Gas Storage Valuation in Incomplete Markets

Nils Löhndorf

Luxembourg Centre for Logistics and Supply Chain Management, University of Luxembourg, 1511 Luxembourg, Luxembourg

David Wozabal

TUM School of Management, Technical University of Munich, 80333 Munich, Germany

Abstract

Natural gas storage valuation is an important problem in energy trading, yet most valuation approaches are based on heuristics or ignore that gas markets are incomplete. We propose an exact valuation model for incomplete gas markets based on multistage stochastic programming. Market incompleteness structurally changes the problem of storage valuation and asset backed trading and the resulting model requires analysis of a combined control problem of storage operation and futures trading that takes risk preferences into account. As the problem is subject to the curse of dimensionality, we reduce the stochastic process to a scenario lattice and solve the resulting problem using stochastic dual dynamic programming. We show that the intrinsic value of storage corresponds to the value under perfect risk aversion and that the rolling intrinsic value, which is popular among practitioners and has been found to be near-optimal when markets are complete, is an inconsistent price rule in incomplete markets. Our results inform managerial decisions on risk management and asset pricing for natural gas highlighting the importance of explicitly modeling risk preferences.

Keywords: multistage stochastic programming, asset pricing, natural gas industry, Markov processes

1. Introduction

Over the past decade natural gas has increasingly replaced other fossil fuels in the United States. The U.S. Energy Information Administration projects that this trend is likely to continue and that natural gas consumption will make up 35% of total energy consumption by 2040 (EIA 2016). This will not only increase trading activities in natural gas markets but also lead to a growing importance of storage facilities which are used to buffer variability in demand. Gas storage facilities are scarce, as they require natural geological formations like depleted gas fields, salt domes, or aquifers. Due to its scarcity, storage capacity is an important asset in energy

trading.

The value of storage derives from the possibility of buying and injecting gas at times of low prices and withdrawing and selling when prices are high. To do so, owners of storage trade on commodity exchanges, such as the New York Mercantile Exchange (NYMEX) in the United States, which hosts spot and futures markets for natural gas. Often, the ownership of gas and storages is separated, and contracts are used to transfer usage rights from storage owners to third parties. Such contracts are sold by owners of physical storage to energy traders that either require storage capacity for the purpose of risk management or use it for arbitrage. Pricing storage contracts is therefore a problem of practical importance in energy trading, which is corroborated by the existence of numerous commercial software products that offer storage valuation tools (e.g., MSCI 2015, KYOS 2015, Lacima 2017). Existing approaches for gas storage valuation are either based on heuristics or assume that gas markets are complete.

The two most common heuristics popular among practitioners are *basket-of-spread options valuation* and *rolling intrinsic valuation*. These approaches define the value of storage as the sum of its intrinsic and its extrinsic value. While the *intrinsic* value of a gas storage can be easily obtained by finding the optimal schedule of injections and withdrawals based on current futures prices, its *extrinsic* value captures the additional value of flexibility to adapt this schedule to future price changes (Gray and Khandelwal 2004a,b). Both approaches make simplifying assumptions to approximate the *extrinsic* value of storage.

The rolling intrinsic approach simulates the possible evolution of the futures curve over time and approximates the extrinsic value as the expected gains from changes in the schedule of injections and withdrawals that result from repeated re-optimization as new information arrives (see Eydeland and Wolyniec 2003, Gray and Khandelwal 2004a,b). Lai et al. (2010) derive a dual upper bound for the value of gas storage in complete markets and find that the rolling intrinsic value is near-optimal. Secomandi (2015) provides additional theoretical support for this finding.

The basket-of-spread options approach, by contrast, approximates the extrinsic value by the value of a portfolio of (time) spread options. Although the results in Lai et al. (2010) indicate that this approach is suboptimal, the computed storage value is *hedgeable* whereas the rolling intrinsic value is not (Gray and Khandelwal 2004b).

Another family of models that is popular in the literature are spot-only models (e.g. Boogert and de Jong 2008, Chen and Forsyth 2007, Thompson et al. 2009, Carmona and Ludkovski 2010, Secomandi 2010, Bjerksund et al. 2011, Boogert and de Jong 2012, Felix and Weber 2012, Malyscheff and Trafalis 2017). Models from this family treat future contracts as derivatives that can be replicated by spot trading and thereby do not have to be modeled as separate assets. While spot-only models are able to exploit seasonal spreads and keep the decision problem

low-dimensional, traders are often suspicious of these approaches, as they ignore the dynamics of the futures curve and are usually inconsistent with observed futures prices (Lai et al. 2010).

In particular, spot-only models require the assumptions that gas markets are complete in order to view futures as derivatives. This in particular implies the span of all traded instruments encompasses all future states of the world. While this assumption is plausible in liquid markets for a product that can be stored at no cost, such as foreign currency, it is hardly tenable for gas markets which are plagued by low liquidity and where storage and transport of the underlying is subject to various physical and institutional limitations.

In an incomplete market, the value an individual agent ascribes to storage may vary, as the value of storage equals the price at which the agent is indifferent between holding and selling the storage (Klöppel and Schweizer 2007). The storage value may thus depend on individual expectations of future price dynamics, specific institutional and organizational constraints, as well as individual risk aversion regarding possible future cash flows. Carmona and Ludkovski (2010) point out that computing this value requires analysis of a combined control problem of storage operation and portfolio optimization that takes the agent's risk preferences into account.

The first contribution of this work is the formulation of a joint stochastic programming model for futures trading and storage operation, which not only captures the multi-dimensional nature of the aforementioned control problem but also allows modeling individual risk preferences. The multistage stochastic programming model yields an optimal asset-backed trading strategy and is summarized in Section 2.2. In particular, to circumvent unbounded speculation on futures curve movements, the model requires that futures contracts are always hedged back-to-back by physical storage operation. In Section 2.3, we show that the intrinsic value is a special case of our model under perfect risk aversion and argue why rolling intrinsic valuation is inconsistent in incomplete markets.

As the valuation problem is subject to the curse of dimensionality, the second contribution is the development of an efficient approximation strategy that is capable of solving the problem to near-optimality. In a first step, we approximate the continuous Markov process of futures prices as a scenario lattice, which provides an efficient discretization that scales well in the number of decision stages. In the second step, we solve the pricing problem using a Markovian version of *stochastic dual dynamic programming* (SDDP) (Mo et al. 2001, Löhndorf et al. 2013, Löhndorf and Shapiro 2019) and adopt the change-of-measure approach for coherent risk measures to model risk preferences (Shapiro et al. 2013, Philpott et al. 2013).

Since the quality of our approximation strategy depends on how close the scenario lattice approximates the price process, the lattice generation method is central to our approach. Discretizing a continuous Markov process to a lattice has been first proposed in Bally and Pagès (2003) and applied to SDDP independently in Bonnans et al. (2012) and Löhndorf et al. (2013).

Bally and Pagès (2003) propose a stochastic gradient algorithm to construct scenario lattices approximating Brownian motions based on the Wasserstein distance between lattice nodes and unconditional distribution of the discrete-time process. In Section 3, we extend the work of Bally and Pagès (2003) by proposing an estimation procedure that ensures that the first moments of all conditional distributions of the lattice are correct, which is important to ensure that the computed value of storage is unbiased.

The overall approximation strategy has two advantages: first, a scenario lattice can be constructed by a computationally inexpensive learning algorithm. Second, SDDP eliminates the need to discretize the high-dimensional resource state by training approximate value functions that are accurate in those regions of the state space that can be reached by the optimal policy. We refer to this two-step approach as *approximate dual dynamic programming.*

To study the quality of our approximation, we conduct extensive numerical experiments, which are summarized in Section 4. In Section 4.2, we present a numerical comparison of our solution with rolling intrinsic valuation as well as the dual upper bound of Lai et al. (2010) in a complete market setting. The comparison shows that our approach produces near-optimal results despite the high dimensionality of the underlying dynamic program. In Section 4.3, we study the storage value in an incomplete market. We demonstrate numerically how the value of storage converges to the intrinsic value as risk aversion increases, and that the rolling intrinsic value corresponds to expected profits under a perfectly risk-averse strategy. Furthermore, we show that rolling intrinsic valuation significantly underestimates the value of gas storage in an incomplete market if the decision-maker is risk-neutral or only mildly risk-averse.

We believe that our results provide energy traders with a new approach for storage valuation and asset-backed trading in incomplete markets. The model allows traders to choose a statistical model of futures price dynamics that matches the empirical realities of markets for natural gas, and can be easily tailored to take risk preferences as well as physical and institutional limitations into account.

2. Model Formulation

In this section, we show how storage valuation in incomplete markets can be formulated as a multistage stochastic program and demonstrate that the intrinsic value is a special case of the model.

2.1. Assumptions

We consider the problem of a price-taking energy trader who manages a gas storage contract, which grants the buyer the right to inject, withdraw, and store gas over a finite time horizon.

The trader operates in an incomplete gas market where the traded instruments are insufficient to replicate the payoff of the storage contract.

The injection, withdrawal, and storage capacity is limited and defined by the contract. We assume that injection and withdrawal limits are independent of storage level, since level dependent limits are usually not part of the provisions in storage contracts. Following Lai et al. (2010), we assume that each physical injection and/or withdrawal incurs a marginal cost in addition to an in-kind fuel loss while storing natural gas incurs no holding cost.

Denote F_t, $t \in [T]$, as the random futures prices defined on the probability space $(\Omega, \mathcal{F}, \mathbb{P})$, with $T \in \mathbb{N}$ as the number of discrete time periods and $[T] = \{1, \ldots, T\}$. In an incomplete market, the value an individual agent ascribes to storage equals the price at which the agent is indifferent between holding and selling the storage, a concept that is also referred to as *indifference pricing* (Klöppel and Schweizer 2007, Carmona 2009). More specifically, the indifference value corresponds to the utility of cash flows $C_1(F_1, x_1), \ldots, C_T(F_T, x_T)$ that can be generated by making optimal operational and trading decisions $x_1, \ldots, x_T$ given uncertain futures prices, $F_1, \ldots, F_T$. As $C_t(F_t, x_t)$ are random variables, their utility depends on individual risk preferences, which we measure by using a dynamic utility measure that is equivalent to a dynamic risk measure with signs reversed (Klöppel and Schweizer 2007, Ruszczyński 2010). The indifference value can thus be expressed as

$$\max_x \mathcal{U}\left(C_1(F_1, x_1), \ldots, C_T(F_T, x_T)\right) \tag{1}$$

where $\mathcal{U}$ models the risk preferences.

In the following, we assume that risk preferences of the agent in (1) can be expressed by the nested conditional value-at-risk (CVaR), e.g. see Ruszczyński and Shapiro (2006). CVaR has generated a lot of attention in the recent literature and is widely used in many applications including investment and trading in energy, e.g. see Conejo et al. (2016).

