the-money only for a few hours

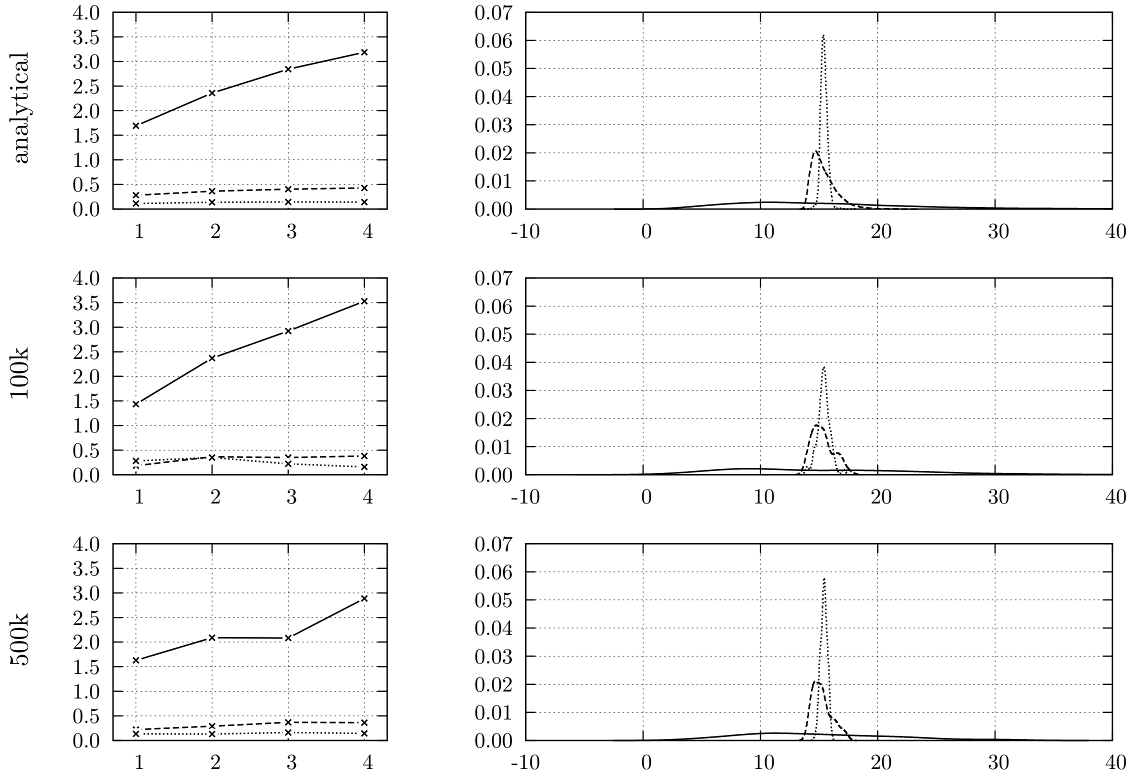

Fig. 5.17: Impact of analytical and numerical Delta hedging (dashed lines) and Delta-Gamma hedging (dotted lines) on the standard deviation of daily payoffs (left, solid lines) and gaussian kernel density of the cumulated payoff (right, solid lines) of a power plant option with $t_{up} = t_{down} = 1$, $\kappa = 50$ and $Res_{pos} = 0.0$. Top: analytical hedging, middle: numerical hedging with 100k inner Monte-Carlo realisations, bottom: numerical hedging with 500k inner Monte-Carlo realisations.

per day. The remaining standard deviation of the Delta-Gamma hedged option is due to the fact that the implemented model allows hedging only in daily granularity. This can obviously only serve as a rough approximation for continuous hedging which would theoretically yield zero standard deviations.

The graphs in the middle of figure 5.17 provide results based on 100k inner Monte-Carlo realisations for the calculation of numerical sensitivities. Both daily standard deviations and the probability density of the unhedged option are not identical to the analytically hedged results as this second backtesting uses different Monte-Carlo realisations than applied in the first example. It can be observed that especially the hedge effect of Gamma is not as good as in the case of analytical hedge parameters. However the overall hedge efficiency of the numerical sensitivities is obvious – even though the chosen amount of 100k inner realisations is rather insufficient to produce stable sensitivities.

Therefore the graphs at the bottom of figure 5.17 provide additional results based on 500k inner Monte-Carlo realisations. Again the results base on different realisations which explains slightly different curve shapes but now the hedge efficiency of the numerical sensitivities is in the range of analytical sensitivities. This confirms the validity of the backtesting hedging approach as well as the quality of the numerical sensitivities for the fully flexible power plant option.

5.4.3 Results for technically constrained options

Similar to figure 5.17 figure 5.18 provides results of the backtesting hedging framework for a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$) without external reserve requirements.

The hedge efficiency of the numerical sensitivities is clearly visible in case of 500k inner Monte-Carlo realisations. As can be expected the numerical sensitivities are superior both in terms of reduced daily standard deviations and a sharper peak structure of the probability density distribution than analytical sensitivities for a fully flexible option, which deviate from appropriate numerical values (compare e.g. figure 5.8). However the overall difference is minor as the deviation between numerical and analytical daily sensitivities is rather small, being especially true for Delta that leads to the majority of the hedging effect.

Therefore the backtesting framework is also applied in combination with more heavily constrained options whose numerical sensitivities deviate more from theoretical results. We use two option specifications that have already been discussed in sections 5.1.1 and 5.1.2 (figures 5.1 and 5.5), i.e. $t_{up} = 24$, $t_{down} = 1$, $\kappa = 50$ and $t_{up} = 1$, $t_{down} = 24$, $\kappa = 50$. Associated backtesting results based on both analytical and numerical sensitivities (500k inner realisations) are provided in figures 5.19 and 5.20. Both visualisations clearly show the on average higher hedging efficiency of the numerical results, which serves as an indication for the applicability of our numerical Deltas and Gammas as estimators of the real first and second order derivatives.

This finding is further supported by figure 5.21 showing the impact of Delta hedging and Delta-Gamma hedging on the normalised standard deviation of the cumulated option payoffs in waterfall diagrams. In both cases it can be observed that the hedge

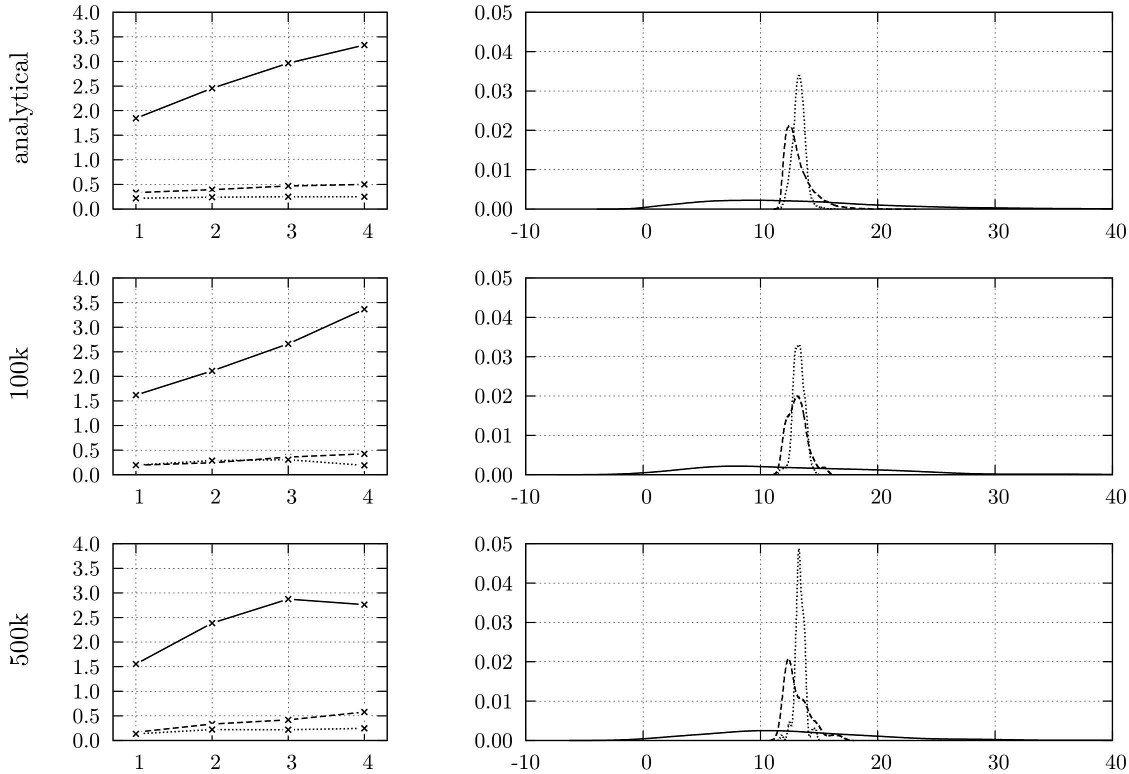

Fig. 5.18: Impact of analytical and numerical Delta hedging (dashed lines) and Delta-Gamma hedging (dotted lines) on the standard deviation of daily payoffs (left, solid lines) and gaussian kernel density of the cumulated payoff (right, solid lines) of a power plant option with $t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$ and $Res_{pos} = 0.0$. Top: analytical hedging, middle: numerical hedging with 100k inner Monte-Carlo realisations, bottom: numerical hedging with 500k inner Monte-Carlo realisations.

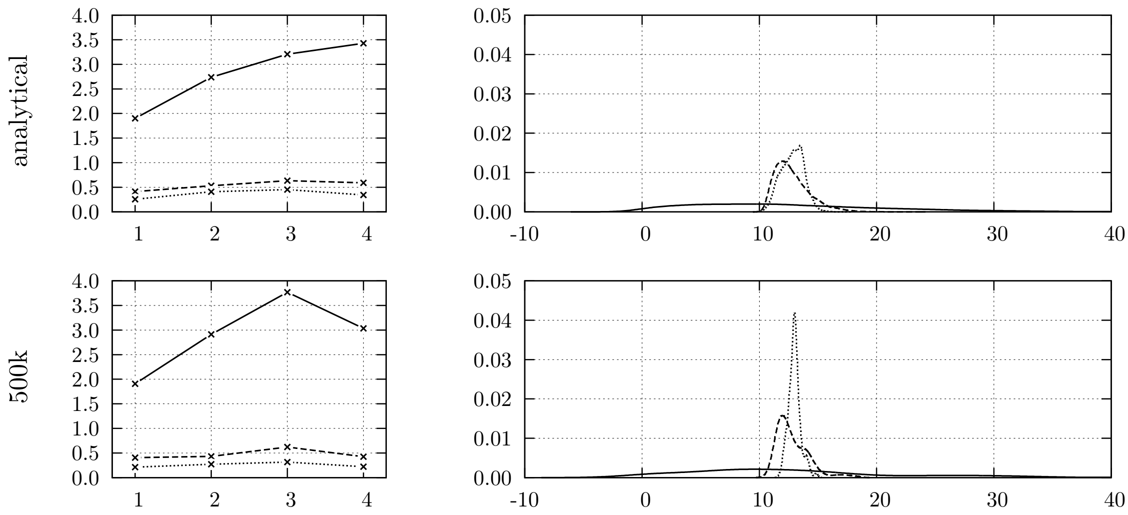

Fig. 5.19: Impact of analytical and numerical Delta hedging (dashed lines) and Delta-Gamma hedging (dotted lines) on the standard deviation of daily payoffs (left, solid lines) and gaussian kernel density of the cumulated payoff (right, solid lines) of a power plant option with $t_{up} = 24$, $t_{down} = 1$, $\kappa = 50$ and $Res_{pos} = 0.0$. Top: analytical hedging, bottom: numerical hedging with 500k inner Monte-Carlo realisations.

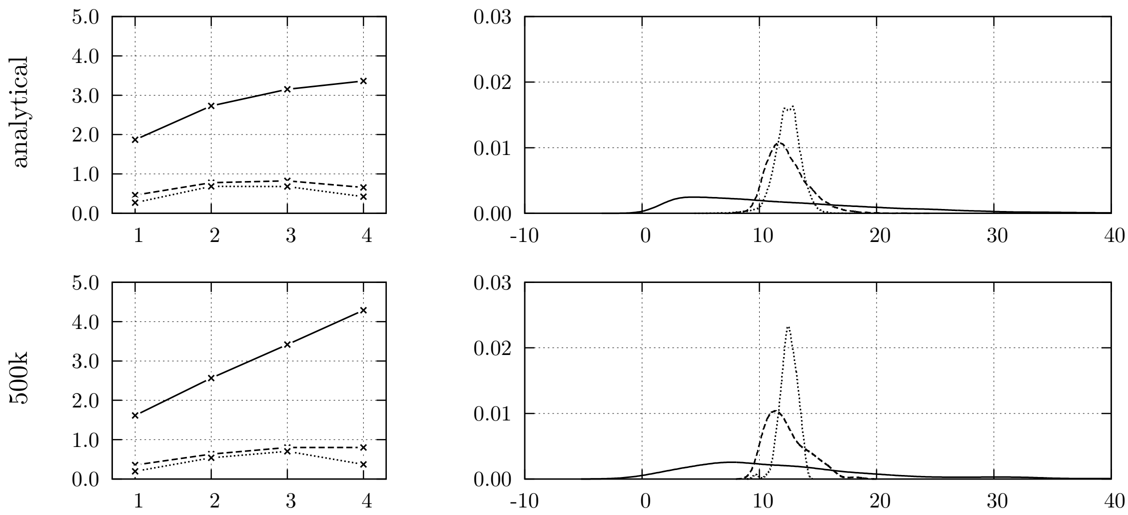

Fig. 5.20: Impact of analytical and numerical Delta hedging (dashed lines) and Delta-Gamma hedging (dotted lines) on the standard deviation of daily payoffs (left, solid lines) and gaussian kernel density of the cumulated payoff (right, solid lines) of a power plant option with $t_{up} = 1$, $t_{down} = 24$, $\kappa = 50$ and $Res_{pos} = 0.0$. Top: analytical hedging, bottom: numerical hedging with 500k inner Monte-Carlo realisations.

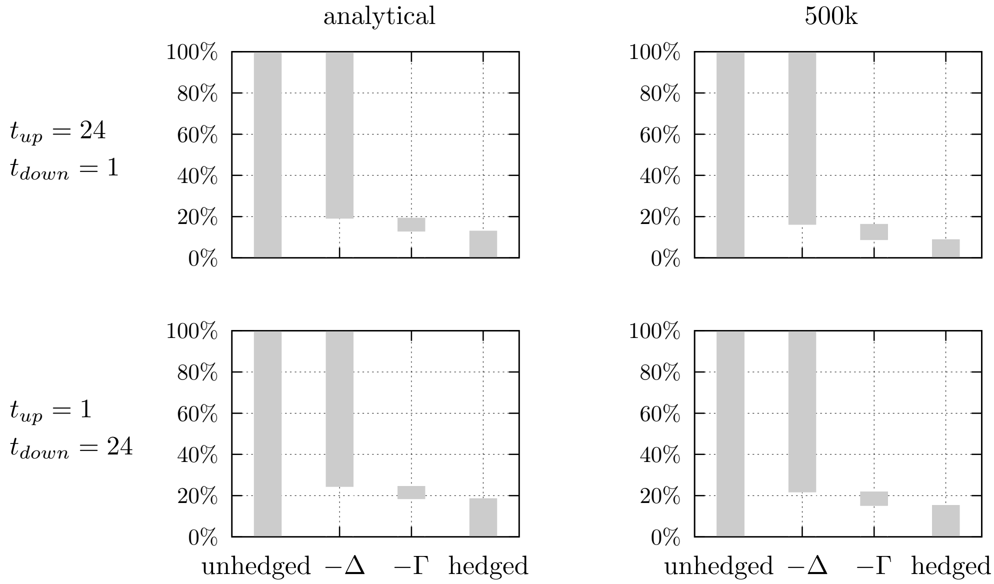

Fig. 5.21: Percentage decrease of the normalised standard deviation of the cumulated option payoff due to Δ hedging and additional Γ hedging ($\kappa = 50$, $Res_{pos} = 0.0$). Left: analytical hedging, right: numerical hedging with 500k inner Monte-Carlo realisations.

efficiency of the numerical Delta exceeds the efficiency of the analytical Delta and even though the impact is smaller in absolute terms, the numerical Gamma is also superior to the analytical reference values.

5.5 Backtesting via single step probability distributions

5.5.1 Implemented method

As explained in section 5.4 we use a more numerical cost efficient method in addition to the backtesting hedging framework to assess the quality of numerical sensitivities also for power plant options subject to reserve requirements. This approach utilises the fact that the step wise evolution of the value of a power plant option dV can be expressed in terms of its sensitivities via a Taylor series (e.g. Riley et al. (2006)). A truncation of this series after the second order sensitivity yield the famous Δ-Γ-approximation (e.g. Deutsch (2004)) which is also commonly used as the basis to derive a simplified Value-at-Risk (e.g. Jorion (2011)):

$$\begin{aligned} dV(\mathbf{X}(t_0)) &= V(\mathbf{X}(t_0) + d\mathbf{X}(t_0)) - V(\mathbf{X}(t_0)) \\ &\approx \sum_{j=1}^{m} \frac{\partial V}{\partial X_j}(t_0) dX_j(t_0) + \frac{1}{2} \sum_{j=1}^{m} \frac{\partial^2 V}{\partial X_j^2}(t_0)(dX_j(t_0))^2 \\ &= \sum_{j=1}^{m} \Delta_j(t_0) dX_j(t_0) + \frac{1}{2} \sum_{j=1}^{m} \Gamma_j(t_0)(dX_j(t_0))^2. \end{aligned} \quad (5.1)$$

The value difference on the left hand side of equation (5.1) can easily be determined by numerically deriving expected option payoffs based on both the initial daily forward price curve $\mathbf{X}(t_0)$ and the same forward curve after a random single step price evolution $\mathbf{X}(t_0+\Delta t)$ and taking the difference. The terms on the right hand side depend only on already known random price evolutions and numerical sensitivities $\Delta_j(t_0)$ and $\Gamma_j(t_0)$ which can been derived ex-ante by separate Proxy Simulation Scheme simulations. Therefore a comparison of probability distributions of both sides of equation (5.1) provides an elegant and pragmatic way to backtest available numerical sensitivities in addition to the full backtesting hedging framework as used in section 5.4. The method also allows to separate the effects of Delta and Gamma by either taking only the first term or both terms on the right hand side into account. In addition it is also a cost efficient alternative to assess the applicability of expected dispatches by using those on the right hand side instead of $\Delta_j(t_0)$.

For this analysis we use three already discussed option setups and 500 random evolutions of a 7 day price forward curve (i.e. $m = 7$) with 2000 Monte-Carlo payoff realisations for the derivation of the left hand side of equation (5.1). The value distributions of both sides are visualised and complemented by a comparison of associated curve descriptors including mean, standard deviation, skew, kurtosis and minimum and maximum value of all discussed alternatives. Similar to section 5.4 standard Gaussian kernel density estimation is used to derive smooth value distributions.

5.5.2 Results for a fully flexible option

The analysis is initiated with the reference example of a fully flexible power plant option without reserve requirements as discussed in sections 4.4 (figure 4.5) and 5.3.1 (figure 5.13). In this setup all numerical sensitivities have already been confirmed by a comparison to available analytical results and there is no additional value in replicating this conclusion via an alternative method. However it it valuable to compare numerical Deltas to expected dispatches especially for this most flexible option.

The graph at the top in figure 5.22 shows the distribution of 7 days cumulated value changes of the fully flexible power plant option in combination with associated Δ-approximations and Δ-Γ-approximations corresponding to only the first and both terms on the right hand side of equation (5.1). The graph at the bottom shows similar distributions of the flexible option and its value replication by using expected dispatches instead of Δ. Here no further Γ-hedging is applied in order to not mix the impact of sensitivities with the impact of the expected dispatch.

While only analytical Delta is obviously not sufficient to cover all characteristics of the single step value changes the combination of Delta and Gamma leads to a good alignment of the distribution curves, which is verified by low deviations of the associated curve descriptors on the right hand side of figure 5.22. As expected, due to the inherent flexibility of the analysed option a replication of value changes via expected dispatches is much worse than the replication via analytical Delta. Interestingly while the distributions show huge differences for positive value changes only minor deviations can be observed in the region of the left tail of the density curve. This makes the expected dispatch a sufficient approximation for Delta in this special setup as long as only the downside risk of the option value is requested.

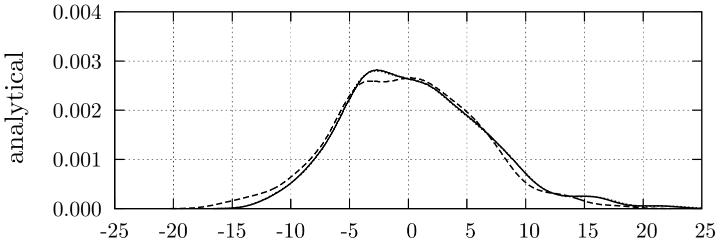

	option	Δ	Δ-Γ
mean	0.63	-0.14	0.64
SD	5.88	5.82	5.92
skew	0.56	0.11	0.59
kurt	0.44	0.06	0.51
min	-11.88	-15.71	-12.00
max	21.79	17.85	22.40

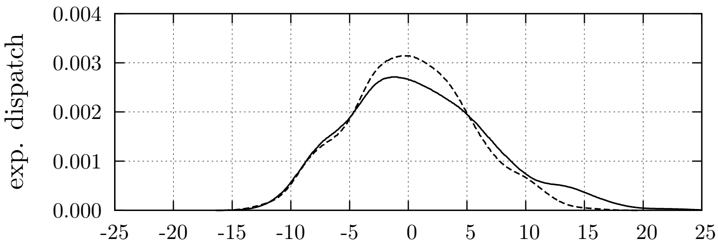

	option	exp. dispatch
mean	0.98	0.15
SD	6.08	5.06
skew	0.52	0.13
kurt	0.09	-0.26
min	-11.58	-12.46
max	22.88	15.59

Fig. 5.22: Comparison of gaussian kernel densities of 1 day value changes (solid lines), Delta or expected dispatch approximations (dashed lines) and Delta-Gamma approximations (dotted lines) for an option with $t_{up} = t_{down} = 1$ and $\kappa = 50$. Top: analytical sensitivities, bottom: expected dispatch.

5.5.3 Results for options subject to technical constraints

Similar to figure 5.22 figure 5.23 shows the distribution of cumulated value changes of a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$) in combination with associated Δ-approximations, Δ-Γ-approximations and approximations on the basis of expected dispatches.

Again the expected dispatch alone is clearly not sufficient to replicate the observed option value changes. The same is true for both analytical and numerical Deltas which lead to similar deviations between the original distributions and their Δ approximations. However both analytical and numerical Δ-Γ approximations lead to rather good replications of the option value changes with a slight superiority of the numerical sensitivities. This is due to the fact that the absolute deviation between numerical and analytical sensitivities is still minor for this option setup (compare figure 5.8).

5.5.4 Results for options subject to technical constraints and reserve requirements

Figure 5.24 shows the distribution of cumulated value changes of a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 50$) subject to positive reserve requirement of $Res_{pos} = 0.3$ in combination with associated Δ-approximations, Δ-Γ-approximations and approximations on the basis of expected dispatches.

Similar to already discussed examples the expected dispatch is not able to replicate the original distribution curve. However in contrast to the setup without reserve requirement (figure 5.23) the results now show also significant differences between approximations based on analytical and numerical results. While analytical sensitivities lead to relevant deviations between curve descriptors (especially mean, skewness and

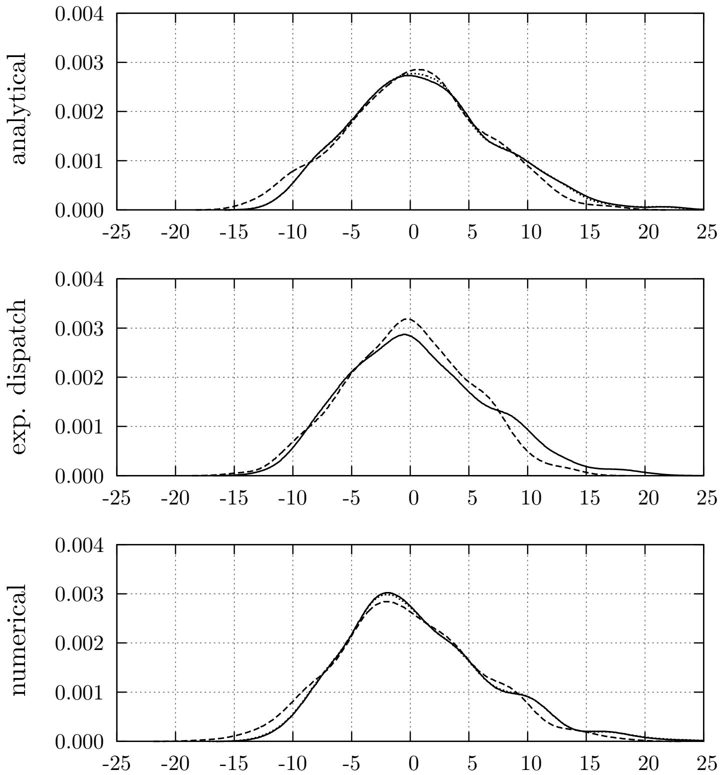

	option	Δ	Δ-Γ
mean	1.22	0.39	1.18
SD	6.08	5.86	6.02
skew	0.48	0.05	0.46
kurt	0.12	-0.25	0.10
min	-11.17	-13.81	-11.19
max	22.36	17.34	21.86

	option	exp. dispatch
mean	0.67	-0.13
SD	5.91	5.21
skew	0.47	0.04
kurt	-0.01	-0.28
min	-12.27	-14.53
max	20.48	14.47

	option	Δ	Δ-Γ
mean	0.59	-0.23	0.59
SD	6.09	5.98	6.14
skew	0.68	0.22	0.69
kurt	0.42	-0.05	0.48
min	-12.15	-17.35	-12.68
max	23.47	18.55	23.84

Fig. 5.23: Comparison of gaussian kernel densities for 1 day value changes (solid lines), Delta or expected dispatch approximations (dashed lines) and Delta-Gamma approximations (dotted lines) for an option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$. Top: analytical sensitivities, middle: expected dispatch, bottom: numerical sensitivities.

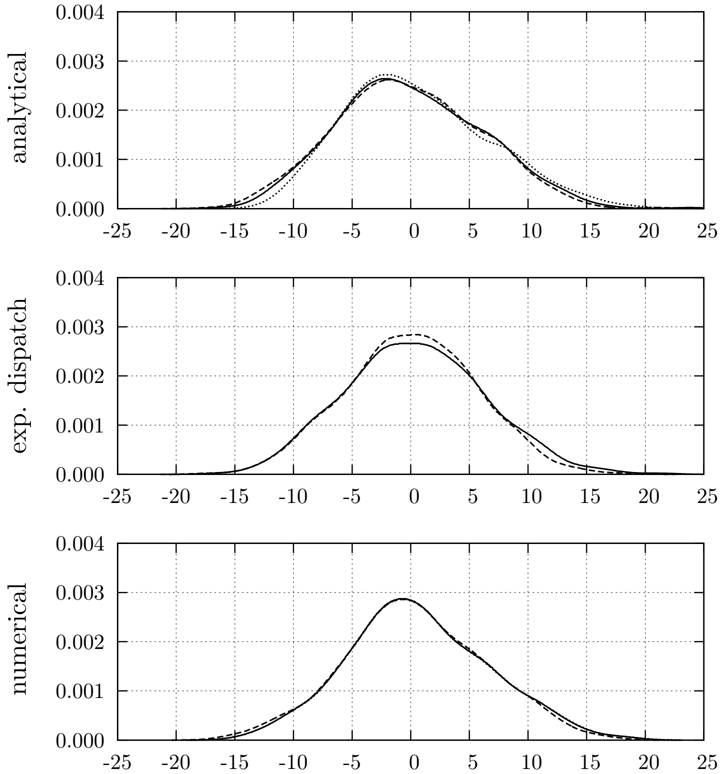

	option	Δ	Δ-Γ
mean	0.06	-0.26	0.57
SD	6.09	6.04	6.15
skew	0.29	0.14	0.57
kurt	-0.12	-0.27	0.27
min	-15.78	-16.55	-12.67
max	23.79	20.92	27.28

	option	exp. dispatch
mean	0.38	0.07
SD	5.93	5.57
skew	0.19	0.02
kurt	-0.04	-0.10
min	-16.17	-17.05
max	21.25	18.25

	option	Δ	Δ-Γ
mean	0.68	0.38	0.69
SD	5.98	5.95	6.00
skew	0.23	0.06	0.24
kurt	-0.15	-0.17	-0.13
min	-14.48	-15.86	-14.43
max	18.46	16.96	18.58

Fig. 5.24: Comparison of gaussian kernel densities for 1 day value changes (solid lines), Delta or expected dispatch approximations (dashed lines) and Delta-Gamma approximations (dotted lines) for an option with $t_{up} = 12$, $t_{down} = 8$ and $\kappa = 50$ subject to positive reserve requirements of $Res_{pos} = 0.3$. Top: analytical sensitivities, middle: expected dispatch, bottom: numerical sensitivities.

kurtosis) of the original distribution and its Δ-Γ approximations, numerical sensitivities lead to a very good approximation of the distribution curve which is also proven by very low deviations of the curve descriptors.

This example makes clear that in case of existing reserve requirements analytical sensitivities are no longer sufficient to describe the risk profile of the constrained power plant option. This is also true when the option is evaluated on a stand-alone basis and no additional portfolio effects are taken into account. In addition the good approximation of distribution curves for all discussed option setups supports the relevance, applicability and high quality of the Proxy Simulation Scheme sensitivities as derived in this thesis.

6 Sensitivities of option portfolios

This chapter extends the single option analyses of chapter 5 by a general analysis of portfolio effects and intra-option dependencies appearing when a stand-alone power plant option subject to reserve requirements is split in more and more almost identical smaller options. This approach leads to complex result structures that can be explained by the existence of distinct dispatch regimes depending on the level of realised daily electricity prices.

After this introduction to constrained option portfolios we continue with a detailed discussion of the impact of an increasing external reserve requirement on both value and sensitivities of three portfolios now containing different power plant options with unequal technical characteristics. By comparing these results with associated portfolio results without reserve requirements we are finally able to derive values and the sensitivities Δ *and* Γ *of various reserve requirements in the context of the assessed portfolios. This analysis is one of the main result of this thesis and to the knowledge of the author has not been presented in academic literature before.*

6.1 Transition from single option to portfolio of multiple identical options

6.1.1 Stand-alone power plant option

Figure 6.1 shows numerical results of a technically constrained power plant option ($t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.3$, $P_{max} = 1.0$, $\kappa = 50$) subject to a positive reserve requirement of $Res_{pos} = 0.2$. As already discussed in section 5.2.1 (figure 5.8) the reserve requirement leads to a significant decrease of daily option values on weekdays and even more on weekend days. Delta is generally increased and also this impact is more prominent on weekend days. In contrast Gamma is heavily decreased on all days of the valuation tenor due to the generally reduced load flexibility. Especially on weekend days these effects are a direct result of the must-run condition that is imposed on the option by the positive reserve requirement.

In the following parts of this section the stand-alone power plant option is split in multiple, identical options which are on a portfolio level always subject to the same reserve requirement of $Res_{pos} = 0.2$. The goal of this analysis is to provide a comprehensive introduction to observable portfolio effects and dispatch dependencies between different options (intra-option coupling effects) which are also relevant in later analyses of portfolios with different power plant options.

