Noname manuscript No.
(will be inserted by the editor)

A Rolling Optimisation Model of the UK natural gas market

Mel T. Devine · James P. Gleeson · John Kinsella · David M. Ramsey

Received: date / Accepted: date

Abstract Daily gas demand in the UK is variable. This is partly due to weather patterns and the changing nature of electricity markets, where intermittent wind energy levels lead to variations in the demand for gas needed to produce electricity. This uncertainty makes it difficult for traders in the market to analyse the market. As a result, there is an increasing need for models of the UK natural gas market that include stochastic demand. In this paper, a Rolling Optimisation Model (ROM) of the UK natural gas market is introduced. It takes as an input stochastically generated scenarios of demand. The outputs of ROM are the flows of gas, i.e., how the different sources of supply meet demand, as well as how gas flows in to and out of gas storage facilities. The outputs also include the daily System Average Price of gas in the UK. The model was found to fit reasonably well to historic data (from the UK National Grid) for the years starting on the 1st of April for both 2010 and 2011. These results allow ROM to be used to predict future flows and prices of gas and to investigate various stress-test scenarios in the UK natural gas market.

Keywords Rolling optimisation, UK natural gas market, stochastic demand scenarios

1 Introduction

Traditionally, the main driver of UK natural gas demand has been temperature. However, the British government has committed to having 15% of its energy generated by renewable sources of energy in 2020 with wind energy being seen as a major contributor to this renewable energy (DECC, 2011;

M. T. Devine
Mathematics Applications Consortium for Science and Industry (MACSI), Department of Mathematics and Statistics, University of Limerick, Limerick Ireland.
email: mel.devine@ul.ie

J. P. Gleeson
MACSI, Department of Mathematics and Statistics, University of Limerick, Limerick Ireland.
email: james.gleeson@ul.ie

J. Kinsella
Department of Mathematics and Statistics, University of Limerick, Limerick Ireland.
email: john.kinsella@ul.ie

D. M. Ramsey
Department of Mathematics and Statistics, University of Limerick, Limerick Ireland.
email: david.ramsey@ul.ie

Table 1 List of acronyms

UK	United Kingdom
ROM	Rolling Optimisation Model
SND	Seasonal Normal Demand
UKCS	United Kingdom Continental Shelf
LNG	Liquified Natural Gas
BBL	Balgzand Bacton Line
IUK	Interconnector United Kingdom
LRS	Long-Range-Storage
MRS	Medium-Range-Storage
SRS	Short-Range-Storage
SAP	System Average Price
GAMS	Generic Algebraic Modeling System
mcm	Million Cubic Meters
MAPE	Mean Absolute Percentage Error

DUKES, 2011). This is leading to increased amounts of electricity generated by wind energy across the UK (DUKES, 2011).

Natural gas is the single largest source of energy used to generate electricity in the UK (DUKES, 2011). As wind energy is an intermittent source of energy, the increased levels of wind power in electricity markets has led to, and will continue to cause, variability in the amount of gas required to generate electricity. This is because when wind energy is available to be harnessed for electricity power generation, gas is not needed as much. However, when wind energy is not available, demand for gas to generate electricity is increased. With the number of wind farms increasing, along with the volatile nature of weather, this has led to an increase in uncertainty in the demand for natural gas in the UK. This paper describes a simplified model of the UK natural gas market that incorporates the stochastic nature of demand. To our knowledge, a mathematical model of this particular market has not previously appeared in the academic literature.

There are two main types of natural gas market models found in the literature. The first are complementarity-based equilibrium models. For example, following earlier in Gabriel et al (2003), Gabriel et al (2005a) established existence and uniqueness results for their complementarity-based equilibrium model. They then applied this model to the North American natural gas market (Gabriel et al, 2005b). Similarly, Egging et al (2010), describe a global gas market model which is based on Gabriel et al (2005a). Another global gas market model is the World Gas Trade Model, (The Baker Institute World Gas Trade Model, 2005). Other examples of complementarity-based equilibrium models of natural gas market include the GASTALE model (Boots et al, 2004; Egging and Gabriel, 2006; Lise and Hobbs, 2008), as well as the GASMOD model (Holz et al, 2008). Both of these models are applied to the European gas market, while (Abada et al, 2012)'s GaMMES model is applied to the Northwestern European Natural Gas Market. In Abrell and Weigt (2012) a mixed-complementarity based model is used to study the interactions between Europe's electricity and natural gas markets. Because complementarity-based equilibrium models have multiple objective functions, Nash-Cournot competition can be incorporated. This allows for imperfect competition amongst the different players, as well as endogenously determined demand. A recent extension of this type of modelling can be found in Siddiqui and Gabriel (2012).

The second type of natural gas market model seen in the literature is based on cost minimisation. There are a number of previous models that have used this approach. For example, Bopp et al (1996) and Butler and Dyer (1999) developed cost minimisation models for optimising the flows of gas. Other examples include the EUGAS model (Perner, 2002) and the MAGELAN model (Seeliger, 2006). These two models are used to model the European and global natural gas markets, in Perner and Seeliger (2004) and Lochner and Bothe (2009), respectively. Unlike complementarity-based equilibrium models, Nash-Cournot competition cannot be included in the formulation of a cost minimisation model, as there is only one objective function. As a result, these models assume that none of the players have market power and that demand is determined exogenously.

It is shown by Devine (2012) that when Nash-Cournot competition (i.e., market power) is removed from a complementarity-based equilibrium model, it reduces to a simpler cost minimisation model. According to a report made by Ofgem[1] for the UK Parliament, the assumption of a fully competitive market is reasonable with respect to the UK natural gas market (HCBEC, 2007). They base this assumption on the fact that there are many players in the UK market and hence "there are limited signs of one source being able to influence market price to a significant extent". As a result, the model described in this paper does not include any form of market power; hence, a cost minimisation model is introduced.

The model developed here is referred to as the Rolling Optimisation Model and is introduced in Section 3. The outputs of the model are the flows of gas in the UK gas market, as well as the daily System Average Price (SAP); the average price of all gas traded on a given day in the UK gas market. The flows of gas detail the amount of gas in the different type of storage facilities seen in the UK, as well as how the different sources of supply meet demand. Descriptions of these storage facilities and sources of supply are given in Section 3.2. The goal of this work is to ensure, by estimating the parameters of the model, that these outputs are qualitatively similar to the corresponding actual flows and prices observed in the UK gas market. Once this is done, the model can be used to forecast future flows and prices as is done in Section 4.

ROM replicates the daily decisions made in the UK gas market under uncertain information about future demand. In order to capture this uncertainty, ROM takes as an input exogenously determined demand scenarios simulated from the stochastic process described in Section 2. This process is based on the statistical analysis of historical data supplied by the UK National Grid[2] and generates multiple demand time series, each of length D days. Each generated time series represents a demand scenario in this model. The demand on the first day (d_1) is the same for each of these time series. This means that the stochastic demand on the first day of the model is scenario-independent. For the next five days the values of demand show increasing uncertainty, i.e., variation from scenario to scenario. For more than five days ahead, only seasonal averages are known. ROM decides the amount of gas to be injected to (or withdrawn from) the different storage facilities of the model, as well as how the various sources of supply meet this scenario-independent demand on the first day (d_1). This is done at minimum cost, whilst also ensuring all possible future demands (over the set of scenarios) are also met at minimum expected cost.

Once the initial optimisation problem of ROM is solved, a new set of demand scenarios are developed using the process described in Section 2, but now with demand time series that have shifted forward by one day. For example, if the first set of demand time series used started on the 1st of April, then the second set would be generated with the same time series model but starting on the 2nd of April. The length (D) of the time series remains the same. Using these updated demand scenarios and the updated amount of gas in

[1]Ofgem (Office of the Gas and Electricity Markets) regulates the gas and electricity markets in the UK. See: http://www.ofgem.gov.uk/About%20us/Pages/AboutUsPage.aspx

[2]The UK National Grid is the owner and operator of national transmission system throughout Great Britain.

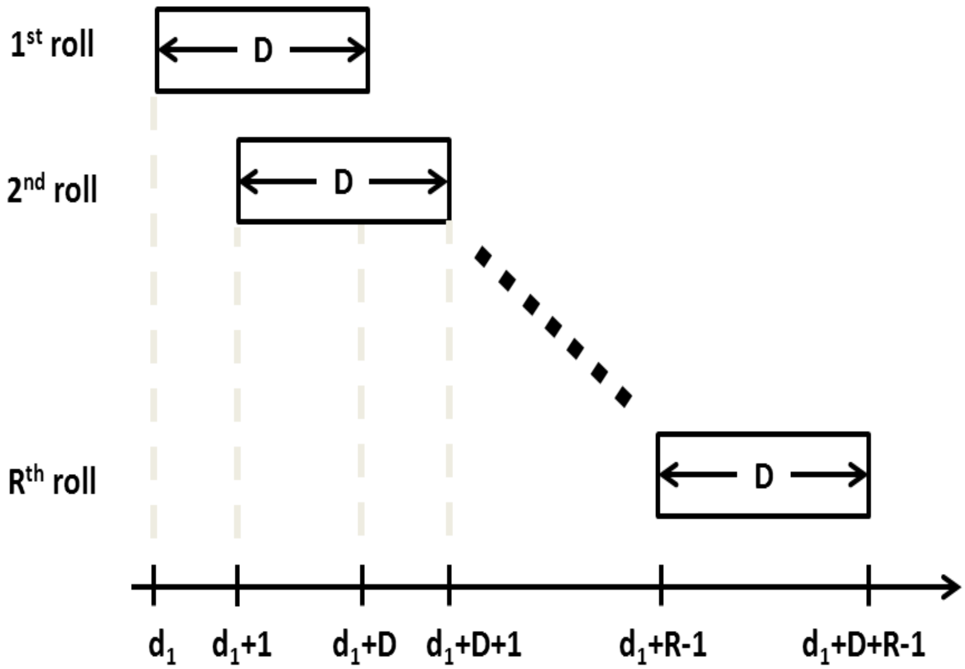

Fig. 1 Rolling horizon of the sets of demand scenarios used in ROM.

storage (obtained from the results of the previous optimisation), the optimisation problem is solved again so that the new scenario-independent demand on the first day (2nd of April) is now met at minimum cost, whilst again ensuring all possible future demands are met at minimum expected cost. Once this demand is met, the demand scenarios are updated again in a similar manner using demand series that have shifted forward one day once more. Following this, the optimisation problem is solved again and so on until R optimisations have been performed. Throughout this paper, solving one optimisation according to ROM is defined as one 'roll' of the model. Fig. 1 describes the rolling horizon of the sets of demand scenarios used in this model. This idea of using a rolling optimisation model to model a natural gas market has previously been unseen in the literature.

Each roll of the model represents a single day's decisions of how demand in the UK gas market is met. When one day's demand is met a new day arrives (i.e., the model moves forward to the next optimisation) where new decisions have to be made on how to meet demand again. It is envisaged that the model will be used by traders in the UK gas market on a day-to-day basis, whereby one roll of the model is solved each day. It is anticipated that the stochastic demand scenarios needed for this roll would be simulated using the daily demand information that is made available by the UK National Grid. This information includes actual demand for the given day, predicted demand for the subsequent five days ahead and assumed seasonal levels thereafter.

While ROM is specifically applied to the UK gas market in this work, it should be noted that in principle the model may be applied to any gas market where market power is not an important factor. The Generic Algebraic Modeling System (GAMS) (GAMS, 2012) is used for model development and programming. The remainder of this paper is organised as follows. Firstly, in Section 2, a stochastic process describing UK gas demand is developed. Secondly, the Rolling Optimisation Model is introduced in Section 3. Some of the applications of ROM are examined in Section 4, while the paper concludes in Section 5 with a summary and conclusions.

2 Stochastic process describing UK gas demand

In this section a stochastic process that captures the uncertainty of gas demand in the UK is introduced. It is supported by time series analysis of historic data on demand for natural gas in the UK (Devine, 2012). In particular, the following relationships were fitted to auto-regressive models of order one (AR(1) models):

1. The difference between the natural log of actual demand and the natural log of seasonal normal demand,
2. The difference between the natural log of actual demand and the natural log of predicted demand.

Actual demand is the time series of the daily figures published by the UK National Grid for system gas demand in the UK. Seasonal normal demand (SND) is the daily time series for an average gas demand year. In other words, it is the daily gas demand one would expect in an average year. It is calculated by the UK National Grid and is published approximately one year in advance[3]. Fig. 2 shows examples of actual demand and SND for the year starting in October 2008.

Predicted demand is the time series of the predicted daily demand given by the UK National Grid. There are five types of predicted demand supplied by the UK National Grid, one- to five-day ahead predictions. One-day ahead predictions are the daily demands predicted one day beforehand. Similarly two- to five-day ahead predictions are the daily demands predicted two to five days beforehand.

Consider the following discrete-time process with S scenarios:

$$Demand_{t,t} = ActDem_t, \tag{1}$$

$$\begin{aligned} \ln(Demand^s_{t,t+d}) = & \ln(ActDem_{t+1}) + \\ & \mu_d + \gamma_d(\ln(Demand^s_{t-1,t}) - \ln(ActDem_t)) + \sigma_d \varepsilon_t, \quad d = 1..5, \end{aligned} \tag{2}$$

$$\begin{aligned} \ln(Demand^s_{t,t+d}) = & \ln(SND_{t+d}) + \\ & \mu_d + \gamma_d(\ln(Demand^s_{t,t+d-1}) - \ln(SND_{t+d-1})) + \sigma_d \varepsilon_t, \quad \forall\ d > 5, \end{aligned} \tag{3}$$

where $Demand^s_{t,t+d}$ is the demand for gas on day $t+d$, simulated on day t and associated with scenario s. These simulated demands are used as inputs to the Rolling Optimisation Model described in detail later in this paper. Actual demand and SND on day t is represented by $ActDem_t$ and SND_t respectively. The noise term ε_t is a Gaussian white noise process with a mean of zero and variance of one, and the parameters are determined by fitting to historical data..

