ISSN 1755-5361

University of Essex

Department of Economics

Discussion Paper Series

No. 658 July 2008

Evaluating the Economic Cost of Strategic Storage of Natural Gas

João Miguel Ejarque

Evaluating the Economic Cost of Strategic Storage of Natural Gas

João Miguel Ejarque*

July 22, 2008

Abstract

The European Comission wants to implement a single market for gas. One of the components of this market is a regulated provision for "security of supply" which consists of rules for the implementation and use of a given reserve stock of gas. We investigate the impact of this policy on the profitability of a storage operator, using data from Denmark and Italy. Keeping storage capacity constant, the costs of the strategic stock are around 20% of the value of the storage market for Denmark, and 16% for Italy. This cost is due to the inability to extract arbitrage profits from the captive stock.

JEL Classification: D92, E32, Q48

*Correspondence to: João Miguel Ejarque, Department of Economics, University of Essex, Colchester CO4 3SQ, UK. Email: jejarque@essex.ac.uk.

1 Introduction

The European Commission wants to implement a single market for gas. One of the components of this market is a regulated provision for "security of supply" which consists of rules for the implementation and use of a given reserve stock of gas. This paper investigates the impact of this policy on the profitability of a storage operator, and therefore indirectly on the incentive to invest in storage.

In order to implement the creation of strategic storage the policy maker regulates the use of the existing stocks owned and managed by commercial operators. This regulation takes the form of a penalty on the "improper" use of the last x% of total capacity in any storage facility. Proper uses are determined by the regulator. The policy analysis here is the comparison of optimal behavior and profits in a market with and without the penalties.

Regulatory intervention is justified in the presence of externalities. Just like the European Commission, this paper assumes that market failures and externalities exist in the market for natural gas, such that it is important for the policy maker to complement the commercial supply of gas storage with an extra stock to be released in case of emergency. In this line Wright and Williams (1982,a,b) note that one of the reasons why storage is needed is because political economy motives will force the prices down (and therefore increase demand) during peak season or during disruptions. In this way, (some) storage arises as a secondary distortion to correct an original, polically motivated one. On the other hand Le-Coq and Paltseva (2008) suggest that storage can induce Moral Hazard as agents consume gas and use their infrastructure in a less than prudent way. These issues are absent from the current analysis.

Another maintained assumption in this paper is that the storage part of the market for natural gas can be unbundled from production and distribution activities and function in a competitive environment. Storage facilities are investments with large initial fixed costs, but casual observation reveals that nevertheless a large number of agents participate in the market.[1]

These assumptions allow us to write a model of optimal storage of gas where the agent (a storage firm) buys gas and stores it when prices are low, and sells gas when prices are high. The model is calibrated to match Italian and Danish storage data. Both of these countries have restrictions on the use of stored gas. Italy sets apart around 38% of its stored gas on a permanent basis, while Denmark imposes that stored gas must at all times be enough to cover the following sixty days of normal consumption. While large, Italy's storage capacity is smaller relative to total consumption than Denmark's, which comfortably meets its sixty day restriction.[2]

We want to know the impact of these policies on the behavior of storage operators and on market outcomes. It is useful to compute a back of the enve-

[1] This is true in germany and the UK. Denmark and Italy data are both monopoly storage countries for most of the sample but they are linked to the international market.

[2] Possibly a better way to implement a strategic storage policy is for the regulator to be a client and buy gas for its specific purpose, storing it in existing commercial sites. This may be the least distortionary way of pursuing the policy goal, but we do not examine it here.

lope example to understand the potential magnitude of the problem. Suppose a country has around 40% of its total storage capacity reserved for strategic purposes. Suppose also that summer prices are on average around half of winter prices (as is the case for the UK National Balancing Point price). Consider a storage site with unit capacity that fills up with gas at the beginning of time and then buys (low) and sells (high) every half year the 60% available capacity it has to trade. Therefore in period one the firm buys the entire stock at price 1/2. Then one semester later sells the fraction $(1-\lambda)$ at price one, and then alternates buying and selling the fraction $(1-\lambda)$ at prices 1/2 and 1. The value of this firm is then given by:

$$V(\lambda) = -\frac{1}{2} + (1-\lambda)\frac{\beta}{1-\beta^2}\left(1 - \frac{\beta}{2}\right)$$

With a six-month discount factor of $\beta = 0.96$, we obtain $V(0) = 5.87$ and $V(0.4) = 3.32$, a difference of 44%. As a rule of thumb in this example the percentage difference is identical to the fraction λ if we ignore the initial expenditure:

$$\frac{V(0) - V(s)}{V(0) + 1/2} = \lambda$$

This is the cost of the forced savings imposed on this market. To have an idea how this number compares with normal precautionary behavior, we can think of the required reserve ratio imposed on commercial banks by a central bank, or on the fraction of income families save due to precautionary reasons. These are numbers far smaller than 40%. This back of the envelope exercise suggests a potential inefficiency is being imposed on the gas market.[3]

Clearly, the actual cost of this restriction is not given by these numbers and a more accurate solution to this problem requires a better model. That is the purpose of the model developed in this paper. Our calibration for Italian and Danish data yields a cost of regulation of around 16% and 20% of discounted net present value of profits for the respective storage markets. These costs are due to the loss of seasonal arbitrage profit associated with the captured stock.

