Dealing with European gas storage contracts: Trading strategies unraveled

Christian Huurman and Andras Szell[1]

Abstract:

Within the deregulated energy industry, gas is the fastest growing energy commodity amongst its 'peers' and takes a unique place in the storage trading landscape: In contrast, electricity is virtually non-storable, oil is kept underground and its storage strategy can be best described as 'laissez-faire'. Gas storage, on the other hand, is often actively managed to secure supply, smoothing it over time to the demand curve, or simply to exploit market arbitrage opportunities. That is, the price risk associated with holding storage inventory is hedged on today's spot and forward markets, or speculated upon by 'virtual storage' owners. In this study we consider some strategies that are typically used by gas traders, which allow for the value extraction from the gas storage facility. It is challenging to choose a valuation methodology that honours both financial (e.g. bid-ask spread, price behavior of all tradable contracts) features, and physical (ratchets and inventory level) features. In an attempt to do this, various authors have proposed a Least-Squares Monte Carlo approach to incorporate both. This often implied that a simplistic single-factor approach was adopted, as it assumes that any shock in the spot price is propagated along the forward price curve. In this paper we will model both gas spot and forward contracts following a realistic multi-factor approach. We apply this multi-factor trading strategy to the most liquid European gas market, the National Balancing Point (NBP). We then compare it with the storage value metrics that would have been obtained when the aforementioned spot-only trading strategy would have been adopted. As both stylized facts of gas prices (volatility, seasonality, mean reversion) and type of storage asset (depleted field, salt cavern) are known to drive the storage value, we consider various price and storage scenarios. Appealing results are for example the empirical evidence that a wider choice of available spot and forward contracts results in the opportunity to both lock in more intrinsic value across the considered storage types. Firstly, it is shown that an increase in the availability of tradable short-term contracts (day-ahead, weekend, balance-of-month contracts) results in a significant increase in both the intrinsic and extrinsic storage value component, hence also the overall value. This result held across all considered storage cycles (10-300 days of injection and withdrawal). Secondly, the extrinsic value-to-capacity multiple was typically double the intrinsic value-to-capacity component for fast storage cycles, which among others implies that one can pocket a considerable amount of additional profit when the bid-ask spread moves over time. In addition, it makes the claim made by Boogert and De Jong (2008) flawed, who say that the rolling intrinsic strategy is not very valuable. When the true tradable contracts in a multifactor-framework are considered this strategy *is* very valuable. Finally, we discuss the implications of our results for gas portfolio management.

Keywords: gas storage, spot and forward prices, seasonal behavior, mean-reversion, multi-factor models, trading strategies, Monte Carlo simulations, risk premiums, energy portfolio management

[1] Both authors are affiliated with FEA, a brand of MSCI Ltd. Any view of the authors expressed in this paper is their personal view only, and does not represent their employer's view unless explicitly stated. Corresponding author: Tel: +44(0)2076182085. All errors pertain to the authors.