For random profits $X_1, \ldots, X_T$, let $\text{CVaR}_{t,\alpha}(X_t) = \alpha^{-1} \int_0^\alpha G_{X_t|X_{t-1}}^{-1}(t) dt$ where $G_{X_t|X_{t-1}}$ is the distribution function of X_t conditional on X_{t-1} and define the random variables

$$\mathbb{U}_{t,\alpha,\lambda}(X_t) = \lambda \text{CVaR}_{t,\alpha}(X_t) + (1-\lambda)\mathbb{E}[X_t|X_{t-1}]$$

for $0 \leq \lambda \leq 1$ and $0 < \alpha \leq 1$. For a sequence of random variables, $X_1, \ldots, X_T$, the nested CVaR is defined as

$$X_1 + \mathbb{U}_{2,\alpha,\lambda}\left(X_2 + \mathbb{U}_{3,\alpha,\lambda}\left(X_3 + \cdots\right)\right), \tag{2}$$

i.e., by a convex combination of expectation and CVaR that recursively includes other convex combinations of expectation and CVaR. For $\lambda < 1$, the nested CVaR is a time-consistent

extension of conventional CVaR to a dynamic setting (Shapiro 2009).

Note that we are maximizing the CVaR, i.e., viewing it as an acceptability functional based on the $\alpha\%$ of the worst (lowest) outcomes of a profit distribution. Correspondingly, typical α's would be close to 0 and not close to 1 as is the case when CVaR is used as a risk measure.

Of course, nested CVaR is only one of many possible ways to model an agent's risk preferences and utility functions that yield other indifference values are conceivable. Nevertheless, unlike many other risk measures, a practical advantage of the nested CVaR is that it keeps optimization problems computationally tractable. Furthermore, by varying parameters α and λ the nested CVaR can be fit to a wide range of risk preferences.

We assume that the agent pursues an asset-backed trading strategy to value gas storage. Many companies trading in energy commodities have been adopting an *asset-backed trading models* (Boudier et al. 2015), which require that futures positions are fully backed by physical assets. In the case of gas storage this implies that futures bought for delivery in t_1 are always mirrored back-to-back with futures sold for delivery in $t_2 > t_1$ in such a way that the storage can be used to store the gas bought at t_1 and deliver it back to the market at t_2. In this way, the projected storage level at the end of the contract is always zero, which restricts the policy to arbitrage seasonal price differences using storage rather than speculate on price shifts of single future contracts over time. On a technical level, limiting speculative trading that is not connected to storage operation, guarantees that the optimization problem is bounded and that the value of storage is not distorted.

In line with the literature, we assume that rewards are generated in monthly increments (Eydeland and Wolyniec 2003, Lai et al. 2010) and that the dynamics of the futures curve follows a multivariate geometric Brownian motion (MGBM), e.g., Clewlow and Strickland (2000), Eydeland and Wolyniec (2003), Lai et al. (2010). At time $t \in [T]$, let $F_t = (F_{tt}, \ldots, F_{tT})$ be the vector of random prices, with F_{tt} being the spot price and F_{tj} being the price of the future contract with maturity in period $j = t + 1, \ldots, T$. Then, the price process in continuous time is given by the following system of stochastic differential equations,

$$\frac{dF_{tj}}{F_{tj}} = \mu_j dt + \sigma_j dZ_{tj}, \ dZ_{tj} dZ_{tk} = \rho_{jk}, \ 1 \leq j, k \leq T, \tag{3}$$

where dZ_{tj} are correlated increments of standard Brownian motions.

For our problem, we are interested in a discrete-time version of (3) in monthly time increments. To this end, we assume that the contract starts at the beginning of a month at time $t = 0$, and that $t = 1$ marks the beginning of the next month and so forth. In the following, we will be concerned with the finite set of log-normal distributed random vectors $F_1, \ldots, F_T$ for $t = 0, 1, \ldots, T$ resulting from (3). Note that, in contrast to Lai et al. (2010), we allow for

non-zero drifts $\mu_j \neq 0$, resulting in processes that are not martingales. The reason is that, as opposed to risk-free pricing, for indifference pricing, physical measures that govern the evolution of the *true* price processes are required.

2.2. Multistage Stochastic Programming Formulation

We model the decision problem of a storage contract owner as a multistage stochastic programming problem,

$$\max_{x_1 \in \mathcal{X}_1(x_0)} C_1(F_1, x_1) + \mathbb{U}_{2,\alpha,\lambda} \left[\max_{x_2 \in \mathcal{X}_2(x_1)} C_2(F_2, x_2) + \cdots + \mathbb{U}_{T,\alpha,\lambda} \left[\max_{x_T \in \mathcal{X}_T(x_{T-1})} C_T(F_T, x_T) \right] \cdots \right],$$

(4)

where $C_t(\cdot, \cdot)$ is the immediate profit, $\mathcal{X}_t(\cdot)$ is the feasible set in $t = 1, \ldots, T$, and (F_1, x_0) is the deterministic initial state. The probability functionals $\mathbb{U}_{t,\alpha,\lambda}$ in (4) are conditioned on the collection of random vectors $(F_1, F_2, \ldots, F_{t-1})$. Furthermore, we assume relatively complete recourse, i.e., that the nested optimization problems are feasible for every x_{t-1}. Note that we formulate the problem in such a way that $\mathcal{X}_t$ depends on previous decisions $x_1, \ldots, x_{t-1}$ only by dependence on x_{t-1} and is independent of F_t. Note that the recursive structure of the above problem guarantees that decisions are non-anticipative, i.e., that x_t only depends on information available up until time t.

Next we define the decision variables that are contained in x_t as well as the feasible sets $\mathcal{X}_t$. Denote by f_{tj} the futures contract position that matures in period j. Trading decisions in period t of futures maturing in month $j > t$ are denoted by $y_{tj} \in \mathbb{R}$, where positive values indicate sales and negative values indicates that the contract j is bought in period t. Accordingly, y_{tt} denotes spot trading decisions.

To enforce an asset-backed trading strategy, we assume that futures positions are always balanced back-to-back with storage operation. This requires us to model planned future storage operation at every period t for every future point $j > t$. To that end, we denote by s_{tj} the storage level and by w_{tj}, i_{tj} injections and withdrawals as planned at period t for future periods $j \geq t$. When a futures contract matures, it must be either fulfilled physically through storage operation or cleared in the spot market, i.e., at time t

$$f_{tt} - y_{tt} = i_{tt} - w_{tt}, \ t \in [T],$$

(5)

where f_{tt} is the futures position that goes into delivery in period t. These decisions lead to a physical storage balance at time $t + 1$ which is given by

$$s_{t+1,t+1} = s_{tt} + d^i i_{tt} - d^w w_{tt}, \ t \in [T],$$

(6)

where d^i and d^w are in-kind losses of storage.

Similarly, we model the planned future storage balance in period t until maturity by the constraints

$$s_{t,j+1} = s_{tj} + d^i i_{tj} - d^w w_{tj}, \ t \in [T-1], \ j = t, \dots, T. \tag{7}$$

and make sure that all tradable contracts are balanced with planned future injections and withdrawals at all times, i.e,

$$f_{tj} = i_{tj} - w_{tj}, \ t \in [T-1], \ j = t+1, \dots, T. \tag{8}$$

and all contracts are subject to balance constraints

$$f_{t+1,j} = f_{tj} - y_{tj}, \ t \in [T-1], \ j = t+1, \dots, T, \tag{9}$$

which track changes in futures contract positions resulting from purchases or sales in t. The constraints (7), (8), and (9) ensure that at time t positions f_{tj} with $j > t$ permit a physically implementable operational schedule, which ensures that open future positions are backed by the storage asset and are not purely speculative.

Furthermore, we require that all trades are made back-to-back, i.e., that the sum of positive and negative positions cancel out at every point in time. This is enforced by requiring that

$$s_{t,T} = 0, \ t \in [T-1] \tag{10}$$

ensuring that there are no open positions at maturity of the storage contract.

Storage operational limits are given by the following set of constraints,

$$i_{tj} \leq \bar{i}, \ t \in [T], \ j = t, \dots, T, \tag{11}$$

$$w_{tj} \leq \bar{w}, \ t \in [T], \ j = t, \dots, T, \tag{12}$$

$$s_{tj} \leq \bar{s}, \ t \in [T], \ j = t, \dots, T, \tag{13}$$

where $\bar{i}$, $\bar{w}$ are the upper limits for injections and withdrawals per period and $\bar{s}$ is the capacity of the storage.

The immediate reward is given by

$$C_t(F_t, x_t) = \sum_{j=t}^{T} \gamma_j F_{tj} y_{tj} - \gamma_t (c^i i_{tt} + c^w w_{tt}), \tag{14}$$

where γ_t is the discount factor for period t and c^i and c^w are marginal physical injection and

withdrawal costs. Note that, since we are using deterministic discount rates, the cash-flows generated from futures and forward contracts are equivalent. To keep the model parsimonious, we therefore avoid the explicit modeling of margining and treat gas futures as forwards.

The problem in $t+1$ only depends on a resource state R_t defined by storage levels as well as futures contract positions. We define this resource state as the set of time-coupling decision variables in $t-1$ that enter the right-hand side of constraints in t,

$$R_t = \{s_{t-1,t-1}\} \times \{f_{t-1,j}\}_{j=t}^T, \ t = 2, \ldots, T. \tag{15}$$

Given that F is a Markov process, the dynamic programming equations for problem (4) can be written as

$$V_t(F_t, R_t) = \max_{x_t \in \mathcal{X}_t(R_t)} \{C_t(F_t, x_t) + \mathcal{V}_{t+1}(F_t, R_{t+1})\}, \ t = 1, \ldots, T, \tag{16}$$

where

$$\mathcal{V}_{t+1}(F_t, R_{t+1}) = \mathbb{U}_{t+1,\alpha,\lambda}\left[V_{t+1}(F_{t+1}, R_{t+1})\right], \ t = 1, \ldots, T, \tag{17}$$

with $\mathcal{V}_{T+1} \equiv 0$. We refer to functions $V_t(\cdot, \cdot)$ as value functions and $\mathcal{V}_t(\cdot, \cdot)$ as expected value functions of the problem.

Remark 1. *The value functions, $V_t(F_t, R_t)$, are concave in R_t if the optimal values of*

$$\max \{C_t(F_t, x_t) + \mathcal{V}_{t+1}(F_t, R_{t+1} | x_t \in \mathcal{X}_t(F_{t-1}, R_t)\} \tag{18}$$

are concave in R_t, which can be shown by induction in t going backwards in time. Since all constraints as well as the objective function are linear, the maximization problem at stage T is linear. The resource state R_T appears only on the right-hand side of the constraints, such that the objective value and therefore the value function V_T is a concave function in R_T. As $\mathbb{U}_{t,\alpha,\lambda}$ is monotone and concave it preserves concavity and the expected value function $\mathcal{V}_T$ is also concave in R_T. By backward induction, it follows that all expected value functions $\mathcal{V}_t$ are concave functions in R_t.