On the first day of the stochastic process, equation (1) states that the simulated demand ($Demand_{t,t}$) equals actual demand for that day. It also shows that the simulated demand is scenario-independent on this first day (note: there is no superscript s for this day). For the next five days, the simulated demand is actual demand for that day plus some randomly simulated noise with a dependence on the error in

[3] See: http://www.nationalgrid.com/uk/Gas/Data/misc/

Table 2 Parameters associated with equations (2) and (3).

d	μ_d	γ_d	σ_d^2
1	-0.007	0.22	0.0015
2	-0.009	0.28	0.0028
3	-0.010	0.37	0.0032
4	-0.012	0.65	0.0031
5	-0.012	0.68	0.0031
>5	-0.0027	0.92	0.0020

predicting demand for the previous day [4]. The parameters for equation (2) are given in Table 2 and were found from AR(1) models fitted to the differences between the natural log of actual demand and the natural log of predicted demand for historical time series from October 2008 to March 2012, see (Devine, 2012). For $d > 5$, equation (3) indicates that the simulated demand follows a discrete mean-reverting process where Seasonal Normal Demand (SND) is the mean. The parameters associated with equation (3) are also given in Table 2 and were found from an AR(1) model fitted to the difference between the natural log of actual demand and the natural log of SND for historical time series from October 2007 to March 2012, see (Devine, 2012). The data for each of these time series was obtained from the UK National Grid's Data Item Explorer [5].

This overall stochastic process replicates the information that those in the UK gas market have on a given day whereby today, the market knows demand exactly as in equation (1). Looking at the next five days ahead, the market has can predict (with some error) what the demand is going to be, but with some uncertainty, as in equation (2). Beyond this, all they know is seasonal normal demand and that if demand is low (or high) on one day, it is likely to be around that level again the next day, as described by equation (3).

3 Model

The Rolling Optimisation Model involves solving a sequence of linear programming problems, each time with updated inputs. Each roll of the model consists of P sources of supply, SO storage facilities, S different demand scenarios and a time horizon of D days. On each day within the time horizon, P sources of supply provide gas that is used either to meet demand or injected to storage. The model constrains the daily amount of gas each source P can supply. On each day, the SO storage facilities either inject gas coming from the sources of supply to its facility or withdraw gas from its facility to be used in meeting demand, or do nothing. ROM constrains the maximum amount of gas injected to, or withdrawn from, storage on any given day. It also constrains the maximum and minimum amount of gas allowed to be held in storage on any given day.

Each of the S demand scenarios is a time series generated (via using a Monte-Carlo simulation) using the stochastic process for demand developed in Section 2. Each scenario s has a probability, $Prob^s$, associated with it. As explained in Section 2, all these demand scenario time series have an identical value on the first day. As a result, the first day of each roll is scenario-independent.

ROM involves an imaginary central planner choosing how the different sources of supply and storage facilities meet demand on this scenario-independent first day, whilst ensuring all possible future demands

[4] When simulating stochastic demand for the first day of this process, the error from the previous day is assumed to be zero.

[5] See: http://marketinformation.natgrid.co.uk/gas/DataItemExplorer.aspx

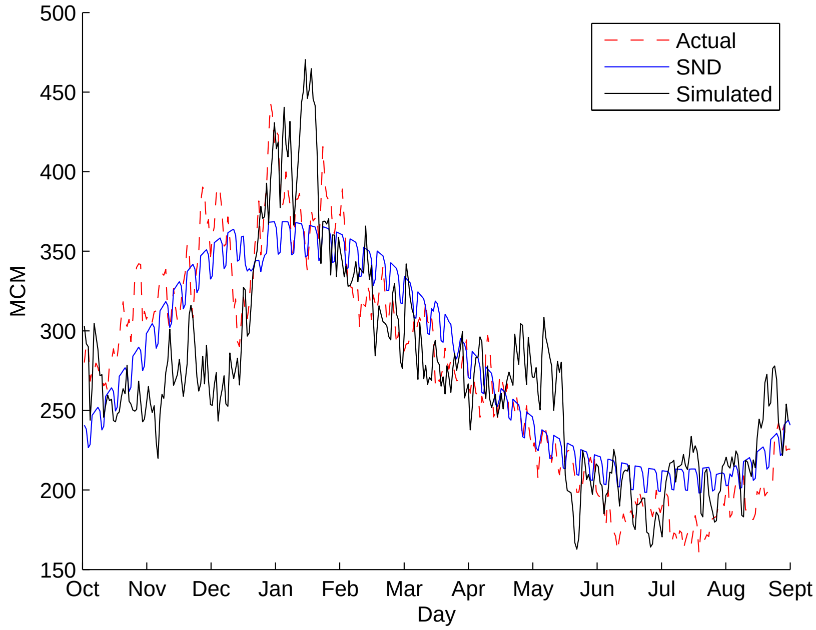

Fig. 2 Actual demand (red line), seasonal normal demand (blue line), and a simulated demand path (black line), for the gas year '08 - '09.

are also met at minimum expected cost. Throughout this paper, the index p (for producers) runs from 1 to P; the index so (for storage operators) runs from 1 to SO; the index s (for scenarios) runs from 1 to n while the index d runs from 2 to D, unless otherwise stated. The index d_1 represents the scenario-independent first day. Each optimisation (or roll) of ROM may be defined as follows:

$$\begin{aligned} \min \quad & \sum_{p=1}^{P}(c_p Q_{p,d_1}) + \sum_{so=1}^{SO}(a_{so,d_1} I_{so,d_1} + b_{so,d_1} W_{so,d_1}) + \\ & + \sum_{s=1}^{S} Prob^s \Big[\sum_{d=2}^{D}\Big(\sum_{p=1}^{P}(c_p Q^s_{p,d}) + \sum_{so=1}^{SO}(a_{so,d} I^s_{so,d} + b_{so,d} W^s_{so,d})\Big)\Big] \end{aligned} \tag{4}$$

subject to:

$$Demand_{d_1,d_1} = \sum_{p=1}^{P} Q_{p,d_1} + \sum_{so=1}^{SO}(W_{so,d_1} - I_{so,d_1}), \quad (\lambda_{Demand_{d_1}}), \tag{5}$$

$$Demand^s_{d_1,d} = \sum_{p=1}^{P} Q^s_{p,d} + \sum_{so=1}^{SO}(W^s_{so,d} - I^s_{so,d}), \quad (\lambda_{Demand^s_d}), \tag{6}$$

$$0 \le Q_{p,d_1} \le Q_p^{max}, \quad (\lambda_{Q_{p,d_1}}), \tag{7}$$

$$0 \le I_{so,d_1} \le I_{so}^{max}, \quad (\lambda_{I_{so,d_1}}), \tag{8}$$

$$0 \le W_{so,d_1} \le W_{so}^{max}, \quad (\lambda_{W_{so,d_1}}), \tag{9}$$

$$0 \leq Q^s_{p,d} \leq Q^{max}_p, \quad (\lambda_{Q^s_{p,d}}), \tag{10}$$

$$0 \leq I^s_{so,d} \leq I^{max}_{so}, \quad (\lambda_{I^s_{so,d}}), \tag{11}$$

$$0 \leq W^s_{so,d} \leq W^{max}_{so}, \quad (\lambda_{W^s_{so,d}}), \tag{12}$$

$$MinStor_{so} \leq IntStor_{so} + I_{so,d_1} - W_{so,d_1} \leq MaxStor_{so}, \quad (\lambda_{Stor_{so,d_1}}), \tag{13}$$

$$MinStor_{so} \leq IntStor_{so} + I_{so,d_1} - W_{so,d_1} + \sum_{e=2}^{e=d} (I^s_{so,e} - W^s_{so,e}) \leq MaxStor_{so}, \quad (\lambda_{Stor^s_{so,d}}). \tag{14}$$

Tables 3 and 4 name the output variables and parameters associated with the model, respectively, while Table 5 describes the input variables that are updated before each roll of the model. The variables in the parentheses, alongside constraints (5) - (14), represent the Lagrange multipliers associated with that constraint. Some of these Lagrange multipliers are also used in the discussion on the results of the model. In particular $\lambda_{Demand_{d_1}}$ and $\lambda_{Demand^s_d}$ are the Lagrange multipliers associated with the demand constraints (5) and (6), and hence represent the marginal cost of meeting demand. As a result, they are used throughout this paper to represent the price of gas. The objective function (equation (4)) minimises

Table 3 Outputs associated with ROM.

Variable	Explanation
Q_{p,d_1}	Amount supplied by source p for the first day
I_{so,d_1}	Amount injected by storage facility so on the first day
W_{so,d_1}	Amount withdrawn by storage facility so on the first day
$Q^s_{p,d}$	Hypothetical amount supplied by source p for day d associated with scenario s
$I^s_{so,d}$	Hypothetical amount injected by storage facility so on day d associated with scenario s
$W^s_{so,d}$	Hypothetical amount withdrawn by storage facility so on day d associated with scenario s
$\lambda_{Demand_{d_1}}$	Marginal price of gas associated with the first day
$\lambda_{Demand^s_d}$	Hypothetical marginal price of gas associated with day d and scenario s

Table 4 Parameters associated with ROM.

Parameter	Explanation
c_p	Cost associated with supply source p
$a_{so,d}$	Unit cost of injection for storage facility so on day d
$b_{so,d}$	Unit cost of withdrawal for storage facility so on day d
Q^{max}_p	Daily maximum production rate for producer p
I^{max}_{so}	Daily maximum injection rate for storage facility so
W^{max}_{so}	Daily maximum withdrawal rate for storage facility so
$MaxStor_{so}$	Maximum storage capacity for storage facility so
$MinStor_{so}$	Minimum storage capacity for storage facility so
$Prob^s$	Probability associated with scenario s

Table 5 Parameters that change from roll to roll in ROM.

Parameters	Explanation
$Demand_{d_1,d_1}$	Gas demand on the first day
$Demand^s_{d_1,d}$	Hypothetical gas demand on day d associated with scenario s
$IntStor_{so}$	Initial amount of gas held by storage facility so

the total cost of production, the cost of injection to storage and the cost of withdrawal from storage for the scenario-independent day d_1, as well as the total expected cost for all other days. As a result, the (imaginary) central planner in this model is assumed to be risk-neutral. This total expected cost is calculated over all possible demand scenarios with the weight associated with each scenario being determined by the probability $Prob^s$. The costs of production (c_p) are assumed not to vary with the level of production [6]. Long-term gas contracts are also not taken into account in the model. In reality, both of these assumptions simplifications. The reason for these simplifications, in the context of the UK natural gas market, is that information regarding costs and contracts are not freely available in the public domain.

Equation (5) ensures that real demand on the first day is met, while equation (6) ensures that (hypothetical) demand is met on every day $d > d_1$ for each scenario s. Equations (7) and (10) provide upper and lower bounds on the production rate for producer p. Equations (8), (11), (9) and (12) provide upper and lower bounds for the injection and withdrawal rates for storage facility so. Equation (13) constrains the total amount of gas in storage facility so at the end of day d_1. It ensures that the total amount of gas stored cannot exceed $MaxStor_{so}$ or go below $MinStor_{so}$. Equation (14) provides similar constraints for each $d > d_1$ and for each scenario s. The subscript e represents all days from day 2 to day d (Note: the total number of days is D).

The decisions made in ROM mimic those made in the real world UK gas market. Each day decisions are made in the real market on how today's (exactly-known) demand is met, as in equation (5). These decisions must take into account all possible demands, as in equation (6), and are based on the information that is known today about demand. Once these decisions are made in the real market, a new day arrives, which brings improved information about demand. This is replicated in the model, as after each optimisation the inputs (i.e., demands and amount of gas in storage) of the model are updated before a new optimisation problem is solved in the next roll. Investment decisions are not taken into account in ROM as the model is (typically) run over the relatively short timescale of one year, for example, see Sections 3.2, 3.3 and 4.

The outputs listed in Table 3 represent the outputs associated with one roll of ROM [7]. Hence, the total output of ROM is the collection of these variables from each roll. In Sections 3.2 and 3.3, only the output variables associated with the scenario-independent first day (d_1) are used for comparison with actual data. The reason for this is that the decision variables associated with d_1 (i.e., Q_{p,d_1}, I_{so,d_1}, W_{so,d_1} and $\lambda_{Demand_{d_1}}$) replicate actual daily decisions of how daily known demand is met. The rest of the variables presented in Table 3 are associated hypothetical decisions of how possible future demands might be met. Tables 4 and 5 display the parameters associated with each roll of ROM. The parameters in Table 4 do not change from from roll to roll while those in Table 5 may (see Section 3.1).

[6] While none of the c_p change with the level of production, it should be noted that in Sections 3.2, 3.3 and 4 the different sources of supply in the UK are broken up into multiple tranches, each with a different cost of production. This implicitly allows the cost of each source of supply to change as the level of production changes.

[7] The outputs of ROM include all Lagrange multipliers. However, only those associated with demand constraints are analysed in any detail.