The current paper is a measurement and policy evaluation exercise on natural gas storage, and to our knowledge there is no close exercise in the literature. The model draws on work by Byers (2006), Thomson, Davidson and Rasmussen (2007), Hall and Rust (2000) and Chaton, Creti, and Villeneuve (2007a,b). These last authors study a problem where storage agents face the possibility of a large disruption and therefore build up a precautionary stock up to a given level and then optimally maintain it. They do not study seasonal arbitrage nor do policy evaluation which is the main focus of this paper. Casassus, Collin-Dufresne and Routledge (2005) study a general equilibrium model of oil as an intermediate input with storage aimed at replicating the asset pricing properties of commodities, but again the seasonal component is absent from the continuous time framework.

[3] A different way to think of this is that if storage equals normal savings for the gas market, a λ of 0.4 implies we force people to increase savings by a factor of 2/3. A λ equal to 0.2 implies a forced increase in savings of a factor 1/4.

2 Data

The model is calibrated to match Italian and Danish data. Figure 1 shows the large positive monthly Italian gas consumption and imports, and closer to zero, production and storage flows (measured as consumption less imports and production). The data is in millions of cubic meters from January 2002 to December 2006, from the OECD.[4] Production is flat and steadily declining while imports and consumption are increasing and seasonal. The seasonality of imports shows that a big part of the adjustment to the large consumption swings is met this way rather than by storage. In Figures 1 and 2, storage flows are positive when injection is taking place, which is generally from April to October.

Figure 1: Italian Gas Data, OECD

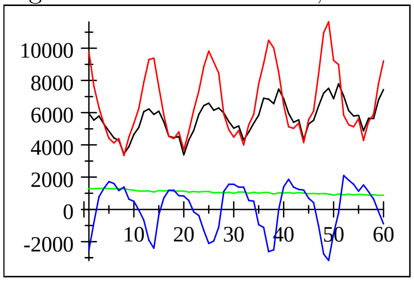

One of the key assumptions used in the paper is that of a competitive market. It is useful to look at Danish data to support this assumption:

Figure 2: Danish Gas Data, OECD

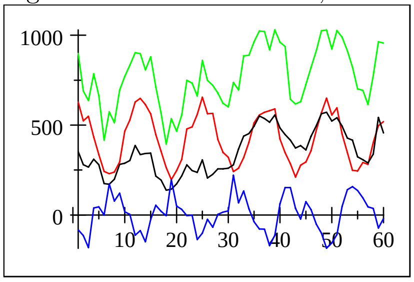

The key difference is that Denmark exports gas, while Italy imports it. Gas production is here the upper most series with consumption closely following its cycle. Exports are also seasonal and increase a little over the sample. Denmark explores its share of the Ormen Lange field, which in turn is connected to the U.K. market by pipeline. Since the U.K. market has a well functioning

[4] The year 2002 has a significant (unaccounted for) difference between flows computed as stock differences, and flows computed as consumption less production and imports. They should be the same as Italy exports almost zero gas.

spot market for gas, and is further connected to the continental one via the interconnector, and since Denmark also has pipelines to the continent, it is safe to assume that competitive spot prices affect the marginal unit of gas stored in Denmark. The final logical step here is that, while we do not have information on spot trading in either Denmark or Italy, the behavior of flows in and out of storage is broadly similar in both countries, justifying the use of a competitive assumption to look at the data.

We will make use of this monthly data from the OECD as well as of daily data available directly from storage operators STOGIT (Italy), and DONG and Energinet (Denmark). There is one additional reason to look at both Danish and Italian data: they both have strategic storage requirements. These requirements are different from each other which suggests a policy comparison in itself.

It is useful to isolate the annual cycles in this data. Figure 3 shows the annual cycle of Italian and Danish gas consumption. These are monthly relative values, January through December, and months are numbered one through twelve on the horizontal axis.[5] Consumption has a single peak in January, and the two series are similar. Denmark has its lowest consumption in July (month 7).