Since $R \mapsto \mathcal{V}_t(F_t, R)$ is concave, it can be approximated from above by the minimum of L supporting hyperplanes, $\hat{\mathcal{V}}_t$, which is defined by scalar intercepts $a_l(F_t, \hat{R}_{t+1}^l)$ and slope vectors $b_l(F_t, \hat{R}_{t+1}^l)$, whereby $\hat{R}_{t+1}^l$ denotes a feasible final resource state, i.e.,

$$\hat{\mathcal{V}}_t(F_t, R_{t+1}) = \min_{l \in [L]} \left\{ a_l(F_t, \hat{R}_{t+1}^l) + b_l(F_t, \hat{R}_{t+1}^l)^\top (R_{t+1} - \hat{R}_{t+1}^l), \ l \in [L] \right\}, \ t \in [T-1]. \tag{19}$$

Combining (14) with (19), (16) can be approximated by a linear program, i.e.,

$$
V_t(F_t, R_t) \approx
\begin{cases}
\max & \sum_{j=t}^{T} \gamma_j F_{tj} y_{tj} - \gamma_t(c^i i_{tt} + c^w w_{tt}) + v_t \\
\text{s.t.} & v_t \le a_l(F_t, \hat{R}_{t+1}^l) + b_l(F_t, \hat{R}_{t+1}^l)^\top (R_{t+1} - \hat{R}_{t+1}^l), \ l \in [L], \ t \in [T-1], \\
& (5), (6), (7), (8), (9), (10), (11), (12), (13), \\
& i_{tj}, w_{tj} \ge 0, y_{tj} \in \mathbb{R}, \ t \in [T], \ j = t, \dots, T, \\
& s_{t+1,j} \ge 0, f_{t+1,j} \in \mathbb{R}, \ t \in [T-1], \ j = t, \dots, T-1, \\
& v_t \in \mathbb{R}, \ t \in [T-1].
\end{cases}
\tag{20}
$$

Note that the risk preferences modeled by the nested CVaR with parameters α and λ are accounted for in $\hat{\mathcal{V}}_t(F_t, R_{t+1})$ by a change of measure (Philpott et al. 2013). A brief description of the approach is provided in Section 3.3.

2.3. The Intrinsic Value as Special Case

An approach often chosen by practitioners is (rolling) intrinsic valuation. The intrinsic value of storage $I(F_1)$ depends F_1 and can be easily obtained by only solving the first stage of the decision problem. Given observable futures prices $F_{11}, \dots, F_{1T}$ in period $t = 1$, we choose a schedule of injections and withdrawals that respects all operational constraints and maximizes the immediate profits generated by futures trading in $t = 1$

$$
I(F_1) =
\begin{cases}
\max & \sum_{j=1}^{T} \gamma_j \left((F_{1j} - c^w) w_{1j} - (F_{1j} + c^i) i_{1j} \right) \\
\text{s.t.} & s_{1,j+1} = s_{1j} + d^i i_{1j} - d^w w_{1j}, \quad \forall j \in [T] \\
& i_{1j} \in [0, \bar{i}], w_{1j} \in [0, \bar{w}], s_{1j} \in [0, \bar{s}], \quad \forall j \in [T]
\end{cases}
\tag{21}
$$

where $s_{1j} \equiv 0$.

Of course, as time passes and prices change the problem can be reoptimized and additional profit can be made. To estimate the value of this additional (extrinsic) optionality, we can simulate futures curve movements over a rolling horizon, for example by sampling from the discrete-time version of (3).

Since each reoptimization brings additional value, the rolling intrinsic value is always higher than the intrinsic value and, as shown in Lai et al. (2010), is a close approximation for the optimal (extrinsic) value in a complete market setting.

Depending on parameters α and λ, the optimal policy nests the optimal expectation maximizing policy as well as the rolling intrinsic solution as special cases. The case of the risk-neutral planner is trivial: by setting $\lambda = 0$, $\mathbb{U}_{t,\alpha,\lambda}(X_t) = \mathbb{E}(X_t | X_{t-1})$ reduces to the expected value and we obtain a policy that maximizes expected discounted rewards.

To relate our policy to the rolling intrinsic solution, we consider the case of $\lambda = 1$, i.e., the

decision-maker that maximizes only CVaR. By additionally letting α go to 0, we obtain

$$\text{CVaR}_\alpha(X) \xrightarrow{\alpha \to 0} \text{ess inf}(X) =: \text{CVaR}_0(X).$$

Under the nested CVaR objective with $\alpha = 0$ and $\lambda = 1$, the optimization problem becomes

$$\max_{x_1} C_1(F_1, x_1) + \mathbb{U}_{2,0,1}\Big[\cdots + \mathbb{U}_{T,0,1}\big[\max_{x_T} C_T(F_T, x_T)\big]\Big] = \max_{x \in \mathcal{X}} \sum_{t=1}^{T} \text{ess inf } C_t(F_t, x_t), \quad (22)$$

where (22) is the extensive form of the problem described in the last section with $\mathcal{X}$ the feasible set for the policy described by (5)-(13).

The next result establishes that the continuous problem (22) based on the evolution of the log-normally distributed F_t defined in (3) yields the rolling intrinsic policy. Hence, the rolling intrinsic is the perfectly risk-averse limit of the gas storage planning problem described in the previous section.

The two parts of the proposition address *the limiting case* of $\lambda = 1$ and $\alpha = 0$ and what happens as the risk preferences *approach* the limiting case as $\alpha \to 0$.

Proposition 1. 1. *The optimal policy for $\lambda = 1$ and $\alpha = 0$ is a rolling intrinsic policy. In particular, the optimal first-stage objective value equals the intrinsic value of storage.*

2. *The optimal policy for $\lambda = 1$ converges to a rolling intrinsic policy as $\alpha \to 0$ and the optimal value converges to the intrinsic value of storage.*

Proof. In $t = 1$, an optimal rolling intrinsic decision x_1^* maximizes the first-stage reward subject to storage constraints, irrespective of possible gains that might arise because of random price changes in the future. This implies that decisions are taken without considering any value function.

Any policy with lower first-stage reward than x_1^* must earn a positive amount through price changes in later periods almost surely to compensate for the difference if it is to be optimal for (22). Therefore, it is enough to show that for every decision $\bar{x}_1$, there is a positive probability that the value of the associated portfolio of futures contracts decreases in subsequent periods. In this case, the evaluation of the decision $\bar{x}_1$ would necessarily be worse than that of x_1^* according to (22), which entirely focuses on the worst outcomes.

In particular, for a futures position $\bar{f}_{2,t}$ resulting from the trades associated with $\bar{x}_1$, the set of prices in $t = 2$, that do not cause the value of the futures positions, $\bar{f}_{2,t}$, to increase, contains

$$\mathcal{H} = \big\{F_2 = (F_{2,2}, \ldots, F_{2,T}) : \bar{f}_{2,t}(F_{2,t} - F_{1,t}) \leq 0, \ \forall \, t = 2, \ldots, T\big\}. \quad (23)$$

$\mathcal{H}$ is an intersection of half-spaces in $\mathbb{R}^{T-1}$ that intersect the positive orthant. As prices under

(3) are log-normally distributed, $\mathbb{P}(\mathcal{H}) > 0$ showing that the probability that the portfolio value decreases in period 2 is positive. A repetition of the argument yields that there is always a positive probability that the losses for portfolio $\bar{f}_{2,t}$ persist until the end of the planning period independent of the decision taken. Hence, any $\bar{x}_1$ that leads to a lower profit than x_1^* in the first stage is suboptimal for (22). This establishes that any rolling intrinsic is an optimal policy for a perfectly risk averse agent concluding the proof of 1.

To prove 2, we start by noting that the problem in the last period

$$V_T(F_T, R_t) = \max_{x_T \in \mathcal{X}_T(R_T)} C_T(F_T, x_T) \tag{24}$$

does not depend on α. Since CVaR_α converges to CVaR_0 as $\alpha \to 0$, it follows that

$$\lim_{\alpha \to 0} \mathbb{U}_{T,\alpha,1}(V_T(F_T, R_T)) = \mathbb{U}_{T,0,1}(V_T(F_T, R_T)) \tag{25}$$

and hence $\mathcal{V}_{\alpha T}(F_{T-1}, R_T)$ converges, where $\mathcal{V}_{\alpha T}$ denotes the expected value function using $\lambda = 1$ and α for the risk preferences. Continuing, we write for arbitrary F_{T-1} and R_{T-1}

$$V_{\alpha,T-1}(F_{T-1}, R_{T-1}) = \max_{x_{T-1} \in \mathcal{X}_{T-1}(R_{T-1})} \{C_{T-1}(F_{T-1}, x_{T-1}) + \mathcal{V}_{\alpha T}(F_{T-1}, R_T)\} \tag{26}$$

and note that the objective function converges pointwise as $\alpha \to 0$ by the argument above and therefore epiconverges by Shapiro et al. (2009), Theorem 7.27. Since the decisions x_{T-1} are bounded, it follows that the decisions x_{T-1} as well as $V_{\alpha,T-1}$ converge (see Shapiro et al. 2009, Proposition 7.26). Repeating the argument recursively backwards in time, it follows that the optimal decisions as well as the optimal objective converge to the perfectly risk-averse case, which yields a rolling intrinsic policy and the intrinsic value as first-stage objective due to 1. $\quad\square$

This result can be intuitively explained as follows: as we assume a log-normal distribution for future prices, there is positive probability that speculation leads to losses which is avoided by following a perfectly risk-averse policy. The rolling intrinsic policy never speculates on future price movements. Hence, perfect risk aversion leads to the rolling intrinsic policy.

Note that the above proof uses the assumption of normally distributed returns only to argue that is always a positive chance of prices moving in either direction. This clearly is the case for any reasonable continuous model of future prices. Thus, the above result can be generalized substantially beyond the models considered in this paper.

Remark 2. *The indifference value of storage is the acceptability of the cash flows resulting from the optimal operation of the gas storage, because for this value the decision-maker is indifferent between owning the storage and owning the objective value in cash. In indifference pricing, the*

risk preferences in the objective therefore determine a fitting optimal operational policy.

Proposition 1 shows that the rolling intrinsic policy corresponds to the decisions that would be taken by a maximally risk averse agent. For these decisions to be optimal in indifference pricing, the agent would have to use the risk preferences (22) in the objective to calculate the value of storage. However, in the rolling intrinsic valuation the profits are valued using the expectation. This obviously leads to an inconsistency between the rolling intrinsic policy and the risk preferences employed to value the cash-flows. In particular, a risk neutral agent would not take the decisions prescribed by the rolling intrinsic policy in the first place.

Moreover, Proposition 1 implies that a perfectly risk-averse decision-maker values a gas storage at its intrinsic value. This inconsistency implies that the rolling intrinsic value does not correspond to an indifference price storage for any set of risk preferences and thus cannot be used for indifference pricing.

3. Method

In this section, we describe a computational strategy to approximate the expected value function (19). We introduce scenario lattices as efficient mean to discretize stochastic processes and show how lattices can be constructed in an optimal way using stochastic gradient descent. We then use Markovian SDDP to solve the problem.

3.1. Scenario Lattices

The multistage stochastic programming problem defined in the previous section is based on the continuous stochastic process $F_1, \ldots, F_T$ and is therefore an infinite dimensional linear programming problem which cannot be solved in closed form.

