

FACILITATING THE DEVELOPMENT AND
INTEGRATION OF LOW-CARBON ENERGY
TECHNOLOGIES

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
THE DEGREE OF
DOCTOR OF PHILOSOPHY
IN
ENGINEERING AND PUBLIC POLICY

EMILY FERTIG

B.A., GEOSCIENCES, WILLIAMS COLLEGE

CARNEGIE MELLON UNIVERSITY
PITTSBURGH, PA

MAY 2013

© Copyright by Emily Fertig, 2013.

All rights reserved.

Abstract

Climate change mitigation will require extensive decarbonization of the electricity sector. This thesis addresses both large-scale wind integration (Papers 1–3) and development of new energy technologies (Paper 4) in service of this goal.

Compressed air energy storage (CAES) could be paired with a wind farm to provide firm, dispatchable baseload power, or serve as a peaking plant and capture upswings in electricity prices. Paper 1 presents a firm-level engineering-economic analysis of a wind/CAES system with a wind farm in central Texas, load in either Dallas or Houston, and a CAES plant whose location is profit-optimized. Of a range of market scenarios considered, the CAES plant is found to be profitable only given the existence of large and infrequent price spikes. Social benefits of wind/CAES include avoided construction of new generation capacity, improved air quality during peak demand, and increased economic surplus, but may not outweigh the private cost of the CAES system nor justify a subsidy.

Like CAES, pumped hydropower storage (PHS) ramps quickly enough to smooth wind power and could profit from arbitrage on short-term price fluctuations exacerbated by large-scale wind. Germany has aggressive plans for wind power expansion, and Paper 2 analyzes an investment opportunity in a PHS facility in Norway that practices arbitrage in the German spot market. Price forecasts given increased wind capacity are used to calculate profit-maximizing production schedules and annual revenue streams. Real options theory is used to value the investment opportunity, since unlike net present value, it accounts for uncertainty and intertemporal choice. Results show that the optimal investment strategy under the base scenario is to wait approximately eight years then

invest in the largest available plant.

Paper 3 examines long-distance interconnection as an alternate method of wind power smoothing. Frequency-domain analysis indicates that interconnection of aggregate regional wind plants across much of the western and mid-western U.S. would not result in significantly greater smoothing than interconnection within a single region. Time-domain analysis shows that interconnection across regions reduces the magnitude of low-probability step changes and doubles firm power output (capacity available at least 92 % of the time) compared with a single region. An approximate cost analysis indicates that despite these benefits, balancing wind and providing firm power with local natural gas turbines would be more cost-effective than with transmission interconnection.

Papers 1 and 3 demonstrate the need for further RD&D (research, development, and deployment) of low-carbon energy technologies. Energy technology development is highly uncertain but most often modeled as deterministic, which neglects the ability both to adapt RD&D strategy to changing conditions and to invest in initially high-cost technologies with small breakthrough probabilities. Paper 4 develops an analytical stochastic dynamic programming framework in which RD&D spending decreases the expected value of the stochastic cost of a technology. Results for a two-factor cost model (which separates RD&D into R&D and learning-by-doing) applied to carbon capture and sequestration (CCS) indicate that given 15 years until large-scale deployment, investment in the RD&D program is optimal over a very broad range of initial mitigation costs (\$10–\$380/tCO₂). While the NPV of the program is zero if initial mitigation cost is \$100/tCO₂, under uncertainty the program is worth about \$7 billion. If initial mitigation cost is high, the program is worth most if cost reductions exogenous to the program (e.g. due to private sector activity) are also high. Factors that promote R&D spending over learning-by-doing include more imminent deployment, high initial cost, lower exogenous cost reductions, and lower program funds available.

Acknowledgments

I am grateful to several people and organizations who have supported me over the course of completing this dissertation.

I would like to acknowledge the support of the National Science Foundation's Graduate Research Fellowship Program, as well as the John and Claire Bertucci Fellowship and the Steinbrenner Institute Graduate Fellowship. This research was further supported by grants from the Electric Power Research Institute (EPRI) to the Carnegie Mellon Electricity Industry Center, from the Doris Duke Charitable Foundation, the Richard King Mellon Foundation, the Heinz Endowments to the RenewElec program at Carnegie Mellon University, and the U.S. National Science Foundation under Award SES-0949710.

I would like to thank my advisor Jay Apt for his support and excellent guidance, and for shaping my thinking on how to define and approach energy problems that matter. I would also like to thank Granger Morgan for his big-picture guidance on my work and for his foundational role in such a unique and intellectually courageous department. I'm also grateful to the other two members of my thesis committee, Dalia Patiño-Echeverri and Mort Webster, for their investment of time and for the new perspectives they brought to this work.

A substantial portion of this work was completed at the Norwegian University of Science and Technology, and I would like to thank my main advisor there, Gerard Doorman, for his insights and support both in research and in integrating me into the department. I would also like to thank Professors Stein-Erik Fleten, Bjørn Bakken, Ivar

Wangensteen, and Anders Gjelsvik for their time and insights. I'm grateful to the NSF and the Research Council of Norway for their funding support through the Nordic Research Opportunity.

I was fortunate to have been among such an intelligent and interesting group of EPP students, and I would especially like to thank Anu Narayanan, Vanessa Schweizer, Warren Katzenstein, Eric and Rachael Hittinger, Steve Rose, Brinda Thomas, Kim Mullins, Tim Gordon, Kyle Siler-Evans, Shantanu Jathar, Marla Sanchez, and Olga Popova for the good memories, productive research discussions, and nonproductive nonresearch discussions.

I would like to thank Ane Marte Heggedal, co-author of Paper 2 of this thesis, whose intellect and approach to problem-solving helped me grow as a researcher. I would also like to thank my NTNU classmate Christian Skar, whose work ethic and intellectual curiosity also made him a joy to work with.

I am grateful to Celia Wexler and Francesca Grifo at the Union of Concerned Scientists and Katrina Kelner at AAAS/*Science* for their mentorship and encouragement of my interest in science policy.

I would also like to thank my mom Susan for her unconditional support and interest, my dad Doug for setting the example of trying to make the world a better place through his work, and my brother Andrew for his support, mostly in the form of a lifetime of reminders that I am a giant nerd.

Last and far from least, I'd like to thank Tor for his encouragement and care, for staunching the rising tide of mess in the apartment, and for his \LaTeX wizardry in `empack.sty`.

Contents

- Abstract** **i**
- Acknowledgments** **iii**
- Contents** **v**
- List of Figures** **ix**
- List of Tables** **xii**
- Introduction** **xiii**
 - Bibliography xviii
- 1 Economics of compressed air energy storage to integrate wind power: a case study in ERCOT** **1**
 - 1 Introduction 2
 - 2 CAES mechanics and extant plants 4
 - 3 Wind/CAES system model 6
 - 3.1 Physical design 6
 - 4 Siting the CAES plant 7
 - 5 Data and energy markets 10
 - 6 Wind/CAES system cost models 12
 - 6.1 CAES plant 12
 - 6.2 Transmission 12
 - 7 Wind/CAES heuristic dispatch strategies and hourly profit models 14
 - 7.1 Balancing Energy Services (BES) market 14
 - 7.2 Regulation and balancing energy markets 16
 - 8 Results 17
 - 8.1 Balancing energy market — zonal prices 17
 - 8.2 Balancing energy market — price cap of \$300/MWh plus capacity payment of \$100/MWd 18
 - 8.3 Balancing energy market — contract price 19
 - 8.4 Analysis of the price-taker assumption for the zonal price scenario 19
 - 8.5 BES and regulation markets 21
 - 8.6 Carbon price for an economically competitive wind/CAES system 21
 - 9 Discussion and policy implication 22

Acknowledgements	24
A CAES mechanics	25
A.1 Ramp rates	26
A.2 Adiabatic CAES	26
B Extant and planned CAES plants	27
B.1 Wind-CAES ancillary services	28
C CAES geology in Texas	28
C.1 CAES in solution-mined salt caverns	29
C.2 Occurrence of salt formations amenable to CAES	30
C.3 CAES in saline aquifers	30
C.4 Occurrence of saline aquifers amenable to CAES	32
C.5 Depleted natural gas fields	34
D CAES plant cost	34
E Transmission capital cost	35
F Wind-CAES dispatch model	38
G Optimization algorithm	40
H Extended carbon price results	40
Bibliography	41
2 Optimal investment timing and capacity choice for pumped hydropower storage	45
1 Introduction	46
1.1 Prospects for pumped hydropower storage	46
1.2 The effect of increased wind capacity on electricity prices	48
1.3 The use of real options to value investment in pumped storage	49
1.4 Structure of the analysis	50
2 Valuation Framework	52
2.1 Price Model	52
2.2 Pumped Storage Scheduling Model	56
2.3 Real Option Valuation	58
3 Numerical Example	61
4 Results	63
5 Conclusions and Further Work	71
Acknowledgements	72
Bibliography	74
3 The effect of long-distance interconnection on wind power variability	76
1 Introduction	77
2 Data and methods	80
3 Results	82
3.1 Frequency domain analysis	82
3.2 Wind power duration curve	84
3.3 Step change analysis and balancing cost comparison	85
4 Discussion and Conclusion	86
Acknowledgements	88

A	Wind power variability	88
A.1	Data	88
A.2	PSD and Kaimal spectrum approximation	89
A.3	PSD slopes in the inertial subrange	91
A.4	Coefficients of variation	92
A.5	Step change analysis	94
A.6	Correlation of wind power output	95
A.7	Cost comparison with natural gas combustion turbines	97
B	Net load	100
B.1	Frequency domain	100
B.2	Step-change analysis	101
B.3	Correlation analysis	102
	Bibliography	104
4	Optimal investment strategy in a clean energy RD&D program under technological uncertainty	106
1	Introduction	107
2	Previous work	110
2.1	Real options for valuing a single project in the energy sector	110
2.2	Experience and learning curves in the energy sector	111
2.3	Energy R&D in larger climate/economy models	113
2.4	Analysis of energy R&D portfolios	115
2.5	Sources on methods	116
3	Stochastic cost model	117
3.1	Single-factor learning curve	120
3.2	Two-factor learning curve	121
4	Parameter estimates for low-carbon energy technologies	123
4.1	Solar photovoltaics: single-factor learning curve	123
4.2	Carbon capture and sequestration (CCS): two-factor learning curve	124
5	Solution method	132
5.1	Single-factor learning curve	133
5.2	Two-factor learning curve	135
5.3	Two competing technologies	137
6	Results	137
6.1	NPV analysis	137
6.2	Results under uncertainty: single-factor cost model	140
6.3	Results under uncertainty: two-factor cost model	142
7	Discussion and policy implications	153
A	The Bellman equation	156
B	Solution for the deterministic case	157
C	Details of the numerical solution for the stochastic case: single-factor learning curve	159
C.1	Substitution to facilitate numerical solution	159
C.2	Numerical PDE solution: Crank-Nicolson method	160

D	Details of the numerical solution for the stochastic case: two-factor learning curve	163
E	Outline of the solution for two technologies	165
	Bibliography	168
Conclusion		173
	Bibliography	185

List of Figures

1	Economics of compressed air energy storage to integrate wind power: a case study in ERCOT	1
1	Schematic diagram of a CAES plant.	5
2	Sketch of the wind/CAES system with load in Houston.	7
3	Transmission cost model used in the profit optimization.	13
4	Wind/CAES system operation for January 7–14, 2008 with load in Houston.	18
5	Wind-CAES system operation under a \$63/MWh contract price with load in Dallas.	20
A1	Capital cost and 95 % prediction intervals for development of a CAES plant with salt-cavern storage.	36
A2	Capital cost and 95 % prediction intervals for development of a CAES plant with aquifer storage.	37
2	Optimal investment timing and capacity choice for pumped hydropower storage	45
1	(a) Diagram of the investment opportunity: pumped hydropower storage plant in southern Norway and HVDC transmission to Germany. (b) Diagram of reservoirs for the pumped storage plant.	51
2	Structure of the analysis.	52
3	Futures prices for 2010 trades show increasing intraday price difference, which creates arbitrage opportunities for pumped storage.	54
4	Simulated trajectories of expected future short-term price volatility by year for the base scenario.	55
5	NPV, continuation value (expected value of postponing investment), and threshold NPV for the base scenario.	64
6	Continuation value (expected value of postponing investment) and NPVs of each project at the optimal investment time (eighth year of the option lifetime) under the base scenario.	65
7	Continuation value and NPVs of each project at $t = 9$ years for a nominal increase in short-term price volatility of 6 percentage points.	67
8	Reservoir level during two weeks of operation for all five pumped storage capacity alternatives.	69
9	Sensitivity of profit to reservoir size for the 960 MW plant (proposed as the Tonstad upgrade) and 2,400 MW plant (the optimal size in the current study).	70
10	Discount rate sensitivity of option value.	71

3	The effect of long-distance interconnection on wind power variability	76
1	The area spanned by each region and the wind plants it contains.	81
2	(a) PSDs for 2009 wind power output of each region and the aggregate of all four regions plotted on log-log axes. (b) Slopes in the inertial subrange for each region and the interconnected regions for all years of available data and the means over time.	84
3	Duration curve for 2009 wind power output.	85
A1	Kaimal approximation to the PSD of 15-minute wind energy data from a single 63 MW wind plant in Snyder, Texas (ERCOT).	91
A2	Coefficient of variation (standard deviation divided by mean) of 2009 wind power output for each of the four regions, pairs of regions, and the four regions combined.	93
A3	Duration curve for hourly step changes of wind power output for the four regions.	96
A4	Pairwise correlation coefficients between regional wind power outputs tends to decrease as a function of distance between regional wind plant centroids. .	97
A5	PSDs of 2010 CAISO net load with historical wind power output and wind power output amplified by a factor of 10.	101
4	Optimal investment strategy in a clean energy RD&D program under technological uncertainty	106
1	Example cost trajectory of the technology as a function of time, following a GBM. (a) shows a sample path of cost from 0 to 7 years and the 95 % prediction interval for cost from 7 to 15 years. (b) shows the lognormal probability distribution for final cost given knowledge of the cost at 7 years.	119
2	Expected value and 50 % and 90 % prediction intervals for CO ₂ mitigation cost with CCS under base case assumptions for (a) exclusive investment in R&D and (b) exclusive investment in LBD.	130
3	Social benefit of CCS deployment as a function of mitigation cost.	132
4	Optimal RD&D investment strategy as a function of cost C and time t under the current framework.	135
5	Comparison of NPV results and SDP-based results under base case assumptions for the single-factor analysis of solar PV (a and b) and for the two-factor analysis of CCS (c and d).	139
6	Value of investment opportunity in RD&D of solar PV under base case assumptions.	141
7	Comparative statics for the main parameters in the solar PV analysis.	143
7	Comparative statics, continued.	144
8	Optimal RD&D strategy for CCS beginning 15 years before a carbon policy is enacted and CCS is deployed at scale.	145
9	Optimal investment strategy in R&D and LBD of CCS 15 years before deployment for different expected effectiveness and uncertainty parameters for R&D (the corresponding LBD parameters are held constant at their base values). .	147

10	Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the BAU rate of cost reductions.	149
11	Optimal RD&D investment strategy 15 years before the carbon policy/CCS deployment time under different assumptions on the BAU rate of cost reductions.	150
12	(a) Sensitivity of the value of the RD&D program 15 years before the deployment date to the maximum rate of spending. (b) Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the maximum rate of program spending.	151
13	Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the correlation between the effects of R&D and LBD.	152
14	Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different discount rates.	153
A1	Stencil for the Crank-Nicolson method, the numerical PDE solution method employed.	160

List of Tables

1	Economics of compressed air energy storage to integrate wind power: a case study in ERCOT	1
1	Parameters for the wind/CAES dispatch and profit optimization models.	8
2	Decision variables in the wind/CAES profit optimization model.	9
3	Rules for wind/CAES dispatch in ancillary service markets.	16
A1	Physical and cost parameters of transmission lines from Hirst and Kirby (2001).	37
A2	Length, capacity, and total cost of transmission from ERCOT study, and predicted total cost based on Equation (A2).	38
A3	Profit-maximizing CAES expander sizes under the contract price scenario, fractions of wind/CAES system energy output from the CAES plant, and carbon prices to reach cost-parity with a NGCC plant at \$5/MMBTU and \$15/MMBTU gas, for both load centers and all years considered.	40
2	Optimal investment timing and capacity choice for pumped hydropower storage	45
1	Parameters and variables in pumped hydropower storage scheduling.	58
2	Results for a range of growth rates in intraday price differences.	66
3	Sensitivity of base scenario results to investment costs.	68
3	The effect of long-distance interconnection on wind power variability	76
A1	Summary statistics for wind power in the four regions examined in 2009.	88
A2	Summary of data sampling frequencies, gaps, and origins.	89
A3	Log PSD slopes for frequencies from $(24 \text{ h})^{-1}$ to $(2 \text{ h})^{-1}$, in which a single wind plant follows the Kolmogorov spectrum (slope = -1.67).	92
A4	Correlation coefficients between 2009 (2008) hourly wind power outputs.	95
A5	Wind-load correlation coefficients for 2009 (2008) hourly data.	102
4	Optimal investment strategy in a clean energy RD&D program under technological uncertainty	106
1	Parameters and variables in the development/deployment option valuation for a single-factor cost model.	121
2	Parameters and variables in the RD&D investment model for the two-factor cost model.	122
3	Comparison of parameter estimates based on three expert elicitations.	128

Introduction

Working Group I of the Intergovernmental Panel on Climate Change (IPCC) found with “very high confidence that the global average net effect of human activities since 1750 has been one of warming” and that CO₂ emissions are the chief contributor to the greenhouse effect (IPCC, 2007). The electricity sector accounts for about 40 % of total annual CO₂ emissions in the U.S., making it a primary target of climate policies to reduce anthropogenic greenhouse gas emissions (EIA, 2008).

Decarbonization of the electricity sector in the near term will require the large-scale integration of existing low-carbon energy technologies. The most prominent of these technologies, wind and solar power, are variable and intermittent and thus pose challenges as they are deployed at increasing scale. Nevertheless, the U.S. Department of Energy (DOE) found that if these and other challenges are met, 20 % of U.S. electricity demand could be fulfilled by wind power by 2030 (DOE, 2008). Over the longer term, in addition to integration of existing low-carbon energy technologies, more severe cuts in greenhouse gas emissions will require the development of new technologies in order to meet electricity demand at acceptable cost and reliability standards.

This thesis addresses both the near-term and the longer-term technology problems, by analyzing proposed methods of wind power integration (Papers 1–3) and by developing a method for optimizing R&D expenditure in low-carbon energy technologies with different cost and risk characteristics (Paper 4).

Wind power is one of the most developed and least costly low-carbon energy technologies. Both U.S. State Renewables Portfolio Standards and a federal production tax

credit have driven rapid expansion of wind power capacity in the U.S.: between 2006 and 2011, installed capacity grew from 11.6 GW to 47.0 GW (WWEA, 2012). To realize the DOE's scenario of 20 % wind electricity by 2030, total installed capacity would reach 305 GW.

At such large scales, the variability and intermittency of wind power output must be mitigated. Wind power variability occurs at all time scales, from seasonal to second-to-second, and output can all but cease for hours or even days at a time. Forecasting methods are imperfect, so system operators must often respond to steep wind power fluctuations with little warning. Currently, wind power variability is most often mitigated by ensuring that enough quick-ramping natural gas capacity is online to respond quickly to worst-case wind power fluctuations. However, natural gas plants emit significant CO₂ and can do nothing to avoid wind power curtailments at times when production exceeds demand in a balancing area. Using wind power to achieve deeper cuts in greenhouse gas emissions will therefore require other solutions to mitigate wind power variability.

One proposed solution is large-scale energy storage, which could increase the availability of wind power and help avoid both curtailment and forced ramping of baseload generators. Many available storage technologies, such as batteries, flywheels, and thermal energy storage, have unsuitable power-to-energy ratios or are too costly to use for large-scale wind power integration. Storage technologies with appropriate ranges of power and energy capacities to balance large-scale wind are compressed air energy storage (CAES) and pumped hydropower storage (PHS).

Paper 1 of this thesis examines the economic viability of CAES as a method to balance wind in Texas. Using a model of a CAES plant with historical hourly prices from the Electric Reliability Council of Texas (ERCOT), and accounting for transmission costs and the geological constraints of CAES, this chapter shows that CAES was not a cost-effective means of smoothing wind power output under 2007–2009 market conditions. This result holds when externalities due to criteria air pollutants and CO₂ emissions are priced.

Paper 2 examines a similar wind integration problem faced in Europe. Like the U.S., Germany has adopted policies to encourage the expansion of wind capacity. The German Advisory Council on the Environment (SRU) projects that wind capacity could grow from 27 GW in 2010 to 113 GW in 2050 (SRU, 2010). Germany has a single CAES plant, which is currently used for balancing and load shifting, as well as limited PHS, which is already used to its full capacity and has little possibility for additional builds. Norway, which is tied to Germany through limited HVDC transmission capacity, has an extensive and flexible hydropower system that could provide balancing services to aid German wind integration. The SRU found that for Germany to achieve 100 % renewable electricity by 2050, expanding the HVDC interconnection in order to use the Norwegian hydropower system is more cost effective than the domestic alternative of a buildout of CAES plants.

Paper 2 analyzes the decision to build a PHS retrofit to a hydropower plant in southern Norway to export wind-balancing services to Germany, from the perspective of a Norwegian hydropower producer. Since the effect of increased wind capacity on prices in the German electricity market is uncertain and subject to change over time, real options rather than net present value (NPV) is used to evaluate the investment opportunity and optimal investment decision. Unlike NPV, real options accounts for flexible future decision making after new information is revealed and values the ability to wait rather than invest immediately. Paper 2 finds that immediate investment in the PHS retrofit would be profitable, but that the likely increase in price volatility in the German market and the resolution of uncertainty on the price effect of wind promotes postponing the investment.

Aside from large-scale energy storage, a second proposed method of smoothing wind power is the interconnection of wind plants across long distances. The extent to which interconnection smooths wind power output depends on the geographical separation of wind plants and factors such as local topography and atmospheric effects, and it varies according to the frequency of the wind power fluctuations. High-frequency fluctuations (on the order of seconds) may be smoothed within a single large wind plant, since they are

determined by small-scale turbulence that affects only subsets of the turbines within the wind plant and is uncorrelated with the local turbulence affecting other areas. In contrast, since the entire wind plant tends to experience the same larger-scale effects such as fronts and sea breezes, lower-frequency fluctuations (on the order of hours to days) are not smoothed within the wind plant but may be smoothed by interconnection with a distant wind plant.

The power spectral density (PSD) measures the strength of wind power fluctuations at different frequencies and is an appropriate tool to analyze the smoothing effect of interconnection. Katzenstein et al. (2010) performed the first analysis to examine this effect with PSDs and found that connecting up to four wind plants within Texas substantially smoothed wind power output but that interconnecting additional wind plants showed rapidly diminishing returns.

Paper 3 of this thesis both extends the PSD-based analysis and uses traditional time-domain techniques to examine the extent to which longer-distance interconnection, across much of the western and mid-western U.S. would further smooth wind power output. Results show that long-distance interconnection would yield negligible additional benefit compared with interconnection within Texas alone. An approximate cost analysis shows that the smoothing effect of long-distance interconnection could be achieved at lower cost with local gas turbines than with increased transmission capacity to connect the wind plants, even when air pollutant externalities of the gas plants are priced.

Papers 1 and 3 suggest that two commonly proposed methods of smoothing wind power output, long-distance interconnection and CAES, are unlikely to be cost-effective means of enabling drastically expanded wind capacity in the U.S. in the near future. Paper 2 shows that Norway's exceptional hydropower system could provide a cost-effective means of integrating wind power in Germany. Since the U.S. does not have access to a similar resource, the results of Paper 2 cannot be extended to the U.S. The dearth of cost-effective, low-carbon solutions for integrating wind power at scale in the U.S. points to the need for

further technological development in order to achieve more extensive decarbonization of the electricity sector.

In models that analyze future emissions trajectories to achieve climate stabilization, technological change is often represented as deterministic. These models often focus on sophisticated representations of other processes and interactions in the climate-economy system. Nevertheless, their neglect of uncertainty in technological change results in energy technology strategies that do not allow for the abandonment of technologies that do not achieve projected cost reductions, for increased deployment of technologies that exceed expected improvements in cost and performance, or for the ability to invest initially in very early-stage technologies that will likely not reach commercialization but have a small probability of achieving great success.

Paper 4 examines optimal government investment strategies to promote technology development under uncertainty. A real options framework is used to analyze an investment decision in the development and deployment of a technology whose cost evolves stochastically but can be reduced in expectation through R&D spending and promotion of learning-by-doing. Results yield insight into the effect of uncertainty and other parameters on the best ways for the government to promote the development of greenhouse gas mitigating energy technologies in anticipation of future climate policy.

Taken as a whole, this thesis addresses challenges in the decarbonization of the electricity sector, from facilitating the integration of an existing low-carbon energy technology—wind power—to strategies for promoting the development of new low-carbon energy technologies under technological uncertainty. In doing so, it contributes to the design of better policies for sustainability and climate change mitigation in the electricity sector.

Bibliography

- DOE (2008). 20% wind energy by 2030: Increasing wind energy's contribution to U.S. electricity supply. U.S. Department of Energy, Energy Efficiency and Renewable Energy.
- EIA (2008). Emissions of greenhouse gases report. Technical Report DOE/EIA-0573(2008), Energy Information Administration.
- IPCC (2007). Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.
- Katzenstein, W., Fertig, E., and Apt, J. (2010). The variability of interconnected wind plants. *Energy Policy*, 38(8):4400–4410.
- SRU (2010). Climate-friendly, reliable, affordable: 100 % renewable electricity supply by 2050. German Advisory Council on the Environment (SRU).
- WWEA (2012). World Wind Energy Association.
<http://www.wwindea.org/home/index.php>.

Paper 1

Economics of compressed air energy storage to integrate wind power: a case study in ERCOT

Abstract

Compressed air energy storage (CAES) could be paired with a wind farm to provide firm, dispatchable baseload power, or serve as a peaking plant and capture upswings in electricity prices. This paper presents a firm-level engineering-economic analysis of a wind/CAES system with a wind farm in central Texas, load in either Dallas or Houston, and a CAES plant whose location is profit-optimized. With 2008 hourly prices and load in Houston, the economically optimal CAES expander capacity is unrealistically large — 24 GW — and dispatches for only a few hours per week when prices are highest; a price cap and capacity payment likewise results in a large (17 GW) profit-maximizing CAES expander. Under all other scenarios considered the CAES plant is unprofitable. Using 2008 data, a baseload wind/CAES system is less profitable than a natural gas combined cycle (NGCC) plant at carbon prices less than \$56/tCO₂ (\$15/MMBTU gas) to \$230/tCO₂ (\$5/MMBTU gas). Entering regulation markets raises profit only slightly. Social benefits of CAES paired with wind include avoided construction of new generation capacity, improved air

This paper was published as Fertig, E. and Apt, J. (2011). Economics of compressed air energy storage to integrate wind power: a case study in ERCOT. *Energy Policy*, 39:2330–2342.

quality during peak times, and increased economic surplus, but may not outweigh the private cost of the CAES system nor justify a subsidy.

1 Introduction

Renewable energy currently comprises 9 % of the United States' net electric power generation (EIA, 2009a). Twenty-nine states' enactment of Renewable Portfolio Standards (RPS) (Database of State Incentives for Renewables and Efficiency, 2010) and the possibility of a Federal RPS suggest that the nationwide share of renewables in the electricity sector could double by 2020 (Waxman and Markey, 2009).

With high penetration of renewables, variability of power output increases the need for fast-ramping backup generation and increases the need for reliable forecasting. Pairing a variable renewable generator with large-scale electricity storage could provide firm, dispatchable power and alleviate the costs and stability threats of integrating renewable energy into power grids. Although it has been argued elsewhere (e.g., DOE, 2008) that dedicated storage is not a cost-effective means of integrating renewables, the cost savings from constructing a small transmission line with a high capacity factor instead of a large transmission line with a low capacity factor could in some cases be sufficient to justify building a dedicated CAES plant.

Utility-scale electricity storage has not been widely implemented: batteries remain prohibitively expensive and pumped hydroelectric storage is feasible only in locations with suitable hydrology. An emerging large-scale storage technology is compressed air energy storage (CAES), in which energy is stored in a pressure gradient between ambient air and an underground cavern. Two CAES plants are in operation: one in Huntorf, Germany and the other in McIntosh, Alabama, USA. FirstEnergy, the Iowa Association of Municipal Utilities, and PG&E are building new CAES systems, the last with the help of federal funding (Haug, 2006; Leidich, 2010; LaMonica, 2009). The New York State Energy

Research and Development Authority (NYSERDA) has commissioned an engineering study for a possible CAES plant in New York (Hull, 2008), and Ridge Energy Storage has proposed a CAES system in Matagorda, Texas (Ridge Energy Storage, 2005).

Denholm and Sioshansi (2009) compare the costs of (1) a co-located wind farm/CAES plant with an efficiently-used low-capacity transmission line to load and (2) a CAES plant located near load that uses inexpensive off-peak power for arbitrage, with a higher-capacity, less efficiently-used transmission line from the wind farm. Avoided transmission costs for co-located CAES and wind in ERCOT outweigh the higher arbitrage revenue of load-sited CAES at transmission costs higher than \$450 per GWm. Although actual transmission cost data vary greatly, many transmission projects cost more than \$450/GWm and would warrant wind-CAES co-location (Denholm and Sioshansi, 2009).

Greenblatt et al. (2007) model CAES and conventional gas generators as competing technologies to enable baseload wind power. The wind/CAES system had the highest levelized cost per kWh at an effective fuel price (the sum of natural gas price and greenhouse emissions price) of less than \$9/GJ (\$8.5/MMBTU). The wind/CAES system had a lower short-run marginal cost, rendering it competitive in economic dispatch and at greenhouse emissions prices above \$35/tC_{equiv} (\$9.5/tCO₂) the wind/CAES system outcompetes coal for lowest dispatch cost (Greenblatt et al., 2007).

(DeCarolis and Keith, 2006) optimize the use of simple and combined cycle gas turbines, storage, and widely-distributed wind sites to enable large-scale integration of distant wind resources. Diversifying wind sites produces benefits that outweigh the ensuing transmission costs, and smoothing due to wind site diversity renders CAES economically uncompetitive at carbon prices below \$1000/tC (\$270/tCO₂). For a single wind site, CAES is cost effective at \$500/tC (\$135/tCO₂).

Each of the above studies uses simulated wind power data or a power curve applied to measured wind speed data. Denholm and Sioshansi (2009) use hourly electricity price data from Independent System Operators (ISOs), while the other two studies examine the

cost-effectiveness of storage for wind integration and make no assumptions about electricity price. We examine the economic and technical feasibility of a wind/CAES system in Texas, using wind power data from a large wind farm in the central part of the state, hourly electricity prices from the Electric Reliability Council of Texas (ERCOT), and monthly gas prices to Texas electric utilities. The model is further constrained by the underlying geology suitable for a CAES cavern. CAES size, transmission capacity, and dispatch strategy are optimized for profit. This research differs from previous work in that it examines CAES as a means of wind power integration in a specific location and incorporates a multiparameter optimization of the wind-CAES system, transmission, and dispatch strategy.

Section 2 describes the mechanics of CAES and the two CAES plants currently in operation. Section 3 describes the wind/CAES system modeled in the current study, and Section 4 explains how the underlying geology and concerns about transmission congestion influence the siting of CAES. Section 5 provides the sources of the data used in the study and describes the function of ERCOT balancing energy and regulation markets. Section 6 provides the cost models used for the CAES system and transmission lines. Section 7 describes the heuristic dispatch strategies and profit optimization models for the wind/CAES system in the energy and regulation markets, Section 8 presents results, and Section 9 provides discussion and policy implications.

2 CAES mechanics and extant plants

Figure 1 is a schematic diagram of a CAES plant, which is analogous to a natural gas generator in which the compression and expansion stages are separated by a storage stage. In a conventional gas plant, 55–70 % of the electricity produced is used to compress air in preparation for combustion and expansion (Gyuk and Eckroad, 2003). In a CAES plant, air is compressed with electricity from a wind farm or off-peak electricity from the grid, so the heat rate is about 4300 BTU/kWh compared with 6700 BTU/kWh for a high-efficiency

natural gas combined cycle turbine (Klara and Wimer, 2007). All designs demonstrated to date combust natural gas, but conceptual adiabatic designs reheat the expanding air with the stored heat of compression and do not use gas.




Figure 1: Schematic diagram of a CAES plant. In the compression stage, CAES uses electricity to compress air into a pressure-sealed vessel or underground cavern, storing energy in a pressure gradient. The air is cooled between each compressor to increase its density and aid compression. To generate electricity, the air is mixed with natural gas and expanded through combustion turbines.

Two CAES plants are currently operational: one in Huntorf, Germany, and one in McIntosh, Alabama, USA. The Huntorf plant was completed in 1978 and is used for peak shaving, to supplement the ramp rate of coal plants, and more recently to mitigate wind power variability. The McIntosh plant was completed in 1991 and is used for storing off-peak baseload power, generating during peak times, and providing spinning reserve (see Appendices 1 and 2) (Gardner and Haynes, 2007).

In a new, less costly, and more efficient design proposed by the Electric Power Research Institute (EPRI), only the low-pressure turbine is combustion-based; the high-pressure turbine is similar to a steam turbine. This difference partially accounts for the lower heat rate of the EPRI design (3800 BTU/kWh) (Schainker, 2008). This study uses technical parameters of the EPRI design.

A CAES plant could reduce wind power curtailment by storing wind energy in excess

of transmission capacity, thereby deferring transmission upgrades and allowing system operators to avoid curtailment payments to wind farm owners. CAES systems have fast ramp rates that match fluctuations in wind power output. A CAES plant with one or more 135 MW generators starts up in 7–10 minutes and once online ramps at about 4.5 MW per second (or 10 % every 3 seconds) (EPRI, 2004). In the compression phase, a CAES plant starts up in 10–12 minutes and ramps at 20 % per minute, which is fast enough to smooth wind power on the hourly timescales modeled in the current study. The fast ramp rate of a CAES expander compared with that of a natural gas turbine (7 % per minute (Western Governors' Association, 2002)) is possible because the compression stage of the CAES cycle is already complete when the CAES ramps.

A wind/CAES system could act as a baseload generator in place of coal and nuclear plants, or could be dispatched as a peak-shaving or shoulder-load plant. The operating flexibility of CAES also enables a wind/CAES system to provide ancillary services such as frequency regulation, spinning reserve, capacity, voltage support, and black-start capability (Gyuk, 2004). Previous research has shown that pumped hydroelectric storage can decrease the total cost of ancillary services by 80 % and generate significant revenue in a simulated market in Tennessee Valley Authority (TVA) (Perekhodtsev, 2004); a quick-ramping, large-capacity CAES system could provide similar benefit. Here we examine the profitability of CAES in up- and down-regulation markets as well as the balancing energy market.

3 Wind/CAES system model

3.1 Physical design

The wind/CAES system is modeled as a 1300 MW wind farm (the combined nameplate capacity of Sweetwater and Horse Hollow wind farms, 16 km apart in central Texas), a wind-CAES transmission line, a CAES plant, and a CAES-load transmission line.

Pattanariyankool and Lave (2010) observe that the economically efficient transmission capacity from a wind farm is often well below the nameplate capacity of the wind farm. Parameters in the economic optimization include the lengths (L_W and L_C) and capacities (T_W and T_C) of both transmission lines as well as the CAES expander capacity (E_E), compressor capacity (E_C), and storage cavern size (E_S) (Figure 2). The cost and optimal location of the CAES plant are also contingent on the underlying geology, as discussed below. Relevant parameters and variables for the wind/CAES system operation and profit models are shown in Tables 1 and 2.

The wind farm is assumed to already exist so its cost is not modeled.




Figure 2: Sketch of the wind/CAES system with load in Houston. With load in Dallas, aquifers underlie the entire 320 km distance between wind and load.

4 Siting the CAES plant

We assume fixed locations of the wind farm in central Texas and load either 530 km away in Houston or 320 km away in Dallas. The location of a CAES plant is optimized for profit subject to the geological constraints discussed below.

CAES is feasible in three broad types of geology: solution-mined salt caverns, aquifers of sufficient porosity and permeability, and mined hard rock caverns (Succar and Williams, 2008). Due to the disproportionately high cost of developing hard rock caverns, we do not consider them here.

The two operational CAES plants in Alabama and Germany both use solution-mined

Table 1: Parameters for the wind/CAES dispatch and profit optimization models. Subscript i denotes a variable that changes hourly. Costs are adjusted to 2009\$ with the Chemical Engineering Plant Cost Index (Lozowski, 2009).

Parameter	Symbol	Base Value	Unit	Reference
Marginal cost of generating wind power	MC_w	\$0.00	\$/MWh	
Wind energy output	w_i		MWh	ERCOT, 2009b
Hourly zonal electricity price	p_i		\$/MWh	ERCOT, 2009a
Hourly up-regulation price	u_i		\$/MW	
Hourly down-regulation price	d_i		\$/MW	
Cost of gas	g_i		\$/MMBTU	EIA, 2009b
Energy ratio of CAES system	ER	.7	kWh in/kWh out	Schainker, 2008
Heat rate of CAES system	HR	3800	BTU/kWh	Schainker, 2008
Heat rate of CAES as a gas turbine	HR_{gas}	10,000		
Blended cost of capital	dr	.10		
30-year annualization factor	A	$dr/(1 - (1 + dr)^{30})$		
Baseline cost of CAES system	C_{CAES}	1700/2000	\$/kW	Schainker, 2008
Marginal cost of CAES expander	C_E	560	\$/kW	Greenblatt et al., 2007
Marginal cost of CAES compressor	C_C	520	\$/kW	Greenblatt et al., 2007
Marginal cost of storage cavern capacity	C_S	1.5	\$/kWh	Schainker, 2008
Energy CAES system can store, hour i	x_i		MWh	
Energy CAES system can generate, hour i	y_i		MWh	
Energy state of cavern, hour i	s_i		MWh	
Energy discharged from storage, hour i	r_i		MWh	
Total energy sold, hour i	e_i		MWh	

Table 2: Decision variables in the wind/CAES profit optimization model.

Decision variable	Symbol	Unit
Zonal electricity price below which wind energy is stored	p_s	\$/MWh
Zonal electricity price above which CAES is discharged	p_d	\$/MWh
Length of wind-CAES transmission line	L_W	km
Length of CAES-load transmission line	L_C	km
Capacity of wind-CAES transmission line	T_W	MW
Capacity of CAES-load transmission line	T_C	MW
CAES expander power	E_E	MW
CAES compressor power	E_C	MW
CAES storage capacity (expander hours)	E_S	MWh

salt caverns for air storage. These structures are advantageous for CAES due to the low permeability of salt, which enables an effective pressure seal, and the speed and low cost of cavern development. The caverns are formed by dissolving underground halite (NaCl) in water and removing the brine solution. The CAES plant injects and removes air through a single well connecting the salt cavern and turbomachinery. Salt that can house a CAES cavern occurs in two general forms: bedded and domal. Domal salt is purer and thicker than bedded salt and therefore superior for CAES caverns, but specific sites in bedded salt can be suitable for CAES as well (Hovorka, 2009).

Underground storage for CAES is also feasible in an aquifer-bearing sedimentary rock of sufficient permeability and porosity that lies beneath an anticline of impermeable caprock to stop the buoyant rise of air and impede fingering (Succar and Williams, 2008). A bubble in the aquifer, developed by pumping air down multiple wells, serves as the air storage cavern. The Iowa Stored Energy Project (ISEP), a wind/CAES system under construction in Dallas Center, IA, will use an aquifer for underground storage (Haug, 2006).

Domal salt is located in the East Texas Basin, South Texas Basin, and Gulf Coast Basin surrounding Houston, as well as the Delaware and Midland Basins of West Texas. Bedded salt underlies much of the eastern part of the state, from the Gulf Coast to 160–240 km inland (Hovorka, 2009). Aquifers possibly suitable for a CAES cavern underlie the western and central parts of the state, including Dallas (Succar and Williams, 2008).

Aquifer CAES is dependent on highly localized aquifer parameters such as porosity, permeability, and caprock composition and geometry, so generalizing on the geographic extent of suitable aquifers is impossible. Appendix 3 contains further information on CAES geology in Texas.

Siting the CAES near wind enables a high-capacity wind-CAES transmission line that minimizes wind power curtailment due to transmission constraints as well as a lower-capacity CAES-load line that the system fills efficiently. Wind-sited CAES, however, compromises the ability of the CAES system to buy and sell electricity optimally from the grid because the lower-capacity CAES-load line is often congested with wind power (Denholm and Sioshansi, 2009). Siting the CAES near load enables larger CAES-load transmission capacity, thereby increasing the potential for arbitrage. Load-sited CAES can also store and supply slightly more power to the grid because transmission losses are incurred before the CAES. Sullivan et al. (2008) found that “the capacity, transmission loss, and congestion penalties evidently outweighed the cost savings of downsizing transmission lines,” making load-sited CAES economically superior.

5 Data and energy markets

Hourly zonal electricity prices are from the ERCOT Balancing Energy Services (BES) market for 2007, 2008, and 2009 (ERCOT, 2009a). Although most energy in ERCOT is traded bilaterally, 5–10 % is traded on the BES market that ERCOT administers for the purpose of balancing generation and load. BES prices are thus proxies for locational marginal prices (LMPs) of electricity (Denholm and Sioshansi, 2009). ERCOT is currently divided into four pricing zones: West, North, South, and Houston. Sweetwater and Horse Hollow wind farms are located in ERCOT West, which experiences frequent negative prices due to wind power congestion that a large CAES system would help relieve. We use ERCOT Houston prices if the CAES plant is sited in Houston and ERCOT North prices if

the CAES is in Dallas. ERCOT plans to switch its primary energy market to nodal pricing within the next few years, allowing prices to better reflect local market conditions (ERCOT, 2008).

In addition to the BES market, ERCOT administers hourly markets for up-regulation and down-regulation. A generator bids capacity into a regulation market 24 hours in advance and can edit the bid until an hour in advance. The generator is paid the product of its accepted capacity bid and the market-clearing price of the regulation market, plus the BES price for the additional energy generated or curtailed. Hourly prices for up-regulation and down-regulation in ERCOT in 2008 and 2009 were obtained from a commercial data provider.

Fifteen-minute wind energy output data from Sweetwater and Horse Hollow wind farms for 2008 and 2009 were obtained from ERCOT's website and summed to produce hourly data (ERCOT, 2009b). To approximate 2007 power output from the two wind farms, system-wide ERCOT wind power data was scaled to the appropriate nameplate capacity (in 2008, power output from Sweetwater and Horse Hollow was highly correlated with aggregate ERCOT wind output ($R^2 = 0.96$)). The data were affected by wind curtailment, which occurred on 45–50 % of the days from January to August 2008 at an average amount of 140–150 MW. Since the installed wind capacity in ERCOT at that time was 7100 MW, curtailment of Sweetwater/Horse Hollow would have averaged, at most, approximately 2 % of capacity. Curtailment would decrease the calculated profit of both the wind/CAES system and the standalone wind farm, and generally tend to increase the profitability of the former (since the extra energy could be sold when prices are high and not only when the wind farm produced it). Our analysis does not account for this effect, which we believe to be small.

Monthly natural gas prices for the electric power industry in Texas in 2007–2009 are from the United States Energy Information Administration (EIA, 2009b).