It is proven by Devine (2012) that the production costs of ROM can be shifted by adding a constant β to the value of each cost parameter c_p without affecting the optimal flows according to the model (Q_{p,d_1}, $Q^s_{p,d}$, I_{so,d_1}, $I^s_{so,d}$, W_{so,d_1}, and $W^s_{so,d}$). Similarly, the production and storage costs can be both scaled by multiplying c_p, $a_{so,d}$ and $b_{so,d}$ by a positive constant α, again without affecting the optimal flows. When this is done, the prices of ROM ($\lambda_{Demand_{d_1}}$ and $\lambda_{Demand^s_d}$) become similarly scaled and shifted by the same values of α and β, respectively. The main assumptions required for these results are the cost minimisation, perfect competition structure of the model and the linearity of costs. These results allow the prices produced by ROM to be calibrated to actual SAPs, as is shown in Sections 3.2 and 3.3. Appendix A illustrates the results for a small theoretical example.

3.1 Update rules after each roll

After each optimisation (or roll) the inputs of ROM are updated as follows:

1. The actual demand and SND time series move forward one day. For example, if for the first optimisation the actual demand and SND time series were from the 1st October 2008 to the 30th of September 2009, then they would now be from the 2nd of October 2008 to the 1st of October 2009.
2. A new set of demand scenarios are developed using the updated actual demand time series, updated SND time series and the methodology described in Section 2.
3. Using the injections to, and withdrawals from, storage for the first day in the previous optimisation, the initial amount of gas in storage, for each storage facility so, is updated as follows[8]:

$$IntStor_{so} = IntStor_{so} + I_{so,d_1} - W_{so,d_1}. \tag{15}$$

After each optimisation, the exactly known demand on the first day, $demand_{d_1}$, as well as the (hypothetical) demand for the following days changes. This reflects what happens in the real UK gas market: for each new day, those in the market have improved information on demand to make their decisions with. Fig. 3 describes the inputs and outputs of each roll of the model for an example starting on the 1st of April 2012 with 365 rolls. It should be noted the optimal results for one roll affect the results for subsequent rolls. For example, if there was a large amount of withdrawals from storage in December, then the amount of gas storage would become low for January thus affecting the amount of gas that could be withdrawn in January.

3.2 ROM fitted to the year beginning in April 2010

Having introduced ROM in the previous section, the model is now calibrated by fitting it to the UK gas market for the year starting on the 1st of April 2010. The aim of this analysis is to estimate the parameters of the model that produce results that best fit actual data. These results include the flows of gas used to meet demand, the amount of gas in storage and the daily SAPs. Each of these are described in further detail below. In order to obtain these results, ROM was formulated with 18 sources of supply ($P = 18$), 3 storage facilities ($SO = 3$) and 3 scenarios ($S = 3$). The model was run over a horizon of $D = 365$ days and for 365 rolls.

[8] For the first roll of the model, the initial amount of gas in storage is a parameter typically determined using actual storage data.

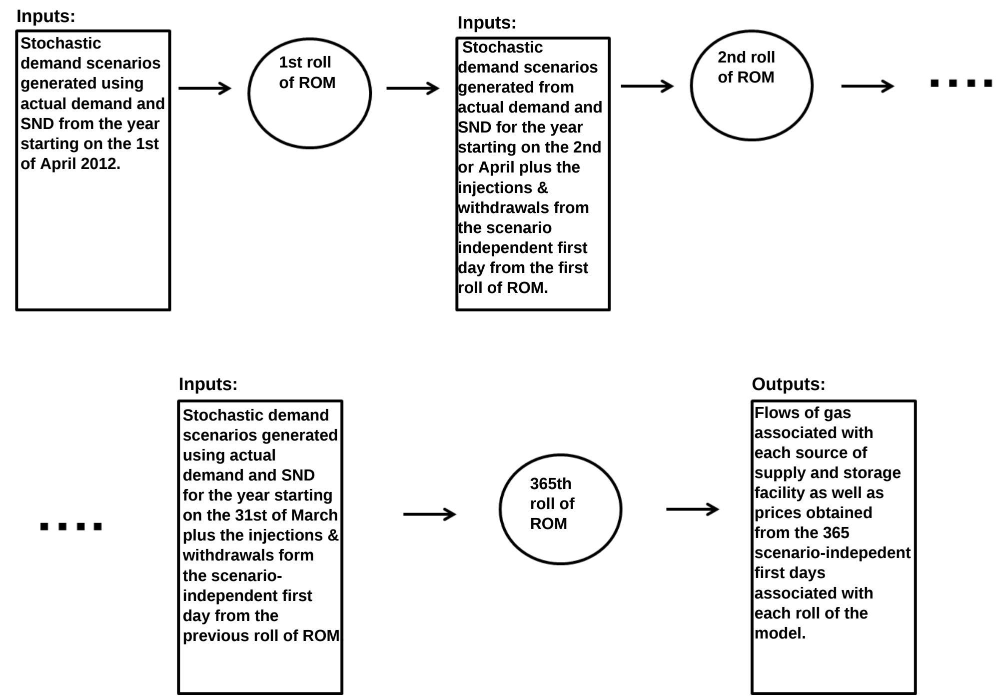

Fig. 3 Inputs and outputs associated with ROM.

The sets of three demand scenarios were generated using the stochastic process for demand described in Section 2. Actual demand and SND for the year starting on the 1st of April 2010 were used for the first roll (or optimisation) of the model. For the second roll of the model, the stochastic process for demand moved forward one day, i.e., actual demand and SND for the year starting on the 2nd of April 2010 was used to develop the scenarios. For subsequent rolls, SND moved forward in a similar pattern. As each of the 3 scenarios were randomly generated for each roll of the model, they were assigned equal probabilities, i.e., $Prob^s = \frac{1}{3}\,\forall\, s$. As well as this, there is no a priori information available to suggest that one scenario should be more or less likely than another, hence they are weighted equally

The 3 storage facilities were chosen to represent the 3 different types of storage facilities in the UK, namely Long-, Medium- and Short-range-storage (LRS, MRS and SRS respectively). LRS captures the seasonal variation in demand in the UK gas market whereby gas is injected in the relatively low-demand summer and withdrawn in the relatively high-demand winter. There is only one LRS in the UK and that is the Rough storage facility (DUKES, 2010). MRS captures the weekly/monthly variations in gas demand in the UK, as well as the seasonal variation in demand. There are six MRS facilities in the UK: Hornsea, Holehouse Farm, Hatfield Moor, Humbly Grove, Aldbrough and Holford (DUKES, 2012) [9]. SRS is also known as LNG storage as it involves storing gas by freezing it into its liquid form[10]. Prior to May 2011, there were three SRS facilities in the UK: Avonmouth, Glenmavis and Partington. However, in May 2011, the Glenmavis and Partington facilities stopped offering commercial services (UK National Grid, 2011c). This has meant that SRS has become insignificant to the UK gas market. As a result and for the sake of brevity, the amount of gas in SRS, as modelled by ROM, is not analysed in detail in this paper; see Devine (2012) for further details.

[9] The Holford MRS facility only became operational in 2012.
[10] See: http://www.nationalgrid.com/uk/Gas/lngstorage/What/

The five sources of supply in the UK gas market are

1. UK Continental Shelf (UKCS),
2. Norwegian imports,
3. LNG imports,
4. Balgzand Bacton Line (BBL) pipeline,
5. Interconnector UK (IUK) pipeline.

The largest source of supply is the UKCS. This is the UK's indigenous gas supply that is taken from the seas surrounding Britain. Norwegian imports represent gas imported to the UK from Norway through pipelines. LNG is natural gas in liquid form. It is obtained by cooling natural gas to -161 degrees Centigrade. This liquifies the gas and hence makes it easier to ship around the world. The UK imports LNG from places such as Algeria, Trinidad and the Middle East[11]. The BBL pipeline links England and the Netherlands and hence represents imports into the UK from that destination[12]. Similarly, the IUK pipeline is a pipeline that allows gas flows from Belgium to England[13]. In contrast to the BBL pipeline however, the IUK pipeline also can cater for exports from the UK to Belgium.

In ROM these five sources of supply were split up into multiple tranches giving 18 sources of supply in total. Each tranche had a separate cost associated with it. Splitting the five sources of supply into 18 tranches provides greater variability to price of gas modelled in ROM. For example, if $P = 5$ instead of $P = 18$, then there would only five price levels associated with gas in ROM, which is not comparable with actual prices as detailed in Fig. 6. Initially, each source of supply had three tranches. However, where necessary, the number of tranches was increased in order to improve the results of the model. The costs (c_p) and maximum capacities (Q_p^{max}) associated with the different tranches are given in Table 6. The maximum capacities are in units of million cubic meters (mcm). The costs of the model are assigned no particular unit because one can up- (down-) scale them in line with prices, without affecting the optimal flows of the model. The capacities of these tranches, as well as the costs associated with them, were determined using a simulated annealing algorithm (further details below).

Table 7 displays the total maximum capacities for the five different sources of supply (i.e., the sum of the tranches' capacities). The total maximum capacities for Norway, the BBL pipeline and the IUK pipeline were obtained from the UK's Department of Energy and Climate Change (DUKES, 2010). The total maximum capacities for the UKCS and LNG were estimated using the UK National Grid's 10 year statement (UK National Grid, 2011c). This statement provides peak supply availability for the gas years beginning on the 1st of October 2009 and 1st of October 2010, neither of which are applicable to the analysis of this section, as this analysis is fitted to data for the year beginning on the 1st of April 2010. As a result, the maximum UKCS and LNG capacities displayed in Table 7 are the average of the October 2009 and October 2010 values.

The total maximum capacity of LNG was also assumed to be only at 70% of its true maximum capacity. This is in accordance with the levels assumed by the UK National Grid (UK National Grid, 2011a,b). 100% LNG capacity would correspond to a situation where a LNG ship is constantly offloading gas into the UK gas network (i.e., when one ship finishes another always automatically begins offloading; this assumption is highly unrealistic and is not supported by the available data). While the size and

[11] See: http://www.nationalgrid.com/uk/GrainLNG/needs/
[12] See: http://www.bblcompany.com/
[13] See: http://www.interconnector.com/

Table 6 The costs and maximum capacities (in mcm) of the different tranches for each sources of supply for the year starting in April 2010, obtained by simulated annealing.

Tranche	c_p	Q_p^{max}
UKCS_1	58.464	48
UKCS_2	57.611	61
UKCS_3	56.552	64
UKCS_4	69.985	10
Norway_1	60.985	21
Norway_2	64.322	54
Norway_3	48.375	30
Norway_4	58.637	23
LNG_1	55.242	38
LNG_2	58.865	36
LNG_3	71.458	12
BBL_1	64.189	11
BBL_2	55.084	16
BBL_3	59.352	14
IUK_1	61.626	11
IUK_2	69.543	16
IUK_3	70.977	27
IUK_4	60.336	18

Table 7 The total maximum capacity (in mcm) for each source of supply for the year starting in April 2010.

Source of supply	Total maximum capacity
UKCS	183
Norway	128
LNG	86
BBL	41
IUK	72

Table 8 Short- and long-run costs associated with long-, medium- and short-range storage for ROM, obtained by simulated annealing.

	LRS	MRS	SRS
$a_{so,d}$ (Short-run)	0.069	0.069	2
$a_{so,d}$ (Long-run)	0.664	0.664	2
$b_{so,d}$ (Short-run)	0.064	0.019	2
$b_{so,d}$ (Long-run)	0.019	0.664	2

number of tranches of the different sources of supply were varied in the model's development, the total maximum capacity of each source of supply was fixed at levels determined from the references above.

The 3 storage operators represent the aggregate of the 3 different types of storage facilities in the UK gas market, namely long- (LRS), medium- (MRS) and short-range-storage (SRS). Table 8 presents the long- and short-run unit costs of storage injection and withdrawal, $a_{so,d}$ and $b_{so,d}$ respectively. Short-run costs are the storage costs associated with the first 26 days of ROM, while the long-run costs are those associated with all other days. As before, the units of these costs are arbitrary. When $d \leq 26$ the injection and withdrawal costs take the short-run values, while when $d > 26$ (called the lag) they take the long-run costs. The long- and short-run costs for MRS encourages gas to be injected quickly and withdrawn quickly, as the short-run costs are cheaper than the long-run costs. As mentioned earlier, MRS is used to capture short term (i.e., weekly/monthly) variations in demand and prices in the UK gas market. The long- and short-run costs for MRS encourage this behaviour. Similar to MRS, the long-run and short-run

Table 9 Parameters associated with long-, medium- and short-range storage (in mcm) for ROM for the year starting in April 2010.

	LRS	MRS	SRS
$IntStor_{so}$	440	295	56
$MaxStor_{so}$	3300	810	180
$MinStor_{so}$	440	169	39
I_{so}^{max}	43	43	35
W_{so}^{max}	43	43	35

injection costs for LRS also encourages gas to be injected quickly. In contrast, the long-run and short-run withdrawal costs for LRS discourage this behaviour, as the long-run costs are cheaper than the short-run costs. These costs allow LRS to capture the seasonal variation in the demand for natural gas. For SRS, the long- and short-run costs are the same.

Simulated annealing was used to calibrate the parameters in Table 8, the short-run lag of 26, as well as the parameters in Table 6[14]. Simulated annealing is an optimisation algorithm that locates a good approximation to the global optimum given a large search space (Kirkpatrick et al, 1983), (Metropolis et al, 1953). The optimum of interest here is the minimum error between the actual data and the results obtained from ROM. This error is detailed in the following equation:

$$Error = \sum_{l=1}^{L} \sqrt{\frac{\sum_{r=1}^{R} (Actual_r^l - Simulated_r^l)^2}{R}} \tag{16}$$

where $Actual_r^l$ and $Simulated_r^l$ are the actual and simulated flows from source l on roll r respectively and where $R = 365$ represents the total number of rolls. The $L = 11$ different sources of supply appearing in the sum of equation (16) were UKCS, Norway, LNG, BBL, IUK, as well as the injections and withdrawals into and out of LRS, MRS and SRS. In total, $LR = 4015$ actual observations were used to calibrate 49 parameters.