A popular approach to overcome this problem is to discretize the price process and reduce it to a scenario tree (Pflug 2001, Dupačová et al. 2003, Pflug and Pichler 2012). To mitigate exponential growth of the number of nodes, scenario trees typically require very low branching factors or contain nodes where the tree does not branch at all, which effectively leads to deterministic problems at these nodes.

The complexity of a scenario tree can be reduced substantially if the data process is Markovian, because the conditional distribution of a random variable F_{t+1} in $t+1$ no longer depends on the entire price history $(F_1, \ldots, F_t)$ but only on F_t. This implies all scenarios that lead to F_t can be represented by the same node at time t and thus can be merged without loss of information. We refer to such a *recombining* scenario tree as *scenario lattice*, in line with the terminology often used in mathematical finance (e.g., binomial lattices for pricing American options).
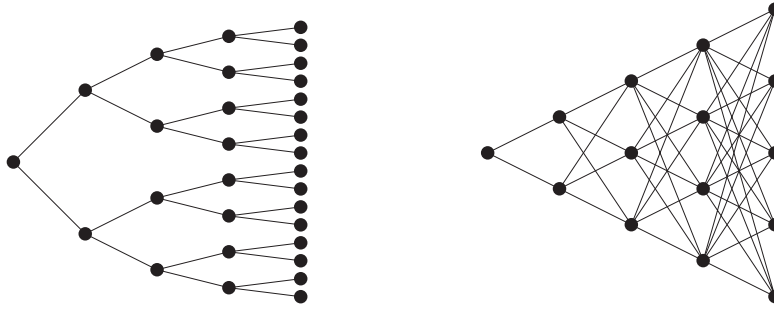
Figure 1: A tree with 31 nodes representing 16 scenarios on the left and lattice with 15 nodes representing 120 scenarios on the right.

Formally, a lattice is a graph organized in a finite number of layers. Each layer is associated with a discrete point in time and contains a finite number of nodes. Successive layers are connected by arcs. A node represents a possible state of the stochastic process, and an arc indicates the possibility of a state transition between the two connected nodes. Each arc is associated with a probability weight, and the weights of outgoing arcs of a node add up to one. A scenario tree differs from a scenario lattice by the additional requirement that every node in t has only one predecessor in $(t-1)$.

Denote N_t as the number of nodes of the lattice in t, and $\bar{F}_{tn}$, $n \in [N_t]$, as the state of the lattice process at node n in t and $\bar{F}_t = \left\{ \bar{F}_{tn} : n \in [N_t] \right\}$ as the set of all possible states at time t. Assuming that all state transitions between nodes on consecutive stages have positive probabilities, the set of possible scenarios on the lattice is given by

$$\bar{F}_1 \times \bar{F}_2 \times \cdots \times \bar{F}_{T-1} \times \bar{F}_T. \tag{27}$$

As the number of stages grows, the additional nodes needed to construct a lattice are those of the newly added stages. In contrast, the number of nodes in a scenario tree where each node has more than one successor grows exponentially in the number of stages. See Figure 1 for an illustrative comparison.

While methods for scenario tree reduction are well studied in the extant literature, methods to construct scenario lattices for general stochastic optimization problems have received almost no attention. A notable exception is Bally and Pagès (2003) who refer to a lattice as *vector quantization tree*. The authors propose a first-order learning algorithm aimed at minimizing the distance between the unconditional distribution of a Markov chain and the nodes of the lattice. However, Bally and Pagès (2003) estimate transition probabilities using maximum likelihood, which can lead to biased trading decisions, as we will demonstrate in Section 4. Our algorithm overcomes this problem by ensuring that the conditional expectations of the lattice coincide with those of the price process.

3.2. Optimal Lattice Discretization

Our objective is to construct a scenario lattice such that the optimal value for the problem on the lattice is close to the optimal value of the problem using the true process. From stability theory of stochastic programs, it is known that minimizing this error is closely linked to the *Wasserstein* distance between the discrete and the continuous process (Dupačová et al. 2003, Bally and Pagès 2003, Pflug and Pichler 2012). While there exists a rich literature on scenario tree generation using the Wasserstein distance (see Kaut and Wallace 2007, for a survey), we are not aware of comparable literature on scenario lattices.

We start by discussing the stage-wise discretization of the price process F_t in t. We first recall the Wasserstein distance for distributions. Loosely speaking, the Wasserstein distance is defined as minimum cost of transporting mass from one distribution to the other, whereby cost are measured by the distance between the respective outcomes (see Villani 2003, for a comprehensive introduction).

Definition 1. *Let $\xi : (\Omega, \mathcal{A}) \to \mathbb{R}^n$ and $\tilde{\xi} : (\tilde{\Omega}, \tilde{\mathcal{A}}) \to \mathbb{R}^n$ be two random variables with distributions P and $\tilde{P}$, respectively. The Wasserstein distance of order $(r \geq 1)$ between ξ and $\tilde{\xi}$ is defined as*

$$W_r(\xi, \tilde{\xi}) = \begin{cases} \inf_{\pi} & \left(\int\limits_{\Omega \times \tilde{\Omega}} \|\xi(\omega) - \tilde{\xi}(\tilde{\omega})\|_r^r \, \pi\,(d\omega, d\tilde{\omega}) \right)^{\frac{1}{r}} \\ s.t. & \pi(A \times \tilde{\Omega}) = P(A) \quad \forall A \in \mathcal{A}, \\ & \pi\,(\Omega \times B) = \tilde{P}(B) \quad \forall B \in \tilde{\mathcal{A}}, \end{cases} \tag{28}$$

where the infimum is taken over all probability measures π on $(\Omega \times \tilde{\Omega}, \mathcal{A} \otimes \tilde{\mathcal{A}})$.

For $r = 1$, the following dual representation holds

$$W_1(\xi, \tilde{\xi}) = \sup_{\|f\|_{Lip} \leq 1} \left\{ \int f \, dP - \int f \, d\tilde{P} \right\},$$

where $\|f\|_{Lip}$ is the Lipschitz constant of f, i.e., the supremum is over all Lipschitz functions with a Lipschitz constant smaller than 1. Based on this result, W_1 can be directly used to bound the error in two-stage stochastic optimization problems in a straight-forward manner. In particular, for a well behaved two-stage stochastic optimization problem,

$$\Pi_P^* = \begin{cases} \max_{x_1} & C_1(x_1) + \mathbb{E}_P(V(x, \xi)) \\ s.t. & x_1 \in \mathcal{X}_1, \end{cases} \qquad V(x, \xi) = \begin{cases} \max_{x_2} & C_2(x_2, \xi) \\ s.t. & x_2 \in \mathcal{X}_2(x_1) \end{cases} \tag{29}$$

with randomness ξ with distribution P, the recourse function Q is Lipschitz (see Kall and Mayer

2006, Theorem 2.3 for the case of linear problems), i.e., there is a $L > 0$ such that

$$|V(x, \xi_1) - V(x, \xi_2)| \leq L \, \|\xi_1 - \xi_2\|_1, \quad \forall x \in \mathcal{X}. \tag{30}$$

If the distribution P is either continuous or contains a large number of atoms, one might want to approximate problem (29) by replacing ξ by a *simpler* version $\bar{\xi}$ with distribution $\bar{P}$ to make it computationally tractable. The error induced by this step can be bounded as follows using the Wasserstein distance

$$|\Pi_P^* - \Pi_{\bar{P}}^*| \leq |\mathbb{E}_P(V(x_{\bar{P}}^*, \xi)) - \mathbb{E}_{\bar{P}}(V(x_{\bar{P}}^*, \bar{\xi}))| \leq L \, W_1(\xi, \bar{\xi}), \tag{31}$$

with $x_{\bar{P}}^*$ the first-stage optimal solution for (29) with the expectation of the second stage revenue taken with respect to $\bar{P}$. It is therefore desirable to select a discrete distribution $\bar{P}$ that is closest to P under the Wasserstein distance to keep the approximation error small.

For our continuous Markov process, this means that we want to find nodes of the lattice, $\bar{F}_t = (\bar{F}_{t1}, \ldots, \bar{F}_{tN_t})$, and corresponding unconditional probabilities such that the Wasserstein distance of the resulting distribution is as close as possible to the real continuous distribution of F_t. In this paper, we focus on the case $r = 2$. We define as d_t the dimension of the process F_t in period t, i.e., $\bar{F}_{tn} \in \mathbb{R}^{d_t}$ for all $n \in [N_t]$.

It follows from Graf and Luschgy (2000), Lemma 3.1 that the problem of finding an optimal discretization of F_t with a fixed number of nodes N_t in terms of the Wasserstein distance is equivalent to solving the N_t-center problem, i.e.,

$$\inf_{\bar{F}_t} D(\bar{F}_t) := \inf_{\bar{F}_t} 2^{-1} W_2^2(\bar{F}_t, F_t) = \inf_{\bar{F}_t} 2^{-1} \int_{\mathbb{R}^{d_t}} \min_{1 \leq n \leq N_t} ||\bar{F}_{tn} - F_t||_2^2 \, \mathbb{P}(dF_t). \tag{32}$$

For a given $\bar{F}_t$, we define the Voronoi partition $\Gamma = (\Gamma_1, \ldots, \Gamma_{N_t})$ of $\mathbb{R}^{d_t}$ by

$$\Gamma_i = \left\{ F \in \mathbb{R}^{d_t} : i = \arg\min_{1 \leq j \leq N_t} ||F - \bar{F}_{tj}||_p \right\}.$$

Unfortunately, problem (32) is an NP-hard, non-convex optimization problem and even finding a local optimum can be computationally expensive for lattices with many nodes.

3.2.1. Learning lattice nodes stochastic gradient descent

The problem of optimally choosing center points can be overcome by learning the nodes of the lattice in an online fashion from simulated sample paths. This approach has been first proposed by Bally and Pagès (2003) who discretize Brownian motions using an online version of the well-known k-means clustering algorithm that is based on stochastic gradient descent.

In particular, denote $(\beta_k)_{k=1}^{K}$ as a sequence of stepsizes with $0 \leq \beta_k \leq 1$, $k \in [K]$. To find an optimal discretization using stochastic gradient descent, we draw random sequences $(\hat{F}_t^k)_{t=1}^{T}$ from the price process, for $k \in [K]$, and define

$$\bar{F}_{tn}^k = \begin{cases} \bar{F}_{tn}^{k-1} + \beta_k \left(\hat{F}_t^k - \bar{F}_{tn}^{k-1} \right) & \text{if } n = \arg\min_m \left\{ ||\hat{F}_t^k - \bar{F}_{tm}^{k-1}||_2^2, \ m \in [N_t] \right\}, \\ \bar{F}_{tn}^{k-1} & \text{otherwise,} \end{cases} \tag{33}$$

with $\bar{F}_{tn}^0 \equiv 0$, for $n \in [N_t]$, $t = 2, \ldots, T$, $k \in [K]$. From Robbins and Monro (1951) it follows that if the sequence $(\beta_k)_{k=1}^{K}$ satisfies $\sum_{k=1}^{\infty} \beta_k = \infty$ and $\sum_{k=1}^{\infty} \beta_k^2 < \infty$, then the resulting nodes are almost surely local minimizers of (32).

