Decisions and Uncertainties in the US Energy System: Electrofuels and Other Applications

Submitted in partial fulfillment of the requirements for

the degree of

Doctor of Philosophy

in

Engineering and Public Policy

Evan D. Sherwin
B.A. Physics, B.A. Applied Mathematics, University of California, Berkeley
M.S. Machine Learning, Carnegie Mellon University

Carnegie Mellon University
Pittsburgh, PA, 15213

May, 2019

Committee Members

Inês M. L. Azevedo (Chair)
Professor, Department of Engineering and Public Policy, Carnegie Mellon University

Max Henrion
CEO, Lumina Decision Systems
Adjunct Professor, Department of Engineering and Public Policy, Carnegie Mellon University

Kenneth Gillingham
Associate Professor, School of Forestry & Environmental Studies, Yale University

J. Zico Kolter
Assistant Professor, Computer Science Department, Carnegie Mellon University

M. Granger Morgan
Professor, Department of Engineering and Public Policy, Carnegie Mellon University

Russell M. Meyer
Senior Economist, Oracle

For my grandfather, Hans Felix Weinberger (Carnegie Tech M.S. 1948, Sc.D. 1950), who devoted his career to applying mathematics to address great societal problems of our time.

Acknowledgments

First and foremost, I would like to acknowledge my primary advisor, Inês M. L. Azevedo, for her firm but gentle guidance throughout the PhD process. From the very beginning, she was training me to become an independent researcher capable of defining and executing my own projects, gradually taking away the training wheels while constantly holding me to a high standard. To my delight, this culminated in Inês' willingness to let me strike out essentially on my own to pursue the most interesting project I could think of (see Chapter 4). Inês has always been enthusiastically supportive, providing introductions to leading figures in academia, industry, and government and constantly encouraging me to share my findings at conferences. In all, it has been a privilege to spend the past five years working with my mentor and friend, Inês.

During my first in-person job interview out of college, I was surprised when Max Henrion, CEO of Lumina Decision Systems, asked "You're planning to go to graduate school, right?" From that point on, Max had already become, in a sense, my PhD advisor. Max also gave the 21-year-old Evan a remarkable amount of freedom to shape projects, simultaneously sharing his wealth of knowledge in the fields of decision analysis, energy systems, and, of course, modeling in Analytica. I still marvel at how lucky I was to have Max continue on as my co-advisor during my PhD.

M. Granger Morgan was a something of a mythical figure during my time at Lumina. Judging from conversations with Max, it seemed like Granger had a hand in at least 25% of the interesting research happening in the energy and climate decision-making space. In fact, without Granger's decision to provide startup funding for my first project at Lumina, my career could have easily gone in a very different direction. I would like to acknowledge and thank Granger

both for numerous discussions that helped shape the course of my dissertation and for leading by example in so many ways.

I was thrilled when J. Zico Kolter agreed to take me on as a secondary MS student in Machine Learning. More than anyone else at Carnegie Mellon University, Zico has taught me to identify prudent uses of machine learning while avoiding the magical thinking that often accompanies this rapidly developing field.

Once it became clear that my thesis would involve a substantial econometric component, Inês tasked me with "finding and economist" to serve on my committee. I was delighted when Ken Gillingham agreed to take the job. Discussions with Ken shaped the framing of Chapter 3 and helped it overcome the traditional bounds of benefit-cost analysis, which is of questionable value for equity-based programs.

Russell M. Meyer was a natural addition to my committee. His hands-on experience with the data set used in Chapter 3 and his expertise in energy utility program evaluation made the process of getting useful insights out of inherently messy real-world data much easier. Conversations with Russ also convinced me to simplify Chapter 3 by dropping a planned analysis of a separate online survey that we conducted for a distinct but related purpose. I think the analysis is much sleeker as a result.

I would like to thank my committee, above, for their valuable feedback along the way. I am particularly grateful that you did not hesitate to point out the limitations of some of the items in

my thesis proposal. Most of all, I would like to thank you for being flexible when data availability issues forced me to completely rethink the final chapter of my thesis.

In addition, I would like to acknowledge W. Michael Griffin for sharing his wealth of expertise in, among other things, techno-economic analysis and the fossil energy industry.

I gratefully acknowledge David W. Keith, Geoffrey Holmes, and Carbon Engineering for extensive discussions that fundamentally shaped the structure of Chapter 4. I also gratefully acknowledge Nathan L. Lewis for pushing me to justify all of my assumptions, including through a 90-minute live stress test of my model.

For Chapter 2, I acknowledge and thank Lynn H. Kaack for sharing the substantial task of data collection and harmonization, as well as Jay Apt, M. Granger Morgan, Alex Davis, W. Michael Griffin, Ed Rubin, Mitchell Small, Gabrielle Wong-Parodi, Daniel Armanios, and Stephen Feinberg for their thoughtful feedback and advice. The final figures in the supplementary information benefited from editorial changes and suggestions from Steven J. Davis.

For Chapter 3, I acknowledge Pedro Ferrerira, Akshaya Jha, Edson Severnini, Severin Borenstein, Alex Davis, Carol Edwards, Luke Lavin, and Peter Kuhn for their thoughtful feedback and advice. In particular, I acknowledge Pacific Gas & Electric Company, including employees Brian Smith, J.P. Dolphin, Catherine Izard, Susan F. Norris, Mary O'Drain, Thuong-Tina Nguyen, Hali Burbige, Perla Barrientos, and Brett Searle, for sharing data with us and for continuing to answer questions and provide additional assistance throughout the process.

For Chapter 4, I acknowledge Inês M. L. Azevedo, David W. Keith and Geoffrey Holmes of Carbon Engineering, Jennifer L. Wilcox, Sean T. McCoy, Ken Caldeira, Klaus S. Lackner, Hadi Dowlatabadi, Selma Brynolf, Mahdi Fasihi, Noah Deich and Matt Lucas of Carbon 180, Adam Brandt, Gabrielle Wong-Parodi, Jeremy F. Keen, Liza Reed, Rebecca E. Ciez, Michael Whiston, Fan Tong, Peter Reinhardt, Ahmed Abdulla, Sarah Deutz, Verena Beckert, and Lonnie Chrisman for their thoughtful feedback and advice.

In addition, I owe a great debt of gratitude to Engineering and Public Policy (EPP) server system administrator Rob Smith and Lumina's Lonnie Chrisman for fixing server-related software issues that could have prevented me from finishing my dissertation in time.

Throughout the course of my PhD, I benefitted greatly from discussions with many fellow students and other colleagues. The names are too many to list. I would, however, like to acknowledge fellow members of the CMU Energy Club reading group, Jeremy F. Keen, Brian Sergi, Luke Lavin, Sinnott Murphy, Erin N. Mayfield, Peter Tschofen, Nichole Hanus, and Priya L. Donti. These weekly discussions and debates have shaped me as a researcher in ways that are difficult to quantify.

EPP would not be half the place it is without its incredible staff, who can make even a routine reimbursement request a delightful experience. It has been an absolute pleasure working with EPP's fantastic staff.

I would also like to acknowledge my friends outside of CMU. These past five years would not have been nearly as fun without board games with Pittsburgh's unofficial Goshen College alumni community, music-filled dinners with Manar and Simon Weiss, dinnertime conversations with the servers at Noodlehead, talking about life and the history of Pittsburgh with Darrell at the Hideaway, and, of course, classes and performances with the Steel City Improv Theater community.

Special thanks to my yoga instructor, Christina Sible, for teaching me numerous techniques that have made me a healthier and happier person.

