
A Stochastic Programming Framework for the Valuation of Electricity Storage

Pedram Mokrian, Moff Stephen
Stanford University

Pedram Mokrian , Ph.D. Candidate
Department of Management Science and Engineering
Terman Engineering Center, 3rd Floor
Stanford University, Stanford, CA 94305-4026
tel: 650-804-5129 fax: 650-723-1614
pmokrian@stanford.edu

Abstract

It is often assumed in practice that electricity cannot be stored, even though in reality the storage of electricity is technically feasible. The issue is not whether electricity can be stored, but whether it can be stored efficiently and economically. In this paper, we present some of the current storage technologies and discuss their applications within power markets. We focus on the use of storage for intraday arbitrage and develop several models for optimizing the operation of a storage facility over a 24 hour period. The optimization problem has been framed as a linear program, a multi-stage stochastic program and a dynamic program. Two specific storage technologies, namely compressed air energy storage (CAES) and the sodium sulfur (NaS) battery, are modeled and analyzed within the separate optimization frameworks. We conclude with a sensitivity analysis of the model parameters and present an overall cost-benefit analysis of using electricity storage for arbitrage.

1 Introduction

Electricity, which is generated, purchased and distributed in its own unique market, currently plays a far more significant role in the fundamental aspects of daily life than ever before. As industrialized and developing countries alike thirst for more and more power, there will be a strong need for innovation in energy efficiency and generation and transmission asset utilization to keep up with the ever increasing energy demands. The IEA forecasts that the U.S. alone will require an additional 350 GW of new generation capacity over the next 25 years [1]. This anticipated increase in the generation capacity must be addressed by taking into account increasing fuel costs, tighter environmental restrictions, and an aging transmission and distribution system, which already causes losses of some \$100 billion annually due to reliability issues [2].

The future of the power industry's generation portfolio will inevitably include a strong mix of renewable technologies. The recent trends in new renewable generation facilities have shown positive growth, especially in the wind sector, where installed capacity within the U.S. has grown by over 300% over the past six years. One of the significant drawbacks of wind technology is the intermittency of the generation due to the variability of wind. This variability leads to increased procurement costs for various reliability services, increased transmission congestion costs, and additional strain on the remaining generation units that must be called upon in the generation voids created by wind [3]. This problem may be minor for the current wind generation levels which account for a fraction of a percentage of the total capacity in the U.S., but will become more drastic as wind generation levels increase [4].

Energy storage may help mitigate all of the listed problems, and in many cases may be a complete solution. This helps motivate the fact that currently 2.5% of the total electric power delivered within the U.S. is cycled through a storage facility. Although this number may appear significant at first, it pales in comparison to the 10% and 15% of total power cycled through storage facilities in Europe and Japan respectively [5]. Drivers for the construction of such storage facilities range from coping with increased peak energy demand to the need for improved energy security and reliability [6].

This paper focuses on the valuation of energy storage technologies through arbitrage profits within large scale power markets. Although there are countless storage applications in fields such as distributed generation and on-site reliability, they will not be considered in this analysis; however, several results may easily be extended to these areas. We begin by introducing the arbitrage profit valuation framework along with an overview of the existing literature on the valuation of electricity storage. We follow with a brief outline of

storage technologies and their applications, and present a mathematical model for such facilities. From there, we incorporate the storage model into a series of mathematical programs that seek to optimize the operation policy of a given storage plant by maximizing intraday arbitrage profits. We conclude with a summary of the findings and the arbitrage valuation results from the analysis.

2 Arbitrage and Electricity Storage

Irrespective of the market structure, pricing within power markets is a complex process marred by price spikes, jumps, high volatility, and seasonal effects, having the combined effect of making the market incomplete from a financial standpoint. Prices determined in electricity spot markets are extremely volatile and are highly dependent on time of day and seasonal effects. The daily fluctuations provide a clear potential for arbitrage profits by purchasing electricity during off peak hours (typically at night) when the price is low, storing that power, and then selling the stored power during peak hours at a significantly higher price. It is claimed, however, that the inexistence of electricity as a tangible, storable asset, and the inadequate relationship between markets eradicates almost all arbitrage arguments. This is a topic that we will specifically address within this paper.

Energy storage technologies enable an offset in time between power generation and power consumption. This ability to store energy has a profound impact not only on the physical characteristics of the power grid, but also on the financial and investment strategies of power market participants. Although the impacts of storage may be well understood, the viability and economic feasibility of such systems remains a topic of debate. The existing literature on the economics of storage facilities covers a broad range of topics, and in this section we will briefly summarize several of the results.

The notion of electricity storage to increase the efficiency of the generation profile to meet peak demand was first introduced in the 19th century. In an article in “Science” magazine published in 1889 [8], the generation issue of coping with the changing daily demand profile is addressed: “[generators] are called on for a maximum service for a short time, which is followed by a smaller demand during the rest of the night. It is patent that such a method of production cannot be economical, for the plant must be idle, or working to but a fraction of its capacity, most of the time.” Although storage was suggested as a solution to this dilemma, history dictates that this issue was addressed through the introduction of peak generation plants.

As mentioned earlier, it is not the technical feasibility of storage that has limited its deployment in power

systems, nor the lack of applications. Schaber et al. [12], Harty et al. [13], Dell and Rand [14] and the EPRI Handbook of Energy Storage [5] all provide an in depth look at the current storage technologies and their applications and are excellent references for an overview of storage technologies. It is the understanding of the benefits of storage and the economic valuation of such systems that has largely impeded their use.

In terms of specific benefits, Schaber et al. [12] argue that energy storage is required for the widespread renewable energy to become reality, in order to match the intermittent renewable supply with customer demand. Makansi [19] furthers this notion by claiming that frequency regulation and wind power reshaping are the highest value markets for storage. The coupling of wind generation and intermittent renewable generation resources with storage is discussed in Black and Strbac [15], Barton and Infield [17] [18], McDowall [16] and the EPRI Handbook [7], amongst others. This is a topic that we will discuss in the concluding section of this paper.

The literature that addresses the valuation of storage facilities includes Lu et al. [9], Deb [10], Schoenung et al. [11], Graves et al. [20] and EPRI [5]. In Lu [9], the authors present bidding strategies for pumped hydro facilities that derive additional value opportunities by simultaneously bidding into the spot market and committing the unit for synchronous reserve when it is in the pumping mode because it can readily reduce its pumping power and, consequently, reduce the overall system load. The analysis derives an optimal strategy through the use of Lagrange multipliers; however the optimization is carried out over an estimated weekly market clearing price curve, which is a deterministic solution.

Deb [10] echoes this concept by suggesting that the generator simultaneously bid into various markets to maximize profits. The paper presents several case studies that demonstrate the profit opportunities that exist for a pumped hydro facility when it participates in real time and multiple ancillary services markets. The paper also draws attention to the need for a structural model based on Monte Carlo simulation; however, the analysis that is presented is once again deterministic. Both of the papers present interesting ideas in terms of off-peak and spinning reserve profit opportunity.

EPRI [5] provides a valuation for storage by simply assuming a fixed on-peak and off-peak price of electricity that lasts for a specified period of time each day. This approach, although valid for small scale applications, ignores the volatility of prices in actual markets. Graves et al. [20] support the fact that using average peak and off-peak prices does not account for the variability in prices and thus leading to significant errors in the optimal strategy. They also discuss the use of a linear program for determining the optimal operation policy.

All of the models thus far in the literature present optimal strategies for storage operations that rely on deterministic prices. Where the volatility is specifically mentioned, the models once again optimize over a given historical price profile. The literature thus far present a rough estimate of the operational profits of a storage facility, yet none of them model what the plant would do in an actual market setting once it is constructed using forward looking, dynamic strategies.

This paper provides insight into such strategies and attempts to value storage, solely from an arbitrage perspective, using a profit optimizing operations strategy for a storage facility within an arbitrary power market. Although the scope of the valuation is limited to arbitrage, the models may be easily extended to participate in the ancillary markets for additional revenue opportunities. Schoenung et al. [11] aggregate the economic value of storage for various roles in current power systems from multiple reports and sources. The paper is an excellent reference for estimated dollar value ranges for each application as a function of the storage size and type.

3 Overview of Energy Storage Systems

Energy can be stored in potential, pneumatic, kinetic, electrochemical, electromagnetic, electrical, chemical or thermal form. The appropriate technology is very much dependent on the application due to the range in characteristics of the different technologies. For high-quality (non-thermal), large-scale energy storage, the most prevalent energy storage technologies are pumped hydro, compressed air energy storage, flywheels and batteries. In this section, we briefly introduce the various applications of electricity storage within power markets, and outline several value opportunities for such systems. We begin with a figure from the Electricity Storage Association summarizing the various storage technologies and their applications.