1. Introduction

The 1990s witnessed the start of a deregulation process in the European power and gas industry, which gradually resulted in the unbundling from sales and network services traditionally owned by state-owned entities, and a landscape where prices are now based on the market rules of supply and demand. As electricity is virtually non-storable, except for hydro-based power, and oil storage is typically done by *not* pumping it out of the ground, we'll focus on gas in order to shed light on the challenge to value an energy storage asset in a realistic way. Roughly three different types of underground gas storage facilities can be considered and we ranked them in terms of how quickly an in-and-out cycle can be completed, also referred to as deliverability: (1) salt caverns, which can be constructed in salt dome formations (2) natural aquifers, which can used for gas storage in certain sedimentary rock formations, and (3) depleted natural gas fields, which share similar geological characteristics as aquifers. Given the size and nature of the natural gas market, its strong seasonal character, and the large volatility of its spot price, a substantial market is now being developed for renting and managing the aforementioned storage capacity types. The theory of storage proposed in the 1930s and applied to commodities to explain the differences between spot and forward prices by analyzing the rationale for holding inventories (back then, the term convenience yield was introduced), has led to a bulk of literature on the relationship between commodity (spot, forward) prices and holding inventory. Therefore it is surprising that only a few studies have looked at the implications of the theory of storage to gas in its current market environment, where the Samuelson effect[2] and price seasonality is particularly pronounced in gas prices, and executing 'timing' injections and withdrawals within the physical constraints to realize contract forward arbitrage, is key to obtain storage value. The aforementioned challenge to incorporate the spot-forward price dynamics into a trading strategy that optimizes the storage value around the myriad of costs and constraints lies in the balance to propose a model that is tractable from a daily decision support management perspective on one hand, and honours both financial features and physical constraints on the other hand. We kindly refer to the eminent work of Géman (2005), who gives an excellent overview of the three types of methodologies to solve the value optimisation problem: (1) select the hedges from the forward price curve to lock in the maximum 'current' storage value, known as the 'intrinsic' strategy, (2) applying real option theory by considering the facility as a basket of calendar spread options and delta hedging the underlying forward contracts, (3) stochastic optimisation through a (forest of) tree(s) or (Least-Squares) Monte Carlo Simulations. The latter method has been adopted by several authors who study gas storage valuations. For instance, Boogert and De Jong (2008) apply the LSMC method on the storage contract, which can be seen as a basket of co-dependent American style options. Both they and Neumann and Zachmann (2009) propose a single-factor mean-reverting spot price model when valuing a storage contract on respectively the Dutch and German markets. It is interesting to note that Boogert and De Jong (2008) suggest that future developments to their framework should be adopting a multi-factor price model. In this study we attempt to do this by following the multi-factor approach proposed by Gray and Khandelwal (2004) to model the spot-forward price dynamics of typically traded spot and forward contracts, and adopt the 'rolling' extension of the intrinsic strategy (ad1), hereafter the rolling intrinsic (RI) to extract the value for various storage type scenarios. Please note that as one rolls from one day to the next, the expected value must be estimated via Monte Carlo simulation. The RI strategy can be easily

[2] Samuelson (1965) shows that when spot prices are mean-reverting, the return volatility of forward contracts rises when the contract rolls to its expiration date.

extended to incorporate various 'spot' and forward contract maturities, which we will show in this study. From our empirical study we disclose that the extended RI strategy can actually capture a considerable amount of extrinsic value when the tradable contracts available at European gas markets are considered in a multifactor-framework, which places Boogert and De Jong's claim that the RI strategy is not very valuable into perspective. Thirdly, the RI strategy is employed by numerous storage operators in Europe and the US[3], and the reported empirical evidence has several practical consequences for gas portfolio management.

The remainder of the paper is structured as follows. In Section 2 we discuss the methodology. Section 3 provides the data and deliverability scenarios. In Section 4 we present and discuss the results.

2. Valuation methodology

The value of a storage asset originates from trading on physical spot and forward contracts. In order to estimate this value, one has to consider a realistic physical model of the storage, and model for spot and forward gas prices. The final value is determined by the "trading strategy".

As mentioned in the previous section, a realistic storage model contains the physical volume limits, and limits on injection and withdrawal speeds as a function of time.

Traditionally, gas spot prices are modeled by lognormal mean reversion processes, like the one factor Schwartz (1997) model, or the two-factor Schwartz and Smith (2000) model. Here, we consider the following two-factor mean reversion model:

$$d\log S_t = a\left[\theta_t - log\left(\frac{S_t}{F_{t,T}}\right)\right]dt + {\sigma_t}^S\, d{W_t}^S,$$

$$dF_{t,T_p} = F_{t,T_p}{\sigma_t}^1 d{W_t}^1,$$

where F_{t,T_p} denotes the "prompt" month forward price at time t. ${W_t}^S$ and ${W_t}^1$ are correlated Brownian motions (later referred as prompt-spot correlation), ${\sigma_t}^S$ and ${\sigma_t}^1$ are piecewise constant volatility values, θ_t is determined by the risk neutral measure and a denotes the constant mean reversion rate.