6 Wind/CAES system cost models

6.1 CAES plant

Equation (1) shows the estimated total capital cost of a large CAES system in a salt cavern. The cost model begins with the EPRI estimate for a 346 MW expansion/145 MW compression/10 storage – hour CAES plant (C_{CAES}), plus incremental costs per MW of expander capacity (C_E), compressor capacity (C_C), and storage cavern capacity (C_S) (Greenblatt et al., 2007; Schainker, 2008). The model is then adjusted upward by a factor of 2.3 to conform to recent industry estimates (Gonzales, 2010; Leidich, 2010). The cost of a CAES plant larger than 1 GW is adjusted from \$1700/kW for a 2 GW plant, after estimates for the anticipated Norton plant. The cost of a smaller CAES plant is adjusted from \$2000/kW for a 500 MW plant. Costs are inflation-adjusted to 2009\$ with the Chemical Engineering Plant Cost Index (CEPCI) (Lozowski, 2009). The cost of CAES with aquifer storage is modeled as 30 % higher, which reflects the difference in average capital cost per kW generation capacity between CAES plants in the two geologies according to data on a possible CAES system in New York (Swensen and Potashnik, 1994), EPRI reports, and data from extant and upcoming plants (Haug, 2004; Hydrodynamics Group, 2009; Marchese, 2009) (see Appendix 4).

$$\begin{aligned} \text{Cost of CAES} = & C_{CAES} \cdot 2000 + C_E \cdot (E_E - 2000) + C_C \cdot (E_C - 1500) \\ & + C_S \cdot (1000 \cdot E_E \cdot E_S - 2 \times 10^7) \end{aligned} \quad (1)$$

6.2 Transmission

Equation (2) models the capital cost of transmission as a function of lengths in km (L_W and L_C) and capacities (T_W and T_C) in MW.

$$\text{Cost of transmission} = 14266 \cdot (L_W \cdot T_W^{0.527} + L_C \cdot T_C^{0.527}) \quad (2)$$

Figure 3 shows a plot of the transmission cost model in dollars per GWm as a function of MW capacity. The model was derived by fitting an exponential curve to transmission costs from planning studies and reflects an economy of scale in which the cost per GWm decreases as power capacity increases (Hirst and Kirby, 2001). Although transmission costs vary widely and are highly dependent on terrain, land use patterns, and other site-specific factors, this function provides a cost estimate that is consistent with past projects (see Appendix 5).

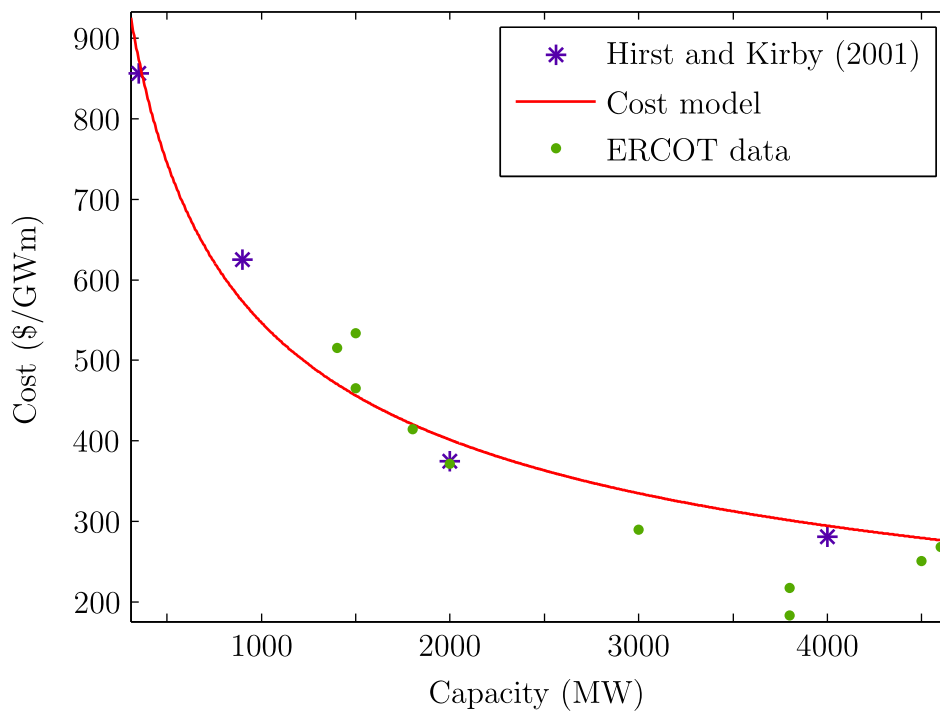


Figure 3: Transmission cost model used in the profit optimization. Data are from Hirst and Kirby (2001) and an ERCOT transmission planning study that assessed the costs of wind integration (ERCOT, 2006). The model fits the ERCOT data with $R^2 = 0.72$.

7 Wind/CAES heuristic dispatch strategies and hourly profit models

7.1 Balancing Energy Services (BES) market

In the BES market, the wind/CAES system is operated to maximize profit based on p_i , price of electricity at hour i for the ERCOT zone in which the CAES system is located. If p_i is less than the marginal cost of generating wind power (MCW, taken as 0), the model stores wind energy up to capacity and curtails the excess. If p_i is greater than MCW but less than the storage threshold price p_s , the system stores wind energy to capacity and sells the excess. If p_i is greater than p_s but less than the dispatch threshold price p_d , the system sells wind power and leaves the CAES system idle. If p_i is greater than p_d , energy is generated from the CAES plant. The prices p_s and p_d are decision variables in the profit optimization, while MCW is an economic property of the wind farm. Since the amount of wind power produced by the wind/CAES system is equal to that produced by the standalone wind farm, the production tax credit for wind power and the sale of renewable energy credits does not affect the difference in profitability between the two and was not included in the analysis. Appendix 6 contains further description of the model. Equation (3) shows the total amount of energy delivered by the wind/CAES system in the hourly energy market in hour i .

$$e_i = \begin{cases} 0 & \text{if } p_i < MCW \\ \min(T_W, T_C, w_i - x_i) \text{ if } w_i > x_i, \text{ else } 0 & \text{if } MCW < p_i < p_s \\ \min(T_W, T_C, w_i) & \text{if } p_s < p_i < p_d \\ \min(T_W + y_i, T_C, w_i + y_i) & \text{if } p_d < p_i \end{cases} \quad (3)$$

Yearly profit, including annualized capital costs, is shown in Equation (4). Revenue is

calculated as the product of electricity sold and the current zonal price, summed over all hours of the year. Operating cost is the cost of gas used by the CAES system. Costs of the CAES system and transmission lines are modeled according to Equations (2) and (3) and are annualized with a 10 % discount rate and 30-year project lifetime:

$$\Pi = \sum_i (p_i \cdot e_i - g_i \cdot r_i \cdot HR) - A \cdot (\text{CAES cost} + \text{transmission cost}) \quad (4)$$

Profit is maximized for three electricity price scenarios: hourly BES prices, the prices capped at \$300/MWh with a \$100/MWd capacity payment, and a constant contract price equal to the mean BES price for the year. The price-cap scenario simulates the case in which a price cap plus capacity payment, instead of price spikes, signals the need for investment in new capacity, and is meant to generalize our results beyond the current ERCOT case. Since ERCOT currently has no capacity market, the value of \$100/MWd is based on the PJM capacity market clearing prices of \$40.80 to \$237.33/MWd for 2007–2009 (mean: \$159.68/MWd), and the observation that prices in the PJM capacity market for these years overrepresented the need for additional capacity and did not provide a cost-effective means of promoting system reliability (Wilson, 2008).

For the contract price scenario, the price-threshold dispatch strategy is infeasible so profit is maximized with the constraint that the capacity factor of the CAES-load transmission line be 80 %, which is approximately representative of a baseload generator. The constraint on transmission capacity factor is not meant to simulate an actual contract; it is imposed only to determine the size and cost of a CAES plant for a wind/CAES system acting as a baseload generator. For all scenarios, we compare results using data from 2007 to 2009.

A simulated annealing algorithm was used to optimize yearly profit (Equation (4)) with decision variables of p_s and p_d (determined monthly), T_C , T_W , L_C , L_W , E_C , E_E , and E_S (see Appendix 7) (Goffe et al., 1994).

7.2 Regulation and balancing energy markets

A separate model allows the wind/CAES system to bid into the up-regulation and down-regulation markets in addition to the BES market. During the morning ramp, the average down-regulation price is greater than the average up-regulation price; during the evening ramp down, the opposite is true. We define a bidding strategy based on four progressively greater daily time thresholds, h_1 through h_4 , as described in Table 3.

Table 3: Rules for wind/CAES dispatch in ancillary service markets. Parameter h denotes the hour of the day, while h_i , $i = 1, 4$, denote thresholds that are decision variables in the optimization. All of the parameters h have integer values from 1 to 24. HR_{gas} denotes the heat rate of CAES when run as a natural gas generator.

Condition	Market into which system bids	Hourly marginal profit
$h_1 < h < h_2$	Down-regulation.	$di \cdot \min(E_C, T_C) + 0.2 \cdot p_i \cdot \min(E_C, T_C)$
$h_3 < h < h_4$	Up-regulation.	$u_i \cdot E_C + 0.2 \cdot E_C \cdot p_i d_i \cdot HR +$ $\max(E_E - (w_i + y_i), 0) \cdot HR_{gas} \cdot g_i / 1000$
$h < h_1,$ $h_2 < h < h_3,$ or $h > h_4$	BES.	$e_i \cdot p_i d_i \cdot g_i \cdot HR / 1000$

When bidding into the BES market, the system uses the same strategy as in the BES-only scenario above with p_s equal to the 33rd percentile price and p_d equal to the 67th percentile price. Since up-regulation and down-regulation procurements in ERCOT are on the order of 1 GW, we fix the CAES expander and compressor capacities at 450 MW to adhere to the price-taker assumption. We assume that the system bids 450 MW into the up-regulation or down-regulation markets and is deployed 90 MW (consistent with average regulation deployment as a fraction of procurement in ERCOT (2010)). When up-regulation is deployed, any wind energy generated up to 90 MWh is transmitted to load, and the CAES plant provides the remainder. If the CAES cavern is depleted, the CAES acts as a natural gas-fired generator with a higher heat rate. When down-regulation is deployed, the CAES cavern stores 90 MWh. If the cavern is full, the compressor is run and exhausted to the ambient air. The 49 decision variables correspond to the four time thresholds optimized monthly and the capacity of the wind-CAES transmission line.

8 Results

8.1 Balancing energy market — zonal prices

Using 2008 zonal prices with load in Houston, the profit optimization results in a CAES with an unrealistically large 24 GW expander that dispatches infrequently (Figure 4). The optimal price thresholds for the dispatch strategy, p_s and p_d , were such that the wind/CAES system stored wind energy 91 % of the time, sold only wind energy 6 % of the time, and discharged the CAES 3 % of the time. This system would earn \$900 million in the BES market, and a standalone wind farm with a single wind-load transmission line would earn \$245 million. Lower expander capacities result in less energy sold during price spikes and therefore lower profit despite the additional cost of expander power. Due to the nature of the objective function, the heuristic optimization algorithm may have failed to find a larger CAES system that could generate even more profit; however, 24 GW is an unrealistically large plant and the profit-generating price spikes are of unpredictable magnitude and frequency, such that a larger CAES system that generates more profit under this strategy is not a valuable result. The economically optimal location for the CAES plant is close to load in Houston, enabling a shorter and less costly high-capacity transmission line from CAES to load. Air storage is in a solution-mined salt cavern, the less expensive of the two geologies considered (see Figure 5).

For all other zonal price scenarios (load in Houston for 2007 and 2009, and load in Dallas for 2007–2009), no CAES system could capture annual revenue that compensates for its annualized capital cost, so the optimal size of all CAES components is 0. The higher cost of building CAES in an aquifer near Dallas or the wind site instead of in a salt cavern near Houston contributes to the unprofitability of CAES with load in Dallas. These results suggest that the profitability of CAES in Houston given 2008 data is due to anomalous price spikes.

A profit-maximizing energy trader would not use constant storage and discharge




Figure 4: Wind/CAES system operation for January 7–14, 2008 with load in Houston.

threshold prices as an operations strategy: a high BES price in the morning, for example, could cause the trader to anticipate an even higher afternoon peak and wait to discharge the storage, and the same price at night could motivate the trader to discharge the storage immediately in anticipation of falling prices and increased wind power output to refill the storage. The current dispatch algorithm would likely generate less profit than a strategy applied by an energy trader.

8.2 Balancing energy market — price cap of \$300/MWh plus capacity payment of \$100/MWd

For 2008 prices with load in Houston, the optimal CAES expander size is 17 GW and the system generates \$300 million in profit, compared with \$245 million for a standalone wind

farm. For 2007 and 2009 prices, the optimal CAES expander size is 6 GW and 3 GW respectively, and generates negative profit. With load in Dallas, the optimal CAES size for all years is zero. Once again, the higher cost of building CAES in an aquifer rendered the Dallas CAES system unprofitable.

8.3 Balancing energy market — contract price

With a contract price and a set capacity factor of 80 % for the CAES-load transmission line, no wind/CAES system generated more profit than a standalone wind farm. The highest profit generated for a system with load in Houston was \$110 million (with a 300 MW CAES expander and 480 MW CAES-load transmission line), compared with \$245 million for the standalone wind farm. The highest profit for a system with load in Dallas was \$70 million in 2008 (for a 260 MW expander and 460 MW CAES-load line), compared with \$210 million for the standalone wind farm. The optimization algorithm convergence characteristics for some scenarios indicate that there are a number of combinations of the decision variables that have approximately the same profit. This gives these results an uncertainty of approximately 10 %; even accounting for this uncertainty, in all cases the lower capital costs of the smaller CAES-load transmission line do not compensate for the cost of the CAES system, and using CAES to smooth power from the wind farm is not profitable.

8.4 Analysis of the price-taker assumption for the zonal price scenario

The profit-maximizing CAES expander in the zonal price scenario would shift the ERCOT generation supply curve outward and reduce prices during times of high demand. To account for this effect, we examined supply curves for Wednesdays in each season of 2008, which we take to be representative of average days. In the region of the supply curve

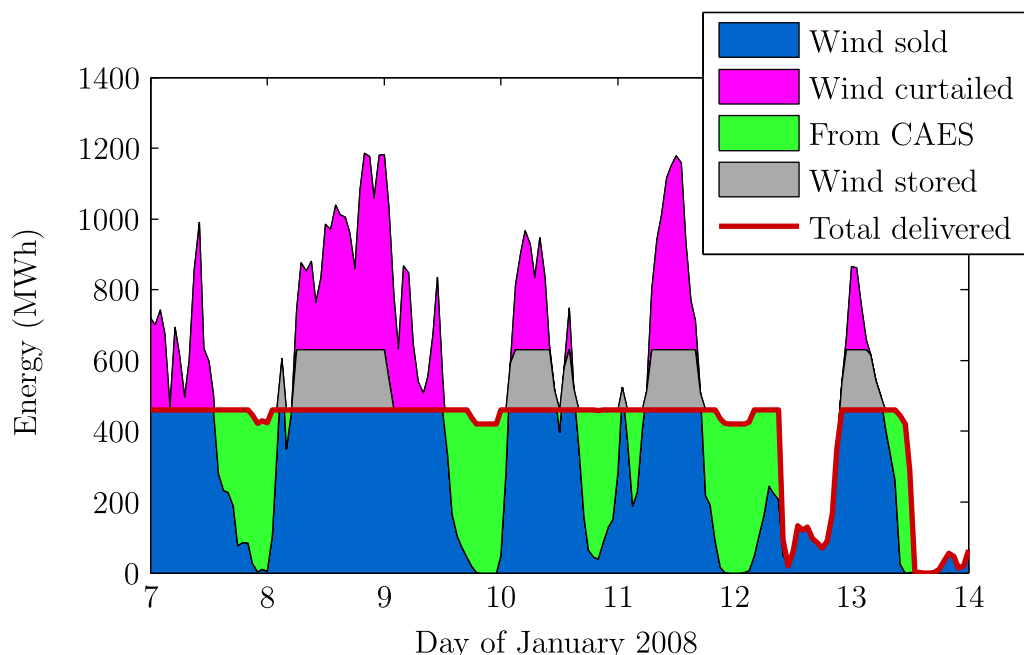


Figure 5: Wind-CAES system operation under a \$63/MWh contract price with load in Dallas. This scenario represents the wind/CAES system acting as a baseload generator, with a 1300 MW wind farm and 260 MW (expansion) CAES plant filling a 460 MW transmission line with 80 % capacity factor.

between first percentile load and 99th percentile load, the maximum price decrease caused by an additional 24 MW generator with low marginal cost is less than \$30/MWh. The optimization for the zonal price scenario was re-run with prices decreased by \$30/MWh when the CAES expander comes online and calculated annual profit decreased to \$700 million, still well above that of a standalone wind farm (\$245 million).

Daily supply curves, including those for days with price spikes on the order of \$1000/MWh, tend to have maximum bids of less than \$200/MWh. This implies that the price spikes are due to factors not directly represented by the bid stacks and ERCOT's economic dispatch algorithm. Possible alternative explanations include strategic bidding by electric power producers and outages of generators and transmission, which may remain largely unaffected by the presence of an additional large generator.

8.5 BES and regulation markets

For 2008 data with load in Houston, a wind/CAES system would maximize profit by bidding into the down-regulation market for 4–7 hours in the early morning in July through November and 0–3 hours the rest of the year. The system would only bid into the up-regulation market for 1–3 hours in the early evening in October through December, and from 9 am until midnight in September. This strategy results in an annual profit of \$100 million, in contrast to an annual profit of \$80 million if the system bids into the BES market alone under the given strategy. With load in Dallas, bidding patterns are similar and entry into regulation markets allows an identically-sized system to earn \$50 million, while bidding into the BES market alone generates a profit of \$20 million. Using 2009 wind and price data, participating in the regulation markets results in a loss of \$40 million (load in Houston) or \$70 million (load in Dallas), in contrast to a loss of \$50 million (Houston) or \$90 million (Dallas) if the system bids into the BES market alone under the given heuristic. In all cases, profit in the regulation and BES markets falls far short of that of a standalone wind farm.

8.6 Carbon price for an economically competitive wind/CAES system

We assessed the carbon price at which the profit-maximizing wind/CAES systems under the contract price scenarios would be economically competitive with a natural gas combined cycle (NGCC) generator producing the same amount of energy per year with a capital cost of \$900/kW and heat rate of 6800 BTU/kWh. At a natural gas price of \$5/MMBTU, the wind/CAES system with 2008 data and load in Dallas (Houston) would be cost-competitive with NGCC at a carbon price of \$230/tCO₂ (\$200/tCO₂); at a gas price of \$15/MMBTU, the wind/CAES system would be cost-competitive at \$56/tCO₂ (\$28/tCO₂). The lower cost of building air storage in a salt cavern renders the Houston

system more competitive. For the smaller profit-maximizing CAES systems of 2007 and 2009, the carbon prices to break even with NGCC are much higher — \$180–\$410/tCO₂ at \$15/MMBTU gas, and \$360–\$580/tCO₂ at \$5/MMBTU gas (Appendix 8). The 2008 results are similar to those of DeCarolis and Keith (2006), who used a different method and found that CAES paired with a single wind farm was cost-competitive at carbon prices above \$140/tCO₂ (2004\$). Since the NGCC could be sited closer to load than the wind farm, accounting for transmission costs would raise the carbon price at which a wind/CAES system is cost competitive.

9 Discussion and policy implication

Given 2007–2009 wind power output, electricity prices, and gas prices, a profit-maximizing owner of a 1300 MW wind farm in central Texas providing power to Dallas or Houston would not build a CAES system. The only profitable wind/CAES system under the zonal price scenario generates its revenue during large price spikes, which cannot be forecasted or expected to occur regularly, and thus provide uncertain revenue with limited power to attract investment (Wilson, 2008). Although such a system could have profitably captured the price spikes of 2008, a risk-averse firm might set future electricity price expectations closer to 2007 or 2009 levels, and thus decide not to build. Modifying the ERCOT supply curve to account for the presence of an additional large generator does not change this result.

With a \$300/MWh price cap and a \$100/MWd capacity payment, a wind/CAES system would be profitable given 2008 data and load in Houston. This result does not account for the additional fuel cost if the system were deployed when the cavern was depleted and the CAES plant was forced to run as a natural gas turbine. Since ERCOT does not currently have a capacity market (and since the system under this scenario is unprofitable given 2007 or 2009 data, or load in Dallas), this result does not support

investment in CAES.

Under the third pricing scenario, selling at a constant price equivalent to the mean BES price, a wind/CAES system is unprofitable. The cost savings of the smaller CAES-load transmission line with an 80 % capacity factor does not compensate for the capital cost of CAES. Allowing the wind/CAES system to bid into regulation markets raises its profit, though not enough to justify pairing CAES with a wind farm. There are currently no rigorous predictions of whether increased wind power penetration would raise ancillary service prices enough to change this result.

While a wind/CAES system in ERCOT would not be economically viable at the firm level, pairing CAES with wind has social benefits that could outweigh private costs. Sioshansi et al. (2009) calculated the net social benefit of large-scale energy storage for arbitrage in PJM (the sum of the changes in consumer and producer surplus due to increased off-peak prices and decreased on-peak prices) as \$4.6 million for a 1 GW/16 hour storage device, with negligible marginal benefit for more storage hours. This calculation was based on data from 2002, when PJM had an average load about 50 % greater than ERCOT's 2008 average load (Biewald et al., 2004). Although more detailed analysis would be necessary to assess the change in economic surplus due to the wind/CAES systems of contract price scenarios, for example, their smaller size and operation in a smaller market both suggest that the benefit would be less than that calculated by Sioshansi et al. (2009). The increase in economic surplus is thus unlikely to compensate for the private deficit and thus does not warrant a subsidy.

A wind/CAES system displacing a natural gas plant would also have human health benefits resulting from improved air quality. Gilmore et al. (2010) analyzed the air-quality effects of a 2000 MWh battery in New York City that charges for 5 hours off-peak and discharges for 4 hours on-peak. When the battery was charged with wind power and used to displace a simple-cycle gas turbine, the resulting social benefit due to reductions in particulate matter ($PM_{2.5}$) and CO_2 (assuming \$20/t CO_2) was \$0.06/kWh. The large

population density of New York City compared with Dallas or Houston, the different generation mixes in ERCOT and NYISO, and different atmospheric circulation patterns prohibit a direct extension of these results to ERCOT. A detailed study of the air quality benefits of storage in ERCOT is warranted to assess whether these benefits are large enough to justify a subsidy.

Pairing a CAES plant with a wind farm, either to produce smooth, dispatchable power or to store wind power and capture large upswings in hourly electricity prices, is not economically viable in ERCOT at the firm level. Further, our results suggest that current CAES technology is not a competitive method of wind power integration in ERCOT under plausible near-future carbon prices and does not produce social benefit that outweighs private costs, unless air quality benefits are shown to be substantial.

Acknowledgements

The authors thank Lee Davis, Horace Horn, Lester Lave, Gary Leidich, Jeremy Michalek, Sompop Patanariyankool, Rahul Walawalkar, and Sean Wright for useful comments and conversations. This work was supported in part by a grant from the Alfred P. Sloan Foundation and EPRI to the Carnegie Mellon Electricity Industry Center, the US National Science Foundation under the Graduate Research Fellowship Program, Carnegie Mellon University's Steinbrenner Institute Graduate Fellowship, the Institute for Complex Engineered Systems at Carnegie Mellon University, the Doris Duke Charitable Foundation, the Department of Energy National Energy Technology Laboratory, and the Heinz Endowments for support of the RenewElec program at Carnegie Mellon University. This research was also supported through the Climate Decision Making Center (CDMC) located in the Department of Engineering and Public Policy. This Center has been created through a cooperative agreement between the National Science Foundation (SES-0345798) and Carnegie Mellon University.

Appendix A CAES mechanics

The compression stage of CAES begins with the intake of air at ambient pressure and temperature. A motor, drawing electricity from the grid, wind farm, or other source, runs a series of progressively higher-pressure compressors and intercoolers to bring the air to its storage pressure and temperature. By cooling the air after each compression stage, the intercoolers reduce the power necessary for compression and the aftercooler reduces the required storage volume for a given mass of air (Gyuk and Eckroad, 2003).

The compressed air is stored in an underground cavern. Above-ground CAES designs have also been explored but are only cost-effective for systems storing less than approximately 100 MWh (Gyuk and Eckroad, 2003). Since this study examines CAES paired with large-scale wind, above-ground air storage is not considered further. Underground air storage is feasible in solution-mined salt caverns, aquifer-bearing porous rock, or mined hard-rock caverns. CAES geology is discussed further in Section 4.

When air is released from the cavern, the pressure must be throttled down to inlet pressure of the first expansion turbine. The expansion phase of the McIntosh-type CAES cycle consists of a high-pressure then a low-pressure combustion turbine. Before entering the high-pressure turbine, the air is heated in a recuperator, a heat exchanger that captures the exhaust heat from the low-pressure turbine. The turbines drive the generator, producing electricity that is sent to the grid and thus completes the CAES cycle.

Between the high- and low-pressure turbines, air is chilled to 60 °F and 1 atmosphere, allowing the system to operate with consistent efficiency even in hot weather (Gyuk and Eckroad, 2003).

The compressed air is stored in an underground cavern. Above-ground CAES designs have also been explored but are only cost-effective for systems storing less than approximately 100 MWh (Gyuk and Eckroad, 2003). Since this study examines CAES paired with large-scale wind, above-ground air storage is not considered further. Underground air storage is feasible in solution-mined salt caverns, aquifer-bearing porous

rock, or mined hard-rock caverns. CAES geology is discussed further in Section 4.

When air is released from the cavern, the pressure must be throttled down to inlet pressure of the first expansion turbine. The expansion phase of the McIntosh-type CAES cycle consists of a high-pressure then a low-pressure combustion turbine. Before entering the high-pressure turbine, the air is heated in a recuperator, a heat exchanger that captures the exhaust heat from the low-pressure turbine. The turbines drive the generator, producing electricity that is sent to the grid and thus completes the CAES cycle.

Between the high- and low-pressure turbines, air is chilled to 60 °F and 1 atmosphere, allowing the system to operate with consistent efficiency even in hot weather (Gyuk and Eckroad, 2003).

A.1 Ramp rates

Aspects of CAES that make it well-suited for leveling wind power output are high ramp rate and quick startup time (Schainker, 2007). In its compression phase, a CAES plant starts up in 10–12 minutes and ramps at 20 % per minute. In its generation phase, CAES starts up in 7–10 minutes and ramps at 200 % per minute. These parameters allow a CAES system to store or supplement wind power output such that the wind/CAES system delivers highly consistent power.

A.2 Adiabatic CAES

Although not yet demonstrated, the concept of adiabatic CAES would eliminate the use of fossil fuel in CAES. Rather than dissipating the heat of compression, as in the current CAES designs, adiabatic CAES would store the heat and subsequently use it to re-heat the air before the expansion stage. The efficiency of the system would be approximately 0.8 (kWh generated per kWh stored). EPRI has estimated the capital cost of an adiabatic CAES plant at \$1000/kW (EPRI estimates \$600–\$750/kW for the second-generation CAES design modeled in this paper). Although adiabatic CAES is likely not cost-effective

at current natural gas prices and under current greenhouse gas regulations, that could reverse under higher gas prices and stricter limits on greenhouse emissions. Industry experts affirm that the technology required to build a viable adiabatic CAES demonstration plant are within reach (Bullough et al., 2004).

Appendix B Extant and planned CAES plants

Two CAES plants are currently operational: one in Huntorf, Germany and one in McIntosh, Alabama. At least three others, in Iowa, Ohio, and Texas, are in planning or construction stages.

The oldest operating CAES plant, in Huntorf, Germany, was completed in 1978. It is used primarily for peak shaving, as a supplement to hydroelectric storage facilities, and as a means to supplement the ramp rate of coal plants. The system was originally designed to provide black-start services to nuclear plants and as a source of inexpensive peak power. The original two hours of storage were sufficient for these purposes, but the plant has since been modified for four storage hours (Gyuk and Eckroad, 2003). Aside from its original functions, it now helps mitigate power fluctuations from wind plants in North Germany (Succar and Williams, 2008).

The Alabama Electric Cooperative owns the McIntosh CAES plant, and completed it in 1991 after 30 months of construction (Gyuk and Eckroad, 2003). After initial problems with the underground storage were addressed, the McIntosh plant reached 91.2 % and 92.1 % starting reliability and 96.8 % and 99.5 % running reliability over 10 years for the generation and compression cycles respectively (Succar and Williams, 2008).

ISEP, slated to come online in 2011, will consist of a 268 MW CAES plant paired with 75–100 MW wind transported from as far as 320 km away (Succar and Williams, 2008). The underground storage will be developed in a saline aquifer in an anticline at a depth of approximately 900 m. The site was the third studied thoroughly after an initial screening

of 20 possibilities.

The Norton CAES plant will be a 2700 MW facility with air storage in an inactive limestone mine 670 m underground. Hydrodynamics Group (2009) and Sandia National Laboratories conducted tests to ensure that the limestone formation would hold its pressure seal and structural integrity at CAES operating pressures. Although the project has encountered siting problems, construction of the plant is now slated to move forward (Succar and Williams, 2008).

B.1 Wind-CAES ancillary services

In addition to ancillary services described in the paper, a wind-CAES system could provide reactive power support, either in an ancillary services market or to compensate for fluctuations in wind power output. ERCOT, however, requires local reactive power support from all generators with capacities greater than 20 MVA, so this service is not traded on the ancillary services market (ERCOT, 2009c). Furthermore, wind turbines with power electronic converter interfaces have a certain amount of built-in static VAR compensation, perhaps rendering VAR support from the CAES system unnecessary (EPRI, 2004).

As discussed previously, the two extant CAES plants primarily serve functions of peak shaving, arbitrage, black start, and supporting the ramp rate of coal plants. Future CAES plants, such as ISEP, will firm and shape wind power to reduce the need for spinning reserve to fill in gaps in wind power generation. The flexibility of CAES operation gives it a broad range of options over which to find the most profitable mode of operation.

Appendix C CAES geology in Texas

CAES is feasible in three broad types of geology: solution-mined salt caverns, aquifers of sufficient porosity and permeability, and mined hard rock caverns. Due to the disproportionately high cost of developing hard rock caverns, they are not considered in

this study. Succar and Williams (2008) estimate ranges of each type of CAES geology in the United States. While this map provides a broad indication of possible locations for CAES development, it is not definitive because siting a CAES plant depends largely on local geological characteristics and preexisting land use patterns.

C.1 CAES in solution-mined salt caverns

The two currently operational CAES plants, in McIntosh, Alabama and Huntorf, Germany, both use solution-mined salt caverns for air storage. These structures are advantageous for CAES due to the low permeability of salt, which enables an effective pressure seal, and the speed and low cost of cavern development. The caverns are formed by dissolution of underground halite (NaCl) in water and subsequent removal of the brine solution. The CAES plant injects and removes air through a single well connecting the salt cavern and turbomachinery.

A layer of water, left over from the solution mining process, remains on the bottom of the cavern and suspends particulates. Particulate matter does not reach the turboexpander inlet to cause corrosion or other problems (Davis, 2009).

While the cost of the salt cavern is relatively independent of the cavern's depth, the operating pressure range of the salt cavern depends on depth: 0.3 psi/ft (6.41 kPa/m) is an approximate lower bound, and 0.7–0.85 psi/ft (15.0–18.2 kPa/m) is an approximate upper bound (Swensen and Potashnik, 1994). The lower bound ensures that the cavern pressure does not deviate excessively from the surrounding lithostatic pressure and cause inward stress to the cavern walls. The upper bound must be less than the pressure that would cause upward force on the casing pipe to exceed the downward force of soil friction on the pipe. The pressure range of the salt cavern constrains the inlet pressure of the high-pressure expansion turbine, which cannot exceed the lower bound on cavern pressure less losses accrued between the cavern and HP expander.

C.2 Occurrence of salt formations amenable to CAES

Salt that can house a CAES cavern occurs in two general forms: bedded and domal.

Domal salt is more pure and massive than bedded salt and therefore superior for CAES cavern development, but specific sites in areas of bedded salt can be amenable to CAES as well. Domal salt occurs primarily in the Gulf Coast and East Texas Basin (Hovorka, 2009). Salt domes are formed when denser lithologies overlie salt beds and the salt begins to buoyantly rise to form diapirs, domes, and other intrusive structures in the overlying rock. The upper regions of salt domes often have concentrations of impurities that form a cap rock that protects the rest of the dome from dissolution in near-surface meteoric water. The salt caverns of both extant CAES plants, in McIntosh and Huntorf, were solution-mined in domal salt.

Bedded salt is originally deposited in restricted marine basins that undergo cyclic flooding and evaporation to form repetitive evaporite sequences containing halite interbedded with limestone, dolomite, anhydrite, polyhalite ($\text{K}_2\text{Ca}_2\text{Mg}(\text{SO}_4)_4 \cdot 2(\text{H}_2\text{O})$), and mudstone. The Bureau of Economic Geology at University of Texas at Austin performed a detailed characterization of bedded salt in the Midland Basin (Hovorka, 2009). Results of the study indicate that the Salado Formation, the dominant halite-bearing unit of the Midland Basin, contains thick and laterally-homogenous bedded salt that thins toward the east. The study provided a map of salt in Texas that provides a good indication of general areas that are likely to harbor the right conditions for a solution-mined CAES cavern (but cannot be interpreted as indicative of sites where CAES is feasible without further study.)

C.3 CAES in saline aquifers

Underground storage for a CAES plant is also feasible in an aquifer-bearing sedimentary rock. The rock must be sufficiently permeable and porous to allow water displacement and air cycling, and lie beneath an anticline of impermeable caprock to stop the buoyant rise of

air and impede fingering (Succar and Williams, 2008). A bubble in the aquifer, developed by pumping air down multiple wells, serves as the air storage cavern. The ratio of the total amount of air in the bubble to the amount that cycles over the course of CAES operation is typically between 5 and 30 (Swensen and Potashnik, 1994). This large amount of cushion serves to keep the bubble at a relatively constant size (Succar and Williams, 2008) and isolate the air/water interface from the wells that serve as conduits to the aboveground turbomachinery. The use of multiple wells instead of a single one ensures that the pressure gradient surrounding each well during CAES operation does not exceed the fracture pressure of the host rock. The Iowa Stored Energy Project (ISEP), a wind/CAES system under construction in Dallas Center, IA, will use an aquifer for underground storage.

The native pressure in the reservoir is approximately equal to the hydrostatic pressure of the aquifer. The operating pressure range of the reservoir is relatively narrow; the total mass of air in the storage bubble is typically 5–30 times the cycling air mass, such that the removal of the cycling air causes a relatively small drop in reservoir pressure. Since water has approximately 50 times the viscosity of air and flow rate is inversely proportional to viscosity, water in the aquifer does not significantly encroach on the bubble over the time-scale of plant operation. The storage reservoir is thus not pressure-compensated, and its function can be modeled as a salt cavern to good approximation (Succar and Williams, 2008).

The total turboexpander volume flow rate during power generation divided by the number of wells is given by Q in Equation (A1) (Swensen and Potashnik, 1994).

$$Q = K \cdot (P_w^2 - P_c^2)^n \quad (\text{A1})$$

P_w is the flowing wellhead pressure, P_c is the static wellhead pressure, and K and n are constants dependent on reservoir properties and well size.

Increasing the number of wells increases P_w but leaves P_c relatively fixed. This raises

the turboexpander flow rate (Q times the number of wells) and therefore the turboexpander inlet pressure. A high turboexpander inlet pressure reduces the specific air consumption (kg/kWh) of the generation phase, which lowers the heat rate and energy ratio and reduces the operating cost. The cost of drilling more wells, however, can offset the reduced operating cost. The number of wells and turboexpander inlet pressure can be optimized to produce the lowest cost per kWh of electricity generation. The optimal number of wells and turbine inlet pressure depend on aquifer parameters such as permeability, porosity, thickness, and depth, which constrain the bulk flow of air through the turbomachinery.

C.4 Occurrence of saline aquifers amenable to CAES

An early study on the use of aquifers for CAES was based on the success in storing natural gas in porous formations, and on the assumption that the techniques of storing air and natural gas are identical. The resulting map of possible aquifer CAES sites covered most of the central United States (Allen, 1985).

In 1994, Energy Storage and Power Consultants (ESPC) screened non-potable aquifers and depleted gas reservoirs in New York as potential sites for CAES (Swensen and Potashnik, 1994). To evaluate aquifers, ESPC first eliminated all geological groups, formations, and members solely associated with potable aquifers. The remaining sites were assessed for adequate thickness and porosity, and areas with land use incompatible with a CAES facility were eliminated. ESPC's search generated three possible sites for air storage in an aquifer, each with depths of 460–910 m and permeability of 100 mD. The report concluded with an enumeration of the process to further assess the aquifer sites for CAES cavern development and the associated costs, which included further searching of public and private records for relevant data, conducting seismic tests, developing a test well, modeling the reservoir to evaluate compatibility with CAES, securing permits, and testing air cycling facilities for the selected reservoir. The process was estimated to take two years and cost \$2,975,000 (1993\$). Although the results of this study cannot be directly applied

to CAES in Texas, they are illustrative of the processes and costs involved in characterizing and choosing an aquifer CAES site.

Following EPRI (1982), Succar and Williams (2008) assembled a table of suitable aquifer characteristics for CAES. It bears noting that the New York ESPC study chose a 3000-foot deep aquifer as a possible CAES site and that the Iowa Stored Energy Project will use an aquifer 880 m deep, both of which fall into the “unusable” depth range of this table (above 760 m). In addition, all three sites in the New York study have permeabilities of 100 mD, on the border between “unusable” and “marginal” in the table. These discrepancies underscore the importance of individual site testing and the difficulty of generalizing parameters for aquifer CAES sites.

Although specific sites for aquifer CAES in Texas have not been extensively examined, the Texas Bureau of Economic Geology (BEG) has evaluated aquifers for use in carbon capture and sequestration (CCS) at depths of 800–3000 m based on the criteria of injectivity and trapping (Hovorka, 1999). Injectivity is a measure of the formation’s ability to receive fluid and is determined by depth, permeability, formation thickness, net sand thickness, percent shale (injectivity declines above 50 % shale), and sand-body continuity (a measure of the possible size of the storage). Trapping is a measure of the formation’s ability to hold the injected fluid in place, and is determined by the thickness and continuity of the top seal, hydrocarbon production from the interval, fluid residence time, flow direction, solubility of the injected fluid in the fluid it displaces, rock/water reaction, and porosity.

CAES requires adequate injectivity and caprock for trapping, but also deliverability of air from the formation to the wells. Unlike CAES aquifers, CCS sites do not require an anticline: flat caprock structures are superior for CCS because they enable faster migration and dissolution of CO₂. The high viscosity of CO₂ under storage conditions and the low permeability in deep aquifers indicate that CO₂ flow behavior will be different than air in CAES (Succar and Williams, 2008). In addition, ideal CCS aquifers are at least 800 m deep

to keep CO₂ in its supercritical state. Depth requirements for aquifer CAES storage are less stringent, though the depth of the formation influences the operating pressure range of the air storage and thus the turboexpander inlet pressure. Although CAES is technically feasible at depths as shallow as 140 m (Succar and Williams, 2008), aquifers at these depths often contain potable water and are hence illegal to disturb (Swensen and Potashnik, 1994).

Studies of aquifers for CCS storage are poor wholesale proxies for CAES siting studies. Nevertheless, CCS studies provide data and analyses that yield limited insight into the siting of CAES facilities. The BEG compiled a database on possible CCS aquifers nationwide, including the Paluxy, Woodbine, Frio, Jasper, and Granite Wash formations of Texas that can be found in its online database (Texas Bureau of Economic Geology, 2009).

C.5 Depleted natural gas fields

Energy Storage and Power consultants screened depleted gas fields in New York for possible conversion to CAES caverns. ESPC chose to evaluate only those between 460 and 1520 m deep and with uncomplicated reservoir and caprock geology, and exclude fields with measurable oil production, more than 20 producing wells, or sensitive surface land use. The sites were further restricted by agreement with host utilities and the New York State Energy Research and Development Authority (NYSERDA). With these constraints, no depleted natural gas fields were found suitable for CAES cavern development (Swensen and Potashnik, 1994).

Appendix D CAES plant cost

The total cost of a CAES plant consists of its capital and operating costs. The capital cost includes the plant's turbomachinery (high and low pressure expanders, compressor, and recuperator), underground storage facility, and the balance-of-plant (including site preparation, building construction, and electrical and controls).

If the underlying geology is suited to a solution-mined salt cavern, storage cavern capital costs include the cost of drilling the wells, the leaching plant, cavern development and dewatering, the brine pipe (to transport the solution away from the site), and water. Development costs associated with aquifer CAES include the cost of drilling multiple wells, the gathering system, the water separator facility, and the electricity used to run an air compressor to initially create the air-storage bubble in the aquifer (Swensen and Potashnik, 1994).

CAES cost data from planning studies and extant plants was regressed against expander capacity (Figures A1 and A2) (Swensen and Potashnik, 1994; Haug, 2004; Schainker, 2008; Hydrodynamics Group, 2009). The capital cost of a CAES plant with salt-cavern storage is close to linear with expander capacity ($R^2 = 0.94$). The capital cost of a CAES plant with aquifer storage is more variable, due to the high site-specificity of the underground storage cost ($R^2 = 0.78$). The data were plotted with a 95 % prediction interval, which defines the range in which 95 % of future observations are expected to fall.

The marginal cost per kWh of energy storage in an aquifer is \$0.10–\$0.20, which reflects the cost of electricity required to expand the bubble such that the generation phase produces an additional kWh. The marginal cost to expand a solution-mined salt cavern to produce an additional kWh is \$1–\$2 (Schainker, 2008).

Appendix E Transmission capital cost

Transmission capital cost was first modeled as a linear regression of cost per GWm on MW capacity. This line had a negative slope and thus produced a parabolic function for total cost, in which extremely high-capacity transmission lines had costs that were near zero or negative. The optimization thus resulted in profits that were artificially high.

The transmission cost model (in \$/GWm) used in this research as a function of length

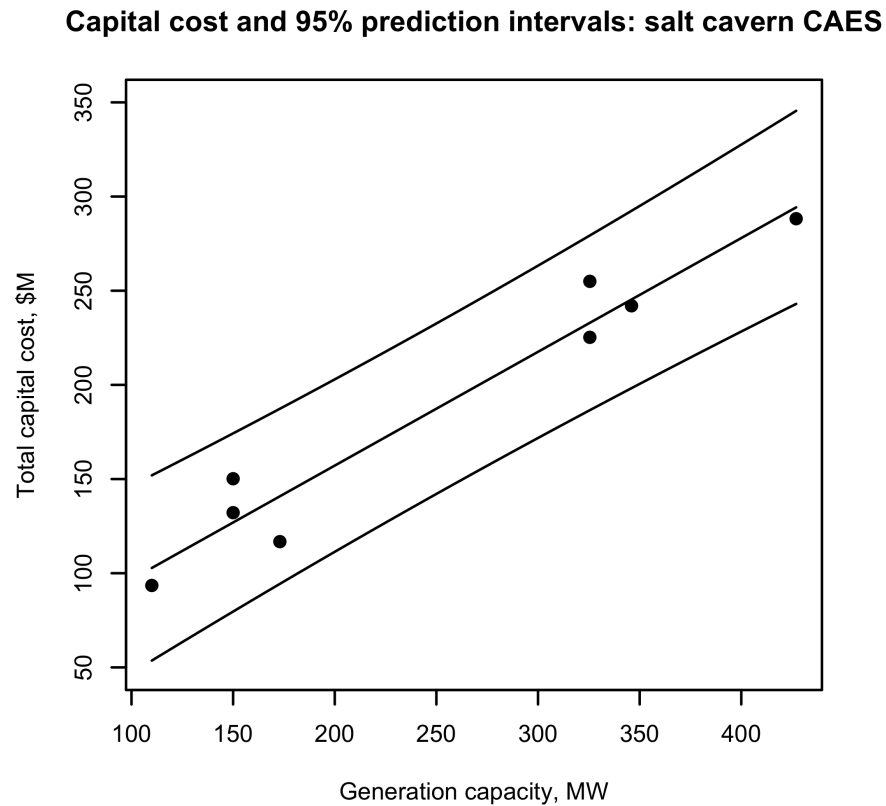


Figure A1: Capital cost and 95 % prediction intervals for development of a CAES plant with salt-cavern storage.

in km (L) and capacity in MW (T) is reproduced in Equation (A2).

$$Cost_T = 14266 \cdot L \cdot T^{0.527} \quad (A2)$$

The model is of the same form as the transmission cost model in Pattanariyankool and Lave (2010) but generates lower cost predictions. Pattanariyankool and Lave (2010) derived their cost model from a regression of inflation-adjusted, log-transformed data from transmission projects across the United States.