Figure 3: Consumption Cycle

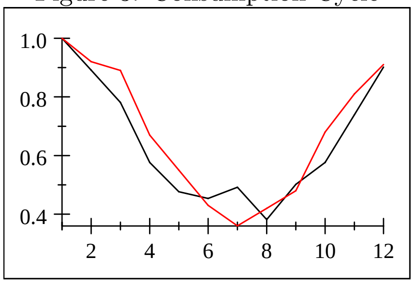

Figure 4 shows the annual cycle of Italian and Danish gas storage flows, normalized. We can see that the largest withdrawal months are January and February, and the highest injection month is July.[6]

Figure 4: Cycle of Storage Flows

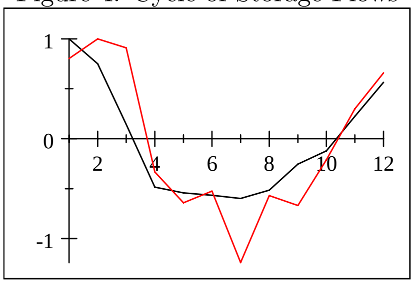

[5] Average the 5 observations for each month. Divide the 12 averages by the max average.
[6] Stocks peak in late October (October 24, 2006 for Denmark) bottom late March or April.

We want the model to generate this pattern of flows. The model takes as driving shock the price of gas to generate endogenous quantity flows. We will make use of the fact that competitive spot prices behave similarly to quantities in our construction of the price process. In Figure 5 we can see the cyclical behavior of the National Balancing Point price data from the UK. One measure averages monthly prices in the years 2002, 2003 and 2004, another averages prices in the years 2005, 2006, 2007, and the third one has the overall average price for each month using all six years. The first three years are much smoother than the last three years. But the *relative* variations in price are not that different. Given that the UK market is the most liquid market for gas in Europe it is not surprising to see this price vary over the calendar year. The fourth line in this graph (highest at 1 and lowest at 12) shows the seasonal cycle of storage flows for the UK. It is normalized to be between 0.5 and 1, but its pattern resembles that of the Danish data. Importantly, it also resembles the price data.

Figure 5: NBP Gas Prices, 2002 to 2007

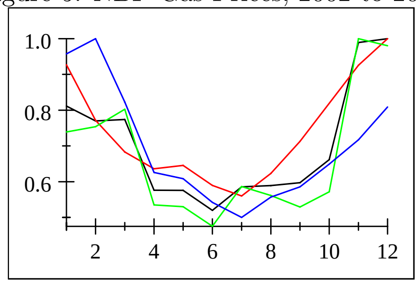

If we are successful, quantitatively and qualitatively, in generating this cyclical behavior then we can use the model to experiment with different policies to evaluate their impact.

3 Model

This is a model of physical storage operation. The main features of this problem are studied in the large literature on inventories. The problem of the storage operator (henceforth labelled interchangeably as firm or agent) is to decide how much to buy (inject, or add to inventory) and/or sell (withdraw, or take from inventory) each period, in order to maximize the expected present value of profits.[7]

We consider a period to be one month since we aim to study annual seasonal variations. In each period the firm has profits given by:

$$\pi_t = P_t x_t - h\left(x_t, S_t\right) - k(S_{t+1})$$

In this expression P_t denotes the realization of the random spot price of gas in month t and x_t denotes purchases or sales of gas in month t. The firm cannot both buy and sell gas in a given month. Therefore x denotes net purchases each period. If x takes a positive value the firm is selling gas and therefore stocks are falling.

The function $h(x, S)$ specifies arbitrarily large costs of exceeding physical injection and withdrawal limits ($\bar{x}$). This is a practical way of imposing this technical constraint on the computational problem.

$$h\left(x_t, S_t\right) = M_p \times \left(x_t > \bar{x}\right) + M_n \times \left(x_t < -\bar{x}\right)$$

The last element of the profit function, $k\left(S_{t+1}\right)$, is the penalty function imposed by the regulator for the misuse of strategic gas:

$$k\left(S_{t+1}\right) = \left[k_0 + k_1\left(\lambda - S_{t+1}\right)\right] \times \left(\lambda > S_{t+1}\right) \times \left(1 - IN_t\right)$$

Since we normalize the capacity of the facility to $\bar{S} = 1$, we define the fraction of gas regulated by the value $\lambda \bar{S} = \lambda$. The policy parameters are k_0 and k_1. The regulation takes the form of a penalty function for the use of gas in excess of $(1 - \lambda)$, and this penalty function is a two part (linear) tariff with a fixed and a proportional component. We will look at two specifications for λ. In the case of Italy this is simply a constant, while in the case of Denmark, this varies over the year and must be enough to always cover the expected gas consumption of the following sixty days.

The timing of $k\left(S_{t+1}\right)$ is important here. The penalty is not always enforced, since it is lifted in case of extreme events detailed in the policy description. Such events are modelled by setting the indicator function IN_t to take the value one. This is also the event when prices increase sharply in our model.[8] Now, if a catastrophe state does not occur today, you need to worry about how much stock you leave for next period. But if a catastrophe state occurs, the penalty

[7] Even though gas storage operators are subject to third party access regulations, selling the storage service to a third party should be equivalent to buying and selling gas.