We model the forward curve by the market standard lognormal multi-factor model:

$$dF_{t,T} = F_{t,T}{\sigma_t}^T d{W_t}^T,$$

where T denotes the time of maturities ($T = 1,2, \dots N$), t the time of observation, ${\sigma_t}^T$ the volatility of the forward contract $F_{t,T}$, and ${W_t}^T$ the standard Brownian motion with calibrated correlation matrix between ${W_t}^1$, ${W_t}^2$, etc.

[3] FEA is one of the leading sources for pricing models and risk management analytics designed specifically for the natural gas industry.

In this paper we focus on the Rolling Intrinsic (RI) trading strategy, which involves the continuous application of the simple intrinsic strategy[4]. On every day of the storage operation, aside from fulfilling prior commitments i.e. taking/making delivery of the agreed amounts – one also calculates a new intrinsic value and its necessary intrinsic hedge based on the updated market prices. If the newly calculated value, plus any potential gain/loss from marking-to-market all outstanding hedge positions, is greater than that which would result from taking to terms the latter, an action is taken in favour of the new hedge. One notable feature of this strategy (due to its 'intrinsic' nature) is that at the end of each day, the storage value is always greater than or equal to that from the day before and it is fully-hedged.

A considerable improvement of the RI model is when one considers spot trading as well above the forward deals. In this "extended" RI strategy one can gain extra profit from the more volatile nature of the spot prices. Here, the actual intrinsic forward portfolio is updated and completed with spot trades if the additional injection or withdrawal on spot increases the value of the portfolio.

With the previously defined price models it is possible to simulate the RI strategy in a realistic way: the two-factor spot process enables the consistent simulation of the connection between spot and forward price movements, while the multi-factor model provides practical simulation of the colourful dynamics of the forward curve.

Various authors have reported computational and performance issues when they applied their valuation algorithm to a storage contract. In general, the performance is a function of three things: the number of simulations, the length of the storage contract and the storage contract constraints. The runtime is linear in the number of simulations. The runtime is less than quadratic in the length of the contract (but not linear) in the length of the storage contract. But the main constraints that consume much of the runtime are the injection/withdrawal ratchets, as they can pose a nonlinear optimisation problem to the model if level-dependent. The optimisation algorithms we use have been optimized in aforementioned areas.[5]

3. Storage contract, scenarios and data

In this section we'll introduce the six deliverability scenarios (subsection 3.2.1), two forward price volatility scenarios (3.2.2) that we'll apply to the following standard storage contract, which is identical to the one studied by Boogert and de Jong (2008) in terms of contract length, minimum, maximum, start and end volume[6]:

[4] Again, we kindly refer to the seminal paper of Gray and Khandelwal (2004), who examine these trading strategies.

[5] In particular, we use multi-threading to which decrease runtime in proportion to the number of processors that are available. In addition, we have tried to maintain as many possible optimisation algorithms and as many possible options for tuning the optimisation algorithms (e.g. in the model parameters), and often, these can be fine-tuned on a case-by-case basis for increasing the performance to some degree.

[6] Please note that we evaluate an NBP storage contract while they examine a TTF storage contract. Hence, our results are expressed in pence or pence/therm rather than Euro or Euro/MWh.

Nomination	Daily
Time to maturity	T = 1 Year
Start date	1 July 2011
End date	30 June 2012
Value date	30 June 2011
Min. volume	0
Max volume	250.000 therms
Start volume	100.000 therms
End volume	100.000 therms
Max withdrawal & max injection	Please see section 3.2.1

Table 1: Characteristics of our standard storage contract

The closing subsection 3.2.3 provides a market description and the dataset.

3.2.1. Deliverability scenarios

Deliverability tends to vary greatly by storage field type, with salt caverns being the fast-cycle types (typically between 10-40 days) versus aquifers and depleted fields on the other side of the spectrum (typically between 180 – 300 days). Please note that the actual injection and withdrawal profiles could be a function of the inventory level, which is then commonly referred to as a ratchets schedule. In this study we have chosen not to include such a level-dependent schedule (as this poses a nonlinear optimisation problem to the model), but have recognized the following six types of storage injection and withdrawal cycles: 10 days, 30 days, 60 days, 180 days, 240 days and 300 days. When you divide the storage capacity by these numbers of days you'll get the range of injection and withdrawal rates (per day) as listed in table 2.

