Estimating the Value of Electricity Storage in PJM: Arbitrage and Some Welfare Effects☆

Ramteen Sioshansi[a,*], Paul Denholm[b], Thomas Jenkin[b], Jurgen Weiss[c]

[a]*Industrial, Welding, and Systems Engineering Department, The Ohio State University, United States*
[b]*National Renewable Energy Laboratory, United States*
[c]*Point Carbon North America, United States*

Abstract

Significant increases in prices and price volatility of natural gas and electricity have raised interest in the potential economic opportunities for electricity storage. In this paper, we analyze the arbitrage value of a price-taking storage device in PJM during the six-year period from 2002 to 2007, to understand the impact of fuel prices, transmission constraints, efficiency, storage capacity, and fuel mix. The impact of load-shifting for larger amounts of storage, where reductions in arbitrage are offset by shifts in consumer and producer surplus as well as increases in social welfare from a variety of sources, is also considered.

Key words: Energy storage, arbitrage, social welfare
JEL: D60, D62, Q41, Q42

1. Introduction

The emergence of wholesale electricity markets in many regions of the United States, together with significant increases in prices and price volatility of natural gas and electricity, have raised the interest in and potential economic opportunities for electricity storage plants. Storage can take advantage of the differences in hourly off- and on-peak electricity prices by buying and storing electricity at times when prices are low, and then selling it back to the grid when the price of energy is greater. Storage also can provide capacity and ancillary services (such as spinning and non-spinning reserves or frequency regulation) as an alternative or complement to energy arbitrage. Large-scale deployment of energy storage, which smoothes the load pattern by lowering on-peak and increasing off-peak loads, will result in a similar smoothing of the price pattern and reduce arbitrage opportunities. Despite its effect of reducing the value of arbitrage, this load smoothing by larger-scale storage can have significant external welfare effects.

In this paper, we analyze four aspects of the economic value of electricity storage deployed in the PJM region.[1] First, in section 2 we examine the basic relationship among storage efficiency, storage energy capacity, and the arbitrage value of energy storage. Second, in section 3 we evaluate the accuracy of theoretical energy storage dispatch and the value of arbitrage using perfect foresight compared to a 'real' value capture that considers the uncertainty of future electricity prices. Third, in section 4 we evaluate the regional and temporal variation in the value of energy arbitrage, examining the impact of transmission constraints, natural gas price variations, and fuel mixes on energy

☆The authors would like to thank seminar participants at the U.S. Department of Energy, several colleagues, and an anonymous referee for helpful comments and suggestions. This work was supported by the U.S. Department of Energy under Contract No. DE-AC36-99GO10337 with the National Renewable Energy Laboratory. Any opinions and conclusions expressed in this paper are those of the authors and do not necessarily represent those of the U.S. Department of Energy or the National Renewable Energy Laboratory.

*Corresponding author
Email addresses: sioshansi.1@osu.edu (Ramteen Sioshansi)

[1]Parts of the work in this paper extend some of the analysis and ideas in Jenkin and Weiss (2005).

storage economics. Finally, in section 5 we consider the impact of larger storage devices, examining how the use of energy storage can decrease on-peak and increase off-peak hourly prices diminishing the value of arbitrage, while generating welfare effects for consumers and generators. We also examine the potential for energy storage to help insulate consumers from energy price spikes. While the focus of this work is related to energy arbitrage, energy storage can provide additional societal benefits including improved use of existing generation and transmission and distribution (T&D) assets, benefits from deferred investment in new generation capacity and T&D, and helping to integrate renewable energy resources.

2. Arbitrage Value of Small Amounts of Electricity Storage in PJM: The Impact of Hours of Storage and Efficiency

One of the best understood and studied applications of energy storage is the use of 'small device' energy arbitrage—the ability to buy low and sell high, where the device is assumed to be small enough that its charges and discharges do not affect the price of electricity. This type of 'price-taker' analysis often assumes perfect optimization of small devices facing known prices. Examples of this type of analysis that have been applied to wholesale electricity markets include Graves, Jenkin, and Murphy (1999), Walawalkar, Apt, and Mancini (2007), and Figueiredo, Flynn, and Cabral (2006). Other recent studies of electricity storage that cover a broader range of applications include Eyer, Iannucci, and Corey (2004) and EPRI (2003).

A storage device captures arbitrage value by storing low-cost energy and then reselling that energy during higher-priced hours. A storage device is characterized by its power capacity (MW), its energy capacity (MWh), and roundtrip efficiency. The energy capacity of a storage device may also be rated by the number of hours of full power output, which is the convention used in this paper. Some storage devices have energy capacities of less than an hour, such as flywheels and batteries designed primarily for ancillary services such as frequency regulation or spinning reserves. Larger devices used for energy arbitrage such as pumped hydroelectric storage (PHS), compressed-air energy storage, or certain large batteries may store enough energy to accommodate a full day's peak demand period of eight hours; and, in some cases, have been built with more than 20 hours of discharge capacity.[2]

We first estimated the historical annual value of arbitrage for a small storage device in PJM from 2002 to 2007. The PJM Interconnection is a regional transmission organization serving about 51 million people in the eastern U.S. with a 2007 peak demand of about 139 GW. PJM operates a series of centralized multi-settlement markets for energy, ancillary services, and capacity on a day-ahead and real-time basis. For each year, the operation of the storage device was optimized to maximize arbitrage profits against hourly load-weighted average marginal energy price data obtained from PJM. The optimization was conducted two weeks at a time, assuming perfect foresight of future hourly electricity prices during each two-week period. This use of a two-week optimization horizon allows for both intra- and inter-day arbitrage opportunities, including greater charging during weekends, because hourly electricity prices often tend to be lower than during the week. Optimizing over a two-week period also reflects the fact that a storage operator would not be realistically expected to make dispatch decisions in anticipation of prices many weeks in the future. To ensure energy stored in the device at the end of each two-week period has 'carryover value,' each optimization was done with a 15-day planning horizon to determine the dispatch of each two-week period. Otherwise, the operator would fully discharge the device by the end of each two-week period, which would not reflect actual device operation. Because of the price-taking assumption, the model is a linear program which we formulate

[2]See Denholm and Kulcinski (2004) for further discussion of this topic.

in GAMS 21.7 and solve using cplex 9.0. Appendix A.1 discusses the formulation of our model in greater detail.

We assumed for these initial calculations an 80% roundtrip efficiency, which is at the upper range of actual storage devices currently available (such as the Bath County PHS plant in PJM, which has an 80.3% efficiency ASCE (1993)). We discuss the sensitivity of our results to storage efficiency in more detail later in this section. As a result 10 hours of charging is required for each eight hours of discharging.[3] The value of arbitrage for each two-week period was then summed over the year to provide annual values for the value of arbitrage in a $/kW-year basis. Figure 1 shows hourly energy prices during a sample one-week period in 2006 and the optimal hourly operation of a 12-hour storage device.[4] As expected, the dispatch pattern follows the prices with energy stored when prices are low and sold when prices are higher. For a price-taking device, if the device is dispatched to charge or discharge in a given hour, it is optimal to dispatch it to the lesser of its power or available energy capacity. It is important to note that hourly charge and discharge patterns in different days are similar but not the same, showing that inter-day and weekend effects matter.[5] Figure 2 shows the historical value of storage as a function of hours of storage. The value of arbitrage in PJM has varied significantly during this time period ranging from about $60/kW-year in 2002 (for a 12-hour device) to more than $110/kW-year in 2005. We discuss the reasons for these observed changes in value in section 4.