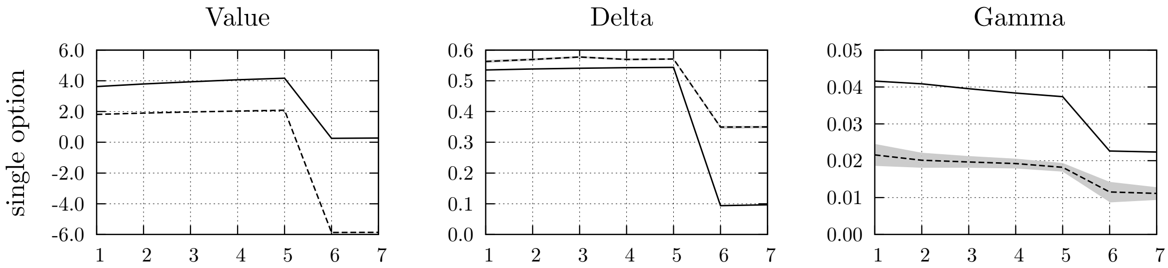

Fig. 6.1: Numerical results of a technically constrained stand-alone power plant option with $t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.3$, $P_{max} = 1.0$, which is subject to a positive reserve requirement of $Res_{pos} = 0.2$ (dashed lines) in comparison to analytical results of a fully flexible option (solid lines). $n = 3$m.

6.1.2 Split of original option into two parts

In a first step the single option (figure 6.1) is split in two options of each half the size of the original option, i.e. $t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.15$, $P_{max} = 0.5$. Both options are identical apart from the strike prices $\kappa^{1,2}$ showing a total deviation of 0.02: $\kappa^1 = 50 - 0.01 < \kappa^2 = 50 + 0.01$. This deviation does not lead to relevant changes of numerical results in case the options are evaluated stand-alone, but in case of a portfolio evaluation it allows the mixed-integer optimisation solver to discriminate the two options. This is not only a numerically driven artificial modification but effectively a rough replication of existing physical or technical constraints which have an impact on dispatch decisions for real power plants. Even in case that a real generation portfolio contains two identically constructed assets these will typically have slightly different run time parameters which leads to a minor preference of one asset over the other when highly flexible serving of reserve requirement is requested.

Figure 6.2 shows associated numerical results of both resulting options in comparison to analytical results for fully flexible options of the same size and figure 6.3 show the same results aggregated on the portfolio level and therefore directly comparable to figure 6.1.

On the portfolio level a value increase with respect to the the original option can be observed for all days. This effect is especially large on weekend days. Delta is generally decreased while Gamma is increased for all days. The higher values can be explained by a lowered must-run condition of the reserve requirement in this setup. As long as a single option is obliged to serve the reserve requirement it is forced to stay online in all hours of the valuation tenor. This leads to a negative impact on the overall value as the option runs at minimum load P_{min} in all out of the money hours instead of being shut down. In contrast each of the split options is able to serve the requested reserve requirement of $Res_{pos} = 0.2$ independently. This enables to shut down one of the smaller options if appropriate as long as the other option stays online, and thereby allows to reduce the negative impact of out of the money hours. In addition the halved minimum load of $P_{min} = 0.3/2$ further reduces the negative value contribution of the reserve serving option by a factor of 2 which can directly be observed by comparing daily values on weekend days in figures 6.1 and 6.3.

This explanation is also supported by the individual option results as provided in

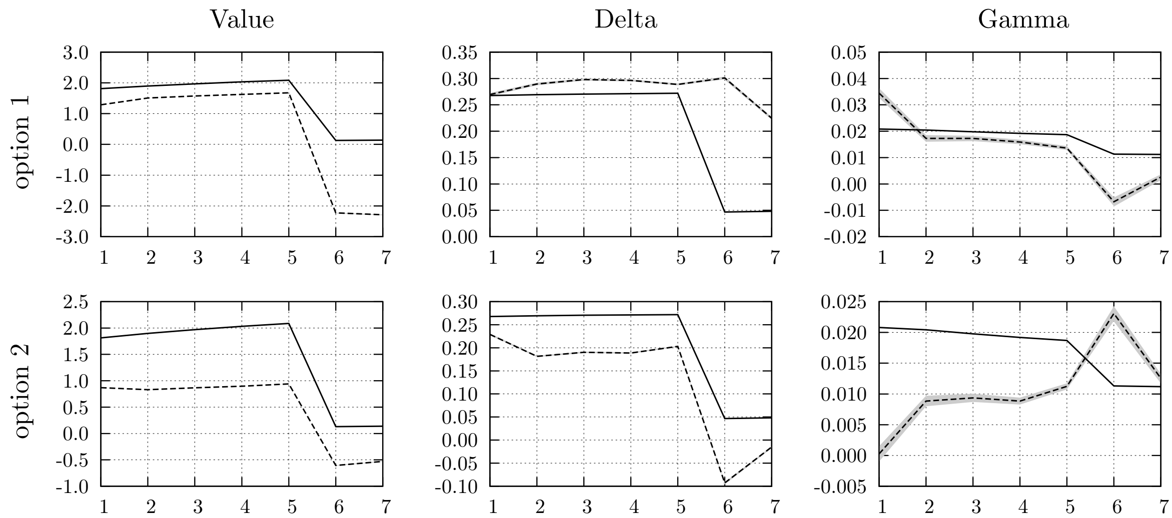

Fig. 6.2: Numerical results for a portfolio of two technically constrained power plant options each having half the size of the option as shown in figure 6.1 ($t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.3/2$, $P_{max} = 1.0/2$) again subject to positive reserve requirement of $Res_{pos} = 0.2$ (dashed lines), in comparison to analytical results of fully flexible options (solid lines). $n = 3$m.

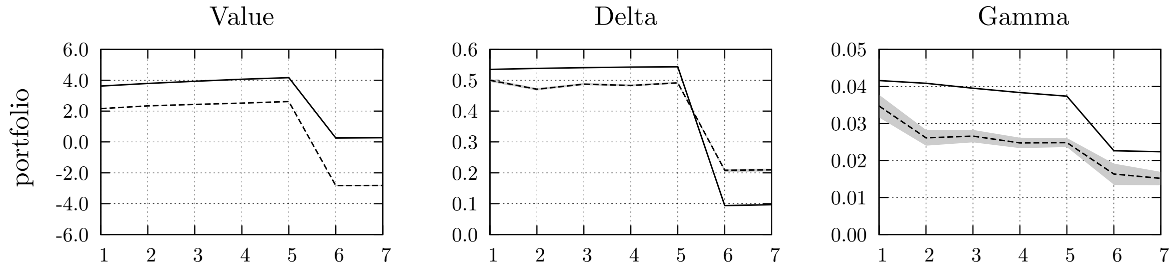

Fig. 6.3: Numerical portfolio results of the two options being individually shown in figure 6.2 (dashed lines) in comparison to analytical results of fully flexible portfolio (solid lines). $n = 3$m.

figure 6.2. On weekdays the slightly cheaper option 1 has significantly higher daily values and Deltas than option 2. This is due to the fact that the costlier option 2 will always serve positive reserve requirement when the overall price level allows both options to stay online at the same time. This is also the explanation for the general deviation of Gamma on weekdays. The load flexibility of option 2 is more often limited to the interval [0.15, 0.3] which leads to an overall reduced Gamma with respect to option 1. Please note that the average Gamma of option 2 on days 2 to 4 (that are not affected by boundary conditions like day 1) is approximately half the size of the associated Gammas of option 1, which is consistent with the relative sizes of load flexibility intervals of the options ([0.15, 0.3] of option 2 and [0.15, 0.5] of option 1). This indicates that both options often run in parallel on these days.

On the first day of the valuation tenor it can be observed that Gamma of option 1 is significantly increased with respect to subsequent weekdays while Gamma of option 2 is decreased. This is a combined effect of the left boundary condition of the mixed integer solver at the first hour of day 1 and the chosen hourly price structure (figure 3.4). In the applied setup option 1 will always be dispatched in the first hour of the first day and will therefore on average be online for more hours on the first day than on other weekdays. Consequently option 1 is less restricted by exercise decisions of option 2 on the first day than on subsequent weekdays. This can be interpreted as a general increase of flexibility of option 1 on the first day which results in the observed uplift of Gamma. At the same time the dispatch of option 2 will always be aligned to the already set dispatch of option 1 on day 1 which results in an on average lower flexibility and a decrease of Gamma for option 2. On the portfolio level the uplifting effect on option 1 dominates leading to a minor increase of Gamma on the first day. Without further insights into intra-option dependencies this can also be interpreted as a result of the open left boundary condition of the mixed integer solver that allows to dispatch both options on the first day independently of historical dispatch decisions in contrast to all other days of the valuation tenor.

Distinct dispatch regimes Interestingly on day 6 (i.e. the first day of the weekend) the Gamma of option 1 drops to a negative value while the Gamma of option 2 is significantly increased and accompanied by a negative Delta. Especially the negative option Gamma is remarkable as both options would be expected to keep a residual flexibility also on weekend days as the applied reserve requirement of $Res_{pos} = 0.2$ is comparably low.

These observations can be explained on the basis of a detailed analysis of hourly dispatches of both options on days 5 to 7. This analysis reveals that the variety of realised hourly dispatches can be assigned to not more than three predominant different dispatch regimes (1, 2 and 3) with up to three sub-regimes (a, b and c) each. The occurrence of a certain distinct dispatch (sub-)regime is defined by the absolute level of the realised daily base load price and historic dispatches on preceding days. All major dispatch regimes are shown in figure 6.4.

Regime 1 occurs for realised daily prices below a threshold around 29.09 where both options are out of the money for all hours of day 6. It is characterised by

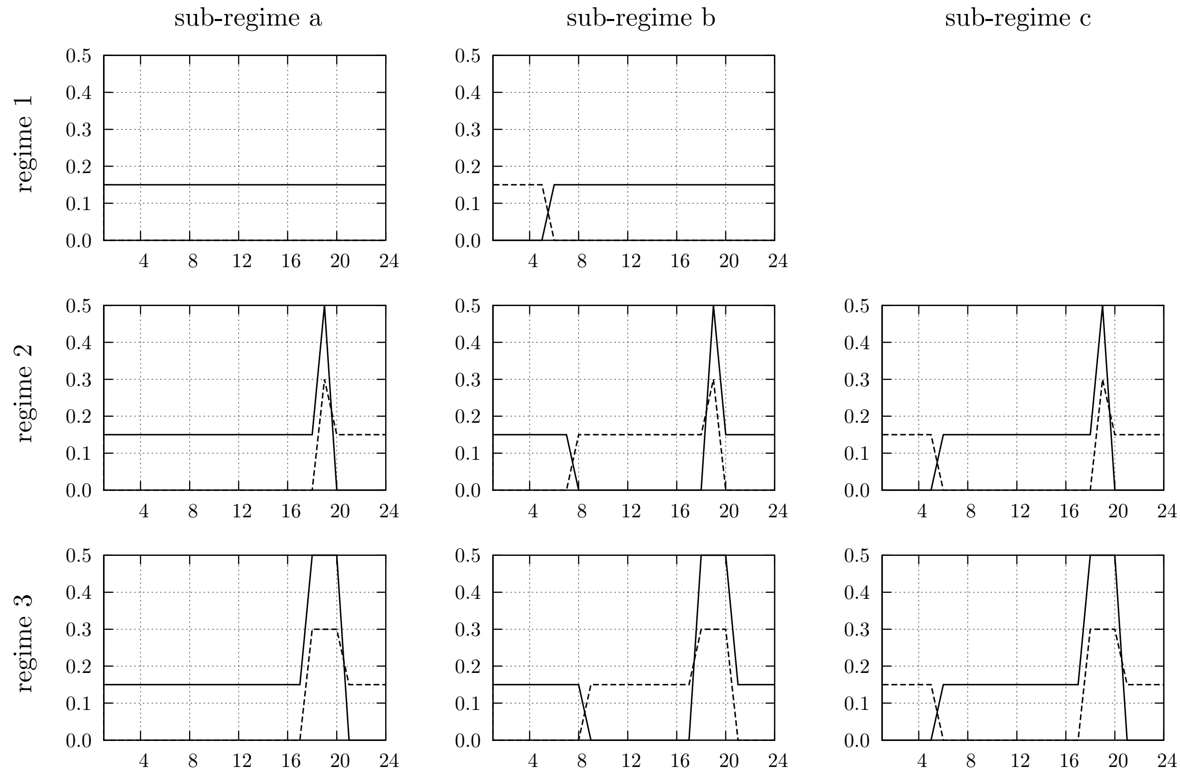

Fig. 6.4: Hourly dispatch of option 1 (solid lines) and option 2 (dashed lines) for all three major dispatch regimes and sub-regimes a, b and c occurring on day 6.

option 1 typically running on minimum load in order to serve reserve requirement and option 2 being offline as many hours as possible. In sub-regime 1a option 1 runs on minimum load in all hours and option 2 is offline during the entire day. Sub-regime 1b is disturbed by historic dispatch decisions on day 5 which require option 2 to keep running during the first hours of day 6 until it is shut down and option 1 starts serving reserve requirements.

Regime 2 occurs for daily prices between 29.09 and 32.05. At these price levels both options are in the money for exactly one evening hour (h19) and each optimal dispatch solution requires both options to be online in this hour in order to generate the maximum possible payoff. In contrast it is most economical to only keep one of the options online for all remaining hours of day 6 as both options are deep out of the money. Therefore all sub-regimes of regime 2 imply that only one option runs at minimum load before hour 19 (typically option 1) and the second option is started for hour 19 where both options run at maximum effective load (i.e. option 1 at maximum load and option 2 at maximum load less reserve requirement). After hour 19 the second option cannot be shut down immediately due to technical constraints. It continues running at minimum load and serving the reserve requirement while the first option is shut down after hour 19. Sub-regime 2a is defined by option 1 running at minimum load until hour 19 where it is running at maximum load before being shut down. Option 2 is running in hour 19 at maximum effective load and continues to run at minimum load afterwards. Sub-regime 2c is identical to sub-regime 2a apart from a disturbance

Dispatch (sub-)regime	rel. frequency	daily dispatch option 1	daily dispatch option 2
1a	65%	3.60	0.00
1b	29%	2.85	0.75
2a	42%	3.20	1.05
2b	30%	2.30	1.95
2c	23%	2.45	1.80
3a	49%	4.05	1.50
3b	34%	3.30	2.25
3c	10%	3.30	2.25

Table 6.1: Dispatch and relative frequency of all major dispatch regimes and sub-regimes occurring on day 6.

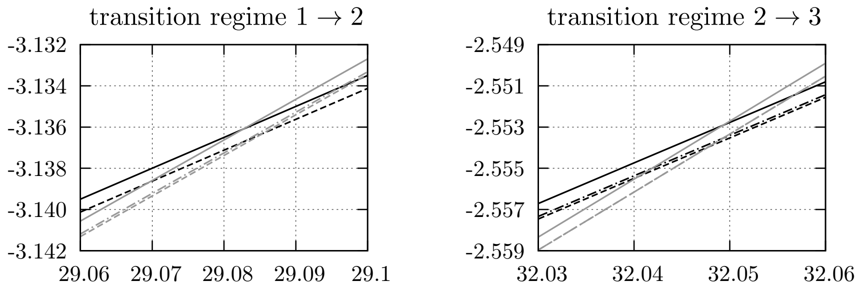

Fig. 6.5: Cumulated daily payoff of different regimes over realised price of day 6 where black lines indicate lower regimes and gray lines higher regimes, sub-regimes are indicated by line type (a: solid line, b: dashed line, c: dashed-dotted line).

of historic dispatch decisions during the first hours of day 6. Sub-regime 2b is identical to 2c but both options have changed roles due to historical dispatch decisions.

Regime 3 occurs for daily prices above 32.05 where both options are in the money in multiple consecutive evening hours. The associated dispatch (sub-)regimes are therefore very close to regime 2 with the only difference that both options now overlap for all in the money hours instead of only one.

Relative frequency and portfolio impact of dispatch regimes The relative frequency of the sub-regimes strongly depends on the structure of the applied hourly prices. Table 6.1 provides an overview about relative frequency and associated daily dispatches of both options. As the individual option dispatches of different sub-regimes show significant deviations it is surprising that all three regimes can be separated by rather well defined thresholds in terms of the realised daily base load price. This is only possible because the cumulated daily payoffs of all sub-regimes of regime 2 exceed the associated payoffs of regime 1 in the same small interval between 29.08 and 29.09 of realised daily price. Similarly the payoffs of all sub-regimes of regime 3 exceed associated payoffs of regime 2 between 32.04 and 32.05. Both regime switches are explicitly shown in figure 6.5.

Figure 6.6 finally provides a comprehensive overview about daily dispatches and daily payoffs of both options and the associated portfolio for 1000 individual Monte-

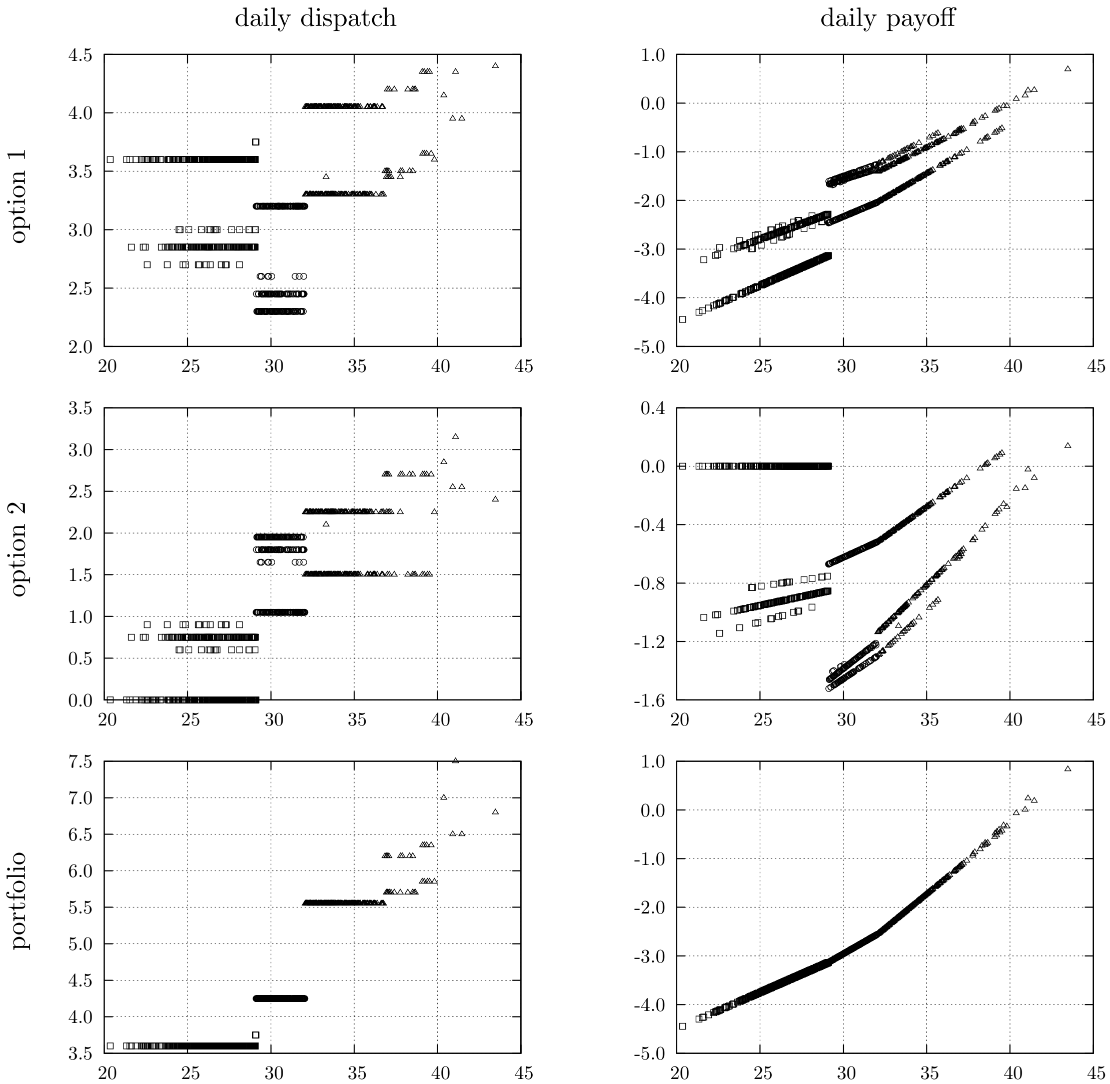

Fig. 6.6: Dispatch and Payoff realisations for both individual options (compare figure 6.2) over realised daily prices of day 6 and the same picture for the portfolio of both options. The figure shows 1000 individual Monte-Carlo realisations where regime 1 is indicated by cubes, regime 2 by circles and regime 3 by triangles.

Carlo realisations of day 6. The two sub-regimes of regime 1 are clearly visible as the individual option results are concentrated on two lines for realised prices below 29.09. By comparison to table 6.1 it can be conducted that the upper dispatch line of option 1 denotes sub-regime 1a and the lower line sub-regime 1b while this order is inverted for option 2. The dispatch line of sub-regime 1b is accompanied by two side lines containing a minor subset of result realisations with slightly different dispatch structures than allocated to sub-regime 1b. These side lines correspond to variations of sub-regimes 1b where the shut-down of option 2 and the start-up of option 1 is taking place one hour earlier or one hour later due to historical dispatch decisions on day 5. On the portfolio level it can clearly be observed that regime 1 is characterised by a daily dispatch of 3.60 independently from the realised sub-regime.

The three sub-regimes of dispatch regime 2 lead to three associated dispatch lines in figure 6.6 for realised prices between 29.09 and 32.05. On portfolio level regime 2 is defined by a cumulated daily dispatch of 4.25. In contrast to dispatch regime 2 the three sub-regimes of regime 3 only lead to two main dispatch lines in figure 6.6. This is due to the fact that sub-regimes 3b and 3c show identical daily dispatches for both options (table 6.1). Regime 3 is characterised by a cumulated daily dispatch of 5.55. For prices above 37.00 the distinctive regime structure breaks as these higher prices lead to an increasing number of in the money hours and therewith a significant increase of dispatch variations for small price differences.

Each dispatch line on the left hand side in 6.6 corresponds to one payoff line on the right hand side. However the allocation of sub-regimes to payoff lines is exactly inverted with respect to the order of dispatch lines, i.e. while sub-regime 1a corresponds to the upper dispatch line it corresponds to the lower payoff line and vice versa for sub-regime 1b. This is due to the fact that both options are deep out of the money on day 6 which leads to a cumulated payoff decrease for an increased daily dispatch. On portfolio level the payoff is a monotonically increasing function of the daily realised price which is consistent with the expectation, that the value of a call option portfolio should increase in line with underlying price realisations.

Explanation of negative sensitivities The discussed dispatch regimes allow to explain the observed negative Delta of option 2 on day 6 as well as the huge Gamma variations for both options between day 5 and day 6. In case of comparably low prices on day 6 the status quo for both options is dispatch regime 1 where the value contribution of option 2 is zero with a probability of 65%. As the initial forward price for day 6 is 30.00, positive initial price perturbations will increase the likelihood of dispatch regime 2 to realise while negative initial price perturbations will not change the dispatch regime. A higher probability of regime 2 is equivalent with a higher probability of negative payoff realisations of option 2. Therefore small initial price lifts will lead to an overall value decrease of option 2 while small initial price drops will not change the option value. In combination these dependencies lead to a negative first derivative for option 2 with respect to the initial forward price of day 6 and therefore the observed negative Delta.

The explanation for the negative Gamma of option 1 on day 6 is more complex.

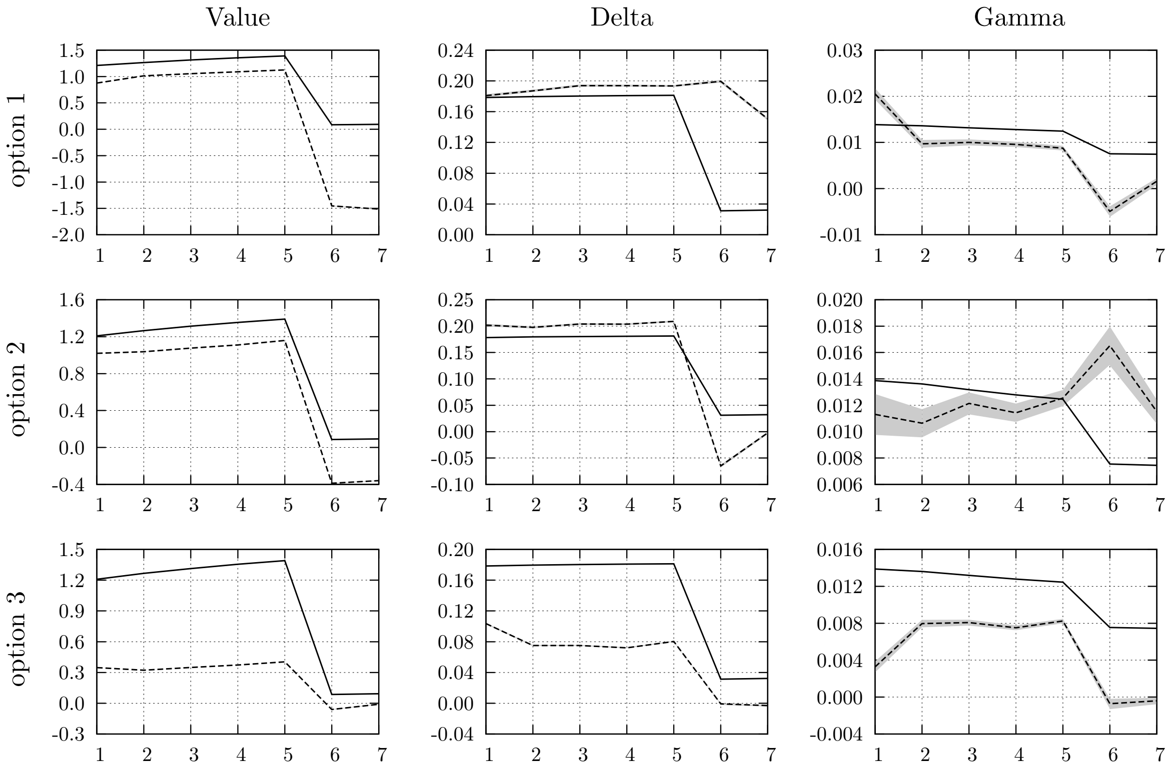

Fig. 6.7: Numerical results for a portfolio of three technically constrained power plant options each having a third of the size of the option as shown in figure 6.1 ($t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.3/3$, $P_{max} = 1.0/3$) subject to a positive reserve requirement of $Res_{pos} = 0.2$ (dashed lines), in comparison to analytical results of fully flexible options (solid lines). $n = 3$m.

Again the status quo for low realised prices is dispatch regime 1 where option 1 usually serves the reserve requirement and its dispatch is equivalent to the dispatch of a minimum load linear product. The higher the initial forward price is, the more likely is an additional exercise of option 2 on day 6. This exercise decision is taken in order to guarantee an optimal utilisation of the rare in the money hours by both options being online in parallel. In comparison to this most prominent value driver the question which options has to serve reserve requirement is always secondary. Therefore it can be argued that for increased initial forward prices of day 6 option 1 is more and more dependent on the exercise decision of options 2. This is consistent with a slightly increased Gamma of option 2 and the negative Gamma of option 1, which is effectively an indication for the huge loss of flexibility due to the strong dependence on option 2.

6.1.3 Split of original option into three and five parts

Similar to figure 6.2 figure 6.7 shows individual option results for a portfolio of three similar options whose size is one third of the original option as shown in figure 6.1. Again the only difference of the three options is a deviation of the strike prices by 0.01 in order to discriminate the options in the course of reserve allocation. The structure of results for options 1 and 2 is very close to the associated structure in the two option

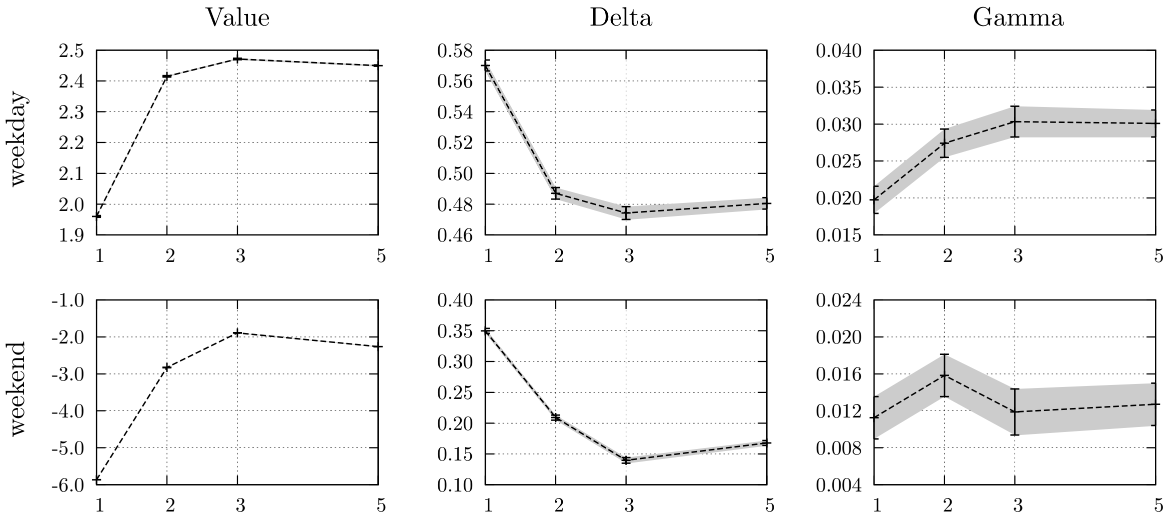

Fig. 6.8: Dependence of average daily portfolio values, Deltas and Gammas (weekdays and weekend) on the number of split options. $n = 3\text{m}$.

portfolio (figure 6.2). Option 3 servers as a wildcard in enabling other options to run at full load whenever option 3 is online in parallel. This is due to the fact that option 3 has the highest strike price and will always serve reserve requirement when being online. Therefore option 3 is characterised by comparably low values and sensitivities. However on a portfolio level the overall portfolio effect of this third option and the general decrease of all option sizes with respect to the two options portfolio lead to an additional increase of the portfolio value, a decrease of Delta and an increase of Gamma (denoting an increase of the portfolio flexibility). This is clearly visible in figure 6.8 which provides aggregated portfolio results for all analysed split portfolios separated by weekdays and weekend days.

Apart from cumulated results of the already discussed portfolios with one, two and three options figure 6.8 contains in addition results for a portfolio with five options each of a size of one fifth of the original option. Obviously this last split of options leads to negative marginal portfolio effects as the overall value decreases, Delta increases (which is an indication for an increased impact of the reserve driven must-run condition) and Gamma slightly decreases at least on weekdays. These effects arise because in contrast to the other portfolios the options are so small that they are no longer able to serve reserve requirement individually. Therefore, in this specific setup at least two options are required to stay online at the same time which obviously reduces the overall flexibility and the value of the portfolio with respect to the setups with three or less options. Associated results for the individual options are provided by figure 6.9. Due to the explained strong link of two options each the result structures of options 1 and 2 are similar to the result structure of option 1 of the three options portfolio while options 3 and 4 of the five options portfolio match option 2 of the three options portfolio. Option 5 is again a wildcard which enables one other option (usually the cheapest option 1) to run at maximum load in case the most expensive option 5 serves part of the reserve requirement.