Six deliverability scenarios		
Injection/ withdrawal (days)	Injection rate (per day)	Withdrawal rate (per day)
10	25,000	25,000
30	8,333	8,333
60	4,167	4,167
180	1,389	1,389
240	1,042	1,042
300	833	833

Table 2

3.2.2. Price volatility scenarios

In this section we investigate various volatility and trading scenarios for RI storage pricing. We show that additional spot trading on the storage increased the extrinsic value considerably during the observation period. In order to do this we calculate RI values using forward trades only and then using forward and spot trades altogether.

High and low forward volatility scenarios were created, where high means about 60% implied volatilities and low means about 20% implied volatilities of the forwards. The evolution of the forward curve depends heavily on the correlation structure: with correlation parameters equal to one, the shape of the forward curve is fixed, hence no extrinsic value can come from the RI forward trading. In contrast, zero correlations would results in plenty of RI opportunities. Here, we fixed the overall correlation values at 97%, which is a conservative overestimation of the overall correlation structure.

The spot trade means day-ahead (DA) deal now. Calibrating to the historical data set, we use 11 days for expecting time to mean revert, 45% model volatility and 70% prompt-spot correlation value.

Rolling Intrinsic, (M) forwards only, LOW forward volatilities					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	303,578	6657 (1.57%)		232,500	71,078
30	303,332	6551 (1.78%)		232,500	70,832
60	276,000	5203 (1.54%)		232,500	43,500
180	213,246	3558 (1.42%)		197,929	15,317
240	160,115	2462 (1.31%)		148,447	11,668
300	128,162	2033 (1.35%)		118,786	9,336
Rolling Intrinsic, (M) forwards and spot (DA), LOW forward volatilities					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	2,298,276	6657 (1.57%)		323,125	1,975,151
30	1,482,574	6551 (1.78%)		262,707	1,219,867
60	971,584	5203 (1.54%)		247,602	723,982
180	385,217	3558 (1.42%)		200,809	184,408
240	287,761	2462 (1.31%)		150,607	137,154
300	230,217	2033 (1.35%)		120,514	109,703
Rolling Intrinsic, (M) forwards only, HIGH forward volatilities					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	1,164,775	6657 (1.57%)		232,500	932,275
30	1,163,923	6551 (1.78%)		232,500	931,423
60	834,097	5203 (1.54%)		232,500	601,597
180	400,940	3558 (1.42%)		197,929	203,011
240	300,437	2462 (1.31%)		148,447	151,990
300	240,451	2033 (1.35%)		118,786	121,665
Rolling Intrinsic, (M) forwards and spot (DA), HIGH forward volatilities					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	3,664,981	6657 (1.57%)		323,125	3,341,856
30	2,580,088	6551 (1.78%)		262,707	2,317,381
60	1,790,676	5203 (1.54%)		247,602	1,543,074
180	686,837	3558 (1.42%)		200,809	486,028
240	512,645	2462 (1.31%)		150,607	362,038
300	410,487	2033 (1.35%)		120,514	289,973

Table 3

From Table 3 it can be clearly seen that additional spot trading provided considerably more extra profit compared to the forward only RI valuation. In the low volatility scenario, the forward only strategy hardly gave any additional values, which is consistent with the evidence reported by

Boogert and De Jong (2008). However, even in this low volatility case the additional spot deals almost doubled the intrinsic value even for a very slow storage scenario. In the high volatility scenario, the forward only strategy performs well already. The additional spot trade even doubled the profit.

3.2.3. NBP market

As in the United States, most European hubs developed at pipeline intersections. Good examples are the National Balancing Point in the UK and the Hub Holland in The Netherlands (latter market is used in eth empirical study of Boogert and De Jong (2008). Whereas on the US Nymex market, only day-ahead (DA), Balance of Month (BOM) and monthly (M) futures are being traded, the European markets provide provide more trade choice variety: The most liquid European reference prices NBP and TTF, which stem from respectively the UK and Dutch hubs provide also quotes for short-term maturities like weekend (WE) and next-week weekdays (WDNW). Depending on time of trade, BOM contracts may change the length of their delivery period. The other short-term contracts have a fixed delivery period length, DA – 1 day, WE – 2 days and WDNW – 5 days. Please note that not all contracts trade on all weekdays (only WD trade on a weekend). In addition, we base our contract structure on the current rules followed at the Intercontinental Exchange ('ICE')[7], which offers a mixed set of forward contracts for physical delivery at the NBP. There is a variety of 10 to 12 consecutive monthly contracts depending on the observation month. In addition, long-term contracts are also available on NBP, such as quarterly and seasonal contracts, but they are excluded from our empirical analysis.