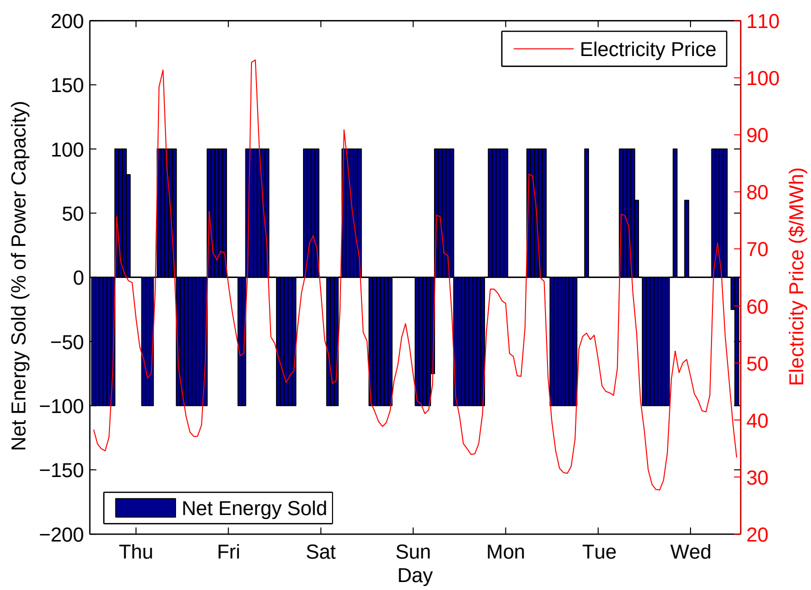

Figure 1: Electricity prices and the optimal hourly dispatch of storage device during one week

Figures 3 and 4 provide additional insight into the relationship between storage value and size. Figure 3 illustrates that most of the arbitrage value in storage comes from intra-day arbitrage, with more than 50% of the total capturable value derived from the first four hours of storage.[6] Additional value is provided by longer-term storage, including the ability to perform inter-day arbitrage as well as charge more during the weekend and discharge during the following week; 8 hours of storage captures about 85% of the potential value, while 20 hours of storage captures about 95%

[3] We assume throughout our analysis that storage devices have the same input and output power capacity.

[4] For reasons of clarity figure 1 only shows prices and dispatch for one week, although the optimization is done with a two-week horizon as described above.

[5] Graves, Jenkin, and Murphy (1999) observed similar operational behavior.

[6] In all cases the optimization horizon remains 15 days.

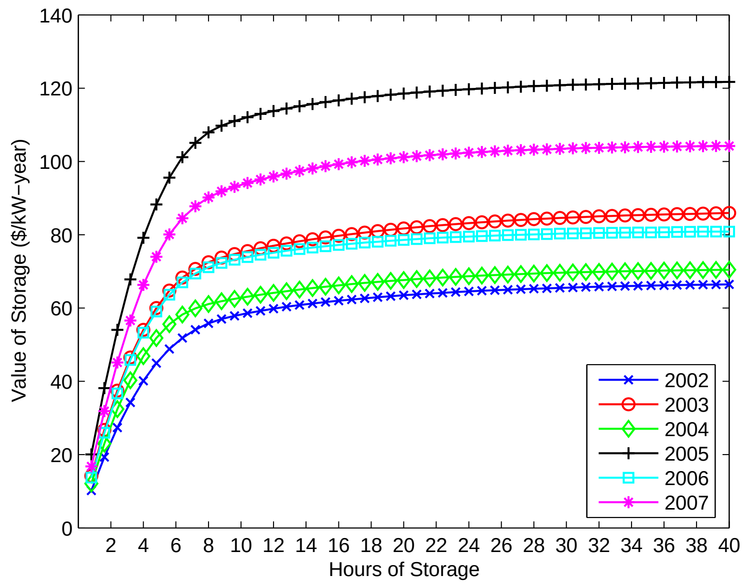

Figure 2: Annual arbitrage value of a price-taking storage device as a function of hours of storage

of potential value.[7] Figure 4 illustrates the marginal value of each additional hour of storage, which falls roughly linearly to about 8 hours, with additional storage providing relatively little incremental arbitrage opportunity.

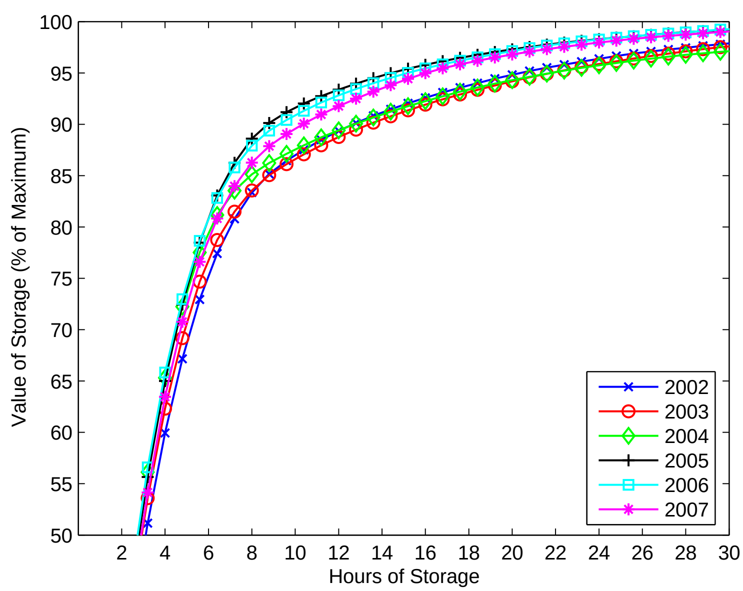

Figure 3: Potential annual arbitrage value captured, as a percentage of maximum theoretical value

The data in figures 2 through 4 can be used to evaluate the optimal size of a storage device for each technology, which will depend on the fixed and variable cost characteristics of the device, as well as efficiency. There is no universal optimal size of storage, because it will depend on the technology and planned applications. Even if only arbitrage is considered, the marginal cost of the next incremental hour of storage can be expected to vary widely by

[7] See also Graves, Jenkin, and Murphy (1999).

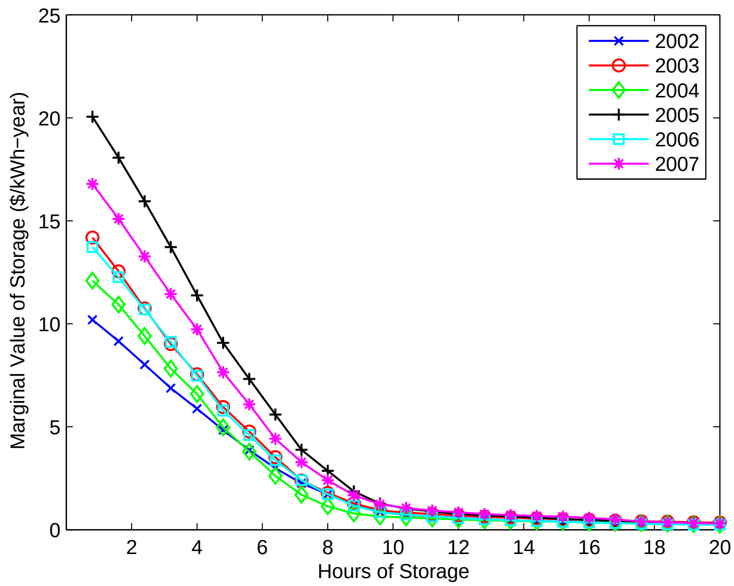

Figure 4: Annual marginal arbitrage value of energy capacity of a storage device

technology, although technology costs and cost structure are not addressed in this paper.[8]

As noted earlier, the storage efficiency of 80% assumed in figures 2 through 4 is at the upper range of actual storage devices currently available. The efficiency of a modern PHS device is in the range of 65-85%, while large batteries (such as sodium-sulfur and vanadium redox) have efficiencies of about 65-75%, as discussed by ASCE (1993) and Denholm and Kulcinski (2004). Figure 5 illustrates the relationship between storage capacity[9] and storage value for systems with a range of efficiencies using average hourly PJM prices in 2006. Efficiency can have a significant impact on the arbitrage value of storage. For example, increasing the efficiency of a 20-hour device from 70% to 80% results in a more than 30% increase in arbitrage value from \$60/kW-year to \$80/kW-year. The reason for this multiplier effect is that a more inefficient device not only needs to charge more hours (for a given number of hours discharged), but these added hours are typically more expensive.[10] Figure 5 also shows the number of hours of storage for each efficiency level at which 90% of the potential maximum value can be captured, showing that between 9 and 10 hours of storage is sufficient for the range of efficiencies examined.