Table 9 provides the initial ($IntStor_{so}$), maximum ($MaxStor_{so}$) and minimum storage ($MinStor_{so}$) levels for each of the 3 facilities. The initial amount of gas in storage is only a parameter for the first roll of the model. For subsequent rolls, $IntStor_{so}$ is determined using the amount of gas injected (I_{so,d_1}) to and withdrawn (W_{so,d_1}) from storage from the previous roll of the model, as described in Section 3.1. The maximum storage levels, daily injection and daily withdrawal rates were obtained from the UK's Department of Energy and Climate Change (DUKES, 2010), while the initial levels of gas in storage were obtained from the actual levels on the 1st of April 2010, as recorded by the UK National Grid. While the theoretical minimum level of gas in storage is zero, actual stock levels suggest that this is never the case. As a result the minimum levels of gas in storage were estimated from the minimum actual storage levels observed in the UK gas market from the 1st of April 2010 to 31st of March 2011. This data was again obtained from the UK National Grid. Table 9 also displays the daily maximum injection (I_{so}^{max}) and withdrawal rates (W_{so}^{max}).

The remainder of this section qualitatively compares the outputs of ROM with actual flows and prices. As described in the previous section, the outputs of the model were obtained from the production rates (Q_{p,d_1}), injection rates (I_{so,d_1}) and withdrawal rates (W_{so,d_1}) on the scenario-independent first days, which were obtained from each roll of the model. For each roll of the model, the outputs were found to be

[14]The parameters in Table 7 were not obtained from simulated annealing as they were obtained from either the UK National Grid or the UK's Department of Energy and Climate Change.

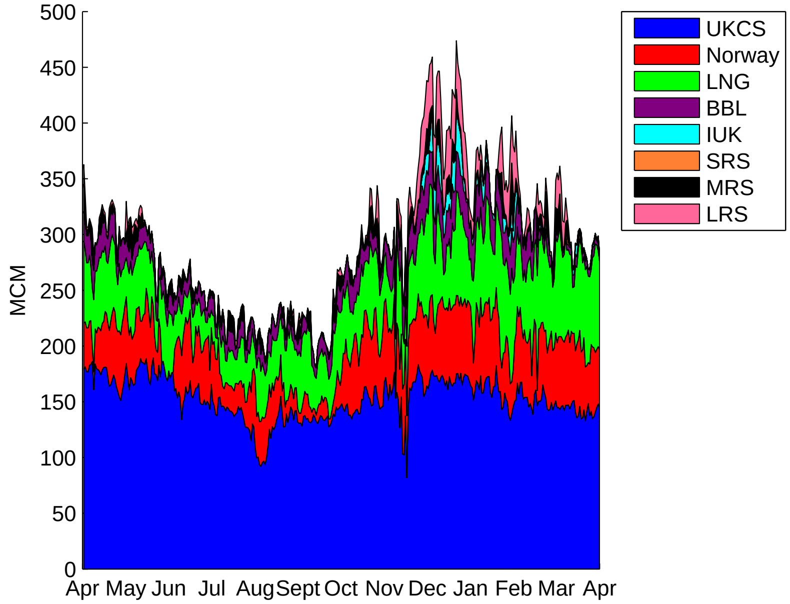

Fig. 4 Actual demand profile from the UK gas market starting on the 1st of April 2010.

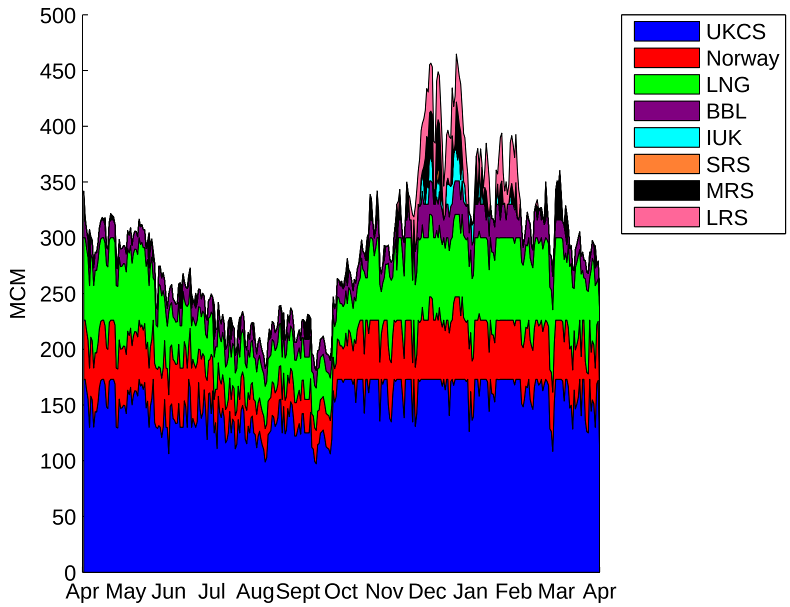

Fig. 5 Demand profile obtained from ROM for the UK gas market starting on the 1st of April 2010.

optimal. Fig. 4 displays the actual demand profile for the UK gas market for the year beginning on the 1st of April 2010. It indicates how the different sources of supply meet demand for each day of the year. The data for this graph was obtained from the UK National Grid's website. Fig. 5 shows a similar plot, but for the flows produced by ROM.

Table 10 Actual and simulated average daily flows for the UK gas market for year beginning on the 1st of April 2010 as well as the MAPE for the difference between them.

Source of supply	Actual	ROM	MAPE
UKCS	152.6	152.3	5.8%
Norway	47.0	47.2	4.1%
LNG	56.0	58.4	4.7%
BBL	19.6	19.2	1.8%
IUK	3.0	1.9	0.7%
SRS	0.1	0.0	0.06%
MRS	6.4	5.2	0.2%
LRS	8.7	6.9	1.4%

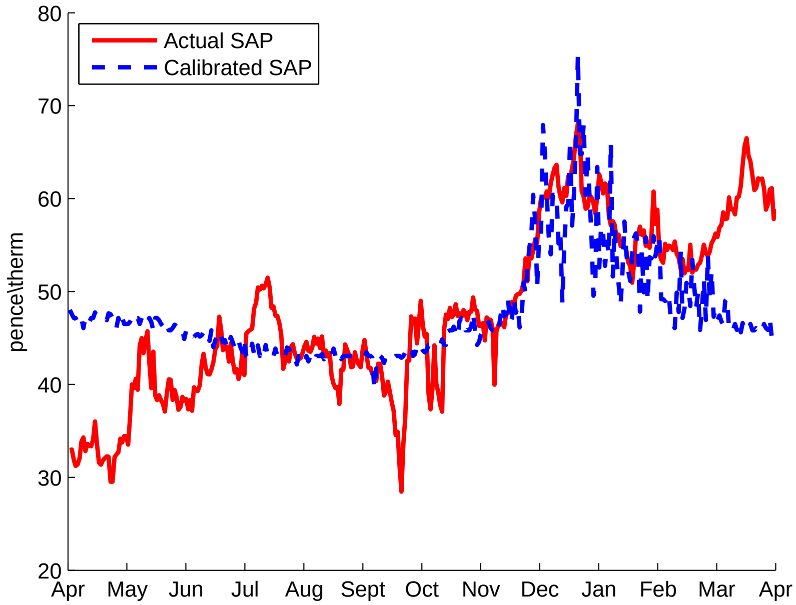

Fig. 6 Actual and calibrated ROM System Average Prices starting on the 1st of April 2010 ($MAPE = 11.3\%$).

Figs. 4 and 5 plus Table 10 indicate that the order and relative size of the actual flows obtained from the ROM are similar to the observed data[15]. The Mean Absolute Percentage Errors (MAPE) given in Table 10 also show that differences between the actual and simulated are relatively small [16] The largest source of supply is clearly the UKCS for both the actual data and the results produced from ROM. Norwegian and LNG imports are the next two largest sources of supply and are at roughly the same level, while imports through the BBL and IUK pipelines are the smallest sources of supply. In terms of storage, the actual data and ROM also both indicate that LRS is the largest storage supplier, while SRS is the smallest. Fig. 5 shows that many of the results, particularly in the winter months, are flat or oscillate around a flat base. The reason for this is that the maximum capacity constraints for many of the tranches are binding, particularly in the high demand periods, thus meaning that there is little variance in the results produced.

[15] The values for UKCS exclude gas that is injected to storage.

[16] The Mean Absolute Percentage Error for source l is $MAPE^l = \sum_{r=1}^{R} \frac{|Actual_r^l - Simulated_r^l|}{\frac{1}{R}\sum_{r=1}^{R}\sum_{l=1}^{L} Actual_r^l}$.

Fig. 6 displays the daily actual SAP, as well as SAPs produced by ROM. As mentioned above, the costs of ROM (and hence $\lambda_{demand_{d_1}}$) can be both scaled and shifted without affecting the optimal flows of the model. As a result, the modelled prices shown in Fig. 6 were obtained using the following formula:

$$CalSAP = \alpha \lambda_{demand_{d_1}} - \beta, \ \ \forall r, \tag{17}$$

where α is the positive scaling parameter and β is the shifting parameter. The values $\alpha = 6.9$ and $\beta = 359.1$ were obtained by minimising the L_2-norm of the difference between actual SAPs and SAPs produced by ROM. Fig. 6 indicates that the calibrated prices provide a reasonable fit to the actual data, particularly for the winter peak. The MAPE associated with this graph is 11.3% [17]. In the summer months, the calibrated SAPs fail to capture the variation in actual SAPs. Fig. 6 also shows that the prices produced by ROM are highly seasonal. The reason for this is that these modelled SAPs are obtained from Lagrange multipliers associated with demand constraints. Thus, when demand is high (low) in the winter (summer), so is the predicted SAP. In reality, other factors affect natural gas prices in the UK. For example, many gas pricing contracts in the UK and particularly in Europe are indexed linked to oil prices (Honoré, 2010). Demand in the summer is likely to be met to a very high degree by these long term contracts, which means that using ROM with fixed costs, the calibrated SAP would be fairly constant, while the actual SAP would vary along with the oil price. The spot market would be most heavily used in the winter, which may well explain why the calibrated SAP fits the actual SAP in this period.

Figs. 7 and 8 display the actual and simulated amount of gas in storage for long- and medium–range-storage, respectively, for the year beginning on the 1st of April 2010. As above, the actual data was obtained from the UK National Grid, while the simulated results were obtained from the withdrawal (W_{so,d_1}) and injection (I_{so,d_1}) rates from the scenario-independent first days from each roll of ROM.

Fig. 7 indicates that the amount of gas in LRS, as derived by ROM, is similar to the actual amount. The MAPE associated with this graph is 5.9%. The upward and downward slopes of both time series in the plot are particularly similar, which suggests that the withdrawal and injection rates defined for the model are correct. The time when LRS starts injecting and withdrawing is also almost identical for both sets of time series.

Fig. 8 shows the actual and simulated amount of gas in MRS. It too indicates that the results obtained from ROM provide a reasonably good fit to the actual data, as it captures the weekly and monthly fluctuations in the amount of gas in MRS. However, the rate of withdrawals from MRS differ between the actual and simulated results. In the actual data, withdrawals from MRS are spread over across the winter months. In the simulated data however, the majority of withdrawals from MRS occur in December. The reason for this, as Fig. 4 shows, is the period of relatively large demand period that occurred in December 2010. The difference suggests that ROM either underestimated the cost of withdrawals from MRS or the actual market underused MRS during the December demand peak. At this point, it must also be noted that the vertical scale of Fig. 8 is relatively small in comparison to that of Fig. 7. Hence, MRS is of lesser importance than LRS.

To summarise this section, ROM was applied and fitted to the UK gas market for data starting on the 1st of April 2010. Figs. 4 - 8 show that the flows and prices produced by the model fit reasonably well to actual data. However, there were some areas where ROM failed to capture the characteristics of the

[17] The MAPE associated with Figs. 6 - 8 is $MAPE = \sum_{r=1}^{R} \frac{|Actual_r - Simulated_r|}{\frac{1}{R}\sum_{r=1}^{R} Actual_r}$.

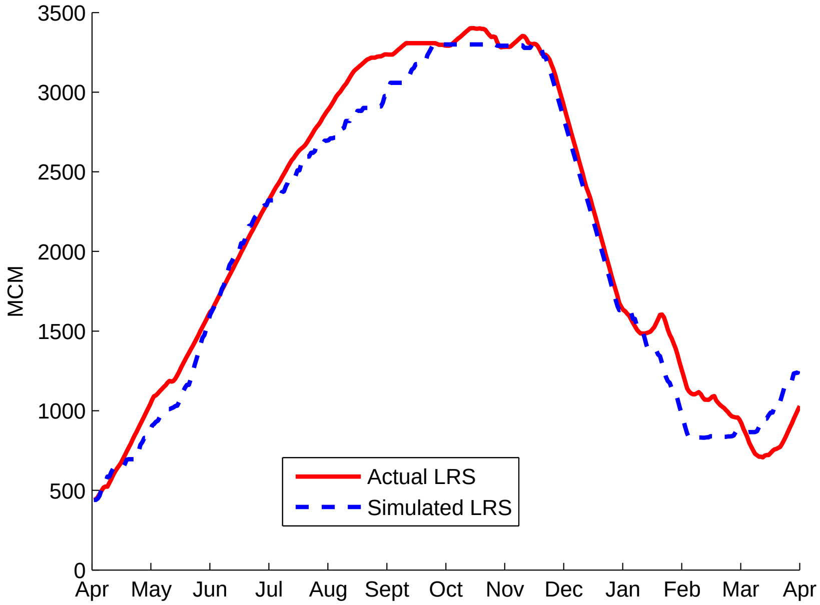

Fig. 7 Actual and simulated amount of Long-Range-Storage starting on the 1st of April 2010 ($MAPE = 5.9\%$).