Warm thanks to Chandralekha Singh, Jeremy Levy, and Akash and Ishan Levy-Singh for providing me with a home away from home and a family in Pittsburgh.

I have the unique privilege of being able to talk life and academia with my brother, Stuart A. L. Sherwin, and dear childhood friends Sahar H. El Abbadi and Adam J. B. Rothman. Our regular conversations have been an important part of these past five years.

Finally, I am profoundly grateful to Catherine J. Weinberger and Mark S. Sherwin for their unwavering support and their depth of insight. I was incredibly lucky to have parents with PhDs in economics and physics, respectively, but I think this is only part of the reason they were able to see things in me long before I saw them in myself.

This material is based upon work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-1252522. This work was funded in part by the Center for Climate and Energy Decision Making (SES-0949710 and SES-1463492), through a cooperative agreement between the National Science Foundation and Carnegie Mellon University. We acknowledge and thank Pacific Gas and Electric Company, and the Wharton Customer Analytics Initiative for providing us with data. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation, Pacific Gas and Electric Company, or the Wharton Customer Analytics Initiative.

Abstract

Achieving a global warming limit of 2°C is likely only possible if humanity ceases to emit greenhouse gases (GHG) well before the end of this century. This can only be accomplished through, among other things, a massive transformation of a deeply unpredictable global energy system on which billions of people depend. This thesis aims to illustrate three methodologically distinct approaches that could be integrated into a framework for energy decision-making capable of guiding thoughtful and equitable planning for robust reductions in GHG emissions in the face of deep, largely irreducible uncertainty. Although the primary object of study is the US energy system, all three analyses aim to draw generalizable conclusions that are useful in other contexts.

Chapter 2 attempts to characterize the predictability and volatility of the US energy system by analyzing errors in past US government projections and historical fluctuations in the price, production, and consumption of key energy quantities. This work finds that the period from 2005-2014 contained a disproportionate number of the largest projection errors and inter-year fluctuations in almost all of the 17 quantities examined. This indicates that the US energy system itself was more volatile and harder to predict in this period than in previous decades.

Chapter 3 uses observational residential electricity consumption data to estimate the effect of a low-income electric subsidy on electricity demand, and the externality costs associated with increased electricity generation and higher peak demand. This work finds that the externality costs are on the order of 11% of total subsidy disbursements, with no significant change in this number if intra-day estimates are used instead of time-invariant estimates. Decarbonization of the electric power system will likely eliminate most emissions from power plants, leaving only

capacity costs of roughly 5% of subsidy disbursements. Thus, policy makers considering low-income subsidies as a means of ensuring that low-income households do not disproportionately bear the burden of an energy transition can use such estimates of price responsiveness to estimate any adjustments in peak capacity requirements that may result from increased demand.

Chapter 4 uses an optimization-based techno-economic model to characterize the decision space for deep decarbonization of liquid-dependent sectors such as aviation and long-distance road transportation. With today's technology electrofuels, synthetic hydrocarbons produced using CO_2 captured from the atmosphere and hydrogen from electrolysis of water, are likely a more expensive mitigation strategy than continuing to use petroleum-based fuels and offsetting the resulting emissions with direct air capture (DAC) of CO_2 with sequestration (DACS). However, if DAC and electrolyzer manufacturers are able to meet near-term cost targets, electrofuels may be competitive with DACS if the cost of petroleum fuels rises substantially or if sequestration costs are higher than anticipated. Several decades into the future, electrofuel costs may fall as low as $2.70 per gallon of gasoline equivalent, potentially achieving cost parity with petroleum fuels. Electrofuel cost is most sensitive to the capital cost the DAC, electrolyzer, and renewable electricity systems, confirming their importance as priorities for research, development, and deployment (RD&D). However, without the operational flexibility afforded by storage or supplementary natural gas or grid electricity interconnections, costs could rise by more than 80%. This points to some less intuitive RD&D priorities, such as metallic phase change materials capable of storing heat above 900°C and low-cost, seasonal CO_2 storage.

As a whole, this work aims to characterize the depth the uncertainties posed by the task of energy transition while synthesizing insights from analysis of historical data and modeling based on engineering knowledge and expert judgment to gain policy-relevant insights into pathways toward deep decarbonization of the energy system. I hope this represents a small step toward a decision-making paradigm capable of addressing the deep uncertainties we face while using the wealth of data and insight at our disposal to chart a thoughtful course ahead.

Table of Contents

List of Tables

List of Figures and Illustrations

List of Abbreviations

AMI	Advanced metering infrastructure
AC	Alternating current
AEO	Annual Energy Outlook
BEA	US Bureau of Economic Analysis
BCF	Billion cubic feet of natural gas
$CaCO_3$	Calcium carbonate
CaO	Calcium oxide
$Ca(OH)_2$	Calcium hydroxide
CARE	California Alternate Rates for Energy
CDF	Cumulative distribution function
CO	Carbon monoxide
CO_2	Carbon dioxide
DAC	Direct air capture (of carbon dioxide)
DACS	Direct air capture with sequestration
DC	Direct current
EIA	Energy Information Administration
FEAA	Federal Energy Administration Amendments
FPL	Federal poverty level
GGE	Gallon of gasoline equivalent
GDP	Gross domestic product
GHG	Greenhouse gas
Gt	Gigaton
H_2O	Water
HS	Historical simulation
IFFS	Intermediate Future Forecasting System
K_2CO_3	Potassium carbonate
KOH	Potassium hydroxide
kg	Kilogram
kW	Kilowatt
kWh(e)	Kilowatt-hours (electric)
kWh(H_2)	Kilowatt-hours (Hydrogen)
kWh(fuel)	Kilowatt-hours (of liquid fuel)
kWh(th)	Kilowatt-hours (thermal)
LCOE	Levelized cost of electricity
LEED	Leadership in Energy and Environmental Design
LIRA	Low-Income Ratepayer Assistance
LNG	Liquefied natural gas
mcf	Thousand cubic feet (of natural gas)
Mbbl	Million barrels of oil
MMBTU	Million British thermal units
Mt	Megaton
MW	Megawatt
MWh	Megawatt-hour
NEMS	National Energy Modeling System
NPV	Net present value

O&M	Operations and maintenance
PEM	Proton exchange membrane (electrolyzer)
PG&E	Pacific Gas and Electric Company
PoT	Peaks-over-threshold
PPM	Parts per million
PtXOpt	Power-to-X Optimization Tool
PV	Photovoltaic
RCT	Randomized controlled trial
RD&D	Research, development, and deployment
RWGS	Reverse water-gas shift
SB	Senate bill
SI	Supporting information
$Si_{49}Mg_{30}Ca_{21}$	A metallic phase change material
$t(CO_2)$	Ton of carbon dioxide
US	United States
VaR	Value-at-risk
WACC	Weighted average cost of capital

1. Chapter 1. Introduction and background

Achieving a global warming limit of 2°C is likely only possible if the global economy ceases to emit greenhouse gases (GHG) well before the end of this century [1]. This can only be achieved through a massive transformation of the global energy system on which billions of people depend and which accounts for the bulk of global GHG emissions [2].

Efforts to transition the global energy system to a low-carbon or carbon-negative future within the span of several decades poses great challenges. The future of the energy system has been notoriously difficult to predict more than a few years out [3,4]. This uncertainty is compounded by policy and technology uncertainty, highlighted by the past decade's rapid, largely policy-driven advances in low-carbon energy technologies such as solar photovoltaics and lithium-ion batteries [5–7].