[8] It is possible to study scenarios of regulatory uncertainty, whereby (k_0, k_1) are random variables with given distributions. We do not examine this case at the moment.

is lifted so that you are allowed to start next period with a smaller stock than the normal strategic requirement. Here the nature of shocks and the penalty interact to affect the decision of how much of the strategic stock to sell. If the penalty for illegal selling is very large the firm will never sell too much to be in a position of violating the constraint two periods ahead, by an inability to build the stock back up.

The sale or purchase of gas each period determines the evolution of inventory. This is given by

$$S_{t+1} = \min \begin{cases} \bar{S} \\ \max \begin{cases} S_t - x_t \\ 0 \end{cases} \end{cases}$$

where $0 \leq (S_{t+1}, S_t) \leq \bar{S}$, and x takes values in the interval $[-\bar{x}, \bar{x}]$. This bound on flows is not necessarily symmetric around zero.

We are now ready to write the dynamic programming problem in terms of the Bellman equation:

$$V(P_t, S_t) = \max_{x_t} \{\pi_t + \beta E_t V(P_{t+1}, S_{t+1})\}$$

where E_t denotes the expectation operator given time t information and β is the discount factor.

The firm chooses x_t to maximize the value of entering the period with a given stock S_t and facing a price realization P_t. The decision x_t is taken after observing P_t and knowing all parameter values and the stochastic process for prices, but not knowing future price realizations. This is a well behaved problem. It is a bounded problem with both x and S bounded above and below. Profits are continuous in x and S. Therefore there is an optimum and it is unique.

The model ignores financial hedging. This is not a large limitation since futures contracts are effective against short term price variability. If agents contract in the summer period for deliveries in winter, they effectively trade away short term future variability, but cannot trade seasonal variability. These contracts are therefore written on the conditional expectation of market conditions, which is the low frequency seasonal variability behind large capacity storage. We consider the model as representing the economics of storage over and above hedging.[9]

3.1 The stochastic process for prices.

The price process is modelled as a Markov process with three components: temporary shocks, catastrophe shocks, and seasonality.

Seasonality is modelled as a deterministic change in the average price from month to month. Prices are low in summer months and high in winter months. We specify a given value for the seasonal price in each month and the transition from month to month (and average price to average price) occurs with

[9] Hedging may be important in our context of large disruptions. The insurance market does trade in catastrophe bonds.

probability one. Specifically, this process has a 12×12 transition matrix with ones in the non main first diagonal, a one in the bottom left corner, and zeros elsewhere: the element $a_{i,i+1} = 1$, $i = 1 : 11$, on row (i) and column (i+1) is the probability of going from month (i) this period to month (i+1) next period. The bottom left element, $a_{12,1} = 1$, is the probability of going from state twelve (December) this period, to state one (January) next period.

The temporary process is a discretization of a first order autoregressive process included to generate small price variations each period. This process has an unconditional mean of one and multiplies the seasonal price such that observed prices are realizations around the seasonal average. For a variety of reasons storage operators may not hedge their entire gas stock with futures contracts, implying they are subject to some short run price variability. By changing the standard deviation of this shock we can proxy for the degree of hedging the storage facility is able and/or willing to obtain.

The final component of the price process is a catastrophe shock. In this paper this consists of the very small probability of a very large increase in price (relative to the normal seasonal average).

3.2 Discussion

The optimal decision has a simple "bang-bang" benchmark: buy as much as possible when prices are low until storage is full, then wait inactive (if necessary) until prices rise enough to sell, and then sell as much as possible until storage is empty. After that, wait again until prices fall enough to buy. Actual behavior will deviate from this pattern for two reasons.

First, in case the price process is not monotonic. If there are short periods of high prices during the buying season, the optimal decision is to sell in those short periods. With a daily frequency in the model we would be able to replicate the data almost exactly. Daily data would then aggregate into a smooth monthly cycle. Solving the model with a monthly frequency implies any non monotonicity will deliver an entire month of selling during the buying season or vice versa, which clearly we do not see. Nevertheless, a weekly calibration of this model aggregates into a reasonable monthly pattern, suggesting the model can indeed be used for its purpose of policy evaluation.

Second, deviations from this pattern can occur for precautionary reasons. Here the catastrophe shock is important. Perhaps counter intuitively, a shock during the high winter months (November, December, January, February) has little or no effect on optimal behavior, and that shock is not what should worry policy makers. The reason is that storage is typically full at the end of October and then withdrawal starts. *The problem is not whether we will freeze in the winter, because storage is usually full in those months.* The problem is the occurrence of a large shock in late March or April when stocks are at their lowest. Unless we have the possibility that prices in March will be, for example, ten times higher than very high prices in January, withdrawal will take place at full capacity during the winter months, since that is when prices are normally high. Still the model has scope for a change in behavior because prices are observed at the beginning of the period, and then withdrawal (selling) decisions are taken. If there is no price spike in November, perhaps a full capacity withdrawal will not be optimal. But if the probability of a large disruption is very small (as is the spirit of these disaster scenarios) it is unlikely that optimal behavior will change significantly.