The model used in the current study was derived from a consultant report on transmission planning that contained the cost estimates presented in Table A1. The declining capital cost per GWm as a function of MW represents an economy of scale due to

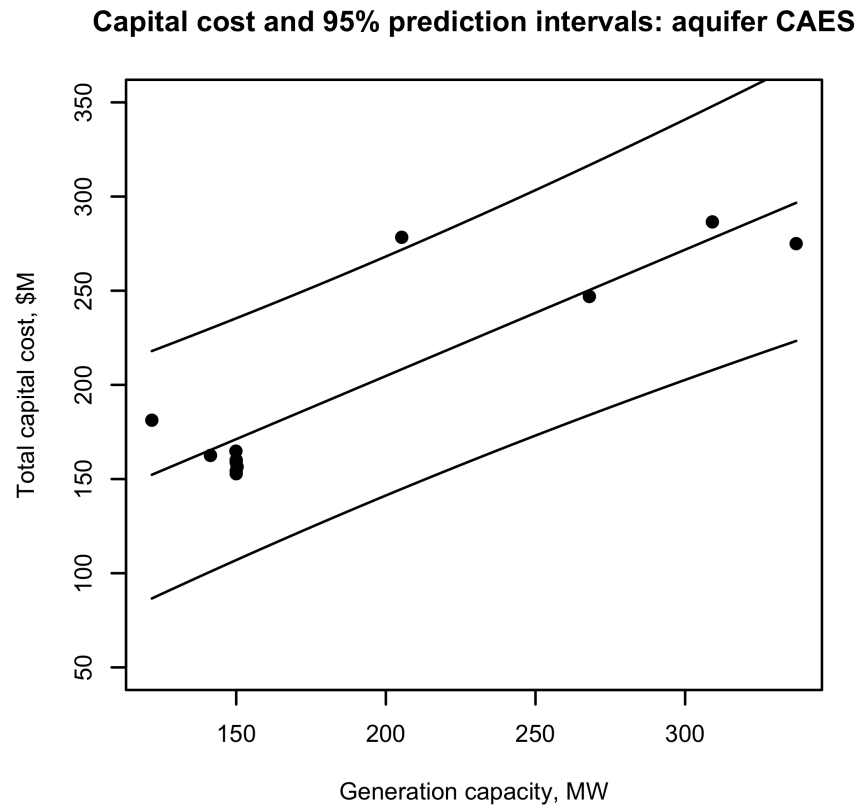


Figure A2: Capital cost and 95 % prediction intervals for development of a CAES plant with aquifer storage.

decreasing corridor widths per MW and to fixed costs of transmission line construction.

Table A1: Physical and cost parameters of transmission lines from Hirst and Kirby (2001).

Voltage (kV)	Capacity (MW)	Capital cost (\$/GWm)	Corridor width (m)
230	350	856	30
345	900	625	38
500	2000	375	53
765	4000	281	61

Predictions generated by the model in Equation (A2) were tested against data from an ERCOT study on transmission costs associated with wind integration. Table A2 presents data from the ERCOT study on the costs and lengths of transmission lines needed to transport a given nameplate capacity of central Texas wind power to load. The model fit the ERCOT data with $R^2 = 0.72$. The mean ratio of predicted to actual cost was 1.04 with

a standard deviation of 0.22.

Table A2: Length, capacity, and total cost of transmission from ERCOT study, and predicted total cost based on Equation (A2).

MW	Cost (\$millions)	Predicted cost (\$millions)	Length (km)
1400	381	344	528
1500	190	184	272
1500	320	270	400
1800	258	257	346
2000	376	397	506
3000	320	358	368
3800	960	1,516	1380
3800	860	1,144	1040
4500	1,130	1,202	1000
4600	1,520	1,498	1232
500	12	13	32

As a counterpoint to the model in which transmission cost is directly proportional to length, Mills et al. (2009) analyzed 40 transmission planning studies and found that cost per kW of transmission capacity is independent of length. Mills et al. (2009) note that the absence of observed length-dependency in the transmission data could be due to inconsistencies among the methodologies of the transmission studies analyzed; the fact that transmission costs are compared in unadjusted nominal dollars for different years could also have obscured other trends in the data. Mills et al. (2009) also note that projects involving greater transmission lengths tend to integrate more wind capacity; this trend reduces the apparent cost per kW of transmission projects involving large cumulative lengths of transmission lines. Mills et al. (2009) examine cost estimates for projects in all areas of the United States, which have large variation in siting difficulties. For a quantitative framework with which to analyze transmission siting difficulty, see Vajjhala and Fischbeck (2007).

Appendix F Wind-CAES dispatch model

If $p_i < MC_W$, the CAES system stores an amount of wind energy equal to the minimum of w_i (wind energy generated), E_C (compressor power), T_W (wind-CAES transmission

capacity), and the amount of energy the CAES cavern is capable of storing

$(E_E \cdot E_S \cdot ER - s_i)$. Any wind energy produced in excess of this amount is curtailed.

If $MC_W < p_i < p_s$, the system stores as much wind power as possible and sells the excess. The model first calculates x_i , the amount of energy that the CAES system can store in hour i , as the minimum of T_W and the amount of extra energy the cavern can store in its current state $(E_E \cdot E_S \cdot ER - s_i)$. If $w_i < x_i$, the system stores the entire output of the wind farm. If $w_i > x_i$, the system stores x_i and sells the remainder of the wind energy that does not exceed the transmission capacities of either line, and curtails any additional wind power.

If $p_s < p_i < p_d$, the system sells as much wind energy as possible directly into the grid. If the wind energy output does not exceed either transmission capacity, w_i is transmitted to load. If wind energy output exceeds T_W , wind generation is curtailed to the minimum of T_W and $T_C + x_i$ and the amount of energy sold is equal to the lower transmission capacity. If wind energy output exceeds T_C but not T_W , the CAES-load line is filled and the excess wind energy up to x_i is stored.

If $p_i > p_d$, the system supplements the wind energy output by discharging the CAES until the CAES-load transmission line is full or the storage cavern is emptied. The model first calculates y_i , the energy that the CAES system can produce in hour i , as the minimum of the expander capacity over the time interval (E_E) , and the total amount of energy that the storage cavern can supply in its current state (s_i/ER) . If wind energy output does not exceed the capacity of either transmission line, the system sells all of the wind energy and supplements it by discharging the storage up to y_i or the capacity of the CAES-load transmission line. If the wind-CAES transmission line is the smaller of the two and wind energy output exceeds the capacity of this line, the model fills the wind-CAES line, curtails the rest of the wind power, and sells the transmitted wind power supplemented with y_i up to the capacity of the CAES-load line. If the CAES-load line is the smaller of the two and wind energy output exceeds the capacity of this line, the system transmits wind energy up

to the CAES-load line capacity, stores wind energy up to x_i , and curtails the rest.

Appendix G Optimization algorithm

A simulated annealing algorithm after Goffe et al. (1994) was used to optimize the profit function. The temperature at each iteration was 85 % of the temperature of the last, and the initial temperature was 1000. Gradient-descent algorithms were impractical because the optimization function has zero local gradient with respect to the storage and discharge threshold prices p_s and p_d : if, for example, the zonal price data contains values of \$76.52 and \$76.58 but nothing in between, all p_s or p_d values between those two prices will generate the same profit (all other parameters being equal) and there is a plateau in the profit function.

Appendix H Extended carbon price results

Table A3 shows optimal expander capacities, fractional energy outputs from CAES, and break-even carbon prices (compared with an NGCC plant) for a wind/CAES system with load in Houston or Dallas and with data from 2007–2009. Sizes of CAES and transmission paired with a Texas wind farm are optimized for profit.

Table A3: Profit-maximizing CAES expander sizes under the contract price scenario, fractions of wind/CAES system energy output from the CAES plant, and carbon prices to reach cost-parity with a NGCC plant at \$5/MMBTU and \$15/MMBTU gas, for both load centers and all years considered.

	CAES expander (MW)	Fraction of energy from CAES	Carbon price (\$/tCO ₂), \$5/ MMBTU gas	Carbon price (tCO ₂), \$15/ MMBTU gas
Dallas 2007	300	0.16	380	210
Dallas 2008	460	0.16	230	56
Dallas 2009	200	0.12	580	410
Houston 2007	300	0.17	360	180
Houston 2008	480	0.17	200	28
Houston 2009	200	0.13	540	370

Bibliography

- Allen, K. (1985). CAES: the underground portion. *IEEE Transactions on Power Apparatus and Systems*, PAS-104:809–813.
- Biewald, B., Steinhurst, W., White, D., and Roschelle, A. (2004). *Electricity Prices in PJM: A Comparison of Wholesale Power Costs in the PJM Market to Indexed Generation Service Costs*. Synapse Energy Economics, Cambridge, MA.
- Bullough, C., Gatzert, C., Jakiel, C., Koller, M., Nowi, A., and Zunft, S. (2004). Advanced adiabatic compressed air energy storage for the integration of wind energy. In *Proceedings of the European Wind Energy Conference*, London, UK.
- Database of State Incentives for Renewables and Efficiency (2010). DSIRE. <http://www.dsireusa.org>.
- Davis, L. (2009). Personal communication with E. Fertig.
- DeCarolus, J. and Keith, D. (2006). The economics of large-scale wind power in a carbon constrained world. *Energy Policy*, 34:395–410.
- Denholm, P. and Sioshansi, R. (2009). The value of compressed air energy storage with wind in transmission-constrained electric power systems. *Energy Policy*, 37:3149–3158.
- DOE (2008). *20 % Wind energy by 2030*. U.S. Department of Energy, Office of Energy Efficiency and Renewable Energy. <http://www1.eere.energy.gov/windandhydro/pdfs/41869.pdf>.
- EIA (2009a). *Annual Energy Review*. Energy Information Administration, U.S. Department of Energy.
- EIA (2009b). *Texas Natural Gas Price Sold to Electric Power Consumers*. Energy Information Administration, U.S. Department of Energy.
- EPRI (1982). *Compressed Air Energy Storage Preliminary Design and Site Development Program in an Aquifer*. Electric Power Research Institute. EPRI Report EM-2351.
- EPRI (2004). *Wind Power Integration Technology Assessment and Case Studies*. Electric Power Research Institute. EPRI Report 1004806.
- ERCOT (2006). *Analysis of transmission alternatives for competitive renewable energy zones in Texas*. Electric Reliability Council of Texas. http://www.ercot.com/news/presentations/2006/ATTCH_A_CREZ_Analysis_Report.pdf.
- ERCOT (2008). *Understanding Texas nodal market implementation*. Electric Reliability Council of Texas. <http://nodal.ercot.com/about/kd/understandingNodal012308.pdf>.
- ERCOT (2009a). *Balancing Energy Services Daily Reports*. Electric Reliability Council of Texas. <http://www.ercot.com/mktinfo/services/bal/>.

- ERCOT (2009b). *Individual Resource Output*. Electric Reliability Council of Texas.
https://pi.ercot.com/contentproxy/publicList?folder_id=39175212.
- ERCOT (2009c). *Protocols, Section 6: Ancillary Services*. Electric Reliability Council of Texas.
- ERCOT (2010). *Historical Day-Ahead Ancillary Services Prices*. Electric Reliability Council of Texas.
https://pi.ercot.com/contentproxy/publicList?folder_id=10001730.
- Gardner, J. and Haynes, T. (2007). *Overview of Compressed Air Energy Storage*. Office of Energy Research, Policy and Campus Sustainability, Boise State University.
- Gilmore, E., Apt, J., Walawalkar, R., Adams, P., and Lave, L. (2010). The air quality and human health effects of integrating utility-scale batteries into the New York State electricity grid. *Journal of Power Sources*, 195:2405–2413.
- Goffe, W., Ferrier, G., and Rogers, J. (1994). Global optimization of statistical functions with simulated annealing. *Journal of Econometrics*, 60:65–99.
- Gonzales, E. (2010). *CAES training at the National Energy Technology Laboratory (NETL)*. U.S. Department of Energy (DOE).
- Greenblatt, J., Succar, S., and Denkenberger, D. (2007). Baseload wind energy: modeling the competition between gas turbines and compressed air energy storage for supplemental generation. *Energy Policy*, 35:1474–1492.
- Gyuk, I. (2004). *EPRI-DOE Handbook Supplement of Energy Storage for Grid Connected Wind Generation Applications*. Electric Power Research Institute (EPRI) and U.S. Department of Energy (DOE). EPRI Report 1008703.
- Gyuk, I. and Eckroad, S. (2003). *EPRI-DOE Handbook of Energy Storage for Transmission and Distribution Applications*. Electric Power Research Institute (EPRI) and U.S. Department of Energy (DOE). EPRI Report 1001834.
- Haug, R. (2004). *The Iowa Stored Energy Plant*. Iowa Association of Municipal Utilities.
- Haug, R. (2006). *The Iowa Stored Energy Plant: Capturing the Power of Nature*. Iowa Stored Energy Plant Agency.
- Hirst, E. and Kirby, B. (2001). *Transmission Planning for a Restructuring U.S. Electricity Industry*. Edison Electric Institute.
- Hovorka, S. D. (1999). *Optimal Geological Environments for Carbon Dioxide Disposal in Saline Aquifers in the United States*. Texas Bureau of Economic Geology.
- Hovorka, S. D. (2009). *Characterization of Bedded Salt for Storage Caverns: Case Study from the Midland Basin*. Texas Bureau of Economic Geology.

- Hull, R. (2008). NYSEERDA awards \$6 million in power delivery R&D projects; efficient, reliable electric grid, reduced emissions are goals. *NYSEERDA Press Release*.
- Hydrodynamics Group (2009). Norton compressed air energy storage. <http://hydrodynamics-group.com/mbo/content/view/16/40>.
- Klara, J. M. and Wimer, J. M. (2007). *Cost and Performance Baseline for Fossil Energy Plants*. Vol. 1, DOE/NETL-2007/1281. U.S. Department of Energy (DOE).
- LaMonica, M. (2009). PG&E to compress air to store wind power. *Wind Watch: Industrial Wind Energy News*.
- Leidich, G. (2010). Personal communication with E. Fertig.
- Lozowski, D. (2009). Chemical engineering plant cost index. *Chemical Engineering*, 116, 79.
- Marchese, D. (2009). Transmission system benefits of CAES assets in a growing renewable generation market. *Energy Storage Association Annual Meeting*.
- Mills, A., Wiser, R., and Porter, K. (2009). *The Cost of Transmission for Wind Energy: A Review of Transmission Planning Studies*. Lawrence Berkeley National Laboratory.
- Pattanariyankool, S. and Lave, L. B. (2010). Optimizing transmission from distant wind farms. *Energy Policy*, 38(6):2806–2815.
- Perekhodtsev, D. (2004). *Two Essays on Problems of Deregulated Electricity Markets*. PhD thesis, Tepper School of Business, Carnegie Mellon University, Pittsburgh, PA.
- Ridge Energy Storage & Grid Services L.P. (2005). The economic impact of caes on wind in tx, ok, and nm.
- Schainker, R. (2007). *Compressed Air Energy Storage: New Results and Plant Configurations*. PowerPoint presentation. The Electric Power Research Institute (EPRI).
- Schainker, R. (2008). *Compressed Air Energy Storage System Cost Analysis*. The Electric Power Research Institute (EPRI). EPRI Report 1016004.
- Sioshansi, R., Denholm, P., Jenkin, T., and Weiss, J. (2009). Estimating the value of electricity storage in PJM: arbitrage and some welfare effects. *Energy Economics*, 31:269–277.
- Succar, S. and Williams, R. H. (2008). *Compressed Air Energy Storage: Theory, Resources, and Applications for Wind Power*. Princeton Environmental Institute.
- Sullivan, P., Short, W., and Blair, N. (2008). *Modeling the Benefits of Storage Technologies to Wind Power*. National Renewable Energy Laboratory, U.S. Department of Energy.
- Swensen, E. and Potashnik, B. (1994). *Evaluation of Benefits and Identification of Sites for a CAES Plant in New York State*. Energy Storage and Power Consultants. EPRI Report TR-104268.

- Texas Bureau of Economic Geology (2009). *Study Areas. Carbon Dioxide Sequestration*.
<http://www.beg.utexas.edu/environqlty/co2seq/co2data.htm>.
- Vajjhala, S. P. and Fischbeck, P. S. (2007). Quantifying siting difficulty: a case study of US transmission line siting. *Energy Policy*, 35:650–671.
- Waxman, H. and Markey, E. (2009). *H.R. 2454: American clean energy and security act of 2009*. United States House of Representatives, 111th Congress.
<http://www.govtrack.us/data/us/bills/text/111/h/h2454pcs.pdf>.
- Western Governors' Association, Northwest Power Planning Council (2002). *New Resource Characterization for the Fifth Power Plan: Natural Gas Combined-cycle Gas Turbine Power Plants*. http://www.westgov.org/wieb/electric/Transmission%20Protocol/SSG-WI/pnw_5pp_02.pdf.
- Wilson, J. F. (2008). *Raising the stakes on capacity incentives: PJM's reliability pricing model (RPM)*. Prepared for The American Public Power Association.
<http://www.publicpower.org/files/PDFs/RPMreport2008.pdf>.

Paper 2

Optimal investment timing and capacity choice for pumped hydropower storage

Abstract

Pumped hydropower storage can smooth output from intermittent renewable electricity generators and facilitate their large-scale use in energy systems. Germany has aggressive plans for wind power expansion, and pumped storage ramps quickly enough to smooth wind power and could profit from arbitrage on the short-term price fluctuations wind power strengthens. This paper considers five capacity alternatives for a pumped storage facility in Norway that practices arbitrage in the German spot market. Price forecasts given increased wind capacity are used to calculate profit-maximizing production schedules and annual revenue streams. Real options theory is used to value the investment opportunity, since unlike net present value, it accounts for uncertainty and intertemporal choice. Results show that the optimal investment strategy under the base scenario is to invest in the largest available plant approximately eight years into the option lifetime.

This paper will be published as Fertig, E., Heggedal, A. M., Doorman, G., and Apt, J. (in press). Optimal investment timing and capacity choice for pumped hydropower storage. *Energy Systems*.

1 Introduction

Wind power capacity in Germany is expected to grow in the coming decades far beyond its current value of 27 GW. The German Advisory Council on the Environment projects that installed wind capacity could reach 113 GW by 2050 (SRU, 2010), and the German Federal Ministry for the Environment estimates that wind capacity could reach 94 GW by 2020 and 147 GW by 2050 (BMU, 2010). Since wind power is nondispatchable, energy storage large enough to absorb excess wind power and release it at hourly and daily timescales can facilitate the integration of wind capacity into electricity grids by shaping power output to better match load and relieving ramping demands on other generators. The storage facilities would profit from providing this service through price arbitrage and participation in ancillary service markets. A certain degree of daily arbitrage revenue, largely due to demand-correlated price fluctuations, is possible in the current market. Increased wind capacity, however, could cause a shift in generator mix that would increase price volatility and therefore arbitrage opportunity (Nicolosi and Fürsch, 2009).

This paper investigates how a pumped hydropower storage system, modeled after the proposed upgrade to the Tonstad hydropower plant in southern Norway, can operate in the German market to maximize profit given expectations of increasing spot price volatility. By optimizing the production schedule of the pumped storage plant for profit, calculating the annual revenue stream, and applying a real options analysis, we find the optimal investment timing and capacity choice for the pumped storage system.

1.1 Prospects for pumped hydropower storage

Storage technologies large enough to perform daily balancing of intermittent renewables include batteries, compressed air energy storage (CAES), and pumped hydropower storage. Batteries have the advantage of flexible siting, but at present cost and scale discourage their use for applications requiring storage capacity on the order of GWh. CAES can

balance wind power at daily timescales; the 290 MW CAES plant in Huntorf, Germany, built in 1978, is increasingly used for this function (Gyuk, 2003). Demonstrated CAES designs, however, use natural gas to reheat the expanding air during the generation phase. While the heat rate of CAES is approximately 4,300 BTU/kWh, which compares favorably with that of the most efficient combined-cycle gas turbines currently deployed at 5,690 BTU/kWh, CAES still produces significant carbon emissions and is vulnerable to price fluctuations in the natural gas market. Pumped hydropower storage, although feasible in fewer areas of the world than CAES, has the advantage of low carbon emissions and independence of fossil fuel prices. Pumped storage plants ramp quickly and have low startup costs in both pumping and generation mode, and their efficiencies are comparable to those of other storage technologies (Hadjipaschalis et al., 2009). They are thus well suited to balancing wind power fluctuations at hourly through weekly timescales.

Germany has 7 GW of pumped storage production capacity with 0.04 TWh reservoir capacity (SRU, 2010), and an additional 4.5 GW hydropower capacity without pumping that utilizes 0.3 TWh of reservoir capacity (Lehner et al., 2005). These resources alone are inadequate to support substantial added wind capacity. In a feasibility study of 100 % renewable electricity in Germany by 2050, the German Advisory Council on the Environment (SRU) found that requiring Germany's electricity system to be entirely independent would require the use of expensive CAES plants in addition to pumped storage to balance wind. The study found that connection with Northern Europe would allow Norway's 27 GW of hydropower production capacity and 84 TWh of storage capacity not only to replace domestic CAES as the most cost-effective method of balancing wind power but would also enable construction of more wind capacity in Germany (SRU, 2010). The need for daily balancing of wind will likely be signaled by increased short-term price volatility in the German market, creating opportunities for Norwegian hydropower storage facilities to profit through arbitrage.

Loisel et al. (2010) analyze CAES and pumped hydropower storage as methods for

future wind power integration in Europe and find that the main roles of the storage systems are wholesale price arbitrage and wind power curtailment avoidance. While the value of storage in the near future may not compare favorably with other wind-smoothing technologies such as gas turbines, Loisel et al. (2010) find that evolving regulatory and market conditions could make large-scale storage economically viable in the future. Deane et al. (2010) review existing and planned pumped storage plants in the European Union, the United States, and Japan, and find that between 2009 and 2017 over 7 GW of pumped hydropower storage will be built in Europe. Of this capacity, 2,140 MW will be located in Switzerland at a cost of 0.75 million euro per MW, 1,430 MW in Austria at 0.74 million euro per MW, and 1,000 MW in Germany at 0.70 million euro per MW. The costs for all proposed hydropower storage plants in the review range from 0.47 million to 2.17 million euro per MW.

Schill and Kemfert (2011) use a game-theoretic Cournot model to examine the strategic use of energy storage in the German electricity market. Results of the study suggest that oligopolistic ownership of storage would result in underutilization, since the owner would have the incentive to withhold storage capacity to maintain short-term price volatility and thus arbitrage opportunities. This effect, however, disappears under the assumption of dissipated ownership of storage assets. Schill and Kemfert (2011) also find that owners of other generators operating in the German market with enough capacity to have market power are unlikely to invest in storage, since its smoothing effect on prices would reduce the profitability of conventional generation. Current energy storage facilities in Germany appear not to be a relevant source of market power.

1.2 The effect of increased wind capacity on electricity prices

Profitability of the pumped hydropower storage facility rests largely on the effect of increased wind capacity on spot price volatility, which is subject to conflicting effects that operate over different timescales. Weigt (2009) finds that given a static generation

portfolio, wind power suppresses average short-term price volatility. Accounting for changes in generator mix due to the increasing share of wind generation over time, Nicolosi and Fürsch (2009) find that increased renewable energy penetration will result in more volatile residual demand (demand minus renewable power output) such that the market share of thermal base load generators shrinks and that of peak load generators grows. Since wind power has negligible marginal cost and peak load power plants tend to have high marginal cost, the shift in generation portfolio will increase volatility in spot prices (Nicolosi and Fürsch, 2009). Employing a bottom-up power system model of Europe, Schaber et al. (2012) find that the standard deviation of electricity prices will increase from 5 euro/MWh to 8 euro/MWh in 2020 with projected increases in renewable energy capacity and without substantial grid extensions. Connolly et al. (2011) examine profitability of a pumped storage system performing daily arbitrage and find that accurate long-term price projections are not necessary to maximize profit, but that profit varies widely from year to year such that realized NPV is highly uncertain.

1.3 The use of real options to value investment in pumped storage

NPV analysis alone is often used to evaluate investment opportunities in energy projects. Real options theory, however, offers the advantage over NPV of comparing the current project value with the expected value of postponing investment, enabling calculation of optimal investment strategy with respect to either time or expected profit. Although delaying investment sacrifices immediate revenue, factors that encourage waiting include an expected rise in future profitability, uncertainty that will resolve in the future, and discounting of capital costs.

Décamps et al. (2006) value the option to invest irreversibly in one of two mutually exclusive projects, and find intervals of output price for which investment versus waiting is optimal. In particular, a waiting region exists at prices above which it is optimal to invest

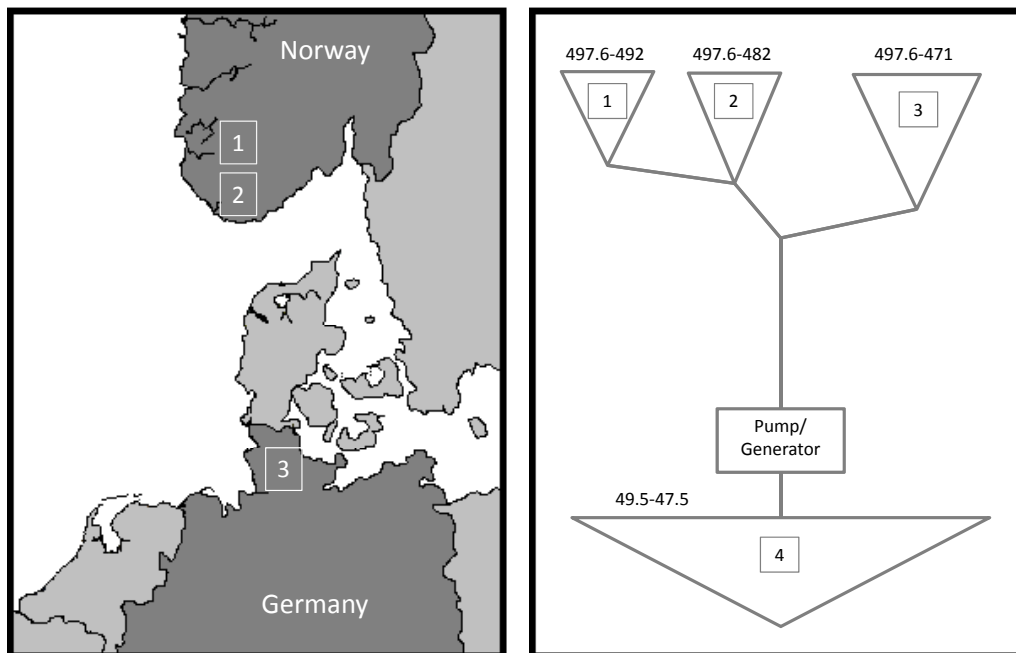
in one project and below which it is optimal to invest in the other, for the purpose of gathering more information before making the investment decision. This solution improves upon that of Dixit (1993), who argues for a single threshold value below which to postpone investment and above which to invest in the most valuable project. The simulation-based method of Longstaff and Schwartz (2001), implemented in the current study, also evaluates an option on a choice among multiple projects and produces a value function in the manner of Décamps et al. (2006).

The only other study of which we are aware that uses real options to value pumped hydropower storage is that of Muche (2009), who calculates the option value of a pumped hydropower storage plant as the arbitrage revenue averaged over a set of simulated price paths given optimal unit commitment for each path. This option value is compared with arbitrage revenue calculated from a single expected price path, which contains no price spikes and shows less short-term price volatility than individual simulated paths. The real option is framed as an option to operate the pumped hydropower storage plant flexibly in response to actual prices, which renders the plant valuable; it is contrasted with an unrealistic scenario of prices that vary deterministically with time and do not spike, resulting in a negative NPV for the pumped storage plant. While our analysis also rests on modeling profit of a pumped hydropower storage plant given simulated price paths, our option valuation technique optimizes investment timing and capacity choice.

1.4 Structure of the analysis

The investment opportunity considered is a unique, previously secured right to construct a pumped storage facility in Norway which operates solely on the German power market, an HVDC connection from the facility to the North Sea coast, and a subsea HVDC connection to Germany. Diagrams of the geography and operation of the facility are shown in Figure 1. The system operates only in the German European Energy Exchange (EEX), which has more volatile electricity prices than NordPool, the Scandinavian power market.

Norwegian power producers are considering connecting a pumped hydropower storage facility with EEX because connection with the Nordic system would require significant investments in the domestic grid in southern Norway. Construction plans for transmission would take a long time to be finalized, and in the meantime the pumped hydropower storage facility could reap the benefits of interconnection with the German system only.



(a) Geographical layout of investment opportunity. (b) Pumped hydropower storage facility.

Figure 1: (a): Diagram of the investment opportunity. A pumped hydropower storage plant, modeled after the planned upgrade to the Tonstad power plant, is located at (1). The facility is connected through an HVDC transmission line to the coast (2), where the subsea HVDC cable starts that traverses the North Sea to Germany (3). (b): The pumped storage facility has upper reservoirs (1), (2) and (3) and lower reservoir (4). The numbers above the reservoirs indicate the meters above sea level between which the water levels can vary. In generation mode water is released from the upper reservoirs through the tunnel, and in pumping mode the flow direction is reversed.

The remainder of this paper is organized as follows: Section 2 presents the valuation framework, which consists of a price model and a pumped storage dispatch model used for calculation of NPVs, which in turn serve as an input to a real option valuation. Section 3 presents a numerical example with data from the Tonstad power plant in southern Norway,

Section 4 presents results, and Section 5 concludes and provides recommendations for future work.

2 Valuation Framework

The valuation framework consists of three parts: a price model, a model for pumped hydropower storage scheduling, and a real options valuation. The price model uses bootstrapping methods to simulate future spot prices that serve as input to the pumped storage scheduling optimization. The optimal production schedule of the pumped storage facility is found by maximizing revenue over successive one-week time intervals. Annual revenue calculated from the production schedules serve as input to the real options valuation, which in turn yields the optimal investment timing and project size. Figure 2 gives an overview of the valuation framework, and a detailed explanation of each component follows.




Figure 2: Structure of the analysis.

2.1 Price Model

Since short-term price volatility rather than mean price drives the arbitrage revenue of a pumped hydropower storage system, the aim of the price model is to capture the change in price volatility over time.¹ An assumed three-year construction lag, ten-year option lifetime, and 40-year project lifetime require each simulated price path in the option valuation to span 53 years. The price path simulation begins by randomly sampling 53

¹Additional revenue streams would be available to the facility through participation in ancillary service markets. Here we consider only arbitrage revenue in the spot market.

years of data, in yearly increments with replacement, from 2003–2010 German EEX historical prices. The pumped storage plant is assumed to be a price taker in the German system. The price in hour i of year y ($2011 \leq y \leq 2063$) of each sample path is then adjusted to reflect the increased price volatility (defined as price standard deviation) due to greater market share of wind power. The adjusted price $\hat{p}_{y,i}$ for hour i in year y is calculated as

$$\hat{p}_{y,i} = \mu_{y,j} + (p_{y,i} - \mu_{y,j}) \cdot \beta_{y,j}, \quad (1)$$

where $\mu_{y,j}$ is the mean historical price for month j in sampled year y , $p_{y,i}$ is the historical price for hour i of sampled year y , and $\beta_{y,j}$ scales the standard deviation of prices in month j of year y .²

$\beta_{y,j}$ is calculated as

$$\beta_{y,j} = \frac{1 + (y - 2010) \cdot (b_j + \sigma_j \cdot \epsilon)}{(1 + k)^{y-2010}}, \quad (2)$$

where b_j is the mean annual percentage-point increase in price volatility in month j since 2010, σ_j represents uncertainty around that value, ϵ is a standard normal random variable, and k is the annual inflation rate taken as 2 %, the annual inflation target for the European Union (European Central Bank, 2011).

Baseline values for b_j and σ_j were calculated from 2010 yearly Phelix futures prices for peak and base load power for 2011 through 2016 (Phelix, 2011).³ Peak load contracts are for delivery from 08:00 to 20:00, Monday through Friday, and base load contracts are for delivery of constant power. These data were used to calculate off-peak prices, since the increase in the difference between peak and off-peak prices (henceforth termed the intraday price difference) is proportional to the increase in price standard deviation. The nominal mean intraday price difference increased by 8 percentage points annually above the 2011

²The mean prices $\mu_{y,j}$ are calculated monthly, since increasing price volatility around a daily mean would ignore that periods of high and low wind can last longer than a day, and adjusting prices around a yearly mean could mask finer-scale fluctuations and artificially suppress or elevate prices that are low or high for longer periods of time due to seasonal or macroeconomic effects.

³Phelix refers to the physical electricity index in the EEX power spot market for Germany and Austria.

value (Figure 3), and this increase was adjusted downward for inflation by 2 % per year to convert to real prices (Equation (2)). We calculate σ in Equation (2) as the annual growth in the standard deviation of the intraday price difference in the futures market, equal to 2.8 percentage points per year. The monthly parameters b_j and σ_j were then calibrated




Figure 3: Futures prices for 2010 trades show increasing intraday price difference (proportional to the increase in standard deviation), which creates increasing arbitrage opportunities for pumped storage systems. Each data point represents the intraday price difference in the yearly futures market for contracts traded on a single day divided by that difference for 2011 futures that day. The uncertainty in this value grows for delivery dates from 2011 through 2016.

according to the average wind capacity factor in Germany, which is near 30 % in the winter but only 15 % in the summer. Given a static generation portfolio, wind power in Germany has historically suppressed price volatility the least in the summer, when capacity factor is lowest (Weigt, 2009). We therefore scale b_j and σ_j to be inversely proportional to wind power capacity factor in month j but have a mean over all twelve months equal to the yearly values calculated from the futures prices.

The Phelix yearly futures markets have low trading volume and almost none more

than three years in advance. When no volume is traded, the price is approximated as the mean of estimates elicited from market participants so does not reflect market conditions as well as prices from actual trades. Low trading volumes for contracts with delivery years in the future contributes to the substantial uncertainty in the change in intraday price differences. Further, yearly futures prices for trades in late 2011 show a weaker increase in intraday price differences than the 2010 prices on which we base our primary results. We address these factors through a sensitivity analysis in Section 4. Figure 4 shows simulated trajectories of the expected real change in intraday price difference for a range of nominal yearly growth rates for intraday prices (with $b = 0.08$ as the base scenario), which form the basis of the sensitivity analysis.




Figure 4: Simulated trajectories of expected future short-term price volatility by year with $b = 0.08$ (calculated from the futures market data) as the base scenario. β scales the real growth in short-term price volatility from 2011 values and b is the nominal annual percentage-point growth in price volatility.

While a bottom-up model of the German power system could be a better source of simulated future prices, such models tend to be tuned to capture other features and therefore underestimate intraday price volatility. No available model adequately captured

the effect of increased wind capacity on intraday prices. As result, the price model used in this analysis accounts for the change in price volatility but not the change in price structure due to an increased share of renewables. The assumed linear growth rate of price volatility also neglects long-lasting shocks due to policy interventions or increased transmission capacity. Nevertheless, the magnitude of intraday price volatility is more important than the functional form of its increase, supporting the use of the linear model. Under the base scenario, the expected real price volatility increase of 50 % by 2020 is near that predicted by a detailed power systems model of Europe in 2020 (60 % above 2010 values) under no substantial increase in transmission capacity (Schaber et al., 2012).

2.2 Pumped Storage Scheduling Model

The pumped storage scheduling model generates the optimal production schedule and expected revenue given hourly spot prices. Operation of the pumped storage system is optimized over a planning horizon of $n = 168$ hours (one week). For hours $i \leq m = 24$, the spot price p_i is known, and for the remaining hours $m < i \leq n$ the spot price is approximated as the mean price for the corresponding hour one week earlier and two weeks earlier and is denoted by \bar{p}_i . Optimizing production over a week instead of a day gives the system the flexibility to store energy on weekends, for example, when prices are lower and produce energy at times of week when prices are higher. Operating the system for a single day before re-optimizing allows new price information to be incorporated on a daily basis without assuming future knowledge of prices beyond a day. Since generators make price-dependent bids, scheduling based on a day of known prices is approximately equivalent to scheduling based on knowledge of the price distribution, which is a weaker assumption than perfect price knowledge and produces nearly the same results.

The following linear program maximizes profit for one week:

$$\max \left(\sum_{i=1}^m p_i \cdot \left[E_i \cdot (1 - h) - S_i \cdot \frac{1}{1 - h} \right] + \sum_{i=m+1}^n \bar{p}_i \cdot \left[E_i \cdot (1 - h) - S_i \cdot \frac{1}{1 - h} \right] \right) \quad (3a)$$

subject to

$$0 \leq R_0 + \sum_{i=1}^k \left(S_i \cdot \eta_S - \frac{E_i}{\eta_E} \right) \leq R_{\max} \text{ for all } k \leq n \quad (3b)$$

$$\sum_{i=1}^n \left(\frac{E_i}{\eta_E} - S_i \cdot \eta_S \right) = R_0 - f \cdot R_{\max} \quad (3c)$$

$$0 \leq S_i \leq S_{\max} \quad (3d)$$

$$0 \leq E_i \leq E_{\max} \quad (3e)$$

Parameters and variables in the pumped storage scheduling model are summarized in Table 1. $E_i \leq E_{\max}$ is the energy produced by the plant at efficiency η_E in hour i , $S_i \leq S_{\max}$ is the energy stored at pump efficiency η_S in hour i , R_{\max} is the capacity of the reservoir in MWh, R_0 is the reservoir level at the beginning of the week in MWh, h is the transmission loss, and f is the target fractional reservoir level at the end of the optimization period. Constraint (3b) ensures that the reservoir level neither exceeds the maximum nor drops below zero and (3c) ensures that the calculated reservoir level at the end of the optimization period is equal to $f \cdot R_{\max}$, in order to avoid an unrealistic plan of completely emptying the reservoir at the end of each period (for sufficiently long planning horizons, this constraint would not be necessary).

Profit earned during a period of $m = 24$ hours is calculated as

$$\sum_{i=1}^m p_i \cdot \left[E_i^* \cdot (1 - h) - S_i^* \cdot \frac{1}{1 - h} \right] \quad (4)$$

in which E_i^* and S_i^* are the optimal amounts of energy produced and stored. We assume maintenance is conducted during times when the opportunity cost of operation is negligible,

and we multiply revenue by a factor of $1 - a = 0.95$ to account for unplanned outages.

Table 1: Parameters and variables in pumped hydropower storage scheduling. Subscript i denotes quantities that may change hourly.

Symbol	Explanation	Value	Unit
p_i	Simulated spot price	-	euro/MWh
\bar{p}_i	Forecasted (rolling-average) price	-	euro/MWh
m	Length of operation period (known prices)	24	h
n	Length of optimization period	168	h
E_i	Energy produced	-	MWh
E_i^*	Optimal energy produced	-	MWh
E_{\max}	Maximum generator capacity	480-2,400	MW
η_E	Generator efficiency	0.9	-
S_i	Energy stored	-	MWh
S_i^*	Optimal energy stored	-	MWh
S_{\max}	Maximum pump capacity	480-2,400	MW
η_S	Pump efficiency	0.80	-
f	Target fractional reservoir level at end of week	0.5	-
R_{\max}	Maximum reservoir capacity	75,000	MWh
h	Transmission loss factor	0.05	-
a	Total unplanned outages of plant and cable	0.05	-

Startup costs, reduced generator efficiency at full power, and head effects are ignored. Implementing approximations of these effects significantly increased computation time but reduced profit by less than 1 %. The location of the pumped storage plant within a large reservoir system would relax reservoir capacity constraints due to greater flexibility in the system as a whole; however, the plant may be subject to occasional restrictions such as forced operation to avoid spillage. These effects are discussed further in Section 4.

2.3 Real Option Valuation

Classical methods of project planning dictate that investment should be made if expected NPV is positive. This investment strategy ignores the value of postponing investment in a positive-NPV project in anticipation of resolution of uncertainty or expected future change in conditions determining the project value. For projects in which these factors are substantial, real options captures the value of flexible decision making and intertemporal

choice. The method is based on comparing the current NPV with the expected value of postponing investment (the continuation value), and investing only if the former exceeds the latter. A holding cost, which we do not incorporate, would decrease the value of postponing investment and accelerate investment timing.

Since the pumped storage investment opportunity is characterized by sunk costs, irreversibility, and a time interval during which investment is possible, we model it as a Bermudan call option.⁴ We calculate the value of the investment opportunity expiring T years in the future as

$$ROV = \max_{t \leq T} (\mathbb{E}[e^{-r \cdot t} \cdot (PV_t - F)], 0) \quad (5)$$

where ROV is the value of the investment opportunity with expiration time T , PV_t is the present value of revenue when the option is exercised at time $t \leq T$, F is the investment cost, r is the discount rate, and \mathbb{E} is the initial expected value operator. If the NPV at the optimal investment time is positive, then the option value equals the NPV and the investment should be made; if the NPV is negative, the option is worthless.

The method used to value the real option to invest in the pumped storage system is Least Squares Monte Carlo (LSM), proposed by Longstaff and Schwartz (2001). Unlike the methods of Dixit and Pindyck (1994) and Décamps et al. (2006), based on partial differential equations, or the lattice method of Cox et al. (1979), LSM is free from the assumption that revenue streams follow a prescribed stochastic process. Since LSM is simulation-based, it can accommodate exotic option characteristics such as multiple project choices, dynamic uncertainties, and disjoint exercise intervals. LSM can be computationally intensive, is not as precise as other methods, and requires time discretization; none of these disadvantages, however, impeded its use in this work.

⁴A Bermudan call option allows the owner to purchase an asset for a given exercise price at given points in time during the option lifetime. Here, the asset value is the present value of revenue to the pumped hydropower storage system, the exercise price is the investment cost, and the option can be exercised yearly. The option valuation method used in this study requires discrete exercise times, although an investment decision could be made more frequently than once per year, this frequency is adequate to yield insight into our problem.

Starting with the option value at the expiration time T (equal to the NPV since postponing investment is worthless), the LSM procedure iterates backward over the option lifetime to find the expected value of investing and that of continuing to hold the option at each time step. The expected value of continuation at time t is found by regressing the maximum values of investing at a later time on basis functions of the simulated present values at time t . If the NPV exceeds the expected continuation value for a given simulation, t becomes the optimal investment time for that simulation, and the backward iteration continues.

Longstaff and Schwartz (2001) argue that the choice of basis function in the regression has little effect on results. Clément et al. (2002) show that LSM converges almost surely under general conditions, and Gamba (2002) extends LSM to value mutually exclusive investment alternatives and shows that the convergence results of Clément et al. (2002) apply. We use basis functions given by

$$\mathbb{E}[Y|X] = \sum_{i=0}^k \alpha_i \cdot X^i \quad (6)$$

in which X is the current present value for a single project (which, in our numerical example, is correlated with the other project values with $\rho > 0.99$), Y is the discounted continuation value, and the α_i are estimated regression parameters. We find that basis functions of higher order than $k = 3$ do not improve results.

The option value at the initial time is calculated as the mean over all simulations of the discounted payoff of investing at the expected optimal time. In summary, the aim of the valuation framework is to conduct a real options analysis for the investment opportunity at hand. Annual revenue from operating the pumped hydropower storage facility serve as input to the real options analysis and is found by first simulating future spot prices (as outlined in Section 2.1) and then calculating profit-maximizing production schedules (as outlined in Section 2.2).

This method could be used to value investment opportunities in other industries or settings in which future revenue is uncertain but can be projected as a random process. Though we assume investment is irreversible and cost is certain, the method could be adapted to relax these assumptions. Larger uncertainty would increase the number of simulations required for convergence, and with computationally intensive revenue calculations, could render the method impractical.