3. Impact of Imperfect Forecasting on Energy Arbitrage Value

One of the limitations of basic arbitrage analysis using historical price data is that it often assumes, as we did in section 2, the optimal operation of the storage device with perfect foresight of hourly energy prices. This approach

[8] To illustrate this idea, consider PHS, which often has more than 20 hours of storage. Part of the reason for the large capacity is that the marginal cost of increasing the size of the reservoir may be small relative to the overall capital costs. Any planned applications beyond arbitrage, such as backup capacity, also may be important.

[9] Because the roundtrip efficiency of the storage device is changed, it is important to note the distinction between storage and discharge hours. As an example, in our simulations, a 50%-efficient device requires two hours of storage for one hour of discharge, whereas an 80%-efficient device requires 1.25 hours of storage for one hour of discharge. As such, all figures use discharge hours on the horizontal axis.

[10] This result does depend on our assumption that a storage device has the same power capacity for charging and discharging. If the charge power capacity is increased, then a more a inefficient device would not need to increase the number of hours it charged for a given discharge.

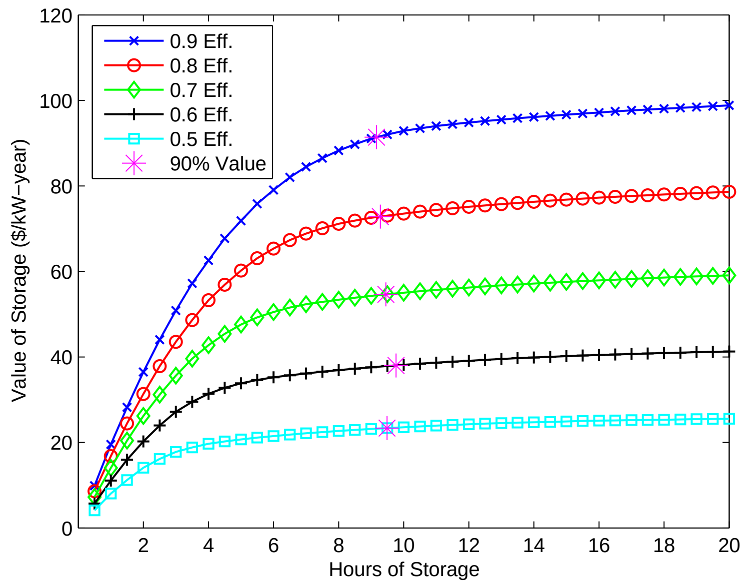

Figure 5: Annual arbitrage value of a storage device with different roundtrip efficiencies in 2006

provides an upper bound on the value of storage. An important question is how close a real operator might come to capturing the theoretical value obtained by perfect foresight. We evaluated the difference between an optimal hourly dispatch and a more realistic approach that does not include any foresight—just knowledge of recent past prices, which is then used to 'guess' the hourly dispatch for the near future. Specifically, we optimized the device in any given two-week period using hourly price data for the *two previous* weeks (which would, of course, be known at that point). Although hourly charge and discharge operations were made using the previous two weeks' price data, the arbitrage value was then estimated using *actual* hourly prices for the two-week period being optimized. In other words, this method 'backcasts' an optimal dispatch for the previous two weeks and applies that to the current two-week period. Each of these two-week estimates were then aggregated to provide estimates for the entire year, and these annual values were compared to the theoretical maximum with perfect foresight of prices.

Figure 6 illustrates the difference between the perfect foresight dispatch and the two-week backcasting approach for a device with 12 hours of storage. In each of the six years evaluated, the backcasting approach captured about 85% or more of the potential arbitrage value. This approach is successful because the hourly operation and value of energy storage is strongly based on historical price and load patterns over a variety of different time-frames, which are to a large extent predictable. The relevant patterns are: (i) the diurnal (or daily) price/load pattern, with fairly predictable hourly off- and on-peak periods, and (ii) the weekday/weekend relationship, with weekends tending to have somewhat lower energy prices. Although the diurnal hourly price patterns differ significantly on a seasonal basis, such differences are largely captured because our backcasting approach only uses a two-week lag.

The simulated two-week backcasting approach does not capture changes in prices that result from nearer-term changes in weather and other short-term load and supply effects, such as generator availability. We would expect it would be relatively straightforward to refine this type of backcasting to substantively increase the value captured *e.g.*, through the use of near-term weather forecasting and the more refined dispatch rules. An example of this is the fact that hourly day-ahead load forecasts are typically within 5% of actual real-time loads—made possible by using historical load patterns and weather forecasting (see a discussion of this fact in PJM (2005)). Such hourly load estimates can, in turn, be used to provide price estimates.

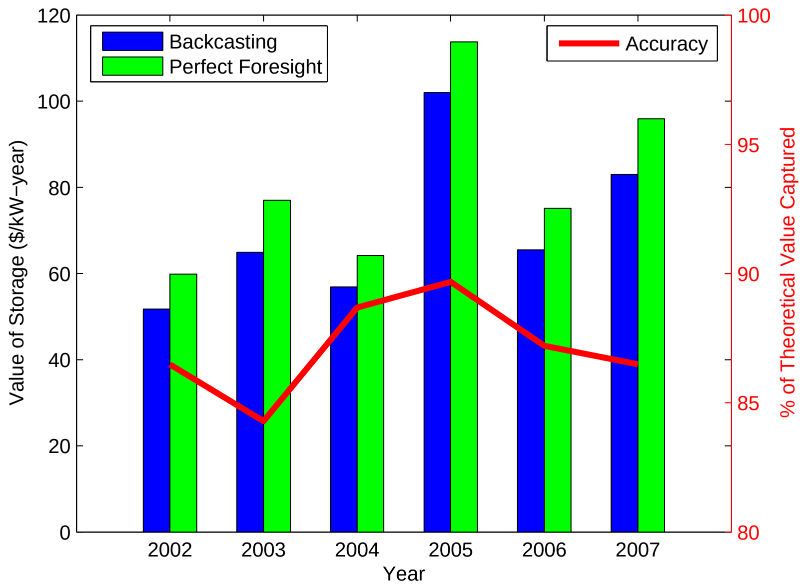

Figure 6: Annual arbitrage value captured by using two-week backcasted dispatch rule versus perfect foresight

We provide this example partly to justify the value of a perfect foresight optimization in obtaining a reasonable estimate of storage value in the price-taking device analysis presented here, and in the large device analysis discussed in section 5. As mentioned above, the no-foresight backcasting approach used here represents a lower bound of value capture that will almost certainly be enhanced by basic forecasting.

4. Variation in the Arbitrage Value of Storage in PJM: The Impact of Temporal and Regional Variation in Fuel and Electricity Prices

As illustrated in the previous sections, the value of storage varies from year to year. In this section we demonstrate that these differences are due partly to variations in fuel price, marginal hourly fuel mix,[11] and transmission constraints. Evaluation of these factors provides an explanation of historical variation in storage value, and can help determine the potential variation in storage value in the future.

Because hourly on-peak electricity prices are often set by natural gas generation, it can be expected that increases in natural gas prices should lead to increases in both hourly on-peak electricity prices and the value of storage. Figure 7 shows the historical monthly price of natural gas sold to electric utilities in the PJM area and the U.S. between 2002 and 2007, based on data from the U.S. Department of Energy's Energy Information Administration.[12] Prices of natural gas have increased from about $3-$4/MMBtu in 2002 to about $6-$8/MMBtu, or more during the past few years. It is important to note that data for 2005 is slightly aberrant—both for the price of natural gas and the value of electricity storage—due partly to the impact of hurricanes on natural gas supplies in the U.S.[13]

Figure 8 illustrates the historical relationship between natural gas fuel prices and arbitrage value. The arbitrage value is derived from figure 2, assuming a 12-hour device. It is evident that the average value of arbitrage for

[11] The marginal hourly fuel mix will depend, primarily, on where the supply curve and load intersect.