3.2.2. Estimating transition probabilities

The learning method in (33) establishes the location of the nodes of the lattice. A straightforward way to estimate the probability of transitions between nodes in t and stage $t+1$ is to directly estimate the probability of transition between the corresponding cells in the partitions Γ_t and Γ_{t+1} (see Bally and Pagès 2003, for a similar approach).

In particular, let us denote the nodes that result from (33) after a predefined $K \in \mathbb{N}$ iterations as $\bar{F}_{tn}^K$ and denote p_t, $t \in [T-1]$, as the $|N_t| \times |N_{t+1}|$ transition matrix between layers t and $t+1$ with elements p_{tnm}. The (conditional) probability of a state transition from $\bar{F}_{tn}$ to $\bar{F}_{t+1,m}$ becomes

$$p_{tnm} = \frac{\sum_{k=1}^{K} \mathbb{1}_{\Gamma_{tn}}(\tilde{F}_t^k) \mathbb{1}_{\Gamma_{tm}}(\tilde{F}_{t+1}^k)}{\sum_{k=1}^{K} \mathbb{1}_{\Gamma_{tn}}(\tilde{F}_t^k)}, \ n \in [N_t], \ m \in [N_{t+1}], \ t \in [T-1], \tag{34}$$

which is the maximum liklihood estimate for the transition probabilities. We refer to this method as *forwards estimation* since probabilities are estimated from forward simulations of the price process.

A problem with forwards estimation is that, for any given node in $t < T$, the (conditional) expectations on the lattice deviate from the conditional expectations of the continuous process, i.e., typically

$$\mathbb{E}(F_t | F_{t-1} = \bar{F}_{t-1,n}) \neq \sum_{m \in [N_t]} p_{t-1,nm} \bar{F}_{tm} \tag{35}$$

This deviation sends a spurious price signal to the optimal policy which introduces a bias into the approximate value function that adds up recursively. In Section 4.2, we will demonstrate that this bias produces significant gaps between optimal and approximate objective value.

To eliminate this bias, we estimate transition probabilities differently and correct the location of the predecessor nodes layer by layer while going backwards in time. We refer to this method

as *backwards estimation*. Correcting the nodes of the lattice going backwards ensures that the conditional expectations of the lattice equal those of the true process.

Since we have to correct the nodes in $t+1$ before we can compute the conditional expectation of the nodes in t, we begin in $T-1$ and proceed backwards in time through the lattice. Assume that nodes $\bar{F}_{t+1,m}$, $m \in [N_{t+1}]$, are fixed and have already been corrected. Before correcting node n in t, we estimate the vector of transition probabilities, p_{tn}, based on a sequence of samples $(\hat{F}_{t+1}^k)_{k=1}^K$ drawn from the continuous distribution resulting from (3) *conditional on* $F_t = \bar{F}_{tn}^K$. In particular, for $n \in [N_t]$, we set

$$p_{tnm} = K^{-1} \sum_{k=1}^K \mathbb{I}_{\Gamma_{tm}}(\hat{F}_{t+1}^k), \ m \in [N_{t+1}]. \tag{36}$$

Using these probabilities, the conditional expectation of $\bar{F}_{t+1}$ at node $\bar{F}_{tn}^K$ is given by

$$\mathbb{E}[\bar{F}_{t+1}|\bar{F}_t = F_{tn}^K] = \sum_{m=1}^{N_{t+1}} p_{tnm}\bar{F}_{t+1,m}, \ n \in [N_t], \ t \in [T-1], \tag{37}$$

whereas its conditional expectation of the MGBM is

$$\mathbb{E}[F_{t+1}|F_t = \bar{F}_{tn}^K] = \bar{F}_{tn}^K \circ (e^{\mu_t}, \dots, e^{\mu_T}), \tag{38}$$

with $\circ$ as point-wise product. To ensure that these two values match for a given node, we correct the relevant elements of the state vector, $\bar{F}_{tn} = (\bar{F}_{tn1}, \dots, \bar{F}_{tnT})$, by setting

$$\bar{F}_{tnj} = \sum_{m=1}^{N_{t+1}} p_{tnm}\bar{F}_{t+1,m}e^{-\mu_j}, \ j = t+1, \dots, T, \tag{39}$$

for all $n \in [N_t]$. Note that, as long as we can explicitly calculate a correction step as in (39), this algorithm can be adapted to any other Markov process for which the conditional expectation $\Psi_{t+1}(F_{tn}) = \mathbb{E}(F_{t+1}|F_{tn})$ is invertible.

The following proposition shows that this correction step also ensures that the conditional expectations of states at later stages correspond to the conditional expectations of the true process.

Proposition 2. *Let $(F_t)_{t\in[T]}$ be the MGBM process defined in (3), and let $(\bar{F}_t)_{t\in[T]}$ be the lattice produced by backward correction (39) with transition probabilities (36). Then for every $t \in [T]$, $s \in T - t$, and $k \in [N_t]$*

$$\mathbb{E}(\bar{F}_{t+s}|\bar{F}_t = \bar{F}_{tk}) = \mathbb{E}(F_{t+s}|F_t = \bar{F}_{tk}), \quad n \in [N_t].$$

18

Proof. By (39) and (36), the nodes and probabilities of the lattice are such that

$$\mathbb{E}(\bar{F}_{t+r}|\bar{F}_{t+r-1} = \bar{F}_{t+r-1,n}) = \mathbb{E}(F_{t+r}|F_{t+r-1} = \bar{F}_{t+r-1,n}), \quad \forall\, n \in [N_{t+r-1}], \,\forall\, r \in [s].$$

Without loss of generality we can assume that $(F_t)_{t \in [T]}$ takes values in $\mathbb{R}$. Note that because the process is GBM, we have

$$\mathbb{E}(F_{t+s}|F_t = x) = e^{s\mu}x,$$

where μ is the drift of the process. Fixing $s = 2$, we get

$$
\begin{aligned}
\mathbb{E}(\bar{F}_{t+2}|\bar{F}_t = \bar{F}_{tk}) &= \sum_{j \in N_{t+1}} p_{tkj} \sum_{m \in N_{t+2}} p_{tjm}\bar{F}_{t+2,m} = \sum_{j \in N_{t+1}} p_{tkj}\mathbb{E}(\bar{F}_{t+2}|\bar{F}_{t+1} = \bar{F}_{t+1,j}) \\
&= \sum_{j \in N_{t+1}} p_{tkj}e^{\mu}\bar{F}_{t+1,j} = e^{\mu}\sum_{j \in N_{t+1}} p_{tkj}\bar{F}_{t+1,j} = e^{\mu}\mathbb{E}(\bar{F}_{t+1}|\bar{F}_t = \bar{F}_{tk}) \\
&= e^{\mu}e^{\mu}\bar{F}_{tk} = \mathbb{E}(F_{t+2}|F_t = \bar{F}_{tk}).
\end{aligned}
$$

The general case for $s > 2$ follows by induction. $\qquad\square$

Note that Proposition 2 works for the more general setting in which the functions Ψ_t are affine. This in particular, covers the case of most price processes used in practice, such as Brownian motion, GBM, Ornstein-Uhlenbeck processes, and jump diffusion models.

In summary, our discretization method consists of two steps. In the first step, we use stochastic gradient descent to approximate the unconditional distributions of F_t by finding a discrete distribution $\bar{F}_t$ with minimal Wasserstein distance distance to F_t. In the second step, the method estimates transition probabilities while going backwards in time to create conditional distributions that are unbiased.

The resulting lattice provides a discretization of the continuous process that can be combined with any type of backwards dynamic programming. However, since the final resource state that is under control of the decision-maker is still high-dimensional, we resort to Markov chain stochastic dual dynamic programming (MC-SDDP) to solve the problem (Löhndorf et al. 2013, Löhndorf and Shapiro 2019).

3.3. Stochastic Dual Dynamic Programming on Lattices

In this section, we briefly summarize how MC-SDDP is used to compute the optimal value of storage based on a scenario lattice as description of the underlying randomness.

The SDDP algorithm goes back to the seminal work of Pereira and Pinto (1991) and its Markovian version to Mo et al. (2001). First risk-averse variants were proposed in Shapiro (2011) and Guigues and Römisch (2012). Convergence proofs of the algorithm can be found

---
Algorithm 1 Markov Chain SDDP

Require: Nodes $\bar{F}$, transition probabilities (p_{tnm}), number of iterations L, sart state R_0

1: $\hat{\mathcal{V}}^0 \leftarrow 0$, $l \leftarrow 1$ ▷ initialize
2: **while** $l \leq L$ **do**   ▷ main loop
3:     $\hat{\bar{F}}_1 \leftarrow F_1$
4:     **for** $t \in [T]$ **do**   ▷ forward pass
5:        $\hat{R}_t^l \leftarrow \mathcal{P}_t(\hat{\mathcal{V}}_t^{l-1}, \hat{\bar{F}}_t^l, R_{t-1}^l)$
6:        Sample $\hat{\bar{F}}_{t+1}^l$ given $\hat{\bar{F}}_t^l$
7:     **end for**
8:     **for** $t \in T : 2$ **do**   ▷ backward pass
9:        **for** $n \in [N_t]$ **do**
10:          $a_{ln} \leftarrow V_t(\hat{\bar{F}}_{tn}, \hat{R}_t^l)$, $b_{ln} \leftarrow \mathcal{D}_t(\hat{\bar{F}}_{tn}, \hat{R}_t^l)$
11:        **end for**
12:        **for** $m \in [N_{t-1}]$ **do**   ▷ generate new cuts
13:          $a_l(\hat{\bar{F}}_{t-1,m}, \hat{R}_t^l) \leftarrow \sum_{n \in [N_t]} p_{tmn} a_{ln}$, $b_l(\hat{\bar{F}}_{t-1,m}, \hat{R}_t^l) \leftarrow \sum_{n \in [N_t]} p_{tmn} b_{ln}$
14:        **end for**
15:     **end for**
16:     $l \leftarrow l + 1$
17: **end while**
---

in Philpott and Guan (2008) for linear problems, Girardeau and Philpott (2014) for convex problems, and Guigues (2016) for risk-averse convex problems.

SDDP is an iterative procedure and alternates between simulating the decision policy (forward pass) and recursively updating the expected value function, $\hat{\mathcal{V}}_t$ (backward pass). We summarize SDDP in Algorithm 1 defining L as the number of iterations. Furthermore, we define $\mathcal{P}_t(\mathcal{V}_t^{l-1}, \hat{\bar{F}}_t^l, \hat{R}_{t-1}^l)$ as a function that returns a primary solution for (20) for lattice node $\hat{\bar{F}}_t^l$ using the value function approximation $\hat{\mathcal{V}}_t^{l-1}$ and the resource state $\hat{R}_{t-1}^l$. Likewise $\mathcal{D}_t(\bar{F}_{tn}, \hat{R}_t^l)$ returns the optimal dual multipliers for the constraints of (20) at $\bar{F}_{tn}$ for which the state $\hat{R}_t^l$ appears on the right hand side.