3 Numerical Example

The investment opportunity considered is a unique, previously secured right to construct a pumped storage facility, modeled after the Tonstad power plant in southern Norway, which operates solely on the German power market through an HVDC connection. Upriver of Tonstad, there are three reservoirs which are small compared to others in the system and serve mainly as short-term storage for the existing power plant. These form the upper reservoir capacity in this analysis (see Figure 1). We assume the pumped storage plant has no inflow and operates in isolation from the rest of the reservoir system. Since the reservoirs already exist and do not require adaptation, their cost is set to zero. The cost of securing the investment option is negligible since there is no permitting fee, the current Tonstad plant owner Sira Kvina already possesses the majority of the necessary land rights, and the cost of assembling an application is approximately 130,000 euro.⁵

The investor can choose among pumped storage plants of capacities between 480 and 2,400 MW, in 480 MW increments (the proposed upgrade to Tonstad is 960 MW, consisting of 2×480 MW pump/generator units). The projects are mutually exclusive with no option to upgrade a smaller project to a larger project. Although the pumped storage

⁵Although the success of obtaining a permit has a small degree of uncertainty, incorporating this uncertainty would require a two-stage option valuation framework that is outside the scope of this analysis. Dixit and Pindyck (1994) have proposed a method for valuing such two-stage options; however, due to the low cost and large likelihood of success in obtaining a permit, the value of the two-stage option would be very close to that of the single-stage option we consider. We therefore choose to ignore uncertainty in the permitting stage.

system is considered an upgrade to the existing Tonstad plant, it would involve the construction of a new water tunnel and the installation of new machinery such that its cost is commensurate with building an entirely new plant. The investment opportunity includes construction of a HVDC transmission line from the power plant to the coast and a subsea HVDC cable between Norway and Germany. We assume that the investor can purchase HVDC transmission of equal capacity to the pumped storage plant and that the cost is proportional to that of the proposed 1,400 MW NorGer cable between Norway and Germany (NorGer, 2009). Since licenses issued by the Norwegian Water Resources and Energy Directorate last five years with the possibility of a five-year extension (NVE, 2010), we assume the investment opportunity lasts ten years beginning in 2011. We later examine the effect of extending the investment window to 20 years. Revenue begins three years after construction costs are incurred, and the lifetime of the power plant and the HVDC cable is set to 40 years.

Table 1 presents parameters and variables for the pumped storage plant used in the numerical example. Capital cost amounts to 0.66 million euro per MW of installed capacity for the pumped storage plant and transmission to the Norwegian coast, and with the subsea HVDC cable the cost becomes 1.71 million euro per MW. We neglect economies of scale with respect to capacity, which could shift decisions toward larger projects. Capital costs include expenses related to financing during the construction period. We assume a before tax, real discount rate of 6 %, in accordance with the Norwegian government's discount rate for high-risk projects (NDR, 2011). Since discount rate is uncertain a sensitivity analysis is presented in Section 4.

The cost excluding the subsea cable is less than the estimates for the planned pumped storage plants in Switzerland (0.75 million euro per MW), Austria (0.74 million euro per MW), and Germany (0.70 million euro per MW). The amount of pumped storage capacity available in each of these countries, however, is less than the 2,400 MW modeled as the largest option for the Tonstad upgrade.

4 Results

Under the base scenario of a nominal growth rate in short-term price volatility of 8 percentage points annually, a profit-maximizing investor following a strategy based solely on NPV would invest immediately in the largest available project since initial expected NPV is 880 million euro. Valuing the flexibility to postpone the investment decision up to ten years in order to resolve uncertainty and reap possible future gains shows that the investment opportunity is worth 1.1 billion euro, 25 % more than if investment were restricted to a now-or-never decision in the first year. At the beginning of the option lifetime, the expected optimal strategy is to wait approximately eight years and then invest in the largest available project, with revenue streams beginning three years later.

The NPV of the maximum-value project and the continuation value at each year of the ten-year option lifetime for the base scenario are shown in Figure 5. As the option nears expiration, the ability to delay investment adds less value to the opportunity, and the continuation value becomes equivalent to NPV at the optimal investment time (approximately the eighth year).⁶ The threshold NPV curve shows the NPV above which investment would be optimal at each timestep. Since only expected NPV can be observed, this curve cannot directly form the basis for an investment decision rule and the optimal strategy must be updated as more information becomes available. The trend of the threshold NPV curve illustrates that investment in earlier years would be optimal only if current NPV is significantly higher than the expected NPV, but as the option nears expiration, the value of holding the option decreases and investment is optimal at lower NPVs. After the optimal investment time, the required NPV for investment is lower than the average NPV in the given year. The threshold NPV and continuation value both drop to zero in the last year of the option lifetime, since in that year the investment opportunity is now-or-never and investment is optimal if the NPV of the maximum-value project is

⁶Continuation values appear to increase with time because they are discounted to the time of investment; if values were discounted to a uniform time they would of course decrease, reflecting the reduced flexibility as the option nears expiration.

positive.

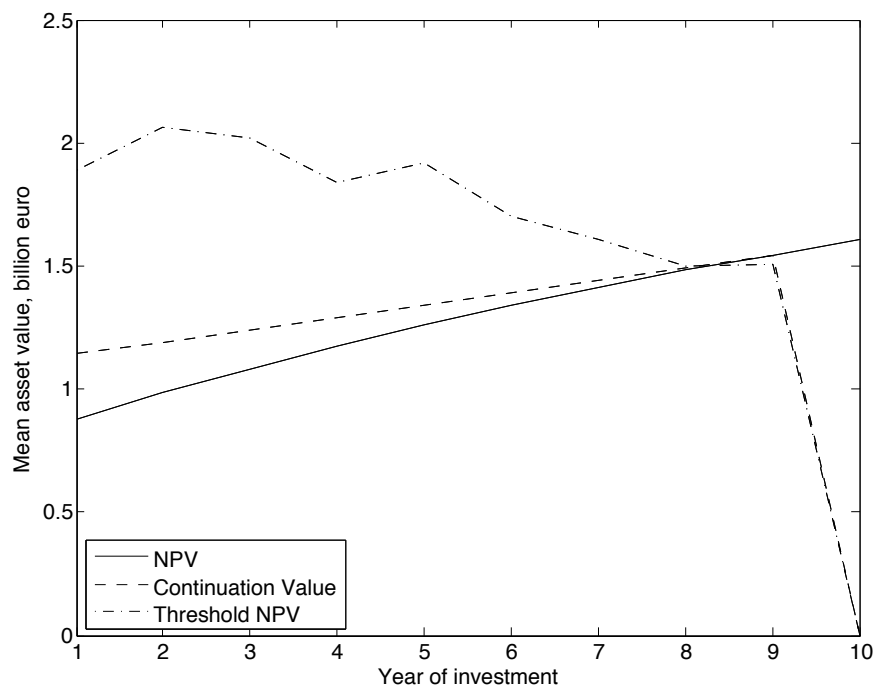


Figure 5: NPV, continuation value (expected value of postponing investment), and threshold NPV for the base scenario. For values above the threshold NPV, investment would be optimal if NPV could be observed. In the eighth year, the expected optimal investment time, the threshold is equal to the mean NPV.

Figure 6 shows the NPV of each project and the continuation value (the expected value of postponing investment) for the base scenario at the expected optimal investment time (the eighth year). All project values are plotted against that of the 2,400 MW project in order to obtain a consistent basis of comparison; since all projects see the same price streams, their present values are almost perfectly correlated. If the present value of the 2,400 MW project is greater than 5.7 billion euro (the value at which maximum NPV and continuation value intersect), investment would be more profitable than continuing to hold the option.

Since profitability of the pumped storage plant is driven by price arbitrage, option value increases with growth in short-term price volatility (b in Equation (2)). Table 2 gives the NPV, option value, mean investment time, and optimal project size for the range of growth rates in short-term price volatility presented in Figure 4.

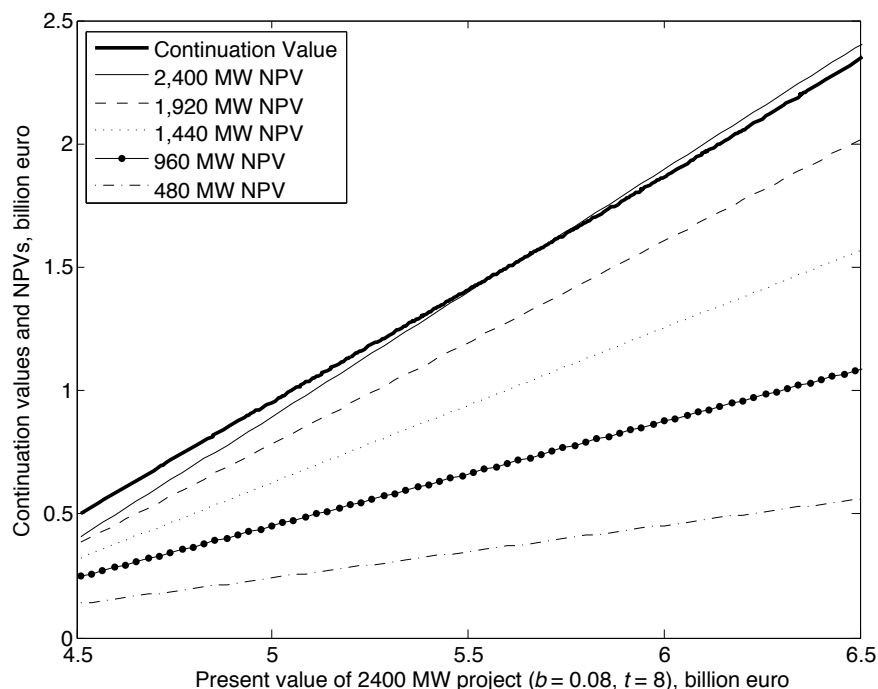


Figure 6: Continuation value (expected value of postponing investment) and NPVs of each project at the optimal investment time (eighth year of the option lifetime) under the base scenario. All project values are plotted against the 2,400 MW present value for a consistent basis of comparison.

Greater increase in short-term price volatility promotes early investment in large projects since revenue grows quickly in early years and is large enough to overcome a high capital cost. With earlier optimal investment times, the option to postpone adds less value relative to the initial NPV. For nominal growth rates in price volatility of 8 percentage points and higher, almost all simulated NPVs are positive and NPV grows linearly with the growth rate; at lower values of the growth rate, the linear relationship does not hold.

In the base scenario, the only realized capacity choice is 2,400 MW such that the investment decision is binary: either invest in the largest project or not at all. If there is little or no increase in nominal short-term price volatility above historical values (0 or 2 percentage points), then arbitrage revenue cannot overcome investment cost and the option is worthless. For a yearly increase in short-term price volatility of between 4 and 6

⁷*b* is the mean nominal yearly increase in price standard deviation. Mean optimal investment time and the most frequent optimal project size are calculated from only those scenarios for which investment in any project during the option lifetime is profitable (excluding scenarios in which no investment occurs).

Table 2: Results for a range of growth rates in intraday price differences (base scenario in bold).⁷

Value of b	0.00	0.02	0.04	0.06	0.08	0.10	0.12	0.14
NPV (million euro)	-530	-380	-190	65	880	1,900	3,000	4,100
Option value (million euro)	0	0	0.77	240	1,100	2,100	3,200	4,200
Mean investment time (year)	-	-	8.5	9.9	8.4	7.1	6.0	5.2
Project size (MW)	-	-	480	2,400	2,400	2,400	2,400	2,400

percentage points, on the other hand, investment can be optimal in a small project if present values are low, a large project if present values are high, and not at all if present values are near their mean, resulting in a discontinuous investment region. Figure 7 shows all NPVs and the continuation value for nominal intraday increase of 6 percentage points in the ninth year of the option lifetime, when investment in the 480 MW project is optimal at low present values, investment in the 2,400 MW project is optimal at high present values, and a waiting region exists at intermediate present values. Value is created through the option to postpone the investment decision until it becomes clear which project is more profitable.

To examine the effect of the assumed functional form for β , the base case was run with (1) β modeled as a discrete-time Brownian motion with drift and (2) no uncertainty in the nominal annual growth in price volatility ($\sigma = 0$). The first case reflects an assumption that is more common in real options analysis but does not fit the futures market data used to construct β in Section 2.1. In the second case, the price paths vary only according to the sampled historical data they contain. Option values using the alternative functions for β are the same as the base case to two significant digits. Optimal investment time for (1) is 8.6 years and for (2) is 8.7 years into the ten-year option lifetime, compared with 8.4 years under the base case, and 2,400 MW remains the optimal capacity choice in all cases. With $\sigma > 0$, short-term price fluctuations create more volatile revenue streams and thus a greater spread in optimal investment time. With investment time curtailed to 10 years, more variability in optimal investment time tends to slightly reduce its mean, counter to the




Figure 7: Continuation value and NPVs of each project at $t = 9$ years for a nominal increase in short-term price volatility of 6 percentage points. Two investment regions occur: at low present values, investment in the smallest project is optimal, and at high present values investment in the largest project is optimal.

common real options effect of delaying investment under greater uncertainty. Results are thus robust to the functional forms for β examined.

The investment cost estimate used in this study, based on communications with Sira Kvina personnel, is nearly 70 % higher (inflation-adjusted) than that on the 2007 permit application for the same sized plant (Sira Kvina, 2007). Since this estimate is uncertain, a cost sensitivity analysis was conducted for the base scenario (Table 3). Reducing cost by 10 % or 25 % substantially increases NPV, accelerates the optimal investment time, and reduces the added value of the option to postpone investment. Raising costs by 10 % or 25 %, on the other hand, delays the optimal investment time such that the high capital cost is maximally discounted. Updating the cost estimate for future economic conditions would allow optimal investment timing to be refined.

In reality, the project could be initiated after the end of ten-year option lifetime. To examine the effect of the assumed investment window, the option lifetime was extended to

Table 3: Sensitivity of base scenario results to investment costs.

Deviation from base scenario investment costs	-25 %	-10 %	0 %	+10 %	+25 %
NPV (million euro)	1,900	1,300	1,100	490	90
Option value (million euro)	1,500	1,700	1,400	850	430
Mean investment time (year)	4.7	7.2	7.8	9.4	10
Project size (MW)	2,400	2,400	2,400	2,400	2,400

20 years with the last ten years of revenue approximated as the mean of the final year of simulated annual revenue. For the longer option lifetime, the option value in the base scenario increases by less than 10 % and the optimal investment time is the tenth year, showing that results are robust to option lifetime. Since short-term price volatility experiences nearly no growth after 2040 (see Figure 4), as the option lifetime progresses there is less incentive to postpone investment and face more heavily discounted revenue.

Considering how power plants of different sizes utilize the reservoir capacity, Figure 8 shows two weeks of operation of the pumped storage facility for the five power plant size alternatives. Larger power plants use a wider range of the reservoir capacity than smaller power plants and are more constrained by capacity limits. When the reservoir level for a large power plant reaches the upper or lower bound on reservoir capacity, expanding the capacity would increase profitability. All power plants profit from both intraweek and intraday price arbitrage, as evidenced by the weekly and daily periodicities in the reservoir levels.

Examining sensitivity of profit to reservoir capacity, Figure 9 shows how the revenue from a 960 MW and a 2,400 MW pumped hydropower storage facility would change with 25 % and 50 % reductions and increases in storage capacity. While expanding the reservoir capacity leads to a small increase in revenue, decreasing the capacity by an equivalent amount results in a greater loss. Larger power plants are more affected by changes in reservoir capacity since they operate more frequently near the reservoir limits, as shown in Figure 8. The optimal generator capacity under the base scenario (2,400 MW) is unaffected, and the investment time is pushed back one year to the ninth year of the option




Figure 8: Reservoir level during two weeks of operation for all five pumped storage capacity alternatives. Reservoir capacity constrains operation, and therefore reduces profit, most frequently for larger plants. Water levels show both daily (24 h) and weekly (168 h) periodicity.

lifetime only for the smallest reservoir capacity. Investment timing and capacity choice are thus robust to reservoir capacity.

Reservoir capacity available to the pumped storage plant depends on the operation of interconnected reservoirs in the system, and consistent operation of the pumped storage plant and existing hydropower plants would create more flexibility in the system as a whole. The assumed reservoir capacity could therefore artificially restrict the operation of the pumped storage plant in a way that constrains profitability of both the pumped storage and the rest of the system. Evaluating the benefit of additional flexibility to the system as a whole is beyond the scope of this analysis.

Figure 10 illustrates the sensitivity of the option value to discount rate. Under the base scenario, the option is worth less than 100 million euro at discount rates of 8 % or more and becomes worthless at a discount rate of 11 %. A 6 % discount rate is common for energy project planning in Norway, but investors in riskier projects use up to 10 %.




Figure 9: Sensitivity of profit to reservoir size for the 960 MW plant (proposed as the Tonstad upgrade) and 2,400 MW plant (the optimal size in the current study). Increasing reservoir size by 50 % results in only a 3 % profit increase for the 2,400 MW plant and a 1 % increase for the 960 MW plant. Decreasing reservoir size by 50 %, in contrast, leads to a 3 % drop in profit for the 960 MW plant and a 9 % drop for the 2,400 MW plant.

The effect of discount rate on optimal investment time is nonlinear for the base scenario, with investment in the tenth (final) year at a discount rate of 2 % but dropping steadily to the seventh year at a discount rate of 5 % and rising again to the final year at discount rates of 7 % and above. Since price volatility increases with time, a low discount rate encourages waiting to capture greater revenue. A mid-range discount rate encourages accelerated investment so that revenue streams are discounted less heavily, and at high discount rates revenue in the future is no longer sufficient to overcome high up-front capital cost so investment is postponed to maximally discount capital cost. The optimal project size shrinks as discount rate increases as well, from 2,400 MW at discount rates from 0 to 7 % steadily down to 480 MW at 10 %, reflecting the dominance of low capital cost over large revenue streams as discount rates increase. In summary, if investors decide the project is risky enough to warrant a higher discount rate, optimal investment is later and

in a smaller project than if the standard 6 % discount rate is used.




Figure 10: Sensitivity of option value to discount rate. The high up-front capital cost and relatively long project lifetime render profitability highly sensitive to choice of discount rate.

5 Conclusions and Further Work

Under a base scenario of growth in price volatility derived from the EEX futures market, a prospective investor in a pumped hydropower storage facility in Norway seeking to profit from arbitrage in the German market should hold the option for approximately eight years and then invest, updating the strategy as more information becomes available. The option to postpone investment provides substantial added value, due to both deterministic growth in price volatility as well as greater uncertainty in revenue combined with the ability to avoid investment if a project proves to be unprofitable. Using the options approach rather than an NPV analysis quantifies the benefit of waiting to invest, and the robustness of the result to doubling the option lifetime shows that the waiting period is not simply an artifact of the options problem framing. Progress toward resolving uncertainty on the price

effect of wind will enable investors to refine optimal investment strategies.

The sensitivity analyses illustrate ways in which resolution of uncertainty could change the optimal investment strategy. While the optimal project size is robust to cost and reservoir capacity, it decreases with decreasing growth in short-term price volatility and increasing discount rates. Optimal investment time is pushed forward with larger growth in short-term price volatility, lower cost, and larger reservoir sizes, whereas higher discount rates have nonlinear effects on investment timing. Consistent across sensitivity analyses was the result that the option to postpone investment created value above that captured in an NPV analysis.

The effect of increased wind power penetration on short-term price volatility will depend on factors such as the viability of base load generators in the changing market, the amount of new transmission capacity integrating European electricity markets, the price of natural gas, and the greenhouse gas regulatory regime. We have encountered neither explicit modeling of these effects nor price scenarios from bottom-up models that sufficiently capture uncertainty, have fine enough time resolution, and have long enough time span for use in this study. Updating the price model used in this work with better information on the listed effects will strengthen the conclusions, and is left for future work. This work serves best as a decision support tool to provide guidance on the optimal investment strategy in a pumped storage system given investor expectations on the mean increase in short-term price volatility.

Acknowledgements

The authors thank Bjørn Bakken, Michael Belsnes, and Atle Harby of SINTEF Energy Research; Stein-Erik Fleten from the Norwegian University of Science and Technology (NTNU); Arne Sæterdal and Rolv Guddal from Sira Kvina Kraftselskap; and anonymous reviewers for helpful input. This work was supported by FINERGY Project 178374, the

Center for Sustainable Energy Studies (CenSES) at NTNU, the National Science Foundation (NSF) Graduate Research Fellowship Program, and the Research Council of Norway and NSF through the Nordic Research Opportunity. This work was also supported in part by grants from the Alfred P. Sloan Foundation and EPRI to the Carnegie Mellon Electricity Industry Center; from the Doris Duke Charitable Foundation, the Department of Energy National Energy Technology Laboratory, and the Heinz Endowments to the RenewElec program at Carnegie Mellon University; and from the U.S. National Science Foundation under Award no. SES-0345798.

Bibliography

- BMU (2010). Bundesministerium für Umwelt, Naturschutz und Reaktorsicherheit (BMU) [German Federal Ministry for the Environment, Nature Conservation, and Nuclear Safety]. Technical Report 12-2010.
- Clément, E., Lamberton, D., and Protter, P. (2002). An analysis of a least squares regression method for American option pricing. *Finance and Stochastics*, 6(4):449–471. 10.1007/s007800200071.
- Connolly, D., Lund, H., Finn, P., Mathiesen, B. V., and Leahy, M. (2011). Practical operation strategies for pumped hydroelectric energy storage (PHES) utilizing electricity price arbitrage. *Energy Policy*, 39(7):4189–4196.
- Cox, J. C., Ross, S. A., and Rubinstein, M. (1979). Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263.
- Deane, J., Ó Gallachóir, B., and McKeogh, E. (2010). Techno-economic review of existing and new pumped hydro energy storage plant. *Renewable and Sustainable Energy Reviews*, 14(4):1293–1302.
- Décamps, J., Mariotti, T., and Villeneuve, S. (2006). Irreversible investment in alternative projects. *Economic Theory*, 28(2):425–448.
- Dixit, A. (1993). Choosing among alternative discrete investment projects under uncertainty. *Economics Letters*, 41(3):265–268.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment Under Uncertainty*. Princeton University Press.
- European Central Bank (2011). Monetary policy [accessed November 2011]. <http://www.ecb.int/stats/prices/hicp/html/index.en.html>.
- Gamba, A. (2002). An extension of least squares Monte Carlo simulation for multi-option problems. In *The Proceedings of the 6th Annual Real Options Conference*, Cyprus.
- Gyuk, I. (2003). *EPRI-DOE Handbook of Energy Storage for Transmission and Distribution Applications*. Electric Power Research Institute (EPRI) and the U.S. Department of Energy (DOE). EPRI Report 1001834.
- Hadjipaschalis, I., Poullikkas, A., and Efthimiou, V. (2009). Overview of current and future energy storage technologies for electric power applications. *Renewable and Sustainable Energy Reviews*, 13(6–7):1513–1522.
- Lehner, B., Czisch, G., and Vassolo, S. (2005). The impact of global change on the hydropower potential of Europe: a model-based analysis. *Energy Policy*, 33(7):839–855.
- Loisel, R., Mercier, A., Gatzert, C., Elms, N., and Petric, H. (2010). Valuation framework for large scale electricity storage in a case with wind curtailment. *Energy Policy*, 38(11):7323–7337.

- Longstaff, F. and Schwartz, E. (2001). Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*, 14(1):113–147.
- Muche, T. (2009). A real option-based simulation model to evaluate investments in pump storage plants. *Energy Policy*, 37(11):4851–4862.
- NDR (2011). Veileder i samfunnsøkonomiske analyser [Guide to economic analysis]. Technical report, Norway Department of Finance.
- Nicolosi, M. and Fürsch, M. (2009). The impact of an increasing share of RES-E on the conventional power market — the example of Germany. *Zeitschrift für Energiewirtschaft*, 33(3):246–254.
- NorGer (2009). Konesjonssøknad, likestrømsforbindelse (1400 MW) mellom Norge og Tyskland, [Application for license for constructing a direct current transmission line (1400 MW) between Norway and Germany]. Technical Report 200704983, Norwegian Water Resources and Energy Directorate.
- NVE (2010). Konesjonshandsaming for vasskraftsaker, [Processing of license applications for hydropower plants]. Technical Report 3/2010, Norwegian Water Resources and Energy Directorate.
- Phelix (2011). Physical electricity index futures prices for Germany and Austria [accessed February 2011]. <http://www.eex.com/en/Download>.
- Schaber, K., Steinke, F., and Hamacher, T. (2012). Transmission grid extensions for the integration of variable renewable energies in Europe: Who benefits where? *Energy Policy*, 43:123–135.
- Schill, W. P. and Kemfert, C. (2011). Modeling strategic electricity storage: The case of pumped hydro storage in Germany. *The Energy Journal*, 32(3):59–87.
- Sira Kvina (2007). Konesjonssøknad, tilleggsinstallasjon i Tonstad kraftverk med mulighet for pumping, [Application for license for additional installation in Tonstad power plant with opportunities for pumping]. Technical Report 3475/200707786, Norwegian Water Resources and Energy Directorate.
- SRU (2010). Climate-friendly, reliable, affordable: 100 % renewable electricity supply by 2050. Technical report, German Advisory Council on the Environment (SRU).
- Weigt, H. (2009). Germany’s wind energy: The potential for fossil capacity replacement and cost saving. *Applied Energy*, 86(10):1857–1863.

Paper 3

The effect of long-distance interconnection on wind power variability

Abstract

This paper uses time- and frequency-domain techniques to quantify the extent to which long-distance interconnection of wind plants in the United States would reduce the variability of wind power output. Previous work has shown that interconnection of just a few wind plants across moderate distances could greatly reduce the ratio of fast- to slow-ramping generators in the balancing portfolio. The current paper finds that interconnection of aggregate regional wind plants would not reduce this ratio further but would reduce variability at all frequencies examined. Further, interconnection of just a few wind plants is found to reduce the average hourly change in power output, but interconnection across regions provides little further reduction. Interconnection also reduces the magnitude of low-probability step changes and doubles firm power output (capacity available at least 92 % of the time) compared with a single region. First-order analysis indicates that balancing wind and providing firm power with local natural gas turbines would be more cost-effective

This paper was published as Fertig, E., Katzenstein, W., Jaramillo, P., and Apt, J. (2012). The effect of long-distance interconnection on wind power variability. *Environmental Research Letters*, 7, doi:10.1088/1748-9326/7/3/034017.

than with transmission interconnection. For net load, increased wind capacity would require more balancing resources but in the same proportions by frequency as currently, justifying the practice of treating wind as negative load.

1 Introduction

Wind power is among the least costly and most developed renewable energy technologies, which renders it well suited to fulfilling the renewable energy targets currently implemented in most U.S. states. Between 2005 and 2010, installed wind capacity in the U.S. increased by a factor of 4.4 and net wind generation by a factor of 5.3 (Wiser and Bolinger, 2011). As wind capacity continues to grow, the variability and intermittency of wind power can create challenges for grid operators. High frequency, second-to-second fluctuations can increase the need for frequency regulation, and lower frequency (hourly to seasonal) fluctuations can change the capacity factors of baseload generators and in severe cases affect reliability. Wind integration studies have suggested that building transmission capacity to interconnect wind power plants could greatly smooth wind power output (Zavadil, 2006; IEA, 2005; EnerNex, 2011; GE Energy, 2010; EERE, 2008), but few explicitly account for the frequency at which the variability occurs.

Katzenstein et al. (2010) performed the first frequency-dependent analysis of the smoothing effect of interconnecting wind plants. Using 15-minute energy output data from 20 wind plants in the Electric Reliability Council of Texas (ERCOT), Katzenstein et al. (2010) find that at a frequency of $(1 \text{ h})^{-1}$ ($2.8 \times 10^{-4} \text{ Hz}$) interconnecting just four wind plants reduces the ratio of high- to low-frequency variability by 87 % compared with a single wind plant, but that connecting additional wind plants yields diminishing returns. At a frequency of $(12 \text{ h})^{-1}$, interconnecting four wind plants reduces this ratio by only 30 % compared with a single plant. Variability reduction was found to depend on factors such as size and location of wind plants as well as the number interconnected.

This result highlights the importance of time scale in characterizing wind power

smoothing and suggests that interconnecting a relatively small number of wind plants could achieve most of the reduction in the ratio of high- to low-frequency variability that would result from interconnecting many more. This ratio is one determinant of the relative requirements for fast and slow ramping sources required to compensate for wind's variability. Katzenstein et al. (2010) limited their study to west-central Texas, where there may be a correlation of weather and wind patterns, and did not examine the effect of wind plant interconnection on the variability of net load (electricity load minus wind power output). Building upon Katzenstein et al. (2010), our work uses frequency-domain analysis to examine both the smoothing effect of interconnecting wind plants across greater distances and the variability of net load under greater wind power penetration (see Appendix 2).

Sørensen et al. (2008) use frequency domain techniques to analyze the reduction in wind power output variability due to interconnecting individual wind turbines within a single offshore wind plant. The smoothing effect is modeled at time scales from minutes to hours and found to be strongest at high frequencies. The analysis and results of Sørensen et al. (2008) are similar to ours despite the difference in scale, highlighting the fractal property of wind energy.

Wind power variability studies utilizing exclusively the time domain include Giebel (2000), Ernst et al. (1999) and Sinden (2007). These studies find that correlation of wind power output decreases predictably as the distance between the wind plants increases but is still slightly positive even for widely separated plants.

Further characterizing geographic smoothing, Degeilh and Singh (2011) introduce a method for selecting from a set of geographically separated wind sites to minimize wind power output variance and show that achieving this objective yields the smallest loss of load probability (LOLP) as well. Kempton et al. (2010) use offshore meteorological buoy data from 2500 km along the U.S. east coast to analyze the effect of interconnecting 11 wind sites and find that interconnection reduces the variance of simulated power output,

slows the rate of change, and eliminates hours of zero production during the five-year study period. Kempton et al. (2010) conclude that the cost of mitigating wind variability with long-distance transmission interconnection has a cost on par with current methods of balancing wind. Dvorak et al. (2012) use mesoscale wind data to find the best locations for four offshore wind plants near the U.S. east coast to reduce variability, hourly ramp rates, and hours of zero power. The latter two studies approximate wind power output using wind speed measurements taken significantly below turbine hub height. Though buoy data are the best available until hub-height met masts become widespread and generate an extensive record, they are often of poor quality and can exaggerate estimated wind power variability (Holttinen, 2005).

In this paper, we analyze the extent to which interconnecting wind plants over broad geographical regions of the United States will reduce variability of wind power output. We use simultaneous wind energy data from four regions (the Bonneville Power Authority (BPA), the Electric Reliability Council of Texas (ERCOT), the Midwest ISO (MISO), and the California ISO (CAISO)) and apply methods suggested in Katzenstein et al. (2010) and Sørensen et al. (2007). The analysis informs the question of whether increasing interregional transmission capacity is an effective means of smoothing wind power output. Section 2 describes the data and methods, Section 3 presents results, and Section 4 discusses implications and concludes.

The observed data show that interconnection of regional wind resources increases the percentage of firm wind power capacity, reduces the coefficient of variation of wind power output, and reduces the likelihood of extreme step changes. Although step changes are one metric for evaluating variability, frequency-domain analysis can help establish the portfolio of generation needed to compensate for variability. If the amplitude of high frequency variations is the same as that of low frequency variations, as much fast-ramping generation must be available as slow-ramping generation. On the other hand, if interconnection is able to reduce the fast fluctuations, much less fast-ramping generation will be required.

Katzenstein et al. (2010) found that interconnecting 4 or 5 wind plants achieves the majority of the reduction in the ratio of high to low frequency fluctuations. Because an asymptote is quickly reached, it is not surprising that we find large scale interconnection does not further reduce this ratio, and that variability reduction at the relevant frequencies could be achieved as effectively by interconnection within regions as between regions. Likewise, inter-regional interconnection does not significantly affect mean step changes in hourly wind power output; the majority of the reduction in mean step changes is achievable through interconnection of wind plants within single regions.

2 Data and methods

We use wind energy output and load data from BPA, CAISO, ERCOT, and MISO (see Appendix A.1). Throughout the analysis, 2009 is emphasized since it was the only year for which data from all four regions were available. When analyzing multiple regions simultaneously, higher-frequency data are summed to hourly, the highest common frequency, and data are adjusted by time zone to coincide. Single missing hourly data points were approximated as the mean of the preceding and following values, and longer gaps were excised. When feasible, analysis of a single region across years characterizes interyear variability. Figure 1 shows a map of the four control regions with wind plant locations and Appendix A.1 contains relevant wind and load statistics.

The principal analytical tool we use is the power spectral density (PSD), which gives a quantitative measure of the strength of wind power fluctuations across a range of frequencies. PSDs of wind power output often contain a peak at $(24 \text{ h})^{-1}$, reflecting daily periodicity (see Figure A1). Wind power PSDs have a negative slope in log-log space: power fluctuations at frequencies corresponding to 10 minutes, for example, are at least a factor of a thousand smaller than those at periods of 12 hours. This property has important practical consequences: if the PSD of wind were flat (white noise), large

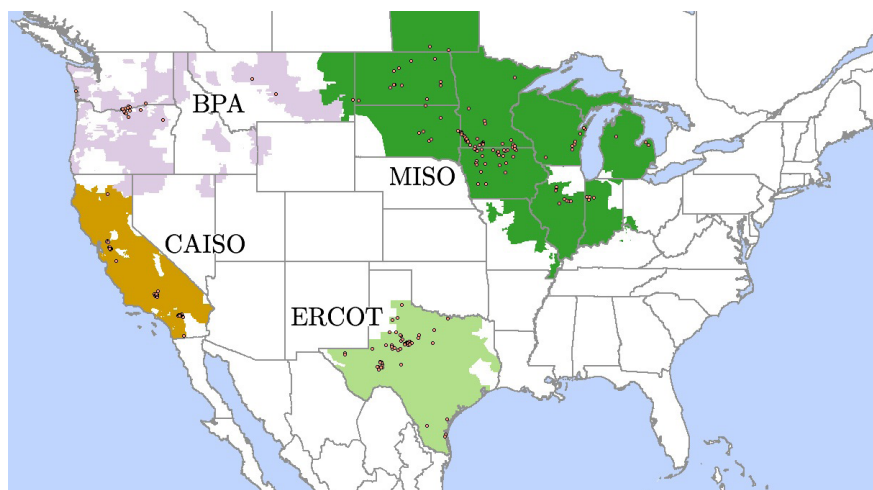


Figure 1: The area spanned by each region and the wind plants it contains.

amounts of very fast-ramping sources would be required to buffer the fluctuations of wind power. The negative slope of the PSD implies that slow-ramping resources such as coal or combined cycle gas plants can compensate for most of wind power’s variability, with less reliance on fast-ramping resources such as batteries and peaker gas plants. The Kaimal spectrum, with a slope of $-5/3$ at frequencies above 24 h^{-1} in log-log space, has been shown to approximate the PSD of power output from a single wind plant (Katzenstein et al., 2010) (see Appendix A.2).

The absolute values of the PSDs, especially at higher frequencies, cannot be directly translated into the wind-balancing resources required at those frequencies. Rare but steep rises or falls in wind power output can increase PSD values at high frequencies such that they no longer reflect general variability patterns. To gain insight into the ideal composition of a wind-balancing portfolio, we observe that the power spectrum for a single, linear ramping generator would be proportional to f^{-2} (Apt, 2007). A generator such as a natural gas plant, sized so that its ramp rate matches wind’s hourly variability, would therefore have nearly twice the capacity necessary to compensate for wind fluctuations observed at daily frequencies. Balancing wind with a portfolio containing fast-ramping resources such as batteries, fuel cells, and supercapacitors, in addition to slower-ramping resources, would avoid the unnecessary expense incurred by building a single type of linear

ramp rate generator that would have excess capacity at low frequencies (Apt, 2007).

3 Results

3.1 Frequency domain analysis

Figure 2(a) shows the PSDs of 2009 wind power output for all four regions and their aggregate as well as a reference Kaimal spectrum (the expected PSD for a single wind plant, here normalized to the summed data). At frequencies higher than $(24 \text{ h})^{-1}$, the aggregated regions show less variability relative to lower frequencies than the reference Kaimal spectrum, as do the individual regions (whose reference Kaimal spectra are not shown). This smoothing pattern is the result of interconnection of individual wind plants within each region, whose power output shows less correlation at high frequencies than at low frequencies.

The similarity of the PSD curves in log-space in Figure 2(a) suggests that variability reduction due to interconnection takes place uniformly across all frequencies examined and that interconnection of regions, unlike interconnection of just a few wind plants, does not reduce the ratio of high- to low-frequency variability for the range of frequencies examined. This effect can be quantified by the slopes of log PSDs in the range of $(24 \text{ h})^{-1}$ to $(2 \text{ h})^{-1}$ (corresponding to the inertial subrange), which reflect the relative variability of power output at the frequencies within that range. Slopes of less than $-5/3$ (the value for a single wind plant) indicate smoothing at higher frequencies due to geographical diversity of interconnected wind plants. The slopes of the log PSDs for individual and interconnected regions are shown in Figure 2(b).

F-tests were used to evaluate the null hypothesis that the PSD slopes in the inertial subrange were the same for individual regions as for combinations of regions. For each group of regions except BPA and CAISO, the null hypothesis failed to be rejected at the 5 % significance level for at least one of the regions tested against the aggregate. The slope

in the inertial subrange for BPA combined with CAISO (-2.61) was significantly different from those of both BPA (-2.51) and CAISO (-2.71). In all cases, although interconnection would reduce variability at all frequencies examined, it would not reduce the slope in the inertial subrange compared with each of its constituent regions. This result indicates that interconnection across regions would not change the proportions of fast- and slow-ramping resources necessary to balance wind power output, and that interconnecting more wind plants within the same region could similarly reduce variability and incur much lower transmission cost.

Figure 2(b) shows that the log slopes of PSD estimates in the inertial subrange can vary between years, implying that the mix of generators, storage, and demand response necessary to compensate for variability of a given amount of wind power can differ from year to year. For the years with data available, these differences are significant at the 5 % level for BPA and CAISO but not for MISO. For ERCOT, differences between years tend to be significant, with the exception of 2007/2010 and 2008/2009.

The PSD slope for BPA wind power output was greater than or equal to that of the other regions for each year examined, indicating comparatively less smoothing due to interconnection of wind plants within BPA. The proximity of the BPA wind plants could expose them to similar weather patterns, limiting the degree of smoothing as a result of interconnection. This effect could also help explain the higher coefficient of variation for BPA than for the other regions.

To summarize, our results suggest that interconnecting multiple wind plants across the four U.S. regions examined would smooth wind power output at all frequencies examined (as quantified by the coefficient of variation; see Appendix A.4). Interconnection would not, however, reduce the ratio of high frequency to low frequency variability in wind power output beyond the reduction found by Katzenstein et al. (2010) for ERCOT wind plants.

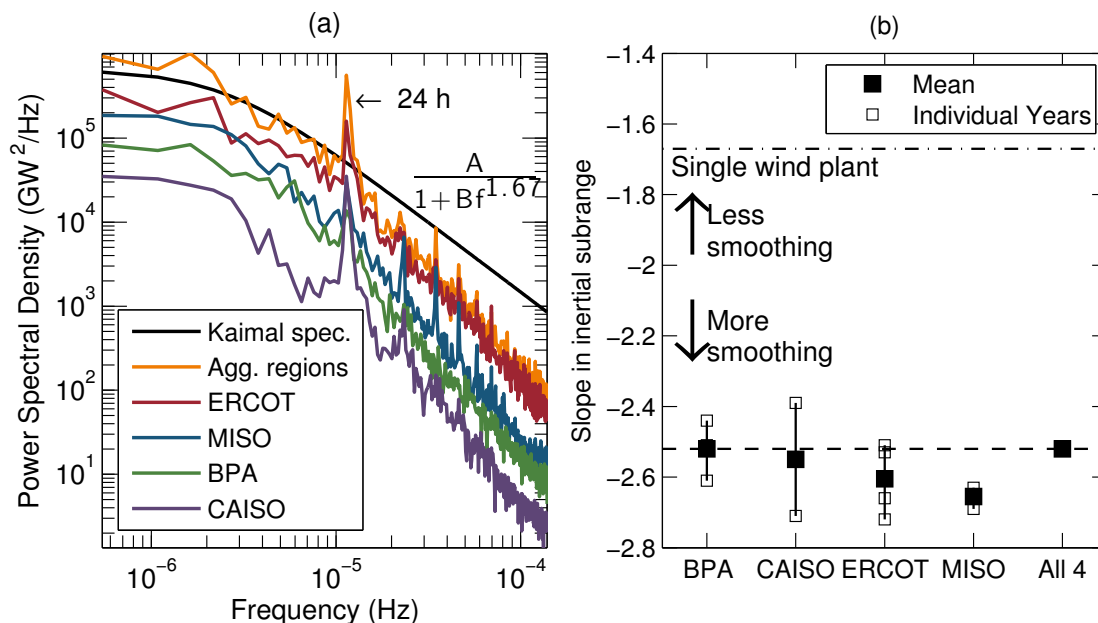


Figure 2: (a) PSDs for 2009 wind power output of each region and the aggregate of all four regions plotted on log-log axes. The displayed Kaimal spectrum equation approximates the PSD for a single wind plant (fitted parameters are $A = 5.84 \times 10^5$ and $B = 2.06 \times 10^9$). In the inertial subrange (frequencies higher than $(24 \text{ h})^{-1}$), the summed power output shows less variability than that of a single wind plant. The legend lists data as they appear from top to bottom. (b) Slopes in the inertial subrange for each region and the interconnected regions for all years of available data and the means over time. The slope for the interconnected regions in 2009 is within the range of slopes for individual regions in other years for which data were available.

3.2 Wind power duration curve

Figure 3 shows a duration curve for 2009 wind power output. Adopting the definition of “firm power” from Katzenstein et al. (2010) as capacity available 79 to 92 % of the time, we find that the interconnected regions have the greatest amount of firm power, with 17 % of installed wind capacity available 79 % of the time and 12 % of capacity available 92 % of the time. MISO, the region with the firmest wind power output as well as the the least likelihood of extreme hourly step changes and lowest slope in the inertial subrange, had 13 % of capacity available 79 % of the time and 6 % of capacity available 92 % of the time. BPA had the least amount of firm wind power, with only 2 % and 0.2 % of capacity available at the ends of the firm power range, consistent with the finding of Katzenstein

et al. (2010) for 2008 data (3 % and 0.5 %). For 2008 ERCOT data, Katzenstein et al. (2010) found 10 % and 3 % of installed capacity available at the limits of the firm power range; for 2009 data, we find 10 % and 4 %. While interconnection of all four regions would at least double the fraction of capacity available 92 % of the time in each region, the gain in firm power (which amounts to approximately 1.5 GW above the sum of that of the individual regions) is unlikely to be sufficient to cover the cost of necessary transmission capacity.

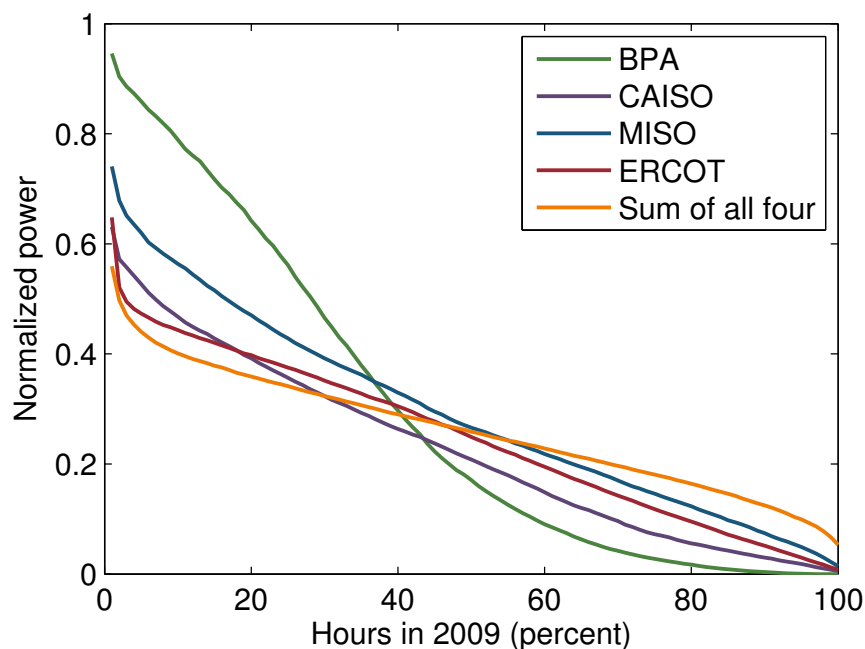


Figure 3: Duration curve for 2009 wind power output. The interconnected regions show the greatest percentage of firm power (capacity available 79 to 92 % of the time) and BPA the least.