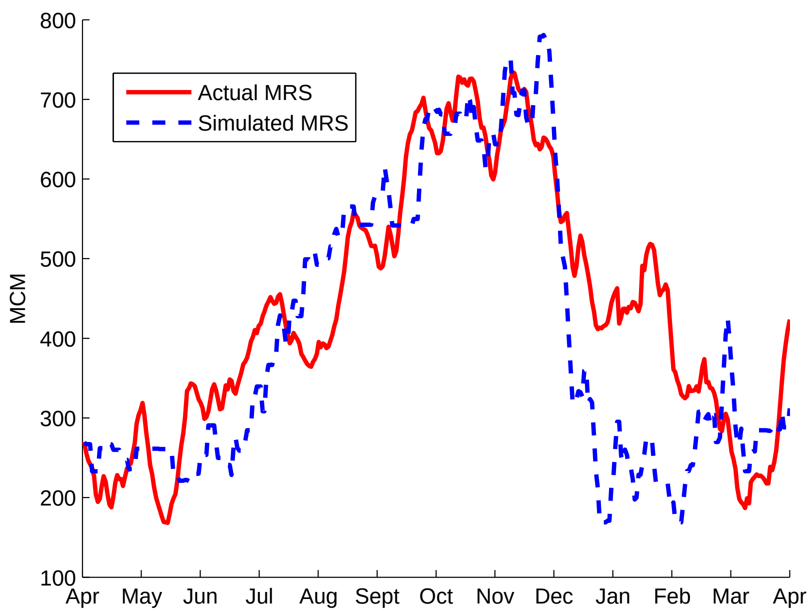

Fig. 8 Actual and simulated amount of Medium-Range-Storage starting on the 1st of April 2010 ($MAPE = 18.8\%$).

Table 11 The costs and maximum capacities (in mcm) of the different tranches for each source of supply for the year starting in April 2011.

Tranche	c_p	Q_p^{max}
UKCS_1	58.464	48
UKCS_2	57.611	61
UKCS_3	56.552	48
Norway_1	60.985	21
Norway_2	64.322	54
Norway_3	48.375	40
Norway_4	58.637	30
LNG_1	55.242	38
LNG_2	58.865	36
LNG_3	71.458	28
BBL_1	64.189	11
BBL_2	55.084	16
BBL_3	59.352	14
IUK_1	61.626	11
IUK_2	69.543	16
IUK_3	70.977	27
IUK_4	60.336	18

actual data. In Section 3.3 below, the model is tested using actual data for the year starting on the 1st of April 2011.

3.3 ROM tested for the year beginning in April 2011

In Section 3.2, ROM was applied to data for the UK gas market for the year starting on the 1st of April 2010. This analysis estimated the parameters of the model that enable a reasonable fit to historical data from the year starting in April 2010. In this section, ROM is tested with similar parameters using data for the year starting on the 1st of April 2011. As in section 3.2, the model considers 3 storage facilities ($SO = 3$), 3 demand scenarios ($S = 3$) and a horizon of $D = 365$ days, for each roll of the model. The number of rolls (or optimisations) was 365. In contrast to Section 3.2 however, the model was formulated with 17 sources of supply ($P = 17$), a decrease of one. This loss of a source of supply takes into account the decreased level of UKCS supply for this time period as detailed below.

For each roll of the model, three demand scenarios were generated using the stochastic process for demand described in Section 2. Actual demand and SND for the year starting on the 1st of April 2011 were used for the first roll. As 3 scenarios were again randomly generated for each roll of the model, they were assigned equal probabilities, i.e., $Prob^s = \frac{1}{3}$ for $s = 1,2,3$. As in Section 3.2, the 17 sources represent the different sources of supply in the UK gas market, broken up into multiple tranches representing varying costs. The 3 storage operators represent the 3 different types of storage facilities in the UK, namely long-, medium- and short-range-storage. Table 11 shows the costs and maximum capacities associated with each tranche, while the total maximum capacities of the five different sources of supply (i.e., the sum of the tranches' capacities) are given in Table 12.

The capacities of the BBL and IUK pipelines, for the year beginning in April 2011, are the same as in Table 7 and are in line with figures supplied by the UK's Department for Energy and Climate Change (DUKES, 2011). The Norwegian total maximum capacity has increased to 145 mcm, again in accordance with values supplied by the UK's Department for Energy and Climate Change (DUKES, 2011). In contrast, the total maximum capacity for the UKCS has decreased to 157 mcm. This is line with fig-

Table 12 The total maximum capacity (in mcm) for each source of supply for the year starting in April 2011.

Source of supply	Total maximum capacity
UKCS	157
Norway	145
LNG	102
BBL	41
IUK	72

Table 13 Parameters associated with long-, medium- and short-range storage (in mcm) for ROM for the year starting in April 2011.

	LRS	MRS	SRS
$IntStor_{so}$	1040	440	39
$MaxStor_{so}$	3500	850	80
$MinStor_{so}$	440	169	15
I_{so}^{max}	43	43	13
W_{so}^{max}	43	43	13

ures supplied by the UK National Grid and the general downward trend seen in UKCS supplies since 2000 (UK National Grid, 2011c; Honoré, 2010). The LNG capacity has increased to 102 mcm. This increase takes into account the improved LNG infrastructure in the UK and is again in accordance with figures supplied by the UK National Grid (UK National Grid, 2011c). As explained in the previous section, the total maximum capacity of LNG was assumed to be at only 70% of its true maximum capacity (UK National Grid, 2011a,b). Table 13 provides the initial ($IntStor_{so}$), maximum ($MaxStor_{so}$) and minimum storage ($MinStor_{so}$) levels for each of the 3 storage facilities. It also displays the daily maximum injection (I_{so}^{max}) and withdrawal rates (W_{so}^{max}). As previously, the initial and minimum amount of gas in storage were obtained from data made available by the UK National Grid. The maximum storage levels, the maximum daily injection rate and the maximum daily withdrawal rate are updated in accordance with the UK's Department of Energy and Climate Change (DUKES, 2011). Most noticeably, they were updated to take into account the de-commissioning of the SRS facilities, Glenmavis and Partington. The costs associated with injection and withdrawal to and from storage remained the same as detailed in Table 8.

As described in the previous section, the outputs of the model were obtained from the supply rates (Q_{p,d_1}), injection rates (I_{so,d_1}) and withdrawal rates (W_{so,d_1}) of the scenario-independent first days, which were obtained from each roll of the model. Fig. 9 displays the actual demand profile for the UK gas market for the year beginning on the 1st of April 2011. The data for this graph were obtained from the UK National Grid's website. Fig. 10 shows a similar plot, but for the flows produced by ROM. Table 14 shows the actual and simulated daily average flows of gas in the UK gas market over the same period.

Figs. 9 and 10 plus Table 14 indicate that the order and relative size of the actual flows obtained from ROM are again similar to the observed data. The Mean Absolute Percentage Errors (MAPE) given in Table 14 also show that differences between the actual and simulated are relatively small[18]. The largest source of supply is again the UKCS for both the actual data and the results produced from ROM. The second largest source of supply is now clearly LNG, while Norwegian supplies are the third largest. Imports through the BBL and IUK pipelines are the smallest sources of supply. In terms of storage, the results obtained from ROM indicate that MRS is the largest storage supplier, while SRS is the smallest. This is in contrast to Section 3.2, where LRS was the largest storage supplier.

[18]The MAPEs in Table 14 were calculated in the same way as Table 10.

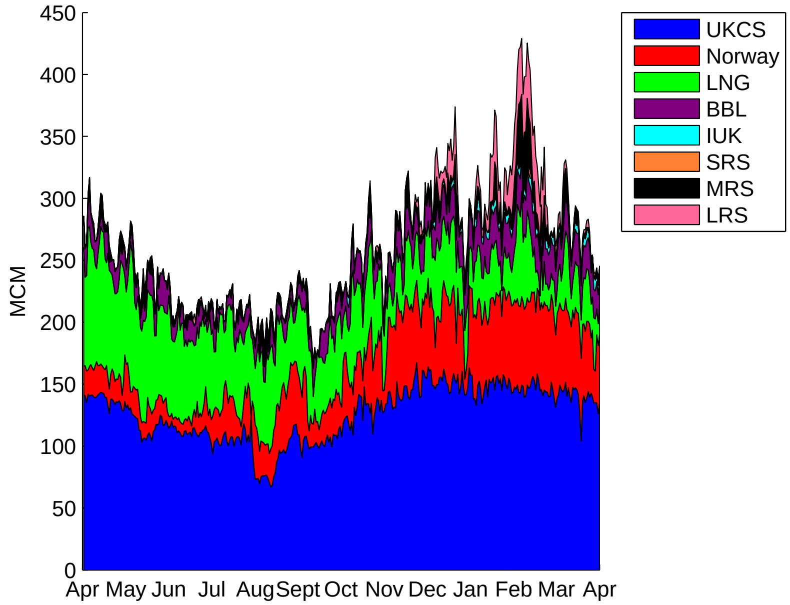

Fig. 9 Actual demand profile from the UK gas market starting on the 1st of April 2011.

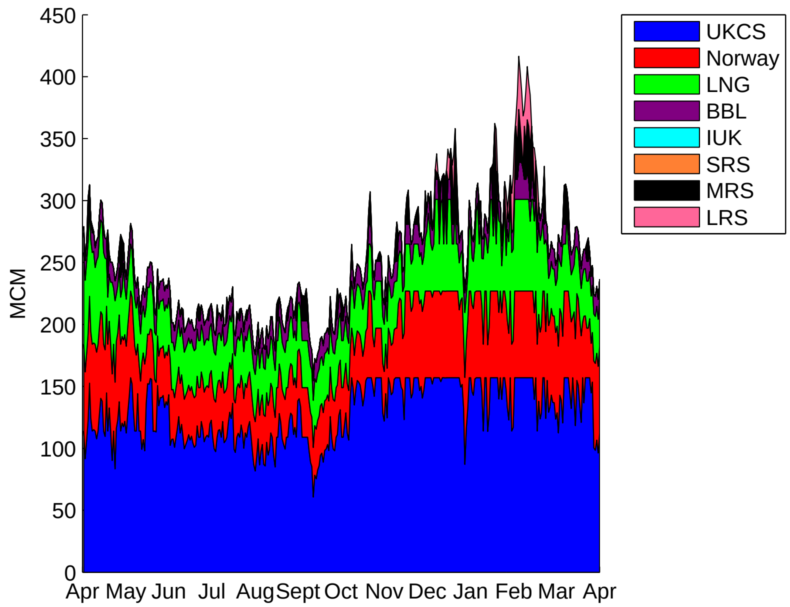

Fig. 10 Demand profile obtained from ROM for the UK gas market starting on the 1st of April 2011.

Table 14 Actual and simulated average daily flows for the UK gas market for year beginning on the 1st of April 2011 as well as the MAPE for the difference between them.

Source of supply	Actual	ROM	MAPE
UKCS	126.9	129.5	5.3%
Norway	55.1	54.7	6.2%
LNG	55.1	45.2	8.4%
BBL	19.8	16.2	2.7%
IUK	2.4	0.0	0.9%
SRS	0.1	0.0	0.04%
MRS	6.5	5.0	2.4%
LRS	6.2	2.2	1.6%

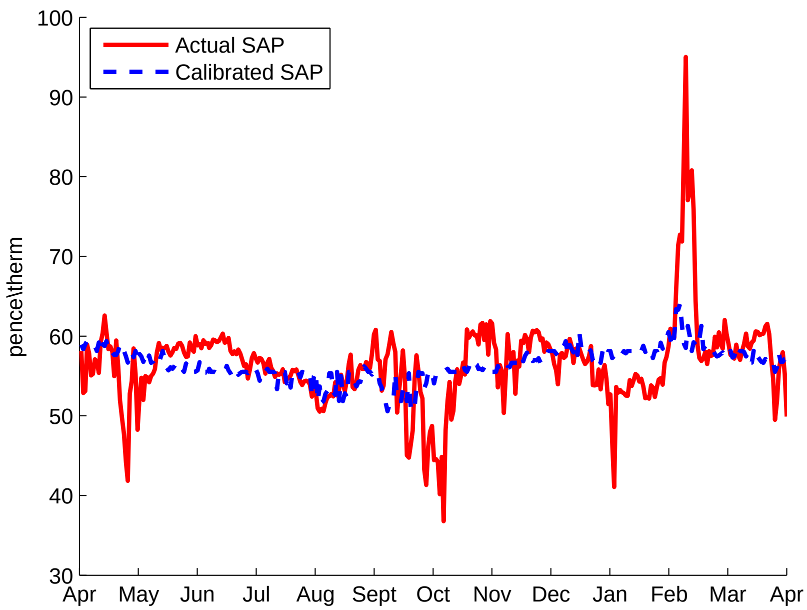

Fig. 11 Actual and calibrated predicted System Average Prices starting on the 1st of April 2011 ($MAPE = 5.8\%$). These prices were calibrated using equation (17) with $\alpha = 6.5$ and $\beta = 323.3$.

Fig. 11 displays the actual daily SAPs, as well as SAPs produced by ROM. In a similar manner to the previous section, the SAPs produced from ROM were calibrated using equation 17. The values of $\alpha = 6.5$ and $\beta = 323.3$ were again obtained ‘by minimising the L_2-norm of the difference between the actual SAPs and SAPs produced by ROM. Fig. 11 indicates that the calibrated prices provide a reasonable fit to the actual data. The MAPE associated with this graph is 5.9% [19]. However, the prices produced by ROM fail to capture the magnitude of some of the price spikes, in particular the large spike seen in February for the actual data.