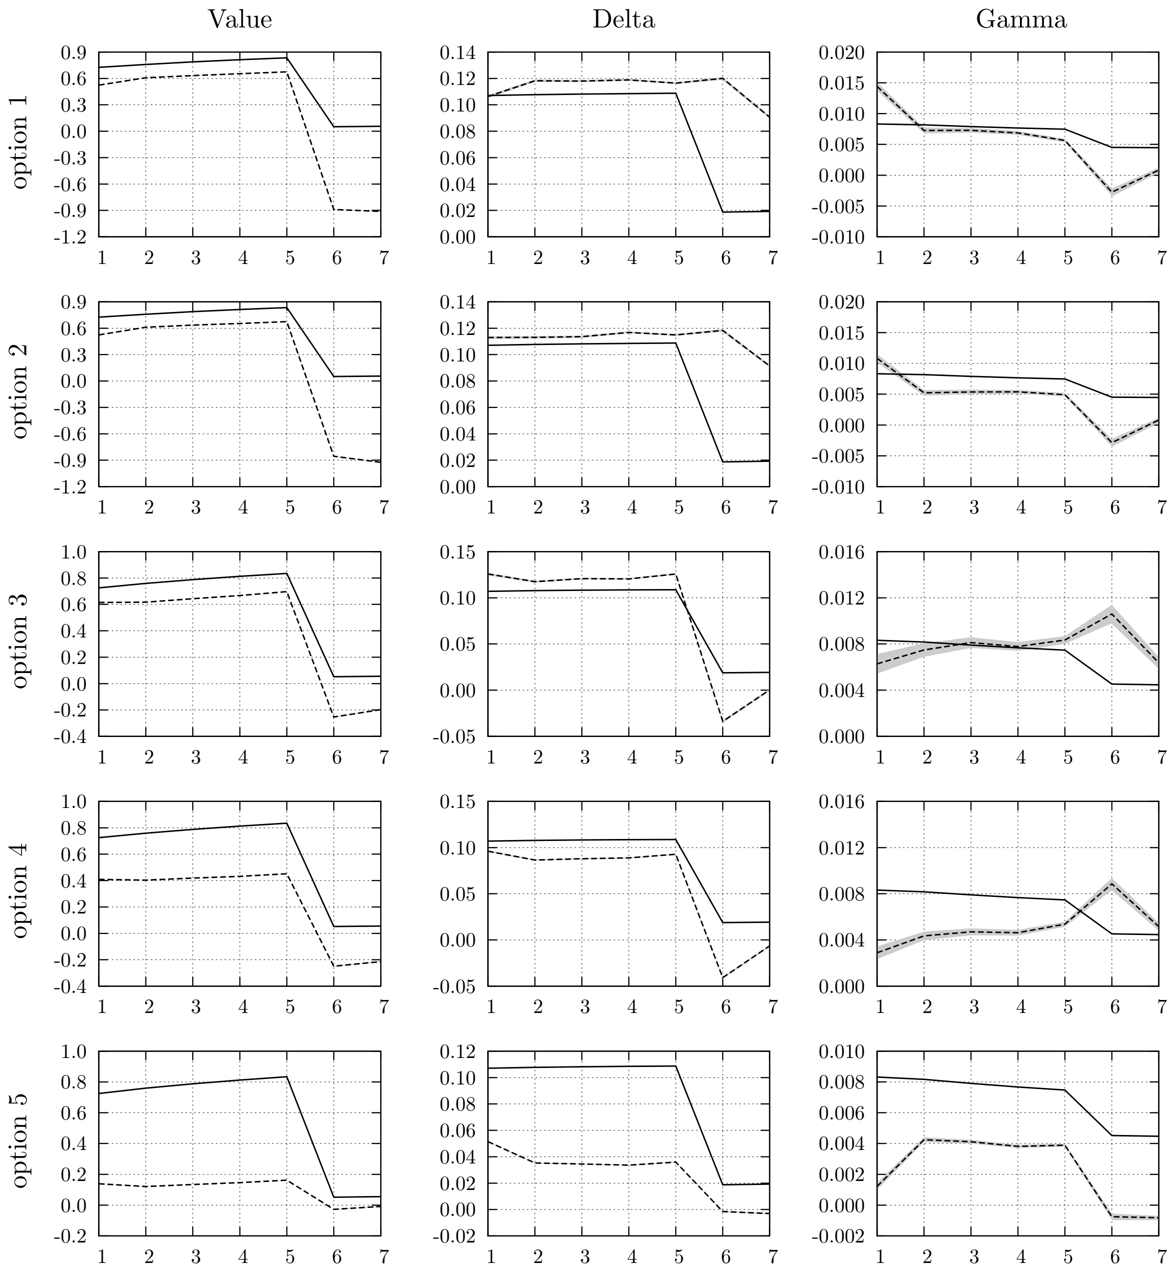

Fig. 6.9: Numerical results for a portfolio of five technically constrained power plant options each having a fifth of the size of the option as shown in figure 6.1 ($t_{up} = 12$, $t_{down} = 8$, $P_{min} = 0.3/5$, $P_{max} = 1.0/5$) subject to a positive reserve requirement of $Res_{pos} = 0.2$ (dashed lines), in comparison to analytical results of fully flexible options (solid lines). $n = 3$m.

Even though the analysed results are very intuitive on the level of the overall portfolio (figure 6.8) these results give no further insights into the complex interdependencies between different options in the portfolio context. The Proxy Simulation Scheme method is an applicable tool to reveal those additional insights and provides a valuable extension to standard risk assessment approaches.

6.2 Portfolios with different options

While section 6.1 provides a detailed overview about portfolio effects and intra option dependencies occurring for portfolios with an increasing number of almost identical options, real physical generation portfolios usually contain power plants of different age and different types and therefore different technical parameters.

In the following sections we consider this fact by deriving values and sensitivities for a range of option portfolios containing at the money and out of the money options with different technical parameters. This highly stylised setup mirrors the current electricity market environment which is characterised by an increasing amount of renewable generation assets and a significant out of the money capacity of old and new gas (and oil) fired power plants. We provide results of various imposed positive reserve requirements for all portfolios in order to explicitly quantify the impact of reserve in the context of different portfolios at the end of this chapter (section 6.7).

In detail we analyse three different option portfolios:

1. *Flexible* option portfolio: Three options with identical load constraints ($P_{min} = 0.3$, $P_{max} = 1.0$) but different strikes and minimum up- and down-times for all options. Option 1 is the most unflexible option ($t_{up} = t_{down} = 8$) which is at the money on weekdays ($\kappa = 50$) and out of the money on weekends. Option 2 is more flexible than option 1 ($t_{up} = t_{down} = 4$) and more expensive to dispatch ($\kappa = 60$). Option 3 is the most expensive option ($\kappa = 70$) but also the most flexible option in terms of up- and down-time constraints ($t_{up} = t_{down} = 2$).

2. *Split flexible* option portfolio: This portfolio is effectively the *flexible* portfolio where all options are split in two parts of identical size. This reduces the minimum and maximum load of all options to $P_{min} = 0.15$ and $P_{max} = 0.5$.

3. *Unflexible* option portfolio: This portfolio is identical to the *flexible* portfolio apart from the fact that all options are significantly less flexible in terms of minimum up- and down-time. Again option 1 is the most unflexible but cheapest option ($t_{up} = 20$, $t_{down} = 12$, $\kappa = 50$), option 2 is more flexible and more expensive ($t_{up} = 18$, $t_{down} = 10$, $\kappa = 60$) while option 3 is the most flexible but most expensive option to dispatch ($t_{up} = 12$, $t_{down} = 8$, $\kappa = 70$).

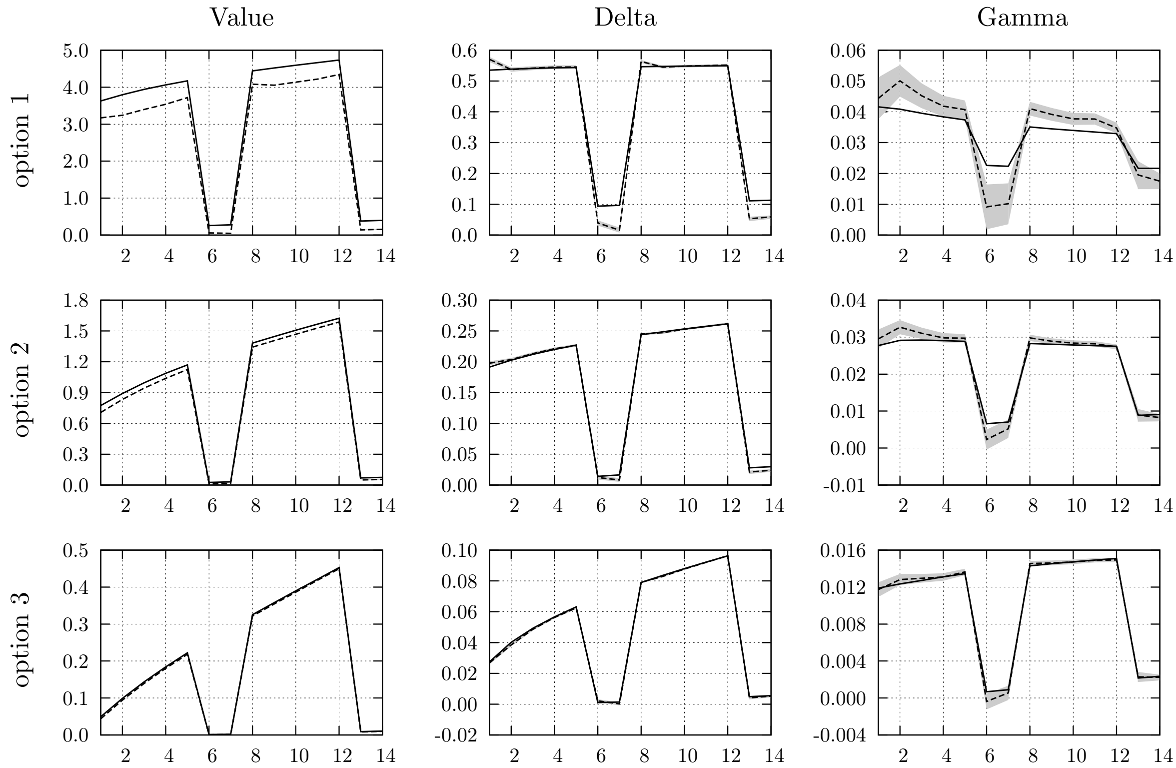

Fig. 6.10: Daily results of all three options of the *flexible* option portfolio (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = t_{down} = 8$, $\kappa = 50$, option 2: $t_{up} = t_{down} = 4$, $\kappa = 60$, option 3: $t_{up} = t_{down} = 2$, $\kappa = 70$. $n = 30$m.

6.3 Flexible option portfolio

6.3.1 Flexible option portfolio without reserve requirement

In order to be able to quantify the impact of reserve on the flexible option portfolio it is mandatory in a first step to derive associated results for the portfolio without any reserve requirements. These values and sensitivities serve as reference results for subsequent analyses with increasing external reserve requirements.

Figure 6.10 shows daily values, Deltas and Gammas of all individual options of the flexible option portfolio for a valuation tenor of 14 days (i.e. two complete weeks according to the daily forwards price structure as provided in figure 3.1). Without any reserve requirements all options can be dispatched independently from each other and the results are similar to single option results as discussed in chapter 5. However it can be observed that numerical results of day 6 and 7 (i.e. the first weekend) are of similar level in case a 14 days valuation tenor is applied. Especially for heavy technical constraints this is different to the 7 days analyses in chapter 5 (e.g. figures 5.1 and 5.5) where results of the last day of the valuation tenor are always increased as an impact of the right hand boundary of the mixed integer optimisation problem. Figure 6.10 shows that while the results of option 1 significantly deviate from analytically results of a fully flexible option, this deviation is lower for option 2 and almost negligible for the

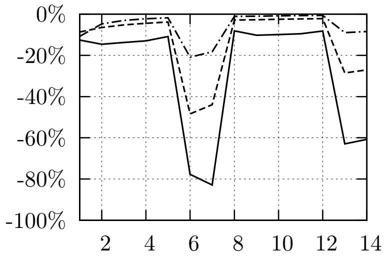

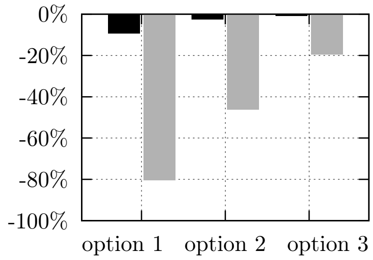

Fig. 6.11: Left: relative drop of daily values of option 1 (solid line), option 2 (dashed line) and option 3 (dashed-dotted line) of the *flexible* option portfolio with respect to analytical results of fully flexible options.
Right: average relative value drop with respect to fully flexible option results for all three options on weekdays within second week (black) and the first weekend (gray). $n = 3\text{m}$.

highly flexible option 3. This is obviously a direct consequence of different technical constraints of the individual options.

Subsequent analyses will often focus on aggregated result figures for weekdays and weekend days (compare figure 6.11). In these cases all five days of the second week will be used in order to generate average weekday figures for value, Delta or Gamma while only days 6 and 7 of the first weekend are used for the calculation of associated weekend figures. This approach guarantees an exclusion of artificial results from either left or right boundary of the mixed integer optimisation problem.

Figure 6.11 shows on the left hand side the relative drop of daily values of all three options of the flexible portfolio with respect to results of fully flexible options with identical strikes. On the right hand side these results are aggregated to associated average results for weekdays and weekend days. In extension to figure 6.10 this visualisation shows that in relative terms not only option 1 is affected by technical constraints but especially on weekends also options 2 and 3 significantly loose value.

Figure 6.12 provides a comprehensive overview about the distribution of the absolute average value drop of the flexible portfolio with respect to associated fully flexible results between all three individual options of the portfolio. The results are again separated by weekdays and weekend days. In absolute terms the overall value drop is significantly higher on weekdays than on weekends while the majority of the effect is generally driven by option 1 as expected. Visualisations like figure 6.12 will be used frequently in the following sections to illustrate the distribution of additional portfolio value changes due to increasing reserve requirements between the individual options.

6.3.2 Flexible option portfolio subject to minor reserve requirement

Figure 6.13 shows daily values, Deltas and Gammas of all individual options of the flexible option portfolio subject to a minor positive reserve requirement of $Res_{pos} = 0.3$.

In comparison to results of the same portfolio without reserve requirements (figure 6.10) a huge value drop can be observed for option 1 which is accompanied by a significant Delta increase and a general Gamma reduction. This effect can be explained

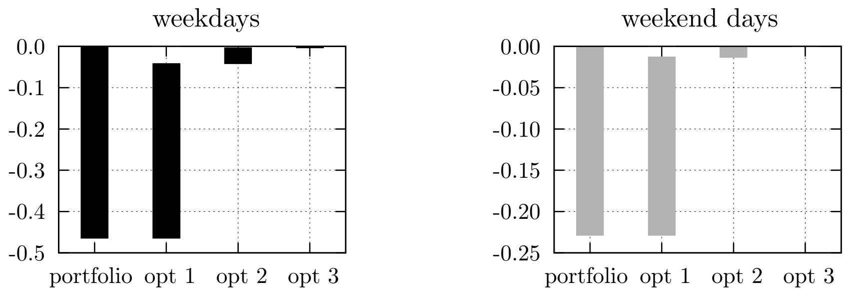

Fig. 6.12: Distribution of the absolute average *flexible* portfolio value drop with respect to fully flexible results between the three individual options (2. week: black, 1. weekend: gray). $n = 3\text{m}$.

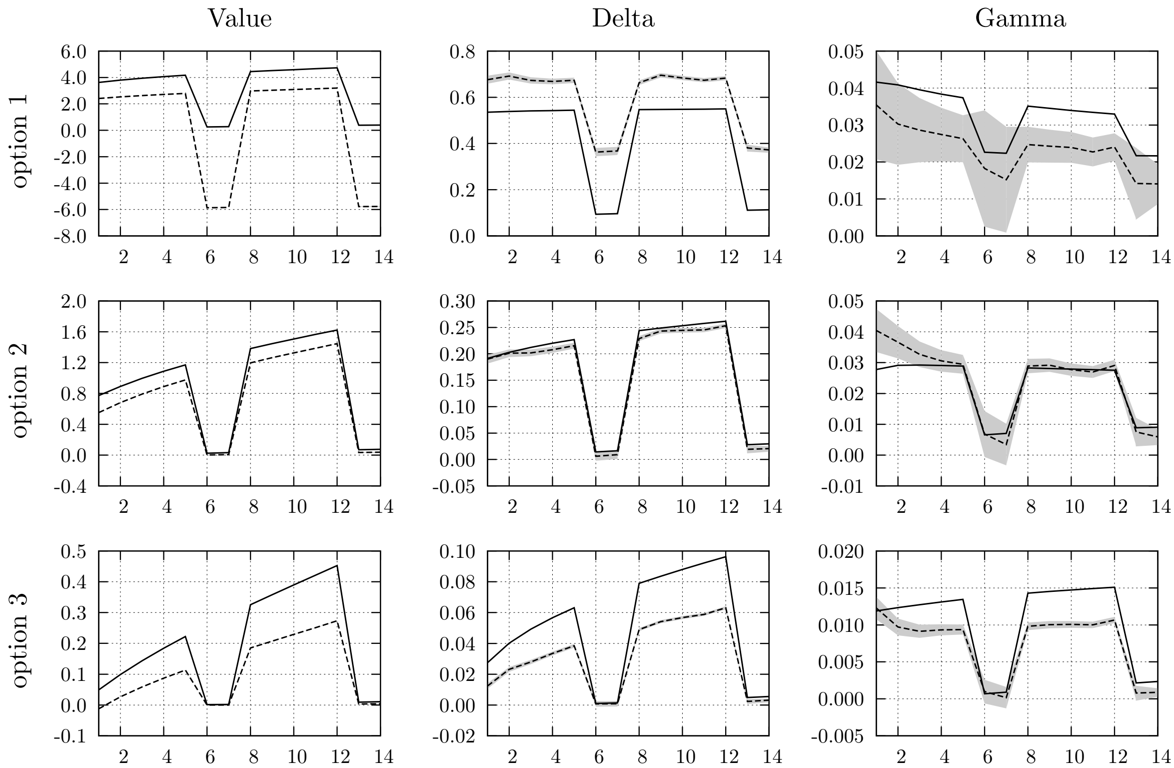

Fig. 6.13: Daily results of all three options of the *flexible* option portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = t_{down} = 8$, $\kappa = 50$, option 2: $t_{up} = t_{down} = 4$, $\kappa = 60$, option 3: $t_{up} = t_{down} = 2$, $\kappa = 70$. $n = 3\text{m}$.

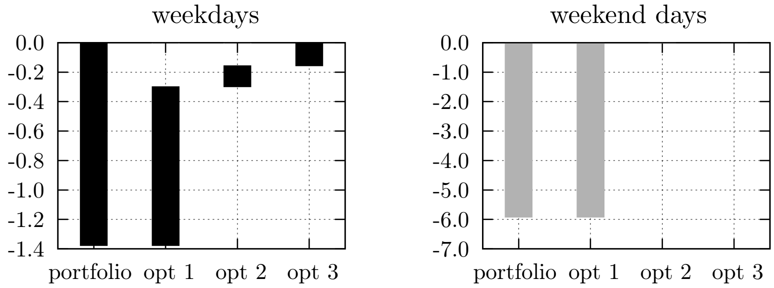

Fig. 6.14: Distribution of the absolute average value drop of the *flexible* portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ with respect to results of the *flexible* portfolio without reserve requirements (2. week: black, 1. weekend: gray). $n = 3$m.

by the observation that running option 1 at minimum load is the most cost effective way to serve the reserve requirement in hours with comparably low prices. This especially includes night hours and midday hours on weekdays enforcing option 1 to run almost everywhere between minimum and maximum load on these days. Thereby a huge share of the reserve-imposed must-run condition is obviously allocated to option 1 which explains the value drop, the increased Delta and the decreased Gamma being consistent with the overall loss of load flexibility. The effect is even more pronounced on weekends because options 2 and 3 are so deep out of the money on weekend days that they will almost never be dispatched and option 1 is forced to serve reserve requirement for almost all weekend hours.

While option 1 typically runs at full load in high priced shoulder hours in order to generate a maximum value contribution option 2 is the best alternative to serve the reserve requirement in these hours. This results in a value reduction especially on weekdays in combination with a very minor reduction of Delta and Gamma.

When option 3 is dispatched it always serves the reserve requirement because it is more economical on a portfolio level to run options 1 and 2 at maximum load in these high priced hours. Therefore option 3 effectively reacts on dispatch decisions of the other two options which limits its flexibility especially on weekdays while it is rather unaffected on weekend days. This intra-option dependency leads to the observed value reduction of option 3 on weekdays which comes along with a significant Delta decrease due to a reduced effective maximum load and a Gamma reduction indicating the generally limited flexibility.

Similar to figure 6.12 figure 6.14 provides an overview about the distribution of the absolute average value drop of the flexible portfolio subject to a minor reserve requirement with respect to associated results of the flexible portfolio without reserve requirements (figures 6.10 and 6.12). Thereby figure 6.14 effectively describes only the distribution of the additional reserve-induced value drop between the three involved options. As expected from the above explanation of figure 6.13 option 1 takes the entire value loss on weekends. This is equivalent to the conclusion that options 2 and 3 do not serve reserve requirement on weekends at all and reflects the negligible impact on weekend results of options 2 and 3 as shown in figure 6.13.

According to figure 6.14 option 1 also serves the majority of reserve requirements on weekdays but options 2 and 3 are significantly affected in addition being consistent

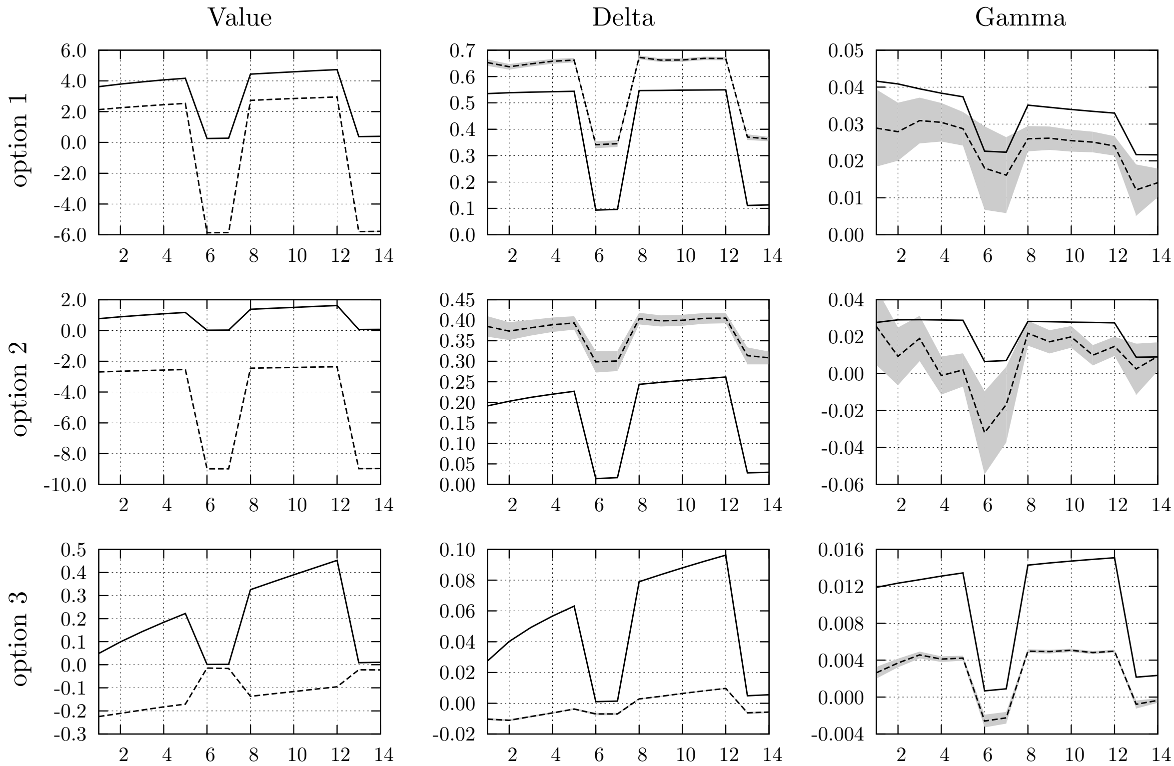

Fig. 6.15: Daily results of all three options of the *flexible* option portfolio subject to a heavy reserve requirement of $Res_{pos} = 1.0$ (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = t_{down} = 8$, $\kappa = 50$, option 2: $t_{up} = t_{down} = 4$, $\kappa = 60$, option 3: $t_{up} = t_{down} = 2$, $\kappa = 70$. $n = 6$m.

with the argumentation above.

6.3.3 Flexible option portfolio subject to heavy reserve requirement

Figure 6.15 shows daily results of individual options of the flexible option portfolio subject to a heavy positive reserve requirement of $Res_{pos} = 1.0$ and figure 6.16 provides the associated portfolio value drop with respect to the same portfolio without reserve requirements and its distribution between all three involved options.

In this setup at least two options are needed in parallel for serving the reserve requirement which is the reason for the huge additional portfolio value drop with respect to the portfolio results without reserve requirements as discussed in section 6.3.1. By comparing figures 6.13 and 6.15 it can be observed that while the results of option 1 are almost not affected by the increased reserve requirement of $Res_{pos} = 1.0$ the results of options 2 and 3 are significantly impacted. Option 2 shows negative values for all days in combination with a highly increased Delta and a generally reduced Gamma being very similar to the effect on option 1 when the reserve requirement is increased from $Res_{pos} = 0.0$ to $Res_{pos} = 0.3$ (compare section 6.3.2). Option 3 shows

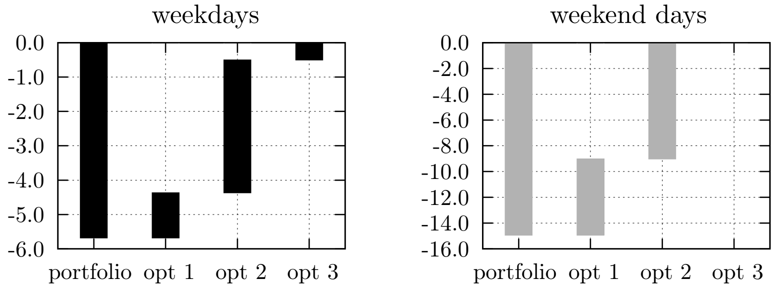

Fig. 6.16: Distribution of the absolute average value drop of the *flexible* portfolio subject to a heavy reserve requirement of $Res_{pos} = 1.0$ with respect to results of the *flexible* portfolio without reserve requirements (2. week: black, 1. weekend: gray). $n = 6$m.

negative values on all days of the valuation tenor which are accompanied by a significant Delta reduction that leads to negative Deltas during the first week and on weekends. Even more remarkable is the negative impact on Gamma of option 3 leading to stable negative Gammas on weekend days (comparable to results of the portfolio analysis with identical options e.g. figures 6.7 and 6.9).

Similar to the setup with minor reserve requirement also in case of heavy reserve requirement option 1 is usually the most cost effective possibility to serve at least a part of the overall positive reserve requirement. This is especially true for weekend days where all options are deep out of the money almost all hours. This leads to rather unchanged results of option 1 in both setups.

Option 2 is the second best alternative for serving the remaining reserve requirement in parallel to option 1. Thereby option 1 usually serves 0.7 of positive reserve requirement in hours with comparably low prices and option 2 is responsible for the remaining 0.3. This imposes a similar must-run condition on option 2 in this setup as valid for option 1 in the setup with minor reserve requirement (compare figure 6.13) which explains the huge drop of all results of option 2. The absolute amount of this drop is larger than observed for option 1 in the minor reserve scenario due to the higher strike price of option 2 leading to higher negative value contributions in hours where the options is subject to the must-run condition.

When option 3 is started in comparably high priced hours it always serves reserve requirement because in the portfolio context it is more economical to increase the load of option 1 to maximum load at the same time. Due to the fact that option 1 is not able to shift all allocated reserve requirement to option 2 (as possible for minor reserve requirement) option 3 is more frequently dispatched in the heavy reserve setup than in the minor reserve setup. This explains the increased absolute value drop of option 3 on weekdays. The shifting of reserve requirement from option 1 to option 3 is more likely to happen for slightly upshifted initial forward prices – meaning that the probability of option 3 to generate negative payoffs is increased for upshifted forward prices. This is consistent with the slightly negative Delta of option 3 as observed in figure 6.15. On weekends this situation is even worse as option 3 is never in the money but only reacts on dispatch decisions of the other two options. Thereby it fully resembles a short option position in the portfolio which is consistent with the negative Deltas and Gammas on weekend days. For heavy reserve requirements all option dispatches follow

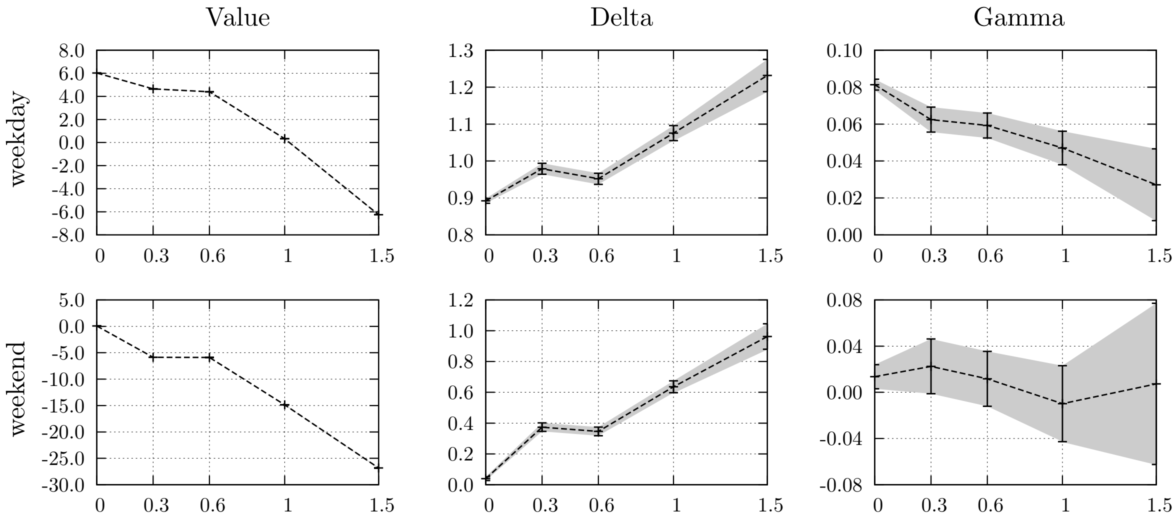

Fig. 6.17: Dependency of average daily results of the *flexible* portfolio on the amount of external positive reserve requirement (Top: weekdays, Bottom: weekend days). Gray area and error bars denote 95% confidence intervals.

a clear pecking order where options 1 and 2 serve reserve requirement for comparably low prices and option 3 is exercised by option 1 whenever the positive value contribution of option 1 running at full load overcompensates the negative value contributions of option 3 serving parts of the reserve requirement.