[12] The natural gas prices for PJM are actually the cost of natural gas sold to electric generators averaged over New Jersey, Maryland, and Pennsylvania.

[13] While hurricanes Katrina and Rita did have a significant effect on natural gas supplies and the price of natural gas, these prices were high absent these events, with average prices at more than $8/MMBtu for more than six months in 2005.

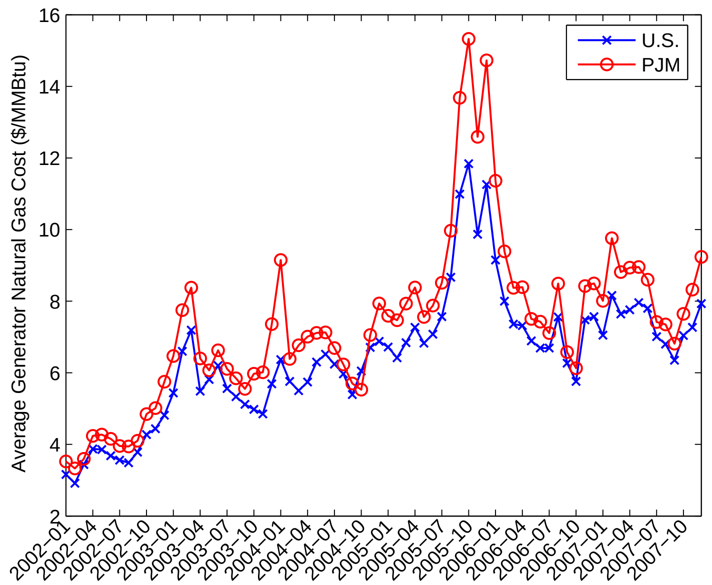

Figure 7: Average monthly price of natural gas paid by electricity generators in the U.S. and in PJM

storage in PJM has increased significantly with increased gas prices, from about \$60/kW-year in 2002 to about \$80-\$100/kW-year, or more in recent years. This 30% to 60% increase in the arbitrage value of storage compared to 2002 is substantive, though significantly less than the more than 100% increase in natural gas prices during the same period. The small increase in the value of storage (relative to the increase in the price of natural gas) can be explained, in large part, by other changes in the PJM market that affected energy prices. The actual value of arbitrage depends on the relationship between off- and on-peak prices, which will depend on the underlying fuel mix of the supply curve and the hourly off- and on-peak loads. In general, storage will be more valuable in regions where nuclear, hydroelectric, and coal are available for off-peak electricity generation.[14]

The relationship between arbitrage values and off- and on-peak price differentials can be observed in figures 9 and 10. The increase in off-peak prices between 2002 and 2005 partially reflects the increase in coal prices, which nearly doubled during this period. Between 2003 and 2004 the arbitrage value decreases despite a small increase in gas prices, due partly to these significant increases in off-peak prices. In contrast, from 2006 to 2007, the on-peak hourly prices and arbitrage value increase despite nearly flat natural gas prices. Explanation of this requires an examination of the actual fuel mix providing off- and on-peak energy.

Figure 11 illustrates the fraction of the marginal fuel mix provided by coal and natural gas during each hour of the year in 2006 and 2007. During this time period, the fraction of the marginal fuel mix derived from coal decreases during on-peak hours with an increase in the percentage of time natural gas sets the margin, resulting in higher on-peak prices. This occurs even with no significant change in the price of natural gas, because natural gas is significantly more expensive than coal.[15]

It is important to note that all previous arbitrage estimates have used load-weighted average hourly PJM prices. Within PJM, the value of arbitrage can be expected to vary in different by location due to transmission constraints;

[14] It is of interest to note that while off-peak prices are largely set by coal, it is never set by lower-cost nuclear power—despite the fact that more than 33% of total generation in PJM in 2007 came from nuclear generation.

[15] It should be noted that other more expensive fuels, such as oil, also contribute to setting the margin during peak periods (primarily in the summer), and increases in the price of oil also drove the increase in on-peak electricity prices.

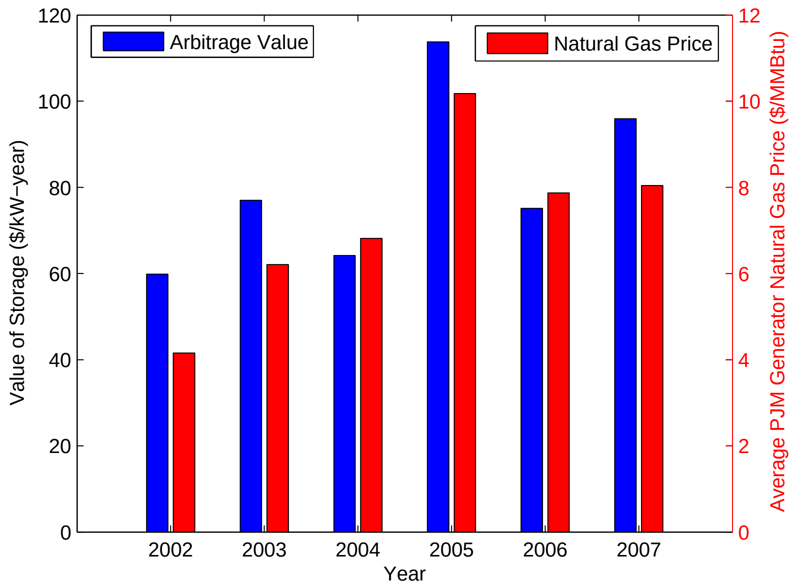

Figure 8: Annual value of arbitrage and annual average price of natural gas paid by generators in PJM

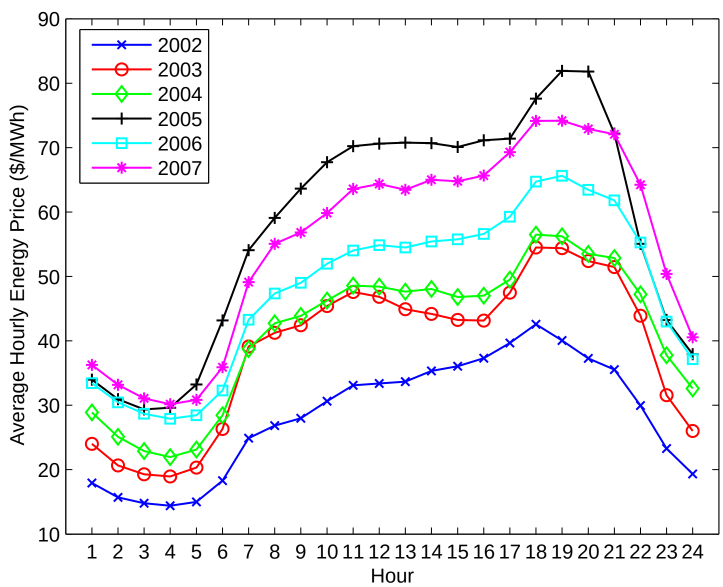

Figure 9: Annual average hourly price of electricity

and, accordingly, the value of arbitrage may be considerably higher at different locations than the arbitrage values calculated using load-weighted average prices. Figure 12 illustrates the variation in annual arbitrage value for different bus locations within PJM in 2006, assuming a device with 80% efficiency and 16 hours of storage capacity. While the average value of arbitrage in PJM for 2006 was $77/kW-year, the value at individual buses can be as high as $105/kW-year, corresponding to an almost $30/kW-year premium.[16]

The analysis presented in these sections represents a 'static' valuation of energy storage arbitrage. Increased

[16]This corresponds to the Bedington bus in PJM (the darkest red point on the map). Similarly, in 2007, while the arbitrage value based on average PJM prices was $99/kW-year, the arbitrage value at the same bus is $137/kW-year—giving a larger premium.