Figs. 12 and 13 display the actual amount of gas in long- and medium-range storage facilities, as well as the amount in storage as predicted by ROM. Fig. 12 shows that the simulated results are similar to the actual data for LRS ($MAPE = 11\%$). As the upward and downward slopes are again relatively similar, it indicates that the injection and withdrawal rates of ROM are correct. The time of withdrawals according to both the actual data and simulated results are also almost identical. In particular, ROM successfully

[19] The MAPE used in Figs.11 - 13 is the same as the one used in Figs.6 - 8.

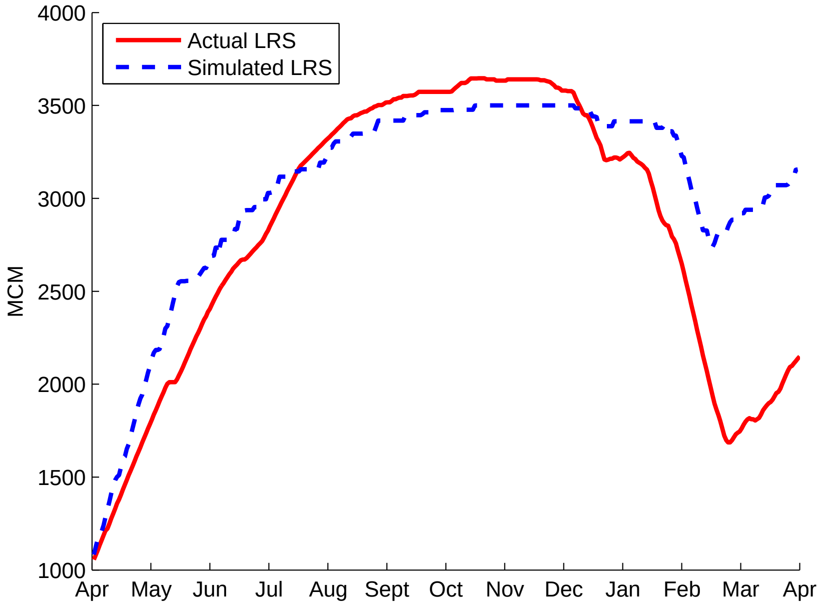

Fig. 12 Actual and simulated amount of Long-Range-Storage starting on the 1st of April 2011 ($MAPE = 11\%$).

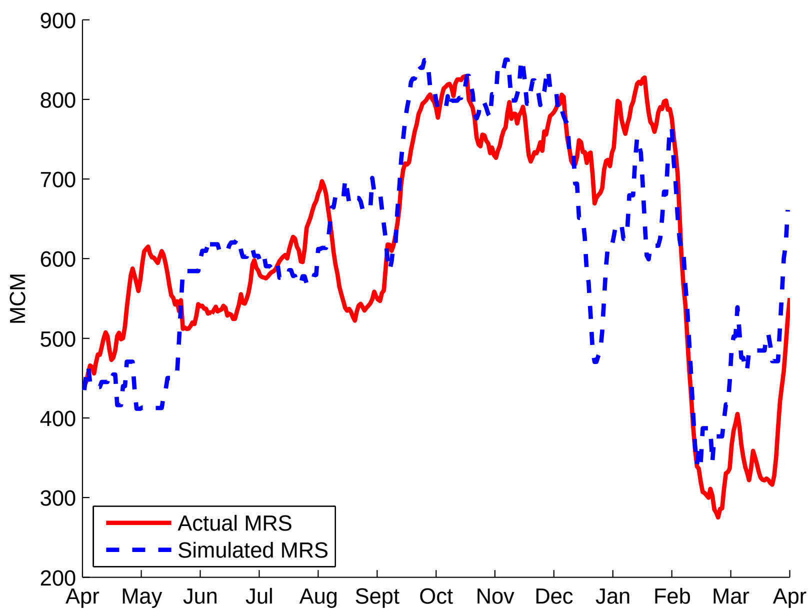

Fig. 13 Actual and simulated amount of Medium-Range-Storage starting on the 1st of April 2011 ($MAPE = 12.9\%$).

models the time when withdrawals stop in winter. However, ROM underestimates the amount withdrawn from LRS. Fig. 13 shows the actual and simulated amount of gas in MRS. It too indicates that the results obtained from ROM are qualitatively similar to the actual data ($MAPE = 12.9\%$), as it captures the weekly and monthly fluctuations in the amount of gas in MRS.

To summarise this section, Figs. 9 - 13 show that the results produced by ROM provide a reasonably good fit to actual data for the year starting on the 1st of April 2011. In Section 3.2 the parameters of the model were fitted for the year beginning on the 1st of April 2010. In this section similar parameters were used again, thus showing the robustness of ROM to changes in the data. Similar high-quality fits were found for the year starting in October 2010. This demonstrates that ROM is not dependent on starting date used.

4 Applications of ROM

In this section, some possible applications of the Rolling Optimisation Model (ROM) are analysed. In particular, the model is used to examine the impact of two stresses that might occur in the UK gas market. The first stress test involves a sudden withdrawal of Liquified Natural Gas (LNG) supplies from the UK. The motivation for this comes from the UK National Grid's Development of Energy Scenarios document (UK National Grid, 2011a), where they state that LNG flows to the UK are "subject to high levels of uncertainty". This is "due to the global nature of LNG and the options to flow to alternative markets". The second stress to the UK gas market examined in this section, is the occurrence of a particularly cold week in January.

In Sections 3.2 and 3.3, the model was applied to parameters for the years beginning on the 1st of April 2010 and 2011, respectively. These analyses are not used in this section as they are both retrospective analyses and were compared with actual data. In this section, ROM is applied to parameters for the year beginning on the 1st of April 2013. This analysis is predictive and is hence suitable for testing possible stresses that may occur in the UK gas market in the future.

4.1 Base case

In order to consider the different stress tests performed in this section, a base case must be analysed first. This was done by applying ROM to parameters for the year beginning on the 1st of April 2013. As in Sections 3.2 and 3.3, the model considers 3 demand scenarios ($S = 3$) and a horizon of $D =$ 365 days for each roll of the model, i.e., 365 rolls. The sets of three demand scenarios were generated using the stochastic process for demand described in Section 2. Seasonal Normal Demand (SND) for the year starting on the 1st of April 2010 was used for both actual demand and SND for the first roll (or optimisation) of the model.

The reason for simulating actual demand from the time series model, using SND as an input, is that actual demand for the year starting in April 2013 was unavailable at the time of writing. As explained in Section 2, the time series model describes the variation around projected SND. SND for the year starting in April 2010 was used as SND for the year starting April 2013 was also unavailable. As mentioned in the introduction, it is anticipated that ROM will be used by those in the UK gas market on a day-to-day basis, whereby one roll of the model will be carried out each day. Each day, those using the model will have available the actual demand for that day, predicted demand for the next five days and SND thereafter.

Table 15 The costs and maximum capacities (in mcm) of the different tranches for each source of supply for the year starting in April 2013.

Tranche	c_p	Q_p^{max}
UKCS_1	58.464	46
UKCS_2	57.611	60
UKCS_3	56.552	35
Norway_1	60.985	21
Norway_2	64.322	54
Norway_3	48.375	43
Norway_4	58.637	30
LNG_1	55.242	38
LNG_2	58.865	36
LNG_3	71.458	28
BBL_1	64.189	15
BBL_2	55.084	20
BBL_3	59.352	18
IUK_1	61.626	11
IUK_2	69.543	16
IUK_3	70.977	27
IUK_4	60.336	20

Table 16 The total maximum capacity (in mcm) for each source of supply for the year starting in April 2013.

Source of supply	Total maximum capacity
UKCS	141
Norway	148
LNG	102
BBL	53
IUK	74

Table 17 Parameters associated with long-, medium- and short-range storage (in mcm) for the Rolling Optimisation Model for the year starting in April 2013.

	LRS	MRS	SRS
$IntStor_{so}$	2164	563	15
$MaxStor_{so}$	3300	1120	80
$MinStor_{so}$	440	169	15
I_{so}^{max}	45	67	15
W_{so}^{max}	45	67	15

When actual demand is unavailable, predicted demand would be the preferred replacement. However, in this example predicted demand is also unavailable for the analysis.

As in Section 3.3, the model was formulated with the same number of sources of supply ($P = 17$), storage facilities ($SO = 3$) and scenarios ($S = 3$). Table 15 shows the costs and maximum capacities associated with each tranche.

The total maximum capacities of the five different sources of supply (i.e., the sum of the tranches' capacities) are given in Table 16. The capacities of Norway, the BBL and IUK pipelines have all increased and are in line with the figures supplied by the UK's Department for Energy and Climate Change (DUKES, 2012). The total maximum capacity for the UKCS has decreased from 157 mcm, in Table 12, to 141 mcm in Table 16. Table 17 provides the parameters associated with each of the 3 storage facilities. The initial amounts of gas in storage are the actual observed levels for the 1st of April 2012, while the minimum levels are the actual minimum levels of storage observed in the UK from April 2010 to April

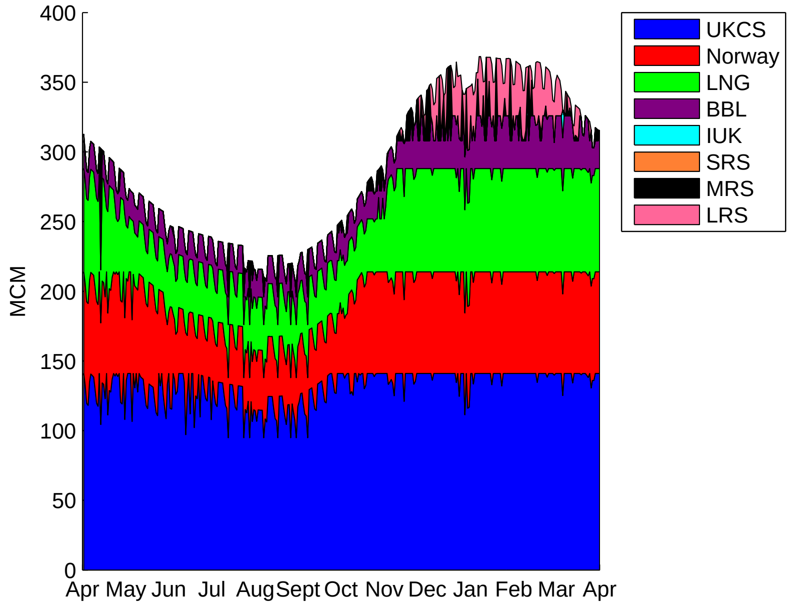

Fig. 14 Demand profile obtained from ROM for the UK gas market starting on the 1st of April 2013.

2012. These values were obtained from data made available from the UK National Grid. The maximum storage levels, the maximum daily injection rate and the maximum daily withdrawal rate are taken from information provided by the UK's Department of Energy and Climate Change (DUKES, 2012)[20]. They are updated to take into account the new MRS facility, Holford, as well as the de-commissioning of the SRS facilities, Glenmavis and Partington. The long- and short-run unit costs of storage injection and withdrawal are the same as in Table 8.

Fig. 14 displays how the different sources of supply meet demand in the UK gas market for the year beginning on the 1st of April 2013, as modelled by ROM. As with Figs. 5 and 10, it predicts that the UKCS will be the largest source of supply in the UK gas market. Imports from Norway and LNG Imports will be the next two biggest sources, at roughly the same level, while imports through the BBL and IUK will be the fourth and fifth largest sources, respectively.

Fig. 18 shows the System Average Prices (SAPs) obtained from ROM for the year starting on the 1st of April 2013. As in Sections 3.2 and 3.3, these prices were calculated using equation 17 with $\alpha = 6.48$ and $\beta = 323.324$. These values of α and β were determined according to the calibration of predicted and actual SAPs (see Section 3.3). As with Figs. 6 and 11, Fig. 18 predicts that prices will follow a similar pattern to demand and be highly seasonal.

Figs. 16 and 17 display the amount of gas in storage, according to ROM, for the year beginning on the 1st of April 2013. Fig. 16 follows a similar pattern to Figs. 7 and 12, whereby gas is injected in the summer and withdrawn in the winter. However, in contrast to Figs. 8 and 13, Fig. 17 does not show much weekly or monthly variation in the amount of gas in MRS. This is because the demand scenarios for this example were developed using the time series model with SND for the year starting in April 2010 as the input for both actual demand and SND. Seasonal normal demand is a relatively smooth time series when

[20] Also see: http://www.eon-uk.com/generation/holford.aspx

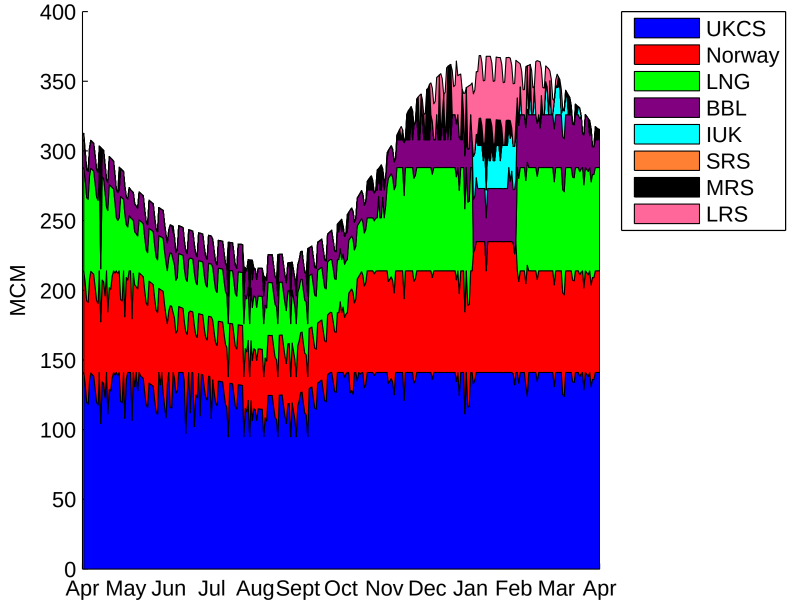

Fig. 15 Demand profile obtained from ROM for the UK gas market when there is a sudden drop in LNG supplies in January.

compared to actual demand (see Fig. 2). As it is unlikely that actual demand for the year starting on the 1st of April 2013 will be as smooth as seasonal normal demand, it is expected that the actual MRS will contain much more weekly/monthly variation.