6.3.4 Portfolio level impact of different reserve requirements on flexible option portfolio

Figure 6.17 provides a comprehensive overview about the dependency of the flexible option portfolio on the amount of applied positive reserve requirement. In addition to already discussed results (sections 6.3.1, 6.3.2 and 6.3.3) the visualisation also contains portfolio results for intermediate and very high positive reserve requirements of $Res_{pos} = 0.6$ and $Res_{pos} = 1.5$. All numerical results are separated by weekdays and weekend days. Figure 6.17 shows 95% confidence intervals which base on sample standard deviations (according to equation (2.6)) of daily results of all individual options. Hereby we follow a conservative approach and assume 100% correlation between results of different days as well as results of individual options.

On a portfolio level all results are relatively intuitive. An increased amount of positive reserve requirement leads to a decreased portfolio value, an increased Delta due to the stricter must-run condition and a decreased Gamma reflecting the increasing loss of load flexibility of the portfolio. The comparably minor result differences between 0.3 and 0.6 reserve requirement are due to the favorable size of the individual options, allowing one option in both cases to serve the entire reserve requirement. This can be an interesting result for practitioners who manage a real portfolio of power plants subject to expected levels of reserve requirements when it comes to the question of sizing new plant investments or decommissionings of old plants of a certain size.

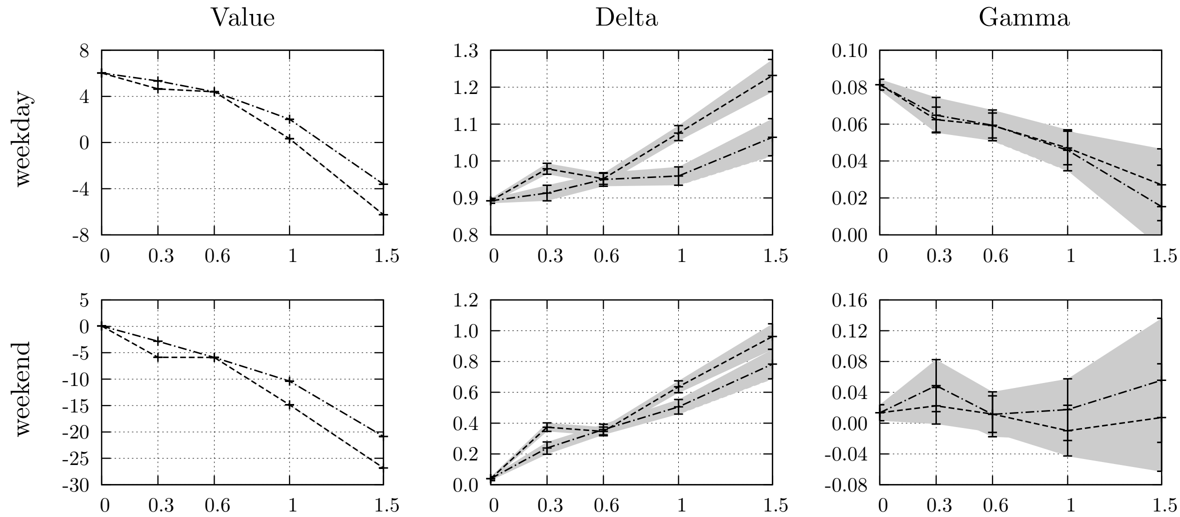

Fig. 6.18: Dependency of average daily results of the *split flexible* portfolio (dashed-dotted lines) on the level of external positive reserve requirement (Top: weekdays, Bottom: weekend days) and comparison to associated results of the *flexible* portfolio (dashed lines). Gray area and error bars denote 95% confidence intervals.

Similar to the final comment in section 6.1 it is again valuable to stress that all intra-option dependencies and coupling effects are only visible in case sensitivities for individual options are made available by an approach like the Proxy Simulation Scheme method. Knowing values and sensitivities only on a portfolio level (according to figure 6.17) is obviously not sufficient to understand all interdependencies and therefore the entire complexity of the underlying portfolio.

6.4 Favorable portfolio effect of split options

Figure 6.18 compares portfolio results of the flexible portfolio (figure 6.17) with associated results of the split flexible portfolio. As explained in section 6.2 the split portfolio is effectively the flexible portfolio where all three options are cut in two parts, each of half the size of the original option. It can be expected that this portfolio setup serves reserve requirement at lower costs due to the ability to allocate it to a larger number of smaller options in order to match the implied must-run condition more cost efficiently than possible with the larger options of the flexible portfolio. This is confirmed by figure 6.18 which clearly shows that the overall portfolio value of the split portfolio is significantly higher for all amounts of reserve requirements apart from 0.0 and 0.6 where both portfolios yield exactly the same average results.

Each option of the split portfolio is able to serve a maximum of 0.35 positive reserve requirement running at minimum load of 0.15, while the options of the flexible portfolio are able to serve 0.7 positive reserve requirement running at minimum load of 0.3. As long as the portfolio is optimised without any reserve requirement all options are dispatched individually from each other and only based on price path realisations. This leads to identical daily values in both portfolios as the two split options with identical constraints will exactly replicate the dispatch of the associated option of the flexible portfolios – resulting in identical portfolio dispatches for identical price realisations.

In case of a minor reserve requirement of $Res_{pos} = 0.3$ this is no longer true. While in the flexible portfolio one large option is required to run at least on minimum load to serve the reserve requirement, in the split portfolio only a single smaller option is sufficient. This significantly reduces the negative impact of the must-run condition as the split option generates only half the amount of negative payoffs in out of the money hours due to its reduced minimum load of $P_{min} = 0.15$. This effect is also the reason for the reduced portfolio Delta with respect to the flexible portfolio and a slightly increased Gamma, the latter indicating that in the split portfolio five remaining options can be dispatched in a flexible way while only two options of the flexible portfolio are not directly affected by the reserve requirement.

This favorable effect vanishes for an increased reserve requirement of $Res_{pos} = 0.6$ which can only be served by two options of the split portfolio in parallel. The combined minimum load of these options is exactly the same as the minimum load of one larger option of the flexible portfolio leading to exactly the same must-run conditions on both portfolios. Therefore no portfolio is more cost effective in serving reserve requirement and both portfolios show identical numerical results for $Res_{pos} = 0.6$.

In contrast the impact of $Res_{pos} = 1.0$ is again similar to $Res_{pos} = 0.3$. Now two options of the flexible portfolio are required to serve reserve requirement and run at least on a combined minimum load of 0.6 in all hours while three smaller options of the split portfolio are sufficient with a combined minimum load of 0.45. This leads to the huge value increase, Delta reduction and Gamma increase with respect to results of the flexible option portfolio in figure 6.18.

The favorable portfolio effect is even more pronounced for a very high reserve requirement of $Res_{pos} = 1.5$ where all options of the flexible portfolio are subject to the must-run condition while one option of the split portfolio can still be dispatched only based on the realised price structure.

6.5 Unflexible option portfolio

6.5.1 Unflexible option portfolio without reserve requirement

Similar to section 6.3.1 a structured analysis of the impact of positive reserve requirement on results of the unflexible portfolio requires a set of reference values of the same portfolio without any reserve requirements. Figure 6.19 provides associated results for all three individual options of the unflexible portfolio (please refer to section 6.2 for technical specifications).

All options show a significant value drop with respect to fully flexible results due to the heavy imposed technical constraints. Delta and Gamma are generally reduced on weekend days where large minimum up-time conditions in combination with rare in the money hours make a start of the options usually uneconomical. On weekdays, daily Delta and Gamma results show the typical spike structure due to technical constraints leading to dispatch flexibility shifts between consecutive days as explained in sections

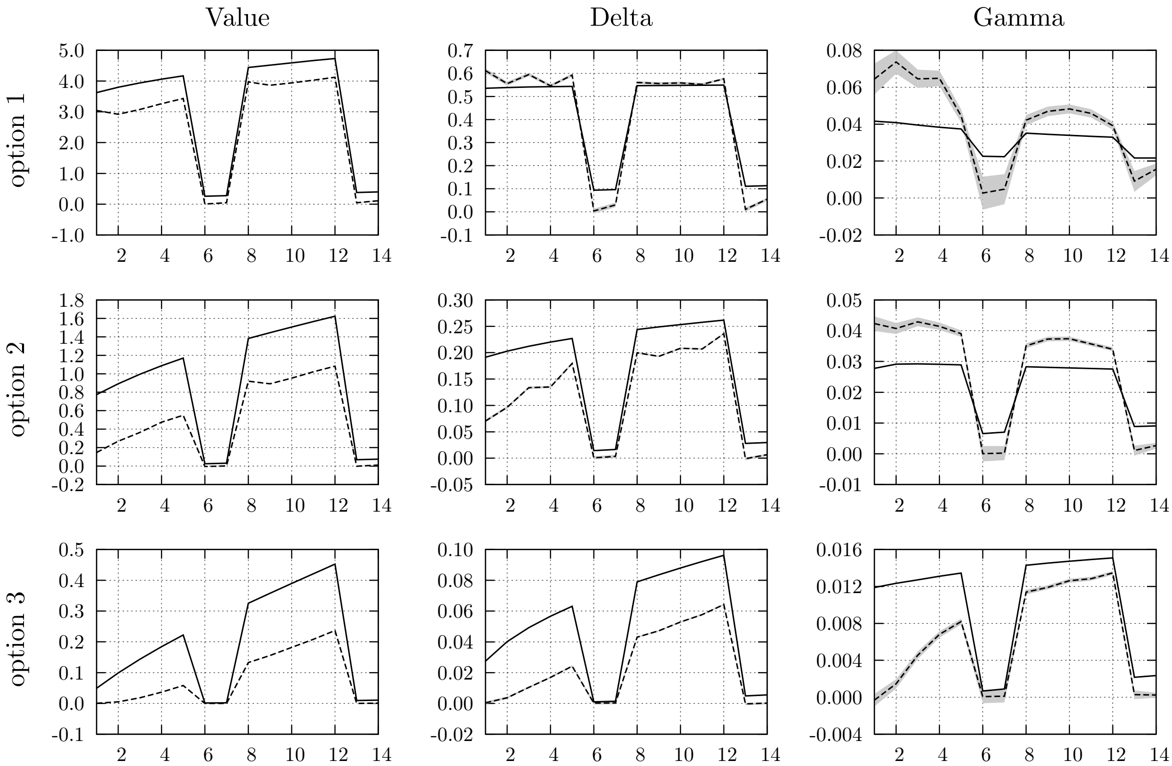

Fig. 6.19: Daily results of all three options of the *unflexible* option portfolio (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = 20$, $t_{down} = 12$, $\kappa = 50$, option 2: $t_{up} = 18$, $t_{down} = 10$, $\kappa = 60$, option 3: $t_{up} = 12$, $t_{down} = 8$, $\kappa = 70$. $n = 3$m.

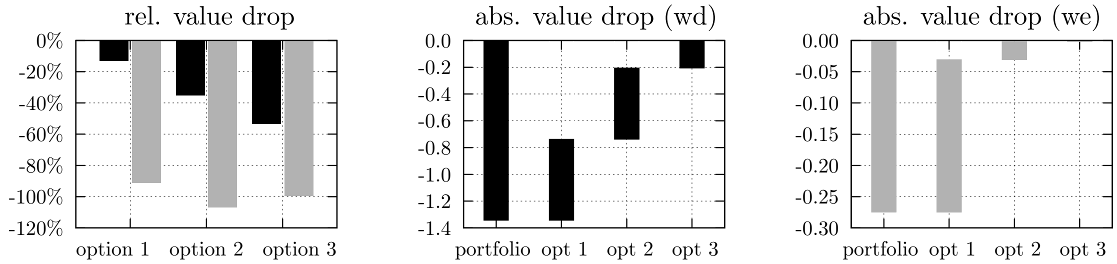

Fig. 6.20: Left: average relative value drop with respect to fully flexible results for all three options of the *unflexible* portfolio on weekdays (black) and weekend days (gray).
Middle and right: distribution of the average absolute value drop of the *unflexible* portfolio with respect to fully flexible analytical results between all three involved options (weekdays: black, weekend days: gray). $n = 3$m.

5.1.1 and 5.1.2. Furthermore options 1 and 2 show the characteristically increased Gamma driven by longer blocks of connected online hours that are on average closer at the money than the shorter blocks of consecutive in the money hours (compare figures 5.4 and 5.7). In contrast to these observations option 3 shows a generally reduced Gamma for all weekdays with very low values especially during the first days of the valuation tenor. A similar impact can be observed for daily Delta results of options 2 and 3 on weekdays where Delta is starting at very low levels (close to zero in case of option 3) and then slightly growing as time to maturity increases. These very low values at the beginning of the valuation tenor are an effect of the technical constraints of options 2 and 3 in combination with their systematically higher strike price. As both options are out of the money for the majority of hours also on weekdays, their large minimum up- and down-time constraints prevent these options from being dispatched at all until evolutions of the price forward curve lead to in the money shifts of individual days. Due to the finite volatility of the daily forward prices such shifts do typically not occur during the very first days of the valuation tenor but have a certain probability at later days. Therefore all numerical results of options 2 and 3 start from a very low level (close to zero) at day 1 and tend to increase in parallel with an increased dispatch probability. This effect is very pronounced for option 3 and also clearly visible in case of option 2.

According to figures 6.11 and 6.12 figure 6.20 provides average relative value drops of all options of the unflexible portfolio with respect to fully flexible results separated by weekdays and weekend days, and in addition the distribution of associated absolute portfolio value drops between all three options. As explained above the relative value drop is close to 100% for all options on weekend days. Interestingly for weekdays it is increasing from option 1 to option 3 and is therefore increasing in parallel to decreasing technical constraints. This is an effect of the increased strike price of options 2 and 3 which pushes these options deeper out of the money. The deeper out of the money an option is the lower is its probability to being dispatched in case of strong technical constraints. This effect overcompensates the slight reduction of minimum up- and down-times from option 1 to option 3 and finally leads to the observed increase of the relative value drop for options with higher strike prices. A decreased moneyness of an option effectively amplifies the relative value reducing impact of technical constraints.

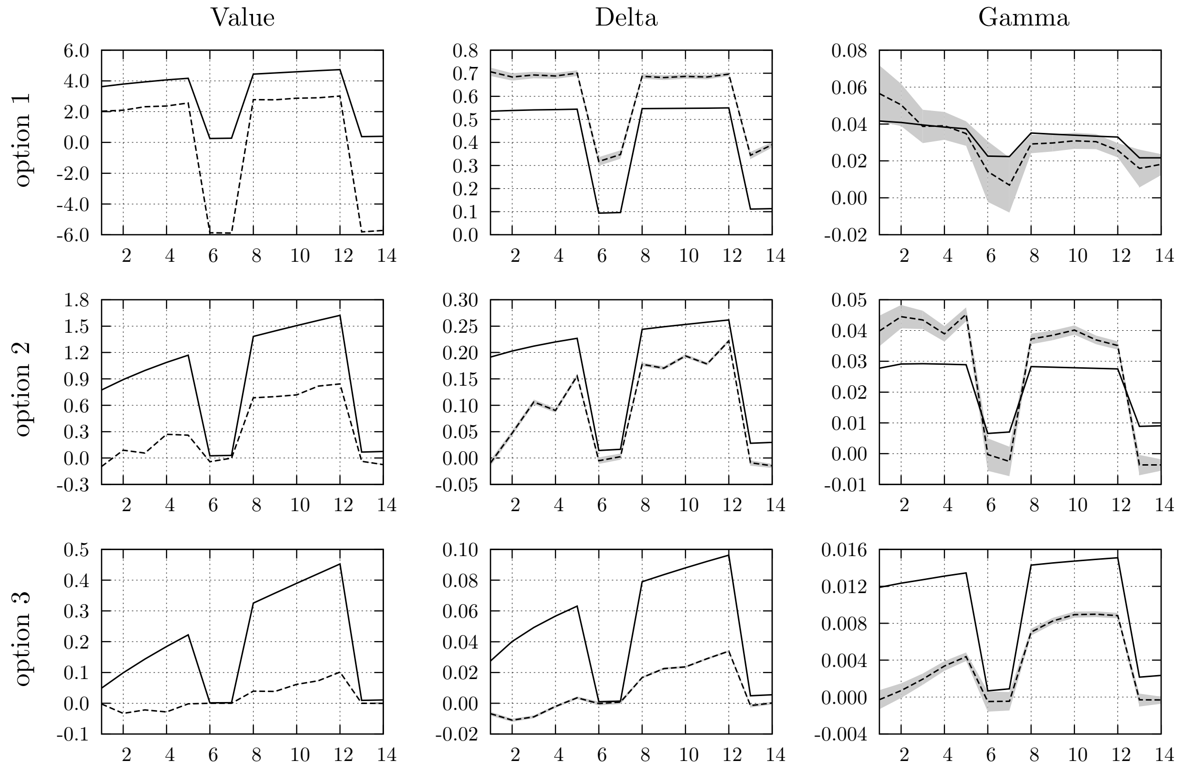

Fig. 6.21: Daily results of all three options of the *unflexible* option portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = 20$, $t_{down} = 12$, $\kappa = 50$, option 2: $t_{up} = 18$, $t_{down} = 10$, $\kappa = 60$, option 3: $t_{up} = 12$, $t_{down} = 8$, $\kappa = 70$. $n = 3$m.

In contrast to this relative view the absolute drop of the unflexible portfolio value is significantly higher on weekdays and the option values are more reduced for stricter technical constraints. A comparison of figures 6.20 and 6.12 makes clear that while only option 1 of the flexible portfolio is significantly impacted by technical constraints all options of the unflexible portfolio are similarly strongly affected.

6.5.2 Unflexible option portfolio subject to minor reserve requirement

Figure 6.21 shows the daily results of all three power plant options of the unflexible portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ in comparison to analytical results of fully flexible options with identical strike prices.

By comparison to figure 6.13 it becomes obvious that the structural impact of minor reserve requirement on the unflexible portfolio is very similar to the same impact on the flexible portfolio. Option 1 is affected by a huge value drop on all days of the valuation tenor in combination with a significant Delta increase and a Gamma reduction. Again this is a result of option 1 being the most cost effective portfolio solution to serve the reserve requirement leading to option 1 running almost every

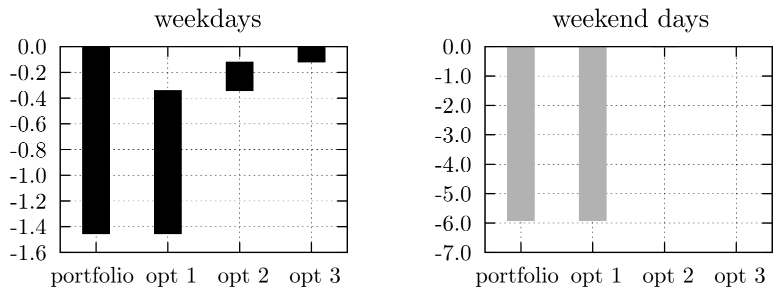

Fig. 6.22: Distribution of the average absolute value drop of the *unflexible* portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ with respect to results of the *unflexible* portfolio without reserve requirements (weekdays: black, weekend days: gray). $n = 3$m.

hour (compare section 6.3.2). The structure of daily results of option 2 is close to results without reserve requirements (figure 6.19) apart from a more pronounced spike structure and a slight value drop. This is due to the fact that option 2 almost always serves reserve requirement when it is running in parallel to option 1 (i.e. in comparably high priced hours) as option 3 is now very rarely started and only then able to take over the reserve requirement. However when option 3 is running it always serves the reserve requirement being the reason for the observed reduction of all results of option 3 even leading to a slightly negative Delta during the first days of the valuation tenor.

Figure 6.22 shows the distribution of the average value drop of the unflexible portfolio due to an increase of positive reserve requirement from 0.0 to 0.3 and its distribution between the three involved power plant options. As expected the distribution for weekdays is close to associated observations for the flexible portfolio (compare figure 6.14) and on weekend days both the overall portfolio value drop and its allocation only to option 1 are even identical for the two portfolios. These identical results arise because in both portfolios options 2 and 3 will almost never run on weekend days due to applied technical constraints in combination with comparably high strike prices. This makes option 1 always the natural choice to serve reserve requirements on weekend days – and both portfolios are effectively reduced to single option portfolios on weekend days eliminating any dependency of the overall value on technical constraints as explained in section 5.2.1. Consequently both portfolio values must exactly match on weekend days.

6.5.3 Unflexible option portfolio subject to heavy reserve requirement

Similar to the impact of minor reserve requirement also the general impact of heavy reserve requirement is similar in the unflexible and the flexible portfolio. Figure 6.23 shows daily results of individual options of the unflexible portfolio subject to $Res_{pos} = 1.0$ which can directly be compared to associated results of the flexible portfolio (figure 6.15). In this setup at least two options are needed to serve the reserve requirement leading to a significant additional value drop, Delta increase and Gamma decrease of option 2 with respect to results of the unflexible portfolio subject to a minor reserve requirement (figure 6.21). As explained for the flexible portfolio (section 6.3.3) option

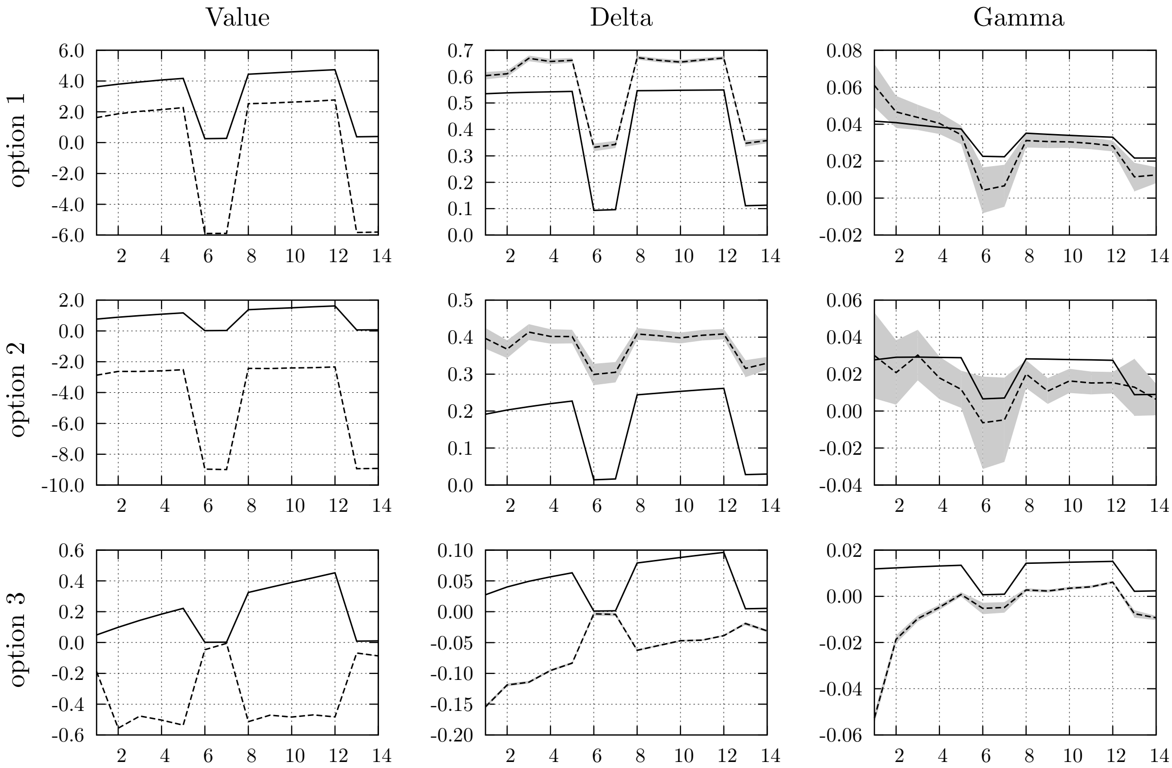

Fig. 6.23: Daily results of all three options of the *unflexible* option portfolio subject to a heavy reserve requirement of $Res_{pos} = 1.0$ (dashed lines) in comparison to analytical results of fully flexible options (solid lines). $P_{min} = 0.3$ and $P_{max} = 1.0$ identical for all options. Option 1: $t_{up} = 20$, $t_{down} = 12$, $\kappa = 50$, option 2: $t_{up} = 18$, $t_{down} = 10$, $\kappa = 60$, option 3: $t_{up} = 12$, $t_{down} = 8$, $\kappa = 70$. $n = 5$m.

3 acts as a call option in the portfolio context which can be exercised by option 1 whenever it is more economically to run option 1 at maximum load and shift reserve allocation to option 3. This leads to negative values, Deltas and partly negative Gamma results for option 3. This effect is more pronounced for the unflexible portfolio than for the flexible portfolio as stricter technical constraints of the unflexible portfolio lead to a higher effective strike price of option 3 pushing this option further out of the money – and thereby amplifying the negative impact on all results. Especially on the first days of the valuation tenor option 3 shows very low Gamma results. This is due to the fact that option 3 will almost never be dispatched when it is in the money as the applied price volatility makes such required high price realisations very unlikely. Therefore option 3 is only dependent on dispatch decisions of the other two options and has no remaining flexibility which is indicated by the negative Gamma results.

Figure 6.24 provides the associated distribution of the portfolio value drop due to the increase of Res_{pos} from 0.0 to 1.0 between all three options of the unflexible portfolio. As expected also these results are very close to corresponding results of the flexible portfolio as provided in figure 6.16.

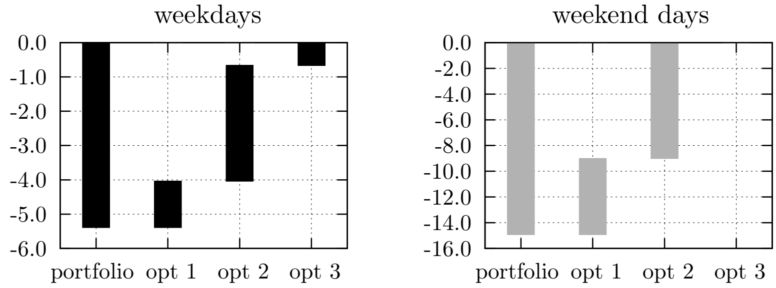

Fig. 6.24: Distribution of the average absolute value drop of the *unflexible* portfolio subject to a heavy reserve requirement of $Res_{pos} = 1.0$ with respect to results of the *unflexible* portfolio without reserve requirements (weekdays: black, weekend days: gray). $n = 5$m.

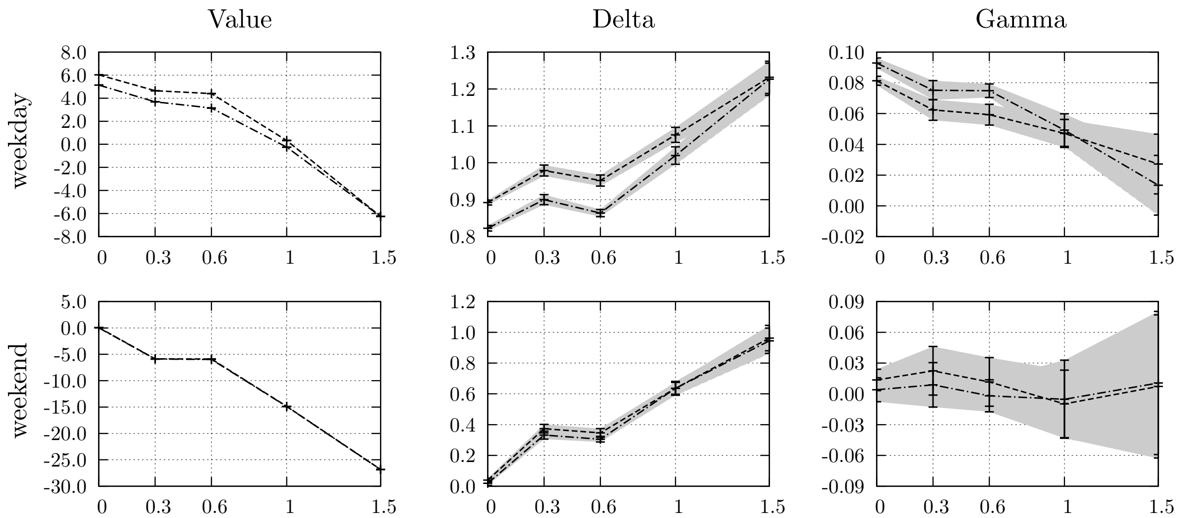

Fig. 6.25: Dependency of average daily value, Delta and Gamma of both the *flexible* portfolio (dashed lines) and the *unflexible* portfolio (dashed-dotted lines) on the level of positive reserve requirement (Top: weekdays, Bottom: weekend days).

6.6 Comparison of portfolio value impacts of reserve requirements

Sections 6.3 and 6.5 provide detailed discussions of the impact of positive reserve requirement on individual options of either the flexible or the unflexible option portfolio. This section complements these analyses by a direct comparison of the reserve dependency of cumulated results of both portfolios in order to identify all general response differences of the two portfolios with respect to applied reserve requirements.

Figure 6.25 shows average daily results separated by weekdays and weekend days of both the flexible and the unflexible portfolio subject to different levels of applied positive reserve requirement between 0.0 and 1.5. As expected for $Res_{pos} = 0.0$ the value of the flexible portfolio exceeds the value of the unflexible portfolio on weekdays due to a generally lower impact of technical constraints on the flexible portfolio. On weekends both portfolio values are almost identical which indicates, that both portfolios are deep out of the money and try to serve reserve requirement as cost effective as possible - which implies using the same strike price driven order of option dispatch. Both observations are consistent with the more detailed analyses of sections 6.3 and 6.5.

For a high reserve requirement of $Res_{pos} = 1.5$ the portfolio values also converge on weekdays due to the fact that in both portfolios all options are needed in parallel to serve the applied level of reserve requirements. According to section 5.2.1 the impact of minimum up- and down-time constraints vanishes once the entire portfolio is forced to run at least at minimum load in all hours. This is the case for $Res_{pos} = 1.5$ consequently leading to the observed identical portfolio values.

Consistently with sections 6.3, 6.5 and the argumentation above Delta is lower for $Res_{pos} = 0.0$ in the unflexible portfolio than in the flexible portfolio while Gamma is higher. Both sensitivities converge for the high reserve requirement in line with the observed convergence of the portfolio values. Weekend day sensitivities are very close in both portfolios which is obviously due to almost identical dispatch patterns as explained above.

Remarkably even though starting from a lower value without reserve requirements on weekdays, the unflexible portfolio looses more value than the flexible portfolio when the reserve requirement is increased from 0.0 to 0.6. This changes for a further increase towards 1.0 and 1.5 where the unflexible portfolio looses significantly less value than the flexible portfolio until both portfolio values finally converge at $Res_{pos} = 1.5$. Figure 6.26 directly shows the absolute value drops of both portfolios on the left hand side, in the middle the same value drops normalised by the applied level of reserve requirement and on the right hand side the differences between absolute and relative value drops of both portfolios. Especially the comparison of normalised value drops clearly demonstrates that the value decrease is higher for the unflexible portfolio until $Res_{pos} = 0.6$ and lower for higher reserve requirements. This effect can be explained by using the concept of online ratios as discussed in the following paragraphs.