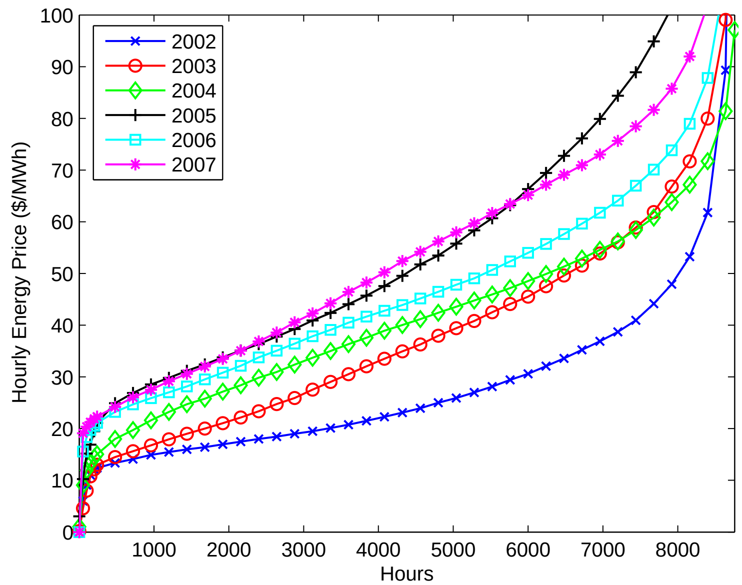

Figure 10: Hourly electricity price duration curve

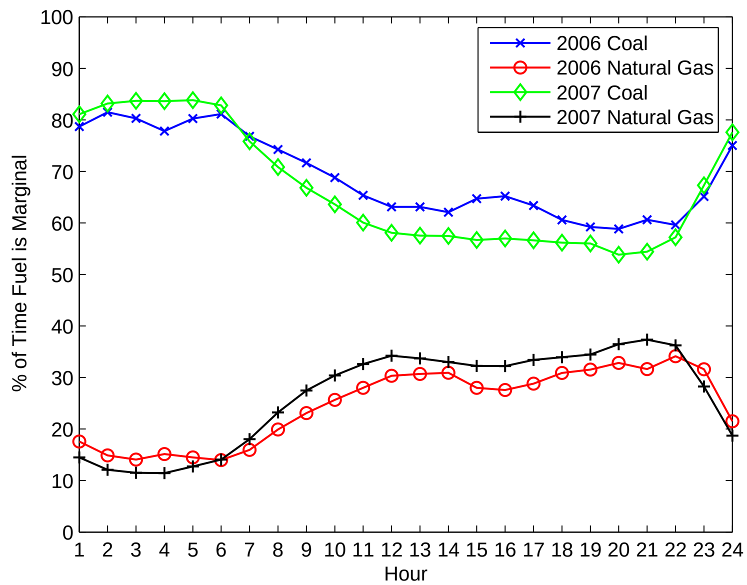

Figure 11: Percent of time in which coal and natural gas were marginal fuel in each hour for 2006 and 2007

transmission capacity potentially can decrease the regional differences in value, while load growth in congested areas (without corresponding increases in transmission capacity or local generation) will tend to increase arbitrage opportunities. Storage may also potentially provide an alternative or complement to transmission to relieve congestion, although the economic evaluation of this application is extremely site specific.[17]

[17] One interesting idea discussed by Eyer, Iannucci, and Butler (2005) is that storage valuation should account for the fact that a relocatable modular storage device might be moved to various locations during its asset life.

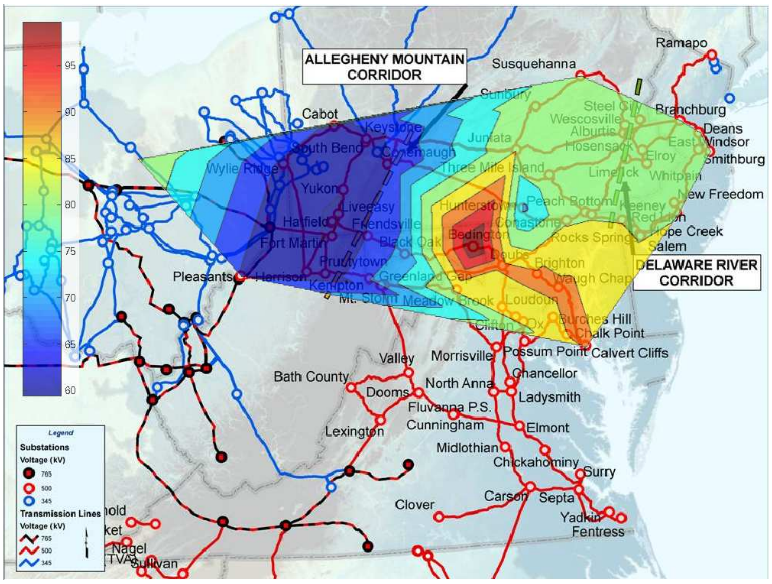

Figure 12: Annual arbitrage value at 47 bus locations within PJM in 2006

5. Impacts of Large-Scale Storage

As the amount of storage in a system increases, the arbitrage value on a \$/kW-year basis will decrease as increasing amounts of on-peak load is shifted to off-peak periods, resulting in lower on-peak prices and higher off-peak prices, thereby reducing the arbitrage value of storage. Figure 13 provides an example of actual hourly price/load data for a single month (June 2006) and an ordinary least squares (OLS) estimate of a linear relationship between price and load, which shows a strong fit. The price/load relationship is driven by where the inelastic load intersects the generation supply curve, and as such the price/load points 'map out' the supply function. The net result of large amounts of storage will be a flattening of the diurnal generation load profile[18] and a corresponding flattening of prices. Theoretically, entry by storage devices should occur until all profitable opportunities to buy inexpensive energy off-peak and sell expensive energy on-peak are arbitraged away.

5.1. Impacts of Large-Scale Storage on Arbitrage Value

We analyze the effects of large-scale storage by modeling the operation of a large storage device,[19] which accounts for the effect that its charging and discharging has on the price of energy. We assume there is a non-decreasing linear relationship between the price of energy and generating load, such as the one shown in figure 13. Because of seasonal differences in fuel costs, generation mix, and loads we assume that each month has a different linear price-load relationship, and estimate the parameters of the function for each month by restricted least-squares[20] using the actual price and load data from that month. We model the storage device's hourly operations by maximizing arbitrage value using the same two-week optimization horizon discussed in section 2, and further assume the storage operator knows the parameters of the price relationship and load for each two-week period with perfect foresight.

[18] It should be noted that absent time-variant retail rates such as real-time pricing or time-of-use rates, the load profile will remain the same because demand does not respond to the use of storage. The generation profile will, however, change in response to charging and discharging of the storage device.

[19] This analysis can be generalized to multiple storage devices that collusively act to maximize total arbitrage value.

[20] The constraint on the OLS estimate, which is always non-binding, is that the price-load relationship be non-decreasing.

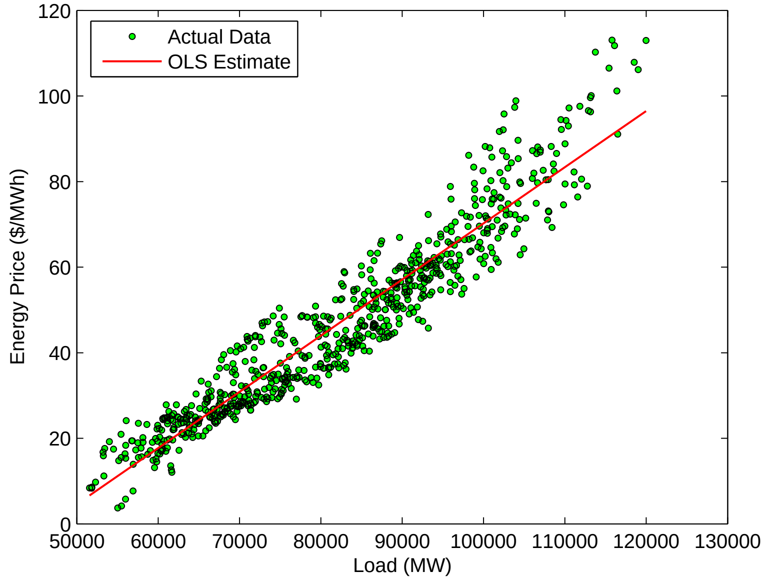

Figure 13: Hourly price-load relationship in June 2006

Thus, our analysis assumes that the storage operator perfectly anticipates hourly electricity prices and the effect that hourly charging or discharging would have on those prices. Because the price-load relationship is assumed to be linear non-decreasing, the resulting optimization is a convex quadratic program, and first-order necessary conditions are sufficient for a global optimum. The model was formulated in GAMS 21.7 and solved using MINOS 5.5. Appendix A.2 gives the explicit formulation of our model and discusses it in more detail.