The United States (US) energy system accounted for 15% of global fossil GHG emissions in 2014 [8]. Thus, decarbonizing the US is a prerequisite for achieving a warming limit of 2°C or less. Thankfully, the US energy system is remarkably well-documented with a wealth of publicly available current and historical data, thanks in large part to the work of the Energy Information Administration (EIA). Thus, the US energy system is an important and convenient object of study, and many lessons learned will generalize, at least in part, to other contexts around the world.

Much of the US energy system consists of long-lived capital assets such as power plants, electricity transmission lines, trucks, automobiles, pipelines, and airplanes, with lifetimes of twenty to fifty years or more. As a result, the energy decisions made now under deep uncertainty will have major repercussions for decades. In addition, such investments may, at least

temporarily, raise the cost of energy, placing a disproportionate burden on low-income households.

There is no clear blueprint for how to smoothly transition to a deeply decarbonized energy system. We have many technological and policy tools at our disposal, but the most cost-effective methods for doing so, particularly for difficult-to-decarbonize sectors such as aviation and long-distance road transportation, will likely require technological advances that are necessarily difficult to predict.

In this context, energy decision-makers, including policy makers, energy producers, and energy transmission and distribution utilities could benefit greatly from decision-making frameworks capable of robustly navigating deep uncertainty.

This dissertation consists of three distinct but related studies that work in different ways toward this the development of such an integrated approach to decision-making for deep decarbonization of the energy system.

Chapter 2, *Estimation of the year-on-year volatility and the unpredictability of the United States energy system*, investigates the historical capacity of US government projections to predict the long-term future of the energy system. This work finds that despite the best efforts of analysts at the EIA to project the likely development of the US energy system under certain constraining assumptions surrounding technological and policy change, i.e. to predict the future under a set of limited constraints, the most extreme errors in their Annual Energy Outlook (AEO) projections have actually increased in frequency in recent years. My coauthors and I believe that this is not because EIA's modeling has somehow gotten worse, but because for various reasons the US energy system itself became more volatile and inherently harder to predict starting in the mid-2000s.

This result motivates an alternative approach to long-term energy investment and policy decision-making that attempts to reduce decision-makers' reliance on the elusive ability to accurately predict the empirical details of the energy system, such as the price and production levels of key primary energy quantities, beyond a few years into the future.

Chapter 3, *Characterizing the relationship between low-income electric subsidies and the intra-day timing of electricity consumption*, looks at feedbacks between equity-based energy policies and energy infrastructure decision-making. An energy transition will likely require substantial expenditures which may increase system costs, at least in the short-run. Electric subsidies are one way to ensure these costs are not disproportionately borne by low-income households, for whom energy expenditures can represent a substantial fraction of total income. However, reductions in price are likely to encourage an increase in consumption, the magnitude of which may vary throughout the day. Such an increase could result in higher emissions of GHG and criteria pollutants as well as a rise in costly peak consumption. An integrated approach to energy transition could benefit from incorporating these indirect effects of equity-based programs into broader decarbonization and resource adequacy planning processes.

Chapter 4, *Electrofuel synthesis from variable renewable electricity: An optimization-based techno-economic analysis*, introduces a modeling framework to guide investment in research, development, and deployment (RD&D) of new energy and climate change mitigation technologies. Direct air capture (DAC) of carbon dioxide (CO_2) could enable decarbonization pathways that allow continued use of long-lived energy assets in two ways. First, DAC facilities can offset emissions from industrial processes and other difficult-to-decarbonize economic sectors. Second, DAC can supply CO_2 as a feedstock for production of electrofuels, such as

synthetic jet fuel or diesel, which can then function as net a carbon-neutral fuel in existing infrastructure for aviation, heavy trucking, and other applications.

The long-term cost of emerging technologies such as DAC and electrofuel production facilities are inherently uncertain. Today's cost estimates are necessarily constrained to engineering estimates and the cost of a small number of commercially deployed systems. If manufacturers can achieve targeted cost reductions for DAC and electrofuels within the next few decades, this could substantially reduce the cost of deep decarbonization by, among other things, forestalling a need to rapidly turn over costly energy-related infrastructure.

The techno-economic model used in Chapter 4 was conceived as a tool for characterizing the sensitivity of the cost of low-carbon electrofuel production to potential advances in component technologies. The model simulates the operation of a renewably-powered electrofuel production facility in enough engineering detail to identify key sensitivities while remaining flexible enough to easily explore alternative designs of components with very different engineering properties, e.g. high-temperature and low-temperature DAC systems. Following the mantra of co-advisor and mentor Max Henrion, I aimed to make the model as simple as possible and no simpler.

I have dubbed the model the Power-to-X Optimization Tool (PtXOpt) because it can simulate production of electrofuels and many other products from variable electricity. This model aims to guide current RD&D priorities both for DAC and electrofuels as a whole and for specific component technologies. Some enabling technologies, such as hundred-megawatt-scale high-temperature electric heating and low-cost seasonal CO_2 storage, do not currently exist commercially, but could likely be engineered with relative ease once there is a clear need.

In conclusion, this work aims to motivate a framework for energy decision-making that is capable of guiding thoughtful and equitable planning for robust reductions in GHG emissions in the face of deep uncertainty.

2. Chapter 2. Estimation of the year-on-year volatility and the unpredictability of the United States energy system*

2.1. Abstract

Long-term projections of energy consumption, supply, and prices heavily influence decisions regarding long-lived energy infrastructure. Predicting the evolution of these quantities over multiple years to decades is a difficult task. Here we estimate year-on-year volatility and unpredictability over multi-decade time frames for many quantities in the US energy system using historical projections. We determine the distribution over time of the most extreme errors (unpredictability) from 1985-2014, and the largest year-over-year changes (volatility) in the quantities themselves from 1949-2014. Our results show that both volatility and unpredictability have increased in the past decade, compared to prior decades. These findings may be useful for energy decision-makers to consider as they invest in and regulate long-lived energy infrastructure in a deeply uncertain world.

2.2. Introduction and background

The United States (US) energy system consists of an enormous interconnected network of long-lived infrastructure, which accounts for a large fraction of national greenhouse gas and air pollution emissions, as well as substantial expenditures. Oil and gas extraction alone contributed $255 billion to GDP in 2015, while transportation of goods and passengers contributed $981

* This chapter is based on the following published work, reproduced with permission from:

Sherwin, E. D., Henrion, M. & Azevedo, I. M. L. Estimation of the year-on-year volatility and the unpredictability of the United States energy system. *Nat. Energy* **3**, 341–346 (2018). https://www.nature.com/articles/s41560-018-0121-4

billion [9]. Present investment decisions relating to energy infrastructure will influence the cost, and environmental and health impacts of the US energy system for decades. Understanding how a national energy system is likely to evolve is a difficult task [3], but critically important for informing long-term energy investment decisions. Understanding historical changes in the projected and actual values of key energy quantities can help decision-makers create robust strategies for a deeply uncertain future.

One way to assess the accuracy of past energy forecasts and projections is to perform retrospective analysis. Early work in this field largely began in the 1980s. One approach, led by Huss (1985) [3-5] and Nelson and Peck (1985) [10–13], focuses on comparing historical errors from different sets of projections, in this case electricity demand projections from the 1970s, primarily for the purpose of model selection. Taking another approach, Landsberg (1985) [14] describes and attempts to explain historical errors from a set of long-term, national US energy projections using anecdotes.