Online ratios Whenever a reserve requirement leads to additional runtime of options which otherwise would have been offline this is typically more expensive in the unflexible portfolio than in the flexible portfolio due to longer minimum up-times of all options of the unflexible portfolio. This can systematically be quantified via an analysis of online ratios especially on weekday hours with comparably low prices, i.e. with hourly curve adjustment factors below 1.0. An online ratio denotes the share of hours with hca factor < 1.0 when the power plant option is running of the total amount of hours with hca factor < 1.0. Especially these hours will typically lead to reduced value contributions of the respective power plant options. Figure 6.27 provides online ratios for all options of both flexible and unflexible portfolio subject to different positive reserve requirements between 0.0 and 1.0. In addition the visualisation shows incremental deltas of these online ratios when Res_{pos} is varied between two discrete values. Table 6.2 provides the same data in numerical form.

Starting from zero reserve requirement and increasing Res_{pos} to 0.3, the weekday online ratio of option 1 increases from 41.5% to 100% in the flexible portfolio, the online ratio of option 2 increases from 1.9% to 5.4% and the online ratio of option 3 from 0.1%

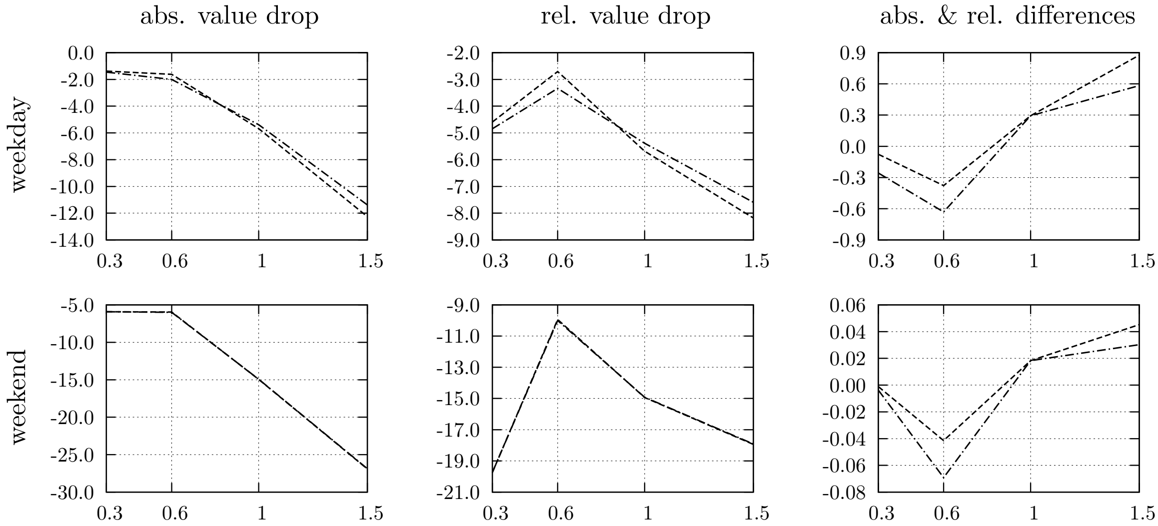

Fig. 6.26: Left: absolute difference between average daily values of both the *flexible* (dashed lines) and the *unflexible* (dashed-dotted lines) portfolio subject to different reserve requirements with respect to the same portfolios without any reserve requirement,
Middle: left graph, normalised by the applied level of reserve requirement,
Right: differences between reserve impact on both portfolios (dashed lines: difference of absolute value impact, dashed-dotted lines: difference of normalised value impact).

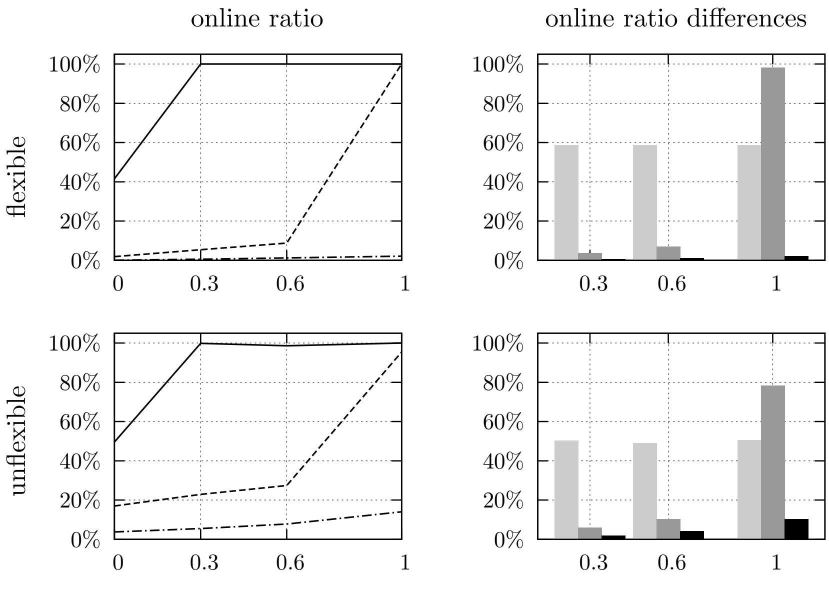

Fig. 6.27: Left: dependency of weekday online ratios of all options (solid lines: option 1, dashed lines: option 2, dashed-dotted lines: option 3) of the *flexible* (Top) and the *unflexible* portfolio (Bottom),
Right: incremental change of the online ratios of all options with respect to the case without reserve requirements in percent points when the level of reserve requirement is increased (light gray: option 1, dark gray: option 2, black: option 3).

to 0.6%. In the unflexible portfolio the online ratio of option 1 increases from 49.6% to 99.8%, the one of option 2 from 17.0% to 22.9% and the one of option 3 from 3.8% to 5.5%. The unflexible portfolio obviously starts with a higher online ratio of rather low priced hours, leading to the comparably lower initial portfolio value only due to higher technical constraints of all options of the unflexible portfolio. Although the cumulative increase of online ratios is higher in the flexible portfolio the increase in the unflexible portfolio is more significant for option 2 and 3 - which leads to higher additional cost due to their higher strike prices. This is the explanation for the relatively higher value decrease of the unflexible portfolio when the reserve requirement is increased from 0.0 to 0.3.

Increasing Res_{pos} further from 0.3 to 0.6, in the flexible portfolio the online ratio of option 1 stays at 100%, the one of option 2 increases from 5.4% to 8.8% and the one of option 3 increases from 0.6% to 1.2%. In the unflexible portfolio the online ratio of option 1 is only marginally changing from 99.8% to 98.6%, the online ratio of option 2 increases from 22.9% to 27.4% and the one of option 3 increases from 5.5% to 7.8%. Thus the online ratios of options 2 and 3 increase in total by 6.8 percent points in the unflexible portfolio while only by 4.0 percent points in the flexible portfolio. This again explains the relatively higher value decrease in the unflexible portfolio than in the flexible portfolio when reserve requirement is increased from 0.3 to 0.6.

At $Res_{pos} = 1.0$ two options are always needed to run in parallel in order to serve the reserve requirement. In the flexible portfolio this condition is almost exclusively met by options 1 and 2 in case of low priced hours, both options having online ratios of 100%. Compared to $Res_{pos} = 0.6$ the online ratio of option 2 jumps by 91.2 percent points from 8.8% to 100%, while option 1 stays at 100% and the online ratio of option 3 increases slightly from 1.2% to 2.1%. In the unflexible portfolio the online ratio of option 1 changes only marginally from 98.6% to 100%, the one of option 2 increases heavily from 27.4% to 95.4% and option 3 doubles from 7.8% to 14.0%. However in total this cumulated increase of 75.6 percent points is significantly below the associated increase of 92.1 percent points in the flexible portfolio which is the reason for the relatively smaller value decrease of the unflexible portfolio when the reserve requirement is increased from 0.6 to 1.0.

Figure 6.26 also shows a very small online ratio effect on weekends but the overall value decrease is negligible with respect to the strong impact of the increasing weekday must-run condition on both deep out of the money portfolios.

6.7 Sensitivities of reserve requirements in different option portfolios

Section 6.6 compares portfolio value changes of all analysed option portfolios driven by an increasing amount of applied positive reserve requirement. This approach can easily be extended towards the sensitivities Delta and Gamma and therewith allows to quantify the impact of an increased reserve requirement on the value and risk profile of each portfolio. Furthermore, by comparing results of a portfolio subject to a given

		positive reserve requirement (Res_{pos})			
	result type	0.0	0.3	0.6	1.0
flexible portfolio	value	722.35	557.03 (-165.62)	527.62 (-194.73)	40.37 (-681.98)
	online ratio option 1	41.5%	100.0% (+58.5%)	100.0% (+58.5%)	100.0% (+58.5%)
	online ratio option 2	1.9%	5.4% (+3.5%)	8.8% (+6.9%)	100.0% (+98.1%)
	online ratio option 3	0.1%	0.6% (+0.5%)	1.2% (+1.1%)	2.1% (+2.0%)
unflexible portfolio	value	616.94	442.40 (-174.54)	376.80 (-240.14)	-29.79 (-646.73)
	online ratio option 1	49.6%	99.8% (+50.2%)	98.6% (+49.0%)	100.0% (+50.4%)
	online ratio option 2	17.0%	22.9% (+5.8%)	27.4% (+10.3%)	95.4% (+78.3%)
	online ratio option 3	3.8%	5.5% (+1.8%)	7.8% (+4.1%)	14.0% (+10.2%)

Table 6.2: Overview of weekday values and associated online ratios of individual options of both flexible and unflexible portfolio. Numbers in brackets denote deltas with respect to results without any applied reserve requirements (i.e. $Res_{pos} = 0.0$).

reserve requirement with results of the same portfolio without any reserve requirements it is also possible to indirectly evaluate the entire impact of the reserve requirement itself. The impact is equivalent to the value and risk profile of the respective reserve requirement in the context of the evaluated power plant option portfolio. This indirect method is used in this thesis to derive both daily values and sensitivities of different amounts of positive reserve requirements in the portfolio context of the *flexible*, the *split flexible* and the *unflexible* portfolio.

Figure 6.28 shows the absolute impact of different positive reserve requirements on daily values, Deltas and Gammas of all three option portfolios and therewith the value and risk profile of the reserve requirements themselves. The indirect method leads to reduced Monte-Carlo convergence levels especially in case of Gamma. However weekday results of the second week can still clearly be distinguished (and have the additional advantage that they are not influenced by boundary effects of the finite valuation tenor). In general the results as presented in figure 6.28 show the same trends as discussed in sections 6.3, 6.4 and 6.5, i.e. the portfolio values and Gammas decrease with increasing positive reserve requirements while Deltas increase. The structural impact on daily results is similar for the flexible and the unflexible portfolio (both with three individual options each) while the split flexible portfolio is usually less affected by applied reserve requirements.

Interestingly the reserve impact on the portfolio Delta is typically positive while the same reserve requirements yields a reduced Gamma. A naive approach to replicate the impact of reserve requirements on a power plant portfolio could be to simply add a *negative* power plant (i.e. a short option position) to the original portfolio – according

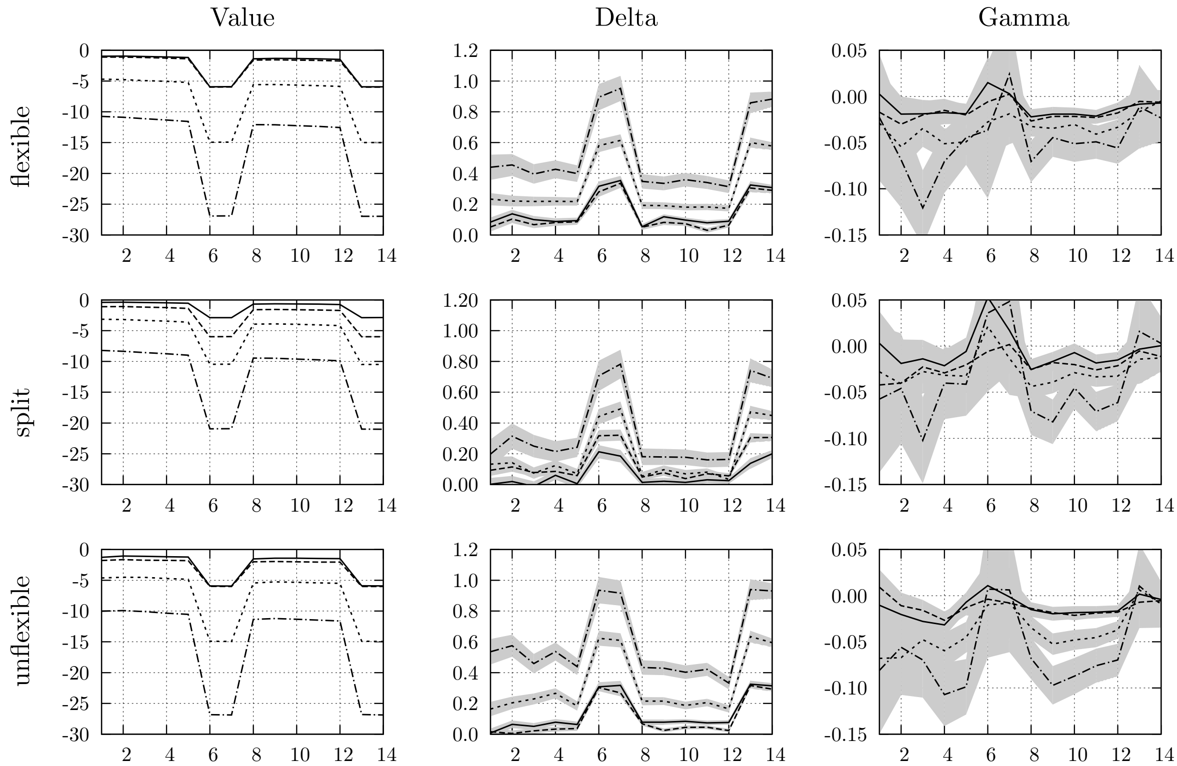

Fig. 6.28: Absolute impact of an increasing positive reserve requirement on daily portfolio values, Deltas and Gammas of the *flexible* portfolio (Top), the *split flexible* portfolio (Middle) and the *unflexible* portfolio (Bottom). Solid lines: $Res_{pos} = 0.3$, dashed lines: $Res_{pos} = 0.6$, dotted lines: $Res_{pos} = 1.0$, dashed-dotted lines: $Res_{pos} = 1.5$.

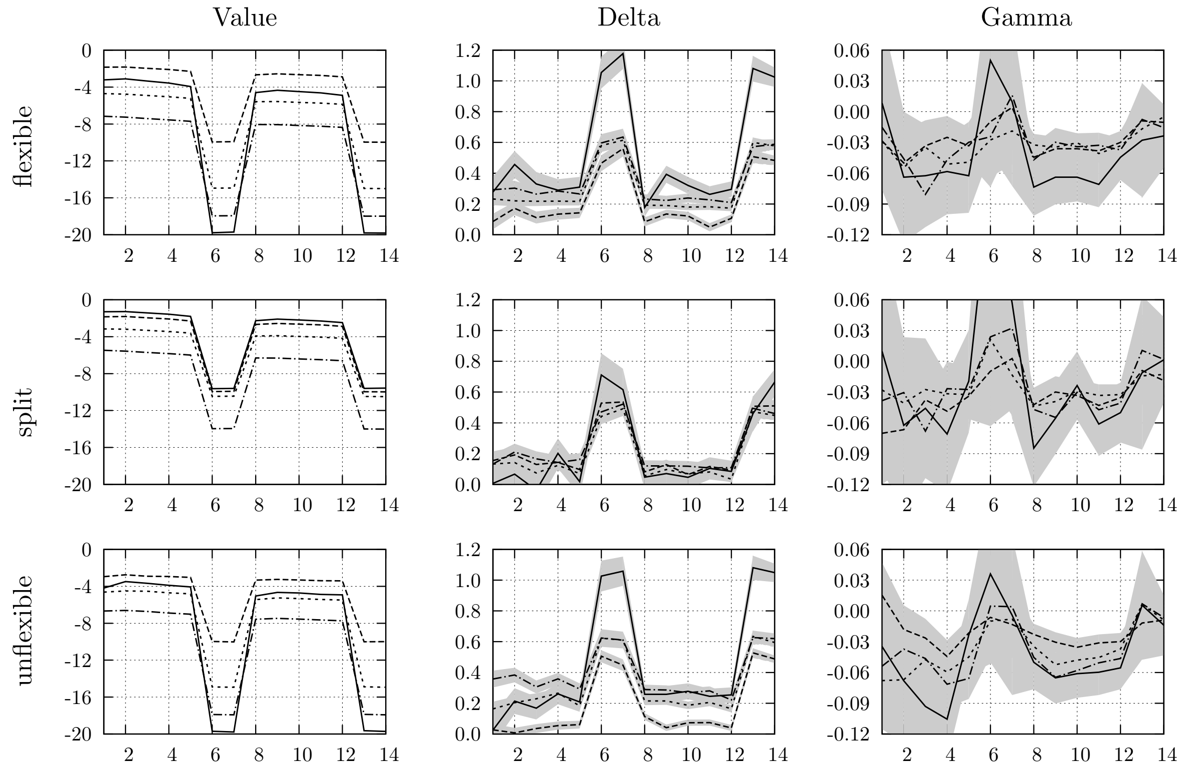

Fig. 6.29: Same results as shown in figure 6.28 but normalised by the level of applied reserve requirement. Solid lines: $Res_{pos} = 0.3$, dashed lines: $Res_{pos} = 0.6$, dotted lines: $Res_{pos} = 1.0$, dashed-dotted lines: $Res_{pos} = 1.5$.

to the logic that a reserve requirement must generally reduce the optionality of the portfolio, which can be replicated by an artificial short option position. However this approach would imply a negative impact on both Delta and Gamma of the portfolio and this is clearly not consistent with our results.

As mentioned before, the exact impact of a reserve requirement on an option portfolio strongly depends on the underlying hourly price structure in combination with the moneyness and further technical constraints of all involved options. Obviously the naive replication of reserve requirements via short options is not applicable in our setup of small portfolios, stylised out of the money options and heavily photovoltaics driven hourly prices. Even though this conclusion cannot be generalised it serves as a strong indication that the impact of reserve on different portfolios can be counterintuitive and it is highly recommended for operators of small physical portfolios to carefully analyse the impact of reserve constraints on their specific portfolio in advance of marketing or hedging decisions. As demonstrated, the Proxy Simulation Scheme approach can be a valuable tool to support such analyses.

Figure 6.29 shows the same results as figure 6.28 but normalised by the amount of applied positive reserve requirement. The negative values can directly be interpreted as specific cost per unit reserve requirement arising in case a portfolio is forced to serve reserve requirements, e.g. when the owner of the portfolio has sold a reserve contract to an external counterparty. In this case the owner of the portfolio needs to

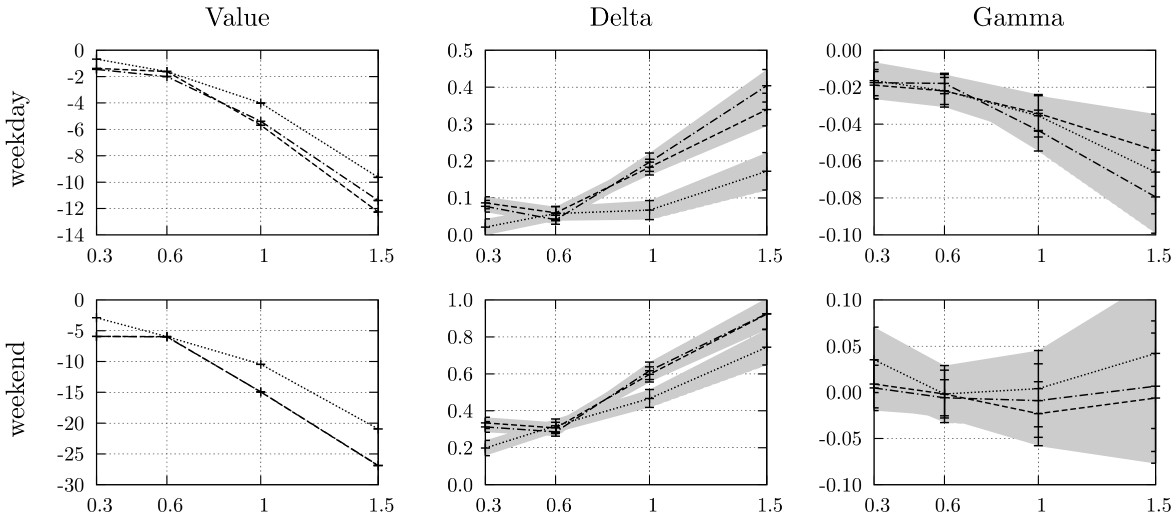

Fig. 6.30: Absolute impact of different positive reserve requirements on average daily results for weekdays and weekend days of the *flexible* portfolio (dashed lines), the *split flexible* portfolio (dotted lines) and the *unflexible* portfolio (dashed-dotted lines).

be compensated at least by the cost of reserve arising for his portfolio in order to not loose money in the transaction. At the same time the owner is able to hedge the sold reserve contract with linear products and options in the electricity market by following a standard Delta-Gamma hedge strategy that bases on sensitivities as provided in figures 6.28 and 6.29.

Remarkably while absolute results show a monotonous dependency on the level of reserve requirements for all portfolios (figure 6.28) this is not true for normalised results. According to figure 6.29 the specific cost of reserve is lowest for $Res_{pos} = 0.6$ in both the flexible and the unflexible portfolio, whereas in the split flexible portfolio the most cost effective reserve level is 0.3. This is due to the fact that in the first two portfolios the reserve imposed must-run condition is similar for 0.3 and 0.6 positive reserve requirement as in both cases at least one of the three options is obliged to run in all hours. As the must-run condition leads to the major share of the portfolio value reduction, the specific cost of reserve is obviously smaller for the higher overall reserve requirement of 0.6. In contrast the more fragmented split portfolio is able to always allocate a different number of smaller options to serve the imposed reserve requirements. This implies a step wise increase of the must-run condition on the portfolio which is consistent with the monotonous dependency of the daily values on the level of applied reserve requirement in figure 6.29.

Figures 6.30 and 6.31 provide a visualisation of average absolute and normalised weekday and weekend day results of different positive reserve requirements. This gives a comprehensive overview about all trends and different portfolio implications as discussed for daily results in figures 6.28 and 6.29 above.

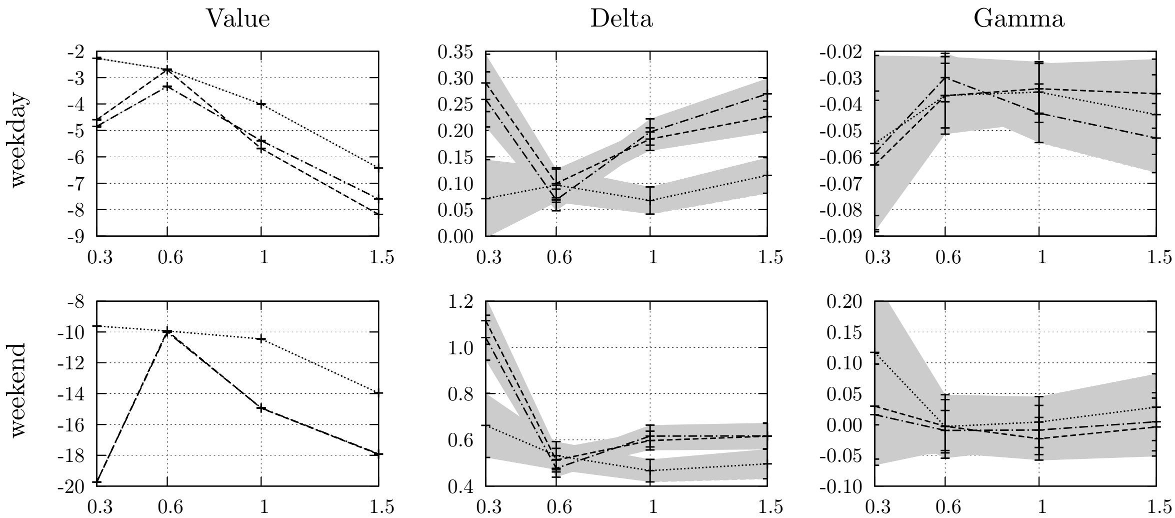

Fig. 6.31: Same results as shown in figure 6.30 but normalised by applied level of positive reserve requirement.

7 Approximate sensitivity derivation

As employed in the previous chapters, the Proxy Simulation Scheme method offers a huge advantage over standard Monte-Carlo approaches in terms of computational cost savings. However, long valuation tenors in combination with complex portfolio setups can still lead to numerical performance issues. In this chapter we discuss the applicability of three alternatives for the derivation of portfolio sensitivities Δ *and* Γ *which explicitly focus on an additional reduction of computation time.*

At first we use a truncated rolling intrinsic valuation tenor instead of solving the mixed integer optimisation problem for all remaining hours in each time step of each Monte-Carlo realisation as before. Secondly we apply a simple heuristic for the determination of the portfolio dispatch being significantly faster than even truncated rolling intrinsic valuation. The third method bases on polynomial approximations of the results of a very small number of Monte-Carlo realisations which leads to surprisingly close sensitivity approximations at negligible computational costs.

7.1 Truncated rolling intrinsic valuation tenor

As discussed in sections 3.2.1 and 5 technical constraints of power plant options lead to potentially strong time coupling effects between dispatch decisions of individual options at different points in time. The length of these temporal interdependencies naturally defines the required length of the mixed integer optimisation problem whose solutions determine the option dispatch realisations. The hourly length of the problem should always exceed the maximum impact length of the time coupling effects in order to yield unbiased dispatch realisations. However at the same time it is recommended to use the shortest applicable mixed integer problem for receiving the maximum amount of numerical cost savings without inducing numerical artifacts.

In order to define this shortest appropriate length we apply a direct numerical approach that still bases on full rolling intrinsic valuation as presented in section 3.2.3. In a first step this standard implementation is used to derive 100 stand-alone hourly dispatch realisations of all options of the flexible and the unflexible option portfolio (compare section 6.2) for a 14 day valuation tenor. In a second step these dispatches are artificially modified for a certain hour $\tilde{l}$ and all preceding hours as follows:

1. If the evaluated power plant option has been running in hour $\tilde{l}$ of an original dispatch realisation, its load is set to zero in hour $\tilde{l}$ and all preceding hours $l < \tilde{l}$.

2. If the evaluated power plant option has been offline in hour $\tilde{l}$ of an original dispatch realisation, it is started in hour $\tilde{l}$ while its load is set to zero in all preceding hours $l < \tilde{l}$.

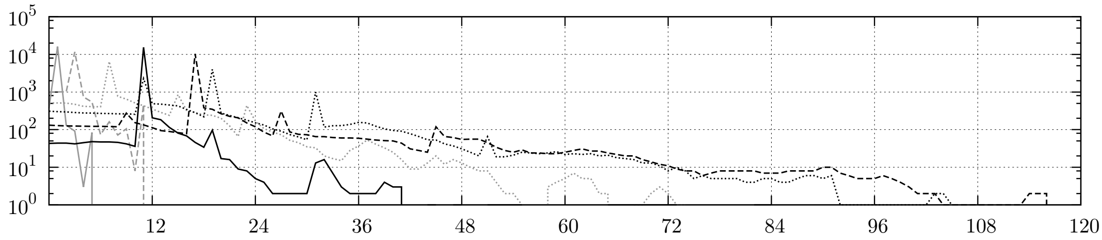

Fig. 7.1: Absolute frequency of hourly lengths of artificial dispatch modifications for stand alone options of the *flexible* portfolio (gray lines; dotted: $t_{up/down} = 8/8$, $\kappa = 50$, dashed: $t_{up/down} = 4/4$, $\kappa = 60$, solid: $t_{up/down} = 2/2$, $\kappa = 70$) and the *unflexible* portfolio (black lines; dotted: $t_{up/down} = 20/12$, $\kappa = 50$, dashed: $t_{up/down} = 18/10$, $\kappa = 60$, solid: $t_{up/down} = 12/8$, $\kappa = 70$).

All dispatch modifications until and including hour $\tilde{l}$ are then allocated as fixed parameters preventing them from any further changes by the mixed integer solver. $\tilde{l}$ covers a time frame of 7 days starting from the first hour of the second day of the valuation tenor, i.e. $\tilde{l} \in [25, 192]$. In a third step the standard valuation framework is used to derive final dispatch realisations of all hours of the 14 days valuation tenor for all combinations of $\tilde{l}$ and original dispatch realisation where exactly the same price evolutions are applied as valid for the original dispatch realisation. By comparing these final dispatches with the original dispatches it can be determined until which hour $l > \tilde{l}$ the applied modification impacts the final hourly dispatch realisation. This length can directly be interpreted as one stochastic realisation of the required length of the time coupling effect.

Impact decay rates of artificial dispatch modifications In total the approach generates $7 \times 24 \times 100 = 16800$ stochastic impact lengths whose absolute frequency is provided for all options of the flexible and the unflexible portfolio in figure 7.1. The distribution curves show distinctive spikes at hourly impact lengths of t^p_{up}, t^p_{down} and associated combinations and multiples. This can be explained by the fact that an artificial start of option p comparably often enforces the option to run t^p_{up} hours before the successive hourly dispatch can be aligned by shutting down the option again (spike at $\tilde{l} = t^p_{up}$). The same argumentation is valid in case of artificial shut-downs that lead to a spike at $\tilde{l} = t^p_{down}$. Smaller spikes are due to temporal interdependencies between consecutive start and shut-down decisions. Obviously the technically more constrained options of the unflexible portfolio show generally longer impact lengths of dispatch modifications and therefore generally longer time coupling effects than the options of the flexible option portfolio.

A more comprehensive visualisation of this effect is given by figure 7.2 which shows the percentage of modified dispatch realisations whose impact length exceeds a certain amount of hours as displayed on the x axis. These curves can be interpreted as impact decay rates of the artificial dispatch modifications. Table 7.1 provides a summary of associated numerical results. The length of time coupling effects of options 2 and 3 of the flexible portfolio is below 12 hours for all realisations while the technical constraints of option 1 ($t_{up/down} = 8/8$, $\kappa = 50$) lead to time coupling even above 2 days of the

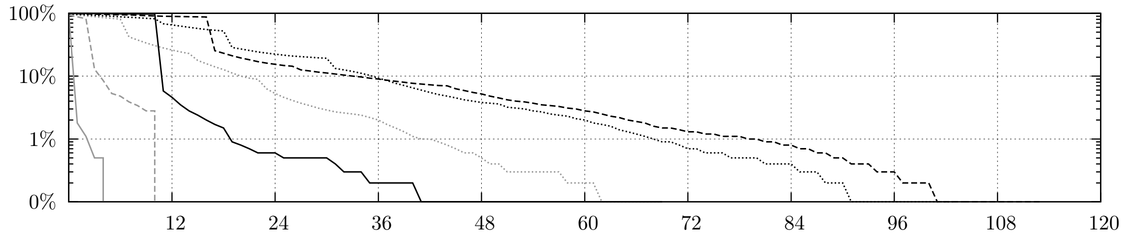

Fig. 7.2: Percentage of modified dispatch realisations with impact length exceeding the number of hours as shown on the x axis (impact decay rate) for all options as shown in figure 7.1.

t_{up}/t_{down}:	8/8	4/4	2/2	20/12	18/10	12/8
κ:	50	60	70	50	60	70
20%	14	2	0	28	19	10
10%	20	3	0	35	33	10
5%	24	5	0	43	48	11
2%	36	10	1	60	65	16
1%	42	10	2	68	80	19

Table 7.1: Number of hours exceeded by x% of the impact length of modified dispatch realisations for different stand alone options.

valuation tenor. In contrast options 1 and 2 of the unflexible portfolio are affected by time coupling for a period up to 4 days and only option 3 leads to comparably moderate impact length below 48 hours.