4.2 Stress test 1: Sudden drop in LNG supplies in January

The first shock to the UK gas market analysed is a sudden drop in LNG supplies in January for a month long period. This shock examines the effect of the complete cessation of LNG supplies in the relatively high-demand winter time. The results of this analysis were obtained by applying the model to the parameters for the year beginning on the 1st of April 2013. For the first 276 rolls of the model (i.e., until the 1st of January), all the parameters were as described in Section 4.1. The maximum capacities of the LNG tranches were set to zero for the next 31 rolls, i.e., for each day from the 1st of January to the 31st of January. After that, the maximum capacities of the LNG tranches returned to the original levels. The market has no prior knowledge of the loss, which is in contrast to Section 4.3 below.

Fig. 15 shows the demand profile predicted by ROM for this case. When compared with Fig. 14, this indicates that IUK and Norwegian supplies, plus withdrawals from MRS, replace the removed LNG supplies in the affected weeks in January. Figs. 16 and 17 show the amount of gas in LRS and MRS, respectively, for both the base case and the first stress test. Fig. 17 shows that there is no dramatic change in the amount of gas in LRS between the two cases. The reason for this is that LRS is withdrawing at (or near) its maximum capacity before the loss of LNG. Thus when the sudden drop in LNG occurs, LRS can do no more but continue to withdraw at its maximum rate. In contrast, Fig. 17 shows that MRS is not withdrawing at its maximum capacity in the base case. As a result, when the sudden drop in LNG occurs, withdrawals from MRS increase sharply for that week.

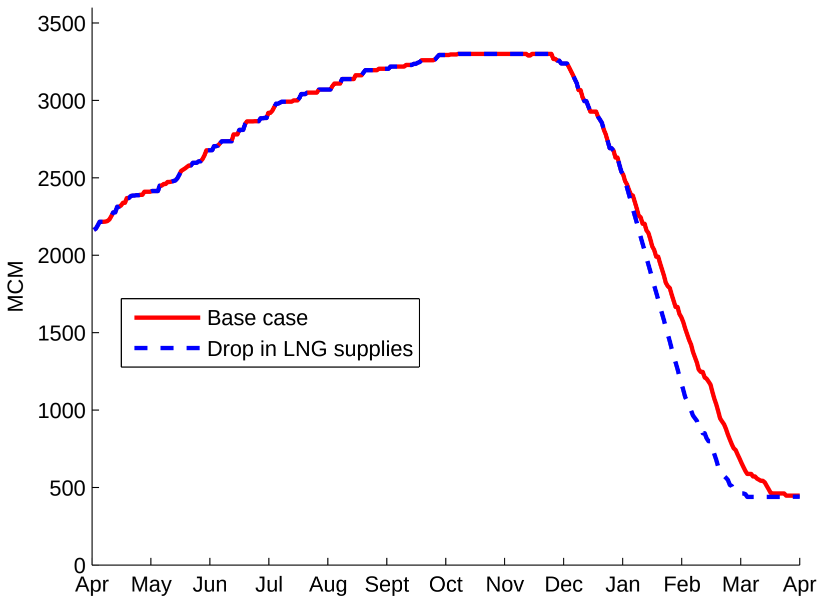

Fig. 16 Amount of gas in Long-Range-Storage as derived by ROM for both the base case and when there is a sudden drop in LNG supplies in January.

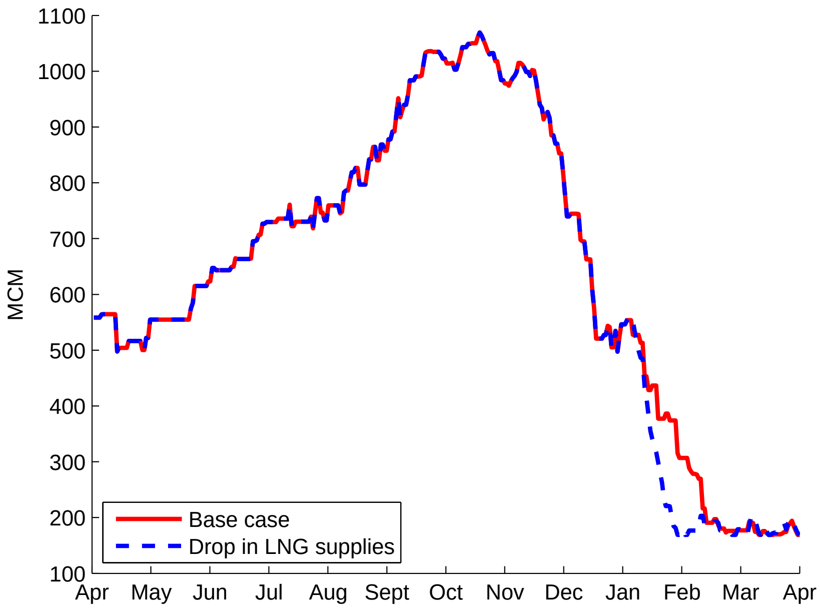

Fig. 17 Amount of gas in Medium-Range-Storage as derived by ROM for both the base case and when there is a sudden drop in LNG supplies in January.

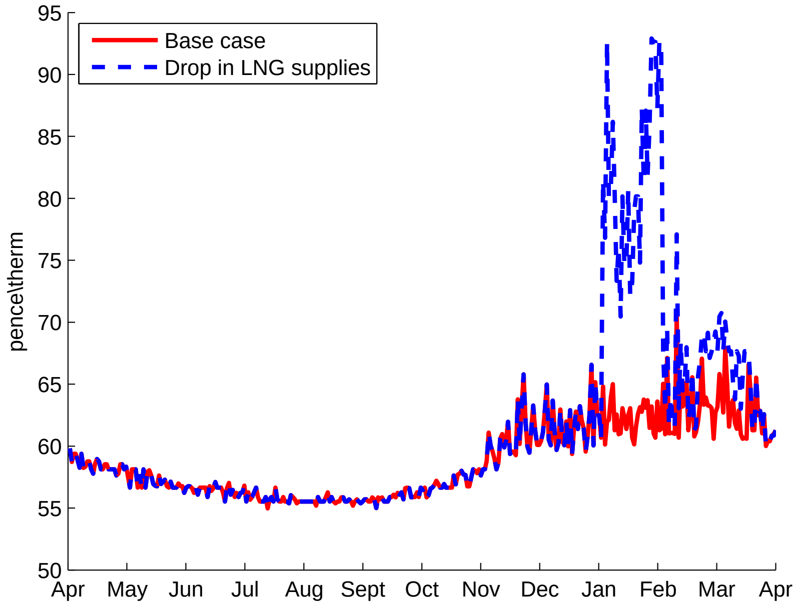

Fig. 18 System Average Prices obtained from ROM for both the base case and when there is a sudden drop in LNG supplies in January.

Fig. 18 displays the SAPs obtained from ROM for the base case and the case with a sudden drop in LNG capacity for a month starting in January. As expected, the SAPs are identical in the two cases up until the 1st of January. When the sudden loss in LNG capacity occurs, the SAPs rise. This increase persists past the 31st of January and are seen again later on in the year, although the second increase is not as severe. The reason for this prolonged increase is the reduced levels of gas available from MRS resulting from the loss of LNG in January. The reduced amount of gas in MRS means that more expensive sources of supply are needed to meet demand in the months following the reduction in LNG supplies, thus increasing SAPs.

4.3 Stress test 2: Particularly cold week in January

The second stress test to the UK gas market examined in this section is a particularly cold week at the start of January. This cold week is modelled by assuming a sharp rise in the level of demand. As in the previous section, the results for these analyses were obtained by applying the model to the parameters for the year beginning on the 1st of April 2013. In contrast to the previous shock, the demand scenarios used in the analysis of this section were developed using a time series consisting of Cold SND for a week in January and SND at all other times. By definition, SND is the demand expected in seasonal normal weather conditions, while Cold SND is the demand expected in particularly cold weather conditions. It is generated by the UK National Grid and is available from their website. Fig. 19 shows the time series used to develop the demand scenarios used in this analysis, along with SND. All the parameters used in this analysis are as described in Section 4.1.

Fig. 20 displays the demand profile obtained from ROM with a particularly cold week in January. When compared with Fig. 14, it indicates that withdrawals from MRS are the main source of supply to

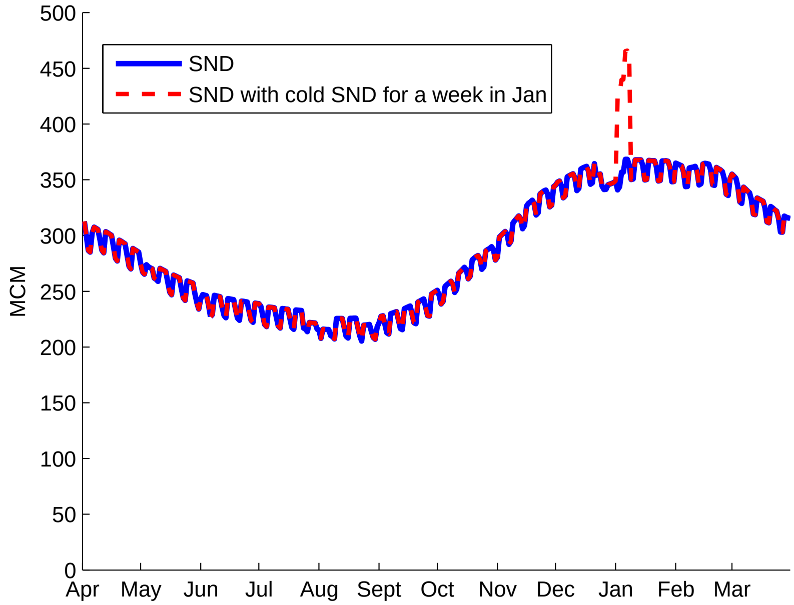

Fig. 19 Time series of seasonal normal demand and seasonal normal demand with a cold week at the start of January.

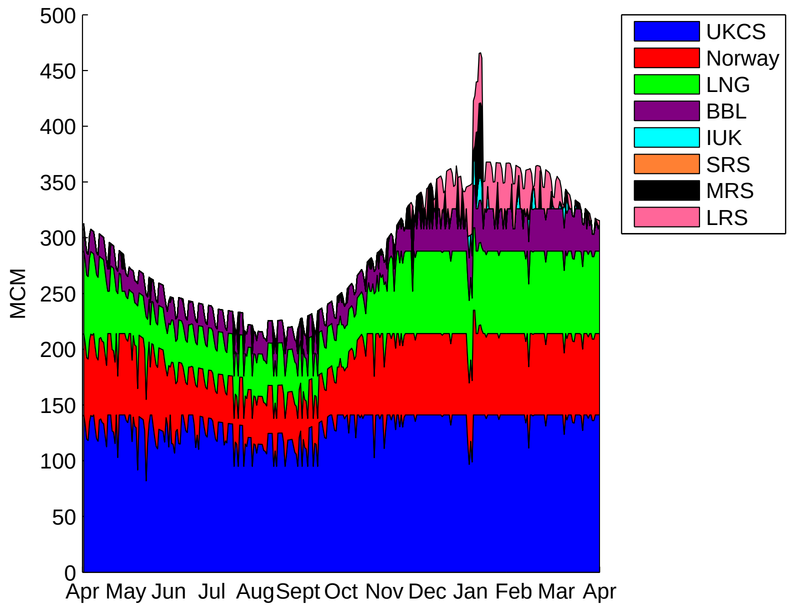

Fig. 20 Demand profile obtained from ROM for the UK gas market when there is a cold week in January.