3.2.3.1 Empirical data

From Graph 1 you can observe that each of the NBP contract maturities have their own dynamics.

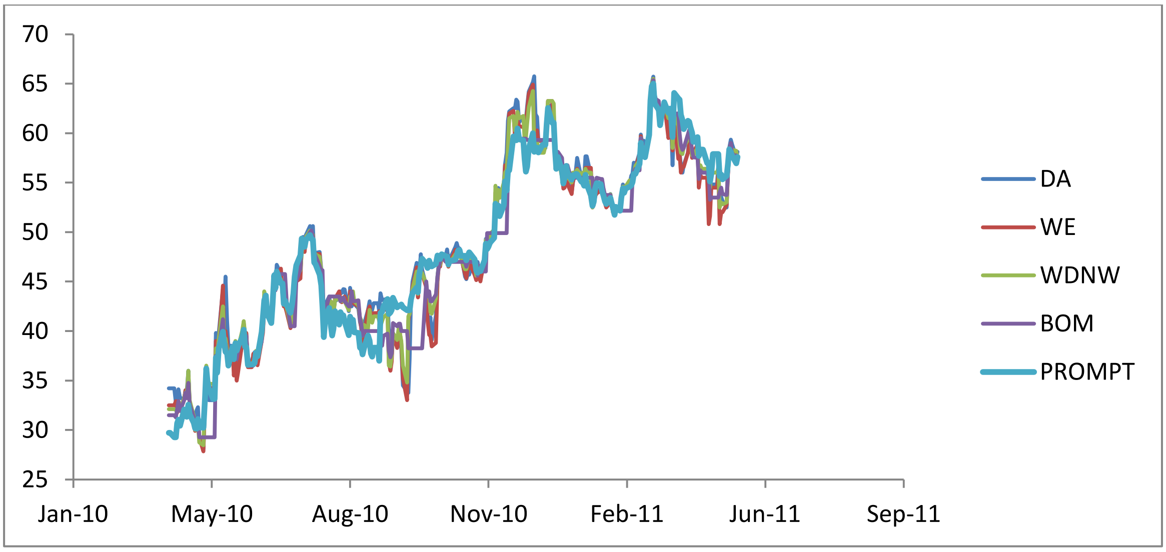

Graph 1: NBP contracts (pence /therm) from 1 April 2010 to 17 May 2011

[7] Please see the www.theice.com for a detailed description of the market and contract conventions.

Summary statistics of the daily NBP contracts are listed in Table 4.

NBP: 1 April 2010 to 17 May 2011					
	DA	WE	WDNW	BOM	Prompt
No of observations	294	294	294	294	294
Mean	48.63	48.12	48.43	47.83	48.30
Standard deviation	9.08	9.04	9.03	9.14	9.15

Table 4

3.2.3.2 RI trade choice scenarios

As mentioned, there is a wide choice of tradable contracts available at the NBP. Obviously, the range of choice affects the opportunity to take away more profit from the spread dynamics. Hence, we have selected three trade choice scenarios: A traditional RI scenario where you only have access to M contracts (hereafter 'M only'). Please note that this scenario is consistent with Boogert and Jong's rolling-intrinsic scenario definition. In addition, we introduce two "extended" RI trade scenarios: One where you can trade both M contracts and DA contracts, and another where you can optimize on arbitrage opportunities in all aforementioned contracts that is, DA-, WE-, WDNW-, BOM- and M contracts. Please note that in the "extended" RI trade scenarios it is assumed that each of the "extended" contracts (DA, WE, WDNW, BOM) mean-revert to the 'rolling-prompt' contract.