Further work in the 1990s and 2000s seeks to explain the historical causes of large projection errors. Huntington (1994) [15] attempted to explain large oil price projection errors in the 1980s. Sohn (2007) [16] computes retrospective errors from the World Input-Output model, created in the 1970s, and discusses reasons for these errors, in largely qualitative terms. Every year since 1996, the Energy Information Administration (EIA) has released a retrospective report, detailing its historical projections for 19-21 key quantities from the Annual Energy Outlook (AEO) [17]. These retrospective reports discuss the largest historical projection errors, proposing explanations for these errors, and guidelines for interpreting EIA projections in light of past errors. In addition, Shlyakhter et al. (1994) [18] and Kaack et al. (2017) [19] attempt to characterize the distribution of projection errors from the AEO.

The early 2000s saw a number of thoughtful review articles comparing numerous sets of past projections [3,20,21]. These articles attempt to convey inherent unpredictability in the energy system and the inadequacy of point projections, with the aim of instilling humility in energy modelers and those who use projections. These articles tend to discuss ways in which projections are useful despite the near-inevitability of large errors. In doing so, they attempt to inform future projection creation processes. Building on these results, Koomey et al. (2003) [4] issued a plea for retrospective analysis of historical energy projections to further inform future projection creation and decision-making practices.

Since then, numerous analyses have sought to assess the historical accuracy of projections, particularly the Annual Energy Outlook and World Energy Outlook. The majority of these analyses seek, in some way or another, to determine historical bias, generally on the basis of mean percent error, or changes in error magnitude over time, using mean absolute percent error and related metrics [22–29].

The existence of retrospective analyses raises the question of the extent to which insights into past errors can help predict future errors. In short, will the future be as difficult to predict as the past was? We note that the AEO's own low and high oil price scenarios began to widen substantially in AEO 2006, suggesting higher uncertainty in at least that quantity (see the supplementary information (SI), Sections 6.1 and 0, and Supplementary Figures **Figure 17** and **Figure 18**). AEO projections are ideal for retrospective analysis because they have been produced every year since 1982 by a stable government organization, using consistent methods (see the SI, Sections 6.3 and 6.4), although the process is not stationary in a strict statistical sense (see the SI, Section 6.5). The stated goal of the AEO is not to forecast the future, but to project the likely development of the US energy system under the policies in place at the time of the

study, and assuming there are no major technological breakthroughs [30]. Still, if the AEO is to guide decision-making, we believe it is important to characterize its historical prediction accuracy.

We attempt to understand whether today's energy system has in fact been harder to predict than in the past by quantifying *unpredictability (the frequency of extreme errors)* through a retrospective analysis of US energy projections from the AEO reference case. We also investigate whether *year-on-year volatility*, the frequency of the largest year-on-year changes in key energy quantities, has changed over time. We find that both unpredictability and volatility have been larger in the past decade than in the prior two decades.

2.3. Methods and data

2.3.1. Experimental design

This study aims to identify historical periods characterized by large fluctuations, *year-on-year volatility*, and *extreme errors*, unpredictability, for key US energy quantities. We use publicly available historical values and projections for seventeen US energy quantities, described below and in the Supplementary Methods, to compute these fluctuations and projection errors. We use several nonparametric methods to compare the prevalence of extreme changes and extreme projections errors for these quantities by decade.

2.3.2. Data

All projection data and observed historical values used in the extreme error analysis come from either the Annual Energy Outlook (AEO) retrospective reports, or from the individual AEO reports themselves. The single exception to this is GDP, which is derived from a combination of AEO projections of GDP growth and US Bureau of Economic Analysis (BEA) values of historical US GDP, described in further detail in Supplementary Methods. For a graphical

representation of these projected and historical values for all quantities examined, see Supplementary Figures **Figure 27**, **Figure 28**, and **Figure 29**.

For our analysis of extreme year-over-year changes, year-on-year volatility, in observed historical values of energy quantities since 1949, we draw data from the EIA's Monthly Energy Review, November 2015 when available [31]. The exceptions are US GDP and inflation, which we draw from the US BEA [32,33]. All prices in the volatility analysis are in nominal dollars.

Data collection and integration are described further in Supplementary Methods.

2.3.3. Statistical Information

We define unpredictability as the frequency of *extreme errors* for one or several projection years (the projection year would be 2000 for a projection produced in 1990, projecting values for the year 2000).

We define an extreme error relative to its *percent error distribution*, the distribution of all percent projection errors for that quantity. We use percent projection error rather than the simple difference between the projected and actual value because we are interested the magnitude of the error relative to the observed historical value. We define an extreme error as being outside a specified percentile of the error distribution. In the baseline analysis, we designate as extreme errors all percent projection errors above the 97.5th percentile, or below the 2.5th percentile. Thus, roughly 5% of all projection values for each quantity are designated as "extreme errors" (small sample size effects can increase or decrease this rate by up to ±0.6%.). We obtain separate *extreme error thresholds* from the percent error distributions for each *projection interval*, short-term (1 to 5 year), medium-term (6 to 10 year), and long-term (11 to 21 year) projections. In the SI, Section 6.8, we test the effect of alternative definitions of extreme error on our results.

In **Figure 1** we demonstrate this method of computing extreme error thresholds, using natural gas production as an example. The vertical axis shows the cumulative distribution function of all projection errors for natural gas production since 1985 and the horizontal axis shows the corresponding percent projection error values. Positive projection errors mean that the projected value was higher than the observed historical value (an over-projection), and vice-versa. We show error distributions separately, by projection interval, in black, blue and red lines. The bounds of the 95% probability interval, 2.5th and 97.5th percentiles, are highlighted in horizontal magenta lines. In Figure 1 we see that in the short-term, cases where natural gas production was over-projected by more than 11% or under-projected below -19% are considered extreme errors. In this case, there is little median drift for short-term and medium-term quantities, which have a median at -1.3% error. There is a median drift of 3.1% for long-term projections. For more on mean and median drift, see the SI, **Table 17** and **Table 18**.

See the SI, **Figure 30**, **Figure 31**, and **Figure 32** for error cumulative distribution functions and extreme error thresholds for all twenty quantities, including three derivative quantities: oil and natural gas imports and energy-related CO_2 emissions. See the SI, Section 6.16 for a related discussion of mean and median drift among the quantities examined.

We estimate the probability that our main results could have occurred by chance using Monte Carlo simulation, representing extreme errors and extreme changes as draws from a Bernoulli distribution and an integer uniform distribution respectively.

For extreme errors, we simulate a set of projections for the seventeen quantities analyzed. We use a Bernoulli distribution to randomly assign each projection as either an over-projected extreme error (probability 0.025), or not (probability 0.975). We similarly assign each projection as either an under-projected extreme error (probability 0.025), or not (probability 0.975). In this

way, 2.5% of all projections are under-projected extreme errors, and 2.5% are over-projected extreme errors. For each quantity, we simulate 71 projection errors in 1985-1994, 126 in 1995-2004, and 181 in 2005-2014. These correspond to the number of projections for sixteen of seventeen quantities. GDP has only 45 in 1985-1994, 124 in 1995-2004, and 181 in 2005-2014. For simplicity, we give GDP the same number of simulated projections as the other quantities. We use Monte Carlo simulation to replicate this process of random extreme error generation 1000 times per simulation.