During the forward pass in line 4 to 7, a sequence of states, $(\hat{\bar{F}}_{tn}^l)_{t=2}^T$, is drawn from the lattice process to generate a sequence of sample resource states, $(\hat{R}_t^l)_{t=2}^T$, by using the current approximation of the expected value function in the specification of (20).

After a forward pass, SDDP performs a *backward pass* in line 8 to 14 adding new hyperplanes to the sets of supporting hyperplanes for each approximate expected value function $\hat{\mathcal{V}}_t^l$, using the dual solution of (20) at every node for the resource states sampled in the forward pass.

Algorithm 1 summarizes SDDP for the case of maximizing the *expected* profits from storage operations, i.e., the valuation problem for a risk-neutral agent. To obtain approximate expected value functions for the nested CVaR as an objective, we use the method proposed by Philpott et al. (2013) for scenario lattices.

The method of Philpott et al. (2013) is based on the idea of reweighing the transition probabilities, which avoids a state space extension required in alternative formulations of the nested CVaR. In general, for a discrete random variable X with possible realizations $X_1, \ldots, X_M$ and corresponding probabilities $p_1, \ldots, p_M$, the risk measure $\mathbb{U}_{t,\alpha,\lambda}$ has the dual representation

$$\mathbb{U}_{t,\alpha,\lambda}(X) = \min_{\xi \in \mathcal{U}} \sum_{i \in [M]} p_i \xi_i X_i, \tag{40}$$

where

$$\mathcal{U} = \left\{ \xi \in \mathbb{R}^M : \sum_{i \in [M]} \xi_i p_i = 1, \xi_i = (1-\lambda) + \lambda \eta_i, 0 \leq \eta_i \leq \frac{1}{\alpha}, \, \forall \, i \in [M] \right\}. \tag{41}$$

Using (17), this enables us to write

$$\mathcal{V}_t(\bar{F}_{tn}, R_{t+1}) = \min_{\xi \in \mathcal{U}} \sum_{m \in [N_{t+1}]} p_{tnm} \xi_m V_{t+1}(\bar{F}_{t+1,m}, R_{t+1}). \tag{42}$$

Hence, $\mathbb{U}_{t,\alpha,\lambda}$ can be viewed as an expectation with a changed probability measure. The supergradients b can be constructed from this representation by noting that

$$\partial_{R_{t+1}} \mathcal{V}_t(\bar{F}_{tn}, R_{t+1}) = \sum_{m \in [N_{t+1}]} p_{tnm} \xi_m^* \partial_{R_{t+1}} V_{t+1}(\bar{F}_{t+1,m}, R_{t+1}), \tag{43}$$

where weights, ξ_m^*, $m \in [N_{t+1}]$, are optimal for R_{t+1} in (42).

As the optimal weights are easily identified by sorting the values $V_{t+1}(\bar{F}_{t+1,m}, R_{t+1})$, $t \in [T-1]$, the above method can be used to construct value function approximations during the backward pass similar to the expectation case.

We adapt this idea to the piecewise-linear approximate post-decision value functions. In this setting, the values corresponding to (42) at sample resource states, $\hat{R}_t^l$, are given by

$$a_l(\bar{F}_{tn}, \hat{R}_{t+1}^l) = \sum_{m=1}^{N_{t+1}} p_{tnm} \xi_m^* \max_{\pi_{t+1} \in \Pi_{t+1}(\hat{R}_{t+1}^l)} \left\{ C_{t+1}(\bar{F}_{t+1,m}, \pi_{t+1}) + \hat{\mathcal{V}}_{t+1}^l(\bar{F}_{t+1,m}, R_{t+2}(\pi_{t+1})) \right\}, \tag{44}$$

$n \in [N_t]$, $t = T-1, \ldots, 1$, $l \in [L]$.

Accordingly, as in (43), the slope vectors of the supporting hyperplanes at $\hat{R}_{t+1}^l$, are given by

$$b_l(\bar{F}_{tn}, \hat{R}_{t+1}^l) = \partial_{\hat{R}_{t+1}^l} a_l(\bar{F}_{tn}, \hat{R}_{t+1}^l), \; n \in [N_t], \; t = T-1, \ldots, 1, \; l \in [L], \tag{45}$$

which can be obtained from the dual solution of (44).

4. Numerical Results

In this section we summarize the results of extensive numerical experiments based on model instances from the extant literature. In particular, we compare the quality of our solution with a known bound for the complete market case and study the influence of the dynamic utility measure on storage value as well as discounted profits for the incomplete market case.

4.1. Setup

For our numerical analysis, we use the same parameters as reported in Lai et al. (2010) who calibrate four driftless versions of the MGBM price process (3), using the information available at the closing of NYMEX Henry Hub on 2006-01-03 (Spring), 2006-06-01 (Summer), 2006-08-31 (Fall), and 2006-12-01 (Winter). The interest rates as reported by the Department of Treasury on these days define the monthly discount factor of the model. The initial state of the MGBM, F_1, is given by the spot price and the futures prices of the first 23 monthly maturities on each of the four days. Implied volatilities of the 23 futures prices are obtained from prices of NYMEX call options on natural gas futures. Storage instances vary in their injection (withdrawal) limits ranging from 0.15 mmBtu/month (0.3 mmBtu/month) for instance A to 0.45 mmBt/month (0.9 mmBt/month) in instance C. If not stated otherwise, model instance A with Spring prices will serve as our base case. The resulting problems have 24 decision stages and cover 2 years of monthly planning. All constant model parameters are summarized in Table A.4 in Appendix A.

In this setup, there are 23 prices, 23 contracts positions to keep track of, and 1 storage state, so that we end up with a high dimensional state space, which is a challenge for conventional solution techniques. For example, consider a lattice with 10 nodes per stage (and a single root node at stage one). If we were to replace the lattice with the equivalent scenario tree, we would end up with with $\approx 10^{23}$ terminal nodes compared with 231 lattice nodes.

Unless stated otherwise, for our numerical experiments, we ran SDDP for 1000 iterations and constructed a lattice with 1000 nodes per stage and backwards estimation as base case. All lattices were built on the basis of $K = 10^6$ price scenarios for the discretization and forwards estimation and 10^3 transition scenarios per node for backwards estimation. Profit estimates for a given policy are obtained by averaging cumulative profits from 10^5 forward passes of the SDDP algorithm, without performing backward passes or updating the value function approximations.

The algorithms have been implemented in Java using FICO Xpress BCL to solve the linear subproblems. All computations were conducted on a Dual-Xeon E5-2650 with 32 cores and 128 GB memory. Computation times are approximately 5 minutes for the base case.

(a) Multivariate GBM

(b) 10-node lattice

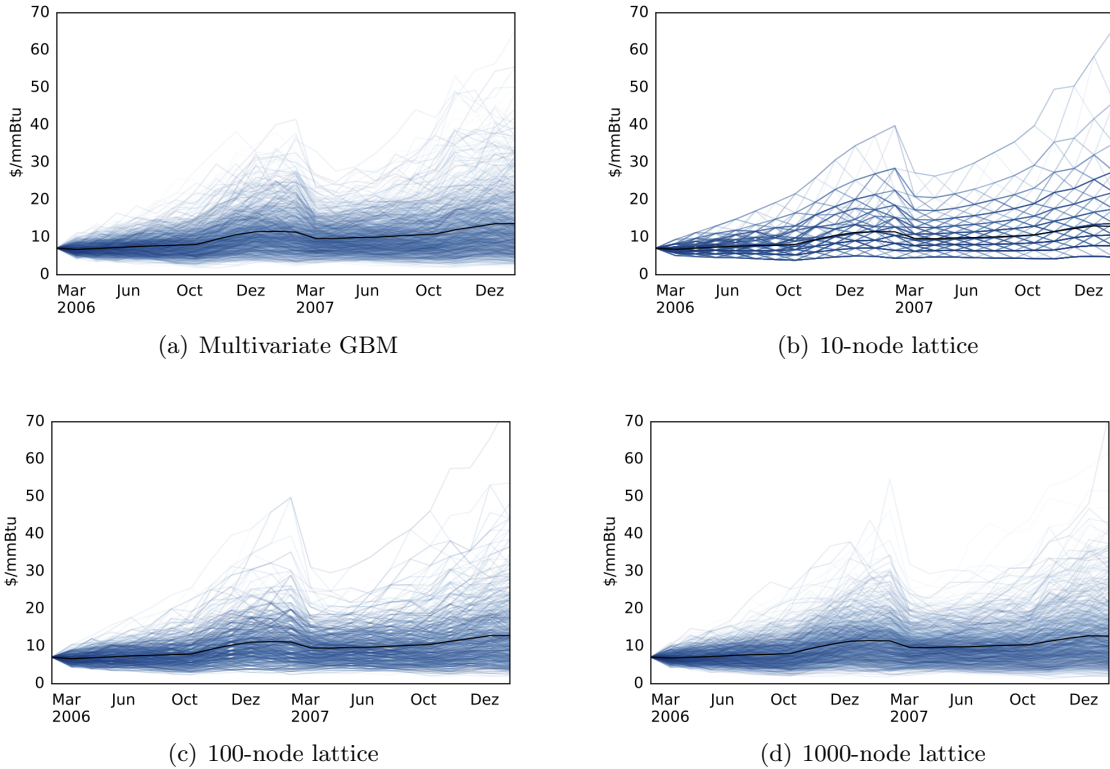(c) 100-node lattice

(d) 1000-node lattice

Figure 2: Time series simulation of the spot price showing 1000 sample paths drawn from the continuous process as well as three different lattices with 10, 100, and 1000 nodes per stage. Darker areas of the graph represent nodes and transitions that are more frequently sampled, while lighter regions are sampled less frequently. The unconditional stagewise means are indicated by a black solid line.

4.2. Approximation Quality

Figure 2 shows the evolution of the front-month prices of the continuous price process along with lattice discretization of different granularity. As can be seen, despite the high dimensionality of the continuous process, our discretization method manages to select representative price states for the unconditional distribution. Important characteristics of the process like mean, variance, skewness and serial dependence remain intact as well.

To assess the quality of the decisions made when the lattice is used, we will compare the optimal objective value with a known upper bound from the literature. Lai et al. (2010) compare heuristics for storage valuation in a complete market and derive a dual upper bound of the storage valuation problem for a number of problem instances. We use this bound and the same model parameters as Lai et al. (2010) to study the optimality gap of the optimal objective value of the stochastic program.