3.1 Applications of Energy Storage Systems

3.1.1 Load Leveling

The term “load leveling” refers to the balancing of night time troughs and afternoon peaks for a more level generation curve, making use of unused off-peak capacity while decreasing the need for peaking capacity. Thus a reduction in high-cost peak generation is met with an increase in the level of low-cost baseload generation, as depicted in Figure 2. In addition to the apparent economic advantage, load leveling also

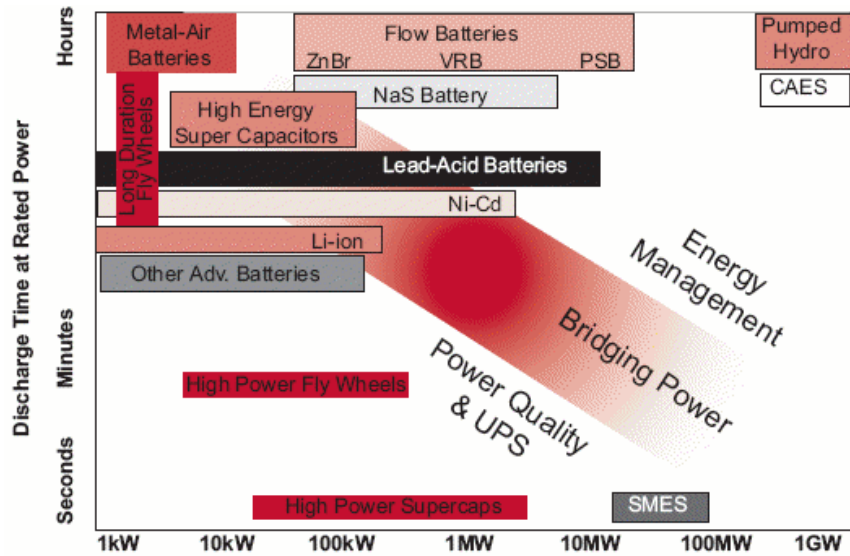



Figure 1: Summary chart of electricity storage technologies

increases generation flexibility, allowing for the use of technologies with greater fuel efficiency and lower emissions.



Reproduced from material developed by Anthony Price, Regenesys Ltd.

Figure 2: Demand profile with electricity storage

3.1.2 Arbitrage

As discussed earlier, as the load profile fluctuates throughout the day, so too does the spot price of electricity, consistent with economic principles for a constrained or limited resource. Arbitrage capitalizes on this economic opportunity through the purchase of inexpensive off-peak power that is sold during peak hours for

a net profit. Alternatively, a company who consumes peak power could use a storage system to purchase inexpensive off-peak power for their own use during peaking hours, thereby reducing their electricity costs.

3.1.3 Spinning Reserve

Spinning reserve is the primary subcategory of reserve power and generally refers to partially loaded plants that can ramp up (or down) to provide additional (or remove excess) generation capacity in order to meet sudden and unpredictable demands or contingencies. Spinning reserve plants have extremely low utilization and quick startup times, which are also characteristics of most types of energy storage systems, thus making energy storage systems a natural candidate for such ancillary services.

3.1.4 System Regulation

System regulation refers to the constant balancing of load fluctuations on the grid to account for power reliability, power quality and transmission and distribution problems. Power reliability refers to the frequency of extended power failures, during which the end-user is completely disconnected from utility power generation for more than about two minutes. Power quality refers to either a very short outage (zero voltage for less than two minutes) or to a non-zero voltage interruption, including voltage inconsistencies and harmonics. Energy storage can increase both power quality and reliability through its use as an uninterruptible power supply (UPS) system in case of power failure, whether caused in generation or transmission, and also through its use as an onsite power generator in remote or isolated distributed generation systems.

In areas where demand has increased beyond the capability of the T&D system, energy storage can store electricity during periods of inadequate transmission capacity to alleviate transmission and distribution overloads down-line of the congested area. Used in this manner, energy storage can defer necessary T&D upgrades for several years.

3.1.5 Storage with Intermittent Renewable Generation

Renewable generation technologies, such as solar and wind, are intermittent in nature and in general provide a stochastic generation curve dictated by the environmental conditions. A storage facility can be used to match generation with demand, and smooth the generation profile accordingly. With regard to utility load leveling, discussed above, energy storage allows constant generation to meet a variable demand; conversely

in this application, energy storage provides the capability for an intermittent (variable) power generator to serve a constant demand, or for a variable generator to meet a different variable demand. Combined with energy storage, renewable generation systems are more flexible and self-sustaining, allowing them to serve remote users and/or break-away sections of the grid.

Energy storage can also aid in transmission and distribution for renewable generation. Many renewable systems are located in rural areas where the electric grid is weak and unable to handle large amounts of energy. Energy storage can transform the intermittent large bursts of energy to a smaller, constant amount of energy that the grid can handle. This is one specific example of the potential transmission and distribution applications discussed in the previous subsection on system regulation.

3.2 Value Opportunities for Energy Storage Systems

Through its various applications, energy storage can be beneficial for all of the stakeholders within the power market structure. In this section of the paper, we examine only a few of these potential benefits. The obvious value opportunity for storage systems from a financial viewpoint is through arbitrage profits from time-differentiated pricing of electricity across peak and off-peak hours. We will briefly discuss the economic value of storage from three other perspectives: the value of storage as a peak generation technology, the value of storage as an ancillary service, and finally the value of storage in terms of the environmental benefits it provides.

3.2.1 Peak Generation

Aside from the time-differentiated profit opportunities, there are several additional benefits of storing power and regenerating it during peak hours when the system is deficient. The most significant impact is through the reduction in the need for peak generation from high marginal cost plants having large fuel costs, such as natural gas facilities, which result in higher average electricity rates for the end user. Fuel cost reductions occur as a result of load leveling, namely using storage during off peak hours to offset a portion of the generation requirements during peak hours.

3.2.2 Ancillary Markets and Reliability Services

Although it is difficult to evaluate the impact of storage services on reliability without taking into account concepts like value of lost load (VOLL), it is clear that any contribution towards system reliability, however marginal, will have a significant impact on social surplus. EPRI has estimated that the annual lost productivity due to short duration power quality events and service disruptions is at least \$53 billion per year, and the total losses from all adverse quality events accounts for \$119 billion per year [6]. Thus an increase in grid reliability, even by only a few percentage points at times when the system is already constrained, can result in a national social benefit on the order of billions of dollars. Once again, the reader is referred to Schoenung et al. [11] for estimated ranges of the value of ancillary services with the important caveat that storage systems are in general highly dependent on location, and the valuation of particular ancillary services will depend on the local market structure and both the need for and value placed on such reliability services.

3.2.3 Environmental Impacts

With increasing concerns regarding global climate change and CO_2 emissions, the movement toward clean technologies is gathering momentum. Unfortunately, many of these technologies, like nuclear, integrated gasification combined cycle (IGCC), and large hydro cannot easily be ramped to meet a variable load, limiting such technologies to baseload generation. Energy storage will become essential in enabling these non-emitting baseload generators to meet peak demand as well, thus significantly reducing CO_2 emissions in regions where baseload generation is primarily a low emission technology.

4 Modeling Electricity Storage

We will now motivate the valuation of a storage facility strictly through arbitrage arguments and discuss the optimal arbitrage policies under uncertainty. In the valuation framework, we seek an optimal operation policy in a power market once a storage plant has been constructed. The issues surrounding optimal sizing of plants and plant location are beyond the scope of this paper. We begin by presenting a mathematical abstraction of the storage system, and the parameters that govern an energy storage model, including all costs.

4.1 Electricity Storage Model

All energy storage systems can be modeled in the same fashion regardless of the form in which the system stores energy. A schematic generalizing the structure of this storage system model is presented in Figure 3, illustrating the basic components of power input, the storage medium, and power output.