We then simulate cross-quantity correlations in two ways. First, we parametrically set a fixed cross-quantity correlation for all pairs of unique quantities for all years. Second, we use cross-quantity correlations derived from projection errors from each individual AEO. We measure cross-quantity error correlation $\rho_{q_i,q_j,k}$, between quantities $q_i \neq q_j$, for a given AEO, k, as:

$$\rho_{q_i,q_j,k} = Corr_t(\varepsilon_{i,t,k}, \varepsilon_{j,t,k}),$$

Where $\varepsilon_{i,t,k}$ is the projection error for quantity q, projection year t, and AEO base year, k. $Corr_t()$ is correlation, operating over projection years, t.

This analysis does not model serial correlation which, if included, could increase the simulated probability of our results occurring by chance. There are two types of serial correlation of concern. The first is serial correlation between errors from the same AEO for successive projection years. The second is serial correlation between errors for the same projection year from AEO reports from successive years. See the SI, Section 6.12 for approximate bounds on the effects of serial correlation.

The two key results we examine are first, the increase in the frequency of extreme errors in 2005-2014 relative to 1995-2004 for 15 of the 17 quantities, and second, that extreme errors occur in

2005-2014 for 10 of the 17 quantities. We use Monte Carlo simulation to estimate the probability of each of these events, P(increase in 15 quantities in '05/'14 v. '95/'04), and P(All + extreme errors for 10 quantities in '05/'14), respectively. We estimate these probabilities as the fraction of the 1000 Monte Carlo iterations in a given simulation in which each respective condition is met.

Figure 5 shows P(increase in 15 quantities in '05/'14 v. '95/'04) using cross-quantity correlation derived from projection errors in each AEO report for which there are at least four projection errors (AEO 1982-2011). The middle column of **Table 1** shows P(increase in 15 quantities in '05/'14 v. '95/'04) using parametric cross-quantity correlation. Note in **Figure 5** that for correlations derived from all AEO reports, P(increase in 15 quantities in '05/'14 v. '95/'04) is less than 5%.

We see instances in which all over-projected extreme errors occur for two to ten quantities, when using empirical correlations from errors in AEO 1982-2011. Thus, the probability of the observed ten quantities occurring by chance is less than 0.5%.

In the vast majority of simulations, we find that no iteration produces ten or more quantities in which all over-projected extreme errors occur in 2005-2014. This occurs only in one of 1000 iterations using correlations from AEO 2004. **Figure 6** shows the maximum number of quantities in which all over-projected extreme errors occur in 2005-2014 for simulations using cross-quantity correlations from each AEO report.

Table 1. Probability of observed results from a frequency of extreme errors analysis under uniform cross-quantity correlation between all unique pairs of quantities. Probability is computed as the fraction of the 1000 Monte Carlo iterations in a given simulation for which the desired condition holds true. The odds of the first outcome occurring by chance are below 5% for correlations less than 75%, while the second outcome does not occur by chance with probability at least 5% even with 99% correlation.

Cross-quantity correlation	P(increase in 15 quantities in '05/'14 v. '95/'04)	P(All + extreme errors for 10 quantities in '05/'14)
0%	0.1%	0.0%
10%	0.0%	0.0%
50%	1.6%	0.0%
75%	5.7%	0.0%
90%	13.4%	0.0%
99%	30.5%	0.4%

The right column of **Table 1** shows the probability of observing ten or more quantities for which all over-projected extreme errors occur in 2005-2014, P(All + extreme errors for 10 quantities in '05/'14), using parametric cross-quantity correlation. The measured probability is zero in all cases except 99% cross-quantity correlation, in which case it is less than 0.5%.

These results suggest that the most extreme errors have indeed become larger for many quantities in the period from 2005-2014. We estimate that it is unlikely, but not inconceivable that we could observe our results by chance. We find that accounting for both types of serial correlation described above may increase the probability of an increase the frequency of extreme errors in 2005-2014 relative to 1995-2004 for 15 of the 17 quantities to above 5%. Adding serial correlation does not increase the probability that all extreme errors occur in 2005-2014 for 10 of the 17 quantities to above 5% unless the values of one or both types of serial correlation are consistently at or above 99%. Spearman serial correlations in both directions described above

have a median value of 75%, with a standard deviation of 36%, well below the 99%+ level required. See the SI, Section 6.12 and **Table 19** for further details.

Similarly, we simulate the probability that our key volatility results could have occurred by chance using Monte Carlo simulation, assigning an upward and downward extreme change for each quantity using an integer uniform distribution over the years between 1950 and 2014. By definition, there is one upward and one downward extreme change for each quantity over the full study period. Our key results are that in 2005-2014, there are nine of a total of thirty-four extreme changes, eight of which are downward. For each quantity, we randomly select two years from an integer uniform distribution between 1950 and 2014, an upward and a downward extreme change. Note that because the data go to 1949, year-over-year changes begin in 1950. Also, because of sampling with replacement, there is a 1.5% chance of both the upward and downward extreme change occurring in the same year. This will slightly bias our results toward a higher probability of multiple extreme changes in the same decade. We consider cross-quantity correlation between pairs of unique quantities, using both a constant parametric correlation between all quantities, and correlations derived from the historical values of the quantities.

Table 2 shows the probability of both key results occurring by chance, the percentage of iterations in which there are nine or more extreme changes in 2005-2014, P(9 extreme changes in '05/'14), and the percentage in which there are eight or more downward extreme changes in 2005-2014, P(8 downward extreme changes in '05/'14). In Table 2, we see that using historical Spearman correlations, both of the baseline results occur with greater than 10% probability, meaning that it is not unlikely that they occurred by chance. The historical correlations are roughly analogous to a uniform correlation level of 50%.

Table 2. Probability of observed results from extreme change analysis. Using historical correlations, or cross-correlations above 50%, the probability of each of the two outcomes occurring by chance is well above 5%.

Rank Correlation	P(9 extreme changes in '05/'14)	P(8 downward extreme changes in '05/'14)
0%	2.1%	0.1%
10%	7.3%	0.9%
50%	18.5%	9.4%
90%	25.6%	15%
99%	27.5%	15.6%
Historical correlations	27.3%	14.9%

2.4. Analysis and results

2.4.1. Year-on-year volatility and unpredictability

We measure changes in the year-on-year volatility and unpredictability for seventeen key US energy-related quantities — the price, consumption, and production of oil, natural gas, and coal; electricity price and sales; residential, commercial, transportation, and total energy consumption; gross domestic product (GDP), and inflation. We measure *year-on-year volatility (volatility)* by computing year-over-year changes in observed historical values, Δh_t, from 1949 to today for each quantity, computed as:

$$\Delta h_t = \frac{h_t - h_{t-1}}{|h_{t-1}|}, \quad (1)$$

where h_t is the historical value of an energy quantity in year t. We identify the years in which the single most positive and most negative changes occur for each quantity as *extreme changes*. Periods with more extreme changes across different quantities are more *volatile*. We compare alternative definitions of volatility in the SI, Section 6.6 and in **Figure 19** and **Figure 20**.

This metric allows comparison of the relative volatility of different quantities, including those with particularly large or small historical variation, for annual-resolution data. Similar metrics are used in the finance and energy forecasting literatures (see the SI, Section 6.7). We define *unpredictability* as the prevalence of extreme errors within a time period. We define a projection error, $e_{i,t}$, from a projection, $p_{i,t}$, made in year i for year t, as:

$$e_{i,t} = \frac{p_{i,t} - h_t}{|h_t|}, \quad (2)$$

where h_t is the historical value in year t.