3.3 Step change analysis and balancing cost comparison

Step changes of wind power output were calculated as the difference between power output in consecutive hours as a fraction of installed capacity. Interconnection of all four regions was found to produce negligible additional reduction in mean step changes compared with that achieved in a single region. BPA and ERCOT have the highest likelihood of large

hourly step changes and MISO and the aggregated regions the lowest (see Appendix A.5).

We wish to evaluate the cost-effectiveness of smoothing wind power output with increased transmission capacity between the regions with the greatest wind variability, BPA and ERCOT. We calculate the length of high-voltage transmission with cost equivalent to that of a peaking gas turbine sized to mitigate negative 99th percentile step changes in BPA and ERCOT, plus a combined-cycle gas turbine providing firm capacity equivalent to what the interconnected regional wind power output could provide. We find that the cost of the gas turbines would only cover 490 to 740 miles of transmission capacity (630 to 960 miles if emissions damages are included), whereas BPA and ERCOT are separated by 1400 miles (see Appendix A.7 for details of the cost calculation). This first-order analysis suggests that local gas is a more cost-effective method of balancing low-probability step changes and providing firm power than increased transmission capacity.

4 Discussion and Conclusion

Frequency domain analysis shows that fluctuations in wind power are not white noise. Fluctuations in aggregate regional wind power output are between three and five orders of magnitude stronger at daily frequencies than at hourly frequencies (see Figure 2). The relative strength of low-frequency fluctuations of wind power output yields the important result that wind power can be balanced to a large extent by slow-ramping generators such as coal plants and combined-cycle natural gas plants.

Interconnection of wind plants within a single region would further reduce the ratio of fast- to slow-ramping generators necessary to balance wind power output, since across short distances wind's high-frequency fluctuations cancel each other more effectively than its low-frequency fluctuations. Our work demonstrates that interconnection of aggregate regional wind power output would provide no further reduction in the ratio of high- to low-frequency fluctuations, and therefore the ratio of fast- to slow-ramping generators in

the balancing portfolio, than the reduction obtained from interconnecting wind plants within a region.

Nevertheless, benefits of interconnecting aggregate regional wind plants include variability reduction at all frequencies examined (as measured by the coefficient of variation), reduction in the likelihood of extreme step changes in wind power output, and doubling of the fraction of wind capacity available 92 % of the time compared with the maximum of the single regions.

BPA is the region that would benefit most from interconnection with other regions. However, BPA is also the only region with substantial hydropower capacity, including pumped storage. Hydropower is a low-emissions technology that ramps quickly enough to follow fluctuations in wind power output, and may be a more successful and cost-effective method for integrating BPA wind power than long-distance interconnection.

Net load (load minus wind generation) shows the same relative proportions of high and low frequency fluctuations regardless of wind capacity, such that the proportion of balancing resources required to compensate for wind variability will be roughly constant as wind capacity grows (see Appendix 2). This finding supports the treatment of wind power as negative load.

A first-order analysis shows that for BPA and ERCOT, the cost of mitigating wind's low-probability step changes and providing equivalent firm power is considerably lower with natural gas turbines than with interconnection of aggregate regional wind plants.

The availability of higher resolution data over a longer time span would refine these conclusions, although the consistency of the findings and their similarity across 2008 and 2009 argue for the robustness of the principal conclusions.

Table A1: Summary statistics for wind power in the four regions examined in 2009.

	BPA	CAISO	ERCOT	MISO
Installed wind capacity (MW) ^a	2,100	2,200	8,400	6,600
Maximum hourly wind (MWh)	2,300 ^b	1,900	6,000	5,400
Capacity factor	0.29	0.22	0.24	0.29
Average load (MW)	6,200	26,600	35,100	62,600
Average wind/load	0.10	0.025	0.067	0.028

^a 2009 yearly average installed wind capacity.

^b In BPA, installed wind capacity increased sufficiently during 2009 such that the maximum power output, which occurred near the end of the year, exceeds the average installed capacity.

Acknowledgements

The authors thank Mark Handschy for useful comments and conversations. This work was supported by the National Science Foundation (NSF) Graduate Research Fellowship Program, grants from the Alfred P. Sloan Foundation and Electric Power Research Institute (EPRI) to the Carnegie Mellon Electricity Industry Center, the Doris Duke Charitable Foundation, the R.K. Mellon Foundation, and the Heinz Endowments to the RenewElec program at Carnegie Mellon University, and the U.S. National Science Foundation under Award no. SES-0949710 to the Climate and Energy Decision Making Center.

Appendix A Wind power variability

A.1 Data

Table A1 gives regional wind and load statistics for 2009 and Table A2 summarizes the time spans, sampling frequencies, dropouts, and origins of the data (BPA, 2011; CAISO, 2011; ERCOT, 2011; MISO, 2011).

Table A2: Summary of data sampling frequencies, gaps, and origins. Wind and load data dropouts occurred simultaneously for all but CAISO 2010. No 2008 MISO load data or 2007–8 CAISO wind data were obtained. Gaps refer to the number of data points missing.

Region	Year(s)	Sampling frequency (min)	Sample gaps (longest consecutive)
BPA	2008	5	27 (12)
BPA	2009	5	48 (12)
BPA	2010	5	-
CAISO	2007	60	4 (1)
CAISO	2008	60	288 (6)
CAISO	2009	60	-
CAISO	2010	1	925 (789) wind; 553 (18) load
ERCOT	2007 - 2010	15	-
MISO	2008-2009	60	-

A.2 PSD and Kaimal spectrum approximation

We use the periodogram approximation to the power spectral density:

$$P(f) = \frac{1}{NF_s} \left| \sum_{n=0}^{N-1} x_n e^{-j2\pi fn/F_s} \right|^2 \quad (\text{A1})$$

in which f is frequency in Hz, N is the number of samples, F_s is the sampling frequency in Hz, and x_n is the time series.

To analyze smoothing of wind power output in the frequency domain, we first estimate a PSD of a single wind plant to serve as a standard for comparison. Apt (2007) showed that fluctuations in wind power output are not white noise, which has equal power at all frequencies; rather, they are much stronger at low frequencies (signals with PSDs of that character are termed red noise). As discussed in the text, this has the important implication that the majority of wind’s variability can be balanced by slow-ramping sources. As an example, we have examined a sample of 59 coal plants, and determined their ramp rates. The slowest-ramping one in our sample is a 600 MW coal plant with a ramp rate of 0.2 % per minute. That is, it can cycle 50 % of its power with a characteristic frequency of $\sim 7 \times 10^{-5}$ Hz. Referring to the PSD of wind power output (Figure A1 below

or Figure 3(a) in the main text), we see that ramping a coal plant at that or lower frequencies can compensate for the great majority of wind’s fluctuations.

The PSD of output from a single wind turbine follows the Kolmogorov spectrum, with slope $-5/3$ on a log-log plot, at frequencies corresponding to 30 s to 2.6 d. After Kaimal et al. (1972) and later Katzenstein et al. (2010), parameters A and B in (A2) are fit to the PSD of output from a single wind plant, such that at low frequencies the slope approaches zero and at high frequencies the slope approaches $-5/3$. The power spectrum described by (A2) is termed the Kaimal spectrum, in which S is the power spectral density:

$$S(f) \approx \frac{A}{1 + Bf^{\frac{5}{3}}} \quad (\text{A2})$$

The parameter B determines the corner frequency, at which the PSD slope transitions from 0 to $-5/3$. The parameter A scales the amplitude. To obtain a standard for comparison for the remainder of the analysis, we fix B as the value for the wind plant in Snyder, Texas, which was the ERCOT wind plant found to best conform to the Kaimal spectrum. We scale A such that the integral of the PSD is the same as that of the Kaimal spectrum, giving the time-domain signals the same variance. Figure A1 shows the PSD and Kaimal fit for Snyder. To reduce noise, the year of data was divided into 16 segments, the PSD calculated for each segment, and the final PSD obtained by taking the average over the segments.

The two distinct regions of slope, at frequencies above and below the corner frequency of about $(24 \text{ h})^{-1}$, result from atmospheric properties at different spatial scales. Boer and Shepherd (1983) show that the energy spectrum of wind has a slope of $-5/3$ at high wavenumbers (spatial frequencies), for which its flows exhibit isotropic turbulence and transient behavior. This region, which corresponds to high-frequency fluctuations in wind speed, is not sensitive to season. At low wavenumbers (corresponding to low frequencies) where the PSD is flat, the energy spectrum varies seasonally and is largely determined by

topography and thermal effects. Over the range of frequencies examined in this paper, wind power output inherits these properties.

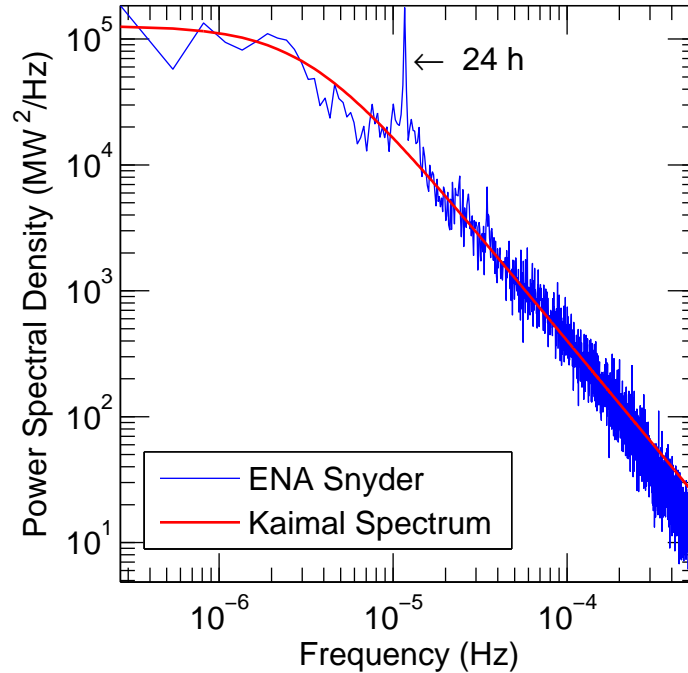


Figure A1: Kaimal approximation to the PSD of 15-minute wind energy data from a single 63 MW wind plant in Snyder, Texas (ERCOT). The peak at the frequency corresponding to 24 hours is highlighted, showing the strong daily periodicity in the wind power output. The fitted parameters are $A = 1.27 \times 10^5$ and $B = 1.47 \times 10^9$. The PSD was calculated with 16 segment averaging to reduce noise.

A.3 PSD slopes in the inertial subrange

Table A3 shows slopes of the log PSDs in the inertial subrange for wind power output of individual and interconnected regions. Slopes of less than $-5/3$ (the value for a single wind plant) indicate reduction in the ratio of high- to low-frequency variability due to geographical diversity of interconnected wind plants. Although each individual region has a substantially lower PSD slope in the inertial subrange compared with a single wind plant, in pairwise linear hypothesis tests (F-tests), no set of interconnected regions has a significantly lower PSD slope than each of its constituent regions. This result implies that interconnection of two or more regions would not reduce proportions of fast- and

slow-ramping resources necessary to balance wind power output.

Table A3: Log PSD slopes for frequencies from $(24 \text{ h})^{-1}$ to $(2 \text{ h})^{-1}$, in which a single wind plant follows the Kolmogorov spectrum (slope = -1.67).

Region	2007	2008	2009	2010
BPA	-	-2.44	-2.51	-2.61
CAISO	-	-	-2.71	-2.39
ERCOT	-2.72	-2.53	-2.51	-2.66
MISO	-	-2.63	-2.69	-
BPA and CAISO	-	-	-2.61	-2.43
BPA and ERCOT	-	-2.48	-2.52	-2.63
BPA and MISO	-	-2.46	-2.64	-
CAISO and ERCOT	-	-	-2.54	-2.40
CAISO and MISO	-	-	-2.70	-
ERCOT and MISO	-	-2.51	-2.50	-
BPA, CAISO, and ERCOT	-	-	-2.55	-2.44
BPA, CAISO, and MISO	-	-	-2.66	-
BPA, ERCOT, and MISO	-	-2.49	-2.50	-
CAISO, ERCOT, and MISO	-	-	-2.51	-
BPA, CAISO, ERCOT, and MISO	-	-	-2.52	-

A.4 Coefficients of variation

Figure A2 shows the coefficients of variation (the standard deviation divided by the mean, abbreviated CV) of 2009 wind power output of each region and their aggregate. Since regions with larger mean wind power output tend to have more wind plants with greater geographic dispersion, the CV generally decreases as mean wind power output increases. The reduction in CV due to pairwise interconnection ranged from 3 % for CAISO (through interconnection with BPA) to 48 % for BPA (through interconnection with MISO; see Figure A2). Most reductions were similar to those expected for uncorrelated wind power output, with the exception of BPA and CAISO, whose CV was 16 % higher than that expected for zero correlation. The correlation coefficient for these two regions was $\rho = 0.32$, likely due to similar east-west moving fronts and sea breeze effects (see Appendix A.6). While interconnection of these two regions would substantially reduce wind variability in

BPA, it would leave CAISO wind variability nearly unchanged.

Interconnection of all four regions would reduce the coefficient of variation by the greatest amount in BPA (58 %) and the least in ERCOT (28 %). The observed correlation of the four interconnected regions is positive and the coefficient of variation is 19 % above that expected for uncorrelated wind power output. Appendix A.6 contains further analysis of the correlation of wind power output between regions.

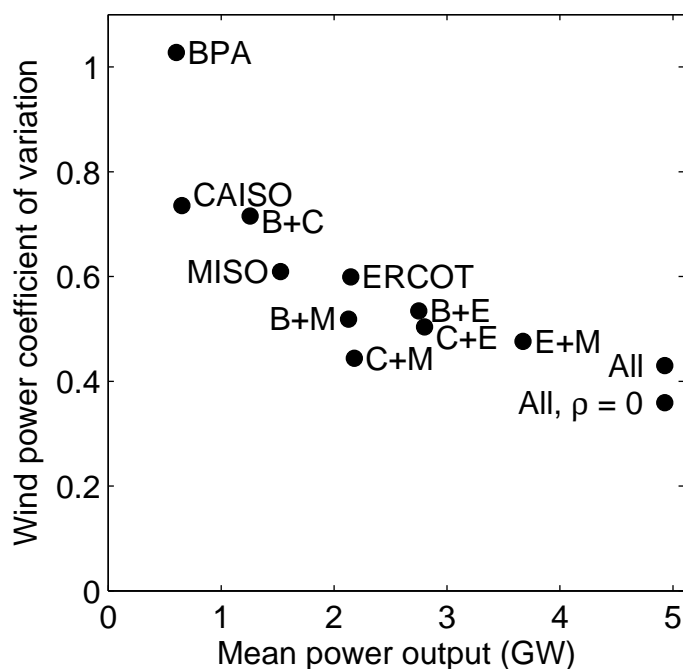


Figure A2: Coefficient of variation (standard deviation divided by mean) of 2009 wind power output for each of the four regions, pairs of regions, and the four regions combined. Pairs of regions are denoted by the first letters of their acronyms separated by the symbol +. “All, $\rho = 0$ ” refers to the expected coefficient of variation for the aggregated regions if the individual regions were pairwise uncorrelated. Positive correlations between regional wind power outputs raise the coefficients of variation of summed wind power above that expected for uncorrelated power output. By this measure BPA had the most volatile wind power output, with its standard deviation approximately equal to its mean.

A.5 Step change analysis

Methods

Step changes are calculated as the fractional changes in wind power output or load over a time interval:

$$\Delta P = \frac{P(t + \Delta t) - P(t)}{P_c} \quad (\text{A3})$$

$P(t)$ is the average power output over the time interval centered at t with length Δt . P_c is the installed wind power capacity.

The step change in wind power must be matched by an opposite change in load or interchange or by activation of balancing and regulation resources. We examine hourly step changes ($\Delta t = 1$ h) since that time scale is important to scheduling and ancillary service markets. While step changes by themselves are a crude measure of the strain wind power imposes on a system, operations personnel at RTOs have indicated that step changes are useful indicators of variability and they enable comparison with other wind variability studies (Wan, 2004; Sørensen et al., 2007). While step changes as a fraction of load would also yield insight into the system response required to balance wind, step changes as fractions of installed capacity are akin to average ramp rates and thus facilitate comparison with other generators.

Results

Mean hourly step changes were 3.0 to 4.3 % of installed capacity (positive) and 2.8 to 4.0 % of installed capacity (negative) for the individual regions. For the four interconnected regions, mean hourly step changes were 2.8 % (negative) and 3.0 % (positive) of installed capacity, indicating negligible benefit from interconnection. Wan (2004) found that individual wind plants in Iowa, Minnesota, and Texas had mean hourly step changes of 4.1 % to 6.1 % of maximum capacity, and that connecting four wind plants in Texas reduced mean step change to 3.7 %. The combined results imply that

interconnecting more than four wind plants would produce diminishing marginal reductions of mean step changes. Mean negative step changes as a fraction of nameplate capacity were consistently less than mean positive step changes, suggesting that wind-balancing resources on average will be required to ramp down more quickly than they ramp up.

Figure A3 shows a duration curve for 2009 hourly step changes (as fractions of installed capacity) for each region and their aggregate (after Sørensen et al. (2007)). For clarity, the plot is cropped to show only positive step changes, though the portion of the plot showing negative step changes is roughly symmetrical. BPA and ERCOT have the highest likelihood of large hourly step changes, and MISO and the aggregated regions tend to have the lowest probability of extreme hourly step changes. Nevertheless, ERCOT's coefficient of variation is the lowest of the four regions (see Figure A2); although wind power output in ERCOT shows low variability in general compared with the other regions, its worst-case fluctuations tend to be more extreme. Appendix B.2 contains a step change analysis of net load.

A.6 Correlation of wind power output

Correlation coefficients for two wind power output time series X and Y were calculated as $\rho_{X,Y} = \frac{1}{n-1} \sum_{i=1}^n \left(\frac{X_i - \bar{X}}{s_X} \right) \left(\frac{Y_i - \bar{Y}}{s_Y} \right)$ in which n is the number of hourly data points, \bar{X} and \bar{Y} are the sample means of X and Y , and s_X and s_Y are the sample standard deviations of X and Y . Negative correlation between wind power output of two regions indicates that connecting the regions could result in a smoother supply of wind power, if the variations tend to be out of phase.

Table A4: Correlation coefficients between 2009 (2008) hourly wind power outputs.

	BPA	CAISO	ERCOT	MISO
BPA	1	0.32	0.04 (0.16)	-0.06 (-0.07)
CAISO		1	0.02	-0.23
ERCOT			1	0.24 (0.25)
MISO				1

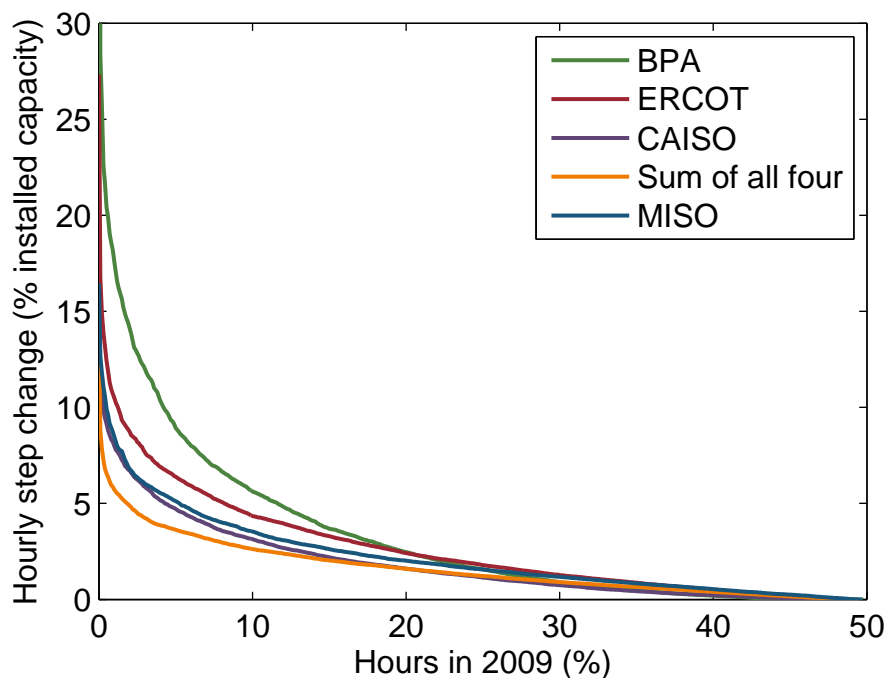


Figure A3: Duration curve for hourly step changes. BPA and ERCOT are the most likely to have large hourly step changes, and MISO and the interconnected regions are the least likely. For clarity, the curve shows only positive step changes; the negative portion of the curve is roughly symmetrical.

Table A4 shows correlation coefficients for wind power output data from 2009 (and 2008 if available). Figure A4 shows pairwise correlation coefficients plotted against the distance between the centroids of the wind plants of each region. All correlation coefficients are different from zero at the 1 % significance level and are highly significantly different from -1 , the value for perfect smoothing. BPA paired with CAISO and ERCOT paired with MISO showed the strongest positive correlation, likely because fronts generally pass from west to east, creating similar conditions along north-south axes. CAISO and MISO showed the strongest negative correlation. Nevertheless, connection of the two regions did not result in smoothing of PSDs at frequencies above 24 h^{-1} , even when CAISO wind power output was scaled to simulate the same installed capacity as MISO.

The correlations are plotted with curves from Giebel (2000), derived from simulated European wind power data from sites up to 4,500 km apart, and Katzenstein et al. (2010), derived from wind power output from 21 Texas wind plants up to 500 km apart. Each time

series consists of one year of hourly data (although that of Giebel (2000) is linearly interpolated from three-hourly data). Compared with pairwise correlations from European wind sites, BPA+CAISO and ERCOT+MISO are more highly correlated than average but well within the scatter of European data, while CAISO+MISO is an outlier compared with European data. The exponential fit for all pairs of ERCOT wind farms did not have the benefit of longer-distance data (Katzenstein et al., 2010), possibly explaining its difference from the European curve.

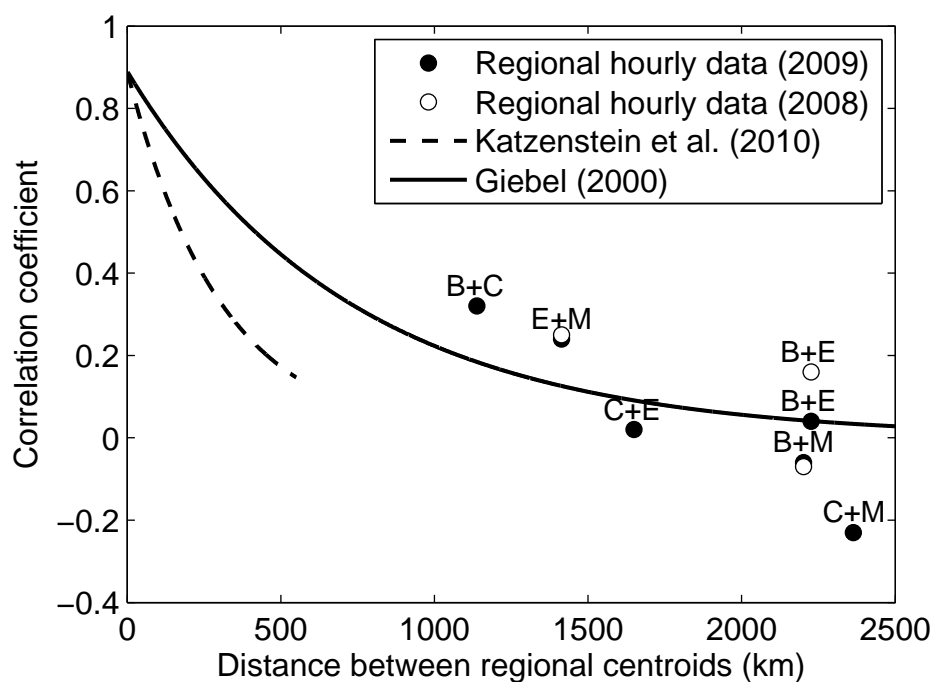


Figure A4: Pairwise correlation coefficients between regional wind power outputs tends to decrease as a function of distance between regional wind plant centroids. Pairs of regions are denoted by the first letters of their acronyms separated by the symbol +. Compared with European data, CAISO and MISO wind power is exceptionally anticorrelated, while correlation coefficients for the other pairs of regions fall within the spread of European data.

A.7 Cost comparison with natural gas combustion turbines

We wish to evaluate the cost-effectiveness of interconnecting wind plants to smooth power output. BPA and ERCOT are the regions with the most volatile wind power output, and the correlation coefficient between them is low (see Table A4). We perform a coarse

calculation of the cost of mitigating their 99th percentile negative step changes with a peaking natural gas plant and generating the extra baseload power that the connected regions could provide with a natural gas combined cycle plant. We then calculate the approximate length of transmission with cost equivalent to that of the extra generation capacity and compare this length with the actual distances between regions.

The 99th percentile negative hourly step change in 2009 for both BPA and ERCOT wind power output was equal to 14 % of regional installed capacity, or about 300 MW and 1.2 GW respectively. We assume that a simple cycle gas plant is built to balance the step changes and that its overnight capital cost is \$1000/kW, fixed operating and maintenance (O&M) is \$7/kWh · yr, and variable O&M is \$15/MWh (EIA, 2010). With 100 h/yr of operation, and amortizing costs with a 10 % discount rate over a 40-year lifetime, the total cost of the gas plant would be \$330 million for BPA and \$1.3 billion for ERCOT, for a total cost of \$1.6 billion. (The O&M costs contribute less than 10 % to the total cost.)

The amount of wind capacity available 92 % of the time is 340 MW in ERCOT, almost none in BPA, and 630 MW in the combined regions. Connecting BPA and ERCOT could therefore replace about 300 MW of baseload capacity. If this baseload capacity were provided by a natural gas combined cycle plant, it would cost \$430 million (using median cost estimates from Lazard (2011)). The total cost of the baseload and peaker gas plants in the absence of interconnection is therefore approximately \$2 billion.

The weighted-average criteria-air-pollutant damages from natural gas plants amount to \$0.36/MWh and the mean carbon emissions from a gas plant are 0.5 tCO₂/MWh (NAS, 2010). With a carbon price of \$50/tCO₂, the total cost of the gas turbines becomes \$2.6 billion.

Depending on factors such as terrain, right-of-way costs, and permitting requirements, 765 kV single circuit transmission tends to cost between \$2.7 million and \$4.1 million (2010 USD) per mile (AEP, 2008). At these transmission costs, building a gas turbine to mitigate 99th percentile step changes would have capital cost equivalent to high-voltage

transmission of length 490-740 miles (630 to 960 miles if criteria air pollutants and CO₂ costs are included). The geographic centroid of wind plants in BPA is separated from that of ERCOT by 1400 miles (the centroid of BPA wind plants is 750 miles from that of its nearest neighbor, CAISO, and the centroid of ERCOT wind plants is 850 miles away from that of its nearest neighbor, MISO.)

Reversing the calculation, the total cost of an ERCOT-BPA transmission line would need to be less than \$1.4 million per mile (\$1.9 million if emissions damages are included) to make the transmission line more cost-effective than gas plants to compensate for the extreme step changes and provide additional base load power. While transmission costs this low have been reported for flat exurban regions, costs over 1400 miles are likely to be higher.

Since the transmission distances with equivalent cost to the gas turbines are low compared with the distances between regions, this first-order analysis suggests that local gas plants would be a more cost-effective means of mitigating extreme negative step changes and providing extra baseload capacity than interconnection with high-voltage transmission. This analysis neglects the fact that the two individual regions with extra gas capacity would produce more energy than the interconnected regions; accounting for this would further reduce the calculated cost effectiveness of building transmission to mitigate wind variability.

We do not attempt to quantify other benefits of increased transmission capacity such as alleviation congestion or bolstering system reliability independent of wind power smoothing. We also do not quantify the possible disadvantages of interconnection between regions, such as the spreading of faults and the difficulty of coordination between markets (Cepeda et al., 2008). Further, the reliability benefit of increased wind-load transmission capacity has been shown to exhibit decreasing marginal benefit, such that economically optimal interconnection capacity is likely lower than installed wind capacity (Karki and Patel, 2005).

We propose that incorporating the cost of extra transmission capacity to balance wind is important even though the existing high-voltage transmission network contains extensive interconnections between the regions examined in this paper. The Pacific DC Intertie (Path 65), for example, connects the wind-rich area of BPA to southern California. This path was found to be one of the most heavily used in 2007, and was operating at more than 75 % of capacity at a frequency of 18 % in the spring, 23 % in the winter, and 32 % in the summer (DOE, 2009). Given this heavy usage, the path is unlikely to have sufficient excess capacity to allow the wind power output of BPA and CAISO to balance each other to the extent modeled in this paper. Incorporating the cost of excess transmission capacity is therefore a requisite element of an economic analysis of wind power smoothing due to geographic diversity.

Appendix B Net load

Like wind, electricity load is variable across a broad range of frequencies. Along with predictable daily and weekly periodicity, higher frequency, less predictable fluctuations present challenges to grid operators in maintaining system balance. Analysis in both the time and frequency domains can yield insight into load variability and the effect of wind power on the net load fluctuations which other generators must match. Apt (2007) found that the power spectra of wind and load were similar at frequencies between $(1 \text{ h})^{-1}$ and $(5 \text{ min})^{-1}$, providing a theoretical basis for the practice of treating wind as negative load at these timescales.

B.1 Frequency domain

Comparing the PSDs of load and net load (load minus wind) reveals the effect of wind power on the variability characteristics of load across a range of frequencies. For current wind power penetrations in each region, net load has a PSD identical to that of load. To

examine the effect of increased wind capacity, the PSDs of CAISO ten-minute net load for 2010 were calculated with both historical wind power output and wind power output multiplied by 10. Figure A5 shows that amplified wind power output translates the PSD upward in log-log space; that is, variability is introduced consistently at all frequencies such that the ratio between spectral power at any two frequencies is preserved. This result suggests that if wind capacity is substantially increased, more balancing and regulation resources will be required, but in the same proportion of slow-ramping to fast-ramping resources as currently.

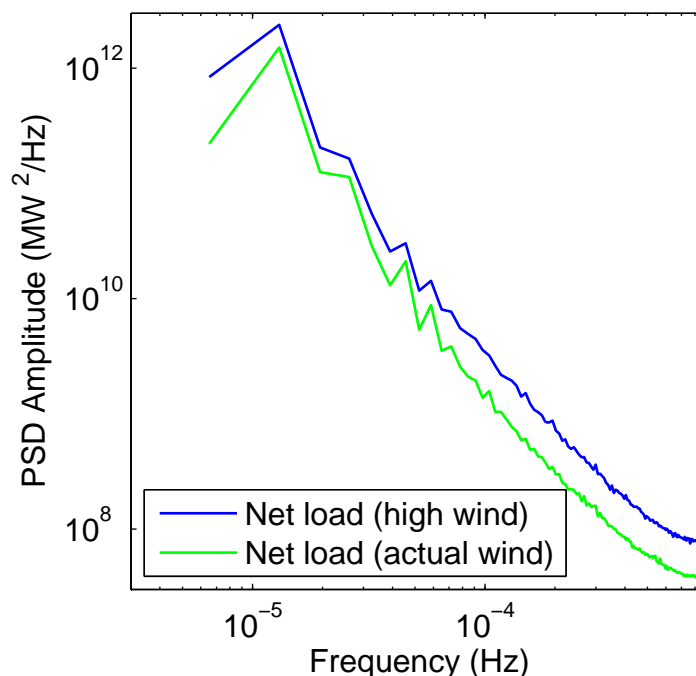


Figure A5: PSDs of 2010 CAISO net load with historical wind power output and wind power output amplified by a factor of 10. Increased wind capacity raises net load variability evenly across frequencies. PSDs are calculated with 16 averaging segments.

B.2 Step-change analysis

Step-changes in net load, when compared with those in load, show the effect of wind power on the ramping requirements placed on the remainder of the generation portfolio. At current penetrations, wind causes increases in maximum and mean net load step changes of

less than 1 % (with the exception of BPA maximum step changes, which increase by less than 3 %). To characterize the effect of expanded wind capacity on net load step changes, wind power output time series were scaled such that wind power fulfilled 50 % of load. Under the high-wind scenario, interconnection mitigated maximum hourly step changes: positive (negative) step changes for the individual regions ranged from 13 to 20 % (15 to 36 %) of total load whereas maximum positive and negative step changes for the aggregated regions were only 11 %. Interconnection had no effect on mean step changes, which consistently rose by 2 % due to the expansion of wind power capacity.

Table A5: Wind-load correlation coefficients for 2009 (2008) hourly data.

Region	$\rho(\text{wind, load})$
BPA	-0.16 (0.00)
CAISO	0.03
ERCOT	-0.28 (-0.18)
MISO	-0.08
BPA and CAISO	-0.02
BPA and ERCOT	-0.21 (0.00)
BPA and MISO	-0.12
CAISO and ERCOT	-0.22
CAISO and MISO	-0.24
ERCOT and MISO	-0.29
BPA, CAISO, and ERCOT	-0.17
BPA, CAISO, and MISO	-0.22
BPA, ERCOT, and MISO	-0.27
CAISO, ERCOT, and MISO	-0.32
BPA, CAISO, ERCOT, and MISO	-0.30

B.3 Correlation analysis

Table A5 shows correlation coefficients between wind and load for each region and the aggregated regions. For 2009 data, ERCOT had the strongest anticorrelation between wind and load followed by BPA, and MISO and CAISO had near-zero correlation.

Anticorrelation between wind and load could exacerbate net load variability as wind power capacity is increased, placing greater demands on balancing and regulation resources.

Further, wind power that is anticorrelated with load will tend to displace baseload generators, which are the marginal plants during times of low load, rather than peak load generators, which are more expensive but easier to ramp. Displacing coal base load may also provide greater reductions in greenhouse gas emissions than displacing load-following resources. Connecting regions does not raise the correlation between wind and load; for most pairs of regions, the correlation coefficient for the summed data is between those for the individual regions, and connecting CAISO and MISO results in a correlation coefficient substantially less than that of either region (2009 data). Interconnecting all four regions results in lower wind-load correlation than in any of the individual regions. Like the frequency domain and step change analyses, the wind-load correlations show that interconnection may not mitigate the negative effects of wind power variability on the remainder of the system.

Bibliography

- AEP (2008). American Electric Power: Transmission Facts.
- Apt, J. (2007). The spectrum of power from wind turbines. *Journal of Power Sources*, 169:369–374.
- Boer, G. J. and Shepherd, T. G. (1983). Large-scale two-dimensional turbulence in the atmosphere. *Journal of the Atmospheric Sciences*, 40:164–184.
- BPA (2011). Bonneville Power Administration: Balancing authority load & total wind generation.
- CAISO (2011). California Independent System Operator (CAISO) Open Access Same-Time Information System (OASIS).
- Cepeda, M., Saguan, M., and Pignon, V. (2008). Generation adequacy in transmission interconnection in regional electricity markets.
- Degeilh, Y. and Singh, C. (2011). A quantitative approach to wind farm diversification and reliability. *Electrical Power and Energy Systems*, 33:303–314.
- DOE (2009). National electric transmission congestion study. Technical report, U. S. Department of Energy.
- Dvorak, M. J., Stoutenburg, E. D., Archer, C. L., Kempton, W., and Jacobson, M. Z. (2012). Where is the ideal location for a US east coast offshore grid? *Geophysical Research Letters*, 39.
- EERE (2008). 20% wind energy by 2030: Increasing wind energy’s contribution to U.S. electricity supply. Technical Report DOE/GO-102008-2567, U.S. Department of Energy: Energy Efficiency and Renewable Energy.
- EIA (2010). Updated capital cost estimates for electricity generation plants. Technical report, U. S. Energy Information Administration: Office of Energy Analysis.
- EnerNex (2011). Eastern wind integration and transmission study. Technical Report NREL/SR-5500-47086, National Renewable Energy Laboratory.
- ERCOT (2011). Electric reliability council of Texas (ERCOT) planning and operations information, demand and energy.
- Ernst, B., Wan, Y.-H., and Kirby, B. (1999). Short-term power fluctuations of wind turbines: Analyzing data from the German 250-MW measurement program from the ancillary services viewpoint. In *Windpower ’99 conference*. National Renewable Energy Laboratory.
- GE Energy (2010). Western wind and solar integration study. Technical Report NREL/SR-550-47781, National Renewable Energy Laboratory.

- Giebel, G. (2000). *On the benefits of distributed generation of wind energy in Europe*. PhD thesis, Carl von Ossietzky University of Oldenburg.
- Holttinen, H. (2005). Hourly wind power variations in the Nordic countries. *Wind Energy*, 8(2):173–195.
- IEA (2005). Variability of wind power and other renewables: Management options and strategies. Technical report, International Energy Agency.
- Kaimal, J., Wyngaard, J., Izumi, Y., and Cote, O. (1972). Spectral characteristics of surface-layer turbulence. *Quarterly Journal of the Royal Meteorological Society*, 98:563–589.
- Karki, R. and Patel, J. (2005). Transmission system adequacy evaluation considering wind power. In *Canadian Conference on Electrical and Computer Engineering*, pages 490–493.
- Katzenstein, W., Fertig, E., and Apt, J. (2010). The variability of interconnected wind plants. *Energy Policy*, 38(8):4400–4410.
- Kempton, W., Pimenta, F. M., Veron, D. E., and Colle, B. A. (2010). Electric power from offshore wind via synoptic-scale interconnection. *Proceedings of the National Academy of Sciences*, 107(16):7240–7245.
- Lazard (2011). Levelized cost of energy analysis — version 5.0.
- MISO (2011). Midwest Independent Transmission System Operator (MISO) market reports: Forecasted and actual load report.
- NAS (2010). Hidden costs of energy: Unpriced consequences of energy production and use. Technical report, National Academy of Sciences.
- Sinden, G. (2007). Characteristics of the UK wind resource: Long-term patterns and relationship to electricity demand. *Energy Policy*, 35:112–127.
- Sørensen, P., Cutululis, N. A., Viguera-Rodriguez, A., Jensen, L., Hjerrild, J., Donovan, M. H., and Madsen, H. (2007). Power fluctuations from large wind farms. *IEEE Transactions on Power Systems*, 22(3):958–965.
- Sørensen, P., Cutululis, N. A., Viguera-Rodriguez, A., Madsen, H., Pinson, P., Jensen, L., Hjerrild, J., and Donovan, M. (2008). Modelling of power fluctuations from large offshore wind farms. *Wind Energy*, 11:29–43.
- Wan, Y. (2004). Wind power plant behaviors: analyses of long-term wind power data. Technical Report NREL/TP-500-36551, National Renewable Energy Laboratory.
- Wiser, R. and Bolinger, M. (2011). 2010 wind technologies market report. Technical report, U.S. Department of Energy, Energy Efficiency and Renewable Energy.
- Zavadil, R. (2006). Final report — 2006 Minnesota wind integration study: Volume I. Technical report, EnerNex Corporation.

Paper 4

Optimal investment strategy in a clean energy RD&D program under technological uncertainty

Abstract

Energy technology development is vital to climate policy analysis and despite being highly uncertain it is most often modeled as deterministic. This simplification neglects the ability both to adapt RD&D strategy to changing conditions and to invest in initially high-cost technologies with small breakthrough probabilities. This paper develops an analytical stochastic dynamic programming method for valuing and informing strategy for a government energy RD&D program under technological uncertainty. Cost of the developmental technology is modeled as stochastic but decreasing in expected value with RD&D spending. Both a single-factor model (with RD&D aggregated) and a two-factor model (which separates R&D and learning-by-doing) are developed. Results of the two-factor model applied to carbon capture and sequestration (CCS) indicate that given 15 years until a carbon policy and large-scale CCS deployment, investment in the RD&D program is optimal over a very broad range of initial mitigation costs (\$10–\$380/tCO₂). While the NPV of the

A portion of this paper was published as Fertig, E. and Apt, J. (2012). Optimal investment strategy in low-carbon energy R&D with uncertain payoff. In *Proceedings of the United States Association of Energy Economics (USAEE) North American Conference*, Austin, TX, USA.

program is zero if initial mitigation cost is \$100/tCO₂, under uncertainty the program is worth about \$7 billion. The value of the program is sensitive to the assumed rate of cost reductions exogenous to the program: at high initial mitigation costs the program is worth most with large exogenous cost reductions. Factors that promote R&D spending over promotion of learning-by-doing include more imminent deployment, high initial cost, lower exogenous cost reductions, and lower program funds available.

1 Introduction

Improvement of low-carbon energy technologies over time is an important factor in projecting future costs of greenhouse gas abatement. Early investment in research, development, and demonstration (RD&D) of low-carbon energy technologies could substantially reduce abatement cost, in turn altering the socially optimal strategy for climate change mitigation. Historically, the outcome of RD&D of energy technologies has been uncertain: some developmental technologies achieve large cost reductions and are commercialized successfully, while others never reach commercialization despite extensive investment.

In spite of the uncertainty in cost trajectories of developmental energy technologies, most research treats technological change in the energy sector as exogenous and deterministic (Section 2 reviews the literature). A smaller set of studies models technological change as endogenous, with cost decreasing as a deterministic function of spending or cumulative installed capacity. Much of this work contains sophisticated models of uncertainty in other aspects of the climate and economy and neglects uncertainty in technological change due to computational limitations. Thus, these existing models fail to capture the effect of uncertainty on the value and optimal strategy for RD&D, and for the exogenous case, neglect flexible decision making in RD&D trajectories as well.

This paper presents a framework for valuing a developmental low-carbon energy

technology with uncertain cost that RD&D investment tends to drive down. The framework both quantifies the effect of uncertainty on the value of the investment opportunity and allows for decisions on RD&D strategy to be updated continuously. Although other aspects of the climate and economy are ignored or highly simplified, the model provides insight into optimal RD&D policy under different expectations on cost uncertainty and RD&D effectiveness.

The analysis takes the perspective of a decision maker overseeing a U.S. government program in low-carbon energy RD&D who anticipates that a carbon price will be enacted T years in the future. The decision maker has the opportunity to invest in an RD&D program to reduce the cost of a single CO₂-mitigating energy technology. Over the course of the program, from 0 to T years, the investment strategy can be updated in response to changes in cost, which result both from the program and from exogenous factors such as R&D spillovers and changes in input costs. When T years have passed, if a simple NPV analysis indicates that the technology is cost effective, it will be deployed; otherwise, it will be abandoned. Although RD&D investment lowers cost in expected value, cost is stochastic and may never fall sufficiently low to justify deployment regardless of RD&D spending. The effect of this uncertainty is quantified in the analysis.

The method of this work is based on continuous time stochastic-dynamic programming (SDP), which is similar to real options when it is used to value an investment opportunity at a specified or discretionary future time under uncertainty. While most previous work that applies real options/SDP to the energy sector models an investment opportunity with uncertain cost or payoff, this work characterizes optimal investment in RD&D that lowers the expected value of uncertain cost. Previous work on learning and experience curves in the energy sector informs the functional form of cost reductions as well as parameter estimates on RD&D effectiveness and uncertainty. This work can be situated among decision support tools for low-carbon energy policy, many of which are based on integrated assessment models. Although the latter type of study captures complex relationships

among other parameters in climate science and economics, technological learning is almost always modeled either as exogenous or as a deterministic function of installed capacity. These assumptions neglect the value of exploratory investment in technologies with low expected value but a small probability of great success, the possibility of negative learning, and the value of flexible future decision making once uncertainty is resolved. This paper, though it abstracts from the larger climate-economy system, isolates the effect of uncertainty in technological learning to account for these phenomena. An extension of this work will examine options to invest in multiple technologies, capturing the hedging effect of a portfolio of assets of uncertain value.

This paper analyzes two investment opportunities in RD&D programs. The first is an opportunity to invest in solar photovoltaics (PV), the cost of which is modeled after a single-factor learning curve under uncertainty; the second is carbon capture and sequestration (CCS), with cost modeled after a two-factor learning curve that disaggregates RD&D into R&D and learning-by-doing (LBD). The remainder of this paper is organized as follows: Section 2 reviews relevant literature. Section 3 describes the general stochastic cost model, Section 4 presents parameter derivations for solar PV and CCS, and Section 5 presents the solution method for the RD&D optimization problem based on analytical stochastic dynamic programming. Section 6 presents results and contrasts them with a classic NPV analysis and Section 7 concludes.