4 Calibration

Before we use the model we need to define values for its parameters.[10] Our discount rate implies a required rate of return on investment which we set at 10%. This is relatively high but we assume higher returns are required for investments with very large fixed costs such as gas storage sites. The monthly discount factor is $\beta = 0.90^{1/12}$. The cost function parameters (M_n, M_p) are arbitrarily large and so are easy to implement. The parameters (k_0, k_1) are also unclear from the data. We do not know how much the legal process from a March 2006 illegal gas sale is costing the storage companies in Italy, so we set (k_0, k_1) high enough so that the constraint is never violated.

The total storage capacity is the numeraire parameter, $\bar{S} = 1$. The strategic storage lower bound is a constant $\lambda = 0.378$ for Italy. The maximum gas stock over the STOGIT sample is 529.2 million gigajules, recorded on in October 10, 2006. Of these, 200 million are permanent strategic stock. The lower bound for Denmark is given by a function of the cycle, $\lambda(n)$. The Danish government has a security of supply goal of 60 days because this is an estimate of how long it would take to repair a major failure in the underwater pipeline coming from Norway.[11] This is a time varying strategic stock requirement, as 60 days of summer consumption differ from 60 days of winter consumption. It is not clear that 60 days is optimal in any way, as private agents may have efficient ways of dealing with scarcity which will be undiscovered if supply is secured at all costs.

Using the OECD data we compute the maximum stock level over the entire sample and compare it to the realized *consumption* of the current and following month. Then we average these for each month over the years in the sample. We also compute this function $\lambda(n)$ for Italy to compare this policy variant to the status quo. One other experiment is to consider only a restriction of the 30 day-ahead consumption, rather than 60 days.

Table 1 has the measure of how the 60-day ahead consumption compares with the maximum stock (capacity).[12]

Table 1: Seasonal Strategic Constraint (60 days)

	J	F	M	A	M	J	J	A	S	O	N	D
DK	0.53	0.51	0.45	0.35	0.28	0.22	0.22	0.25	0.31	0.40	0.46	0.51
IT	0.99	0.89	0.72	0.55	0.49	0.50	0.46	0.46	0.56	0.69	0.87	0.99

The (deservedly) much vaunted Italian system cannot cope with this policy

[10] Any fixed payment per period from the policy maker is indistinguishable from a fixed operating cost and we set it all to zero.

[11] From the website of ENERGINET: "The two storage facilities are dimensioned and placed such that they can supply the firm gas market with natural gas for a period of approximately 60 days which is the estimated time required to repair the gas pipelines in the North Sea."

[12] This takes the data on consumption for the current plus next month, and divides it by the max over the series of the "opening stock level". For DK it uses data from January 2003 to December 2006. The 2002 and 2007 seems to be less reliable and affects this computation more than others. For consistency the same period was used in Italian data. Each monthly value in the table is then the average over four years (four values) of the same month.

measure, much less a 90 days constraint being vented in the European Union. At the start of January, 99% of Italy's total storage capacity would not be allowed to be used as prevention for the catastrophe shock. This does not leave much gas to handle the seasonal cycle.

We now need to define injection and withdrawal capacities. For this purpose we use the daily storage data from STOGIT, DONG and Energinet. Everything is measured relative to total capacity. The maximum level of the stock recorded in the Italian data is of 329.2 MGJ, and the minimum is -45.07, recorded in March 17, 2006. Negative values are recorded when stocks fall below the maximum volume of strategic storage which is 200 MGJ. Therefore, while some of the strategic stock was sold, it was by no means the entire amount.

Injection and withdrawal flows are not symmetric. This is true of the 1461 days of observation from STOGIT: 919 days (or 63% of observations) have positive flows (injections), 6 days have a zero flow, and 536 (or 37%) have a negative flow (withdrawal). To have an idea of what this implies in terms of capacity, we take the daily data on net flows and compute the *moving sum* of the past 30 days of flows. This is not a calendar month total net flow, but a daily measure of the flows in the preceding 30 days, and we have 1432 observations of these sums for Italy, while we have only 48 months in STOGIT data. We then compute the maximum and the minimum of these moving sums. This maximum and minimum are then divided by the maximum daily stock over the entire sample (529.2 in the case of Italy). This maximum stock attained over the sample is a proxy for maximum capacity which in the model is normalized to one. We use the same procedure on the Danish data. The ratios we compute then give us a measure of maximum injection and withdrawal capacity in a given month as a fraction of total storage capacity.[13]