Definition of truncated rolling intrinsic approaches Based on this analysis a length of 2 days of the mixed integer problem should be appropriate for all stand-alone options of the flexible options portfolio while a shorter tenor of 1 day should still be sufficient for options 2 and 3. The same 2 days tenor could be questionable for a valuation of options 1 and 2 of the unflexible portfolio but should be sufficient for option 3, whereas a 1 day tenor is clearly not appropriate for the unflexible portfolio.

For a systematical assessment of these expectations we evaluate the flexible portfolio, the split flexible portfolio and the unflexible portfolio with and without positive reserve requirements by using the following three truncated rolling intrinsic approaches:

1. 1 day rolling intrinsic without taking historical dispatch decisions into account,

2. 1 day rolling intrinsic taking historical dispatch decisions into account,

3. 2 days rolling intrinsic taking historical dispatch decisions into account.

Obviously the first approach offers the largest numerical cost advantage due to a significantly reduced mixed integer optimisation problem. However as this approach does not take historical dispatch decisions into account it systematically overestimates the flexibility and therefore the value of the evaluated option portfolios.

Results of truncated rolling intrinsic approaches This is confirmed by figures 7.3 and 7.4 showing a comparison of average portfolio results as derived by full rolling intrinsic and the three truncated rolling intrinsic approaches. Both figures provide results of the flexible, the flexible split and the unflexible portfolio subject to three different levels of positive reserve requirements including $Res_{pos} = 0.0$, $Res_{pos} = 0.3$ and $Res_{pos} = 1.0$. 95% confidence intervals are indicated by gray areas or error bars. Figure 7.3 shows average results for weekdays and figure 7.4 shows associated results for weekend days. Especially for the unflexible portfolio both figures show clearly that 1 day rolling intrinsic without taking historical dispatch decisions into account tends to overestimate portfolio values with respect to full rolling intrinsic results. In contrary 1 day rolling intrinsic taking historical decisions into account leads to partly significant underestimations for the unflexible portfolio. This can be interpreted as a general information bias which occurs as future price information are only taken into account for the next 24 hours and neglected afterwards, leading to an insufficient utilisation of the entire available market information. The bias is no longer visible when the rolling intrinsic tenor is extended to 2 days effectively doubling the involved market information. This still highly simplified setup yields surprisingly low deviations to numerical results of full rolling intrinsic valuation even for the unflexible portfolio. However this observation can be understood with the help of figure 7.2 which shows, that a 48 hours rolling intrinsic tenor covers more than 90% of the time coupling effects of the unflexible portfolio without reserve requirements. This obviously limits the induced approximation error and leads to the close alignment between fully flexible results and results of the 2 days truncated rolling intrinsic approach.

Truncated results of the flexible portfolio show very similar behavior on weekdays as discussed for the unflexible portfolio as long as no reserve requirement is applied. For weekend days or in case of positive reserve requirements all truncated methods yield similar results which are very close to results of the full rolling intrinsic valuation. On weekend days the reason for this observations lies within the limited dispatch variations as all power plant options are deep out of the money and are typically offline as long as they are not required to serve reserve requirement. This simple dispatch behavior can obviously be covered even by a 1 day rolling intrinsic without taking historical dispatch decisions into account. According to figure 7.2 only option 1 is affected by a truncated rolling intrinsic valuation of 1 day which could lead to result deviations between the different truncated valuation approaches. However, as discussed in section 6, in case of reserve requirements option 1 usually serves the major part of the portfolio reserve requirement leading to a strong limitation of dispatch flexibility of this option. The resulting dispatch pattern can obviously be replicated even by a heavily truncated rolling intrinsic tenor of 1 day.

Apart from averaged portfolio values it is also interesting to compare the daily structure of results as calculated via full and truncated rolling intrinsic. From the impact decay analysis (figure 7.2) it can be expected that those structures should be comparable for the flexible and the flexible split portfolio as both portfolios contain options with similar technical constraints. Indeed the truncated valuation approaches lead to similar result structures for these portfolios even in case of reserve requirements as can be seen in figure 7.5 providing a comparison of daily results of the flexible portfolio

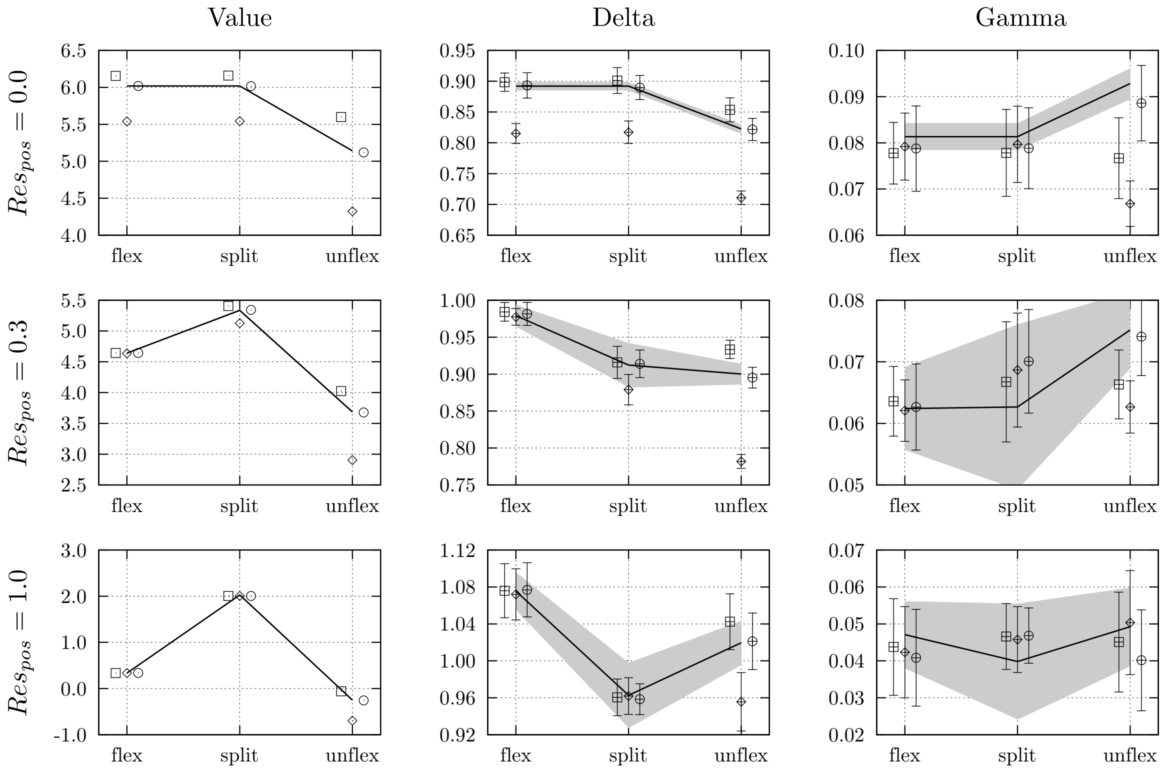

Fig. 7.3: Comparison of average daily results for weekdays of all portfolios as derived with full rolling intrinsic Monte-Carlo (solid lines, gray areas indicate 95% confidence intervals) with results from 1 day rolling intrinsic without history (rectangles), results from 1 day rolling intrinsic taking historic dispatches into account (diamonds) and results from 2 days rolling intrinsic taking historic dispatches into account (circles). $n = 3$m realisations for all simplified Monte-Carlo simulations.

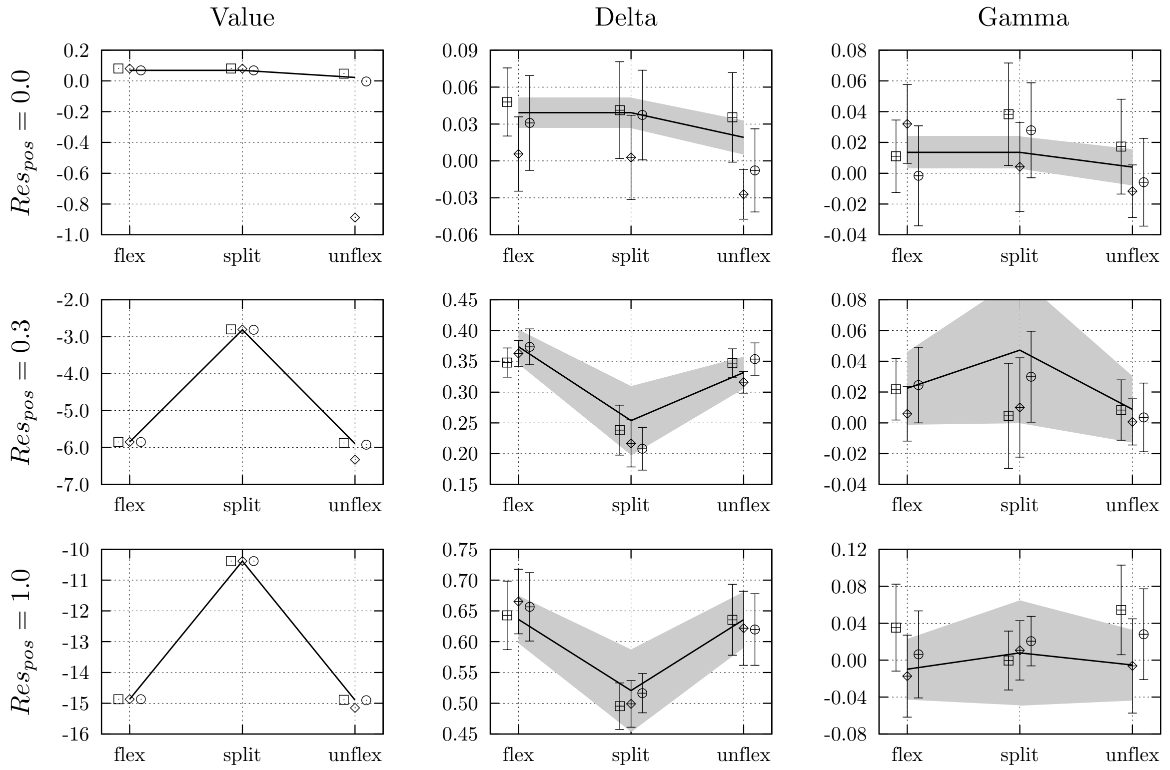

Fig. 7.4: Weekend day results corresponding to figure 7.3.

based on full rolling intrinsic and all three truncated rolling intrinsic approaches.

This good alignment is no longer generally valid for results of the unflexible portfolio as demonstrated by figures 7.6 and 7.7. Similar results can only be observed for the 2 days rolling intrinsic approach while the two 1 day tenors lead to significant numerical artifacts. Especially the combination of a 1 day valuation tenor and utilisation of historic dispatches leads to significant daily result spikes as the approach bases its dispatch decisions on incomplete information about the future and thereby does not take technical constraints correctly into account. This leads to suboptimal dispatch decisions being beneficial for the respective day but then affecting the entire hourly dispatch of the subsequent day as technical constraints are fully considered retrospectively. The lengths of minimum up- and down-times allow a correction of this dispatch decision only during the second day after the original decision. This interdependency yields the characteristic 2 days spike structures which can be observed both without and with positive reserve requirement in figures 7.6 and 7.7.

7.2 Heuristic approach via rules of thumb

The great advantage of a heuristic approach to derive the dispatch of the option portfolios lies in its superior numerical performance. This is due to the fact that the heuristic approach replaces the time consuming solution of a mixed integer problem by a comparably simple numerical decision algorithm. Obviously this routine should be able to

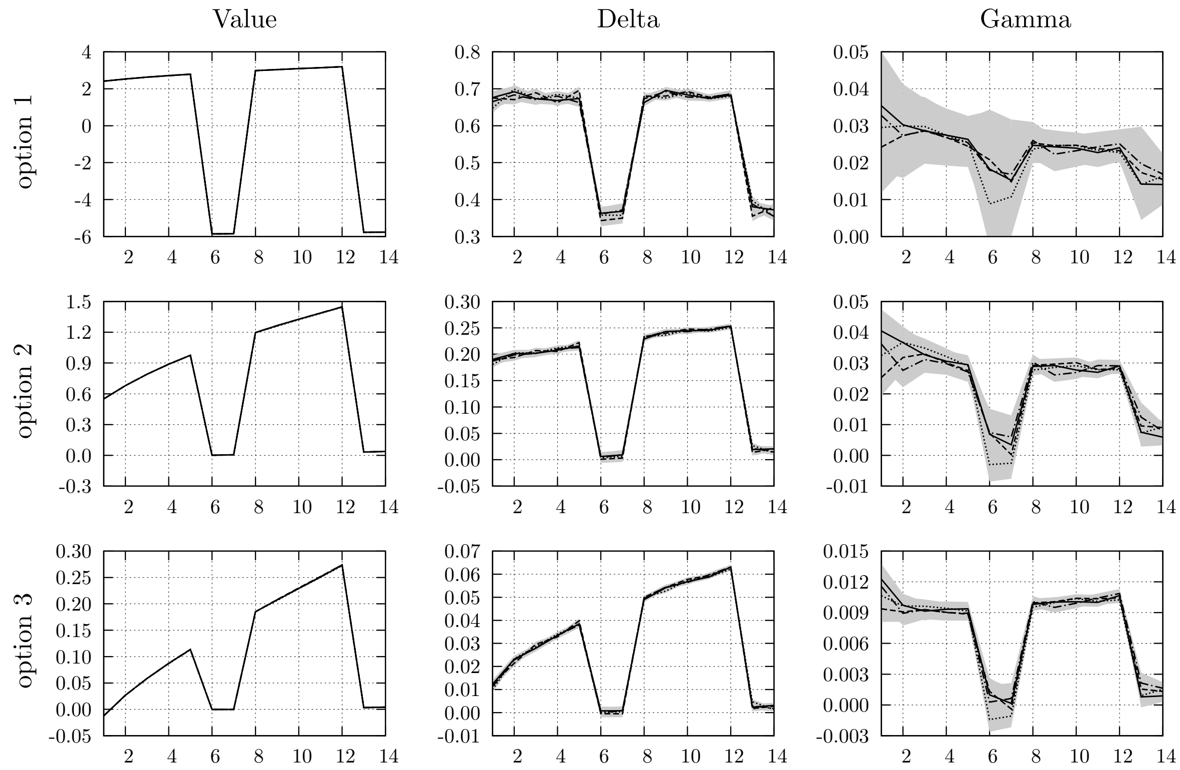

Fig. 7.5: Daily results of the *flexible* option portfolio subject to minor reserve constraints ($Res_{pos} = 0.3$) based on full rolling intrinsic (solid lines), 1 day rolling intrinsic without history (dashed lines), 1 day rolling intrinsic with history (dotted lines) and 2 days rolling intrinsic with history (dashed-dotted lines). Gray areas indicate associated 95% confidence intervals. $n = 3$m realisations for all simplified simulations.

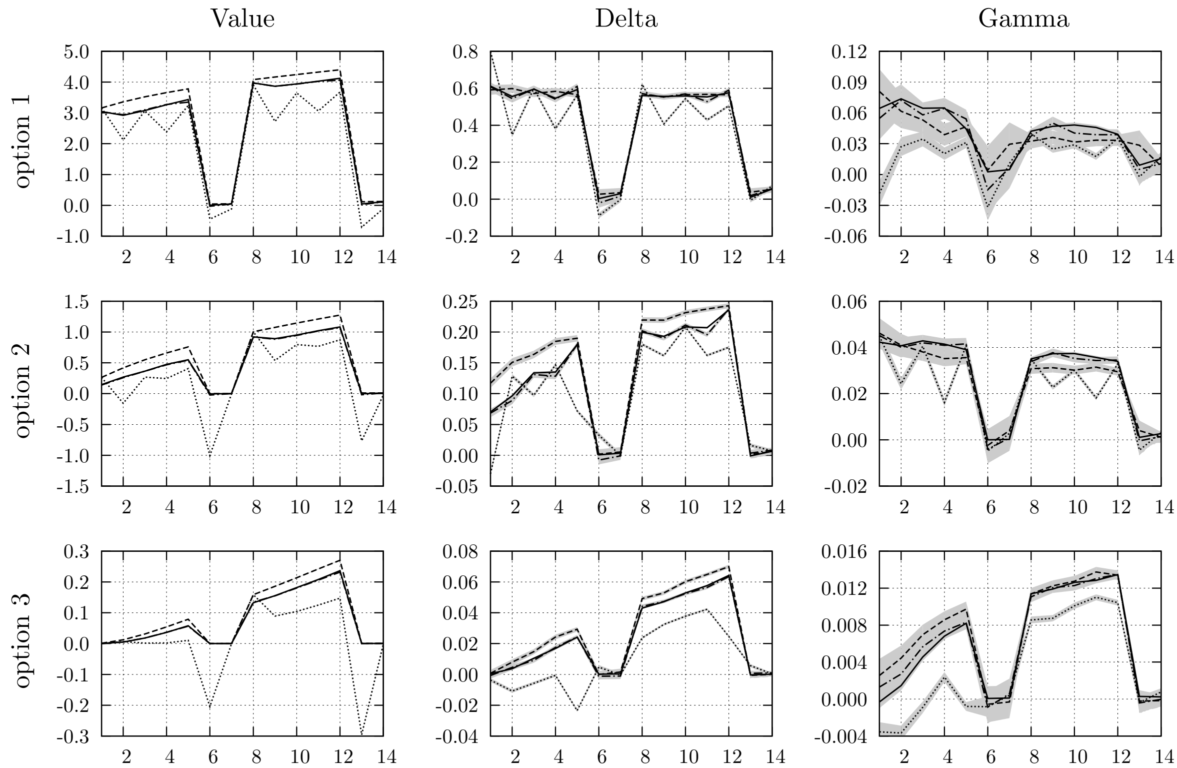

Fig. 7.6: Results of the *unflexible* option portfolio without reserve constraints (similar to figure 7.5).

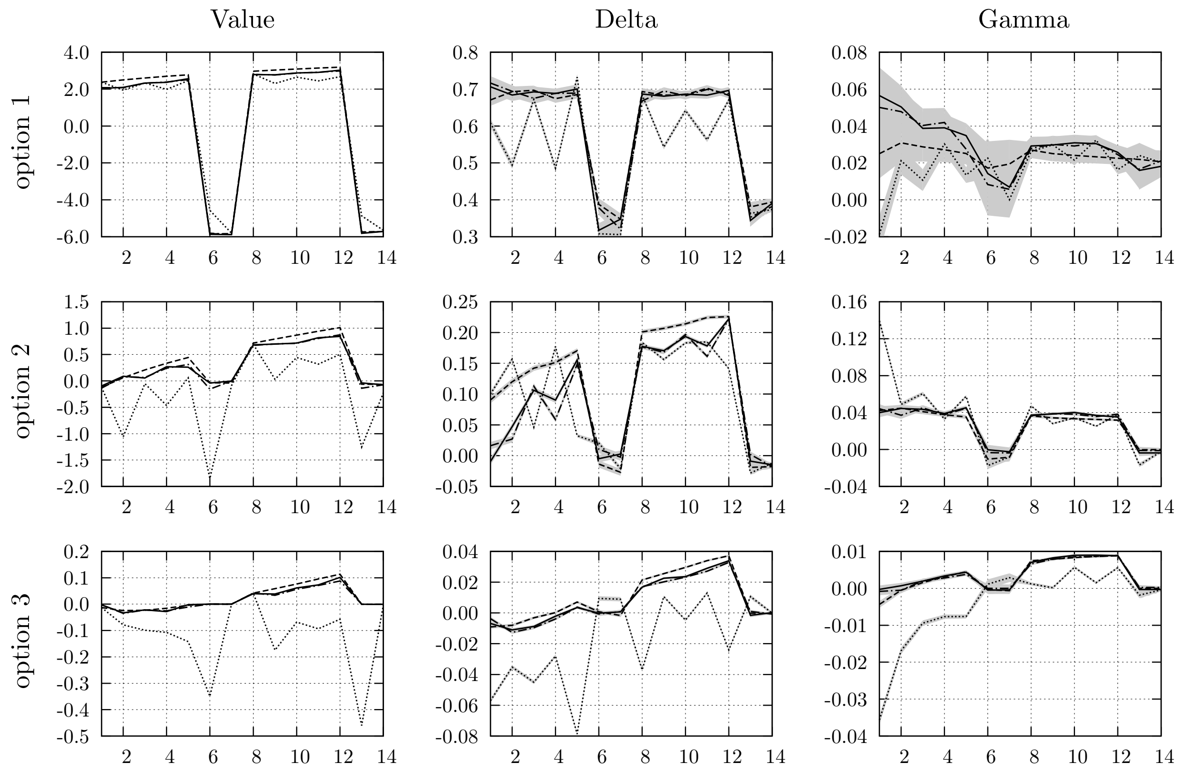

Fig. 7.7: Results of the *unflexible* option portfolio subject to a minor reserve constraint of $Res_{pos} = 0.3$ (similar to figure 7.5).

cover the most significant value drivers appropriately on the basis of a limited number of decisions, in order to derive a dispatch which is not necessarily the optimal dispatch but ideally close.

In case of a portfolio of power plant options subject to technical constraints and reserve requirements the heuristic approach can be very complicated when it tries to cover all constraints in parallel. In addition this procedure would potentially be very close to solving the mixed integer optimisation mip problem and would therefore not be a great invention with respect to the typical approach as used in this thesis. Instead a natural and more straight forward approach is to define the reserve serving power plant options either ex-post or ex-ante and thereby to decouple the treatment of technical constraints and portfolio constraints as imposed by the reserve requirement.

Definition of dispatching heuristics In order to come up with a simple but nonetheless sensible heuristic we chose the latter method and propose three slightly different heuristics that allocate the reserve requirements to individual power plant options before the remaining options are "freely" dispatched. A structural flowchart of this reserve allocation algorithm is provided in figure 7.8. The heuristic effectively bases only on a very small number of steps:

1. For day $q = k$ collect all plants which are generally able to serve reserve requirements, i.e. plants that have already been running in the last hour before day k or can be started in the first hour of day k.

2. Order these plants ascending by the absolute difference between their strike and the base load electricity price of day k. In case of a 2 day heuristic the average base load price of days k and $k + 1$ is applied.

3. Based on this order allocate as much reserve requirement to these individual plants as possible, i.e. the entire range between maximum and minimum load if requested.

4. If all reserve requirement has been allocated, the last affected option is subject to a must-run condition but may have some spare dispatch flexibility. It will be dispatched optimally on an hourly basis.

5. All other options will be "freely" dispatched according to the second part of the heuristic approach.

A flowchart of the second part of the heuristic approach is provided in figure 7.9. It essentially performes the following steps:

1. For day $q = k$ collect all plants which do not serve reserve requirements.

2. Start with the first plant p and the first hour of day k.

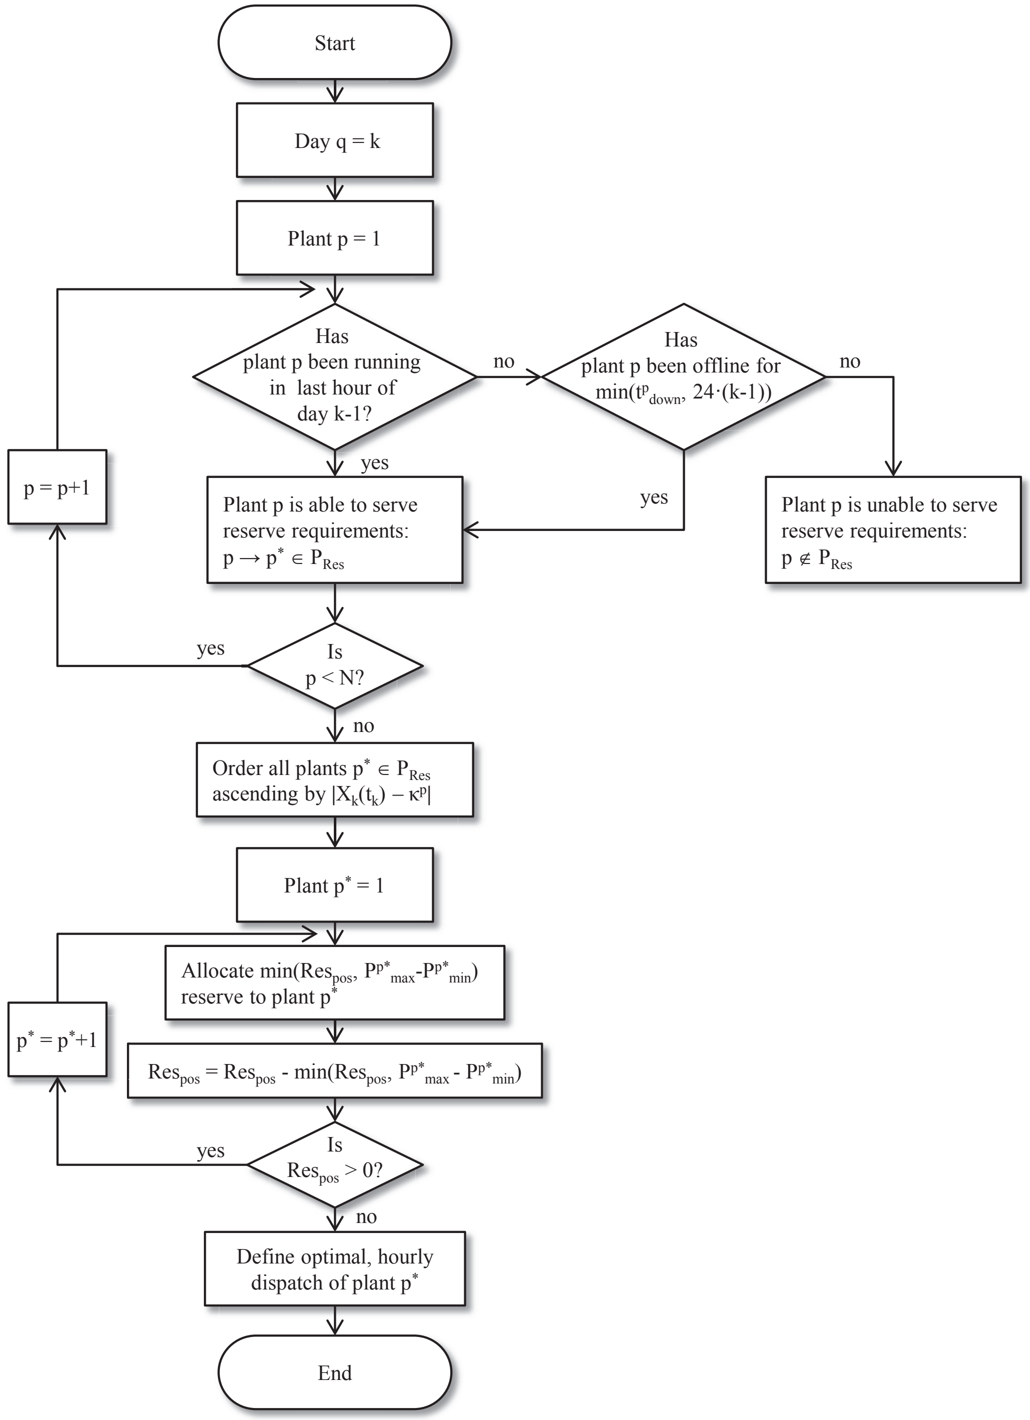

Fig. 7.8: Flowchart of implemented heuristic approach to ex-ante allocation of reserve constraints to individual power plants.

3. If the plant has been offline in the hour before it will be started in case this is economically, i.e. whenever the power plant shows a positive value contribution for an hourly block of t^p_{up} and as long as the start should not be delayed by one hour.

4. If the plant has been online in the hour before it will be shut down in case this is economically, i.e. whenever the power plant shows a negative value contribution for an hourly block of t^p_{down} and as long as the shut down should not be delayed by one hour.

5. Otherwise the status of power plant p remains unchanged.

6. Proceed to next hour of day k and perform the same decision until the last hour of day k is reached.

7. Proceed to the next power plant and start again with the first hour of day k.

The three heuristic approaches all base on the same algorithmic structure as presented in figures 7.8 and 7.9. However they differ in both the time period which is applied to derive the absolute differences between strike prices and the base load electricity price (i.e. the amount of future price information that is taken into account for reserve allocation) and in the usage of historic dispatch information:

1. No historical dispatches are taken into account and a 1 day time period is applied to derive the absolute differences between strikes and the base load electricity price in course of reserve allocation.

2. Historical dispatches are taken into account and a 1 day time period is applied to derive the absolute differences between strikes and the base load electricity price in course of reserve allocation.

3. Historical dispatches are taken into account and a 2 days time period is applied to derive the absolute differences between strikes and the base load electricity price in course of reserve allocation.

Results of dispatching heuristics Similar to figures 7.3 and 7.4 figures 7.10 and 7.11 show average results of the flexible, the split flexible and the unflexible portfolio based on all three heuristics in comparison to associated results of the full rolling intrinsic approach.

The heuristics work rather well for both the flexible and the split portfolio - which are not heavily affected by technical constraints - as long as no reserve requirements are applied on the portfolio level. In contrast it shows significant deviations for the unflexible portfolio even without any reserve requirements indicating, that the chronological dispatching approach (figure 7.9) may still be too simplistic for this more complex optimisation problem.