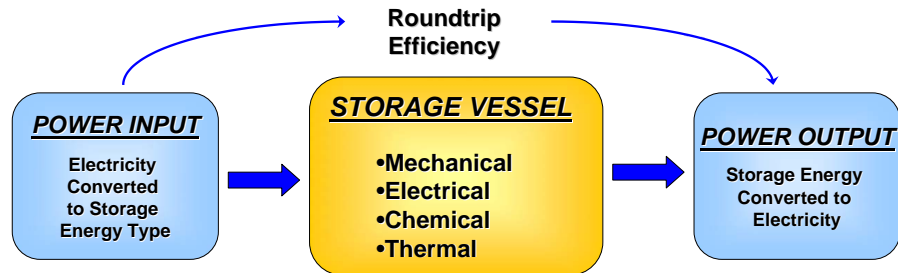


Figure 3: A generalized schematic of a storage system

Energy storage technologies are characterized by four distinguishing physical attributes:

- **Power Rating** [MW]: the maximum power output. The maximum power input is assumed to be equal to the maximum power output unless otherwise stated; if different, the maximum power input is the **Conversion Rate**.
- **Energy Capacity** [Joules or MWh]: the amount of energy that can be stored within the system. The amount of time a storage system can output at a given power is the Discharge Time and is directly proportional to the energy capacity for a constant power output.
- **Efficiency**: the ratio of energy discharged by the system to energy input into the system, and can be split into: **Conversion Efficiency** refers to losses experienced when converting the power input into a storage medium. **Storage Efficiency** refers to time-based losses during storage.
- **Reaction Time**: the time necessary for a system to “turn on” and begin charging or discharging or to switch between charging/discharging modes.

Using these physical characteristics, we can define the following set of variables:

Storage Constraints

- t period of time under consideration (hours)
- \bar{q}^D maximum quantity that can be sold (Discharged) in a single period (MWh)
- \bar{q}^R maximum quantity that can be bought (Recharged) in a single period (MWh)
- \bar{S} maximum storage capacity as determined by the system's energy capacity (MWh)
- γ_S storage efficiency (fraction of stored electricity maintained over one period)
- γ_C conversion efficiency (fraction of purchased electricity that gets stored)

4.2 Dynamic Model for Operations

We now assume that an energy storage facility exists and that the objective of the firm owning this facility is to maximize profits simply by exploiting arbitrage opportunities in the real-time electricity spot market. The decision of the plant operator in each period is to choose a quantity of energy that it will either purchase and put into storage, or discharge to the grid. The resulting decision variables are:

Decision Variables

- q_t^D quantity of electricity sold (Discharged) at time t (MWh)
- q_t^R quantity of electricity purchased (Recharged) at time t (MWh)
- S_t amount of electricity in storage at time t (MWh)

The amount of electricity in storage, S_t , is not a variable in and of itself since it is a direct result of the history of the recharge/discharge cycles. The basic operation of the plant may now be described by the storage dynamics:

Storage Dynamics

$$S_t \leq \gamma_S \cdot S_{t-1} + \gamma_C \cdot q_t^R - q_t^D \quad \forall t$$

Prior to establishing any financial models, several additional exogenous parameters must be defined:

Additional Parameters

- p_t price of electricity at time t
- r interest rate
- T number of time periods

The time parameter, t , in this model appears subtle, yet significant complexity arises with the discretization of time into distinct time periods, and decision periods (or stages). Each time period may be on the order

of minutes, hours, days or weeks, with the total number of periods summing to an hour, a day, a week or month, respectively. In general, electricity market prices are provided with an hourly increment, and due to the periodic nature of prices, a daily or weekly total decision horizon is typically sufficient. The estimation challenges will be further discussed when we later examine price processes.

With the time period specified, the next issue to be addressed is the length and frequency of decision stages. The time period represents the discretization of the samples of the prices, however each decision stage may be a duration of time that involves several time periods. With hourly prices, the decision stage may be hourly, whereby the operator makes a storage decision for each hour in a day, or it may be several hours, where a day is broken into say 3 or 4 stages of 8 or 6 hours each respectively. The decision would be to specify a storage policy over the duration of each stage.

The market structure typically restricts the decision stage; for example, in a fixed day ahead market, the storage facility would need to forecast its operation policy for the entire 24 hour period of the following day without any realizations of actual prices to update the policy. This issue will be further addressed when we discuss the optimal policies in a stochastic setting.

4.3 Arbitrage Model for Profits

We are interested in determining precisely how much profit can be earned by following an arbitrage strategy. The firm must choose how much power to buy or sell over time at the prevailing prices, subject to a variety of costs and constraints that are related to the physical aspects of the storage facility. The cost parameters are:

Storage Costs

$C_D(q_t^D)$ cost of discharging q_t^D units

$C_R(q_t^R)$ cost of recharging q_t^R units

$C_S(S_t)$ cost of storing S_t units for one period

Such costs arise when we consider the physical nature of the charging and generation costs, or costs associated with maintaining the energy within the storage facility. A simple example of this is the fuel generation costs of a CAES facility, where the plant must use a natural gas turbine to convert the compressed air back into electricity. Given these variables and parameters, we can now establish the storage plant's profit maximizing

objective function:

$$\sum_{t=1}^T [p_t(q_t^D - q_t^R) - C_D(q_t^D) - C_R(q_t^R) - C_S(S_t)] \cdot e^{-rt}$$

This objective as stated is nonlinear, since the cost functions have not been described and are unrestricted.

The problem is easily made linear if the cost functions are restricted to being affine with zero intercept (i.e. linear). The resultant linear program (LP) is:

$$\begin{aligned} & \text{maximize} && \sum_{t=1}^T [(p_t - C_D) \cdot q_t^D - (p_t + C_R) \cdot q_t^R - C_S \cdot S_t] \cdot e^{-rt} \\ & \text{subject to} && 0 \leq q_t^D \leq \bar{q}^D && \forall t \\ & && 0 \leq q_t^R \leq \bar{q}^R && \forall t \\ & && 0 \leq S_t \leq \bar{S} && \forall t \\ & && S_t \leq \gamma_S \cdot S_{t-1} + \gamma_C \cdot q_t^R - q_t^D && \forall t \end{aligned}$$

Here C_D , C_R and C_S represent the slopes of the respective cost functions. Since most available cost data is presented as a variable cost, this linear cost assumption is reasonable.

One of the main drawbacks of the linear programming model is that it assumes that all future prices are known with certainty, which is clearly unrealistic considering the volatility of electricity spot prices. In fact, it would be a significant mistake to use the LP model as a means of estimating the expected profits of a storage facility by looking back on realized prices since what is optimal in hindsight is not the same as what is optimal looking forward.

The LP model does, however, provide the absolute highest achievable profit which may then be used as a benchmark for other models, and may be used as a means of assessing the value of information for the uncertainties. To deal with the uncertainty in prices, we will introduce stochastic programming and dynamic programming models.

5 Stochastic Profit Models

The linear programming (LP) model has been introduced as a means of finding the optimal solution to a deterministic storage problem. For any particular price path, this solution is an upper bound of achievable profits given perfect information. Under particular market structures it may be possible to come close to achieving this upper bound in day ahead markets using LP's on very accurate forecasts. However, in general,

the real time price for the spot rate of electricity exhibits significant volatility and is therefore difficult, if not impossible, to predict accurately.

This motivates the use of probabilistic models for the derivation of arbitrage value, in particular the use of dynamic programming and stochastic programming. The LP approach can also be used in a stochastic setting by setting the prices equal to the expected values of the price realizations over time. This approach to the storage problem with uncertainty can be used to quickly find a lower bound on profits, but otherwise it is unsophisticated and performs poorly relative to the alternative techniques that we will now discuss.

5.1 The Dynamic Programming Model

In general, a dynamic programming (DP) approach to a problem discretizes the problem's state space and uses a backwards recursion to derive the optimal value and optimal policy at every state. The state variable describing the problem at a particular stage is defined in such a way that it completely describes the process. Given the state of the process at the beginning of a stage, we make a decision which transforms the process to the ending state at the end of the stage. The objective is to maximize (or minimize) the expected objective over all stages.

Dynamic programming algorithms are often limited in their application due to the "curse of dimensionality", which describes the rapid expansion in model size with the number of states. To maintain a tractable formulation, the dimension of the state space must stay small (typically no more than 3 or 4 dimensions) and the discretization of each dimension must be relatively coarse. This rules out the use of processes that track past histories (e.g. ARMA processes) and limits the ability of a DP to properly address continuous state and action spaces without great computational expense. For this reason, we restrict ourselves to a Markovian price process for the DP model.

Nevertheless, the solution to a DP is quite useful in that it provides an optimal contingency plan for every realizable state of a system, requiring only one pass through the DP to obtain the optimal policy for all possible realizations of the random process. The DP approach also allows for the consideration of many decision stages, since the problem size grows linearly in the number of stages. This makes dynamic programs particularly useful for solving problems in which many decisions must be made.

The storage problem involves determining the optimal profit maximizing policy depending on both the current storage level and the realized price throughout the hours of the day. Thus we are dealing with a

24-stage, 2-dimensional model with a state space that spans all possible storage levels and all possible spot prices. The action space is simply the amount of electricity to be charged or recharged and is subject to the standard capacity constraints. At each stage, every action generates an immediate profit or loss and determines the storage level that will be obtained in the subsequent stage. The expected profit of each action is therefore the sum of the immediate profit or loss and the expected value of the next state, which we can compute since we have already established the expected values for all future states. The optimal choice at each stage is clearly the action that yields the highest expected profit.