The bulk of the analysis focuses on the CCS case study. Results illustrate the value of accounting for uncertainty in planning a government RD&D program: although the NPV of the program under base case assumptions is 0 if the mitigation cost 15 years before deployment exceeds \$100/tCO₂, valuing uncertainty and flexible decision making shows the program to be worth about \$7 billion. Initial investment in RD&D is optimal at mitigation costs up to \$380/tCO₂ due to the possibility of an unlikely breakthrough, even though the program would most likely be deemed unsuccessful and terminated a short time later. Results are highly sensitive to the rate of business-as-usual (BAU) cost reduction in

absence of the government program, which has opposite effect on program value at low and high initial mitigation costs. Factors that promote R&D spending over LBD include more imminent deployment, high initial cost, lower BAU cost reductions, and lower program funding.

2 Previous work

This work relates to four different lines of research: (1) real options for investment in the energy sector under technological uncertainty, (2) learning curves that express cost reduction in energy technologies as a function of experience or research, (3) decision support tools based on climate-economy models (including integrated assessment models) that incorporate technological change in the energy sector, and (4) R&D portfolio optimization in the energy sector.

2.1 Real options for valuing a single project in the energy sector

Most of the studies using real options to value energy projects analyze a single technology, and relatively few account for uncertainty in technological learning. An early study in the latter category is Pindyck (1993), who proposes an SDP-based method for project valuation and illustrates it with an example from the nuclear power industry in the early 1980s.

Pindyck (1993) accounts for both technical and market uncertainty using a functional form that is the starting point for the project valuation method in this work. While Pindyck (1993) models the remaining construction cost of a technology with a constant option to abandon the project or continue investing in construction, I model RD&D investment to drive down the cost of a technology in advance of a specific deployment date. Bednyagin and Gnansounou (2011) use an alternate option valuation method based on Black-Scholes theory to value a multi-stage investment opportunity in R&D of fusion energy. These studies yield insight into the effects of technical uncertainty on optimal investment strategy

but only value R&D in a single project, neglecting the effects of portfolio diversification.

Similar work outside the energy sector includes Alvarez and Stenbacka (2001), who use real options with uncertain parameters represented as geometric Brownian motions to find optimal investment timing in an unspecified existing technology with an embedded option to update the technology to future versions that may be superior.

An extensive body of work uses real options to value investment in the energy sector under cost, price, or regulatory uncertainty but neglecting technical change (e.g. Blyth et al. (2007), Fleten et al. (2012), Bøckman et al. (2008), Cheng et al. (2011), Fernandes et al. (2011), Reinelt and Keith (2007), Patiño-Echeverri et al. (2007), Siddiqui et al. (2007), and Yang et al. (2008).

2.2 Experience and learning curves in the energy sector

Experience and learning curves model the cost of a technology as a function of the capacity deployed and are characterized extensively for energy technologies.¹ Neij (2008) aggregates many estimates for experience curves in the energy sector and finds them to be broadly similar to bottom-up estimates obtained by synthesizing different sources of cost reductions or through expert elicitation. Junginger et al. (2008) present an extensive literature review of energy technology development studies, criticize the experience curve approach, and examine the success of policy interventions to promote technological learning. Grübler (2010) presents a case of negative learning-by-doing in the French nuclear scale up. Similarly, Rubin et al. (2007) find that costs initially increased due to unreliability and underperformance in early system designs for flue gas desulfurization, selective catalytic reduction, combined cycle gas, and liquid natural gas production.

Kouvaritakis et al. (2000) introduced the two-factor learning curve, which models cost reduction as a function of both cumulative capacity and knowledge stock obtained through

¹Though learning and experience curves are often used interchangeably, learning curves represent cost reductions achieved by a single organization through repetition of tasks, while experience curves represent changes in the performance metrics of an entire industry over time.

R&D. Miketa and Schrattenholzer (2004) use two-factor learning curves to optimize energy R&D expenditures and find that up to a point, competition between energy technologies for market share and R&D funding need not lead to the lock-in of certain technologies at the expense of others. Klaassen et al. (2005) derived a two-factor learning curve for the wind industry in Europe and found robust learning rates of 5.4 % (LBD) and 12.6 % (public R&D), though the effect of public R&D may have been overestimated due to the neglect of private R&D investment.

A substantial body of literature critically assesses learning curves. Van der Zwaan and Sagar (2006) question the dichotomy between R&D and learning-by-doing since the stages can alternate and overlap, and note that learning curve estimates could be very biased in incorporating only relatively successful technologies. Jamasb and Köhler (2007) find that single-factor learning curves often underestimate R&D-driven cost reductions in early-stage technologies compared with two-factor learning curves, and that the specification of a “floor cost” for technologies often results in static equilibrium. Ferioli et al. (2009) break down the drivers of learning-by-doing, discuss different future scenarios for learning in CCS, find that wind power learning curves estimated in 1995 poorly predict the current cost of wind power, and show that the learning rate for a single technology can change substantially over time. Pan and Köhler (2007) find that the credibility of learning curves is diminished by their inability to capture either R&D effects or continued cost decreases in mature technologies, and that a logistic curve better fits wind power costs in the UK. Yeh and Rubin (2012) review the principal types of uncertainty in learning curves and suggest better methods for reporting and characterizing these uncertainties for use in policy analyses. Though the method utilized here is identical to the traditional log-linear learning curve under certain assumptions, there are many ways in which the log-linear model does not accurately predict cost reductions due to experience and learning.

2.3 Energy R&D in larger climate/economy models

While most climate/economy models (including integrated assessment models (IAMs)) take technological change as exogenous, there is a growing literature that endogenizes technological learning. Most work assumes deterministic learning curves (e.g. Riahi et al. (2004) and McFarland and Herzog (2006) for carbon capture and sequestration, Rao et al. (2006) and Barreto and Kypreos (2004) for a portfolio of conventional and renewable energy technologies, Miketa and Schramm (2004) for two-factor learning curves for wind and solar power). Pizer and Popp (2008) highlight the gap between empirical research on learning curves and model-based climate policy analyses that incorporate technological change and discuss how the two streams of research could develop in a complementary fashion.

Few studies that use IAMs incorporate uncertainty in technological learning. Gritsevskiy and Nakićenović (2000) use a structured sampling technique to capture technological uncertainty (a full Monte Carlo analysis would be computationally prohibitive) and find that different assumptions on the (uncertain) relationship between technologies produced different emergent energy systems with very similar cost and risk. Popp (2004) modified the DICE model to account for endogenous technological change that depends on both accumulated capacity and depreciating knowledge stock and found that doing so substantially increased the social benefit of a carbon tax. Using a numerical climate-economy model, Bosetti and Drouet (2005) find that uncertainty in technological learning and irreversibility in R&D investment results in an optimal strategy of waiting then making large R&D investments to hedge against uncertainty. Bosetti et al. (2009) model early-stage but promising “breakthrough” technologies with deterministic two-factor learning curves in WITCH, a hybrid climate-energy-economy model, and find that large investment in these technologies is indicated and that they play a crucial role in cost-effective strategies to stabilize atmospheric greenhouse gas concentrations. Bosetti and Tavoni (2009) incorporate uncertainty in R&D effectiveness into the WITCH model (in a

two-stage framework where R&D can result in three possible outcomes) and find that doing so results in higher investments in innovation and lower policy costs. Webster et al. (2012) improved upon the representation of uncertainty in climate policy analyses using IAMs by adopting an approximate dynamic programming framework; results show that a two-stage decision model may not be sufficient to capture the value of investments in factors such as RD&D which is path-dependent and can operate over an extended period of time.

Blanford and Clark (2003) define key areas for further work in R&D strategy in response to climate policy and propose a probabilistic description of technological uncertainty that could eventually be used in conjunction with IAM results (Blanford and Weyant (2005) develop this into a multi-stage decision framework). Both Löschel (2002) and Grubb et al. (2002) review the approaches taken to technological change in different energy, economy, and environment models and find evidence that technological change is induced largely by market conditions and expectations (and therefore by policy) and that it could substantially reduce the long-term cost of addressing climate change. Yeh et al. (2009) identify and critically review principal sources of uncertainty in technological learning and their implications for IAM results. Grübler and Gritsevskiy (2002) present a framework for modeling endogenous technological change in the energy sector based on a broad literature review, and Krey and Riahi (2009) identify attributes of least-cost risk hedging strategies for the energy sector given multiple interacting uncertainties.

Examining the effect of technological learning on optimal climate policy separate from IAMs, Gerlagh et al. (2009) study the relationship between carbon taxes and innovation externalities and emphasize the effect of patent laws. Baker and Peng (2012) assess the expected value of better information gained through expert elicitations on technology development to inform marginal greenhouse gas abatement cost curves and therefore optimal abatement levels. Applying a theoretic framework based on convex optimization, Athanassoglou et al. (2012) find that large increases in R&D investment in solar technology is likely to yield significant returns.

2.4 Analysis of energy R&D portfolios

The literature using real options to optimize R&D investment in multiple energy technologies of uncertain cost is sparser than that which examines a single technology. Two such studies, Siddiqui and Fleten (2010) and Davis and Owens (2003), adapt and extend methods from Dixit and Pindyck (1994). Siddiqui and Fleten (2010) model investment in an unconventional energy technology (UET) and a more established renewable energy technology, where the investor has the option to pay a lump sum to start the UET down a learning curve that follows geometric Brownian motion with negative drift. Davis and Owens (2003) use a closely related method to value investment in a renewable energy technology with uncertainty the cost of non-renewable energy, the remaining cost of developing and switching to the renewable energy technology, and the cost of the developed renewable energy technology.

Many additional studies use methods other than real options to value energy R&D portfolios. Baker and Solak (2011) combine decision analysis and economics to derive marginal abatement cost curves for greenhouse gas emissions and frame the investment decision as a two-stage problem with recourse, in which success of the R&D program is binary and uncertainty is resolved prior to the second-stage decision. Baker and Solak (2011) find that the optimal portfolio composition is robust to the level of damage risk and that portfolio diversification is strong. Shittu and Baker (2010) examine the effect of a rising carbon tax on optimal energy R&D portfolios, neglecting both uncertainty in technological change and the hedging effect of diversification.

Other work uses real options to value R&D portfolios outside the energy sector. Eckhause et al. (2009) use SDP-based real options methods combined with technology readiness levels to model technologies competing for R&D funding, with a decision maker supporting a smaller group of technologies at each stage. Huchzermeier and Loch (2001) develop an SDP-based real options model for investment decisions on a single technology under multiple uncertainties. A relevant empirical example from the management

literature is Girotra et al. (2007), who analyze data on drug development from pharmaceutical companies and find that the negative effect of drug failure on a firm's value is lessened when the firm has other substituting products in development, and that the diversification strategy is most effective under low odds of successful product development, long lead times, and low correlation between the successes of different concepts.

More generally, Sagar and Holdren (2002) describe gaps in the understanding of the global energy innovation system (such as characterization of the allocation, risk level, and payoff horizons for R&D spending) and argue that energy R&D is an appropriate role for governments. Hart and Kallas (2010) perform a detailed qualitative analysis of historical federal RD&D policy for SO₂ and NO_x controls and draw insights into effective greenhouse gas reduction policy. Nemet (2009) examines the cost effectiveness of recent R&D spending in the renewable energy sector and argues that policy makers should more explicitly consider uncertainty in future cost and that better tools should be developed to study the significance of near-term deviations from cost projections.

The work described here extends previous research by simultaneously modeling the R&D investment decision in continuous time, using a continuous distribution to model cost uncertainty, and incorporating literature on experience curves into an SDP-based valuation of the R&D investment opportunity. This work will provide a quantitative decision analysis framework to help inform policy makers on optimal R&D strategy in low-carbon energy technologies under uncertainty.

2.5 Sources on methods

Methods are based largely on Dixit and Pindyck (1994) and Pindyck (1993), who present a continuous-time method for analyzing investment decisions where uncertainties are well-captured by Brownian motions with or without drifts and Poisson jumps. The numerical solution to the resulting partial differential equation (PDE) problem is based on Muthuraman (2008), who derives and proves convergence for a similar method that solves

the free-boundary problems arising from valuing American put options.

3 Stochastic cost model

This section presents the general form of the cost model as well as the specific forms used in this thesis. The general form is an Itô process with a drift and volatility that are (possibly constant) functions of the rate of spending on R&D and learning-by-doing (LBD). Cost is expressed as

$$\frac{dC}{C} = a(\mathbf{I}(C, t))dt + b(\mathbf{I}(C, t))dz \quad (1)$$

in which C is the mitigation cost using the technology, t is time, $\mathbf{I}(C, t)$ expresses the optimal investment rate in RD&D as a function of C and t , z is a standard Brownian motion (normally distributed with mean 0 and variance t), and a and b are well-behaved functions that map the vector \mathbf{I} to a scalar. \mathbf{I} is expressed as a vector to account for the possibility of multiple drivers of cost reduction (as is the case with a two-factor learning curve). This functional form implies that C can never drop below 0, as fits a cost function, and it allows for an analytical solution to the stochastic-dynamic programming problem that optimizes the investment strategy.

Limitations of the Itô process as a cost model include its inability to capture fat-tailed phenomena and its exclusion of discrete stochastic jumps. Further, this process has the Markov property, in which the distribution of outcomes for any future state depends only on the current state and not on any prior state. This requires the assumption that RD&D spending has an instantaneous effect on cost, and that prior RD&D spending cannot affect current cost. While this is not strictly realistic, the model developed in this paper has a temporal separation between the deployment of the technology and the earlier stages of the RD&D program, so the technology cannot be deployed immediately as a result of a spending-induced cost reduction. Further, this model does not account for stocks of

knowledge or capacity, except to the extent that their implicit accumulation over time affects cost. Rather, this analysis focuses on the optimal rate of investment in R&D and the promotion of LBD as determined by the current cost of the technology and the time until deployment.

The value of $a(\mathbf{I}(C, t))$ is the instantaneous drift rate of the technology's cost as a function of RD&D investment \mathbf{I} , which in turn depends on the cost C and time t .² Since RD&D is expected to decrease costs, $a(\mathbf{I}) < 0$. The value of $b(\mathbf{I})$ is the yearly standard deviation in cost as a percentage of current cost. As cost decreases, its standard deviation therefore decreases proportionally. Variability in cost could be due to uncertainty in market conditions such as the cost of factors of production or to uncertainty in the effectiveness of the RD&D program. Recent examples of the former type of uncertainty include the depletion of waste silicon from the electronics industry for use in solar PV manufacturing and the rise in commodities prices for wind turbine manufacturing.

In regions of the state space where a and b are constant, C follows a geometric Brownian motion (GBM) with drift, which is a class of Itô process often used to model economic parameters such as securities prices, interest rates, wage rates, and output prices (Dixit and Pindyck, 1994). For processes that follow a GBM, the distribution of outcomes at any specific time in the future is lognormal.

In financial economics it is sometimes argued that calculating the value of an investment opportunity in a project whose value is driven by a GBM is possible only if the risk of the project is directly traded in a market or if there is a portfolio of assets whose risk replicates that of the project (see, e.g., Brealey (2002)). Since there is neither a market nor a replicating portfolio for the risk inherent in RD&D investment, the following valuation of the RD&D program uses stochastic dynamic programming with a subjective discount rate. Sensitivity analyses on relevant parameters are performed in order to cover a broad range of possible conditions, and trends rather than specific numerical values are

²For the remainder of the analysis the explicit dependence of \mathbf{I} on C and t is dropped, and $\mathbf{I}(C, t)$ is written simply as \mathbf{I} .

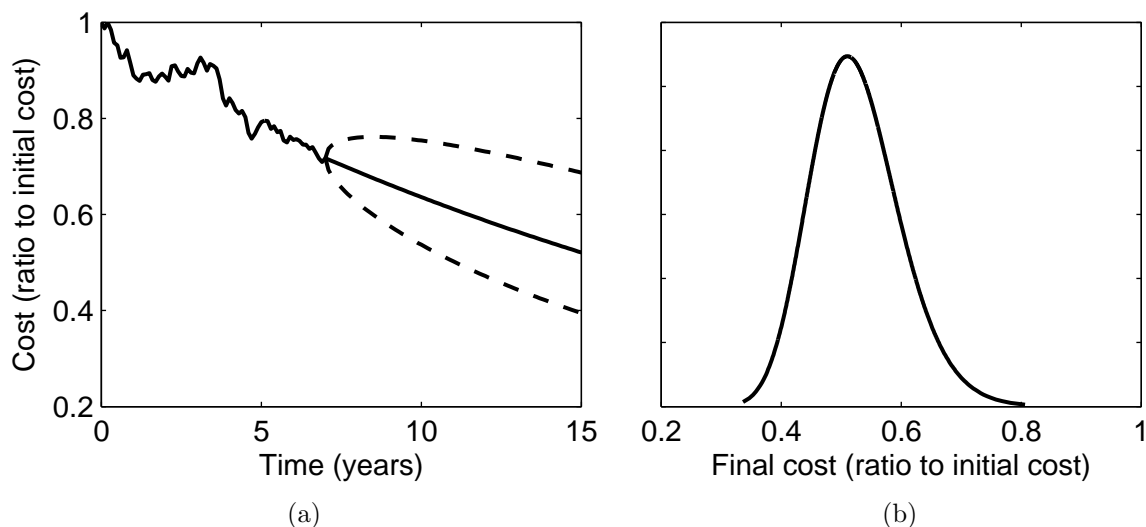


Figure 1: Example cost trajectory of the technology as a function of time, following a GBM. (a) shows a sample path of cost from 0 to 7 years and the 95 % prediction interval for cost from 7 to 15 years. (b) shows the lognormal probability distribution for final cost given knowledge of the cost at 7 years. The 7-year threshold was chosen arbitrarily as an illustration. Here, it is assumed that R&D investments are made at a constant rate over the whole 15-year period, though the model allows for constant updating of investment strategy.

emphasized in the results.

Figure 1 shows an example GBM for the technology's cost C during a 15-year development period assuming constant negative drift ($a < 0$ in (1)) and constant volatility (b in (1)). Figure 1(a) shows a sample path of cost from 0 to 7 years and the subsequent 95 % prediction interval for cost from 7 to 15 years. The 7-year threshold was chosen arbitrarily as an illustration. Figure 1(b) shows the distribution of the final cost at year 15 given knowledge of the cost in year 7 as shown in (a) (the mode is equal to $C_t e^{-\lambda q(T-t)}$, in which the C_t is the known cost at time $t = 7$ and the other parameters are defined as in Table 1). All costs are expressed as fractions of the initial cost. Note that no sample paths are generated as part of the solution procedure (as in Monte Carlo analysis, for example). Uncertainty is accounted for analytically, as detailed in Section 5.

3.1 Single-factor learning curve

The first cost model used in this analysis uses a single factor to represent the effect of RD&D on cost. This model is

$$\frac{dC}{C} = -\lambda I dt + \gamma(I)^{\frac{1}{2}} dz + \sigma dw \quad (2)$$

in which the parameter λ determines the effectiveness of R&D spending, I is the scalar rate of investment which can vary between 0 and a maximum rate q , γ is the endogenous volatility of cost, σ is the exogenous volatility of cost, and dw and dz are increments of Brownian motion (parameters and variables are summarized in Table 1).³

Of the two volatility parameters, σ best describes uncertainty in exogenous factors such as input costs and γ best describes uncertainty in the effectiveness of the RD&D program, since the second term on the right hand side of (2) is nonzero only when investment is taking place ($I > 0$). This term introduces the possibility that RD&D could at times cause cost to rise. Examples in which this took place include nuclear fusion technology, in which R&D has caused cost and feasibility estimates to be revised unfavorably over time; CCS, in which LBD increased projected costs by revealing unexpected ammonia losses; and fossil fuel technologies including flue gas desulfurization and selective catalytic reduction in which cost increases were observed during early stages of commercialization (Rubin et al., 2007). The exponent of $\frac{1}{2}$ in this term both ensures that increases in I have diminishing marginal effects on prediction intervals of future cost and facilitates the analytical solution.

Taking the expected value of (2) and substituting $\lambda I dt = \frac{\beta}{x} dx$ results in a traditional learning curve of the form

$$C(x) = C_0 x^{-\beta} \quad (3)$$

in which x is cumulative installed capacity and β represents the effectiveness of

³Although (2) has two Brownian motions and may thus seem to contradict the formulation in (1), note that dz and dw are normally distributed random variables and consequently so are their linear combinations.

Table 1: Parameters and variables in the development/deployment option valuation for a single-factor cost model. Learning parameters are adapted from Drury et al. (2009) and Neij (2008)

Symbol	Definition	Base value	Units
C	Net present CO ₂ mitigation cost	-	billion USD
V	Value of investment opportunity	-	billion USD
I	Rate of R&D/LBD investment	-	billion USD/y
q	Maximum rate of R&D investment	0.4	billion USD/y
λ	Effectiveness of R&D/LBD spending	0.13	(billion USD) ⁻¹
σ	Exogenous uncertainty of C	0.05	y ⁻¹
γ	Endogenous uncertainty of C	-	y ⁻¹
ρ	Correlation between R&D/LBD uncertainty	0.5	-
T	Duration of RD&D program	15	y
P	Social benefit of CO ₂ mitigation with solar PV	50	billion USD
μ	Annual discount rate	0.03	-
t	Time after 2015	-	y
z, w	Standard Brownian motions	-	-
C_l	Lower cost threshold for investment decision	-	billion USD
C_u	Upper cost threshold for investment decision	-	billion USD

R&D/LBD. This substitution implies that the two formulations are equivalent if a set amount of investment (under the proposed model) has the same effect as a set percentage increase in capacity (under the traditional model). Although this equivalence does not hold in general, it illustrates the similarity of the two cost functions and the analogy allows for the estimation of plausible parameter estimates.

3.2 Two-factor learning curve

While (2) combines the effects of R&D and LBD into a single process, the second cost model disaggregates their effects as drivers of cost reductions. This model is expressed as

$$\frac{dC}{C} = -(\lambda_R I_R + \lambda_L I_L + \alpha)dt + \gamma_R (I_R)^{\frac{1}{2}} dz_R + \gamma_L (I_L)^{\frac{1}{2}} dz_L + \sigma dw \quad (4)$$

in which I_R and I_L are the rates of spending on R&D and demonstration projects to promote LBD (there is a budget constraint such that $I_R + I_L \leq q$), λ_R and λ_L represent the expected effectiveness of R&D and LBD spending in reducing cost, α is the exogenous drift

Table 2: Parameters and variables in the RD&D investment model for the two-factor cost model. Learning parameters, with the exception of ρ , are derived from Savitz et al. (2007), Baker and Peng (2012), and Chan et al. (2011).

Symbol	Definition	Base value	Units
I_R	Rate of R&D investment	-	billion USD/y
I_L	Rate of LBD investment	-	billion USD/y
q	Maximum rate of investment in R&D/LBD	250	million USD/y
α	Exogenous drift rate of cost	0.05	y ⁻¹
λ_R	Effectiveness of R&D spending	0.12	(billion USD) ⁻¹
λ_L	Effectiveness of LBD	0.2	(billion USD) ⁻¹
γ_R	Uncertainty in R&D effect	5	%
γ_L	Uncertainty in LBD effect	3	%
σ	Exogenous uncertainty of C	10	%
ρ	Correlation between R&D/LBD uncertainty	0.2	-
T	Duration of RD&D program	15	y
μ	Annual discount rate	2	%
z_R, z_L, w	Standard Brownian motions	-	-
s	Social cost of carbon	27	\$/tCO ₂
c	IGCC/CCS capacity factor	0.8	-
a	Avoided CO ₂ for IGCC/CCS	0.65	t/MWh

rate of cost, γ_R and γ_L are the endogenous volatilities of R&D and LBD, σ represents exogenous uncertainty, and dw , dz_R , and dz_L are all standard Brownian motions, the latter two of which have correlation coefficient ρ ($E[dz_R dz_L] = \rho dt$). See Table 2 for a summary of parameter definitions and values.

The cost function in (4) is analogous to a two-factor learning curve, reviewed in Section 1, which expresses cost as

$$C(x, y) = C_0 x^{-\beta_1} y^{-\beta_2} \quad (5)$$

in which x is installed capacity, β_1 represents the rate of LBD, y is accumulated knowledge gained through R&D (both spending and number of patents have been used as proxies), and β_2 represents the effect of R&D spending.

While most previous research on learning curves takes installed capacity as an independent variable, this thesis instead takes spending as an independent variable. This

requires the assumption that government investment results in LBD, either through subsidies to drive learning through capacity expansions or through the construction of demonstration projects.

4 Parameter estimates for low-carbon energy technologies

Numerical examples and full solutions are presented for both the single factor cost model described in Section 3.1, applied to solar PV, and the two-factor cost model described in Section 3.2, applied to CCS.

4.1 Solar photovoltaics: single-factor learning curve

To evaluate the RD&D investment decision for the single-factor stochastic cost model, parameters are fit to (2) based on previous work on learning curves, cost and deployment projections, and estimates of the carbon mitigation potential of solar PV. The discount rate in the base case is 3 %, which is consistent with a societal perspective.

The carbon price in the base case is \$20/tCO₂, assumed to reflect the social cost of carbon (SCC), and takes effect $T = 15$ years in the future. This carbon price is assumed to motivate large-scale deployment if solar PV has developed sufficiently in T years to mitigate CO₂ emissions at lower cost. At the deployment time, the social benefit of large-scale solar PV is calculated by subtracting the cost of avoided CO₂ emissions with solar PV from the social cost of the emissions. The assumed SCC is low in the range of current estimates, so the results of this analysis indicate a lower bound for the value of the RD&D program.

The learning rate for solar PV is typically estimated near 20 %, which means that a doubling in capacity would cause cost to decrease by 20 % (see, e.g., Neij (2008)). Using this learning rate and a growth rate in capacity of 17 % per year (rounded from Drury et al. (2009)), solar PV would achieve just over three doublings in 15 years and its cost would decrease by 53 %. Neij (2008) reported uncertainty in the learning rate for solar PV

of ± 5 percentage points. Interpreting this as one standard deviation, I set $\sigma = 0.05$.⁴ For the purposes of this illustration, γ and α are set to zero. After 2030, capacity is assumed to grow and cost to decrease deterministically at the same rate until 2050.

Solar PV at high penetrations has the potential to mitigate about 1 GtC/year worldwide, or about 0.25 GtC/year in the U.S. given 800 GW installed capacity by 2050 (Drury et al., 2009). With an SCC of \$20/tCO₂, when solar PV is fully deployed in the U.S. (2050 under current assumptions), the total annual social benefit of mitigating CO₂ with solar PV is \$18 billion. Accounting for the deployment rate and discount rate from 2030 to 2060, this amounts to a net present benefit of $P = \$170$ billion in 2030. The cost C can therefore be understood as the net present cost of a similar mitigation strategy given current PV costs.

The maximum annual spending rate q is set to \$400 million USD. Historical DOE spending on solar PV has almost always been between \$50 and \$100 million (2007 USD), with the exception of the early 1980s when it exceeded \$300 million. DOE's 2013 budget request for solar (both PV and CSP) is \$310 million. This assumption on q requires that λ be 0.13 to yield the appropriate expected reduction in cost under full R&D investment of \$400 million/y.

4.2 Carbon capture and sequestration (CCS): two-factor learning curve

Given the United States' vast coal and natural gas resources, carbon capture and sequestration (CCS) will likely become an important technology in the 21st century if a stringent climate policy is enacted. CCS remains largely undemonstrated and subject to substantial disagreement among experts on its cost and readiness for large-scale deployment, and it could benefit substantially from an expanded RD&D program. CCS technology therefore serves as a relevant case study for the real options valuation with cost

⁴Earlier-stage technologies could have substantially higher values of σ . For reference, σ for the New York Stock Exchange is about 0.2.

based on a two-factor learning curve under uncertainty.

Studies on cost projections for CCS rest either on expert elicitations or on empirical analyses of cost trajectories of analogous fossil fuel technologies such as flue gas desulfurization. The former includes Baker and Peng (2012), Savitz et al. (2007), and Chan et al. (2011); the latter includes Rubin and Zhai (2012). Only the expert elicitations report uncertainty estimates, so they are used as the basis for parameter estimates in this work.

Based on a Monte Carlo analysis of expert elicitations, Chan et al. (2011) derived cumulative distribution functions (CDFs) for the capital cost of a new integrated gasification combined cycle (IGCC) coal plant with CCS. Baker and Peng (2012) presented cost CDFs from their own expert elicitation and that of a National Academy of Sciences study (Savitz et al., 2007) for CCS mitigation cost per tCO₂. In the absence of data from Chan et al. (2011) on the experts' judgments of the cost of a coal plant without CCS, these estimates cannot be converted to CO₂ mitigation costs. I assume that most of the reduction in capital cost in Chan et al. (2011) is in the CO₂ capture system rather than other parts of the coal plant, and that capital cost is the dominant factor in determining mitigation cost, so that a percentage reduction in capital cost approximates the same percentage reduction in mitigation cost per tCO₂.

The approach to valuing the investment opportunity is based in part on Savitz et al. (2007), a National Research Council (NRC) report that examined the potential effectiveness of a U.S. Department of Energy (DOE) program in R&D of CCS. A main finding was that given the anticipation of a carbon policy, the fossil fuel industry would engage in substantial R&D activity, and therefore the purpose of a DOE program to fund CCS would be to accelerate learning by only about three years. Although this effect seems small, the DOE program would have an NPV of about \$3.5 billion and would be worth undertaking. The NRC report assumed a carbon tax of \$100/tC (\$27/tCO₂) specifically to drive private sector investment; Baker and Peng (2012) sought to quantify expected cost reductions resulting from government R&D without considering the private sector. Nevertheless, both

studies found similar business-as-usual (BAU) rates of cost reduction (see Table 3).

In keeping with the finding that cost reductions are likely to occur in the BAU scenario, the value of undertaking the RD&D program is calculated as the difference between the value of the development/deployment option given the availability of the government program and the value of the option under the BAU scenario for a period of three years. This time period is chosen with the expectation that private sector RD&D would achieve similar costs with or without a government program by about three years after the carbon policy is enacted (after Savitz et al. (2007)).

The following subsections describe derivations of individual parameter estimates from previous work. Table 3 shows a comparison of key parameter estimates from different studies. All parameters are very uncertain and are thus treated with sensitivity analyses in the results.

Maximum yearly R&D/LBD spending: q

Chan et al. (2011) elicited recommended levels of government funding for a range of CCS technologies. The experts recommended that \$430 million/y be allocated to capture technologies, split with \$260 million/y allocated to basic and applied research and \$170 million/y to experiments, pilots, and commercial demonstration. Savitz et al. (2007) assumed mean annual spending of about \$220 million from 2001 to 2025. Baker and Peng (2012) assumed \$110 million for a “high” funding trajectory annually from 2010-2020. The base value for q is therefore taken as the rounded mean of these values, \$250 million/y.

BAU rate of cost reduction: α

The BAU drift rate of cost was calculated by taking the mean log cost under the BAU scenarios at the end of the study period and subtracting it from the log of the current estimated mitigation cost for pre-combustion capture of \$51/tCO₂ (Baker and Peng, 2012) (or, in the case of Chan et al. (2011), the log of the mean current cost estimate for IGCC/CCS). The BAU drift rate is obtained by dividing this quantity by the length of the

RD&D program each study assumed (15-20 years). Baker and Peng (2012) and Savitz et al. (2007) yielded similar estimates for α while that of Chan et al. (2011) was an order of magnitude lower. Since the former two studies more explicitly model CO₂ mitigation cost, they are given more weight in the estimate and the base value for α is taken as 0.05.

It is noteworthy that Savitz et al. (2007) assume a carbon tax to drive private sector R&D investment, while Baker and Peng (2012) do not; nevertheless, the two studies imply similar values for the BAU drift rate in cost, α .

Effectiveness of spending on R&D/LBD: λ_R and λ_L

None of the three studies distinguishes between cost reductions due to R&D and those due to LBD so they cannot inform direct estimates of λ_R and λ_L . Only an overall effectiveness parameter can be calculated, equal to the average of λ_R and λ_L weighted by their respective rates of spending.

To measure the overall effectiveness of spending, the drift rate of cost under the R&D scenario for each study was calculated and the drift rate of cost for the BAU scenario (α) was subtracted from it. The result was then divided by the maximum spending rate q to obtain the cost reduction per dollar per year.

Table 3 shows that the implied effectiveness of spending from the three studies varies over more than an order of magnitude. Chan et al. (2011) and Savitz et al. (2007) have similar implied rates of spending effectiveness, while that of Baker and Peng (2012) is almost an order of magnitude higher. Baker and Peng (2012) assumed the least funding of the three studies, so the greater spending effectiveness could reflect diminishing returns to scale in R&D investment. This result could also indicate that experts are keying more on the existence of R&D spending than they are on the specific level; along these lines, Chan et al. (2011) note that their expert elicitation does not account for differences in expert opinions on current cost or differences in the experts' familiarity estimating percentiles. The base estimate for spending effectiveness is taken as the mean of those derived from the

Table 3: Comparison of parameter estimates based on three expert elicitations. [1] is Chan et al. (2011), [2] is Baker and Peng (2012), [3] is Savitz et al. (2007), and [4] is Van den Broek et al. (2009).

	[1]	[2]	[3]	[4]
Annual spending (\$M)	430	110	220	n/a
Investment timeframe	2010-2030	2010-2025	2001-2020	2001-2050
BAU cost drift (α)	0.004	0.06	0.07	0.06
Drift with funding	0.02	0.1	0.08	n/a
Spending effectiveness λ	0.04	0.4	0.03	n/a
Current cost uncertainty	0.56	n/a	n/a	n/a
Final cost uncertainty (no funding)	0.10	0.15	0.11	0.03
Final cost uncertainty (funding)	0.09	0.25	0.10	n/a

three studies, 0.16, and treated with sensitivity analysis.

This overall spending effectiveness is divided into λ_R and λ_L . The only apparent basis for this division in previous work is Lohwasser and Madlener (2010), who estimated a two-factor learning curve for CCS by developing an analogy with flue-gas desulfurization. Lohwasser and Madlener (2010) found a LBD rate of 7.1 % and a learning-by-researching (R&D) rate of 6.6 %, using the number of patents as the independent variable for learning-by-researching. These parameter values are very similar, and no other studies corroborate them. Based on the conjecture that compared with the effect of R&D, LBD is significantly more likely to result in steady, incremental cost decreases and less likely to result in major decreases, the spread between λ_R and λ_L is adjusted upward. Base values are taken as $\lambda_R = 0.12$ and $\lambda_L = 0.2$.

Exogenous uncertainty: σ

Exogenous uncertainty is determined by fitting lognormal distributions to CDFs of future cost for the BAU scenarios from the three expert elicitation studies. Exogenous uncertainty is equal to 10–15 % (standard deviation) per year in all three studies. Assuming an irresolvable cost uncertainty of 20 % that applies at all times, the drift rate becomes 8–12 %. The base value is taken as 10 %.

In Chan et al. (2011), it is noteworthy that the experts indicate similar degrees of

uncertainty for the 2030 BAU case and the 2010 reference case (in the raw data, one out of seven experts gave a point estimate for 2030 cost and a range for 2010 cost). This suggests that the experts may be insufficiently considering the additional uncertainty in input cost developments between 2010 and 2030.

Endogenous uncertainty in effectiveness of R&D/LBD: γ_R and γ_L

Endogenous uncertainty is estimated by fitting lognormal distributions to the CDFs for cost under government funding in the three studies and comparing the standard deviations with the previously calculated values of σ . Both Chan et al. (2011) and Savitz et al. (2007) imply that there is no change in uncertainty (as a percentage of cost) under government funding, while Baker and Peng (2012) imply a 50 % increase. Adopting the value from Baker and Peng (2012), the endogenous uncertainty is equal to about 30 % per year. Assuming $\gamma_R > \gamma_L$ and that the rate of spending is at its maximum, γ_R and γ_L are taken as 0.5 and 0.3.

Using the above parameter estimates (which are summarized in Table 2), Figure 2 shows the expected value and prediction intervals for mitigation cost given exclusive investment in R&D (Figure 2(a)) and LBD (Figure 2(b)). The moderate probability of initial cost increase with LBD is consistent with the negative learning effects found by Rubin et al. (2007).

The higher volatility of the effect of R&D results in a high upper bound on the 90 % prediction interval, though there is low probability mass near this bound. A large stochastic increase could be due to the revelation of negative health or environmental consequences of a component of the technology. If this probability is deemed too large for the case of CCS, it should be noted that at the deployment time the RD&D program is worthless whether cost is significantly above or only slightly above the assumed carbon price of \$27/tCO₂. The consequence of a large stochastic cost increase is that the technology will not be deployed at scale, which is a realistic outcome even if a large




Figure 2: Expected value and 50 % and 90 % prediction intervals for CO₂ mitigation cost with CCS under base case assumptions for (a) exclusive investment in R&D and (b) exclusive investment in LBD. The RD&D program is assumed to last 15 years, after which a carbon policy is enacted and the deployment decision is made. Cost and spending are modeled as real (inflation adjusted).

stochastic increase in cost due to R&D spending is deemed unrealistic.

As a point of clarification, the cost trajectories represent the expected cost if the technology were deployed in year t . They do not portray the cost uncertainty surrounding that expectation, which is assumed to decrease with time until it vanishes at deployment time T .

Correlation between the effect of R&D/LBD: ρ

A positive correlation between the effects of R&D and LBD reflects the possibility of the two factors contributing to the same technical advances or knowledge. An example could be either a result from operations research or trial-and-error in factories producing the same improvement in manufacturing. There is no apparent previous research along these lines to inform estimates of the value of ρ .

If $\rho = 0$, the only realized investment outcomes are full investment in R&D, full

investment in LBD, and no investment (though the strategy can switch continuously). When R&D and LBD are independent, one of them clearly dominates the other in its effectiveness at reducing cost. When their effects interact, split investment is optimal. Since $\rho < 0$ is unlikely, to avoid the “bang-bang” solution, ρ is arbitrarily chosen as 0.2.

Final value of the investment opportunity

At time T , when the CO₂ price takes effect, the RD&D program is terminated, and the deployment decision is made, the final value to society of deploying CCS at scale can be calculated as a function of mitigation cost. A linear demand curve for installed IGCC/CCS capacity is derived from Savitz et al. (2007), who found that a 0.4 ¢/kWh increase in electricity cost due to CCS (equivalent to about \$6/tCO₂) would result in about 75 GW installed capacity by 2025, a 1¢/kWh increase (about \$15/tCO₂) would result in 35 GW, and a 1.8 ¢/kWh increase (about \$27/tCO₂, the assumed carbon price) would be prohibitively expensive.

Since the NRC report (Savitz et al., 2007) found that a government R&D program would provide benefits above private-sector R&D for only about three years, avoided emissions are valued three years into the future. Installed CCS capacity is assumed to grow at the same (linear) rate from 2025–2028 as it did between 2024 and 2025. The total benefit of CCS deployment at the end of the option lifetime T , with mitigation cost $\$/\text{tCO}_2$, is calculated as

$$V(C, T) = (s - C) \sum_{i=1}^y \frac{G_{C,y} h c a}{(1 + \mu)^{i-1}} \quad (6)$$

in which s is the social cost of carbon, y is the year following deployment, $G_{C,y}$ is the cumulative deployment (in GW) of IGCC/CCS in year y , h is the number of hours in a year, c is the capacity factor, a is the avoided CO₂ per MWh, and μ is the discount rate. With the parameter assumptions shown in Table 2, the value of the CCS deployment opportunity at the end of the RD&D investment period is that shown in Figure 3. To

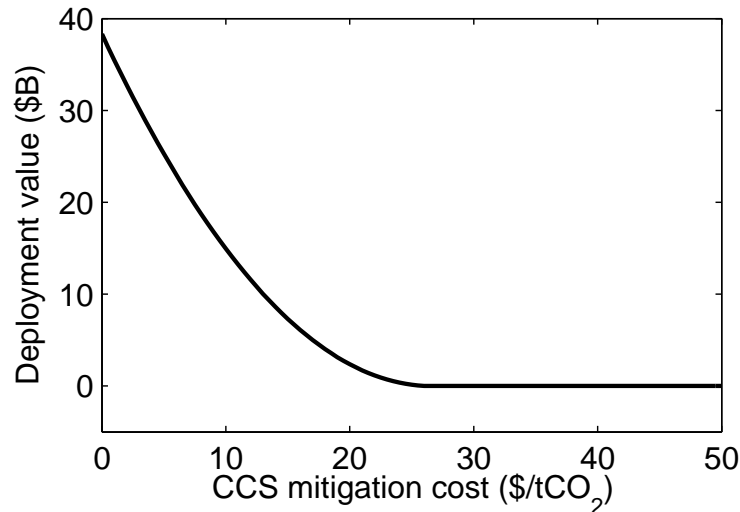


Figure 3: Social benefit of CCS deployment as a function of mitigation cost.

determine the optimal RD&D investment strategy, the development/deployment option is valued with (4) as the cost function. To calculate the value of the RD&D program, the value of the development/deployment option under a business-as-usual case (with $I_R = I_L = 0$) is subtracted from the value under optimal RD&D investment.

5 Solution method

The problem the decision maker faces is to invest in RD&D to maximize the value of the technology less cumulative spending at time T (or minimize wasted investment, if the technology is not cost-effective at T), accounting for uncertainty and BAU cost reductions. This problem can be solved using the Bellman equation of dynamic programming, which expresses the value of the investment opportunity as the maximized sum of the payoff of a current choice plus the expected future payoff of the investment opportunity as a consequence of that choice. In this case, the Bellman equation can be expressed as

$$\mu V dt = \max_{\mathbf{I}} \{ \mathbf{E}[dV] - \mathbf{1}^\top \mathbf{I} dt \}. \quad (7)$$

in which V is the value of the investment opportunity, μ is the discount rate, \mathbf{E} is the expected value operator, $\mathbf{1}$ is a vector of ones of the same length as \mathbf{I} , and \mathbf{I} is the rate of RD&D investment (see Appendix 1 for a full derivation). Equation (7) implies that the appreciation in value at the discount rate μ equals the incremental change in the expected value of the investment opportunity less RD&D expenditures \mathbf{I} , maximized with respect to \mathbf{I} .

5.1 Single-factor learning curve

Proceeding with the solution for a single RD&D factor driving cost reductions, since V is a function of C and t , dV can be expanded with Itô's Lemma as

$$dV = \frac{\partial V}{\partial C}dC + \frac{1}{2}\frac{\partial^2 V}{\partial C^2}(dC)^2 + \frac{\partial V}{\partial t}dt \quad (8)$$

(Note that $(dC)^2 = \sigma^2 C^2 dt + o(dt)$, where $o(dt)$ contains terms that go to zero faster than dt .) Setting $\alpha = \gamma = 0$, substituting (2) and (8) into (7), dividing by dt , and rearranging yields the differential equation

$$\mu V = \max_I \left\{ I \left(-\lambda C \frac{\partial V}{\partial C} - 1 \right) + \frac{1}{2} \sigma^2 C^2 \frac{\partial^2 V}{\partial C^2} + \frac{\partial V}{\partial t} \right\} \quad (9)$$

Since (9) is linear in I , the optimal control I is 0 when its multiplier is negative and the maximum q when its multiplier is positive. This framework neglects changing returns to scale on R&D investment and implies that the optimal rate of investment is either 0 or the maximum feasible amount.

For the deterministic case, where $\sigma = 0$, the maximization problem can be solved exactly (see Appendix 2).

To solve the problem for the stochastic case, first note that boundary conditions for $V(C, t)$ are known. At the deployment time T , the investment decision becomes a classic “now-or-never” decision based on NPV and the value of the investment opportunity is $\max\{P - C, 0\}$, where $P - C$ is the difference in abatement cost using the backstop

technology and the modeled technology (expressed in (10a); note that a simpler model for the deployment value is used for this single-factor model than for the two-factor model).