The operation of the storage device with prices varying in response to generating loads will be largely similar to that with prices fixed, with the storage device charging when prices are low and discharging when prices are high. Figure 14 contrasts electricity prices and the differences in the operation of a 1 GW device with 12 hours of storage over a sample week-long period in 2006, with varying and fixed prices.[21] [22] The prices show the expected smoothing behavior with lower prices on-peak and higher prices off-peak due to changes in the generating load resulting from operation of the storage device. The operation of the storage device also shows changes. We saw in figure 1 that with fixed prices the device is always operated at the lesser of its energy and power capacity when discharging or charging. With varying prices, charges and discharges are sometimes curtailed when the price impacts reduce the marginal arbitrage value to zero. In other cases, such as on Friday morning, the device does not operate at all with varying prices, even though it would with fixed prices. While 1 GW is a large amount of storage, it is worth noting that a number of pumped storage facilities in the U.S. are 1 GW or greater, for example the Tennessee Valley Authority's Raccoon Mountain PHS plant can continuously discharge at 1.6 GW for 22 hours.

Figure 15 summarizes the value of 1 GW of storage, showing the percentage of potential value that can be captured if prices respond to generating loads, compared to assuming the prices are fixed but follow the same linear price-load relationship. Our analysis shows that the value of storage would have been diminished relative to a price taker-by approximately 10% during the past three years, but more than 20% for some earlier years. The reason for the differences across the years stems from the fact that PJM grew between 2002 and 2007 by adding adjacent control

[21] Although figure 14 only shows prices and dispatch for one week, the optimization was done using two-week planning horizons.

[22] To make the two cases comparable, the fixed prices were derived from the price-load relationship using the actual system load.

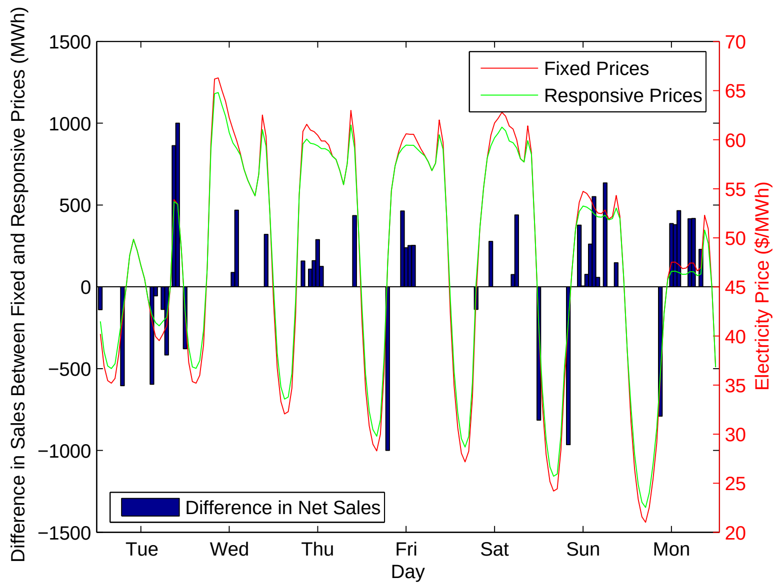

Figure 14: Electricity prices and difference in net sales with fixed and responsive prices

areas. As such, a 1 GW device represents a smaller device relative to the size of the system in 2007 as opposed to 2002. In 2002, for instance, the peak load was 63,761 MW with an average load of 35,470 MW, whereas in 2007 these values were more than doubled to 139,427 MW and 82,667 MW, respectively. Because the off- and on-peak price difference did not change in proportion to the size of the load, a 1 GW device would have a much larger price-shifting effect in 2002 than in 2007. This is reflected in the fact that the value of a large storage device is diminished much more in the earlier years compared to the later ones.

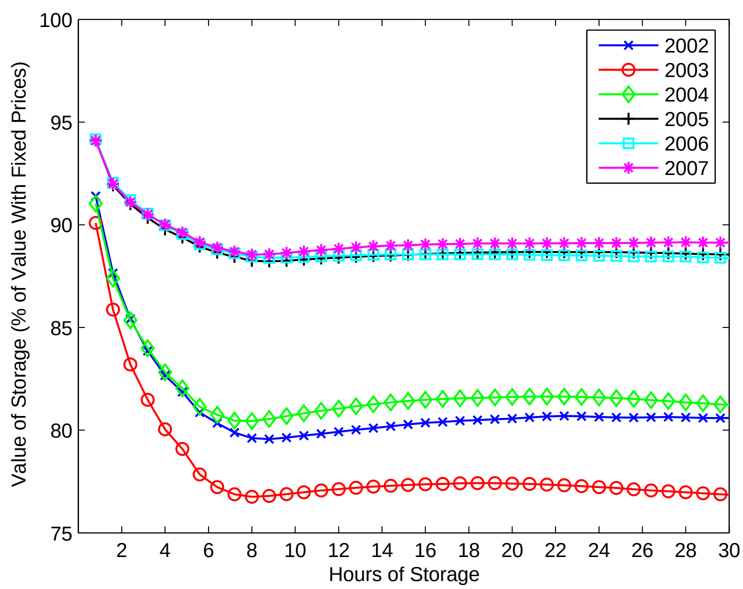

Figure 15: Arbitrage value captured with price responsiveness, as a percentage of value with fixed prices

5.2. External Welfare Effects of Large-Scale Storage

In addition to the arbitrage value captured by the storage device owner, consumers and generators will also benefit and lose from the use of energy storage. Reduced on-peak and increased off-peak prices to consumers can result in consumer surplus gains, stemming from lower energy costs to consumers. Because of the relationship between energy prices and generation, an increase in generation off-peak (due to energy storage) with an offsetting decrease in generation on-peak (due to meeting some load with energy discharged from the storage device) will result in increases and decreases in prices off- and on-peak. However, because consumer demand tends to be significantly lower off-peak, the decrease in consumer surplus from the higher price that is paid off-peak will be more than offset by an increase in consumer surplus on-peak due to decreased generation needs and a corresponding drop in the price of energy. Conversely, generators will generally see their profits decrease from use of a storage device, because the increase in profits off-peak will be offset by the drop in profits on-peak.

Figure 16 demonstrates this effect for a single paired hourly charge/discharge cycle of a storage device. The line represents the marginal cost of electricity generation as a function of generation, which we assume sets the wholesale price of electricity. Without any charging or discharging, the load and energy generated off- and on-peak are given by l_1 and l_2, respectively, and the price of energy would be p_1 and p_2 off- and on-peak, respectively. When the storage device charges off-peak and discharges on-peak, consumer demand remains the same, but the generating load increases to $\tilde{l}_1$ off-peak and decreases to $\tilde{l}_2$ on-peak, with commensurate changes in the energy price off- and on-peak to $\tilde{p}_1$ and $\tilde{p}_2$, respectively. As a result of these changes in prices and generation quantities, there will be changes in consumer and producer surplus. Consumer surplus decreases by the rectangle labeled A off-peak due to the higher price of energy, but increases by the sum of the areas labeled C, D, and E on-peak due to the lower on-peak price of energy. Producer surplus increases by the sum of the areas labeled A and B off-peak due to higher generating loads and an increase in the energy price, and decreases by the areas labeled C and D on-peak. Adding these terms, the effect of the charge/discharge cycle is that the sum of consumer and producer surplus increase by the sum of the areas labeled B and E.