Table 2: Injection and Withdrawal Constraints

$\bar{x}$	Injection	Withdraw
DK	0.3694	0.2426
IT	0.1687	0.2199

It seems to be harder to inject than to withdraw in Italy, but the reverse in Denmark. In fact daily Italian injection capacity is listed as being of 5.122 MGJ per day while withdrawal capacity is listed as being slightly higher at 5.910 MGJ per day. These numbers suggest higher withdrawal and injection than what we observe.[14] One reason we observe less than full capacity injection and withdrawal when we look at monthly data is the time aggregation of short run

[13] If we multiply these numbers by 7/30 we obtain the equivalent values for weekly frequencies. We use a weekly calibration to see if it improves the fit of the model. It does. However, the fit of the model in the monthly experiments presented below was quite reasonable and so we report only this frequency. The weekly calibration is quite time consuming in the computer.

[14] Multiplying by 30 these numbers imply a total injection capacity of 153, and withdrawal capacity of 177 per month but we never see such numbers in monthly data. On the other hand, in daily data the biggest observed injection day is of 6.58 MGJ and the biggest withdrawal is of -6.38, *higher* than the listed capacity limit.

arbitrage which sees injection days during withdrawal months and vice versa. In the end, due to the numerical difficulty of working with daily frequencies, we impose a capacity limit based on what is observed.

4.1 Calibration of the Stochastic Process

The model is a filter that takes exogenous stochastic prices and generates endogenous behavior for quantities (stocks and flows). Our aim is to *replicate* the behavior of stocks and flows. The price process has three components, seasonal, catastrophic, and short run shocks:

$$P_t = P_t^S P_t^C P_t^s$$

Short run variations, P_t^s, are modelled as a log normal process $P_t^s = \exp\{z_t\}$, with $z_t = \rho z_{t-1} + u_t$ and $u_t \sim N(0, \sigma)$.

This shock process depends on the frequency used in the model. Using the gas flows series from the OECD data up to December 2007 we get six cycles. We run an OLS regression of the 72 observations of monthly gas flows (normalized by their max) against a constant, a trend, and eleven monthly dummies. For Italy the R squared is 0.85, the F test for the seasonal dummies is 29, and the residual is relatively small and flat. Nevertheless it has a significant first order serial correlation (ρ) of 0.40. The standard deviation (σ) of this residual is 0.1645. For Denmark this regression procedure yields an R squared of 0.79, an F test of 20, a residual standard deviation of 0.20 and a residual serial correlation of 0.11.

Table 3: Short run shock process

z_t	ρ	σ
	Month	Month
DK	0.11	0.20
IT	0.40	0.16

The seasonal process, P_t^S, has the transition matrix detailed above. We now need to define its support. Our seasonal price is normalized to reach a maximum of 1 but we need to know how this unconditional expectation evolves over the year.

We saw earlier that competitive prices behave over the cycle in a way similar to quantities. We use this fact to justify taking the actual behavior of *quantities* to compute the seasonal cycle for the *price* to be fed to the model. We also want prices to fluctuate within the boundaries of the NBP price. Therefore we use a simple process of normalization for annual average cycles using monthly aggregates. Denote the raw series of monthly net gas storage flows by $\{z_t^0\}$. This can take positive as well as negative values. Construct $z_t^1 = z_t^0 + abs(\min(z_t^0))$. The resulting series $\{z_t^1\}$, is now bounded below by zero. Then construct $z_t^2 = z_t^1 + \max(z_t^1)$, and finally $z_t^3 = z_t^2 / \max(z_t^2)$. The final series $\{z_t^3\}$, is bounded below by 0.5 and above by 1. This is similar in behavior to the NBP price but also to consumption flows. This series is then fed to the model as the seasonal average cycle of the price.

One problem is that the data may include exceptional circumstances, and so we are not aggregating the "normal" cycle. But if we look only at Italian gas consumption in 2002-2004, which did not have extractions from the strategic stock, this is indistinguishable from the 2002-2006 period, which did. The "exceptional" events of 2005 and 2006 have no expression in final consumption, so using the data should not be a problem.

The catastrophe shock is the most wide open of the three processes. The sale of strategic stock which occurred in 2006 in Italy only required a small inversion in the relative price of electricity between Italian prices and European exchange prices to take effect. The agents selling gas took a light view of the 3.5 euro penalty per GJ "illegaly" sold. The opposite case was an accident at the Rough facility in the UK on March 16, 2006 which closed it for the rest of the year. NBP gas prices spiked immediately, but even so this episode - which constitutes a legitimate shock - was not too severe due to an average winter and to the existence of alternative supplies from the continent.[15]

If we compute ratios of winter to summer spot prices, they vary significantly across markets where such prices are available, but a value above 2 is unusual. However a ratio around 2 is not what worries EU policy makers. They worry more about politically motivated disruptions or severe damage to pipelines which are harder (if not impossible) to predict. Accordingly they define security of supply as a minimum amount in storage able to service demand for a given period of time.[16] The European Council Directive 2004/67/EC of 26 April 2004 – "concerning measures to safeguard security of natural gas supply" – is in fact not totally clear about what security of supply means. It is worried about a "major supply disruption" which it defines as:

> "a situation where the Community would risk losing more than 20 % of its gas supply from third countries and the situation at Community level is not likely to be adequately managed with national measures".