We will now apply this dynamic programming routine to the linear programming model. We begin by defining the parameters of the general Markov decision process (MDP) with specifications particular to the storage problem given in parentheses:

MDP Parameters

T	number of stages (each stage $t = 1, \dots, T$ is assumed to be one hour)
P	discrete space of realizable electricity prices ($p_t \in P \subseteq [10, 200]$)
S	discrete space of storage levels ($s_t \in S \subseteq [0, \bar{S}]$)
Ω	finite state space ($\omega_t \in \Omega = P \times S$)
X	finite action space ($x_t \in X \subseteq [\bar{q}^D, \bar{q}^R]$)
	where x_t is the quantity charged or discharged from storage in stage t
$Pr(\omega'_{t+1} \omega_t, x_t)$	conditional probabilities of state transitions
$\Pi(\omega_t, x_t)$	profit function for all $\omega_t \in \Omega$, $x_t \in X$

Within the DP, the profit function is defined as:

$$\Pi(\omega_t, x_t) = [p_t \cdot (q_t^D - q_t^R) - C_D \cdot q_t^D - C_R \cdot q_t^R - C_S \cdot s_t] \cdot e^{-rt} \quad \forall \omega_t \in \Omega, x_t \in X$$

where $q_t^D = -[x_t]^-$ and $q_t^R = [x_t]^+$.

The conditional transition probability distributions reflect the uncertainty in the MDP, and since the only uncertainty here is the price, we can represent these distributions as:

$$Pr(\omega'_{t+1}|\omega_t, x_t) = Pr((p'_{t+1}, s'_{t+1})|(p_t, s_t), x_t) = Pr(p'_{t+1}|p_t) \cdot I_{(s'_{t+1}=s_t-x_t)}$$

where I is a standard indicator function. That is, the only states we can transition to with positive probability are those with storage levels equal to the current storage level minus (plus) the amount currently discharged (recharged). The positive probabilities are given by the Markov price process and are dependent on both the current price and the current time.

The Bellman equation for the defined MDP becomes:

$$V_t(\omega_t) = \max_{x_t \in X} \left[\Pi(\omega_t, x_t) + \sum_{\omega'_{t+1} \in \Omega} Pr(\omega'_{t+1} | \omega_t, x_t) V_{t+1}(\omega'_{t+1}) \right]$$

When we use the same deterministic prices, we get the same storage policy and optimal profit as before, which is of course expected. However, this model allows the additional flexibility of considering uncertain prices through the use of the transition probability function, making the second term in the Bellman equation nothing more than the conditional expectation of the value in the next stage given the current state.

The drawbacks of this approach include the necessary discretizations of the price, storage and action spaces, and the fact that only basic (essentially Markov) price processes can be modeled. The discretizations are particularly problematic when storage and conversion inefficiencies are present in the model. When these parameters are present, it is frequently necessary to round to the nearest state. This encourages the algorithm to optimize over this rounding process, which is clearly undesirable. As evidence of the scope of this problem, situations have been observed in which the profits from the dynamic program exceed those of the linear program over the same price path, which is clearly erroneous since the LP provides an upper bound on profits.

5.2 The Stochastic Programming Model

The stochastic programming (SP) approach is quite different from the dynamic programming approach and can be used to solve an entirely different class of problems. Whereas the DP is useful for solving problems with multiple decision stages and limited state spaces, the SP performs best on problems with just a few decision stages (e.g. three or four), but can handle very large state spaces. This ability to handle large spaces allows the SP to deal with random processes that retain historical values. Thus we can use an ARMA or a non-Markovian stochastic process to model prices.

We implement several versions of the SP to solve the storage problem, including both a three-stage and a four-stage model. The three-stage model has the following formulation:

Three-stage SP Formulation:

$$\begin{aligned}
max_{q_t^D, q_t^R, S_t} \quad & \frac{1}{N} \sum_{i=1}^N \sum_{t=1}^8 [(\tilde{p}_t^i - C_D) \cdot q_t^D - (\tilde{p}_t^i + C_R) \cdot q_t^R - C_S \cdot S_t] \\
& + \frac{1}{MN} \sum_{i=1}^N \sum_{j=1}^M \sum_{t=9}^{16} [(\tilde{p}_t^{ij} - C_D) \cdot q_t^{D^i} - (\tilde{p}_t^{ij} + C_R) \cdot q_t^{R^i} - C_S \cdot S_t^i] \\
& + \frac{1}{LMN} \sum_{i=1}^N \sum_{j=1}^M \sum_{k=1}^L \sum_{t=17}^{24} [(\tilde{p}_t^{ijk} - C_D) \cdot q_t^{D^{ij}} - (\tilde{p}_t^{ijk} + C_R) \cdot q_t^{R^{ij}} - C_S \cdot S_t^{ij}]
\end{aligned}$$

subject to

$$\begin{aligned}
0 \leq q_t^D \leq \bar{q}^D, \quad 0 \leq q_t^R \leq \bar{q}^R, \quad 0 \leq S_t \leq \bar{S}, \quad S_t \leq \gamma_S \cdot S_{t-1} + \gamma_C \cdot q_t^R - q_t^D, \quad \forall t \\
0 \leq q_t^{D^i} \leq \bar{q}^D, \quad 0 \leq q_t^{R^i} \leq \bar{q}^R, \quad 0 \leq S_t^i \leq \bar{S}, \quad S_t^i \leq \gamma_S \cdot S_{t-1}^i + \gamma_C \cdot q_t^{R^i} - q_t^{D^i}, \quad \forall t, i \\
0 \leq q_t^{D^{ij}} \leq \bar{q}^D, \quad 0 \leq q_t^{R^{ij}} \leq \bar{q}^R, \quad 0 \leq S_t^{ij} \leq \bar{S}, \quad S_t^{ij} \leq \gamma_S \cdot S_{t-1}^{ij} + \gamma_C \cdot q_t^{R^{ij}} - q_t^{D^{ij}}, \quad \forall t, i, j
\end{aligned}$$

The models return a single optimal policy for the first eight or six hours for the three and four stage models respectively, and the optimal policies for the next stages are dependent on the simulated price paths. In order to implement these models in practice on actual observed price paths, we use the observed prices in previous stages and solve the reduced problem going forward. For example, in the three-stage problem, the optimal policy that is returned for the first eight hours is carried out. Then the actual prices observed during the first eight hours are used to generate correlated price paths for the two-stage problem going forward. Solving the two-stage problem gives the optimal policy for the hours 9 through 16, and then the prices observed over that period are used to simulate prices for the final one-stage problem. Figure 4 shows the tree structure of sample price paths that are generated, where for clarity we show an example with $L = M = N = 3$. The price process used in the simulations is the topic of the upcoming section.




Figure 4: A sample set of random price scenarios for a 3 stage model

6 Treatment of Electricity Prices

In this section we explore the characteristics of electricity prices, and the corresponding price models which may be used to capture these unique features. Several of the distinguishing features of electricity price processes include high volatility, large spikes, mean reversion to a daily pattern and seasonality.

There are a variety of reasons for the high volatilities, but the fundamental drivers include fuel prices, hydrological conditions, seasonal effects, weather dependence, new generators and transmission upgrades. The biggest drivers are the inelasticity of demand and the “hockey stick” nature of the supply curve. This fundamental nature of the supply versus demand creates an environment that is prone to large electricity price spikes.

Another reoccurring theme among most electricity price models is mean reversion. Typically, the electricity market exhibits extremely high mean reverting tendencies, where spikes generally last for only brief periods. An important caveat regarding this reversion property is that the mean value for electricity prices varies on an hourly basis in a quasi-predictable manner that is correlated with the demand. Hence a model that uses a constant mean value across all hours of a day will overestimate the actual volatility of prices, and will lose the trends that are necessary for temporal arbitrage.

Similar to the predictable pattern of price movements throughout a day, electricity prices fluctuate on a seasonal basis as well according to the local demand profile. This seasonality must be taken into account when modeling the price process in order to achieve proper estimates of the mean and the variance of the hourly prices for each month. Failure to account for the seasonality will affect the true mean of the hourly prices and will introduce an overestimation of the variance of the prices.

6.1 Mathematical Models of Electricity Prices

In light of the characteristics of electricity described thus far, we will introduce several methods for representing the electricity price process. We begin with a very simple process using empirical values, and move from there to more complex financial engineering methodologies.