We define projection length ($l_{i,t}$), with $l_{i,t} = t - i$. Because the nearer future may be easier to predict, we perform our analysis over different *projection intervals*, defined as short-term ($1 \leq l_{i,t} \leq 5$), medium-term ($6 \leq l_{i,t} \leq 10$), and long-term ($11 \leq l_{i,t} \leq 21$). For each projection, $p_{i,t}$, its projection length, $l_{i,t}$, falls in one of these projection intervals, categorizing that projection as short-term, medium-term, or long-term. Some grouping is necessary to ensure adequate statistical power in our analysis. In 6.8, we define projection intervals in different ways and find that our key results are generally robust to different projection intervals. For each projection interval, we define *extreme errors as* those errors located outside the 95% probability interval – i.e. below the 2.5th percentile or above the 97.5th percentile of errors (See Methods and data, Section 2.3, and **Figure 1**).

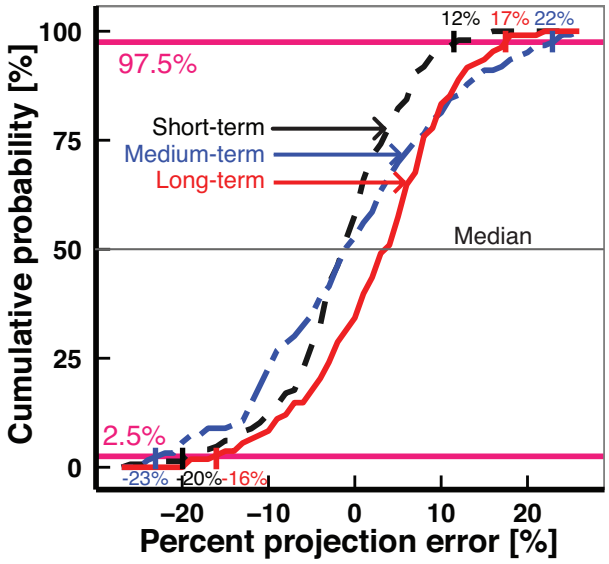

Figure 1. Cumulative distribution functions (CDFs) of the percent projection errors for natural gas production separated by projection interval. The 2.5th and 97.5th percentiles are shown as magenta horizontal lines, whose intersection with the CDFs are the extreme error thresholds. Projections 1-5 years into the future are short-term, 6-10 years are medium-term, and 11-21 years are long-term. Note that the median value for short-term and medium-term projections is close to zero, while the median for long-term projections is closer to 3%.

We measure unpredictability by comparing the frequency of extreme errors, f_τ, over a time-period, τ:

$$f_\tau = \frac{N_{e,\tau}}{N_{p,\tau}} \tag{3}$$

Where $N_{e,\tau}$ is the total number of extreme errors in time-period τ, $N_{p,\tau}$ is the total number of projection values in time-period τ, and τ is a single year, or a set of years.

2.4.2. Increase in volatility and unpredictability

We compare the relative frequency of extreme changes (*year-on-year volatility*) and extreme errors (*unpredictability*) in consecutive 10-year periods.

Our definition of volatility is based on year-over-year percent changes, shown below for natural gas price and oil production in **Figure 2**. Note that the largest decrease in natural gas price and the largest increase in oil production between 1949 and 2014 occur between 2005 and 2014. See the SI, **Figure 21** for plots of the remaining quantities.

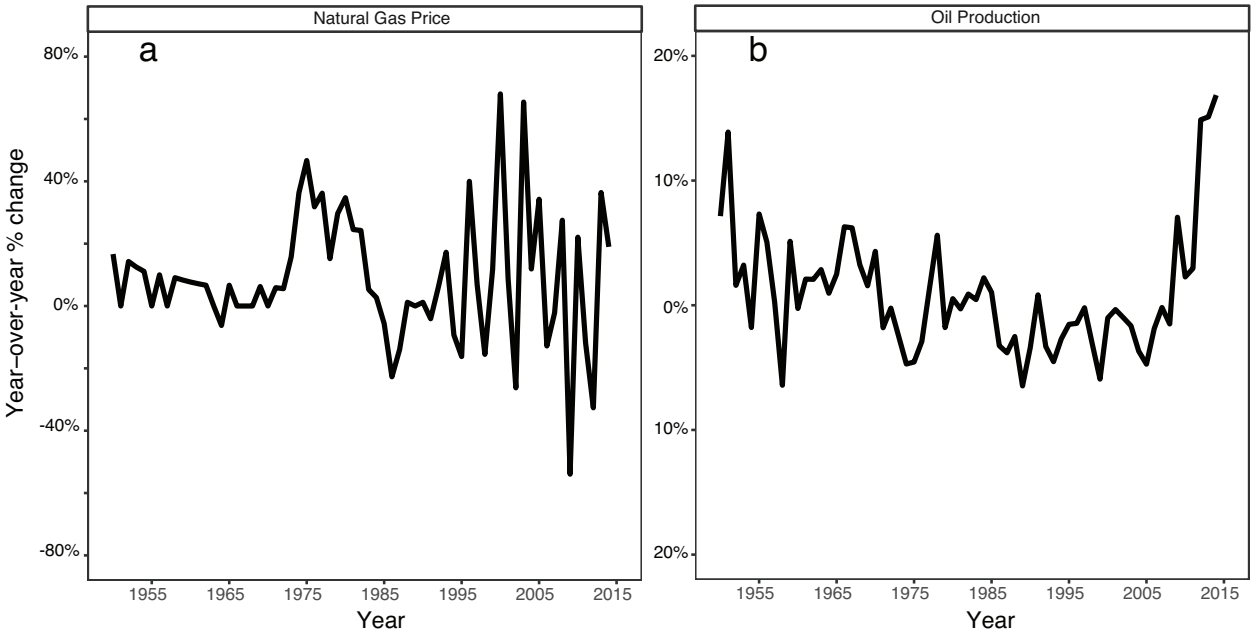

Figure 2. Year-over-year percent changes for two energy quantities. a) natural gas price (in constant 2005 dollars) and b) oil production data show that the largest decrease in natural gas price and the largest increase in oil production both occur between 2005 and 2014.

The last decade (from 2005 to 2014) was more volatile than the preceding three decades: comparable levels of volatility are only seen in the 1950s and 1960s. Figure 3 shows extreme changes for each energy quantity over time. Black triangles indicate the year of the greatest increase, Δh_t, in each quantity since 1949, and red triangles indicate the year of greatest decrease.