The Bellman equation (7) can be solved at each time step by iterating backwards, with the additional boundary conditions:

$$V(C, T) = \max\{P - C, 0\} \quad (10a)$$

$$V(0, t) = Pe^{-\mu(T-t)} \quad (10b)$$

$$\lim_{C \rightarrow \infty} V(C, t) = 0 \quad (10c)$$

$$\left. \frac{\partial V}{\partial C} C \right|_{C \in \{C_l, C_u\}} = -\frac{1}{\lambda} \quad (10d)$$

Equation (10b) states that as the cost of the modeled technology approaches 0, uncertainty collapses as well and the value of the investment opportunity equals the discounted social benefit of deploying the technology. (10c) implies that as the cost of the modeled technology becomes very high, it will never be cost-effective and the investment opportunity is worthless. (10d) arises from the value matching condition across the free boundary: the multiplier of I in (9), $(-\lambda C \frac{\partial V}{\partial C} - 1)$, equals 0 at the boundary whether $I = 0$ or $I = q$. This requirement necessitates the smooth pasting condition expressed in (10d).

The resulting PDE is solved numerically in MATLAB using the Crank-Nicolson method (see Appendix 3). The solution structure is shown in Figure 4. If the cost of the technology C is below the threshold C_l , further investment in RD&D will not lower costs sufficiently to justify the expense. If cost is above C_u , the likelihood that cost will fall sufficiently to justify deployment is too low for RD&D investment to be worthwhile. If cost is between C_l and C_u (the shaded region in Figure 4) investment at the maximum rate is optimal. Because C is stochastic, it could cross each threshold multiple times during the 15-year development period.

The free boundaries C_l and C_u vary with all parameters in (9), including the uncertainty σ . The upper investment boundary C_u tends to increase with σ due to the




Figure 4: Optimal RD&D investment strategy as a function of cost C and time t (representing the year of the RD&D program prior to possible deployment at $T = 15$ years) under the current framework. Above C_u the cost of the technology is unlikely to fall sufficiently by time T , so zero investment is optimal. Below C_l , the cost of the technology is so low that further RD&D investment is not cost-effective, so again zero investment is optimal. Between the two thresholds, investment is optimal at the maximum rate.

larger probability that the cost may stochastically decrease. C_l and C_u are found using an iterative method adapted from Muthuraman (2008) (see Appendix 3 for details on the solution method).

5.2 Two-factor learning curve

For the cost model with two factors corresponding to R&D and LBD (4), the Bellman equation (7) reduces to

$$\mu V dt = \max_{I_R, I_L} \{E[dV] - I_R dt - I_L dt\} \quad (11)$$

in which μ is the discount rate, V is the value of the investment opportunity (as an implied function of C and t), I_R is the rate of R&D spending, and I_L is the rate of spending on

LBD. Applying Itô's lemma and substituting (4) and (8) into (11) yields

$$\begin{aligned} \mu V = \max_{I_R, I_L} \left\{ \left(\frac{1}{2} \gamma_R^2 I_R + \frac{1}{2} \gamma_L^2 I_L + \rho \gamma_R \gamma_L \sqrt{I_R I_L} + \frac{1}{2} \sigma^2 \right) C^2 V_{CC} \right. \\ \left. - (\lambda_R I_R + \lambda_L I_L + \alpha) C V_C - (I_R + I_L) + V_t \right\} \quad (12) \end{aligned}$$

The boundary conditions for this partial differential equation are

$$V(C, T) = \max\{P - C, 0\} \quad (13a)$$

$$V(0, t) = V(0, T) e^{-\mu(T-t)} \quad (13b)$$

$$\lim_{C \rightarrow \infty} V(C, t) = 0 \quad (13c)$$

$$\left(\frac{1}{2} \gamma_R^2 I_R + \frac{1}{2} \gamma_L^2 I_L + \rho \gamma_R \gamma_L \sqrt{I_R I_L} \right) C^2 V_{CC} - (\lambda_R I_R + \lambda_L I_L) C V_C - (I_R + I_L) \Big|_{C \in \{C_l, C_u\}} = 0 \quad (13d)$$

Equation (13b) states that as the cost of the modeled technology approaches 0, uncertainty collapses as well and the value of the investment opportunity equals the discounted deployment value, $V(0, T)$. Equation (13c) implies that as the cost of the modeled technology becomes very high, it will never be deployed and the investment opportunity is worthless. Equation (13d) comes from the value matching condition across the free boundary: the part of the function that depends on I_R and I_L in (9) equals 0 at the boundary whether $I_R = I_L = 0$ or $I_R + I_L = q$.

Figure 4 again conceptually shows the optimal investment rule as a function of cost and time. It can easily be shown that in the region marked “No Investment”, $I_R = I_L = 0$ and in the region marked “Maximum investment”, $I_R + I_L = q$ (see Appendix 4). The solution thus switches abruptly (is “bang-bang”) in terms of whether or not investment is taking place, but admits a range of values for I_R and I_L .

Equation (12) is solved numerically by discretizing in time and applying MATLAB's built-in boundary value problem solver *bvp4c* to iteratively solve the resulting system of

free-boundary ordinary differential equations (ODEs). Each ODE is solved with an initial guess for the free boundaries C_l and C_u and the value of the I_R and I_L for $C_l < C < C_u$, after which the I_R and I_L are updated with the new values for V_C and V_{CC} . This procedure is repeated until the I_R and I_L converge. The locations of the free boundaries are then updated according to (13d), the I_R and I_L are updated, and the procedure is repeated until both the free boundaries and the I_R and I_L converge. This method performs well except at very high volatilities and values of C , for which the I_R and I_L can fail to converge.

Appendix 4 presents details of the solution method.

5.3 Two competing technologies

Analyzing optimal investment in two competing technologies would yield insight into strategies for hedging against uncertainty by investing, at least initially, in competing technologies. While this thesis does not contain a detailed analysis of this case, the solution structure is developed in Appendix 5.

6 Results

As described in Section 3, this analysis takes the perspective of a decision maker within the U.S. federal government who administers a clean energy technology development program. The decision maker has the opportunity to invest in an energy RD&D program for T years before a climate policy is enacted. At this time, the technology is deployed at large scale if it is cost-effective and abandoned otherwise.

6.1 NPV analysis

Under the classical method of NPV-based project planning, a decision maker with this investment opportunity would calculate the RD&D strategy that minimizes the expected cost of the technology less total RD&D spending T years in the future, adopt that strategy if its discounted expected value were positive, and abandon the project otherwise.

To analyze the deterministic case, σ and γ are set to 0. As in the stochastic case, the optimal investment strategy is illustrated schematically by Figure 4. If the initial cost of the technology is above C_u , investment will never be optimal under any circumstances. If the initial cost is between C_l and C_u , the deterministic cost will stay within the investment region (shaded in Figure 4) for the whole 15-year period and constant investment is optimal. If the initial cost is below C_l , over time cost may equal C_l at which point it will be optimal to invest and the cost of the technology will track C_l for the remainder of the investment period. If cost is very low, it will never equal C_l and investment will never be optimal. Note that for the deterministic case, stopping investment after it has started is never optimal due to perfect foresight and the positive discount rate.

Figure 5 shows the value of the investment opportunity and the cost thresholds C_u and C_l for the base scenario for both the single factor and two-factor case studies, under uncertainty and with an NPV analysis. Figures 5(a) and (c) show that if initial cost is very low, the value of the investment opportunity under uncertainty is similar to NPV (partly because uncertainty is proportional to cost) and zero RD&D investment optimal. As initial cost rises, NPV drops off quickly and reaches 0 at a present mitigation cost of \$400 billion for the single factor analysis of solar PV and \$100/tCO₂ for the two-factor analysis of CCS. Accounting for cost uncertainty and the ability to alter RD&D strategy in response to changing costs adds substantial value to the investment opportunity. When initial cost is between \$400 billion and \$700 billion for the solar PV analysis and \$100 and \$400/tCO₂ for the CCS analysis, investment in RD&D is optimal even though NPV is negative due to the possibility of costs stochastically falling and the subsequent opportunity to invest and lower costs further. At higher initial costs, zero investment is optimal since stochasticity and RD&D spending are unlikely to render the technology cost-effective by the deployment date.

Figures 5(a) and (c) show that for moderately high initial costs, accounting for uncertainty captures the value of potential stochastic cost decreases and such that initial

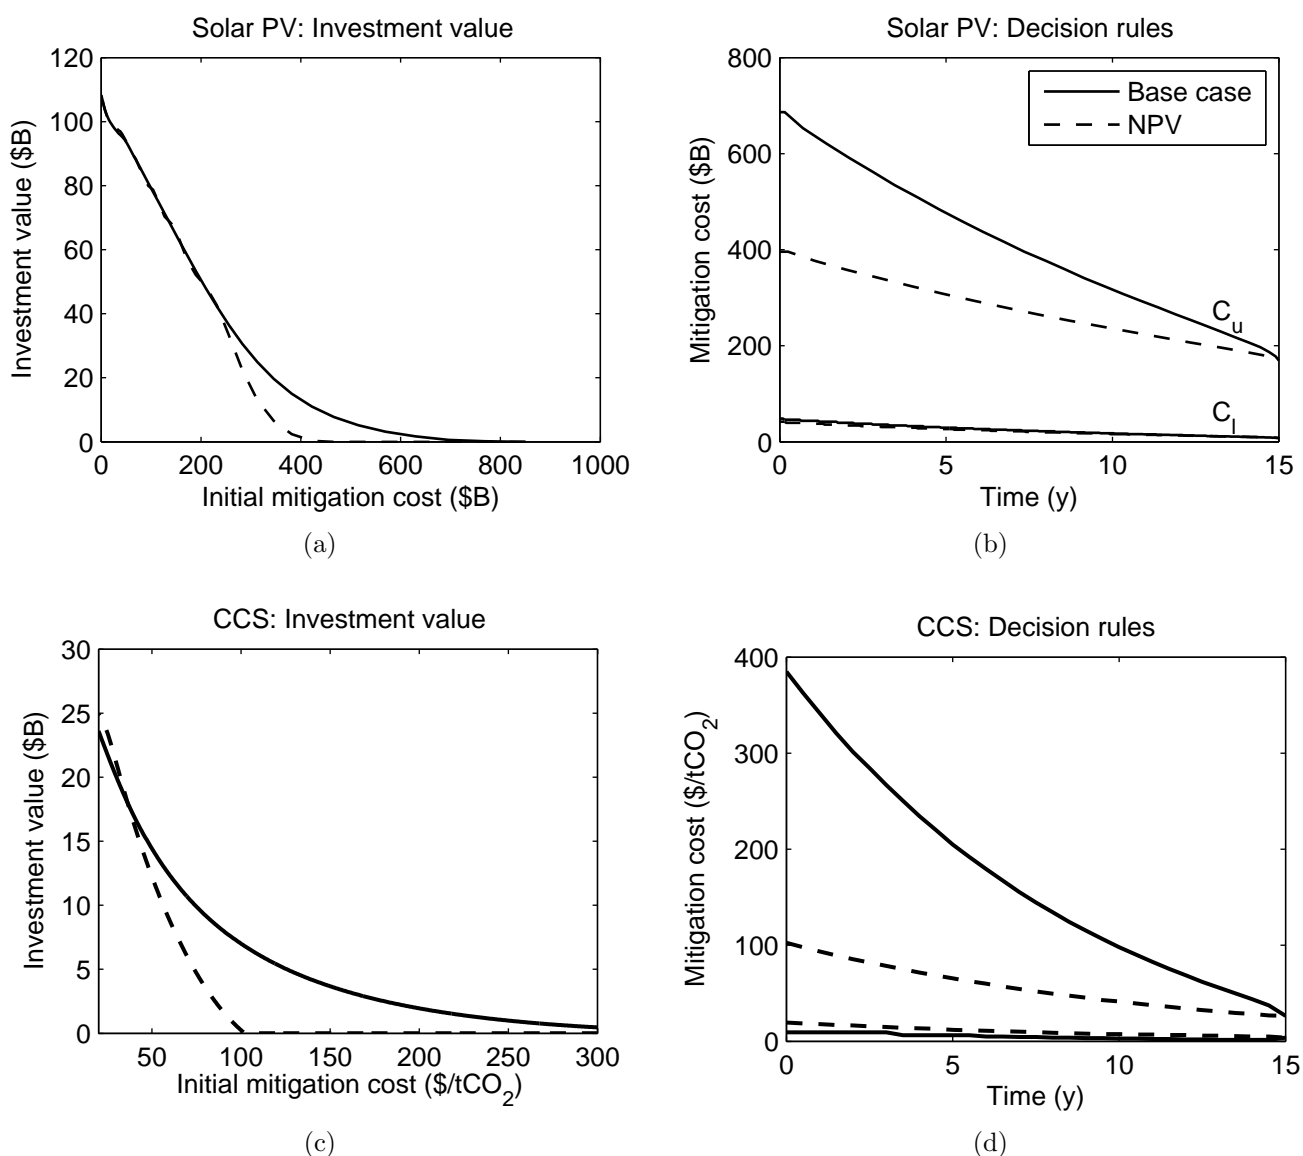


Figure 5: Comparison of NPV results and SDP-based results under base case assumptions for the single-factor analysis of solar PV (a and b) and for the two-factor analysis of CCS (c and d). (a) and (c) show that the NPV is the same as the result under uncertainty at low initial costs, but as cost rises the NPV quickly goes to zero while the value under uncertainty remains positive due to the possibility of cost stochastically decreasing. (b) and (d) show the two thresholds, C_l and C_u , for the RD&D investment decision as shown in Figure 4 (full investment is optimal for the region between the dashed lines for the deterministic NPV analysis and between the solid lines for the stochastic analysis). C_l is largely unchanged by uncertainty and C_u is higher in the stochastic case, indicating that initial investment in an unprofitable project could be optimal due to uncertainty. This effect shrinks as time approaches the deployment time T . The effects of uncertainty are more pronounced for the two-factor case, since this model incorporates endogenous as well as exogenous uncertainty into the cost model, resulting in greater volatility.

investment in a negative-NPV project could be optimal. As the deployment time T nears, this effect diminishes since there is less time for stochastic cost decreases to take place, and C_u converges to that of the deterministic case.

6.2 Results under uncertainty: single-factor cost model

Figure 6 shows the value of the RD&D investment opportunity under uncertainty with base-case assumptions presented in Section 3. With a social benefit of deploying solar PV of approximately $P = \$170$ billion (calculated from a relatively low estimate of the social cost of carbon, $\$20/\text{tCO}_2$), the ability to invest in RD&D adds substantial value to the development/deployment opportunity if initial mitigation cost is high, as discussed in the previous section. At time $T = 15$ years, the value of the investment opportunity simplifies to an NPV analysis and is linear in cost or equal to 0. Holding cost constant, at low initial cost the time derivative of the value function is positive due to the discount rate and at high initial cost the time derivative is negative due to the effect of uncertainty.

Figure 7 shows comparative statics for parameters in the model. Figure 7(a) shows that uncertainty in the future cost of the low-carbon energy technology can add substantial value to the investment opportunity, but only if the initial cost is more than about twice the backstop cost. An uncertainty of 5 % per year, the base assumption, adds little value, though larger uncertainties can add great value even for high initial costs. Figure 7(b) shows that the upper investment boundary rises with higher σ , such that for $\sigma = 16$ % it is initially optimal to invest at rate q even if cost is $\$700$ billion. The lower investment boundary is insensitive to σ .

Figures 7(c) and 7(d) show that the value of the investment opportunity and the upper cost threshold for investment both increase substantially with λ , since this parameter controls the expected rate of cost reduction with respect to RD&D investment.

Figures 7(e) and 7(f) show that for the range of parameters examined, the value of the investment opportunity is less sensitive to q than it is to λ : the ability to spend more does




Figure 6: Value of the R&D investment opportunity for $P = \$170$ billion, $q = \$400$ million/y, $\lambda = 0.13$, $\sigma = 0.05$, and $\mu = 0.03$. At time $T = 15$ years, the deployment decision reduces to a now-or-never investment decision based on NPV. Uncertainty adds value when the initial cost of the technology is high.

not add as much value as an increase in the effectiveness of spending.

Figures 7(g) and 7(h) show discount rate sensitivity. The discount rate for the base case is 3 %, which is appropriate from a social perspective. A higher, corporate discount rate would drastically reduce the value of the investment opportunity and would raise C_i , since R&D investment becomes more expensive relative to future payoff under high discount rates.

6.3 Results under uncertainty: two-factor cost model

Optimal RD&D investment strategy

Figure 8 shows the optimal RD&D strategy for IGCC/CCS beginning 15 years before the carbon policy enactment and deployment opportunity. The cost axis of Figure 8 spans the range of 11 estimates from the literature (compiled by Baker and Peng (2012)), and the optimal investment strategy varies with both time and cost. With a carbon policy and possible CCS deployment anticipated 15 years in the future, an initial cost estimate above \$70/tCO₂ would indicate that investment at the maximum rate \$ q billion/y is optimal, with over 75 % of it directed to R&D and the rest to promotion of LBD through demonstration projects. An initial cost estimate below about \$40/tCO₂ would still merit maximum investment, but with 75 % in promotion of LBD and the remainder in R&D. If the cost estimate stochastically rises, larger amounts of R&D would be indicated on the chance that the program could generate a breakthrough to decrease cost substantially. An anticipated deployment time nearer in the future—in 10 years instead of 15, for example—would also promote more investment in R&D, since the possibility of quick, drastic cost reductions will outweigh the more certain cost reductions of LBD.

Decision support for different expectations on R&D effectiveness and volatility

The optimal investment strategy is highly sensitive to the parameters controlling the effectiveness of R&D/LBD spending: λ_R and λ_L , which represent the expected

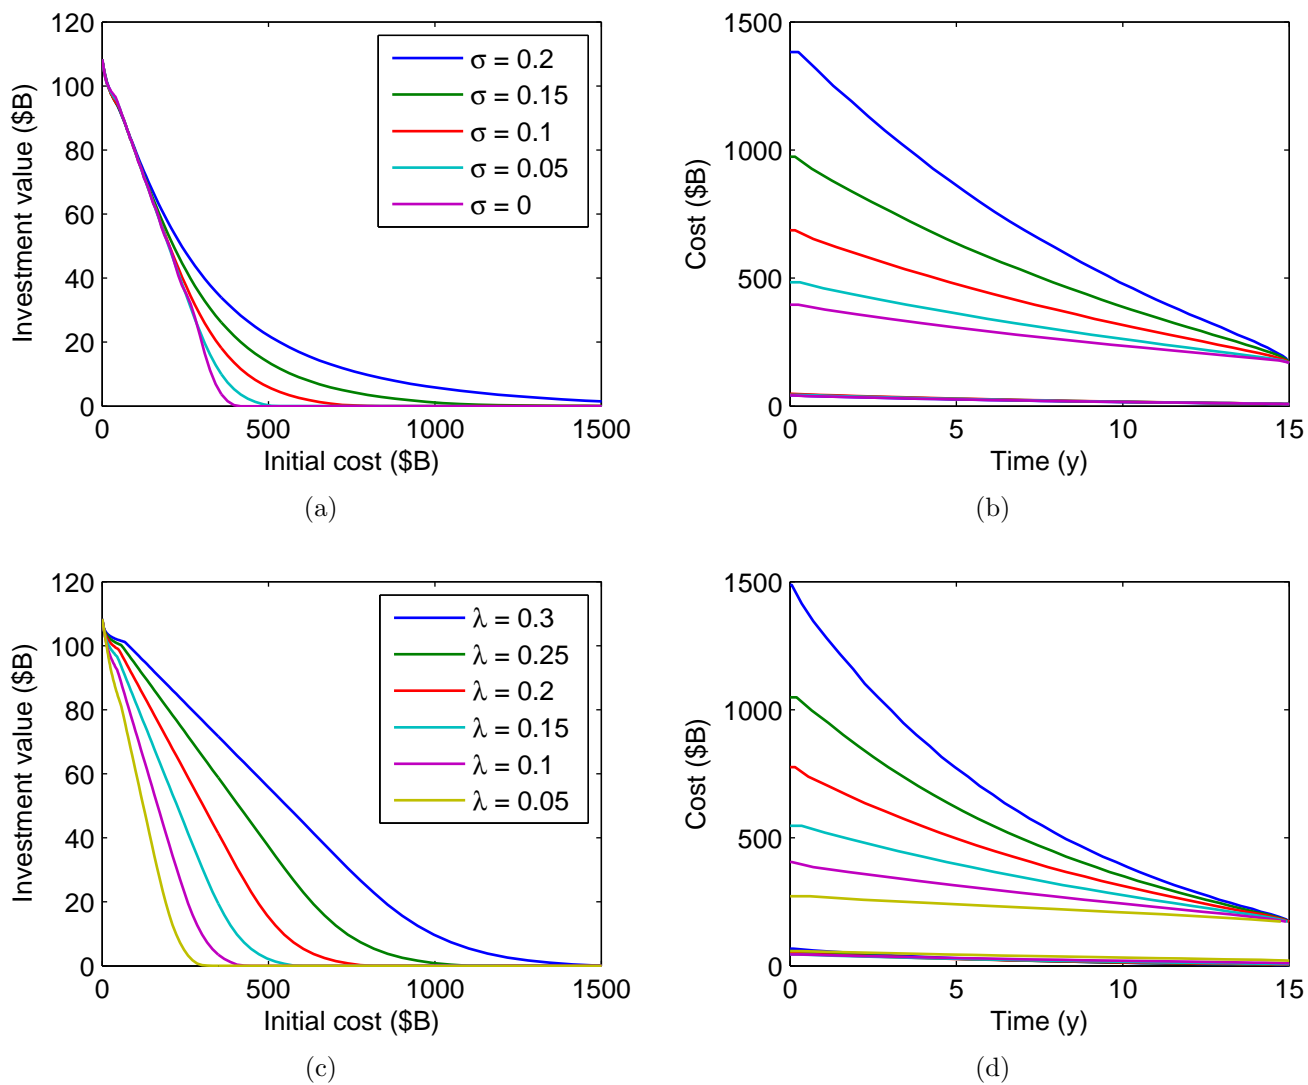





Figure 7: Comparative statics for the main parameters in the solar PV analysis. The value of the R&D investment opportunity with respect to initial cost is shown in the left column. The upper and lower investment boundaries, C_u and C_l , are shown as functions of time in the right column.




(e)



(f)



(g)



(h)

Figure 7: Comparative statics, continued.




Figure 8: Optimal RD&D strategy for CCS beginning 15 years before a carbon policy is enacted and CCS is deployed at scale. In the region marked '0 inv.', CCS is unlikely to reach cost-competitiveness by the deployment time and zero investment is optimal. In all other regions of the figure, full investment at rate $\$q$ billion/y is optimal, partitioned between R&D ('R') and LBD ('L').

effectiveness, and γ_R and γ_L , which represent uncertainty in the effectiveness. Increases in any of these parameters promote greater investment in the corresponding aspect of the RD&D program. The base parameter values, though derived from the literature, are highly uncertain. An appropriate framing for the results is a sensitivity analysis on the main parameters, which can help inform investment decisions to promote technology development under different estimates of the relative effectiveness and volatility of R&D and LBD.

Figure 9 shows the optimal investment strategy under different assumptions on the expected effectiveness and uncertainty associated with R&D (the parameters for LBD are held constant at their base values of $\lambda_L = 0.2$ and $\gamma_L = 0.3$). Over the range of current CCS cost estimates, increasing γ_R to 0.6 (twice γ_L) results in heavy R&D investment being optimal, even when the expected effectiveness of R&D spending is reduced to half that of LBD. Conversely, a lower uncertainty of $\gamma_R = 0.4$, 33 % greater than γ_L , causes LBD to dominate the investment strategy. Consistent over all parameter assumptions presented is the result that full investment in the RD&D program is optimal for the entire range of mitigation cost estimates.

Added value of the RD&D program

Savitz et al. (2007) found that the contribution of a government RD&D program for CCS was to accelerate the technology's development compared with the business-as-usual (BAU) scenario until a short time after the enactment of a carbon policy. Once a policy was in place, the study found, private sector investment would result in cost equivalence under the BAU and government funding scenarios within about three years. Savitz et al. (2007) cited private sector investment as a major driver of BAU cost reductions.

Unlike Savitz et al. (2007), the expert elicitation of Baker and Peng (2012) specifically did not account for private sector R&D; nevertheless, experts predicted BAU cost reductions similar to those estimated by Savitz et al. (2007). To account for the BAU cost




Figure 9: Optimal investment strategy in R&D and LBD of CCS 15 years before deployment for different effectiveness (λ_R) and uncertainty (γ_R) parameters for R&D (the corresponding LBD parameters are held constant at their base values). Increasing the uncertainty associated with R&D spending (and thus the probability of a breakthrough) increases the optimal rate of R&D spending ('R') compared with LBD ('L'). Increasing the effectiveness of R&D spending has a qualitatively similar but lesser effect.

trajectory, the value of undertaking the RD&D program is calculated by subtracting the value of the development/deployment opportunity without the RD&D program (found by setting $q = 0$) from the value with the program. In the former case, changes in cost are determined by the expected rate of BAU cost reductions α and the exogenous uncertainty σ .

Figure 10 shows the value of the RD&D program 15 years before the deployment date under different assumptions on the BAU rate of cost reduction α . The dashed line in the figure represents the value of the program for the deterministic case ($\sigma = \gamma_L = \gamma_R = 0$), equivalent to the difference in NPVs between the deployment opportunity with and without the program, under the base assumption for α . For the stochastic base case, the government program is worth at least \$4 billion over the domain of plausible current CCS carbon mitigation costs derived from the literature. At expected mitigation costs between about \$30/tCO₂ and \$50/tCO₂, the NPV of the RD&D program closely approximates the

value under uncertainty. At mitigation costs below \$30/tCO₂, accounting for uncertainty adds value due to the ability of the RD&D program to drive costs down as the development period progresses if early spending does not decrease costs as expected. If costs are already at commercial levels, RD&D investment may still be optimal in order to increase the amount of CO₂ mitigated and maximize the social benefit of deployment.

Accounting for uncertainty adds substantial value to the RD&D program at initial mitigation costs above \$50/tCO₂. At an initial mitigation cost of \$100/tCO₂, the NPV of the program is 0 while its value under uncertainty is about \$7 billion. The higher value under uncertainty is due to the possibility of stochastic cost decreases and the ability to adjust subsequent RD&D and deployment decisions to updated cost information. The NPV analysis is consistent with Savitz et al. (2007), who found that the government RD&D program had an NPV of \$3.5 billion and that initial CCS costs added 30 % to the cost of energy (which is approximately equivalent to \$30/tCO₂).

Under uncertainty, the assumption on the rate of BAU cost reductions α has an opposite effect at low mitigation costs as it does at high mitigation costs. At costs below about \$40/tCO₂, the program has greatest value assuming no BAU cost reductions; this is due to the fact that cost is already low enough that under substantial BAU cost reductions, the government program provides little value added.

Recent evidence suggests that CCS costs may be higher than previous literature has estimated (see, e.g. Finkenrath (2011)). An important result from Figure 10 is that at costs above about \$80/tCO₂, the program is most valuable under the highest rate of BAU cost reductions, since without the BAU reductions cost would be so high that the government program alone would be unlikely to achieve the reductions needed to render the deployment opportunity valuable.




Figure 10: Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the BAU rate of cost reductions, α (the base assumption is $\alpha = 0.05$).

Sensitivity analyses

Since the evolution of CCS costs over time is highly uncertain, sensitivity analyses were conducted on the parameter assumptions. Figure 11 shows the optimal RD&D investment strategy 15 years before the deployment time under different assumptions on the expected value of the rate of BAU cost reduction α . If no BAU cost reduction is expected to occur ($\alpha = 0$), the majority of the RD&D funds should be directed to R&D, except at very low costs in the estimated current range. If α is high, the large BAU cost reductions enable the government program to immediately enter the phase of LBD promotion except at very high initial cost.

The value of the RD&D program is likewise sensitive to the maximum rate of spending q (Figure 12(a)). If initial mitigation cost is high, increasing the available program funds q adds significant value to the RD&D program, though the effect shows diminishing returns with q . At mitigation costs below about $\$30/\text{tCO}_2$, increasing q from $\$50$ million/y to $\$150$

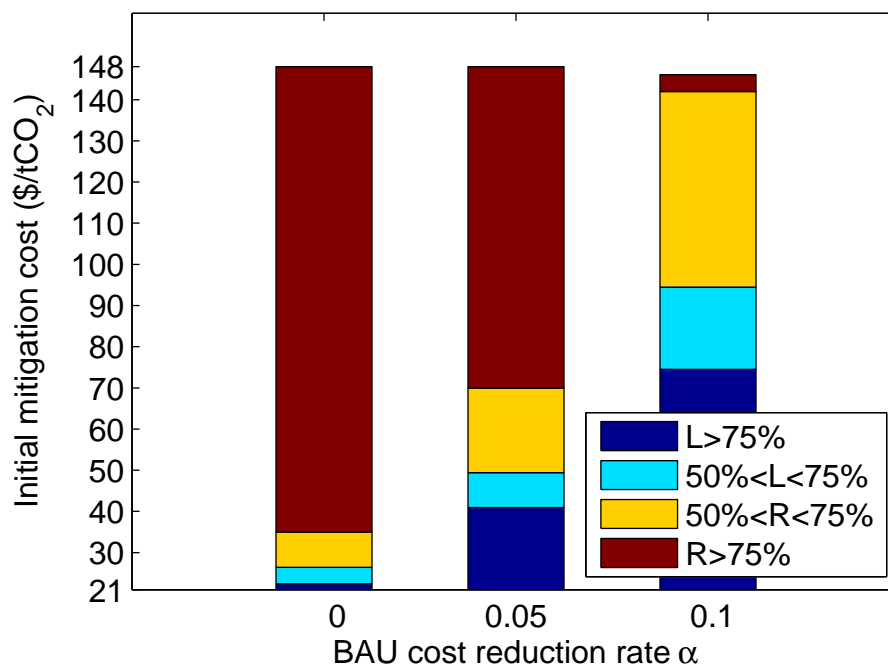



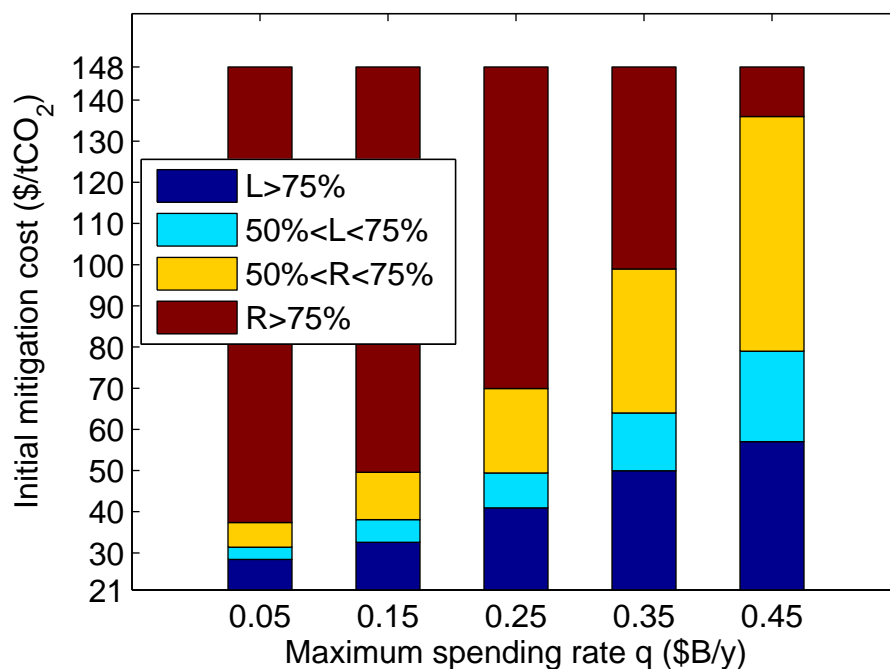
Figure 11: Optimal RD&D investment strategy 15 years before the carbon policy/CCS deployment time under different assumptions on the BAU rate of cost reductions, α (the base assumption is $\alpha = 0.05$).

million/y would significantly increase the value of the program, but increasing q beyond \$250 million/y would add no further benefit. This analysis does not explicitly account for the opportunity cost of investing in other programs, but this result raises the question of the benefit of diversification in the government RD&D investment portfolio. Especially at lower mitigation costs, investment in a separate developmental technology, whose role has little overlap with CCS, could yield greater benefit than further investment in CCS. Quantifying this benefit is left for future work, but Appendix 5 develops the analytical solution for the case of investment in two technologies.

Figure 12(b) shows that if q is low, most of the program funds should be directed towards R&D, except at very low initial mitigation costs. If more funds are available, a greater percentage should be allocated to promoting LBD. This analysis indicates that if only limited program funds are available, they should be directed primarily towards R&D in order to maximize the expected value of the program.



(a)



(b)

Figure 12: (a) Sensitivity of the value of the RD&D program 15 years before the deployment date to q . The program value shows diminishing marginal returns as q is increased, an effect that is more pronounced if mitigation cost is low. (b) Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the maximum rate of spending, q (the base assumption is $q = \$0.25$ billion/y).




Figure 13: Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different assumptions on the correlation between the effects of R&D and LBD, ρ (the base assumption is $\rho = 0.2$).

Figure 13 shows the sensitivity of the optimal RD&D strategy to the correlation between the effects of R&D and LBD, ρ . As correlation increases, the optimal strategy tends toward a more even distribution of R&D and LBD (as expected from examination of (12)). Greater correlation increases the volatility of mitigation cost, thereby slightly increasing the value of the RD&D program (not shown). The parameter ρ is one of the most difficult parameters to estimate, with no apparent basis in the literature, and the selection of 0.2 as a base value is largely arbitrary. However, Figure 13 shows that the mitigation cost at which the decision switches from mostly R&D to mostly LBD is insensitive to the assumption on ρ .

Figure 14 shows the value of the RD&D program 15 years before deployment for different discount rates μ . Since all costs are incurred during the development period and benefit is incurred immediately following the development period, the value of the program decreases with increasing discount rate and is approximately a factor of four lower for a




Figure 14: Value of the government RD&D program 15 years before the carbon policy/CCS deployment time under different discount rates μ (the base assumption is $\mu = 0.02$).

corporate discount rate of 8 % than it is for the base assumption of 2 %. The optimal RD&D investment strategy is insensitive to the discount rate (and to the exogenous uncertainty σ) so these analyses are not shown. A sensitivity analysis for the social cost of carbon was beyond the scope of this analysis since it would require constructing CCS demand curves other than the one derived from Savitz et al. (2007).

7 Discussion and policy implications

Introduction of a climate policy at the U.S. federal level is likely in coming years, and low-carbon energy technology development would both increase the likelihood of such a policy being enacted and reduce the cost to society of complying with it. The long time scales involved, great potential for social benefit, and billion-dollar scale of necessary investments support low-carbon energy technology development as an appropriate role for the federal government.

Although the effectiveness of energy RD&D investment is highly uncertain, decision support models for climate and energy policy usually ignore this uncertainty or account for it only through scenario analysis. In contrast, this work explicitly accounts for the effect of technological uncertainty on the value and optimal strategy for a government RD&D program. The analytic framework captures the value of initial investment in technologies with low expected value but a slim probability of a breakthrough, as characterizes many early-stage energy technologies, as well as the value of flexible future decision making over the course of the program. Although this work neglects important phenomena in the rest of the climate-economy system, its isolation of the effects of technological uncertainty yields insight into optimal RD&D strategy.

A principal finding is that over a very broad range of the parameter space, maximum RD&D investment (assumed to be \$250 million/y in the CCS analysis) is optimal. For the CCS base case and a deployment date 15 years in the future, cost would have to exceed \$380/tCO₂ to discourage initial investment. For initial mitigation costs near \$20/tCO₂, on the extreme low end of cost estimates in the current literature, a point of diminishing returns on RD&D investment is reached at about \$250 million/y, but at higher initial costs, benefit continues to scale almost proportionately with expenditure. The demonstrated substantial benefit to society of investing in energy technologies under a wide range of cost and risk characteristics supports extensive RD&D investment in these technologies on the part of the federal government.

Results are highly sensitive to the expected rate of business-as-usual (BAU) cost reductions in absence of the government program. These reductions could be due to factors such as spillovers from other parts of the economy or input cost decreases. At initially high mitigation costs, the government RD&D program adds substantial value assuming a high rate of BAU cost reduction, which could bring costs down to the point where the added effect of the government program could render the technology cost-effective. In this case, greater investment in LBD is optimal, since the BAU reductions will bring the technology

through the higher-cost, R&D phase.

The results for the two-factor cost model show other circumstances under which investment in R&D, which is assumed to be less effective on average but have a greater chance of extreme outcomes, should be emphasized over investment in promoting LBD. A more immediate deployment date and higher initial cost both promote greater investment in R&D. Under these circumstances, when the technology's prospects appear worse, a breakthrough is needed to render the program successful. Less available program funding also promotes a greater allocation to R&D, since a smaller program would depend more on breakthroughs to achieve substantial progress. As funding level increases, more can be allocated to LBD. This finding is consistent with the observation that demonstration projects tend to be more expensive than basic research programs, and would thus require a relatively large funding commitment before they become an important part of the RD&D portfolio.

An important aspect of the RD&D problem that this analysis did not consider is the possibility of investing in a portfolio of low-carbon energy technologies with different cost and risk characteristics, which may be imperfect substitutes in the energy sector. This analysis motivates further study in this direction by showing that the value of the RD&D program has diminishing returns with the level of funding, suggesting that past a certain program size, investing in two isolated programs could provide greater benefit than investing in a single one twice as large. A principal area of future work is to expand the analysis into an option on RD&D of two or more technologies in order to examine the effect of their interaction on the optimal investment strategy and value of the RD&D program.

This work provides evidence that expanded RD&D investment would be an excellent investment opportunity for the U.S. federal government in anticipation of future climate policy. Private sector RD&D is discouraged by uncertainty on the timing and form of future climate policy and by the high discount rates and the shorter-term nature of investment planning in the private sector. Once a policy is in place, private sector RD&D

could accelerate such that low-carbon energy technologies would reach the same stage of development, with or without a government RD&D program, within a few years. Even accounting for this effect, a government RD&D program that takes place prior to climate policy enactment could yield social benefit amounting to billions of dollars plus positive spillover effects in other countries and sectors of the economy. The best allocation of RD&D funds between riskier, breakthrough-inducing R&D and promotion of more consistently effective LBD depends on factors such as the anticipated rate of cost reduction in absence of the government program and the expected effectiveness and uncertainty of R&D and LBD. This work provides qualitative guidance given expectations on these parameters, and a principal finding that is robust to parameter assumptions is that extensive RD&D investment in a developmental energy technology is indicated over a broad range initial CO₂ mitigation costs.

Appendix A The Bellman equation

The following is adapted from Dixit and Pindyck (1994). To derive (7), we first set up the Bellman equation

$$V(C, t) = \max_I \left\{ -I\Delta t + \frac{1}{1 + \mu\Delta t} \mathbf{E}[V(C', t + \Delta t)|C, I] \right\} \quad (\text{A1})$$

in which C' is the value of C at time $t + \Delta t$. This equation states that the value of the investment opportunity at time t is equal to the discounted expected value of the opportunity at time $t + \Delta t$ less the R&D expenditure incurred over the course of Δt .

Multiplying by $(1 + \mu\Delta t)$ and rearranging, we obtain

$$\mu\Delta t V(C, t) = \max_I \{ -I\Delta t(1 + \mu\Delta t) + \mathbf{E}[\Delta V] \}. \quad (\text{A2})$$

Dividing by Δt and letting it go to zero,

$$\mu V(C, t) = \max_I \left\{ -I + \frac{1}{dt} \mathbf{E}[dV] \right\} \quad (\text{A3})$$

Multiplying both sides by dt yields (7)

Appendix B Solution for the deterministic case

In the deterministic case, cost dynamics are governed by

$$dC = -\lambda I(t)C dt \quad (\text{A4})$$

or equivalently

$$C(t) = C_0 e^{-\lambda Q(t)} \quad (\text{A5})$$

in which $I(t)$ is the instantaneous rate of investment at time t and $Q(t)$ is the total accumulated investment, such that $Q(t) = \int_0^t I(s) ds$. The parameter λ determines the rate at which (the log of) cost decreases with respect to investment $I(t)$.

We wish to maximize the value of the investment opportunity $V(t)$ with respect to the control $I(t)$:

$$V(t) = \max_{I(t)} \left\{ e^{-\mu(T-t)} V(T) - \int_t^T I(s) e^{-\mu(s-t)} ds \right\} \quad (\text{A6})$$

where $V(T) = \max(P - C(T), 0)$.

Although $I(t)$ can take any value between 0 and q , due to the structure of the problem there is a bang-bang effect where $I(t) = 0$ or q for all t (from here forward, I drop the explicit dependence of I on t). To show this, we begin by expanding the expression for dV :

$$dV = \frac{\partial V}{\partial C} dC + \frac{\partial V}{\partial t} dt \quad (\text{A7})$$

Substituting (A4) into the above,

$$dV = -\lambda C I \frac{\partial V}{\partial C} dt + \frac{\partial V}{\partial t} dt \quad (\text{A8})$$

As in the stochastic case, we set up the Bellman equation, which simplifies to

$$\mu V dt = \max_{I(t)} \{ \mathbf{E}[dV] - I dt \} \quad (\text{A9})$$

Substituting (A8) for dV and rearranging,

$$\mu V = \max_{I(t)} \left\{ \frac{\partial V}{\partial t} - \left(\lambda C \frac{\partial V}{\partial C} + 1 \right) I \right\} \quad (\text{A10})$$

To find V , we therefore need to minimize $(\lambda C \frac{\partial V}{\partial C} + 1) I$ with respect to I . Since this expression is linear in I , either $I = 0$ or $I = q$. If $C \frac{\partial V}{\partial C} \leq -\frac{1}{\lambda}$, meaning the rate that the option value increases as a function of cost is sufficiently high, then investment is worthwhile and $I = q$. If the opposite is the case, $I = 0$.

For the region where $I = 0$,

$$V(t) = e^{-\mu(T-t)}(P - C) \quad (\text{A11})$$

and for the region where $I = q$,

$$V(t) = e^{-\mu(T-t)} (P - C e^{-\lambda q(T-t)}) + \frac{q}{\mu} (e^{-\mu(T-t)} - 1) \quad (\text{A12})$$

On the upper boundary C_u (shown in Figure 4), the cost reduction achieved by investing at rate q is such that the investor is indifferent between making the investment and not making it (in either case the option is worthless). Therefore $V(t) = 0$ whether $I = 0$ uniformly or $I = q$ uniformly. Using this observation (and the fact that $C(t) \geq P$ along this boundary, since R&D expenditures can bring C down to P by time T but no

lower), using (A6) we obtain

$$\int_t^T qe^{-\mu(s-t)} ds + e^{-\mu(T-t)} (C_u(t)e^{-\lambda q(T-t)} - P) = 0 \quad (\text{A13})$$

Substituting and simplifying,

$$C_u(t) = \left(P - \frac{q}{\mu} + \frac{q}{\mu} e^{\mu(T-t)}\right) e^{\lambda q(T-t)} \quad (\text{A14})$$

Along the lower boundary C_l , the incremental cost of investing in R&D equals the incremental benefit:

$$qdt = e^{-\mu(T-t)} dV(T) \quad (\text{A15})$$

Since $dV(T) = -dC(T) = -d[C_l e^{-\lambda q(T-t)}]$, substituting and solving for C_l yields

$$C_l(t) = \frac{e^{(\mu+\lambda q)(T-t)}}{\lambda} \quad (\text{A16})$$

Appendix C Details of the numerical solution for the stochastic case: single-factor learning curve

C.1 Substitution to facilitate numerical solution

To solve (9), first the substitution

$$G(x, t) = \left(V(C, t) + \frac{I}{\mu}\right) e^{\mu(T-t)} \quad (\text{A17})$$

is made in (8), in which $x = \ln C$. This substitution yields

$$-\left(\lambda I + \frac{1}{2}\sigma^2\right) \frac{\partial G}{\partial x} + \frac{1}{2}\sigma^2 \frac{\partial^2 G}{\partial x^2} + \frac{\partial G}{\partial t} = 0 \quad (\text{A18})$$




Figure A1: Stencil for the Crank-Nicolson method, the numerical PDE solution method employed. The i index x and the j index t . The solution proceeds backward through time, and the $G_{i,j}$ are found simultaneously for fixed j .