6.1.1 Correlated Random Variables with Gaussian Noise

We begin by describing a very simple Gaussian noise price process. Due to the predictable patterns that electricity prices exhibit throughout a day, the mean varies on an hourly basis, and so our mean value, μ_t , exhibits a time dependency. In its most trivial form such a process may be simulated by:

$$\vec{P} = E(\vec{P}) + \sigma \cdot \vec{\epsilon}$$

where \vec{P} is a 24x1 vector of hourly prices, $E(\vec{P})$ is the vector of the hourly mean prices, σ is a diagonal matrix of discrete hourly volatility, and $\vec{\epsilon}$ is a 24x1 vector of realizations of a standard normal random variable.

In order to alleviate the independence assumption of hourly prices, the cross correlations of hourly prices must be incorporated into the volatility term, σ . Using a Cholesky decomposition, the positive definite covariance matrix of the original hourly prices is decomposed into $C = LL'$, where L is a lower triangular matrix with positive diagonal elements. Thus the price process may now be simulated using:

$$\vec{P} = E(\vec{P}) + L \cdot \vec{\epsilon}$$

It can be shown that simulation of the price process using this approach preserves all of the statistical properties of the original distribution of prices.

6.1.2 Advanced Financial Engineering Techniques using Brownian Motion

The empirical framework may be generalized to a stochastic process that describes the price as a Brownian motion process. Although this methodology has its shortcomings for the modeling of electricity price dynamics as described above, it is a good general starting point to build from. With this assumption in mind, the general framework for the modeling of electricity spot rates will be consistent with models found in interest rate derivatives.

We begin the modeling framework with the simple case of single factor equilibrium models. In the world of financial assets, we refer to the standard geometric Brownian motion (GBM) model of short rates, otherwise known as the Rendleman and Bartter model of the risk neutral process for the short rate. In this case the spot price P follows:

$$dP = \mu P dt + \sigma P dz$$

where dP is the change in price over the time interval dt , μ is the expected drift of the price over this interval, σ is the volatility of the price, and dz is a standard Brownian motion.

Commodity price structures often exhibit strong mean reversion, as is the case with electricity prices. We introduce mean reversion to the Brownian motion model to arrive at the classic Vasicek model:

$$dP = a(\mu - P)dt + \sigma dz$$

where a is the mean reversion coefficient. A more generalized mean reversion model using time dependent mean values, $\mu(t)$, has been introduced by Hull and White [22]:

$$dP = a \left(\frac{\mu(t)}{a} - P \right) dt + \sigma dz$$

The Hull-White model is not a GBM model, and so a simple modification leads to the GBM mean reversion model:

$$dP = a(\mu(t) - \ln P) P dt + \sigma P dz$$

where a is the mean reversion rate, e^μ is the long term level to which P reverts, and σ is the constant volatility parameter.

The models introduced thus far do not account for the spikes that are evident in electricity prices. For this, we may refer to the jump diffusion framework first presented by Merton, and later extended to the electricity price process framework by Deng [23]:

$$\frac{dP}{P} = a(\mu - \ln P)dt + \sigma(t)dz + \kappa P dq$$

The downfall to this model is the difficulty in characterizing the required parameters. In effect, we have four static parameters to calculate, in addition to two parameters, namely the volatility and the drift, which are functions of time and underlying price. If we further assume that the volatility is not a deterministic function, but that it is in fact stochastic in nature, then the complexity of the model in itself becomes far too difficult to manage, or parameterize in a meaningful manner. For this reason we will limit our model to the mean reverting model based on the Hull-White approach. For additional insight into the unique nature of electricity pricing models, the interested reader is referred to [24] and [25].

7 Profit Maximization Through Arbitrage Models

Our analysis of arbitrage has been focused, thus far, on the development of the necessary models for the stochastic price process, and the physical operation of storage facility, given that it is already available. We will now calibrate the generalized models and test the stochastic models under different price processes and compare the resulting profit with the best case control profits (perfect information).

7.1 Storage Model Calibration

The analysis will contrast two storage technologies, compressed air energy storage (CAES) and the Sodium Sulfur (NaS) battery. Of the available storage technologies, pumped hydro and CAES are the two most suitable for exploiting arbitrage opportunities; however, in the United States, there are very few suitable geographic locations remaining for pumped hydro, if any. On the other hand, the Electric Power Research Institute estimates that more than 85% of the U.S. has geological characteristics that will accommodate an underground CAES reservoir [5], making CAES the logical choice for assessment.

Smaller scale storage technologies will be limited to the NaS battery, since storage losses and capital costs of the other technologies render them virtually useless in such applications. We refer the reader back to Figure 1 contrasting the various storage technologies. We begin the analysis by introducing the physical constraints and parameters that are consistent with our two target storage facilities.

Characteristic	NaS Battery	CAES
Power recharge capacity	10 MW	300 MW
Power discharge capacity	10 MW	300 MW
Energy capacity	100 MWh	3000 MWh
Storage efficiency	0.97	1.00
Conversion efficiency	0.85	0.85
Discharge Cost	0	$\$3/\text{MWh} + \$7/\text{MMBtu} \times 4.1 \frac{\text{MMBtu}}{\text{MWh}}$ $= \$31.7/\text{MWh}$
Recharge Cost	0	0
Storage Cost	0	0

7.2 Calibration of Price Processes

To develop policies from our calibrated models, two price processes will be introduced, and the optimization programs will be assessed for each of the developed price routines. The data used for the development of the price processes is hourly marginal locational prices obtained from the New England ISO power market.

7.2.1 Correlated Brownian Motion

The first price process will be based on correlated Gaussian noise using the Cholesky decomposition of the covariance matrix of hourly price data as described in the previous section. This process will be the basis of the main SP model and the Gaussian noise process. In each stage, the random prices are generated using the knowledge of all previous prices; thus, all realizations of hourly prices have the desired mean, variance and correlation characteristics.

7.2.2 Mean-Reverting Markov Stochastic Process

In order to do away with the complete path dependence of the previous method, we generalize the model to a Markovian stochastic process which may easily be used in a DP framework. The chosen mean reverting process is described by a simple modification to the standard Hull-White model. The discrete time representation of this model is:

$$\Delta P = a(\mu_t - P)\Delta t + \sigma\epsilon\sqrt{\Delta t}$$

where ϵ is a standard normal random variable.

Here, the drift terms, μ_t , are time dependent parameters that vary according to the hour of the day. This enables the described Markov process to be free from serial correlations with the previous realizations of the hourly prices, while maintaining the familiar daily trends that are prevalent in electricity market prices.

This price process has been calibrated through a maximum likelihood estimation approach using the same NE-ISO price data. The details of the formulation of the MLE equations and their results will not be provided here, however it is important to note that a significant issue with this approach is the estimation of a large number of parameters (24 hourly mean values, the volatility, and the mean reversion coefficient).

A second version of the stochastic program has been implemented using this mean-reverting Markov process to model prices. This version is generally less accurate than the correlated Gaussian noise approach, since it sacrifices the ability to correlate current prices with historical prices. However, this approach has the advantage that it can sensibly be compared with the dynamic program, since the DP uses this same Markov process.

7.3 Model Results of Arbitrage Profits

7.3.1 LP Optimal Storage Policies

In order to examine the effect of both the technological parameters as well as the hourly price path, we begin by examining the policies dictated by the linear program on a deterministic price path in Figure 5. The significance of this result lies in the contrast between the two policies given storage efficiency. Recall that the CAES system is assumed to have loss-less storage over time, whereas the storage efficiency of the NaS battery is assumed to be 0.97 [5]. Consequently this “lossy” storage system cannot take full advantage of the low prices during off-peak periods since its time dependent losses will negate any profit opportunity by the time the prices increase.




Figure 5: Optimal storage policies for the CAES and NaS battery systems

We now examine the optimal deterministic LP storage policies under a set of three sample price paths denoted by path (a), (b) and (c). The paths, depicted in Figure 6, corresponds to three typical price scenarios:

- **Price Path (a)** - portrays a typical price path throughout a day with the expected low off-peak prices and a significant rise in prices during peak periods.
- **Price Path (b)** - portrays an unusual day in the electricity markets where prices remain relatively low for the duration of the day.
- **Price Path (c)** - portrays the standard low off-peak prices and higher peak-prices; however there is no significant peak in prices. This two-tiered price path is typical to the standard model of price paths used in most power market analyses, including the storage valuation study conducted by EPRI [5]. We will shortly demonstrate that this price path has a tendency of undervaluing the arbitrage profit opportunities.