For each price model, we vary number of nodes, number of state transitions per node, as well as estimation method ceteris paribus. The results of this comparison are summarized in

Season	Nodes (N)	Transitions (K)	Method	Percentage of DUB (SE)		
				Lattice SDDP	Lattice RI	Process RI
Spring	1000	1000	F	1.484 (0.000)	1.048 (0.001)	0.982 (0.000)
			B	0.966 (0.004)	0.964 (0.003)	
Summer			F	1.456 (0.000)	1.048 (0.000)	0.982 (0.000)
			B	0.967 (0.004)	0.965 (0.004)	
Fall			F	1.456 (0.001)	1.043 (0.001)	0.977 (0.000)
			B	0.947 (0.002)	0.945 (0.002)	
Winter			F	1.829 (0.001)	1.170 (0.002)	0.918 (0.001)
			B	0.882 (0.001)	0.861 (0.002)	
Spring	100	1000	F	1.506 (0.003)	1.038 (0.001)	0.982 (0.000)
			B	0.949 (0.002)	0.946 (0.002)	
		100	F	1.596 (0.002)	1.061 (0.001)	
			B	0.982 (0.010)	0.980 (0.010)	
		10000	F	1.461 (0.000)	1.045 (0.001)	
			B	0.967 (0.001)	0.965 (0.001)	
	10000	1000	F	1.418 (0.000)	1.041 (0.000)	
			B	0.974 (0.004)	0.973 (0.004)	

B = backwards estimation, F = forwards estimation, DUB = dual upper bound, RI = rolling intrinsic value, SE = standard error with 5 replications

Table 1: Comparison of discretization quality for several price models and lattice parameters

Table 1. The table compares the objective value of the stochastic program (Lattice SDDP), the rolling intrinsic value computed with the lattice (Lattice RI), and the rolling intrinsic value computed with the continuous process (Process RI) as percentage of the dual upper bound (DUB). We make the following observations:

1. Policies that are based on lattices that have been created using forward estimation yield values that are 42-83% higher than the dual upper bound (DUB) provided in Lai et al. (2010) and are thus biased. The reason for this is that conditional means are not accurately represented by the corresponding lattices, which gives the optimal lattice policy the opportunity to speculate on expected price movements that do not exist with the continuous process (3). By matching the first moments using backwards estimation we are able to avoid this bias. Consequently, with backwards estimation, Lattice SDDP deviates less than 2.5% on average from the DUB for the base case.

2. With backwards estimation, Lattice SDDP and Process RI are almost equally close to the DUB, which shows that our approach yields results that are close to optimal in a complete markets setting. Moreover, there is no significant gap between Lattice RI and Lattice SDDP, which provides additional evidence for the near-optimality of the rolling intrinsic value in a complete market, because the optimum obtained using SDDP is a tight upper bound for the storage value with prices defined by the lattice.

3. The number of nodes and the number of samples K used to train the lattice have little

impact on solution quality. While the results are slightly better with denser lattices, the difference is rather small.

4. Looking at the results for backward estimation, it can be seen that the RI and SDDP essentially produce the same result. The only notable exception is the Winter instance, where Lattice SDDP is 2% higher than Lattice RI. Interestingly, this is also the instance, where the difference to DUB is highest. A possible explanation for this is that the futures curve in this case has a much lower Summer-Winter spread and therefore there exists no obvious arbitrage strategy, which makes the problem more difficult to solve.

In summary, the results provide evidence that our approach produces near-optimal results despite the high-dimensionalty of the underlying optimization problem.

4.3. Valuation in Incomplete Markets

For our next analysis, we drop the assumption of market completeness to study the value of storage under varying degrees of risk-aversion. To that end, we augment the price process estimated in Lai et al. (2010) by non-zero drift terms μ, which we estimate from monthly historical returns (see Table A.3 in the Appendix Appendix A). Of course, our approach is not specific to the multivariate geometric Brownian motion (3). Traders can use any process that they see fit to describe futures dynamics. In particular, it seems reasonable to capture mean reversion properties of the prices and seasonally varying variances. However, a discussion of such a model is beyond the scope of the paper, which is why we use the simple process (3) also in this section.

Following Lai et al. (2010), we use 12 model instances that combine four variants of the price model with parameters of the three different storage systems from Lai et al. (2010). In particular, we are interested in the impact of the risk measure on storage value and decision policy. We consider the following four risk-preferences.

1. **Risk-neutral**: maximize the expected value (EV).
2. **Low risk aversion**: maximize a convex combination of expected value and CVaR with $\lambda = 0.05$ and $\alpha = 0.05$ (Low).
3. **High risk aversion**: maximize only CVaR with $\alpha = 0.05$ and $\lambda = 1$ (High).
4. **Perfect risk aversion**: the rolling intrinsic policy (RI), which yields the intrinsic value as price of the storage.

Table 2 reports the computational results, including the value of storage defined as the objective of the stochastic optimization problem as well as expectation and standard deviation of (discounted) overall profits from following the corresponding policy. We use the standard deviation of the profits as a proxy for the riskiness of the policy. The results shows that a

25

Instance		Storage Value				Expected Profit				Std Dev			
Season	Storage	EV	Low	High	I	EV	Low	High	RI	EV	Low	High	RI
Spring	A	4.89	4.22	3.83	3.76	4.83	4.32	4.19	4.14	3.78	0.43	0.19	0.27
	B	6.10	5.19	4.54	4.31	6.05	5.43	5.14	5.07	4.99	0.84	0.30	0.41
	C	6.72	5.65	4.80	4.40	6.79	5.95	5.60	5.50	5.72	0.98	0.39	0.54
Summer	A	5.86	4.68	4.26	4.26	5.87	4.88	4.66	4.63	5.78	0.74	0.22	0.27
	B	7.89	6.26	5.56	5.29	7.69	6.73	6.15	6.04	8.12	1.50	0.32	0.38
	C	8.75	6.77	5.80	5.48	8.77	7.37	6.53	6.61	9.11	1.93	0.39	0.52
Fall	A	5.10	4.16	3.70	3.53	5.06	4.31	4.12	4.01	6.45	0.52	0.23	0.30
	B	7.63	6.32	5.53	5.10	7.54	6.55	6.20	6.19	7.94	0.80	0.33	0.47
	C	8.88	7.45	6.46	5.94	8.98	7.77	7.32	7.25	8.94	1.11	0.42	0.60
Winter	A	2.28	1.79	1.21	0.93	2.26	1.94	1.69	1.64	2.66	0.58	0.25	0.36
	B	3.24	2.44	1.66	1.17	3.28	2.67	2.35	2.29	3.57	0.77	0.35	0.54
	C	3.64	2.81	1.85	1.27	3.65	3.11	2.65	2.55	3.94	1.00	0.41	0.63

EV = maximize expected profit, Low = maximize nested CVaR with $\alpha = 0.05$ and $\lambda = 0.05$, High = maximize nested CVaR with $\alpha = 0.05$ and $\lambda = 1.0$, I = intrinsic value, RI = rolling intrinsic value

Table 2: Risk-adjusted value of storage compared to expected discounted profits (in \$/mmBtu)

risk-neutral decision-maker (EV) achieves the highest profits but is also exposed to a significant amount of risk. Risk exposure can be controlled by setting the parameters of the risk measure accordingly, with the rolling intrinsic policy (RI) being the least risky policy. These findings are consistent across all problem instances and are in line with our theoretical investigation from Proposition 1.

The effect of the risk measure on the decision policy is illustrated by plotting the distribution of the cumulative profits in Figure 3. While the risk-neutral planner invests in the beginning and postpones realization of profits to the end of the time horizon, the risk-averse agent realizes a larger fraction of profits upfront with only marginal gains over time.

The storage value equals expected rewards only in the risk-neutral case. In all other cases, storage value is less than expected rewards, as the risk-averse decision-maker puts more weight on expected tail loss. As can be seen from Table 2, the storage value is significantly higher if risk aversion is low, where the difference between RI and EV ranges from 15-39%. This shows that perfect risk aversion leads to costly hedging decisions, which makes the rolling intrinsic policy completely unsuitable for a risk-neutral decision-maker. This is in contrast to the common view in the literature that views the rolling intrinsic policy as a close to optimal heuristic for pricing and operating gas storage contracts.
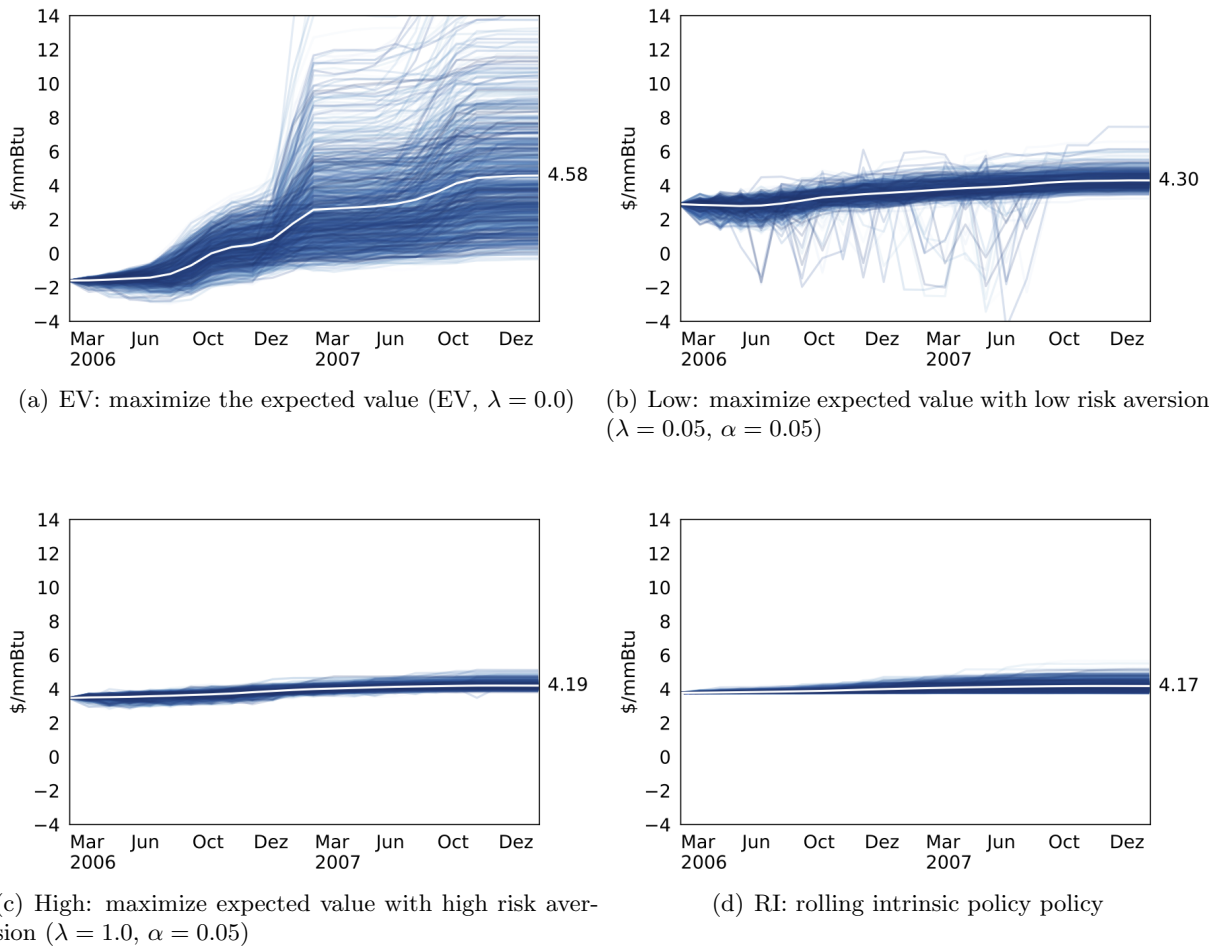
(a) EV: maximize the expected value (EV, $\lambda = 0.0$)

(b) Low: maximize expected value with low risk aversion ($\lambda = 0.05$, $\alpha = 0.05$)

(c) High: maximize expected value with high risk aversion ($\lambda = 1.0$, $\alpha = 0.05$)

(d) RI: rolling intrinsic policy policy

Figure 3: Cumulative profit scenarios for different risk preferences, with the sample average as white solid line.