All the terms are now linear, so the discretized PDE equations can be solved with matrix arithmetic. The boundary conditions (10) become

$$G(x, T) = \max\{P - e^x, 0\} \quad (\text{A19a})$$

$$\lim_{x \rightarrow -\infty} G(x, t) = P \quad (\text{A19b})$$

$$\lim_{x \rightarrow \infty} G(x, t) = 0 \quad (\text{A19c})$$

$$\left. \frac{\partial G}{\partial x} \right|_{x \in \{x_l, x_u\}} = -\frac{1}{\lambda} e^{\mu(T-t)} \quad (\text{A19d})$$

and the terminal free boundary points are $x_l(T) = \ln \frac{1}{\lambda}$ and $x_u(T) = \ln P$.

C.2 Numerical PDE solution: Crank-Nicolson method

To solve (A18), I discretized it in both x and t and implemented the Crank-Nicolson method in Matlab. The stencil for this method is shown in Figure A1, in which $i = 1, m$ indexes x and $j = 1, n$ indexes t .

Partial derivative approximations for the Crank-Nicolson method are as follows:

$$\frac{\partial G}{\partial x} \approx \frac{1}{2} \left(\frac{G_{i-1,j+1} - G_{i+1,j+1} + G_{i-1,j} - G_{i+1,j}}{2\Delta x} \right), \quad (\text{A20})$$

$$\frac{\partial^2 G}{\partial x^2} \approx \frac{1}{2} \left(\frac{G_{i-1,j+1} - 2G_{i,j+1} + G_{i+1,j+1} + G_{i-1,j} - 2G_{i,j} + G_{i+1,j}}{(\Delta x)^2} \right), \quad (\text{A21})$$

and

$$\frac{\partial G}{\partial t} \approx \frac{G_{i,j+1} - G_{i,j}}{\Delta t}. \quad (\text{A22})$$

Note that $i - 1$ and $i + 1$ are (somewhat counterintuitively) reversed, since the implementation is easier if $x_{i+1} < x_i$. Substituting these equations into (A18) and rearranging yields

$$\begin{aligned} (b + a)G_{i-1,j} + (-2b - 1)G_{i,j} + (b - a)G_{i+1,j} \\ = (-a - b)G_{i-1,j+1} + (2b - 1)G_{i,j+1} + (-b + a)G_{i+1,j+1} \end{aligned} \quad (\text{A23})$$

in which

$$a = - \left(\lambda I + \frac{1}{2} \sigma^2 \right) \frac{\Delta t}{4 \Delta x} \quad (\text{A24})$$

and

$$b = \sigma^2 \frac{\Delta t}{4(\Delta x)^2} \quad (\text{A25})$$

Proceeding backwards from $j = n$, at each time step j the $G_{i,j+1}$ are known for all i . The $G_{i,j}$ must be solved for simultaneously, which is accomplished with the matrix equation

$$c_j + A_j \times G_{2,m-1,j} = B_{j+1} \times G_{2,m-1,j+1} + d_{j+1} \quad (\text{A26})$$

in which

$$c_j = \begin{pmatrix} (a + b)G_{1,j} \\ 0 \\ \vdots \\ 0 \\ (-a + b)G_{m,j} \end{pmatrix}$$

$$A_j = \begin{pmatrix} -2b-1 & -a+b & 0 & 0 & 0 & 0 & \cdots & 0 \\ a+b & -2b-1 & -a+b & 0 & 0 & 0 & \cdots & 0 \\ 0 & a+b & -2b-1 & -a+b & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & a+b & -2b-1 & -a+b & 0 \\ 0 & 0 & 0 & 0 & \cdots & a+b & -2b-1 & -a+b \\ 0 & 0 & 0 & 0 & \cdots & 0 & a+b & -2b-1 \end{pmatrix}$$

$$B_{j+1} = \begin{pmatrix} 2b-1 & a-b & 0 & 0 & 0 & 0 & \cdots & 0 \\ -a-b & 2b-1 & a-b & 0 & 0 & 0 & \cdots & 0 \\ 0 & -a-b & 2b-1 & a-b & 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & -a-b & 2b-1 & a-b & 0 \\ 0 & 0 & 0 & 0 & \cdots & -a-b & 2b-1 & a-b \\ 0 & 0 & 0 & 0 & \cdots & 0 & -a-b & 2b-1 \end{pmatrix}$$

and

$$d_{j+1} = \begin{pmatrix} (-a-b)G_{1,j+1} \\ 0 \\ \vdots \\ 0 \\ (a-b)G_{m,j+1} \end{pmatrix}$$

The free boundaries enter through the A_j and the B_{j+1} , setting $I = 0$ or q as appropriate in (A24). They are found by iteratively solving fixed-boundary problems. For the initial guess at the free boundaries, I select $x_l = \ln \frac{1}{\lambda}$ and $x_u = \ln P$. For the next iteration, at each time step the two points are found which satisfy (A19d), and are set as the next

boundary guess. The solution generally converges in fewer than 20 iterations, which take 0.15 seconds each with a mesh size of 0.1 for x and t on a standard Macbook.

This method is based on Muthuraman (2008), who proved that the solution converges and is unique.

Appendix D Details of the numerical solution for the stochastic case: two-factor learning curve

To solve (12) numerically, the state space is discretized and an ordinary differential equation in C is solved at each time step. The optimal R&D investment level \hat{I}_R is found iteratively as a function of cost, as are the thresholds C_l and C_u , which separate the zero-investment and the maximum-investment regions of the state space. The iterative solution procedure proceeds as follows, where the subscript ‘old’ indicates the value for the previous iteration.

```

for  $t = T - \Delta t \rightarrow 0$  do
  while  $|C_l - C_{l,old}| \geq \delta_1$  or  $|C_u - C_{u,old}| \geq \delta_1$  do
    while  $\max |I_R - I_{R,old}| \geq \delta_2$  do
      Solve for  $V$ 
      Update  $\hat{I}_R$  and  $\hat{I}_L$ 
    end while
    Update  $C_l$  and  $C_u$ 
  end while
end for

```

For the ‘‘Update \hat{I}_R ’’ step, \hat{I}_R and \hat{I}_L must be found for each point in the state space as a function of C , $\frac{\partial V}{\partial C} = V_C$, and $\frac{\partial^2 V}{\partial C^2} = V_{CC}$. To do this, (12) is rewritten as

$$\mu V = \max_{I_R, I_L} \{aI_R + bI_L + c(I_R I_L)^{\frac{1}{2}} + d\} \quad (\text{A27})$$

where

$$a = \frac{1}{2}\gamma_R^2 C^2 V_{CC} - \lambda_R C V_C - 1 \quad (\text{A28})$$

$$b = \frac{1}{2}\gamma_L^2 C^2 V_{CC} - \lambda_L C V_C - 1 \quad (\text{A29})$$

$$c = \rho\gamma_R\gamma_L C^2 V_{CC} \quad (\text{A30})$$

$$d = \frac{1}{2}\sigma^2 C^2 V_{CC} + V_t \quad (\text{A31})$$

Due to the budget constraint of $I_R + I_L \leq q$, the possible investment region is the union of the triangle defined by $(0, 0)$, $(0, q)$, and $(q, 0)$ and its interior, in which the coordinates represent (I_R, I_L) . To find the maximum, (A27) is differentiated with respect to I_R and I_L and the result set equal to zero, yielding $\frac{c}{2\sqrt{I_R I_L}} = 0$. Since this equation has no solution in the interior of the triangle, the only maxima must exist on its boundary.

Setting $I_R = 0$ and finding the maximum of (A27) with respect to I_L yields $\mu V = \max_{I_L} \{bI_L + d\}$, which implies $I_L = 0$ or $I_L = q$. Setting $I_L = 0$ yields $I_R = 0$ or $I_R = q$. If $0 < I_R < q$, then I_L lies on the diagonal of the triangle ($I_R = q - I_L$). Therefore, the solution of the maximization problem in (A27) (and (12)) is $\hat{I}_R = \hat{I}_L = 0$ or $\hat{I}_L = q - \hat{I}_R$.

Substituting $\hat{I}_L = q - \hat{I}_R$ into (A27) and analyzing the second derivative shows that if $c > 0$ (in (A30)), then investment can be split continuously between I_R and I_L ; if $c < 0$ for the region of full investment there will be a ‘‘bang-bang’’ solution in which either $\hat{I}_R = 0$ and $\hat{I}_L = q$ or $\hat{I}_R = q$ and $\hat{I}_L = 0$. Since all terms in (A30) but ρ are strictly non-negative, if $\rho > 0$ investment will be split continuously and if $\rho \leq 0$ the ‘‘bang-bang’’ solution will be optimal.

If $I_L = q - I_R$, then (A27) reduces to

$$\mu V = \max_{I_R} \{aI_R + b(q - I_R) + c(I_R(q - I_R))^{\frac{1}{2}} + d\} \quad (\text{A32})$$

Taking the derivative of this equation with respect to I_R , setting it equal to zero, and solving for the maximizing value I_R^* yields

$$I_R^* = \frac{q}{2} - \frac{kq}{2\sqrt{4+k^2}} \quad (\text{A33})$$

in which

$$k = \frac{(\gamma_L^2 - \gamma_R^2)V_{CC}C^2 + 2(\lambda_R - \lambda_L)V_C C}{\rho\gamma_L\gamma_R V_{CC}C^2} \quad (\text{A34})$$

Inspection of (A27) shows that if $aI_R^* + b(q - I_R^*) + c(I_R^*(q - I_R^*))^{\frac{1}{2}} > 0$, then $\hat{I}_R = I_R^*$ and $\hat{I}_L = q - I_R^*$. If $aI_R^* + b(q - I_R^*) + c(I_R^*(q - I_R^*))^{\frac{1}{2}} < 0$, then $\hat{I}_L = \hat{I}_R = 0$.

In the solution procedure, for $C_l < C < C_u$, \hat{I}_R is defined by (A33) and $\hat{I}_L = q - I_R^*$. For $C < C_l$ and $C > C_u$, $\hat{I}_R = \hat{I}_L = 0$.

Due to the value-matching condition, at $C = C_l$ and $C = C_u$, $aI_R^* + b(q - I_R^*) + c(I_R^*(q - I_R^*))^{\frac{1}{2}} = 0$ (where I_R^* is defined by (A33)). In the solution procedure, once the I_R^* are found, C_l and C_u are updated by finding the two values of C where $aI_R^* + b(q - I_R^*) + c(I_R^*(q - I_R^*))^{\frac{1}{2}} = 0$.

Appendix E Outline of the solution for two technologies

This Appendix analyzes the case of investment in RD&D of two technologies. In this case, the portfolio effects of distributed investment can be examined. For simplicity in this example, R&D and LBD are aggregated as they are in the single-factor model in the body of the thesis.

The costs of the two technologies are assumed to evolve according to

$$\frac{dC}{C} = -\lambda_C I dt + \gamma_C (I)^{\frac{1}{2}} dz_C + \sigma_C dw \quad (\text{A35})$$

and

$$\frac{dB}{B} = -\lambda_B J dt + \gamma_B (J)^{\frac{1}{2}} dz_B + \sigma_B dw \quad (\text{A36})$$

Proceeding similarly to the other analyses, the Bellman equation simplifies to

$$\mu V = \max_{I,J} \{E[dV] - I dt - J dt\} \quad (\text{A37})$$

in which I (J) is the rate of R&D spending on the technology with cost C (B) with spending subject to the budget constraint $I + J \leq q$.

Expanding $E[dV]$ in this equation yields

$$\mu V = \max_{I,J} \left\{ aI + bJ + c\sqrt{IJ} + d \right\} \quad (\text{A38})$$

in which

$$a = \frac{1}{2}\gamma_C^2 C^2 V_{CC} - \lambda_C V_C C - 1 \quad (\text{A39})$$

$$b = \frac{1}{2}\gamma_B^2 B^2 V_{BB} - \lambda_B V_B B - 1 \quad (\text{A40})$$

$$c = \gamma_B \gamma_C \rho_z BC V_{BC} \quad (\text{A41})$$

$$d = V_t + \frac{1}{2}\sigma_C^2 C^2 V_{CC} + \frac{1}{2}\sigma_B^2 B^2 V_{BB} + \sigma_B \sigma_C BC V_{BC} \rho_w - \alpha_C V_C C - \alpha_B V_B B \quad (\text{A42})$$

in which

$$E[dz_B dz_C] = \rho_z dt \quad (\text{A43})$$

and

$$E[dw_B dw_C] = \rho_w dt \quad (\text{A44})$$

The parameter ρ_z represents the correlation between the effects of RD&D for the two technologies and $\rho_z > 0$ would imply that the technologies are subject to similar breakthroughs and setbacks. The parameter ρ_w represents the correlation between the input costs for the two technologies and $\rho_w > 0$ would imply that the factors of production are similar.

Following an argument similar to that in Appendix 4, it can be shown that either zero or full investment is optimal, with full investment divided between the two technologies only if $c > 0$ and investment in either but not both of the technologies if $c \leq 0$. My conjecture is that $V_{BC} > 0$, so the solution to the problem as posed here would require $\rho > 0$ for split investment to be optimal.

This result would be strongly dependent on the final boundary condition, which values the program at the deployment time. Introducing a boundary condition with a negative second derivative with respect to the costs of the technologies (which imposes a penalty if the RD&D program fails) would introduce risk-aversion, making $V_{BC} < 0$ and rendering split investment optimal if the endogenous changes in the costs of the technologies are anticorrelated ($\rho_z < 0$).

The final boundary condition could depend only on the cost of the cheaper of the two technologies (if the technologies are assumed to be perfect substitutes), or it could be a convex function of the costs of the two technologies, incentivizing development of both.

The solution of investment in either but not both of the technologies could be avoided if the functional form of (A35) and (A36) with respect to I and J were altered. The exponent of $\frac{1}{2}$ in the second term of the right-hand side gives the solution to (A38) its linear structure when $c < 0$, and changing this exponent would produce a different function with respect to I and J (eliminating the exponent, for example, would make V quadratic with respect to I and J).

Bibliography

- Alvarez, L. and Stenbacka, R. (2001). Adoption of uncertain multi-stage technology projects: a real options approach. *Journal of Mathematical Economics*, 35:71–97.
- Athanassoglou, S., Bosetti, V., and De Maere d’Aertrycke, G. (2012). Ambiguous aggregation of expert opinions: The case of optimal R&D investment. *FEEM Working Paper No. 4.2012*.
- Baker, E. and Peng, Y. (2012). The value of better information on technology R&D programs in response to climate change. *Environmental Modeling and Assessment*, 17:107–121.
- Baker, E. and Solak, S. (2011). Climate change and optimal energy technology R&D policy. *European Journal of Operational Research*, 213:442–454.
- Barreto, L. and Kypreos, S. (2004). Endogenizing R&D and market experience in the bottom-up energy-systems ERIS model. *Technovation*, 24:615–629.
- Bednyagin, D. and Gnansounou, E. (2011). Real options valuation of fusion energy R&D programme. *Energy Policy*, 39:116–130.
- Blanford, G. J. and Clark, L. E. (2003). On the optimal allocation of R&D resources for climate change technology development. Technical Report UCRL-TR-200982, Lawrence Livermore National Laboratory.
- Blanford, G. J. and Weyant, J. P. (2005). A global portfolio strategy for climate change technology development. In *International Energy Workshop*.
- Blyth, W., Bradley, R., Bunn, D., Clarke, C., Wilson, T., and Yang, M. (2007). Investment risks under uncertain climate change policy. *Energy Policy*, 35:5766–5773.
- Bøckman, T., Fleten, S.-E., Juliussen, E., Langhammer, H. J., and Revdal, I. (2008). Investment timing and optimal capacity choice for small hydropower projects. *European Journal of Operational Research*, pages 255–267.
- Bosetti, V., Carraro, C., Massetti, E., Sgobbi, A., and Tavoni, M. (2009). Optimal energy investment and R&D strategies to stabilize atmospheric greenhouse gas concentrations. *Resource and Energy Economics*, 31:123–137.
- Bosetti, V. and Drouet, L. (2005). Accounting for uncertainty affecting technical change in an economic-climate model. FEEM Working Paper 147.05.
- Bosetti, V. and Tavoni, M. (2009). Uncertain R&D, backstop technology, and GHGs stabilization. *Energy Economics*, 31:S18–S26.
- Brealey, R. (2002). *Principles of Corporate Finance*. McGraw-Hill/Irwin Series in Finance, Insurance, & Real Estate.

- Chan, G., Anadon, L. D., Chan, M., and Lee, A. (2011). Expert elicitation of cost, performance, and RD&D budgets for coal power with CCS. *Energy Procedia*, 4:2685–2692.
- Cheng, C.-T., Lo, S.-L., and Lin, T. (2011). Applying real options analysis to assess cleaner energy development strategies. *Energy Policy*, 39:5929–5938.
- Davis, G. and Owens, B. (2003). Optimizing the level of renewable electric R&D expenditures using real options analysis. *Energy Policy*, 31:1589–1608.
- Dixit, A. K. and Pindyck, R. S. (1994). *Investment under uncertainty*. Princeton University Press.
- Drury, E., Denholm, P., and Margolis, R. M. (2009). The solar photovoltaics wedge: pathways for growth and potential carbon mitigation in the US. *Environmental Research Letters*, 4.
- Eckhause, J. M., Hughes, D. R., and Gabriel, S. A. (2009). Evaluating real options for mitigating technical risk in public sector R&D acquisitions. *International Journal of Project Management*, 27:365–377.
- Feroli, F., Schoots, K., and Van der Zwaan, B. C. C. (2009). Use and limitations of learning curves for energy technology policy: A component-learning hypothesis. *Energy Policy*, 37:2525–2535.
- Fernandes, B., Cunha, J., and Ferreira, P. (2011). The use of real options approach in energy sector investments. *Renewable and Sustainable Energy Reviews*, 15:4491–4497.
- Finkenrath, M. (2011). Cost and performance of carbon dioxide capture from power generation. International Energy Agency working paper.
- Fleten, S.-E., Haugom, E., and Ullrich, C. (2012). Renewable energy investments under different support schemes: A real options approach. *European Journal of Operational Research*, 220(1):225–237.
- Gerlagh, R., Kverndokk, S., and Rosendahl, K. E. (2009). Optimal timing of climate change policy: Interaction between carbon examples and innovation externalities. *Environmental Resource Economics*, 43:369–390.
- Girotra, K., Terwiesch, C., and Ulrich, K. T. (2007). Valuing R&D projects in a portfolio: Evidence from the pharmaceutical industry. *Management Science*, 53:1452–1466.
- Gritsevskiy, A. and Nakićenović, N. (2000). Modeling uncertainty of induced technological change. *Energy Policy*, 28:907–921.
- Grubb, M., Köhler, J., and Anderson, D. (2002). Induced technical change in energy and environmental modeling: Analytic approaches and policy implications. *Annual Review of Energy and the Environment*, 27:271–308.

- Grübler, A. (2010). The costs of the French nuclear scale-up: A case of negative learning by doing. *Energy Policy*, 38:5174–5188.
- Grübler, A. and Gritsevskiy, A. (2002). *A model of endogenous technological change through uncertain returns on innovation*, chapter 11, pages 280–319. Resources for the Future.
- Hart, D. and Kallas, K. (2010). Alignment and misalignment of technology push and regulatory pull: Federal RD&D support for SO₂ and NO_x emissions control technology for coal-fired power plants, 1970-2000. MIT-IPC-Energy Innovation Working Paper 10-002.
- Huchzermeier, A. and Loch, C. (2001). Project management under risk: Using the real options approach to evaluate flexibility in R&D. *Management Science*, 47(1):85–101.
- Jamasb, T. and Köhler, J. (2007). *Delivering a Low Carbon Electricity System: Technologies, Economics, and Policy*, chapter Learning curves for energy technology: a critical assessment. Cambridge University Press.
- Junginger, M., Lako, P., Lensink, S., van Sark, W., and Weiss, M. (2008). Climate change scientific assessment and policy analysis: Technological learning in the energy sector. Technical Report WAB 500102 017/NWS-E-2008-14/ECN-E-08-034, Netherlands Research Programme on Scientific Assessment and Policy Analysis for Climate Change.
- Klaassen, G., Miketa, A., Larsen, K., and Sundqvist, T. (2005). The impact of R&D on innovation for wind energy in Denmark, Germany, and the United Kingdom. *Ecological Economics*, 54:227–240.
- Kouvaritakis, N., Soria, A., and Isoard, S. (2000). Modelling energy technology dynamics: methodology for adaptive expectations models with learning by doing and learning by searching. *International Journal of Global Energy Issues*, 14(1-4):104–115.
- Krey, V. and Riahi, K. (2009). Risk hedging strategies under energy system and climate policy uncertainties. Technical Report IR-09-028, International Institute for Applied Systems Analysis (IIASA).
- Lohwasser, R. and Madlener, R. (2010). Relating R&D and investment policies to CCS market diffusion through two-factor learning. Institute for Future Energy Consumer Needs and Behavior. FCN Working Paper No. 6/2010.
- Löschel, A. (2002). Technological change in economic models of environmental policy: a survey. *Ecological Economics*, 43:105–126.
- McFarland, J. R. and Herzog, H. J. (2006). Incorporating carbon capture and storage technologies in integrated assessment models. *Energy Economics*, 28:632–652.
- Miketa, A. and Schratzenholzer, L. (2004). Experiments with a methodology to model the role of R&D expenditures in energy technology learning processes; first results. *Energy Policy*, 32(15):1679–1692.

- Muthuraman, K. (2008). A moving boundary approach to American option pricing. *Journal of Economic Dynamics and Control*, 32:3520–3537.
- Neij, L. (2008). Cost development of future technologies for power generation - a study based on experience curves and complementary bottom-up assessments. *Energy Policy*, 36(6):2200–2211.
- Nemet, G. (2009). Interim monitoring of cost dynamics for publicly supported energy technologies. *Energy Policy*, 37:825–835.
- Pan, H. and Köhler, J. (2007). Technological change in energy systems: Learning curves, logistic curves and input-output coefficients. *Ecological Economics*, 63:749–758.
- Patiño-Echeverri, D., Morel, B., Apt, J., and Chen, C. (2007). Should a coal-fired power plant be replaced or retrofitted? *Environmental Science and Technology*, 41:7980–7986.
- Pindyck, R. S. (1993). Investments of uncertain cost. *Journal of Financial Economics*, 34:53–76.
- Pizer, W. A. and Popp, D. (2008). Endogenizing technological change: Matching empirical evidence to modeling needs. *Energy Economics*, 30:2754–2770.
- Popp, D. (2004). ENTICE: endogenous technological change in the DICE model of global warming. *Journal of Environmental Economics and Management*, 48:742–768.
- Rao, S., Keppo, I., and Riahi, K. (2006). Importance of technological change and spillovers in long-term climate policy. *The Energy Journal, Endogenous Technological Change and the Economics of Atmospheric Stabilisation Special Issue*, pages 25–41.
- Reinelt, P. and Keith, D. (2007). Carbon capture retrofits and the cost of regulatory uncertainty. *The Energy Journal*, 28(4):101–128.
- Riahi, K., Barreto, L., Rao, S., and Rubin, E. S. (2004). Towards fossil-based electricity systems with integrated CO₂ capture: Implications of an illustrative long-term technology policy. In *Proceedings of the 7th International Conference on Greenhouse Gas Control Technologies (GHGT-7), Vancouver, Canada*, pages 921–929. Elsevier.
- Rubin, E. and Zhai, H. (2012). The cost of carbon capture and storage for natural gas combined cycle power plants. *Environmental Science and Technology*, 46:3076–3084.
- Rubin, E. S., Yeh, S., Antes, M., Berkenpas, M., and Davison, J. (2007). Use of experience curves to estimate the future cost of power plants with CO₂ capture. *International Journal of Greenhouse Gas Control*, 1:188–197.
- Sagar, A. D. and Holdren, J. P. (2002). Assessing the global energy innovation system: some key issues. *Energy Policy*, 30:465–469.

- Savitz, M., Cohen, L., Corman, J., DeCotis, P., Espino, R., Fri, R., Hanemann, W. M., Harris, W., Krebs, M., Lave, L., Newell, R., Siegel, J., Smith, J., Surlles, T., Sweeney, J., and Telson, M. (2007). Prospective evaluation of applied energy research and development at DOE (phase two). Technical report, National Academy of Sciences.
- Shittu, E. and Baker, E. (2010). Optimal energy R&D portfolio investments in response to a carbon tax. *IEEE Transactions on Engineering Management*, 57(4):547–559.
- Siddiqui, A. and Fleten, S.-E. (2010). How to proceed with competing alternative energy technologies: A real options analysis. *Energy Economics*, 32:817–830.
- Siddiqui, A., Marnay, C., and Wiser, R. (2007). Real options valuation of US federal renewable energy research, development, demonstration, and deployment. *Energy Policy*, 35(4):265–279.
- Van den Broek, M., Hoefnagels, R., Rubin, E., Turkenburg, W., and Faaij, A. (2009). Effects of technological learning on future cost and performance of power plants with CO₂ capture. *Progress in Energy Combustion Science*, 35:457–480.
- Van der Zwaan, B. and Sagar, A. (2006). Technological innovation in the energy sector: R&D, deployment, and learning-by-doing. *Energy Policy*, 34:2601–2608.
- Webster, M., Santen, N., and Parpas, P. (2012). An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. *Computational Management Science*, 9:339–362.
- Yang, M., Blyth, W., Bradley, R., Bunn, D., Clark, C., and Wilson, T. (2008). Evaluating the power investment options with uncertainty in climate policy. *Energy Economics*, 30:1933–1950.
- Yeh, S. and Rubin, E. (2012). A review of uncertainties in technology experience curves. *Energy Economics*, 34:762–771.
- Yeh, S., Rubin, E., Hounshell, D. A., and Taylor, M. R. (2009). Uncertainties in technology experience curves for integrated assessment models. *Environmental Science and Technology*, 43(18):6907–6914.

Conclusion

This thesis has contributed to the analysis of solutions in the electricity sector for transitioning toward low-carbon sources of energy in service of climate change mitigation. Papers 1–3 contain policy-relevant analyses of methods to integrate wind power, and together they indicate that further innovation is needed to achieve drastic cuts in greenhouse gas emissions in the U.S. electricity sector. Paper 4 follows this line of reasoning by developing a tool to inform government support of innovation through spending on RD&D (research, development, and demonstration) of low-carbon energy technologies. This concluding chapter both traces the thread that connects Papers 1–4 and summarizes the main findings, contributions, policy implications, and further questions raised by each chapter individually.

Paper 1: Economics of compressed air energy storage (CAES) to integrate wind power: a case study in ERCOT

At both state and federal levels in the U.S. as well as in the E.U., policy incentives are resulting in the rapid large-scale deployment of relatively developed low-carbon energy technologies. Wind power is first among these technologies due to its relatively low cost, ample resource base, and negligible CO₂ output per kWh generated. The variability and intermittency of its power output, however, hinders its deployment at scale.

Two of the most widely-discussed solutions for mitigating the variability and intermittency of wind power are large-scale energy storage and long-distance interconnection. The former would absorb excess wind power and release it (less efficiency

losses) when demand increases and the wind dies down, and the latter rests on the hypothesis that low wind power production in one area would usually be offset by higher production in a connected area.

Papers 1 and 2 of this thesis examined the use of large-scale energy storage to mitigate wind power variability. Paper 1 modeled the use of compressed air energy storage (CAES) to smooth power output fluctuations of two large wind plants in western Texas, with the wind/CAES system operating in the Electric Reliability Council of Texas (ERCOT) hourly market. The analysis found that under 2007–2009 market conditions, the only circumstances that rendered the wind/CAES system profitable were the presence of large price spikes, which were present to the required extent only in 2008. Varying the location of the CAES plant—near wind to enable a smaller wind-to-load transmission line, or near load to minimize losses when purchasing and storing off-peak power—did not improve the economic performance of the wind/CAES system, nor did simulated alternative market conditions such as a price cap and capacity payment. Accounting for the air quality benefits of avoided peaker operation likewise did not render the system profitable.

The intellectual contribution of this paper is in its synthesis of an engineering model of a CAES system, multiple electricity market designs, real wind and price data, and geological constraints in a multiparameter optimization of a wind/CAES system with profit as the objective. The specificity of the case study supports the validity of the results for the Texas market. Since the analysis was conducted for a high-wind area with geology suited for CAES, the result that CAES is economically unfeasible is likely valid more generally in the U.S.

Since CAES is among the least expensive low-carbon technologies for achieving wind power integration, this paper points to a broader lack of cost-effective solutions in this area. Given the recent discovery in the U.S. of vast unconventional natural gas resources and the subsequent drop in gas prices, quick-ramping gas plants are currently the most cost-effective solution for smoothing wind variability, even under relatively stringent

simulated carbon prices. Efficiency of gas plants has improved significantly in recent years, and compared with other fossil fuel generators gas plants emit relatively little CO₂ (though the methane emitted from the remainder of the gas supply chain may contribute substantially to lifecycle greenhouse gas emissions). If a more complete decarbonization of the electricity sector is to occur, cost-effective solutions with lower greenhouse emissions than natural gas plants will need to be developed and implemented.

Since traditional CAES uses a small amount of natural gas, it may serve as an interim solution rather than a long-term solution. Adiabatic CAES is a technology in early stages of development that stores the heat of compression rather than dissipating it and therefore requires no natural gas to reheat the air in the generation phase. If this technology becomes cost-effective and feasible at scale, it could serve as bulk energy storage to smooth wind fluctuations in a drastically emissions-reduced electricity sector. Future work could analyze the engineering and economic potential of adiabatic CAES to smooth wind power.

As the ERCOT market continues to change due to increasing wind capacity, price structures could shift. Hourly prices are currently determined primarily by daily, weekly, and seasonal load patterns. Higher wind penetration, however, could shift price structures to a regime that follows wind power output more closely. Future work could analyze whether such a development alter price variability in such a way as to render CAES profitable.

Paper 2: The profitability of pumped hydropower storage to integrate wind: A real options approach

Aside from CAES, the only other energy storage technology with high enough power output and storage capacity to smooth large-scale wind power at a reasonable cost is pumped hydropower storage (PHS). Like parts of the U.S., Germany is rapidly expanding its wind power capacity and is in need of cost-effective energy storage to balance wind. Germany has little possibility to expand its domestic pumped storage capacity, but it does

have the opportunity to use the Norwegian hydropower system to balance wind fluctuations. Expanded interconnection with Norway has been shown to be a more cost-effective method of smoothing large amounts of wind capacity than the domestic alternative of a buildout of CAES plants. Tonstad hydropower plant in southern Norway is being considered for a pumped storage retrofit and HVDC interconnection with Germany. If this plan is realized, the plant could initially operate in only the German market but would eventually be integrated with the Norwegian system as well.

Paper 2 of this thesis analyzed the decision from the perspective of a Norwegian hydropower producer to invest in a PHS facility modeled after Tonstad, which would operate in the German market. Since the continuing rapid expansion of wind power in Germany could result in changes in price patterns in the next decade, real options was used to value the investment opportunity under uncertainty in hourly price volatility. Unlike net present value (NPV), real options allows for flexible future decision making and accounts for the effect of uncertainty on the value of the investment opportunity.

An NPV analysis demonstrated that the investment opportunity would be profitable under base-case assumptions on increasing price volatility due to wind power expansion. However, the real options analysis showed that the effect of uncertainty, as well as the expectation of increasing price volatility, promoted waiting about eight years to invest. This result illustrates the value to a Norwegian hydropower producer of waiting for further information and for the expected rise in hourly price volatility before making the investment decision.

The contribution of this paper is its use of real options, an appropriate tool to analyze investment decisions under uncertainty, to inform a topical decision. The analysis synthesized a model of a PHS system with accurate technical detail, a stochastic price model that incorporates the likely effect of wind on German price volatility, and a regression-based option valuation algorithm.

The specific effect of vastly increased wind power capacity on hourly prices in the

German market is poorly understood and is the principal uncertainty in this paper. Like in ERCOT, the price structure in the German market is currently dominated by daily, weekly, and seasonal load fluctuations, but this could be altered drastically at high wind penetrations. More observational data will help resolve this effect, as will the use of more sophisticated models. Further work could incorporate a bottom-up price model of the German system under increased wind capacity rather than the statistical model applied in Paper 2. This development could enable more specific insight on the long-term effects of wind on price volatility, which is a large determinant of the profitability of the pumped hydropower system.

Further work could also take a more integrated approach to evaluating wind power integration in Germany, such as the possibility of reforming feed-in tariffs to encourage operation of wind plants in a way that is more responsive to the conditions in the rest of the grid. Feedback effects in which access to the Norwegian hydropower system promotes wind capacity expansion could also be examined.

The problem of wind integration in Germany is also strongly connected to the future of the country's nuclear program. The analysis in Paper 2 was conducted before the incident at Fukushima and the subsequent decision to discontinue Germany's nuclear program. The abandonment of nuclear power could elevate the importance of wind integration if Germany is to adhere to its greenhouse gas emissions targets, perhaps incentivizing access to the Norwegian hydropower system.

Paper 3: The effect of long-distance interconnection on wind power variability

In addition to large-scale energy storage, long-distance interconnection of wind plants is another proposed method of smoothing wind fluctuations. This method rests on the hypothesis that if wind power production dies down in one region, wind plants in a distant region are still likely to be producing.

Wind power fluctuations operate at different time scales, with higher-frequency fluctuations (those that operate on the order of hours or faster) causing the greatest impediment to wind integration. Previous work showed that interconnecting wind plants within Texas smoothed higher-frequency fluctuations, but that the effect reached diminishing returns after about four wind plants were connected (Katzenstein et al., 2010). Paper 3 develops this work to examine the effect of interconnecting wind plants across longer distances in the U.S. Results showed that interconnecting the aggregate wind power resources of four areas of the mid-western and western U.S. would provide little further benefit above interconnecting wind plants within a single region. An approximate cost analysis showed that the variability reduction achieved through long-distance interconnection could be obtained at substantially lower cost with local gas plants, even when CO₂ and criteria air pollutant externalities were priced.

The contribution of this work is in its application of a range of techniques to quantify the variability of wind power output and the extent to which interconnecting wind plants could reduce variability. The principal analytical tool, frequency domain analysis, is common in atmospheric sciences but has rarely been applied to wind power output. It is nevertheless well-suited to analyzing wind power fluctuations that operate at different time scales.

This work provides the foundation for future frequency-dependent analyses of wind power. Even if long-distance interconnection of wind plants is unlikely to justify the increased transmission expenditure, future work remains on characterizing the conditions that cause diminishing returns in the smoothing effects of interconnection. Such work could entail more extensive use of atmospheric sciences, especially mesoscale meteorology, to quantify the extent of wind power smoothing at different frequencies and inform system planners on the likely smoothing effects of interconnecting wind plants at different spatial scales.

Paper 4: Optimal investment strategy in a clean energy RD&D program under technological uncertainty

Like Paper 1, Paper 3 highlights the need for improved solutions in the U.S. both for wind power integration and the decarbonization of the electricity sector as a whole. The high cost of CAES and significant emissions from the natural gas-fired part of its cycle, as well as the ineffectiveness of long-distance interconnection at smoothing wind power output beyond the effect achieved by interconnection within a region, illustrate the need for innovation in wind power integration as well as in low-carbon electricity production more generally. While Paper 2 presents a case where existing low-carbon energy technology could be a cost-effective means of integrating wind at large scale to achieve deep cuts in CO₂ emissions, its results are contingent on Germany's ability to use the Norwegian hydropower system. They are thus not generalizable to other areas of the world seeking to integrate wind.

The need for further innovation in order to achieve extensive decarbonization of the U.S. electricity sector motivates Paper 4, which analyzes an investment decision in an RD&D program under uncertainty in technological change from the perspective of the U.S. federal government.

Most analyses that include technological change in the energy sector model it as deterministic. This approach implies that investment decisions on technology development and deployment are made based on deterministic, NPV-type analyses, ignoring the effects of uncertainty on the optimal investment strategy. Including uncertainty in the analysis introduces important phenomena observed in real-life RD&D decision making, such as the ability to invest in developmental technologies with negative expected values but small probabilities of breakthroughs, the ability to accelerate or terminate a program as the success or failure of the program becomes apparent, and the possibility that the program will result in increases in cost estimates.

The small fraction of previous work that incorporates uncertainty in the effect of

RD&D programs often model their outcomes as binary, resulting in either success or failure at a single point in time. In reality, such programs can result in a range of outcomes from complete failure to breakthroughs that far exceed expectations, which can occur at any time during the program. The work presented in Paper 4, while it simplifies many aspects of the climate-economy system that other studies represent more fully, models the RD&D investment decision with uncertain outcomes that are continuous in both time and deployment cost. While the model contains a stylized representation of the RD&D process and makes use of parameters that are difficult to estimate, the generality of the model and the broad range of parameter assumptions explored yield robust qualitative insights into optimal RD&D strategy under uncertainty. The intellectual contribution of this work is in its development of a technique to inform RD&D strategy that draws from real options theory and applied math as well as literature on technological change.

A main finding of this work is that over a broad range of parameter values representing the effectiveness of spending and uncertainty associated with R&D and learning-by-doing (LBD, promoted through demonstration projects), extensive investment in the program is optimal with a decade or more to deployment even if cost of the technology is very high. For CCS, a case study examined in Paper 4, substantial RD&D spending is indicated at 15 years until deployment unless current mitigation cost exceeds \$380/tCO₂, well above the range of estimates in the literature. For technologies that might exceed this cost, such as direct air CO₂ capture, cost may be too high to justify an RD&D program; however, greater uncertainty in the outcome of RD&D of direct air capture compared with CCS could encourage investment at higher costs. Further, the CCS analysis assumed that the additional benefit of the government RD&D program would last until only three years after a carbon policy is enacted and CCS is deployed at scale, after which private sector RD&D would overtake any gains the government program achieved. For technologies further from commercialization, a correct valuation of an RD&D program might account for a longer period of time in which benefits accrue, thereby increasing both the value of deploying the

technology and the cost threshold below which initial investment is indicated.

This analysis also highlights the importance of the expected rate of business-as-usual (BAU) technology development in absence of the government program. BAU cost reductions could be due to changes in input costs or to spillovers from technology development in other parts of the economy or the world. If no BAU cost reduction is expected, the added value of the government RD&D program is greatest at relatively low initial costs and the investment decision is weighted heavily towards R&D as opposed to LBD. For higher rates of BAU cost reduction, the government RD&D program adds substantial value at high initial costs, since BAU reductions are likely to bring high costs down to the point where the added effect of the government program could render them cost-effective. In the latter case, greater investment in LBD is optimal, since the BAU reductions will bring the technology through the R&D phase. This effect illustrates the importance of the broader innovation context in valuing a government RD&D program.

This work supports the inclusion of commercially-focused demonstration projects in the overall RD&D portfolio. For technologies that are sufficiently developed (such as CCS, if current cost were near the low end of the range of estimates) the lower breakthrough possibility but higher expected cost reductions resulting from demonstration projects can play an important role in a government RD&D program, even in the presence of large expected BAU cost reductions. More spending on LBD is also indicated in RD&D programs that have a long enough time frame to foster the incremental effects of LBD in addition to betting on the extreme effects of R&D.

A factor promoting greater R&D investment over LBD is less time until the carbon price is enacted and the deployment decision is made: as the deployment decision nears, if the technology still has relatively high cost its success will depend on a breakthrough rather than a slow and consistent reduction. This result depends on the assumption that a backstop technology exists that is guaranteed to mitigate CO₂ emissions at a known cost, and if the modeled technology exceeds this cost, deployment of it is worthless. It is unclear

that this assumption will hold. This motivates future analysis of scenarios in which there is no viable backstop technology or failure to develop the modeled technology has negative consequences, perhaps corresponding to the opportunity cost of investing in another program. One way to address this problem is to model investment in a portfolio of technologies with a very high backstop cost.

Results show diminishing marginal returns with the size of the RD&D program, especially once the developmental technology reaches lower costs, indicating that investment in a second technology could yield greater benefit than equivalent additional investment in the first. This finding further motivates an analysis of RD&D of a portfolio of technologies.

Other areas for future work include more complex modeling of the RD&D investment decision in the context of future developments in the energy sector. For the case of CCS, for example, such developments could include the cost and political feasibility of large-scale nuclear power and continued political support of the fossil fuel industry, both of which could alter the cost at which CCS technology becomes deployable at large scale.

Future work could also examine the effects of path-dependency on technological development trajectories. While the optimal RD&D strategy as developed in Paper 4 depends only on current cost and time until deployment, other work in this area models the effect of accumulated knowledge stock (which can depreciate over time) and cumulative installed capacity.

For models such as the one developed in Paper 4 to be capable of informing RD&D policy more specifically, better estimates on the uncertainty in the effectiveness of R&D spending should be generated. Many studies that analyze learning curves for energy technologies focus on those that have been relatively successful, possibly underestimating the uncertainty in the programs.

As the representation of technological change and its interaction with the rest of the energy sector and political environment becomes more complex, an analytical solution will

no longer be feasible, necessitating the use of numerical, simulation-based methods. While the latter type of model would sacrifice generality, it could yield more specific insight into RD&D strategy in the full economic and political context of the technology's development. Approximate dynamic programming (ADP) would be an appropriate technique to pursue this research and its application to energy RD&D has been developed by Webster et al. (2012) and Santen (2012). In contrast, the method developed in Paper 4 of this thesis sacrifices complexity but captures overall dynamics of the RD&D problem, producing results that are more generalizable.

Taken as a whole, this thesis has broader implications and raises further questions about wind power integration and innovation in low-carbon energy technologies. Papers 1 and 3 raise the question of whether a CAES plant combined with an interconnected group of wind plants would be a cost-effective source of firm capacity, since the smoothing effect of interconnection could enable the use of a smaller CAES plant. The wind power data used in Paper 1 was from two of the largest wind plants in central Texas, both of which show substantial high-frequency smoothing due to the geographical separation of turbines. Paper 3 and its antecedent, Katzenstein et al. (2010), suggest that connection of the wind plants studied in Paper 1 with distant wind plants is unlikely to result in substantial additional smoothing, such that the size of the CAES plant necessary to provide firm capacity and the general results of the economic analysis are likely to remain unchanged. The additional cost of transmission to the distant wind plants further reduces the likelihood that such a system would be cost-effective, further demonstrating the limitations of current methods of wind integration.

Papers 1 and 2 both highlight the lack of cost-effective large-scale energy storage with flexible siting requirements. If a low-carbon storage technology such as adiabatic CAES becomes cost-competitive, it could change the economic picture of wind integration. Since the future cost and performance of adiabatic CAES is highly uncertain, decisions on its development and deployment could be informed by an SDP-based analysis such as that

presented in Paper 4. The result from Paper 4 that RD&D investment is optimal for a broad range of parameter estimates for CCS suggests that if a similar analysis were run for adiabatic CAES, results may indicate that initial RD&D investment in that technology is optimal as well. More generally, the advent of other low-carbon energy technologies capable of producing firm power could discourage investment in current storage technologies by diminishing the importance of wind integration, reducing the ability of storage plants to capture intertemporal arbitrage revenue, and causing investors to wait to observe innovation trends before investing in energy storage for wind integration.

Factors such as these raise the question of tradeoffs between promoting decarbonization of the electricity sector through RD&D versus deployment of existing technologies. While this thesis suggests that investment in the former is optimal over a broader range of circumstances, it does not examine incentives for the latter such as avoiding lock-in of carbon-intensive energy technologies and possible time preference in emissions trajectories.

This thesis has analyzed obstacles to large-scale wind power integration, generating evidence of the need for innovation in the electricity sector in order to drastically reduce CO₂ emissions. It has also developed a simple framework for informing the strategy and valuation of government RD&D investment in low-carbon energy technologies. In doing so, this thesis has contributed to developing solutions for the decarbonization of the electricity sector in service of climate change mitigation.

Bibliography

- Katzenstein, W., Fertig, E., and Apt, J. (2010). The variability of interconnected wind plants. *Energy Policy*, 38(8):4400–4410.
- Santen, N. (2012). *Technology Investment Decisions under Uncertainty: A New Modeling Framework for the Electric Power Sector*. PhD thesis, Massachusetts Institute of Technology.
- Webster, M., Santen, N., and Parpas, P. (2012). An approximate dynamic programming framework for modeling global climate policy under decision-dependent uncertainty. *Computational Management Science*, 9:339–362.