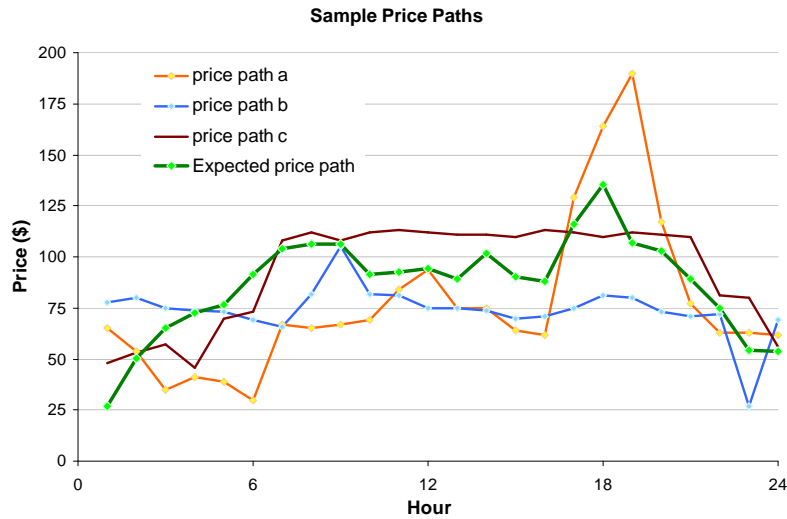


Figure 6: The three different electricity price paths

The second series of plots in Figure 7 provide the optimal deterministic LP policy for a CAES system for each of the three described paths. The resultant arbitrage profits under these scenarios for the CAES facility are \$89,354, \$1,726 and \$22,985 for paths (a), (b) and (c) respectively. It is clear that if prices follow a typical path having low prices and a marked price peak (such as in path (a)), then the opportunities for profits are significant. Moreover, such a scenario may result in profits that are four times greater than a standard two-tiered price approach.

7.3.2 SP and DP Optimal Storage Policies

The same three price paths may be used to give us an interesting perspective on the operation of the optimal policy from the stochastic program. Figure 8 provides a comparison of the optimal policies from a 3 stage and a 4 stage stochastic program. Note that the policy is the same irrespective of the price path within the first stage (the first 8 hours for the 3 stage SP, and the first 6 hours for the 4 stage SP).

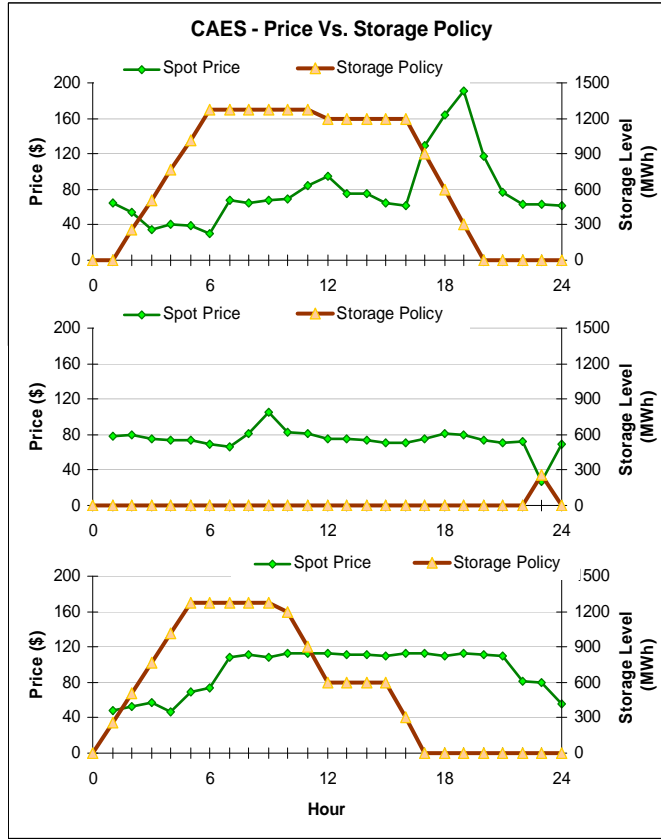


Figure 7: Optimal deterministic LP storage policies for the CAES system under different price paths. From top to bottom: path (a), path (b) and path (c)

Finally, we have the dynamic programming solution for the CAES system under the three different paths (Figure 9). The DP formulation has the following characteristics:

Dynamic Program Parameters

- T stages in hourly increments (24 hours)
- P discretized hourly prices in \$1 increments, ($P \subseteq [10, 200] \Rightarrow 191$ Prices)
- S discretized storage levels:
 - CAES: discretized in 10MWh increments ($S \subseteq [0, 3000] \Rightarrow 301$ Levels)
 - NaS Battery: discretized in 0.5MWh increments ($S \subseteq [0, 100] \Rightarrow 201$ Levels)
- X finite action space of the quantity charged/discharged (MWh)
 - CAES: discretized in 10MWh increments ($X \subseteq [-300, 300] \Rightarrow 61$ Levels)
 - NaS Battery: discretized in 0.5MWh increments ($X \subseteq [-10, 10] \Rightarrow 21$ Levels)
- $Pr(\omega'_{t+1}|\omega_t, x_t)$ conditional probabilities of state transitions (191x191x24 matrix)




Figure 8: Optimal SP storage policies for the CAES system under the different price paths

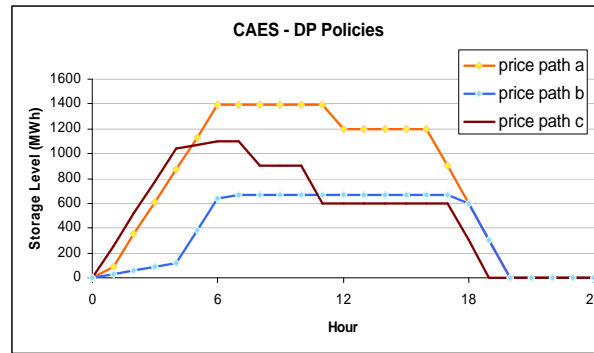


Figure 9: Optimal DP storage policies for the CAES system under the different price paths

Although the policies here are not as extreme as the ones given by the deterministic solution, the DP does provide notably different policies for the different price paths. This is a result of the more frequent decision stages that the DP uses in determining the optimal policy. For example, in the case of price path (b), the DP solution takes a very modest position in the storage level, and does not purchase a large quantity of electricity during off peak periods since it repeatedly updates its state and realizes that the price is not following the standard trend. As a result, the profit levels for the three paths under the different optimization policies reflects the significance of the additional decision stages. A comparison of the policies and their resulting profits for each price path is provided in Figure 10.

Although the previous analysis helps in observing the behavior of the different policies under quite distinct price paths, it is necessary to carry out the policies in a general setting in order to determine the expected arbitrage profits. To do this, we have determined the expected profits under five different policies:

- (i) The linear program on the expectation of the price path
- (ii) The 3 stage SP using a sample of 125,000 scenarios using the Gaussian noise price process