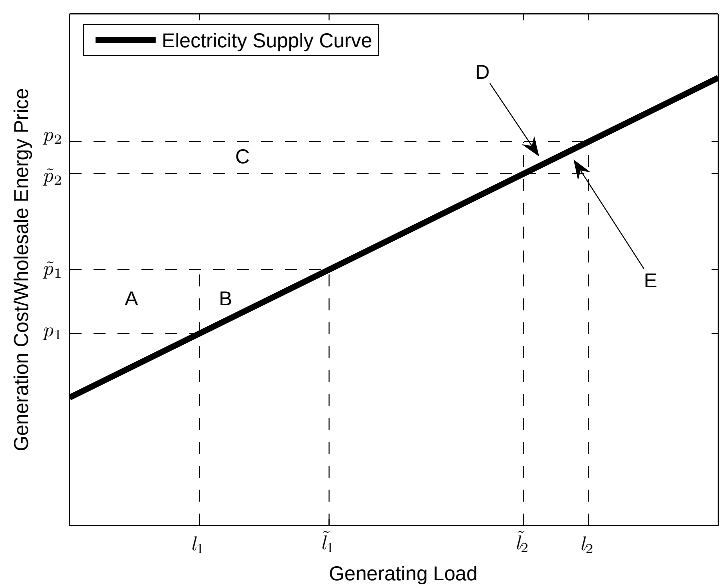

Figure 16: Linear price-load relationship

We analyze the welfare effects of using a large storage device, assuming the same linear relationship between prices and generating loads and that the storage device is operated to maximize arbitrage value (*i.e.*, the storage operator does not consider external welfare effects). The surplus calculations are based on the changes in prices and generation shown in figure 16. In computing producer surplus changes we assume that generators behave competitively and prices reflect the actual marginal cost of generation. Tables 1 through 3 summarize these welfare effects for a 1 GW storage device with 4, 8, and 16 hours of storage with 80% efficiency in 2002 and 2007. Our results show that the external welfare effects for consumers and producers are on the same relative scale as the arbitrage value. Moreover, although there are large wealth transfers from generators to consumers, the fact that the increase in consumer surplus is greater than producer surplus loss shows that there are net social welfare gains stemming from the load-shifting effects of large-scale storage.

Table 1: Social value of storage device with 4 hours of storage ($ million)

Year	Arbitrage Value	ΔCS	ΔPS
2002	26.8	16.8	-14.3
2007	47.3	22.7	-20.2

Table 2: Social value of storage device with 8 hours of storage ($ million)

Year	Arbitrage Value	ΔCS	ΔPS
2002	37.0	21.5	-17.3
2007	64.9	27.3	-23.4

Table 3: Social value of storage device with 16 hours of storage ($ million)

Year	Arbitrage Value	ΔCS	ΔPS
2002	42.1	26.3	-21.7
2007	73.7	34.6	-30.3

5.3. Impacts of Large-Scale Storage on Reducing Consumer Impacts of Electricity Price Shocks

Large-scale storage can also potentially help mitigate the impact of price volatility resulting from supply disruptions, such as that which occurred with hurricanes Katrina and Rita in 2005. While natural gas prices were high that year in general, supply disruptions resulted in transient jumps in natural gas and electricity prices. Figure 17 shows price and load data for two days before and after Hurricane Katrina landed, which are one week apart (August 25 and September 1). While the load pattern was similar for both days, it is clear that the price-load relationship has changed, in that there is a much greater off- and on-peak price difference, with a corresponding increase in the slope of the price-load relationship after the hurricane. As such, one effect of the hurricane was to increase electricity costs for consumers in PJM (as well as other parts of the US in which natural gas-fired generation would set the margin). One potential impact of the load-shifting effect of large-scale storage is that it could help to mitigate the effect of these price shocks by reducing the increase in on-peak electricity prices. While there would be a corresponding increase in off-peak prices when the storage device is charged, there would likely be large net increases in consumer surplus because this price increase is applied to a smaller load than the on-peak price. Moreover, because of the steeper price-load relationship post-hurricane, the changes in off- and on-peak prices after the hurricane will be larger than they would have before.

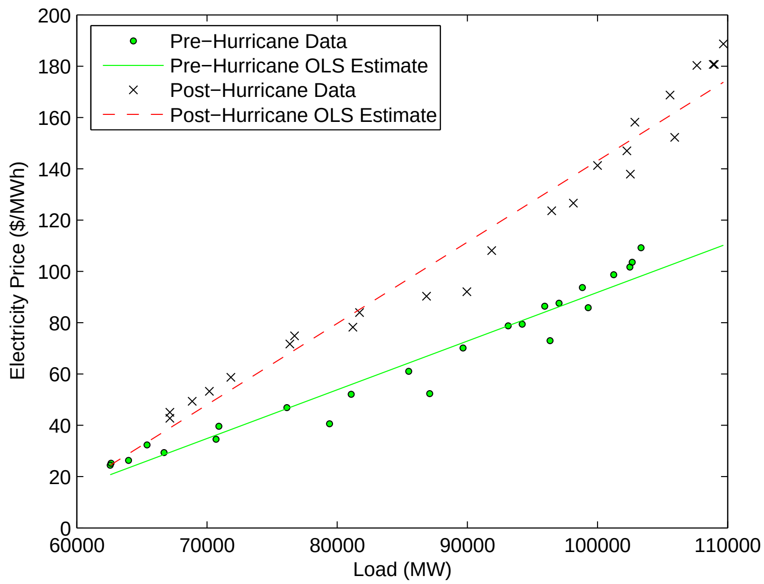

Figure 17: Hourly price-load data and OLS estimate of relationship before and after Hurricane Katrina

We examined this potential benefit of large-scale storage by using the price-load relationship for each of the two days, simulating the operation of a 1 GW storage device with 8 hours of storage. The simulation was done for each day separately, and only optimized over the one day. As before, the storage device is assumed to be a profit-maximizer (*i.e.*, it does not consider consumer or producer surplus changes). Table 4 summarizes the arbitrage value and consumer and producer surplus changes before and after the hurricane. Our results show a more than 70% increase in arbitrage value after the hurricane, stemming from the larger off- and on-peak price difference. The analysis also shows an increase in the consumer surplus change of about 70%, due to the increased price-shifting ability of a large storage device and its ability to partially insulate consumers from higher on-peak prices.

Table 4: Arbitrage value and changes in consumer and producer surplus before and after hurricane ($)

Day	Arbitrage Value	ΔCS	ΔPS
Pre-hurricane	345,000	188,000	-172,000
Post-hurricane	590,000	320,000	-295,000

Our analysis of the load-shifting effects associated with large-scale storage is illustrative to the extent that we have used a linear price-load relationship. This assumption of a linear relationship yields a convex quadratic programming problem, which makes the analysis tractable. Figures 13 and 17 demonstrate that using short time frames of a month or less to fit the linear relationship can provide a good fit to the data.

6. Discussion and Conclusions

Wholesale electricity markets in many regions make it possible to evaluate the potential arbitrage value of energy storage in many parts of the country and around the world. Our analysis shows that there are a number of drivers behind the value of arbitrage including location, fuel price, fuel mix, efficiency, and device size, as well as the hourly load profile. In the case evaluated here, the annual value of arbitrage for a price-taking storage device in PJM with an 80% roundtrip efficiency and a storage time of 12 hours was found to have increased from about \$60/kW-year to

$110/kW-year or more in recent years. These estimates of arbitrage value based on average PJM prices undervalues the actual potential regional value of arbitrage, with our analysis showing that certain buses within PJM have an additional premium of $20/kW-year to $30/kW-year or more due to transmission congestion and losses.

As expected, the marginal arbitrage value (on a $/kW basis) of the next hour of storage drops sharply as a function of energy capacity, with the knee of the curve at about 8 hours. There is no one optimal number of hours of capacity for a storage device, rather, it will vary by technology and applications. Even if only arbitrage is considered, different technologies may have different fixed and variable costs (for both energy and power capacity) and efficiencies.