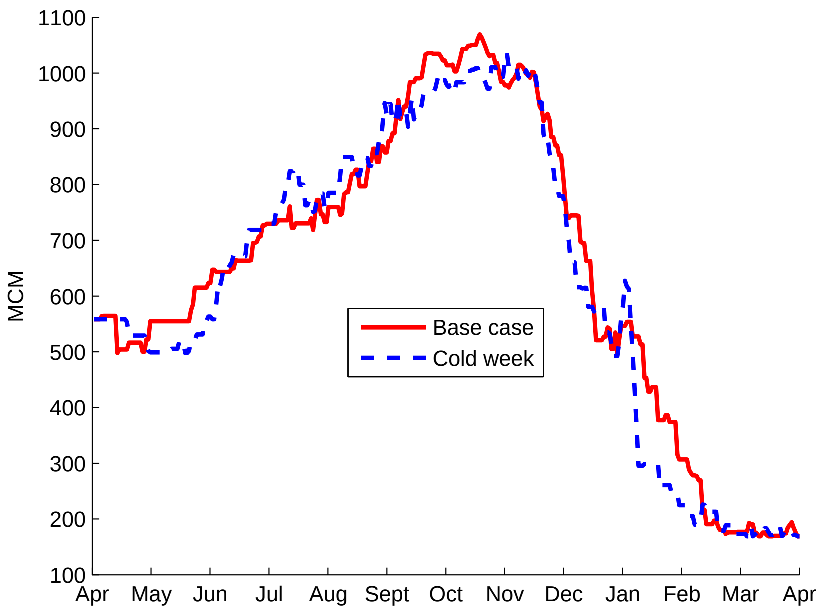

Fig. 21 Amount of gas in Medium-Range-Storage as obtained by ROM for both the base case and when there is a cold week in January.

increase as demand increases in the cold week. The amount of gas in LRS is largely unaffected by the particularly cold week in January. As result, similar results to Fig. 16 were found. Fig. 21 displays the amount of gas in MRS for the base case and the case with a cold week in January. It shows a sharp rise in the rate of withdrawals from MRS once the cold snap arrives. This is similar to Fig. 17, where there was a sudden drop in LNG supplies. However, in contrast to Fig. 17, Fig. 21 also shows a sharp rise in injections to storage before the cold snap. This occurs because the stochastic process for demand has a limited foresight of five days ahead (see Section 2). As a result, the model starts to take into account the increased demand five days before the cold snap. This means that MRS can prepare for the cold week five days before it happens, hence the sharp rise in injections to MRS before the 1st of January.

Fig. 22 displays SAPs obtained from ROM for the base case and the case with a particularly cold week in January. It indicates that the SAPs in the two cases are very similar up until the 27th of December (i.e., five days before the shock). Once information about the cold snap becomes available (via weather forecasts), an increase in the SAP can be seen for the cold week case. The initial rise in prices is as a result of the cost of the increased amount of gas injected to MRS. From the 1st of January the rise in SAP is due to the increased cost of meeting demand.

5 Summary and conclusions

In this paper a rolling optimisation model of the UK gas market was introduced. The aim of this work was to model the flows and prices of gas in the UK natural gas market, whilst incorporating the stochastic nature of demand. Previously, a rolling optimisation model of this market had not been seen in the literature. In Section 2 a stochastic process describing UK gas demand was introduced. This process reflects the daily information that is known regarding demand in the UK gas market: on the first day of

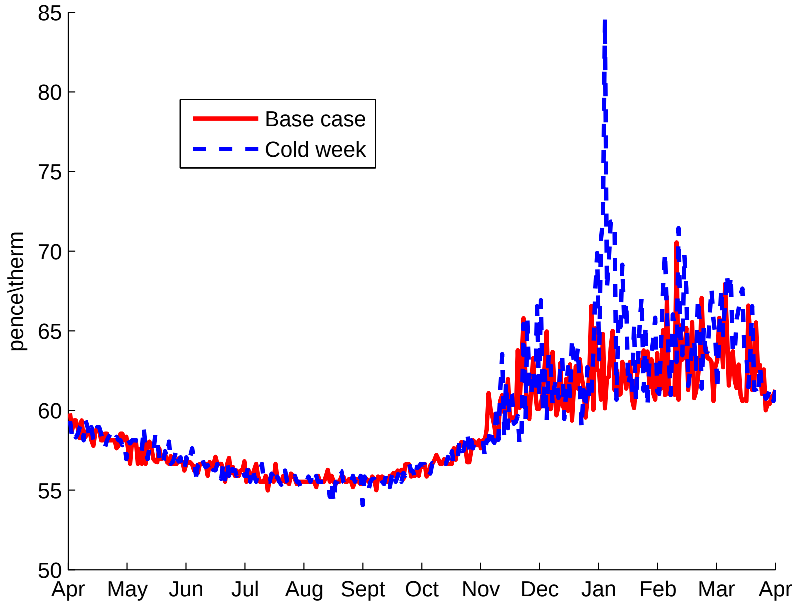

Fig. 22 System Average Prices obtained from ROM for both the base case and when there is a cold week in January.

the process, demand is exactly known. The next five days of the process are based on the relationships between the difference between actual demand and one- to five-day ahead predictions. After the sixth day, the process is based on the relationship between actual demand and SND.

In Section 3 the Rolling Optimisation Model (ROM) is described. It takes as an input simulated demand scenarios generated from the process described above. ROM simulates the daily decisions made in the UK gas market under uncertain information about future demand. The model enables decisions regarding the amount of gas to be injected to, or withdrawn from, the different storage facilities considered in the model, as well as how the different sources of supply meet the exactly known demand on the first day of the stochastic process. This is done at minimum cost whilst also ensuring all hypothetical future demands are also met at minimum expected cost. In Section 3.2, the parameters of ROM were chosen so as to best fit data from the UK gas market for the year beginning on the 1st of April 2010.

Using these parameters, flows of gas simulated by ROM were similar to actual flows. These flows illustrate the amount of gas in the different type of storage facilities seen in the UK, as well as how the different sources of supply meet demand. Using similar parameters, ROM was then tested in Section 3.3 for data from the UK gas market for the year starting on the 1st of April 2011. This analysis again showed that results obtained from ROM fitted reasonably well to actual data.

In Section 4 ROM was used to investigate the effect of various potential stresses on the UK gas market. This was done by applying the model to parameters for the year starting on the 1st of April 2013. The first stress test examined the effect of a sudden and complete cessation of LNG supplies into the UK for a month starting in January. This analysis found that supplies from MRS, IUK and Norway would make up the downfall in LNG supplies. The results also indicated that SAPs rise sharply once the shock to the market occurs with this rise lasting until after the return of full LNG supplies. The second stress test investigated the effect of a particularly cold week in January on the UK gas market. The results were similar to the first stress test. However, in contrast, ROM was able to anticipate this stress and prepared by injecting gas into MRS before the cold snap arrived.

Throughout the paper, the limitations of ROM are discussed. Further works will attempt to address these issues which include

- modelling long-term contracts and non-linear costs,
- development of risk preferences for the (imaginary) central planner,
- modelling the lag between the production and consumption of natural gas, particularly for LNG due to its global nature.
- adapting the stochastic process model of demand to take into account increasing supply of wind energy.

Acknowledgements This work is funded by Science Foundation Ireland under programmes MACSI 06/MI/005 and 09/SRC/E1780. The authors would also like to thank Bord Gáis Energy for their contributions, in particular Gavin Hurley.

A Analysis of Karush-Kuhn-Tucker conditions

In this section, a small algebraic example illustrates the fact that changing the production costs of ROM, without altering the merit order, does not affect the volumes but does change the marginal supply cost. Consider the following problem:

$$\min\ c_1 Q_1 + c_2 Q_2, \tag{18}$$

subject to:

$$Q_1 \leq Q_1^{max}, \ \ (\lambda_1), \tag{19}$$

$$Q_2 \leq Q_2^{max}, \ \ (\lambda_2), \tag{20}$$

$$Q_1 + Q_2 = Demand, \ \ (\lambda_D), \tag{21}$$

where $Q_{1,2}$ represent production levels, $c_{1,2}$ represent the costs associated with them and represent $Q_{1,2}^{max}$ maximum production levels. The variables in the parentheses, alongside constraints (19) - (21), represent the Lagrange multipliers associated with that constraint. The Karush-Kuhn-Tucker conditions for optimality (Bazaraa et al (1993)) associated with this problem are:

$$c_1 + \lambda_1 + \lambda_D = 0, \tag{22}$$

$$c_2 + \lambda_2 + \lambda_D = 0, \tag{23}$$

$$\lambda_1 (Q_1 - Q_1^{max}) = 0, \tag{24}$$

$$\lambda_2 (Q_2 - Q_2^{max}) = 0, \tag{25}$$

$$\lambda_1 \geq 0, \tag{26}$$

$$\lambda_2 \geq 0, \tag{27}$$

as well as constraints (19) - (21). These conditions show that the optimal production levels depend on the production capacities and the demand whereas the price associated with meeting demand (λ_D) depends on costs.

References

Abada I, Gabriel S, Briat V, Massol O (2012) A generalized Nash-Cournot model for the northwestern European natural gas markets with a fuel substitution demand function: The GaMMES model. Netw Spat Econ **13**:1–42

Abrell J, Weigt H (2012) Combining energy networks. Netw Spat Econ **12**:377–401

Bazaraa M, Sherali H, Shetty C (1993) Nonlinear Programming Theory and Algoritims. John Wiley and Sons, New York

Boots M, Rijkers F, Hobbs B (2004) Trading in the downstream European gas market: A successive oligopoly approach. The Energy Journal **25**:73–102

Bopp A, Kannan V, Palocsay S, Stevens S (1996) An optimization model for planning natural gas purchases, transportation, storage and deliverability. Omega - International Journal of Management Science **24**:511–522

Butler J, Dyer J (1999) Optimizing natural gas flows with linear programming and scenarios. Decision Sciences **30**:563–580

Devine M (2012) Rolling optimisation, stochastic demand modelling and scenario reduction applied to the UK gas market. PhD thesis, Department of Mathematics and Statistics, University of Limerick, Ireland.

Egging R, Gabriel S (2006) Examining market power in the European natural gas market. Energy Policy **34**:2762–2778

Egging R, Holz F, Gabriel S (2010) The World Gas Model: A multi-period mixed complementarity model for the global natural gas market. Energy **35**:4016–4029

Gabriel S, Manik J, Vikas S (2003) Computational experience with a large-scale, multi-period, spatial equilibrium model of the North American natural gas system. Netw Spat Econ **3**:97–122

Gabriel S, Kiet S, Zhuang J (2005a) A mixed complementarity-based equilibrium model of natural gas markets. Operations Research **53**:799–818

Gabriel S, Zhuang J, Kiet S (2005b) A large-scale linear complementarity model of the North American natural gas market. Energy Economics **27**:639–665

GAMS (2012) GAMS - A user's guide. GAMS Development Corporation, Washington, DC, USA, available at: http://www.gams.com/dd/docs/bigdocs/GAMSUsersGuide.pdf Accessed 30 November 2012

Holz F, von Hirschhausen C, Kemfert C (2008) A strategic model of European gas supply (GASMOD). Energy Economics **30**:766–788

Honoré A (2010) European Natural Gas Demand, Supply, and Pricing. Oxford University Press

House of Commons Business and Enterprise Committee (2007) Energy prices, fuel poverty and ofgem. eleventh report of the house of commons' business and enterprise committee of the 2007-08 session. Available at: http://www.publications.parliament.uk/pa/cm200708/cmselect/cmberr/293/293ii.pdf Accessed 30 November 2012

Kirkpatrick S, Gelatt C, Vecchi M (1983) Optimization by simulated annealing. Science **220**:671–680

Lise W, Hobbs B (2008) Future evolution of the liberalised European gas market: Simulation results with a dynamic model. Energy **33**:989–1004

Lochner S, Bothe D (2009) The development of natural gas supply costs to Europe, the United States and Japan in a globalizing gas market-model-based analysis until 2030. Energy Policy **37**:1518–1528

Metropolis N, Rosenbluth A, Rosenbluth M, Teller A (1953) Equation of state calculations by fast computing mechanics. The Journal of Chemical Physics **21**:1087–1092

Perner J (2002) Die Langfristige Erdgasversorgung Europas, Schriften des Energiewirtschaftlichen Instituts, Band 60, Munchen.

Perner J, Seeliger A (2004) Prospects of gas supplies to the European market until 2030-results from the simulation model EUGAS. Utilities Policy **12**:291–302

Seeliger A (2006) Entwicklung des weltweiten Erdgasangebots bis 2030-Eine modellgestutzte Prognose, Oldenburg Verlag, Munich.

Siddiqui S, Gabriel S (2012) An SOS1-based approach for solving MPECs with a natural gas market application. Netw Spat Econ **13**:205–227

The Baker Institute World Gas Trade Model (2005) Available at: http://www.bakerinstitute.org/publications/the-baker-institute-world-gas-trade-model-biwgtm. Accessed 30 April 2013

UK National Grid (2011a) UK National Grid's development of energy scenarios. Available at: http://www.nationalgrid.com/NR/rdonlyres/C934D455-1438-4949-A7CD-3FE781FE39E8/47855/DevelopmentofEnergyScenariosTBE2011.pdf Accessed 30 November 2012

UK National Grid (2011b) UK National Grid's future energy scenarios, Available at: http://www.nationalgrid.com/uk/Electricity/Operating+in+2020/UK+Future+Energy+Scenarios.htm Accessed 30 November 2012

UK National Grid (2011c) UK National Grid's gas ten year statement. Available at: http://www.nationalgrid.com/NR/rdonlyres/E60C7955-5495-4A8A-8E80-8BB4002F602F/50703/GasTenYearStatement2011.pdf Accessed 30 November 2012

UK's Department of Energy and Climate Change (2010) Digest of United Kingdom energy statistics (DUKES) Available at: http://www.decc.gov.uk/assets/decc/Statistics/publications/dukes/348-dukes-2010-printed.pdf Accessed 30 November 2012

UK's Department of Energy and Climate Change (2011a) Digest of United Kingdom energy statistics (DUKES) Available at: http://www.decc.gov.uk/assets/decc/11/stats/publications/dukes/2310-dukes-2011-internet-booklet.pdf Accessed 30 November 2012

UK's Department of Energy and Climate Change (2011b) Electricity market reform white paper. Available at: http://www.decc.gov.uk/assets/decc/11/policy-legislation/EMR/2210-emr-white-paper-full-version.pdf. Accessed 30 November 2012

UK's Department of Energy and Climate Change (2012) Digest of United Kingdom energy statistics (DUKES) Available at: http://www.decc.gov.uk/assets/decc/11/stats/publications/dukes/5954-dukes-2012-chapter-4-gas.pdf Accessed 30 November 2012