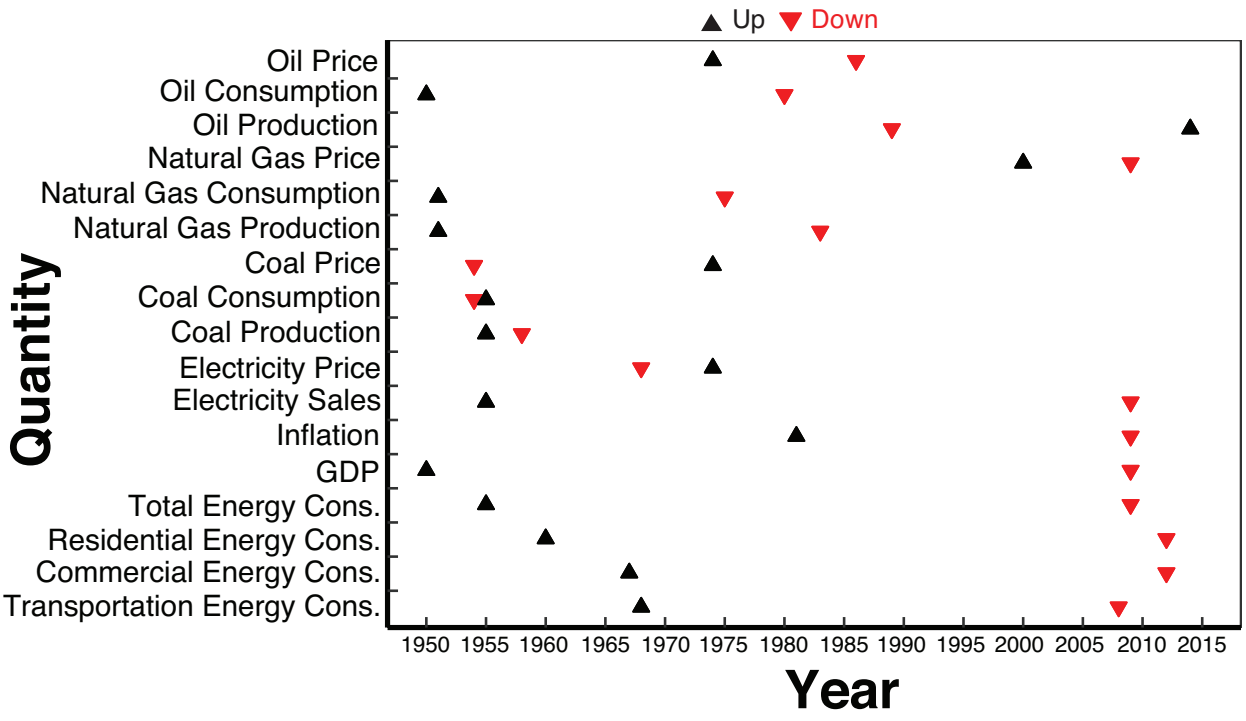

Figure 3. Extreme changes for seventeen energy quantities, from 1949 to 2014. The black and red triangles indicate the largest year-over-year increases and decreases in the quantities. Note the high concentrations in the 1950s, 1960s, and from 2005-2014.

By our definition, there are thirty-four extreme changes, two for each of the seventeen quantities. Nine of these quantities fall between 2005 and 2014, with only seven in the entire thirty-year period from 1975 to 2004. The remaining eighteen fall between 1950 and 1974. Only in the 1950s and 1960s is there a comparable concentration of extreme changes, largely driven by high economic growth rates during that period, a tighter relationship between economic growth and energy consumption, and smaller baseline levels for most quantities.

The few extreme changes that occur between 1975 and 2004 are associated with oil and natural gas, due to major swings in international oil markets in the 1970s and 1980s, and rapid changes in the use and regulatory structure of natural gas in the 1980s and 1990s[34].

Of the nine extreme changes in the decade from 2005 to 2014, eight are abrupt decreases, most likely due to the Financial Crisis and its aftermath. The widespread adoption of horizontal drilling with hydraulic fracturing in shale formations, particularly after 2007, is unquestionably a major factor in the 17% increase in oil production in 2014, and the 54% decline in natural gas prices in 2009.

If we normalize energy production and consumption quantities and GDP, by total US population, the volatility results are similar, with eight extreme changes in 2005-2014 and seven in 1975-2004 (see the SI, Section 6.9).

2.4.3. Unpredictability in recent years

We find that unpredictability, measured as the frequency of extreme errors in AEO projections, has increased in the most recent decade. Figure 4 shows the frequency of over-projected (red) and under-projected (black) extreme errors since 1985 (the first year for which there are AEO projections). The placement of circles along the x-axis corresponds to the year in which extreme errors occur. The size of each circle corresponds to the frequency of extreme errors: for example, the 47% frequency of over-projected extreme errors in coal consumption in 2012 means that 47% of all projections of coal consumption for year 2012 from different AEO reports resulted in over-projected extreme errors. See Supplementary Data 1 for data underlying Figures 3 and 4.

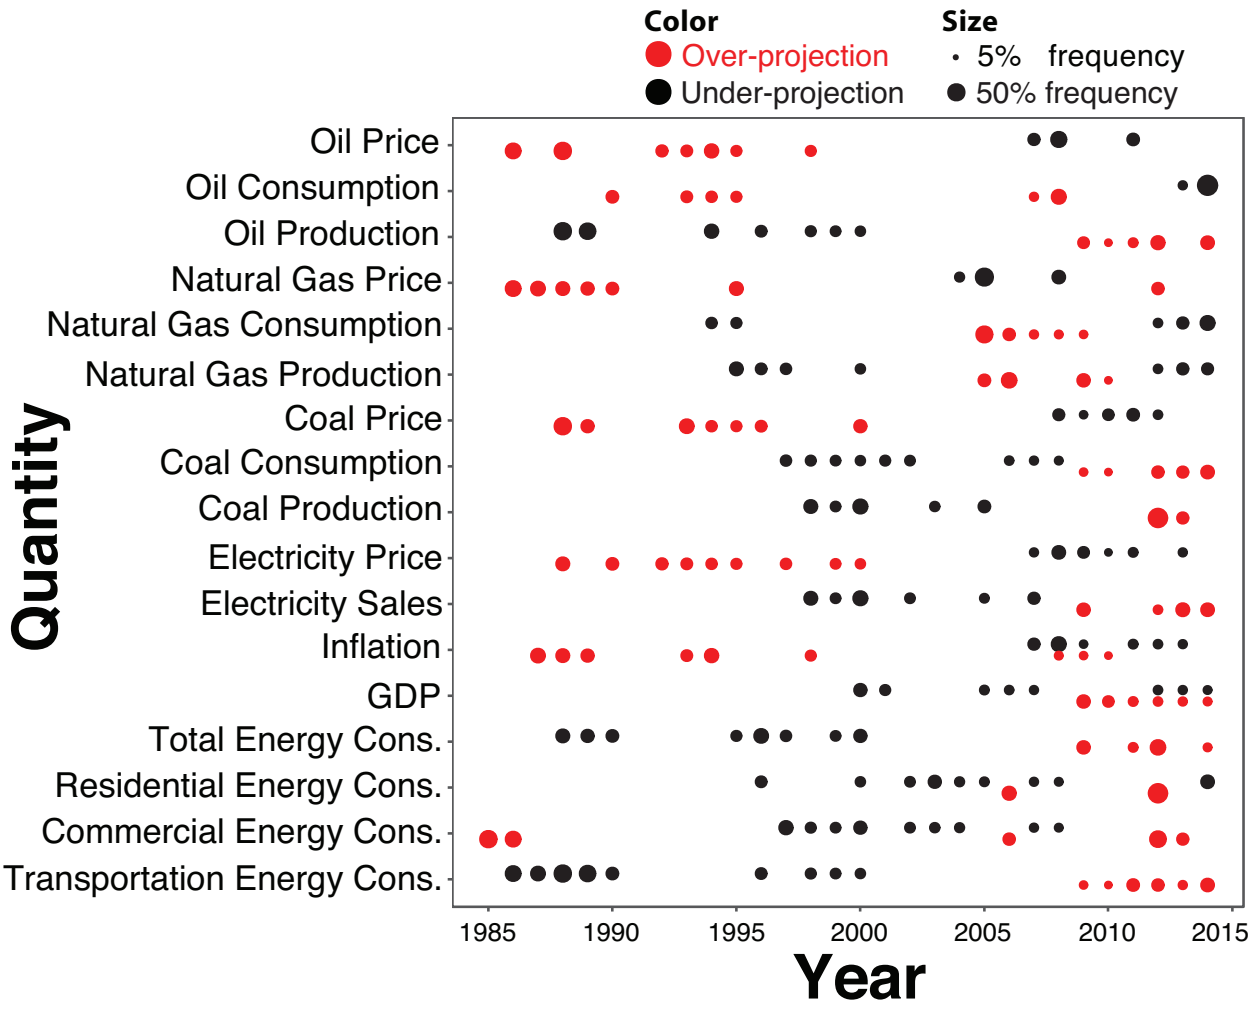

Figure 4. Annual frequency of extreme errors for each quantity. Red and black circles correspond to over-projected and under-projected extreme errors, respectively. The size of each circle corresponds to the frequency of extreme errors in that year. Note that for ten quantities, all over-projected extreme errors occur in 2005-2014.