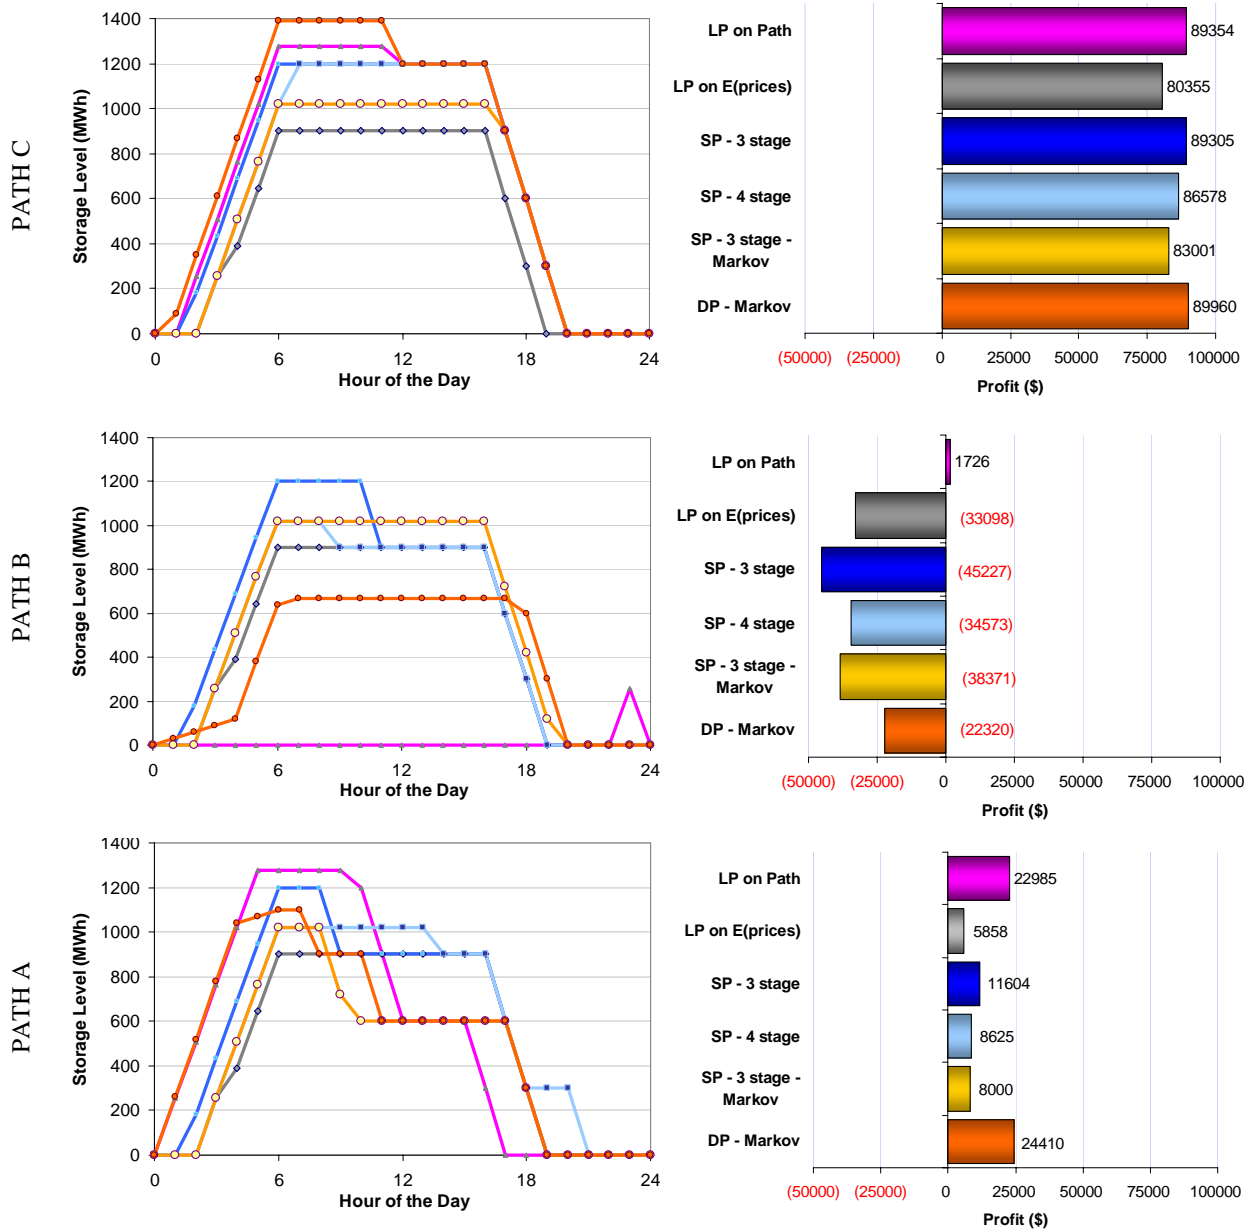


Figure 10: Optimal storage policies for the CAES system using the different optimization schemes for price paths (a), (b) and (c)

- (iii) The 4 stage SP using a sample of 160,000 scenarios using the Gaussian noise price process
- (iv) The 3 stage SP using a sample of 125,000 scenarios using the mean-reverting price process
- (v) The dynamic program using the mean-reverting price process

The models have been implemented using the parameters $L = M = N = 50$ and $K = L = M = N = 20$ for the 3 and 4 stage models respectively. With these values, the problems could be solved in about thirty seconds on a standard laptop. Models (ii) and (iii) demonstrate the effect of introducing an additional decision stage, while the purpose of model (iv) is to provide a comparison between the SP and the DP policies under a single price process, in order to prevent any further discrepancies between the models. The arbitrage profits from these models for the CAES and NaS battery systems are summarized in Figure 11.




Figure 11: Expected daily arbitrage profits for the CAES and NaS battery systems using the different optimization schemes

7.4 Sensitivity Analysis

We now examine the sensitivity of the profits to the model parameters. For the sake of brevity, we provide only the sensitivity observations of significance. We contrast both the stochastic solution in addition to the linear program over the expected prices in each case, and observe that the stochastic program consistently dominates the simple expected LP solution.

We begin by varying the physical parameters of the CAES system, namely the conversion efficiency, the maximum discharge rate (the generator size) and the energy capacity (storage size). As expected, we have relatively linear increases with increases in both the recharge/discharge size and conversion efficiency. The critical observation here is that there is a finite useful size for the storage itself. Here we can see that

for the 300MW CAES facility, approximately 3 hours of storage is sufficient to extract all arbitrage profit opportunities. The bottleneck is the generator size, where an increase in the discharge/recharge rate provides greater profit opportunities than a larger storage tank.

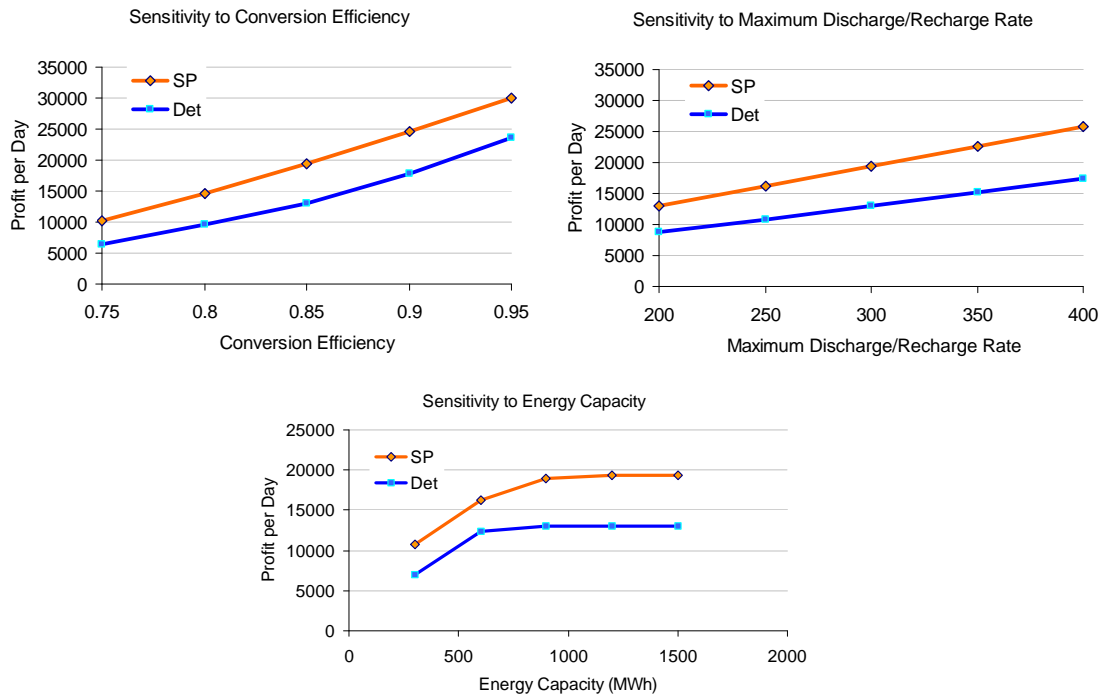


Figure 12: Sensitivity of arbitrage profits to physical parameters of the CAES system

The profit opportunities of the CAES system are also highly susceptible to the prevailing fuel costs. As shown in the sensitivity graph, a 30% increase in the discharge cost leads to a 50% decrease in arbitrage profits. Thus fears of an impending rise in fuel costs may have a detrimental effect on the arbitrage opportunities of CAES facilities.

Finally, we examine the effects of storage efficiency on arbitrage profits. Beginning with the base case of 97% efficiency for the NaS battery [5], we note that there are significant incentives in decreasing the time dependent losses of storage systems, since a modest increase in storage efficiency yields gains of 50% or more in arbitrage profits.