The expression "would risk" makes it an appropriately vague definition. The absence of an exact quantification - losing the supply for how long? - further increases its vagueness. Without a clear indicator of what a large disruption is, we set **as a benchmark** an event that doubles the winter maximum price from 1 to 2 (which implies the winter price would be 4 times the summer price).

Finally, we must define the likelihood of such an event. Since the higher its probability, the less it can be defined as unusual - and frequent disruptions are likely to be well internalized by commercial operators - we set the probability of this disruption to a benchmark of once every five years (59/60) at monthly frequencies. The shock is iid.

[15] Winter 2007 sees a low NBP price due to the introduction of new pipelines from Norway.
[16] Curiously, prices are absent from this definition.

5 Simulations of Policy Impact

We are now ready to examine the behavior of the model with and without the $k(S)$ constraint. We look at both the Italian and Danish cases. The model is simulated at monthly frequency. The artificial sample is of 1200 periods, which is 100 years. In every experiment (every row) the same exact realization of shocks is used.

Table 4: Policy Experiments

	λ	q	$\bar{p}$	V_0	V_1	$(V_1 - V_0)/V_0$	n
IT	0.378	59/60	2	3.545	2.966	-0.163	26
IT	0.378	95/96	2	3.475	2.879	-0.171	17
IT	0.378	59/60	1	3.349	–	–	26
IT	0.378	59/60	3	3.805	3.237	-0.149	26
IT	0.189	59/60	2	3.544	3.355	-0.053	26
DK	$\lambda(n)$	59/60	2	4.963	3.970	-0.200	26
DK	$\lambda(n)$	95/96	2	4.873	3.861	-0.208	17
DK	$\lambda(n)$	59/60	1	4.725	–	–	26
DK	$\lambda(n)$	59/60	3	5.267	4.330	-0.178	26
DK	$\lambda(n)/2$	59/60	2	4.964	4.583	-0.077	26

The first column describes the policy being evaluated. The second column of Table 4 shows the iid probability of a catastrophe shock each month, (q). We look at two values, 59/60, which is once every five years, and 95/96, or once every eight years. The third column shows the value of the catastrophe price. Columns 4 and 5 show the value of the firm, first with $k(S) = 0$, then with an active penalty/policy. Column 6 shows the cost of this policy in terms of the value of the firm. Column 7 shows the number of catastrophes actually occurring in the artificial sample.

Column five deserves special attention because V_1 is the "factual" value when policy is active. This value is of course conditional on our scenario assumptions for prices and probabilities of a catastrophe, which, for example, in the first row are that $\bar{p} = 2$ and $q = 59/60$.

We computed also the pure option value for the firm. This is calculated by imposing the same time series of shocks on the model's optimal decision, but shutting down *the realization* of the catastrophe shocks. In all cases the resulting value of the firm is about 0.25% different from what we show here. This means that the impact of the catastrophe shock on the value of the firm is essentially an option value. It is the fact that high prices may occur in the distant future that creates value.

Figure 6 shows the fit of the model (constrained and unconstrained) for Italy, shown in the first row of Table 4. The fit of the unconstrained model is relatively worse as one would hope.

The first five rows of Table 4 examine Italy. The first row takes as benchmark a shock frequency of one month every five years. This results in

Figure 6: Model Fit, Italy

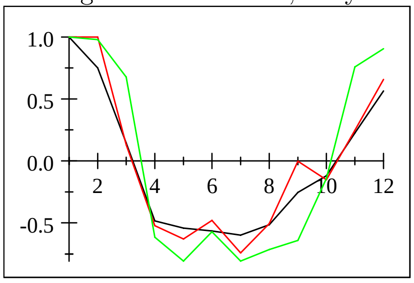

a value loss from being constrained of 16% relative to the unconstrained case. This is less than the value of λ and shows the initial impression from the back of the envelope example used in the introduction is both wrong (because the value loss is about half of λ) and right (because it is still a very large portion of the value of the firm).

Changing the frequency of the shock from once every eight years (row two) to once every five years (row one) has a negligible effect on the value of the firm, again stressing that what really matters is the possibility of using the strategic stock during normal times. In row 3 we compute the unconstrained value for a world without catastrophes. This generates a value difference with respect to the unconstrained value in row one of 5.5%. In row 4 we raise the catastrophe price to 3 times the normal price and the value of the firm increases by about 7%. The value difference from the constraint is roughly the same. Row 5, on the other hand, shows the impact of reducing the strategic stock requirement by a half. This reduces the cost of regulation more than proportionately.