Capturing value through energy arbitrage requires short-term forecasting of hourly electricity prices to appropriately charge and discharge to maximize price differentials. Perfect foresight of energy prices appears to be a reasonable approximation of actual value capture, based on our simple two-week backcasting-based dispatch capturing about 85% or more of the theoretical value. Moreover, this value would be expected to improve significantly if simple forecasting techniques were added.

The observed annual variation and general increase in arbitrage value between 2002 and 2007 is driven by the difference in hourly off- and on-peak electricity prices, which themselves are driven by the underlying cost of fuel and the fuel mix, which in turn will depend on the load. The increase in natural gas prices is the main driver, though changes in the amount of time natural gas provides the marginal generation fuel on-peak and changes in coal and oil prices also will be important.

In PJM and other energy and capacity markets a storage device may also be eligible for capacity payments in addition to the energy arbitrage value estimated above. Such payments are designed to encourage additional capacity where price caps limit energy prices. The value of such payments is highly uncertain, and so we have chosen to mention them here as a potential adder rather than estimate them, although they may be substantial.[23] Another source of value from storage can come from co-optimizing between different markets, such as energy arbitrage and ancillary services (*e.g.*, frequency regulation and spinning reserves)—though these have not been considered in this paper.

The introduction of energy storage on a large scale has the potential to increase off-peak prices and decrease on-peak prices, thereby decreasing the value of energy arbitrage. Arbitrage is not, however, the only important source of value, especially for devices that can shift load and prices. Specifically, despite this decrease in arbitrage value, large-scale storage can potentially provide other social welfare improvements, including improved utilization of the electricity infrastructure, deferred need to build generation, T&D assets, and the ability to reduce congestion. The value of these benefits can be significant, though are extremely site-specific and some of these benefits are hard to quantify. We demonstrated that there can be large shifts in consumer and producer surplus, associated with increases in prices to consumers when the device is charged and decreases in prices when discharged. Because the on-peak load is greater than off-peak load the use of large amounts of storage can lead to significant net increases in consumer surplus, and associated decreases in costs to end users. We also showed that this welfare-shifting effect can provide a partial risk mitigation tool against supply disruptions and dampen increases in consumer energy costs.

Because these external welfare benefits will not necessarily by captured by a private sector investor who relies on arbitrage, such an investor may have a reduced incentive to invest in energy storage due to the diminished value of arbitrage. This raises questions regarding the best ownership structure for large amounts of storage, because it will

[23] As an example, Felder and Newell (2007) recognize the difficulties that arise in using historical capacity market prices to estimate the value of capacity payments, due to market structure problems. Instead, they use the levelized cost of new entry, which they estimate at $58/kW-year, as a proxy value.

impact both investment and operational decisions. In contrast to the decreasing benefits seen by a private owner, a transmission owner or regulated entity may have better incentives to invest in energy storage due to its valuing the external social benefits. A transmission owner, for example, may also benefit from decreased congestion costs and benefits associated with better use of infrastructure assets.[24] As a result, a regulated storage owner may view load shifting—and its many attributes beyond arbitrage—as a benefit to society, while this may not be the case for a merchant operator.[25]

In summary, the recent increases in the price of natural gas suggest a growing potential role for storage in some electric power systems. However, any analysis of energy storage that considers only one or a few attributes (such as energy arbitrage) and neglects the interplay among various sources of value is likely to significantly underestimate the value and social benefits of energy storage. The ability to realize the inherent value of storage will vary markedly with ownership, contract, and market structure. All of these factors need to be considered with cost and other potential alternatives when making any real investment decision.

References

ASCE (1993): "Compendium of Pumped Storage Plants in the United States," Task Committee on Pumped Storage of the Hydropower Committee of the Energy Division of the American Society of Civil Engineers.

Denholm, P., and G. L. Kulcinski (2004): "Life cycle energy requirements and greenhouse gas emissions from large scale energy storage systems," *Energy Conversion and Management*, 45, 2153–2172.

EPRI (2003): "EPRI-DOE Handbook for Energy Storage for Transmission and Distribution Applications," 1001834.

Eyer, J. M., J. J. Iannucci, and P. C. Butler (2005): "Estimating Electricity Storage Power Rating and Discharge Duration for Utility Transmission and Distribution Deferral," A Study for the DOE Energy Storage Systems Program.

Eyer, J. M., J. J. Iannucci, and G. P. Corey (2004): "Energy Storage Benefits and Market Analysis Handbook," A Study for the DOE Energy Storage Systems Program.

Felder, F. A., and S. A. Newell (2007): "Quantifying Demand Response Benefits In PJM," The Brattle Group, Inc.

Figueiredo, F. C., P. C. Flynn, and E. A. Cabral (2006): "The Economics of Energy Storage in 14 Deregulated Power Markets," *Energy Studies Review*, 14, 131–152.

Graves, F., T. Jenkin, and D. Murphy (1999): "Opportunities for Electricity Storage in Deregulating Markets," *The Electricity Journal*, 12, 46–56.

Jenkin, T., and J. Weiss (2005): "Estimating the Value of Electricity Storage: Some Size, Location and Market Structure Issues," in *Electrical Energy Storage Applications and Technologies Conference*, San Francisco, CA.

[24] Some of these ideas are explored in Jenkin and Weiss (2005).

[25] The question of ownership is more complicated than covered here. For example, many of the issues facing a merchant storage operator might be mitigated if, for example, the merchant owner is compensated for the benefits associated with load shifting, such as congestion relief. At the same time some real benefits associated with storage, such as better system utilization, may be difficult to quantify or specify in a contract.

PJM (2005): "PJM System Operator Seminar Spring—2005," http://www.pjm.com/services/training/downloads/20050607-seminar-state-pjm-system.pdf.

Walawalkar, R., J. Apt, and R. Mancini (2007): "Economics of electric energy storage for energy arbitrage and regulation in New York," *Energy Policy*, 35, 2558–2568.

A. Formulation of Models

We describe the models used in our analysis of small price-taking and large storage devices.

A.1. Price-Taking Storage Model

We first define the following notation:

Problem Parameters

- T: number of hours in dispatch horizon
- p_t: energy price in hour t
- κ: power capacity of storage device
- h: number of hours of storage
- η: roundtrip efficiency of storage device

Decision Variables

- c_t: energy charged into storage device in hour t
- d_t: energy discharged from storage device in hour t
- s_t: energy stored in storage device at end of hour t

The problem is formulated as maximizing arbitrage value:

$$\max_{c,d,s} \sum_{t=1}^{T} p_t \cdot (d_t - c_t);$$

subject to the following constraints:

- storage level definition ($\forall\ t = 1, 2, \ldots, T$):
$$s_t = s_{t-1} + \eta \cdot c_t - d_t;$$
- power capacity of device ($\forall\ t = 1, 2, \ldots, T$):
$$d_t, c_t \in [0, \kappa];$$
- energy capacity of device ($\forall\ t = 1, 2, \ldots, T$):
$$s_t \in [0, h \cdot \kappa].$$

A.2. Large-Scale Storage Model

We first define the following notation:

Problem Parameters

- T: number of hours in dispatch horizon
- $p(g)$: non-decreasing linear function giving price/load relationship
- l_t: consumer electricity load in hour t
- κ: power capacity of storage device
- h: number of hours of storage
- η: roundtrip efficiency of storage device

Decision Variables

- c_t: energy charged into storage device in hour t
- d_t: energy discharged from storage device in hour t
- s_t: energy stored in storage device at end of hour t

The problem is formulated as maximizing arbitrage value:

$$\max_{c,d,s} \sum_{t=1}^{T} p(l_t - d_t + c_t) \cdot (d_t - c_t);$$

subject to the following constraints:

- storage level definition ($\forall\ t = 1, 2, \ldots, T$):

$$s_t = s_{t-1} + \eta \cdot c_t - d_t;$$

- power capacity of device ($\forall\ t = 1, 2, \ldots, T$):

$$d_t, c_t \in [0, \kappa];$$

- energy capacity of device ($\forall\ t = 1, 2, \ldots, T$):

$$s_t \in [0, h \cdot \kappa].$$