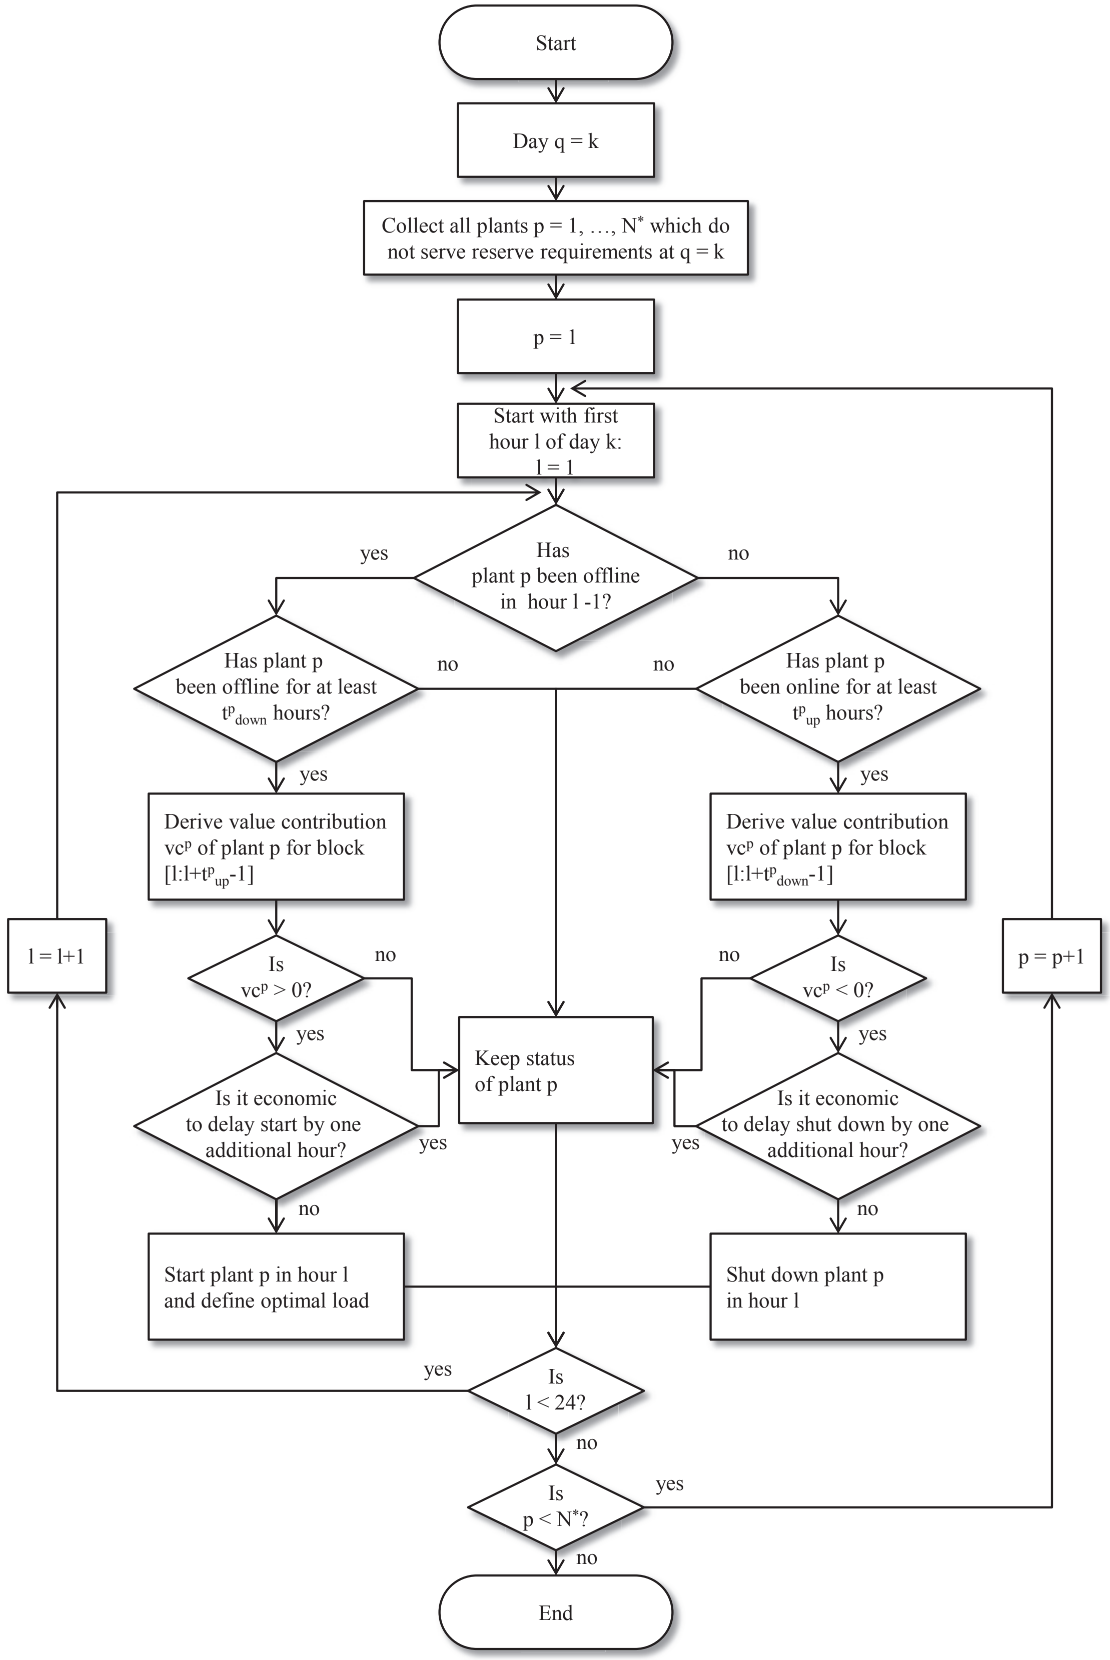

Fig. 7.9: Flowchart of implemented heuristic dispatching approach for all plants not serving reserve requirements.

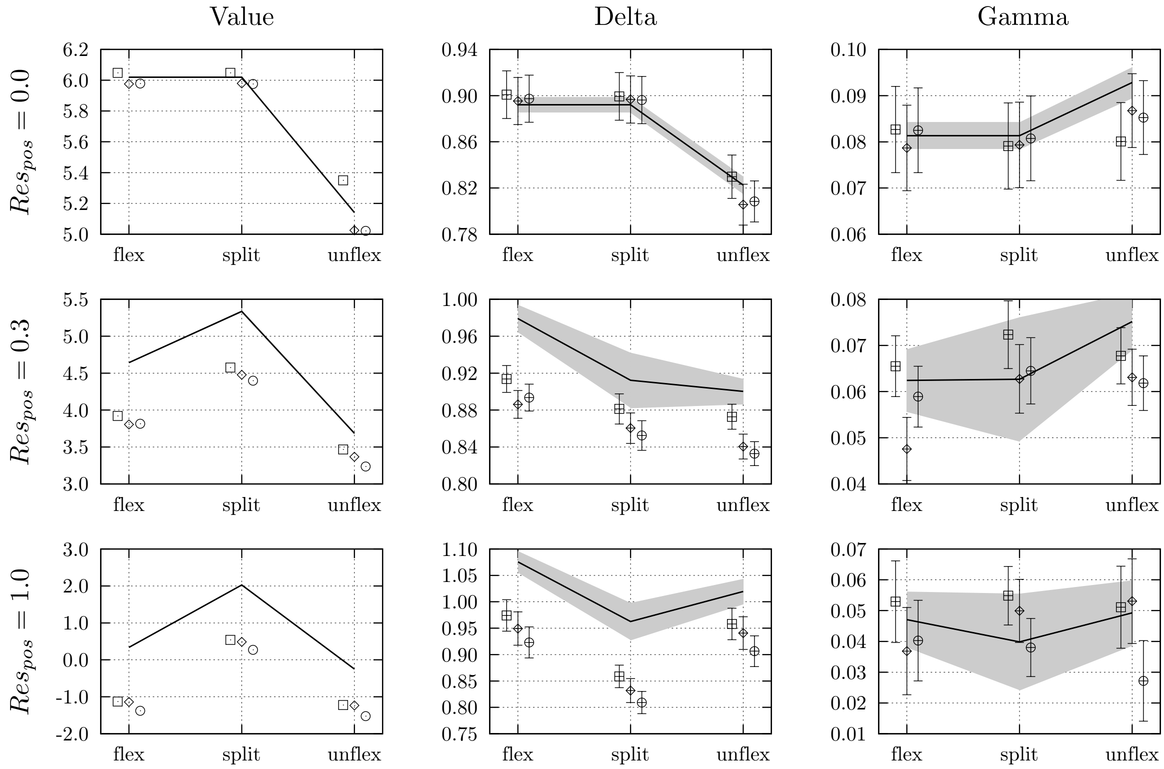

Fig. 7.10: Comparison of average daily results for weekdays (basis is 2. week of valuation tenor) of all portfolios as derived with full rolling intrinsic Monte Carlo (solid lines, gray areas indicate 95% confidence intervals) with results from 1 day heuristic without history (rectangles), results from 1 day heuristic taking historic dispatches into account (diamonds) and results from 2 days heuristic taking historic dispatches into account (circles). $n = 6$m realisations for all heuristics.

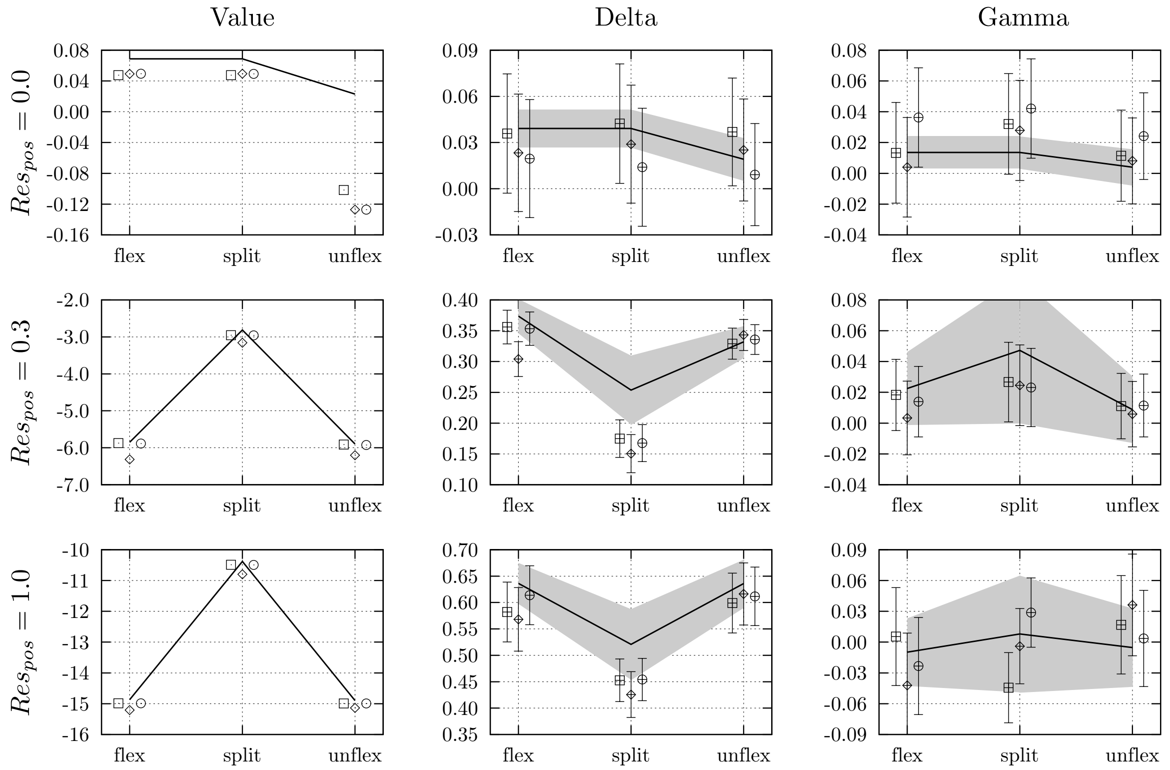

Fig. 7.11: Weekend day results corresponding to figure 7.10.

In case of imposed reserve requirements all heuristics significantly underestimate the portfolio values on weekdays but show a surprisingly good alignment on weekends. This is possible as all options are generally deep out of the money on weekends and reserve is typically served only by the cheapest options in the portfolio as shown in figures 6.22 and 6.24. The same options are also chosen by the heuristic approach trying to minimise the absolute difference between option strike prices and the daily base load forward price. Therefore the heuristics are able to replicate rolling intrinsic portfolio values very well in this special setup.

However portfolio Deltas are generally underestimated for all reserve serving portfolios on weekdays and weekend days and Gamma results also tend to be downwards biased. These deviations are not improved by an extension of the price period in the first part of the heuristics from 1 day to 2 days. The ex-ante reserve allocation approach is obviously not suited to cover all portfolio dependencies in case of imposed positive reserve requirements.

In addition to the analysis of average portfolio results it is again interesting to compare the daily structure of heuristic results with the structure of results based on full rolling intrinsic. Similar to figures 7.5, 7.6 and 7.7 the figures 7.12, 7.13 and 7.14 provide daily results based on all three heuristics of the flexible portfolio subject to minor reserve requirements, the unflexible portfolio without reserve requirements and the unflexible portfolio subject to minor reserve requirements in comparison to full rolling intrinsic results.

In contrast to the shortened rolling intrinsic methods as discussed in section 7.1

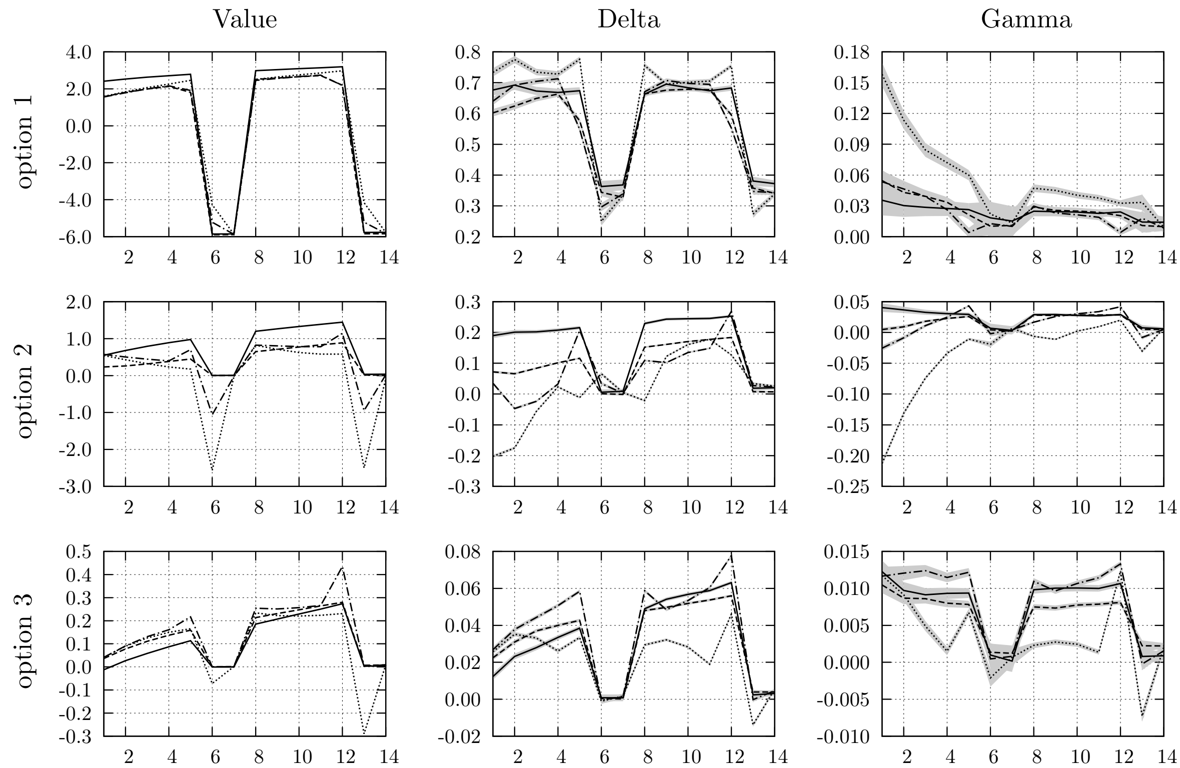

Fig. 7.12: Daily results of the *flexible* option portfolio subject to minor reserve constraints ($Res_{pos} = 0.3$) based on full rolling intrinsic (solid lines), 1 day heuristic without history (dashed lines), 1 day heuristic with history (dotted lines) and 2 day heuristic with history (dashed-dotted lines). Gray areas indicate associated 95% confidence intervals. $n = 6$m realisations for all simplified simulations.

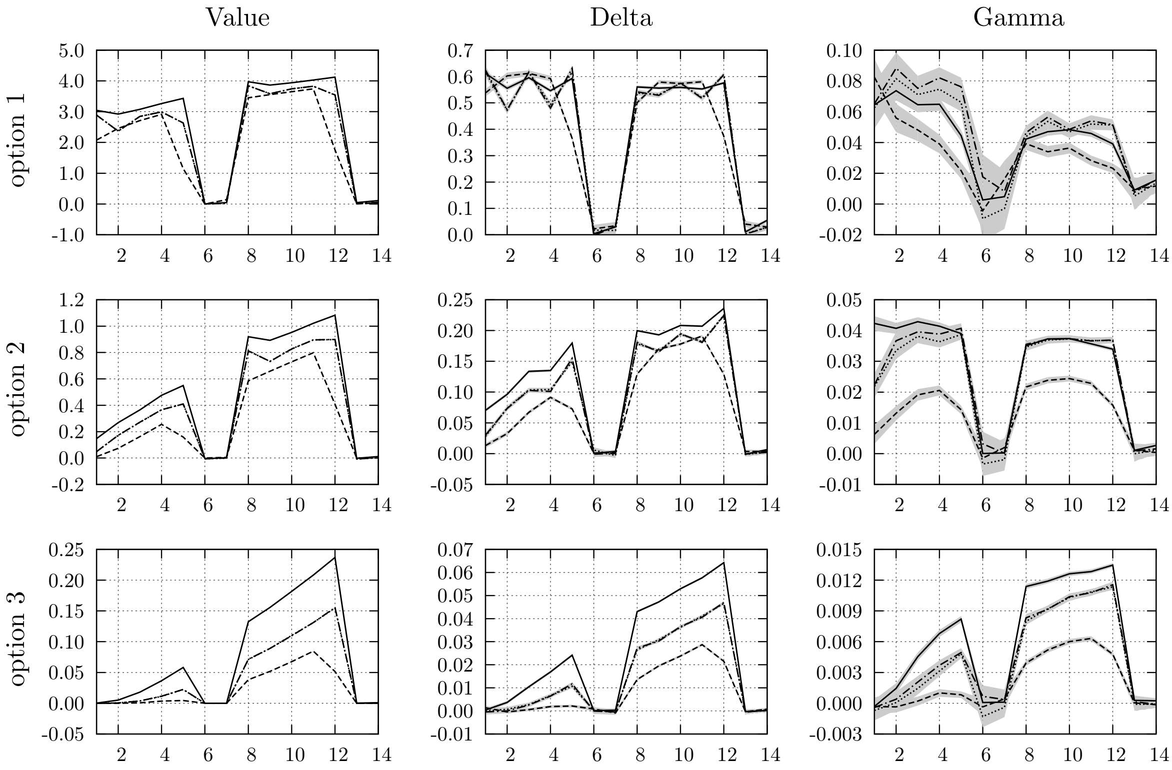

Fig. 7.13: Results of the *unflexible* option portfolio without reserve constraints (similar to figure 7.12).

the heuristic approaches show rather bad alignment with full rolling intrinsic daily results of the flexible and flexible split portfolio. Surprisingly the structure of results is comparably well aligned for the unflexible portfolio without reserve requirements apart from a significant downwards bias as already observed in the discussion of average portfolio results. In case of absent reserve requirements (figure 7.13) almost no difference can be observed between 1 day and 2 day heuristic with history due to the fact the these heuristics only differ in the way they treat reserve requirements. But when the unflexible portfolio is subject to minor reserve requirements (figure 7.14) this alignment entirely breaks and the comparison of result structures shows significant deviations. Also this analysis of the structure of daily results supports the conclusion that the applied simplistic heuristic approaches are not able to cover all intra-portfolio dependencies adequately.

7.3 Polynomial sensitivities

In the preceding sections 7.1 and 7.2 we discussed numerical results of simplified dispatching approaches including truncated rolling intrinsic methods and heuristics, which both offer significant numerical cost savings but are simultaneously affected by result biases for more complex portfolio setups. In this section we present a completely different approach that bases on polynomial approximations of daily dispatches $D_{k,i}$ and

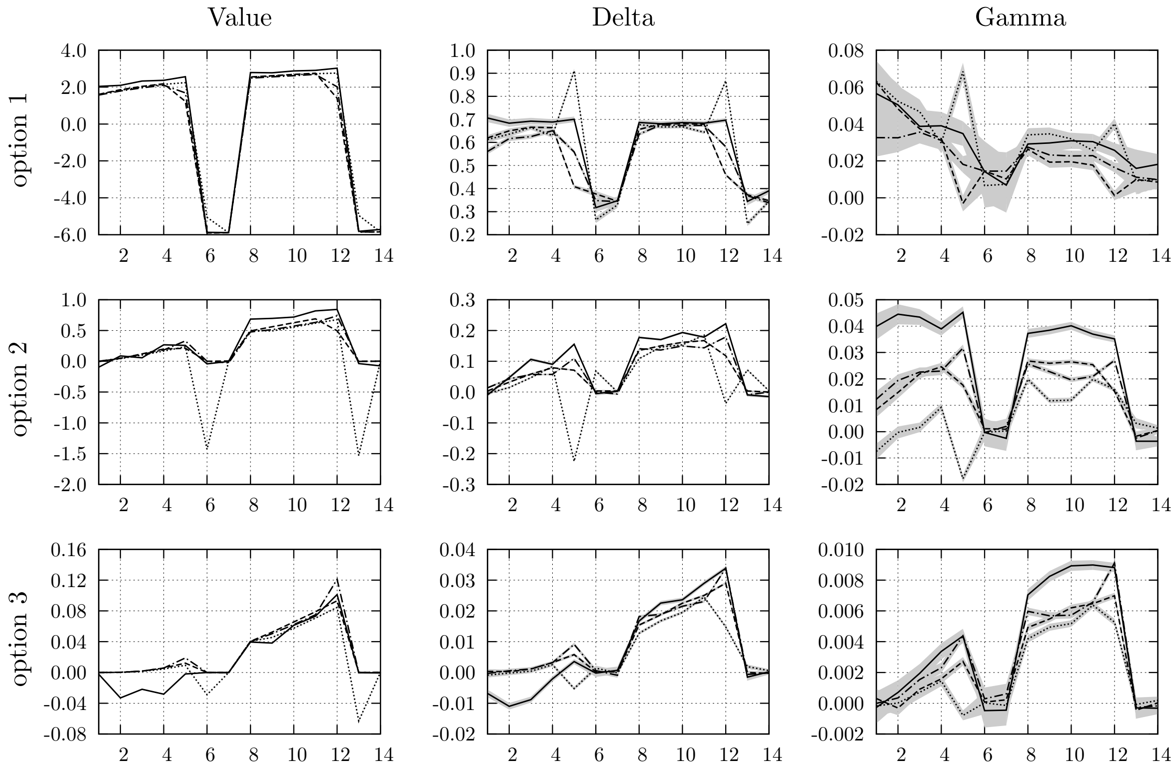

Fig. 7.14: Results of the *unflexible* option portfolio subject to a minor reserve constraint of $Res_{pos} = 0.3$ (similar to figure 7.12).

payoffs $f_{k,i}$ for a small number of Monte-Carlo realisations $i = 1, \ldots, n$. The underlying idea is similar to the method of Fu et al. (2012) with the general difference that they analysed the dependency between payoff and initial forward price instead of payoff and realised (forward) price.

Figure 7.15 shows 1000 realisations of $D_{k,i}$ and $f_{k,i}$ of the flexible portfolio subject to a minor reserve requirement of $Res_{pos} = 0.3$ on a typical weekday, ordered by the associated daily base load price realisation $X_{k,i}(t_k)$. In addition figure 7.15 provides linear least squares fits of the dependency of daily dispatches on daily price realisations and quadratic fits of the dependency of daily payoffs on daily price realisations. The quadratic expression of f_k seems to be reasonable as approximately

$$D_k = aX_k(t_k) + b \tag{7.1}$$

for a huge subset of realised prices and the daily payoff of option p can most simplistically be expressed as

$$f_k(X_k(t_k)) = D_k \cdot (X_k(t_k) - \kappa^p) \sim (X_k(t_k))^2 \,. \tag{7.2}$$

Please note that the observed linear dependency between realised payoff and realised daily price cannot be generalised but is heavily depending on both the applied hourly price shape and the characteristics of the evaluated options.

Expression (7.2) motivates a Taylor expansion of the daily dispatch $f_k(X_k(t_k))$

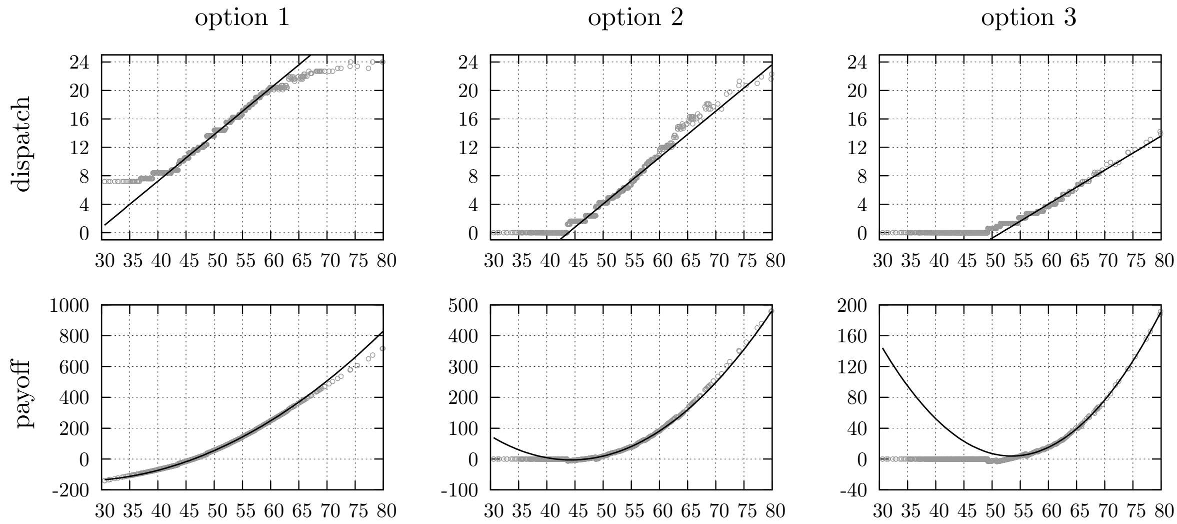

Fig. 7.15: Top: 1000 daily dispatch realisations of day 10 (i.e. $k = 10$) for all options of the *flexible* portfolio subject to minor reserve constraints ($Res_{pos} = 0.3$) over associated daily price realisations and linear function fit.
Bottom: 1000 daily payoff realisations for the same setup and quadratic polynomial fit.

around the average daily realised price $\overline{X}_k(t_k) = X_k(t_0)$ (50 on weekdays, 30 on weekend days)

$$\begin{aligned} f_k(X_k(t_k)) &= f_k(\overline{X}_k(t_k)) \qquad (7.3)\\ &+ \frac{\partial}{\partial X_k(t_k)} f_k(\overline{X}_k(t_k)) \left(X_k(t_k) - \overline{X}_k(t_k)\right) \\ &+ \frac{1}{2} \frac{\partial^2}{\partial (X_k(t_k))^2} f_k(\overline{X}_k(t_k)) \left(X_k(t_k) - \overline{X}_k(t_k)\right)^2 \\ &+ O\left(\left(X_k(t_k) - \overline{X}_k(t_k)\right)^3\right). \end{aligned}$$

The coefficients of this polynomial expression can be estimated via a quadratic approximation of $f_k\left(X_k(t_k) - \overline{X}_k(t_k)\right)$ and can directly be interpreted as estimators of the sensitivities Δ_k and $\frac{1}{2}\Gamma_k$. Associated analytical expressions can be derived from (7.2) as follows:

$$\begin{aligned} \frac{\partial}{\partial X_k(t_k)} f_k\left(X_k(t_k)\right) &= \frac{\partial}{\partial X_k(t_k)} \left(D_k \cdot (X_k(t_k) - \kappa^p)\right) \\ &= \frac{\partial}{\partial X_k(t_k)} \left(a(X_k(t_k))^2 - (a\kappa^p + b) X_k(t_k) - b\kappa^p\right) \\ &= 2aX_k(t_k) - a\kappa^p + b, \qquad (7.4)\\ \frac{\partial^2}{\partial (X_k(t_k))^2} f_k\left(X_k(t_k)\right) &= 2a. \qquad (7.5) \end{aligned}$$

Figures 7.16 and 7.17 show these daily polynomial sensitivities for the flexible and the unflexible portfolio without reserve requirements and subject to positive reserve requirements of 0.3 and 1.0 in comparison to full rolling intrinsic results. Obviously

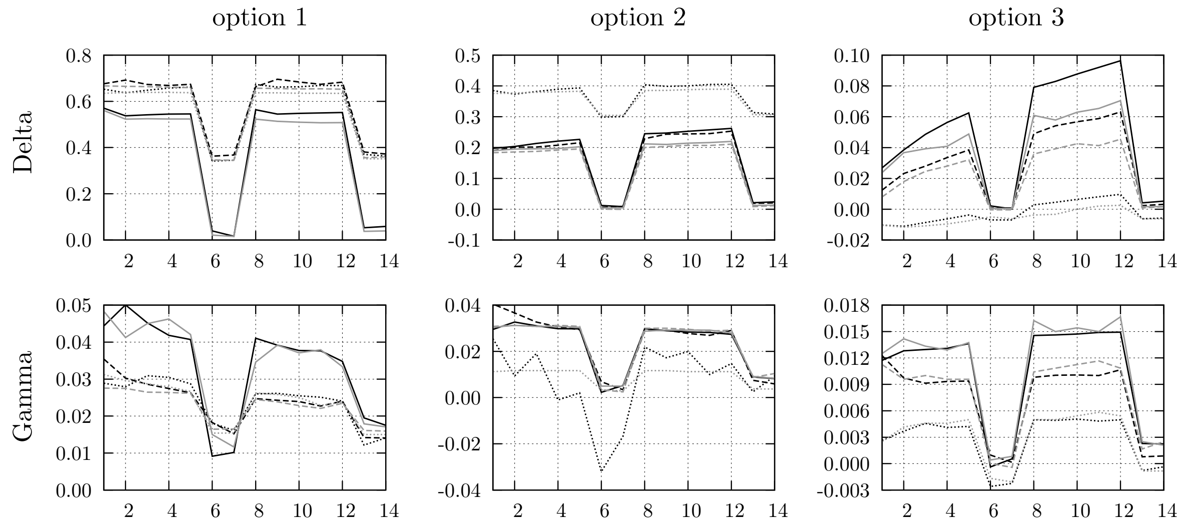

Fig. 7.16: Daily results of the *flexible* option portfolio based on full rolling intrinsic (black lines) in comparison to polynomial sensitivities (gray lines) without reserve requirement (solid), subject to $Res_{pos} = 0.3$ (dashed) and subject to $Res_{pos} = 1.0$ (dotted).

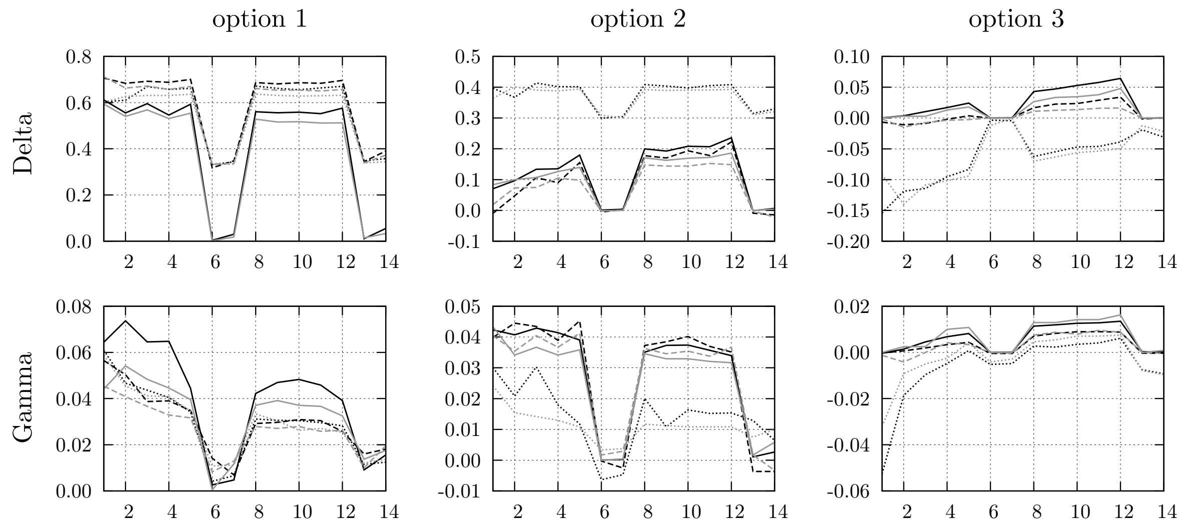

Fig. 7.17: Daily results of the *unflexible* option portfolio based on full rolling intrinsic (black lines) in comparison to polynomial sensitivities (gray lines) without reserve requirement (solid), subject to $Res_{pos} = 0.3$ (dashed) and subject to $Res_{pos} = 1.0$ (dotted).