4. Empirical results

In section 4.1.1 we present the empirical storage valuation metrics for our storage contract and we discuss the findings in light of both our theoretical storage valuation metrics (section 3.2.2.) and Boogert and De Jong's findings. Section 4.2.2 presents the empirical storage premiums for each trade choice and storage type asset. All model parameters have been calibrated to the market.

4.1.1 Storage values

We ran the rolling intrinsic simulation for all six deliverability scenarios (so given the 3 RI trade choice scenarios, 18 runs in total), each time with 100 simulations. The results have been reported in Table 5. We recognized that this number of paths may be considered low. Please note that we've also reported the MC simulation error, which has been calculated as the standard deviation of the simulated distribution, divided by the square-root of the number of simulations. From these metrics one can see that convergence is achieved[8], and the added benefit is that the runtime for our storage contract ranged between just 3 to 4 minutes.

It is a well known fact that a fast storage is worth more than a slow one, as can be seen in the tables. Another trivial observation is that adding more and more trading opportunities increased both the

[8] We've done a lot of testing regarding the convergence properties of the simulation and have found that in most cases, a one-year storage deal converges to within 5% of the MC error with 100 simulations.

intrinsic and extrinsic value of the storage. Adding only one spot trading (DA) hardly increased the intrinsic value; however it increased the extrinsic value considerably. Including all available spot deals, the intrinsic value increased significantly, since one can gain profit from the arbitrage opportunities between them. The extrinsic per intrinsic ratio was not as great in this case, although one could achieve the largest extrinsic profits here.

Another interesting finding was the magnitude of extra profit one can gain from the “extended” RI strategy. While the traditional, forward only method provides very limited extra values, consideration of spot like trades gave very significant extrinsic amounts. We stress that this profit was achieved with a very safe, simple and all time fully hedged trading strategy. Please note that the realistic estimation of the expected profit requires the application of sophisticated price processes whereas simple Monte-Carlo calculations are required for the pricing.

So how did our metrics relate to the empirical evidence provided in similar studies? To our best knowledge only Boogert and De Jong (2008) have shown empirical evidence on value distributions of gas storage contracts, whereby a variety of model parameters (e.g. low-high volatility scenarios) and storage type scenarios are considered. They’ve reported an extrinsic-intrinsic multiple ranging between 1.5 and 2, depending on the volatility case under consideration. We can observe from table 5 that the size of this multiple is easily reached for the fast storage types when the “extended” RI trade scenarios are considered (in the (DA, M) case even between 3 to 5) .

Rolling Intrinsic (M only) - 100 Monte Carlo simulations					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	370,095	6657 (1.57%)		232,500	137,595
30	368,976	6551 (1.78%)		232,500	136,476
60	338,214	5203 (1.54%)		232,500	105,714
180	249,882	3558 (1.42%)		197,929	51,953
240	187,944	2462 (1.31%)		148,447	39,498
300	150,203	2033 (1.35%)		118,786	31,418
“Extended” Rolling Intrinsic (DA, M) - 100 Monte Carlo simulations					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	2,521,530	57,831 (2.29%)		323,125	2,198,405
30	1,576,368	39,757 (2.52%)		262,707	1,313,661
60	1,043,039	25,749 (2.47%)		247,602	795,438
180	416,841	11,149 (2.67%		200,809	216,032
240	309,498	8,683 (2.81%)		150,607	158,892
300	244,505	7,036 (2.88%)		120,514	123,991
“Extended” Rolling Intrinsic (DA, WE, WDNW, BOM, M) - 100 Monte Carlo simulations					
Injection/ withdrawal(days)	Total value (pence)	MC error (as % of total value)		Intrinsic component	Extrinsic component
10	10,899,979	179,145(1.64%)		3,743,342	7,156,637
30	6,897,441	76,110 (1.10%)		3,028,526	3,868,915
60	5,140,092	43,755 (0.85%)		2,837,990	2,302,102
180	4,509,474	66,606 (1.48%)		2,521,395	1,988,079
240	3,297,809	18,031 (0.55%)		2,521,395	776,414
300	3,297,809	18,031 (0.55%)		2,521,395	776,414