Rows six to ten examine Denmark. Here the value loss from the cyclical constraint is around 20%, whereas in Italy it was around 16%. All other qualitative results are replicated here. Figure 7 shows the model fit for Denmark, in the two simulations corresponding to row 6 in Table 4. Here the fit is better when we actually shut down policy. We experimented with the quarterly data calibration and the fit improves significantly for the "factual" constrained benchmark.

Figure 7: Model Fit, Denmark

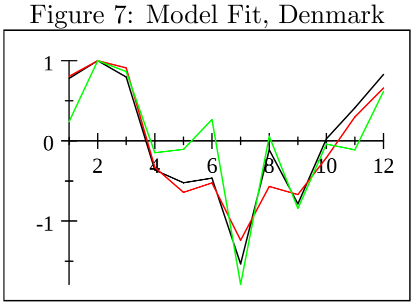

Our final experiment, shown in table 5, is to transform Italy into Denmark. We impose on Italy a cyclical constraint with the pattern of their 60 day-ahead consumption, but normalized so that it is on average 0.378 of the stock.

Table 5: Italy with a cyclical constraint

	λ	q	$\bar{p}$	V_0	V_1	$(V_1 - V_0)/V_0$	n
IT	0.378	59/60	2	3.545	2.966	-0.163	25
IT	$\lambda(n)$	59/60	2	3.546	2.813	-0.207	25

This alternative actually *reduces* the value of the firm because it increases the strategic stock above 0.378 during the winter when prices are high, and reduces the strategic stock during summer when prices are low. But this is actually worse for the firm than the constant level of the constraint. The constant constraint provides less insurance but gives the firm higher profits.

6 Conclusion

In this paper we construct a model of storage of natural gas that replicates the behavior of stocks and flows in the Italian market. The model is then used to evaluate the impact of strategic storage policy on the value of the firm. Our experiments suggest the impact of strategic storage policy on the value of the firm is very large. Our calibration for Italian and Danish data yields a cost of regulation of around 16% and 20% of discounted net present value of profits for the respective storage markets. Furthermore, imposing a danish-type cyclical constraint on the Italian market that preserves the average size of the strategic stock provides higher insurance but at a higher cost per unit insured, as more gas is withheld in the winter when it is more valuable, and in turn more gas is released in the summer when it is less valuable.

The percentage loss is not as sensitive to the catastrophe price as it is to the size of the constraint (37.8% of stocks for Italy) or to the normal variation of the price from summer to winter. The reason is that the biggest part of the loss comes from the inability to exercise the seasonal arbitrage on a significant fraction of the gas in stock.

Finally, the exogeneity of the price and the assumption of a fixed capacity clearly condition the quantitative findings in this paper. Modelling the demand for gas and including it in the model is not an impossible task and is a useful extension to the present exercise. As for endogeneizing capacity, given that investments in gas storage are long projects with very large fixed costs, it is acceptable to assume a constant capacity in the medium run. Endogenizing capacity is a more daunting task than modelling demand.

References

[1] Byers, J., (2006). ”Commodity Storage Valuation: A linear optimization based on traded instruments”, Energy Economics 28, 275–287.

[2] Casassus, J., Collin-Dufresne, P., and Routledge, B. (2005), "Equilibrium Commodity Prices with irreversible Investment and Non-Linear Technologies", NBER Working Paper No. W11864

[3] Chloé Le Coq and Elena Paltseva, (2008) "Common Energy Policy in the EU: The Moral Hazard of the Security of External Supply, SIEPS Working Paper 1.

[4] Chaton, C., Creti, A. and Villeneuve, B. (2007a), "Some economics of seasonal gas storage." INSEE Working Paper, 52.

[5] Chaton, C., Creti, A. and Villeneuve, B. (2007b), "Storage and security of supply in the medium run." IEFE Working paper, 1.

[6] Hall, G. and Rust, J. (2000), “An empirical model of inventory investment by durable commodity intermediaries.” Carnegie-Rochester Conference Series on Public Policy

[7] Natural Gas Information 2005, International Energy Agency, IEA Statistics, OECD.

[8] Natural Gas Information 2007, International Energy Agency, IEA Statistics, OECD.

[9] Thomson, M., Davison, M., and Rasmussen, H. (2007); “Natural Gas Storage Valuation and Optimization: A Real Options Approach”, manuscript, University of Western Ontario.

[10] Wright, B., and Williams, J. (1982a), "The roles of public and private storage in managing oil import disruptions." The Bell Journal of Economics, 341-353.

[11] Wright, B., and Williams, J. (1982b), "The economic role of commodity storage" The Economic Journal, 596-614.