5. Conclusion

We propose a new approach for valuation of gas storage contracts in incomplete markets. We show that in incomplete markets, the rolling intrinsic policy is not optimal for a risk-neutral decision-maker. Furthermore, we argue that the value of storage can only be computed by finding the optimal solution to a stochastic-dynamic control problem that jointly considers futures trading as well as storage operation and explicitly models a trader's risk preferences. In particular, this implies that the rolling intrinsic value is not a consistent pricing mechanism as it would erroneously combine the optimal policy of a perfectly risk-averse decision-maker with the valuation function of a risk-neutral agent.

To compute the storage values, we propose an asset-backed multistage stochastic optimization model to compute the value of storage. As the multistage stochastic program cannot be solved in its extensive form, we reformulate the problem as a high-dimensional dynamic pro-

gram. To deal with the curse of dimensionality, we propose an approximation strategy that first discretizes the high-dimensional price process to a scenario lattice and then solves the resulting problem using stochastic dual dynamic programming.

The discretization method minimizes the Wasserstein distance between the scenario lattice and the continuous price process by learning an optimal quantization of the unconditional distributions of the stochastic futures price process. We find that small errors in the first moment of the conditional distributions of a scenario lattice are magnified in the optimization and introduce a considerable bias into the optimal policy and thereby into the computed storage values. We show that this problem can be mitigated by estimating transition probabilities going backwards in time, which ensures that the first moments of the lattice corresponds to those of the continuous process.

We demonstrate in a numerical study that our approach finds near-optimal solutions and how our approach can be used to obtain the correct value of storage in an incomplete market setting.

Interesting extensions of the models proposed in this paper include the application of our methodology to more sophisticated models of futures curve dynamics and models that explicitly account for illiquidity in the form of price impact of large trades as well as bid ask spreads. In particular, the explicit modeling of market illiquidity would improve the model, since it avoids unbounded solutions by penalizing large trades. While this extension is straightforward from a modeling point of view, estimating its parameters requires order book level data, which is neither easy to obtain nor easy to process. Nevertheless, developing such models would be an important building block to accurately price storage in incomplete markets.

References

Bally, V., Pagès, G., 2003. A quantization algorithm for solving multidimensional discrete-time optimal stopping problems. Bernoulli 9, 1003–1049.

Bjerksund, P., Stensland, G., Vagstad, F., 2011. Gas storage valuation: Price modelling vs optimization methods. The Energy Journal 32, 203–228.

Bonnans, J., Cen, Z., Christel, T., 2012. Energy contracts management by stochastic programming techniques. Annals of Operations Research 200, 199–222.

Boogert, A., de Jong, C., 2008. Gas storage valuation using a Monte Carlo method. The Journal of Derivatives 15, 81–98.

Boogert, A., de Jong, C., 2012. Gas storage valuation using a multifactor price process. Journal of Energy Markets 12.

Boudier, E., Lorenz, S., Marten, I., 2015. A golden period for asset-

backed trading. URL: `https://www.bcg.com/publications/2015/energy-environment-commodity-trading-golden-period-asset-backed-trading.aspx`.

Carmona, R., 2009. Indifference Pricing: Theory and Applications. Princeton Series in Financial Engineering, Princeton University Press.

Carmona, R., Ludkovski, M., 2010. Valuation of energy storage: an optimal switching approach. Quantitative Finance 10, 359–374.

Chen, Z., Forsyth, P., 2007. A semi-Lagrangian approach for natural gas storage valuation and optimal operation. SIAM Journal on Scientific Computing 30, 339–368.

Clewlow, L., Strickland, C., 2000. Energy Derivatives: Pricing and Risk Management. Lacima Publications.

Conejo, A., Morales, L., Kazempour, S., Siddiqui, A., 2016. Investment in Electricity Generation and Transmission: Decision Making under Uncertainty. Springer International Publishing.

Dupačová, J., Gröwe-Kuska, N., Römisch, W., 2003. Scenario reduction in stochastic programming an approach using probability metrics. Mathematical Programming, Series B 95, 493–511.

EIA, 2016. Annual Energy Outlook 2016. U.S. Energy Information Administration. URL: `http://www.eia.gov/outlooks/archive/aeo16`.

Eydeland, A., Wolyniec, K., 2003. Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging. Wiley Finance, Wiley.

Felix, B., Weber, C., 2012. Gas storage valuation applying numerically constructed recombining trees. European Journal of Operational Research 216, 178–187.

Girardeau, P., L.V., Philpott, A., 2014. On the convergence of decomposition methods for multistage stochastic convex programs. Mathematics of Operations Research 40, 130–145.

Graf, S., Luschgy, H., 2000. Foundations of Quantization for Probability Distributions. Lecture Notes in Mathematics, Springer Berlin Heidelberg.

Gray, J., Khandelwal, P., 2004a. Realistic gas storge models II: Trading strategies. Commodities Now September, 1–5.

Gray, J., Khandelwal, P., 2004b. Towards a realistic gas storage model. Commodities Now June, 1–4.

Guigues, V., 2016. Convergence analysis of sampling-based decomposition methods for risk-averse multistage stochastic convex programs. SIAM Journal on Optimization 26, 2468–2494.

Guigues, V., Römisch, W., 2012. Sampling-based decomposition methods for multistage stochastic programs based on extended polyhedral risk measures. SIAM Journal on Optimization 22, 286–312.

Kall, P., Mayer, J., 2006. Stochastic Linear Programming: Models, Theory, and Computation. International Series in Operations Research & Management Science, Springer US.

Kaut, M., Wallace, S., 2007. Evaluation of scenario generation methods for stochastic programming. Pacific Journal of Optimization 3, 257–271.

Klöppel, S., Schweizer, M., 2007. Dynamic indifference valuation via convex risk measures. Mathematical Finance 17, 599–627.

KYOS, 2015. KyStore: Gas Storage. URL: `http://www.kyos.com/wp-content/uploads/2016/10/KYOS-KyStore.pdf`.

Lacima, 2017. Valuation & Optimization Suite: Storage. URL: `http://lacimagroup.com/SiteFiles/lacimagroupcomau/Storage_Fact_Sheet.pdf`.

Lai, G., Margot, F., Secomandi, N., 2010. An approximate dynamic programming approach to benchmark practice-based heuristics for natural gas storage valuation. Operations Research 58, 564–582.

Löhndorf, N., Shapiro, A., 2019. Modeling time-dependent randomness in stochastic dual dynamic programming. European Journal of Operational Research 273, 650–661.

Löhndorf, N., Wozabal, D., Minner, S., 2013. Optimizing trading decisions for hydro storage systems using approximate dual dynamic programming. Operations Research 61, 810–823.

Malyscheff, A., Trafalis, T., 2017. Natural gas storage valuation via least squares monte carlo and support vector regression. Energy Systems 8, 815–855.

Mo, B., Gjelsvik, A., Grundt, A., 2001. Integrated risk management of hydro power scheduling and contract management. IEEE Transactions on Power Systems 16, 216–221.

MSCI, 2015. FEA @Energy Suite. URL: `https://www.msci.com/documents/1296102/1636401/FEA_Factsheet.pdf`.

Pereira, M., Pinto, L., 1991. Multi-stage stochastic optimization applied to energy planning. Mathematical Programming 52, 359–375.

Pflug, G., 2001. Scenario tree generation for multiperiod financial optimization by optimal discretization. Mathematical Programming, Series B 89, 251–271.

Pflug, G., Pichler, A., 2012. A distance for multistage stochastic optimization models. SIAM Journal on Optimization 22, 1–23.

Philpott, A., Guan, Z., 2008. On the convergence of stochastic dual dynamic programming and related methods. Operations Research Letters 36, 450–455.

Philpott, A., de Matos, V., Finardi, E., 2013. On solving multistage stochastic programs with coherent risk measures. Operations Research 61, 957 – 970.

Robbins, H., Monro, S., 1951. A stochastic approximation method. Annals of Mathematical Statistics , 400–407.

Ruszczyński, A., 2010. Risk-averse dynamic programming for markov decision processes. Mathematical Programming 125, 235–261.

Ruszczyński, A., Shapiro, A., 2006. Conditional risk mappings. Mathematics of Operations Research 31, 544–561.

Secomandi, N., 2010. Optimal commodity trading with a capacitated storage asset. Management Science 56, 449–467.

Secomandi, N., 2015. Merchant commodity storage practice revisited. Operations Research 63, 1131–1143.

Shapiro, A., 2009. On a time consistency concept in risk averse multistage stochastic programming. Operations Research Letters 37, 143 – 147.

Shapiro, A., 2011. Analysis of stochastic dual dynamic programming method. European Journal of Operational Research 209, 63 – 72.

Shapiro, A., Dentcheva, D., Ruszczyński, A., 2009. Lectures on Stochastic Programming: Modeling and Theory. MOS-SIAM series on optimization, Siam.

Shapiro, A., Tekaya, W., da Costa, J., Soares, M., 2013. Risk neutral and risk averse stochastic dual dynamic programming method. European Journal of Operational Research 224, 375 – 391.

Thompson, M., Davison, M., Rasmussen, H., 2009. Natural gas storage valuation and optimization: A real options application. Naval Research Logistics 56, 226–238.

Villani, C., 2003. Topics in optimal transportation. volume 58 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI.

Appendix A. Problem Parameters

t	1	2	3	4	5	6	7	8	9	10	11	12
μ_t		0.0092	0.0095	0.0101	0.0101	0.01	0.0097	0.0098	0.01	0.01	0.0099	0.0101
t	13	14	15	16	17	18	19	20	21	22	23	24
μ_t	0.0102	0.0108	0.0113	0.0114	0.0103	0.0107	0.0111	0.0113	0.0116	0.0111	0.0112	0.0109

Table A.3: Drift parameters of the physical measure.

Parameter	Notation	Instance	Value(s)	Unit
Storage limit	$\overline{s}$		1.00	mmBtu
Initial resource state	R_1		0	
Injection (withdrawal) limit	$\overline{i}, \overline{w}$	A	0.15 (0.30)	mmBtu
		B	0.30 (0.60)	mmBtu
		C	0.45 (0.90)	mmBtu
Injection cost	c^i		0.02	US\$/mmBtu
Withdrawal cost	c^w		0.01	US\$/mmBtu
Injection loss	d^i		0.99	mmBtu/mmBtu
Withdrawal loss	d^w		1.01	mmBtu/mmBtu
Interest rates	δ	Spring	4.74	% p.a.
		Summer	5.05	% p.a.
		Fall	5.01	% p.a.
		Winter	4.80	% p.a.
Time horizon	T		24	month

Table A.4: Parameters for the numerical study.