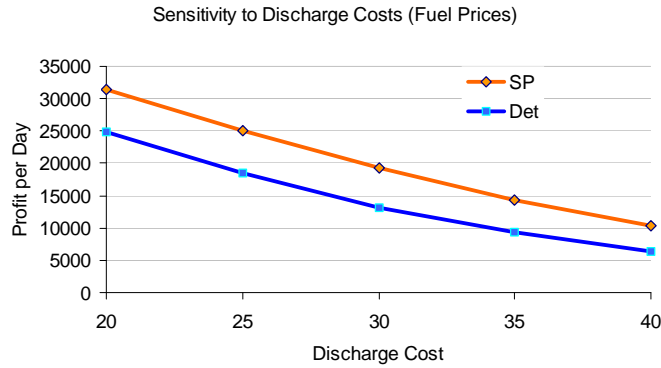


Figure 13: Sensitivity of arbitrage profits to discharge costs of the CAES system




Figure 14: Sensitivity of arbitrage profits to storage efficiency of the NaS battery system

8 Summary and Conclusions

8.1 Cost Benefit Analysis

We conclude this paper by first summarizing the cost-benefit analysis of storage facilities using the results of our arbitrage profit models. For the two storage facilities under consideration, the cost breakdown is:

Cost Component	NaS Battery	CAES
Power Conversion Cost	\$2270/kW	\$450/kW
Storage Cost	\$200/kWh	\$1/kWh
Total Cost	\$43 million	\$138 million
O&M Fixed Cost	\$50/kWyear = \$500,000/year	\$13/kWyear = \$3.9 million/year
Disposal Cost	\$43/kW = \$430,000	0
Life	15 years	30 years

Using these costs we can summarize the profits as:

	Model	Daily Profit	Annual Profit (\$M)	Fixed Costs (\$M)	Net Profit (\$M)
CAES	SP	\$22,500	\$5.63	\$3.90	\$1.73
	DP	\$37,280	\$9.32	\$3.90	\$5.42
NaS Battery	SP	\$1,147	\$0.287	\$0.500	-\$0.213
	DP	\$1,625	\$0.406	\$0.500	-\$0.093

The NaS Battery has a negative net profit, and so is not a feasible solution for arbitrage profit schemes. The CAES facility yields positive net profits; however, when one considers the roughly \$138 million price tag for the construction of such a plant, the finances do not bode well. Nevertheless, we have only considered arbitrage profits alone in this valuation. When coupled with additional revenue and value sources, such as participation in ancillary service markets and reliability services, the CAES facility has the potential of becoming economically viable.

8.2 General Results

The primary difference in our models (LP vs. DP vs. SP) stems from the number of decision stages used in the optimization. The dynamic program outperformed all models due to the fact that the optimization was carried out over 24 stages. Although, the dynamic programming approach is amenable to multi-stage decision problems (eg. 24 hours), and provides an exact solution without simulation requirements, the state-space problem limits the choice of pricing models, and further limitations on the problem size within each dimension make the DP susceptible to discretization errors.

The stochastic programming approach provides flexibility in the choice of price models, allowing for a richer modeling framework for representing such stochastic processes. This flexibility extends to the alleviation of discretization errors within the process. The framework allows for the ability to deal with a large number of uncertainties, coming at the expense of a limited number of decision stages. Moreover, the structure of the problem within specified stages is often appropriate for power markets where participation in the market often occurs hours in advance, rendering most dynamic programming solutions virtually useless.

The selection of a single day (24 hour period) for the optimization has proven to be an accurate reflection of the overall system dynamics. A separate analysis was carried out using an optimization over five separate

days vs. a single consecutive five-day period. Although Lu et al. [9] claim that a horizon of one day is too short to consider the optimal operation strategy, in particular for pumped-storage units, our results showed little difference, providing evidence that a daily approach is a sensible time frame. There is, however, a significant difference when considering the starting point within the 24 hour period. A final test was carried out using different endpoints for the 24-hour interval for the optimization problem, and as expected the results displayed significant discrepancies. The standard practice of starting at midnight was shown to be optimal, once again not affecting the results of the analysis provided in this paper.

Finally, sensitivity analysis results have shown that 5 hours of storage capacity is more than sufficient to capture the majority of arbitrage profits, a significant decrease from the previously claimed lower limit of 20 hours presented in [20]. The limitation in most storage systems from an arbitrage perspective is the discharge/recharge rate, namely the size of the generator. When dealing with storage plants in general, the storage efficiency is a very important parameter to consider for daily arbitrage applications; a result that is once again in direct opposition to the claim that the value of storage is insensitive to the storage efficiency presented in [15]. It has been shown, via the analysis of the NaS battery, that even a small storage loss will almost eliminate all arbitrage profit opportunities.

The framework presented in this paper is amenable to the valuation of storage using multiple revenue sources as suggested in the literature. Furthermore the models have been extended to value an intermittent generation facility (such as wind) coupled with a storage facility. The stochastic nature of not only the spot price, but also the wind generation level is readily modeled and a robust valuation framework has been developed and will be the subject of future work.

9 Acknowledgements

The authors would like to thank Kelsey Lynn, Jim Sweeney and Gerd Infanger for their valuable insights, comments and suggestions. The authors would also like to acknowledge EPRI-EDF and the Global Climate and Energy Project for their support of this research.

References

- [1] “Annual Energy Outlook 2006 with Projections to 2030”, Energy Information Administration Report: DOE/EIA-0383, February 2006
- [2] “Electricity Technology Roadmap: 2003 Summary and Synthesis”, EPRI, 2003
- [3] Lynn, K.B., P. Mokrian and M. Stephen, “Economic Overview of Energy Storage in Modern Power Systems”, Stanford University Report, December 2005.
- [4] “Wind Power Impacts on Electric Power System Operating Costs: Summary and Perspective on Work to Date”, NREL, March 2004
- [5] “Handbook of Energy Storage for Transmission or Distribution Applications”, EPRI, Dec 2002
- [6] Rastler, Dan. “Accelerate Innovation on Development of Break-through Energy Storage Systems - for Electric Utility Markets.” Presentation with Electric Power Research Institute, Palo Alto Research Center, 10 Nov. 2005
- [7] “Energy Storage for Grid Connected Wind Generation Applications”, EPRI-DOE handbook Supplement, December 2004
- [8] “Electric Accumulators”, Science, Vol. 14, No. 354 pp 325-327, Nov. 15, 1889
- [9] Lu, N., J. H. Chow, and A. A. Desrochers, “Pumped-Storage Hydro-Turbine Bidding Strategies in a Competitive Electricity Market”, IEEE Transactions on Power Systems, Vol. 19, No. 2, May 2004
- [10] Deb, R., “Operating Hydroelectric Plants and Pumped Storage Units in a Competitive Environment”, The Electricity Journal, pp 24-32, April 2000
- [11] Schoenung, S.M., J.M. Eyer, J.J. Iannucci, S.A. Horgan, “Energy Storage for a Competitive Power Market”, Annual Reviews, Energy Environ. 1996.
- [12] Schaber, C., P. Mazza, and R. Hammerschlag, “Utility-Scale Storage of Renewable Energy”, The Electricity Journal, pp 21-29, July 2004
- [13] Harty, F. R., F. Depenbrock, P. W. Ward and D. L. Sackett, “Options in Energy Storage Technologies”, The Electricity Journal, pp 21-26, July 1994
- [14] Dell, R.M. and D.A. Rand, “Energy storage - a key technology for global energy sustainability”, Journal of Power Sources, 100 2-17, 2001

- [15] Black, M. and G. Strbac, "Value of storage in providing balancing services for electricity generation systems with high wind penetration", *Journal of Power Sources*, 2005
- [16] McDowall, J. "Integrating energy storage with wind power in weak electricity grids", *Journal of Power Sources*, 2005
- [17] Barton, J.P. and D. Infield, "A probabilistic method for calculating the usefulness of a store with finite energy capacity for smoothing electricity generation from wind and solar power", *Journal of Power Sources*, 2005
- [18] Barton, J.P. and D. Infield, "Energy Storage and Its Use With Intermittent Renewable Energy", *IEEE Transactions on Energy Conversion*, Vol. 19, No. 2, June 2004
- [19] Makansi, J., "Energy Storage: Value is Clear, But Who Will Pay?", *Power Engineering*, 108, 8; *ABI/INFORM Global*, Aug 2004
- [20] Graves, F., T. Jenkin, and D. Murphy, "Opportunities for Electricity Storage in Deregulating Markets", *The Electricity Journal*, pp 46-56, October 1999
- [21] <http://www.eia.doe.gov/oiaf/servicerpt/derivative/tbl3.html>
- [22] Hull, J. and A. White, "Pricing interest-rate derivative securities", *The Review of Financial Studies*, Vol 3, No. 4, pp. 573-592, 1990.
- [23] Deng, S. "Stochastic Models of Energy Commodity Prices and Their Applications: Mean-reversion with Jumps and Spikes," Working Paper, Georgia Institute of Technology, 1999
- [24] "Derivatives and Risk Management in the Petroleum, Natural Gas, and Electricity Industries", *Energy Information Administration*, October 2002
- [25] Eydeland, A., and K. Wolyniec, "Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging", *Wiley*, 2002