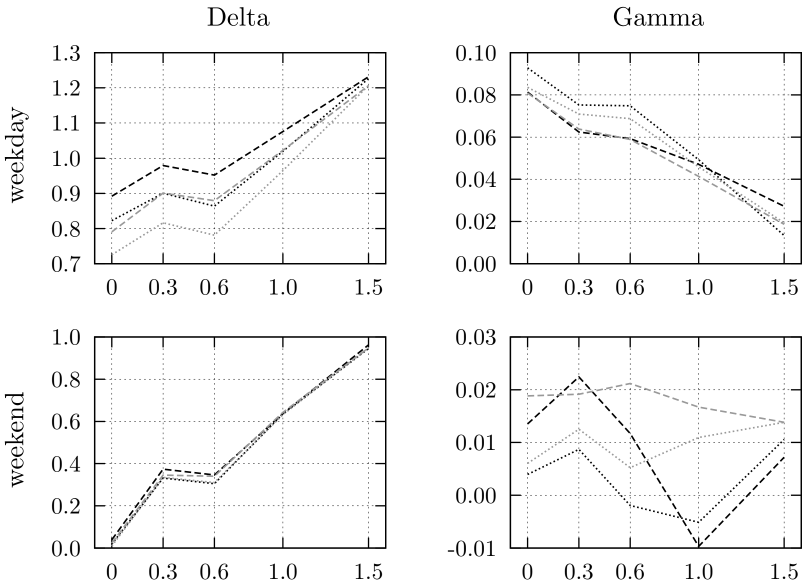

Fig. 7.18: Dependency of average weekday and weekend results on the level of applied reserve requirement based on full rolling intrinsic (black lines) in comparison to polynomial sensitivities (gray lines) for the *flexible* portfolio (dashed) and the *unflexible* portfolio (dotted).

the polynomial sensitivities replicate the temporal structure surprisingly well, although they are slightly downwards biased. Remarkably this rough semi-analytical approach is even able to cover the extreme characteristics of Delta and Gamma of option 3 of the unflexible portfolio in case of a high positive reserve requirement ($Res_{pos} = 1.0$).

The general bias becomes more distinct in a visualisation of average results rather than daily figures as shown in figure 7.18. The reason for this bias lies in the polynomial approach that relies on the simplifying expression (7.2) for the daily payoff f_k which also affects the calculation of sensitivities. Instead of this direct linear dependency on D_k, f_k actually depends on hourly dispatches d_k^l of all hours l of day k and the associated hourly prices $X_k(t_k) \cdot \alpha_k^l$ defining the hourly payoffs:

$$f_k = \sum_{l=1}^{24} d_k^l \left(X_k(t_k) \cdot \alpha_k^l - \kappa^p \right). \tag{7.6}$$

In addition Δ_k does not only depend on f_k but rather on the entire expected payoff of the valuation tenor $\sum_{j=1}^m f_j$:

$$\begin{aligned} \Delta_k &= \frac{\partial}{\partial X_k(t_0)} \mathbb{E} \left[\sum_{j=1}^{m} \sum_{l=1}^{24} d_j^l \left(X_j(t_j) \cdot \alpha_j^l - \kappa^p \right) \right] \\ &= \frac{\partial}{\partial X_k(t_0)} \mathbb{E} \left[\sum_{j=1}^{m} \left(X_j(t_j) \sum_{l=1}^{24} d_j^l \alpha_j^l - \kappa^p \sum_{l=1}^{24} d_j^l \right) \right]. \end{aligned} \tag{7.7}$$

The last term of (7.7) is identical to the daily dispatch, i.e. $D_j = \sum_{l=1}^{24} d_j^l$, but in addition the expression depends on a weighted sum of hourly dispatches $\sum_{l=1}^{24} d_j^l \alpha_j^l$.

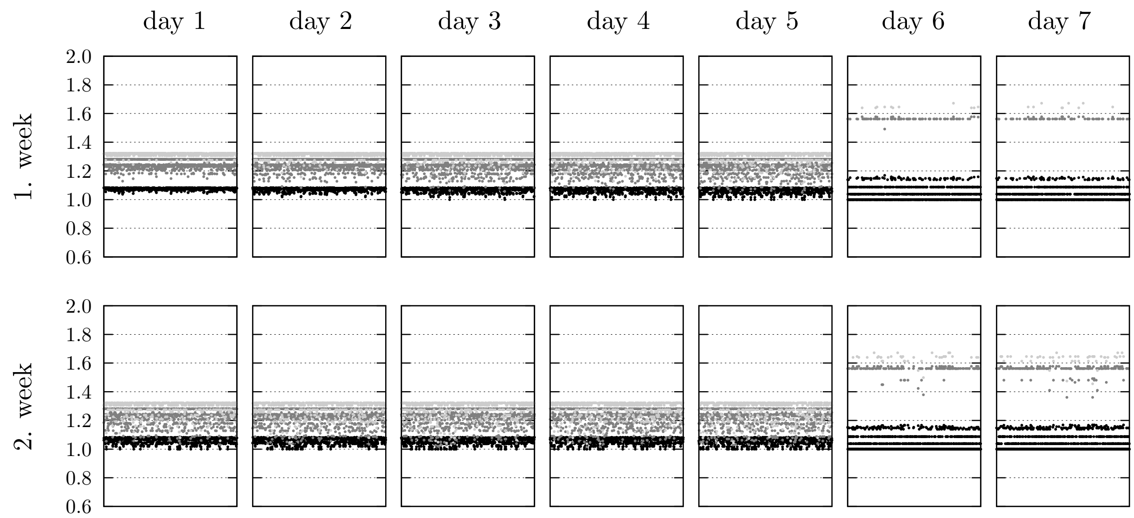

Fig. 7.19: 1000 realisations of $\alpha_{j,i}$ for all 14 days of the valuation tenor and all options (black: option 1, dark gray: option 2, light gray: options 3) of the *flexible* portfolio subject to minor reserve constraints ($Res_{pos} = 0.3$). Compare to figure 7.15.

The formula could be much simplified if our data gave evidence about a linearity between this weighted sum and the cumulated daily dispatch D_j, i.e.

$$\sum_{l=1}^{24} d_j^l \alpha_j^l \approx \alpha_j D_j. \tag{7.8}$$

Here the daily coefficients α_j are estimated via

$$\alpha_j = \frac{1}{n} \sum_{i=1}^{n} \frac{1}{D_{j,i}} \left(\sum_{l=1}^{24} d_{j,i}^l \alpha_{j,i}^l \right) := \frac{1}{n} \sum_{i=1}^{n} \alpha_{j,i}. \tag{7.9}$$

Figure 7.19 provides exemplary realisations of $\alpha_{j,i}$ for all options of the flexible option portfolio subject to minor reserve requirements of $Res_{pos} = 0.3$. Evidently these numerical results show a strong vertical clustering of $\alpha_{j,i}$ which supports the validity of assumption (7.8). Hence it is possible to write expression (7.7) in a more comprehensive form:

$$\Delta_k \approx \frac{\partial}{\partial X_k(t_0)} \mathbb{E} \left[\sum_{j=1}^{m} \left(X_j(t_j) \alpha_j D_j - \kappa^p D_j \right) \right]. \tag{7.10}$$

For exemplary reasons we stick to the above approximation of $D_j = aX_j(t_j) + b$

leading to the following analytical expressions of the sensitivities Δ_k and Γ_k:

$$\begin{aligned}
\Delta_k &\approx \frac{\partial}{\partial X_k(t_0)}\mathbb{E}\left[\sum_{j=1}^{m}\left(X_j(t_j)\alpha_j\left(aX_j(t_j)+b\right)-\kappa^p\left(aX_j(t_j)+b\right)\right)\right] \\
&= \frac{\partial}{\partial X_k(t_0)}\mathbb{E}\left[\sum_{j=1}^{m}\left(\alpha_j a\left(X_j(t_j)\right)^2+\left(\alpha_j b-\kappa^p a\right)X_j(t_j)-\kappa^p b\right)\right] \\
&\approx \mathbb{E}\left[2\alpha_k a X_k(t_k)\frac{\partial X_k(t_k)}{\partial X_k(t_0)}+\left(\alpha_k b-\kappa^p a\right)\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right] \\
&= \mathbb{E}\left[\left(2\alpha_k a X_k(t_k)+\alpha_k b-\kappa^p a\right)\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right] \qquad (7.11)\\
\Gamma_k &= \frac{\partial}{\partial X_k(t_0)}\Delta_k=\mathbb{E}\left[2\alpha_k a\left(\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right)^2\right], \qquad (7.12)
\end{aligned}$$

with $\frac{\partial^2}{\partial(X_k(t_0))^2}X_k(t_k) = 0$. The structure of these expressions is very close to the polynomial sensitivities (7.4) explaining the good replication of the structure of daily results as shown in figures 7.16 and 7.17. However the existence of quadratic and mixed terms of $O\left(X_k(t_k)\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right)$ and the factor α_k lead to structural differences between both expressions, whereby the absolute impact of both disturbances is of different size.

An estimation of the impact of the mixed terms is possible by elaborating the expressions of the sensitivities and taking advantage of main features of the implemented forward price process:

$$\begin{aligned}
\Delta_k &= 2\alpha_k a\mathbb{E}\left[X_k(t_k)\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right]+\left(\alpha_k b-\kappa^p a\right)\underbrace{\mathbb{E}\left[\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right]}_{=1} \\
&= 2\alpha_k a\mathbb{E}\left[X_k(t_0)\left(\exp\left(-\frac{1}{2}\left(\sigma_k^*(t_0)\right)^2 t_k+\sigma_k^*(t_0)\sqrt{t_k}dx\right)\right)^2\right]+\left(\alpha_k b-\kappa^p a\right) \\
&= 2\alpha_k a X_k(t_0)\exp\left(\left(\sigma_k^*(t_0)\right)^2 t_k\right)+\left(\alpha_k b-\kappa^p a\right) \qquad (7.13)\\
\Gamma_k &= 2\alpha_k a\mathbb{E}\left[\left(\frac{\partial X_k(t_k)}{\partial X_k(t_0)}\right)^2\right]=2\alpha_k a\exp\left(\left(\sigma_k^*(t_0)\right)^2 t_k\right). \qquad (7.14)
\end{aligned}$$

Here $dx \sim N(0,1)$. In the context of our chosen volatility term structure (and $t \leq 14/365$) the following assessment holds:

$$\|\exp\left(\left(\sigma_k^*(t_0)\right)^2 t_k\right)-1\| < 3\%. \qquad (7.15)$$

Therefore the implied negative bias of the mixed terms will be rather small.

The impact of the hourly dispatch distribution and therewith the hourly price structure is more significant. Figure 7.19 indicates that typically

$$\|\alpha_k - 1\| < 30\%. \qquad (7.16)$$

α_k is always positive and will lead to an additional and more significant negative bias than the mixed terms.

The combination of both effects is well able to explain the observed bias in general. However the differences are heavily dependent on the model setup and a professional usage of polynomial sensitivities would require a detailed error analysis whenever a minor model variation is implemented. The method is very interesting on the one hand side as it provides a rough estimation of the general structure of sensitivities in extremely short time and sometimes with surprisingly well accuracy. On the other hand the method provides biased results which never leads to a great level of confidence. It can be applied wherever all interdependencies within a model setup are understood in detail and when practitioners clearly understand the size of the bias and are able to correct it accordingly. In case of quickly changing model parameters or highly complex and intransparent problems this simplified approach cannot be recommended.

8 Concluding remark

A proper knowledge of values and sensitivities of both first and second order (Δ and Γ) of their generation assets with respect to the market price of electricity is mandatory for utilities in order to perform adequate risk management and to decide about portfolio adjustments. In this thesis we achieve to derive values and Monte-Carlo sensitivities Δ and Γ for a variety of stylised electricity generation portfolios by utilizing the numerical advantages of the Proxy Simulation Scheme method. We are able to analyse the impact of increasing external reserve requirements on the sensitivities of all portfolios, thereby providing an indirect evaluation of both value and the Δ–Γ risk profile of reserve contracts in different portfolio contexts.

Our results – in combination with the provided detailed backtesting – can serve as a solid starting point for further applications of the Proxy Simulation Scheme method also in the more energy focussed academic community. We see a particularly large potential in an extension of our work in the following directions.

Technical details In this thesis we focus on the evaluation of stylised power plant options, constrained by minimum up- and down-times and minimum and maximum loads. This allows to replicate key dispatch characteristics of real power plants but is only a selection of the entire range of technical constraints of physical generation assets. It would be interesting and highly relevant for practitioners to extend the level of technical details also towards ramp rates, load dependent efficiencies or start-up costs (to name a few) and investigate associated effects.

Portfolio size We present the impact of a simple portfolio fragmentation on our results. However physical generation portfolios can include a significantly larger number of assets with a broad variation of technical constraints and profitability, motivating an enlarged portfolio study with the aim to analyse whether the observed portfolio effect is scalable.

Price shape We apply hourly prices of electricity that are heavily affected by a strong photovoltaic energy production profile, hence showing comparably low midday prices and high prices in shoulder hours. It would be interesting to define generic future price shapes for different European countries and compare the associated impact on both value and sensitivities of reserve requirements in order to assess opportunities of highly flexible assets in the different markets.

Multi-commodity setup Our analyses base on a single-commodity setup with the market price of electricity being the only source of risk. This does only allow to derive electricity price Δ and Γ which is effectively close to the widely accepted industry approach of applying ceteris-paribus sensitivities. However by expanding the model to a multi–commodity setup it would also be possible to analyse fuel or spread sensitivities and compare their hedge efficiencies.

Valuation tenor and granularity We focus on daily results for a valuation tenor up to 14 days, thereby covering only a small part of the typical hedging period of physical generation portfolios which usually includes all sufficiently liquid electricity and fuel forward products. In addition the Proxy Simulation Scheme method can easily be used for longer or shorter valuation tenors and forward price curves with higher or lower granularity – thus opening an entire field of research ranging from quarter hour sensitivities in the intraday electricity market to seasonal fuel sensitivities of generation portfolios.

In any case we would highly recommend to implement the Proxy Simulation Scheme method for the regular evaluation of a small physical generation portfolio in order to receive historic backtesting data and to test the method in the real market environment.

List of Tables

List of Figures

Bibliography

Arbeitsgruppe Erneuerbare Energien - Statistik (AGEE-Stat). Zeitreihen zur Entwicklung der erneuerbaren Energien in Deutschland, 2013.

Bauwens, L., Hafner, C. M., and Pierret, D. Multivariate Volatility Modeling of Electricity Futures. *Journal of Applied Econometrics*, 28(5), 2012.

Benhamou, E. Optimal Malliavin Weighting Function for the Computation of the Greeks. *Mathematical Finance*, 13(1):37 – 53, 2003.

Benth, F. E. and Koekebakker, S. Stochastic Modelling of Financial Electricity Contracts. *Energy Economics*, 30, 2008.

Benth, F. E. and Saltyte-Benth, J. Analytical Approximation for the Price Dynamics of Spark Spread Options. *Studies in Nonlinear Dynamics & Econometrics, ISSN 1558-3708*, 10.2006(3), 2006.

Benth, F. E., Dahl, L. O., and Karlsen, K. H. Quasi Monte Carlo evaluation of sensitivities of options in commodity and energy markets. *International Journal of Theoretical and Applied Finance*, 2003.

Bingham, N. and Kiesel, R. *Risk-Neutral Valuation, Pricing and Hedging of Financial Derivatives.* Springer Verlag, Heidelberg, 2004.

Bjerksund, P., Stensland, G., and Vagstad, F. Gas Storage Valuation: Price Modelling v. Optimization Methods. *The Energy Journal*, 32(1), 2011.

Black, F. and Scholes, M. The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy*, 81(3), 1973.

BMWI. Erneuerbare-Energien-Gesetz (EEG), 2012.

Boogert, A. and de Jong, C. Gas storage valuation using a Monte Carlo method. Working paper, Birkbeck, Department of Economics, Mathematics & Statistics, 2007.

Borchert, J. and Hasenbeck, M. Bewertung und Steuerung von Kraftwerksscheiben. *ZfE Zeitschrift für Energiewirtschaft*, 2009.

Borel, E. Les probabilités dénombrables et leurs applications arithmétiques. *Rendiconti del Circolo Matematico di Palermo*, 27, 1909.

Börger, R. *Energy-Related Commodity Futures.* PhD thesis, Universität Ulm, 2007.

Börger, R. Explicit Approximations of multi-asset option prices including Greeks. *International Journal of Portfolio Analysis and Management*, 1(4/2014):314–329, 2014.

Börger, R. and van Heys, J. Calibration of the Libor Market Model Using Correlations Implied by CMS Spread Options. *Applied Mathematical Finance*, 17(5), 2010.

Boyle, P., Broadie, M., and Glasserman, P. Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control*, 21:1267–1321, 1997.

Brajkovic, J. Real options approach to investment in base load coal fired plant. Research paper, University of Southampton, Energy Institute Hrvoje Pozar, 2010.

Breslin, J., Clewlow, L., Elbert, T., Kwok, C., Strickland, C., and van der Zee, D. Gas storage: rolling intrinsic valuation. *energy risk*, pages 61–65, 2009.

Brigo, D. A Note on Correlation and Rank Reduction, 2002.

Broadie, M. and Glasserman, P. Estimating Security Price Derivatives Using Simulation. *Management Science*, 42(2):269 – 285, 1996.

BSW. Entwicklung des deutschen PV-Marktes, 2013.

Bundesministerium für Wirtschaft und Energie (BMWi). Zweiter Monitoring-Bericht "Energie der Zukunft". 2014.

Burger, M., Graeber, B., and Schindlmayr, G. *Managing energy risk, An Integrated View on Power and Other Energy Markets.* John Wiley & Sons, West Sussex, 2007.

Capriotti, L. and Giles, M. Fast correlation Greeks by adjoint algorithmic differentiation. *Risk*, pages 79 – 83, 2010.

Capriotti, L. and Giles, M. Adjoint Greeks made easy. *Risk*, pages 96 – 102, 2012.

Carmona, R. and Durrleman, V. Pricing and Hedging Spread Options. *SIAM Review*, 45(4):627–685, 2003.

Carr, P., Madan, D. B., and Smith, R. H. Option Valuation Using the Fast Fourier Transform. *Journal of Computational Finance*, 2:61–73, 1999.

Chan, H. J. and Joshi, M. Minimal partial proxy simulation schemes for generic and robust Monte-Carlo Greeks. *Journal of Computational Finance*, 15(2):77 – 109, 2011.

Chan, H. J. and Joshi, M. Optimal limit methods for computing sensitivities of discontinuous integrals including triggerable derivative securities. Working paper, Available at SSRN: http://ssrn.com/abstract=2011690 or http://dx.doi.org/10.2139/ssrn.2011690, 2012a.

Chan, H. J. and Joshi, M. Fast Monte-Carlo Greeks for Financial Products with Discontinuous Pay-Offs. *Mathematical Finance*, 2012b.

Chen, N. and Glasserman, P. Maliavin Greeks without Malliavin calculus. *Stochastic Processes and their Applications*, 117:1689 – 1723, 2007.

Clewlow, L. and Strickland, C. Valuing Energy Options in a One Factor Model Fitted to Forward Prices. *Available at SSRN: http://ssrn.com/abstract=160608*, 1999.

Clewlow, L., Strickland, C., University of Technology, S. S. o. F., Economics, and Group, Q. F. R. *A Multi-factor Model for Energy Derivatives.* Research paper. School of Finance and Economics, University of Technology, Sydney, 1999.

De Vos, K. and Driesen, J. Impact of wind power reserve requirements on system operation and costs. Conference publication, 12. IAEE European Energy Conference, 2012.

Deng, S. and Xia, Z. Pricing and Hedging Electricity Supply Contracts: a Case with Tolling Agreements. Technical report, School of Industrial and Systems Engineering, Georgia Institute of Technology, 2005.

Deng, S., Johnson, B., and Sogomonian, A. Exotic electricity options and the valuation of electricity generation and transmission assets. In *PROCEEDINGS OF THE CHICAGO RISK MANAGEMENT CONFERENCE*, pages 383–392, 2001.

Deng, S.-J., Li, M., and Zhoc, J. Closed-form Approximations for Spread Option Prices and Greeks. *The Journal of Derivatives*, 15(3):58–80, 2008.

Deutsch, H.-P. *Derivate und Interne Modelle – Modernes Risikomanagement.* Schäffer-Poeschel Verlag, Stuttgart, 2004.

Deutscher Bundestag. Bericht des Ausschusses für Bildung, Forschung und Technikfolgenabschätzung: Regenerative Energieträger zur Sicherung der Grundlast in der Stromversorgung. 2012.

Eichhorn, A., Heitsch, H., and Römisch, W. *Stochastic Optimization of Electricity Portfolios: Scenario Tree Modeling and Risk Management.* Handbook of Power Systems II,. Springer-Verlang Berlin Heidelberg, 2010.

EPEX SPOT SE. www.epexspot.com, 2014.

Faria, E. and Fleten, S. *Day-Ahead Market Bidding for a Nordic Hydropower Producer: Taking the Elbas Market into Account.* Computational Management Science. Springer-Verlag Berlin Heidelberg, 2011.

Fournié, E., Lasry, J.-M., Lebuchoux, J., Lions, P.-L., and Touzi, N. Application of Malliavin calculus to Monte Carlo methods in finance. *Finance and Stochastics*, 3: 391 – 412, 1999.

Fries, C. Proxy Simulation Schemes for Generic Robust Monte-Carlo Sensitivities and High Accuracy Drift Approximation with Applications to the LIBOR Market Model, 2006.

Fries, C. *Mathematical Finance: Theory, Modeling, Implementation.* John Wiley & Sons, West Sussex, 2007a.

Fries, C. P. Bumping the Model, Generic robust Monte-Carlo Sensitivities using the Proxy Simulation Scheme Method. 2005.

Fries, C. P. Localized Proxy Simulation Schemes for Generic and Robust Monte-Carlo Greeks. 2007b.

Fries, C. P. and Joshi, M. S. Partial Proxy Simulation Schemes for Generic and Robust Monte-Carlo Greeks. *The Journal of Computational Finance*, 11(3):79 – 106, 2008.

Fries, C. P. and Kampen, J. Proxy Simulation Schemes for generic robust Monte-Carlo sensitivities, process oriented importance sampling and high accuracy drift approximation (with applications to the LIBOR Market Model). *The Journal of Computational Finance*, 10(2):97 – 128, 2007.

Fu, H., Jin, X., Pan, G., and Yang, Y. Estimating multiple option Greeks simultaneously using random parameter regression. *Journal of Computational Finance*, 16 (2), 2012.

Gerstner, T. and Kloeden, P. Interdisciplinary Mathematical Sciences. In *Recent Developments in Computational Finance – Foundations, Algorithms and Applications*, volume 14. World Scientific Publishing Co. Pte. Ltd., 2013.

Giles, M. Vibrato Monte Carlo sensitivities. In *Monte Carlo and Quasi-Monte Carlo Methods*, New York, 2008. Springer.

Giles, M. Monte Carlo evaluation of sensitivities in computational finance. In E.A. Lipitakis, editor, HERCMA 2007 Conference Proceedings, 2007.

Giles, M. and Glasserman, P. Smoking adjoints: fast Monte Carlo Greeks. *RISK*, pages 88 – 92, 2006.

Glantz, M. H. *Currents of Change, Impacts of El Niño and La Niña on Climate and Society.* Cambridge University Press, Cambridge, 2001.

Glasserman, P. *Monte Carlo Methods in Finance Engineering (Stochastic Modelling and Applied Probability).* Springer, New York, 2003.

GLPK. GNU Linear Programming Kit, available under http://www.gnu.org/software/glpk/, 2014.

Grasselli, M. and Pelinovsky, D. *Numerical Mathematics.* Jones and Bartlett Publishers Canada, Mississauga, 2008.

Härdle, W. K. and Trück, S. The dynamics of hourly electricity prices. *SFB 649 Discussion Paper 2010-013*, 2010.

Hauser, E. Kurzfristige Effekte der PV-Einspeisung auf den Großhandelsstrompreis, 2012.

Hurd, T. and Zhou, Z. A Fourier transform method for spread option pricing. *SIAM Journal on Financial Mathematics*, 1(1):142–157, 2010.

Jäckel, P. *Monte Carlo Methods in Finance.* John Wiley & Sons, Ltd., West Sussex, 2003.

Jorion, P. *Financial Risk Manager Handbook.* John Wiley & Sons, New Jersey, 2011.

Kelly, J. GNU gfortran GLPK language bindings, available under http://en.wikibooks.org/wiki/GLPK/Fortran, 2014.

Kroese, D. P., Taimre, T., and Botev, Z. I. *Handbook of Monte Carlo Methods.* John Wiley & Sons, New Jersey, 2011.

Longstaff, F. A. and Schwartz, E. S. Valuing American options by simulation: A simple least-squares approach. *Review of Financial Studies*, pages 113–147, 2001.

Lutz, M. Extracting Correlations from the Market: New correlation parameterizations and the Calibration of a Stochastic Volatility LMM to CMS Spread Options. *Available at SSRN: http://ssrn.com/abstract=1620356 or http://dx.doi.org/10.2139/ssrn.1620356* , 2010.

Lyuu, Y.-D. and Teng, H.-W. Unbiased and efficient Greeks of financial options. *Finance and Stochastics*, 15(1):141 – 181, 2010.

Mauritzen, J. What happens when it's windy in Denmark? An Empirical Analysis of Wind Power on Price Volatility in the Nordic Electricity Market. *Available at SSRN: http://ssrn.com/abstract=1754931 or http://dx.doi.org/10.2139/ssrn.1754931* , 2010.

Montero, M. and Kohatsu-Higa, A. Malliavin Calculus applied to Finance. *Physica A: Statistical Mechanics and its Applications*, 320(C):548–570, 2003.

NASDAQ OMX. www.nasdaqomx.com, 2014.

Øksendal, B. *Stochastic Differential Equations, An Introduction with Applications.* Springer, Heidelberg, 2007.

Palmintier, B. and Webster, M. *Impact of Unit Commitment Constraints on Generation Expansion Planning with Renewables.* Published in: Power and Energy Society General Meeting, 2011 IEEE, 2011.

Papatheodorou, B. Enhanced Monte Carlo Methods for Pricing and Hedging Exotic Options. Master's thesis, Pembroke College, 2005.

Pardo, M. *Pricing daily options in the German power market.* PhD thesis, Universidad Pontificia Comillas, ICAI, Madrid, 2005.

pvXchange. Preisindex, www.pvxchange.com, 2014.

Rebennack, S., Iliadis, N., Kallrath, J., and Pardalos, P. M. *Short term portfolio optimization for discrete power plant dispatching.* Conference publication, Power & Energy Society General Meeting. IEEE, 2009.

Rebennack, S., Kallrath, J., and Pardalos, P. M. *Energy Portfolio Optimization for Electric Utilities: Case Study for Germany.* Energy, Natural Resources and Environmental Economies. Springer-Verlang Berlin Heidelberg, 2010.

regelleistung.net. Internetplattform zur Vergabe von Regelleistung, available under www.regelleistung.net, 2014.

Riley, K. F., Hobson, M. P., and Bence, S. J. *Mathematical Methods for Physics and Engineering.* Cambridge University Press, Cambridge, 2006.

Schoenmakers, J. and Coffey, B. Systematic generation of parametric correlation structures for the LIBOR market model. *International Journal of Theoretical and Applied Finance*, 6(5), 2003.

Schwartz, E. S. The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *The Journal of Finance*, 52(3):923–973, 1997.

Schwartz, E. S. and Smith, J. E. Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Management Science*, 46(7):893–911, 2000.

Shapiro, A., Dentcheva, D., and Ruszczyński, A. *Lectures on stochastic programming, Modeling and Theory.* Mathematical Programming Society Philadelphia, Philadelphia, 2009.

Swider, D. J. *Handel an Regelenergie- und Spotmärkten – Methoden zur Entscheidungsunterstützung für Netz- und Kraftwerksbetreiber.* Deutscher Universitäts-Verlag, Stuttgart, 2006.

Tseng, C. and Barz, G. Short-term generation asset valuation: a real options approach. *Operations Research*, 50(2):297–310, 2002.

Vázquez, M. *Optimizing the Spinning Reserve Requirements.* PhD thesis, University of Manchester, 2006.

Vázquez, M., Barquín, J., and Batlle, C. Electricity forward and volatility curves computation based on Monte Carlo simulation. *9th International Conference on Probabilistic Methods Applied to Power Systems*, 2006.

Vázquez, M., Sánchez-Úbeda, E., Berzosa, A., and Barquín, J. Short-term evolution of forward curves and volatility in illiquid power markets. *Online at http://mpra.ub.uni-muenchen.de/9448/*, 2008.

Vicente-Serrano, S. M. and Trigo, R. M. *Hydrological, Socioeconomic and Ecological Impacts of the North Atlantic Oscillation in the Mediterranean Region.* Springer-Verlag, Heidelberg, 2011.

Wessa, P. Kernel Density Estimation (v1.0.11) in Free Statistics Software (v1.1.23-r7), Office for Research Development and Eduction, URL http://www.wessa.net/rwasp_density.wasp/, 2012.

Winkler, J., Sensfuß, F., Keles, D., Renz, L., and Fichtner, W. *Perspektiven für die langfristige Entwicklung der Strommärkte und der Förderung Erneuerbarer Energien bei ambitionierten Ausbauzielen.* Fraunhofer ISI, 2013.

Wirth, H. *Aktuelle Fakten zur Photovoltaik in Deutschland.* Fraunhofer ISE, 2014.

Yan, H., Luh, P. B., and Guan, X. Scheduling of Hydrothermal Power Systems. Conference publication, American Control Conference, 1994.

Danksagung

Zuallererst möchte ich mich ganz herzlich bei meinem Doktorvater Professor Dr. Rüdiger Kiesel bedanken, der diese Arbeit durch seine Unterstützung erst möglich gemacht hat. Ich habe den vertrauensvollen Kontakt zu ihm sehr geschätzt und möchte ihm an dieser Stelle noch einmal danken, dass er mir auf der EnergyFinance Conference 2013 ein erstes Forum zu Präsentation meiner Ergebnisse bot.

Mein besonderer Dank gilt meinem Kollegen und Freund Dr. Reik Börger, der mich nicht nur anfangs in meinem Entschluss zu promovieren bestärkt hat, sondern mir auch während der gesamten Zeit als interessierter und motivierender Ansprechpartner zuverlässig zur Seite stand.

Ich möchte mich außerdem bei allen Kollegen bedanken, die mir in zahlreichen guten Gesprächen Anregungen und Bestätigung gaben. Hier sind stellvertretend Dr. Uwe Fermum und Dr. Richard Kiefer zu nennen, die mich zudem noch bei der Textdurchsicht unterstützt haben.

Mein größter Dank geht jedoch an meine Frau Amrei, welche die Erstellung dieser Dissertation enger und mit mehr Einsatz als jeder andere begleitet hat und deren bedingungsloser Unterstützung ich mir zu jedem Zeitpunkt sicher sein konnte.