Table 5

4.1.2 Storage premiums

Another way of interpreting the metrics presented in Table 5 is to analyse these in terms of risk-return perspective. In the RI strategy, at any given point, and certainly before the storage asset is leased, the intrinsic value is the only value that is certain. After all, the extrinsic component embedded in the total expected value is risky and cannot be hedged. This implies that any bid for storage that is over the intrinsic is risky in the sense that value might not be recovered. The market is more aware that the "faster" the storage the more likely it is that one is able to capture extrinsic value from spot spikes and dips, and some of that risk inherent in making back your bid (i.e. making more money than you paid) is mitigated. In Table 6 we have reported the empirical metrics in terms of these storage 'bid' premiums, which is calculated as the Intrinsic (Intr) or Extrinsic value (extr), divided by the storage capacity (cap) of the contract, 250.000 therms.

RI strategy	(M)		(DA, M)		(DA, WE, BOM, M)	
Injection/ withdrawal(days)	Intr/ cap (p/therm)	Extr/ cap (p/therm)	Intr/ cap (p/therm)	Extr/ cap (p/therm)	Intr/ cap (p/therm)	Extr/ cap (p/therm)
10	0.93	0.55	1.29	8.79	14.97	28.63
30	0.93	0.55	1.05	5.25	12.11	15.48
60	0.93	0.42	0.99	3.18	11.35	9.21
180	0.79	0.21	0.80	0.86	10.09	7.95
240	0.59	0.16	0.60	0.64	10.09	3.11
300	0.48	0.13	0.48	0.50	10.09	3.11

Table 6

The slow storage facilities ("300", "240", "180") provide very little opportunity to exploit arbitrage opportunities from the spot market, because they're too slow to react to spot price (DA, WE, WDNW) dynamics. Therefore, one might decide to bid just the intrinsic, 10.09p in our case. For a "60"-day service, there's more optionality so might make sense to pay an additional premium on top of the intrinsic. For "30"-day service there's even more optionality, so it is logical to make the bid more competitive. In this framework, a "10"-day service could almost be considered as a "cash point" as it allows to 'cycle' through the storage many times a year, having the price risk hedged through the DA (and/or WE and WDNW and BOM) trade choices.

5. Concluding remarks

In this study we have described how realistic trade strategies can be applied to the valuation of a storage contract, hereby distinguishing the variety in storage type(s) and gas price dynamics. We have provided theoretical (from volatility scenarios) and empirical evidence (from NBP) that the "storage value" is subjective to aforementioned stylised facts and, last but not least, the amount of "risk" a firm is willing to take. We claim that our findings show that the RI strategy is a key valuation approach for traders, storage operators and risk managers who deal with storage contracts.

6. Literature

Boogert, A., and C. de Jong (2008), Gas Storage Valuation Using a Monte Carlo Method, *Journal of Derivatives*, Vol 15 (3), pp 81–98.

Géman, H., 2005, *Commodities and Commodity Derivatives*, John Wiley & sons, Chichester.

Gray, J. and P. Khandelwal (2004), Towards a Realistic Gas Storage Model, *Commodities Now*, June 2004, pp 1-4.

Gray, J. and P. Khandelwal (2004), Towards a Realistic Gas Storage Model, *Commodities Now*, September 2004, pp 1-5.

Neumann, A. and G. Zachmann, 2009, “Expected vs. Observed Natural Gas Storage Usage: Limits to Intertemporal Arbitrage” in *The Economics of Natural Gas Storage*, Springer-Verlag, Heidelberg. Pp 13-31.

Samuelson, P, “Proof that Properly Anticipated Prices Fluctuate Randomly”, *Industrial Management Review*, 6, pp. 41-50.

Schwartz, E., 1997, “The stochastic behavior of commodity prices: Implications for valuation and hedging”, *Journal of Finance*, 52(3), pp 923-973.

Schwartz, E. and J. Smith, 2000, “Short-term Variations and Long-term Dynamics in Commodity Prices”, Management *Science*, Vol 46 (7), pp 893-911.