The high concentration of extreme errors in the last decade consists largely of under-projections for prices and inflation and over-projections for energy production, and consumption. All over-projected extreme errors over the thirty-year study period occur in 2005 to 2014 for ten quantities: production and consumption of natural gas and coal; oil production; electricity sales; total, residential, and transportation energy consumption; and GDP. This means that the largest over-projected errors in this period for the short-, medium-, and long-term for these quantities

were larger than any seen in the preceding twenty years. In the same period, 2005-2014, fifteen quantities, all except total and commercial energy consumption, experience an increase in the overall frequency of extreme errors relative to 1995-2004. See the SI, **Figure 22** and **Figure 23**, and Section 6.10 for analogs to **Figure 3** and **Figure 4** that include three derivative but important quantities, oil and natural gas net imports and energy-related carbon dioxide emissions.

To evaluate the probability of these results occurring by chance, we perform Monte Carlo simulation. In each simulation, we randomly generate 1,000 datasets of projection errors from the seventeen quantities. Drawing from a Bernoulli distribution, each simulated projection error has a 2.5% chance of becoming a positive extreme error, and a 2.5% chance of becoming a negative extreme error. The probability of an event occurring by chance, the p-value, is the fraction of simulated scenarios in which that event occurred. We replicate each simulation using cross-quantity Spearman correlations derived from the projection errors in each AEO report (see Section 2.6 for these and other related correlations, and the SI, Section 6.11 for further discussion of these correlations). We report an upper bound on the highest p-value across all such simulations as the probability of an event occurring by chance.

Our results suggest that under realistic levels of cross-quantity projection error correlation, it is unlikely but possible, with probability below 5% in all cases, that we would observe an increase in extreme errors from 1995-2004 to 2005-2014 for 15 quantities by chance in a time-stationary process (see **Figure 5** and **Table 1**). It is highly unlikely, with probability below 0.5%, that all over-projected extreme errors would occur in 2005-2014 for 10 quantities (see **Figure 6**). In other words, both of these results are, in a certain sense, statistically significant. See Section 2.3 and the SI, Section 6.12 for a further discussion bounding the effects of serial correlation.

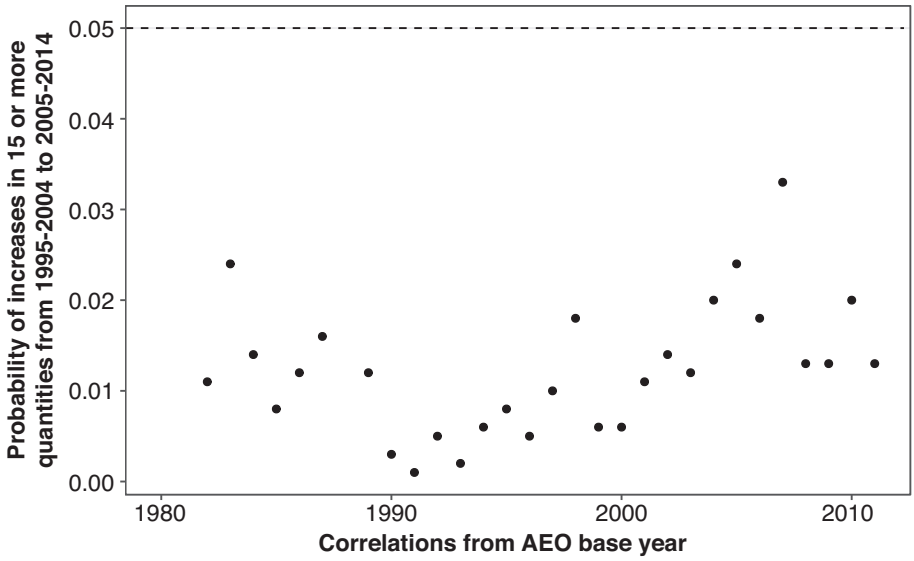

Figure 5. The simulated probability of observing increases in the frequency of extreme errors for at least fifteen of seventeen quantities between in 2005-2014 v. 1995-2004. This uses correlations from AEO 1982-2011. In all cases, this probability is less than 5%, the dashed horizontal line.

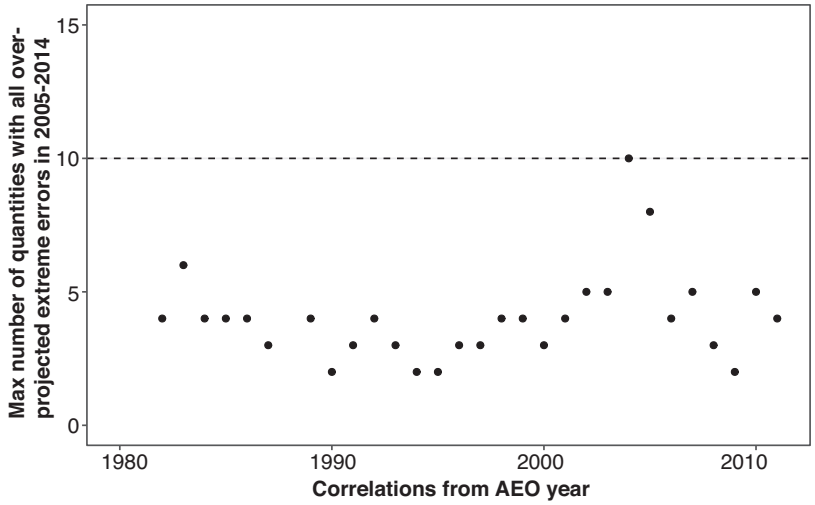

Figure 6. The maximum number of quantities for which all over-projected extreme errors occur in 2005-2014. This is derived from 1000 simulations using empirical cross-quantity correlations from each AEO report. The dashed horizontal line is at the observed value of ten.

Using similar methods, described in Materials and Methods, we compute the probability of the volatility results occurring by chance. Given historical levels of correlation between the

quantities themselves, **Table 2** shows roughly a 10%-20% chance of finding the observed disproportionate clustering of extreme changes in 2005-2014 in a time-stationary process. The large number of under-projected extreme changes is consistent with extreme errors results, especially the large number of quantities for which all over-projected extreme errors occur in 2005-2014.

In the SI, Section 6.8, we test the robustness of these results to alternative definitions of extreme error, and subsets of the data. See the SI, **Figure 24** and **Figure 25** for graphical representations of these results.

2.5. Conclusions and policy implications

We find an increase in both year-on-year volatility and unpredictability for a broad range of quantities in the most recent decade relative to the immediately preceding decades. In the SI, Section 6.13, we demonstrate a case in which considering errors from 2005-2014 makes the difference between profit and loss for a liquefied natural gas export terminal (see Supplementary Figure 10 for the natural gas price scenarios used). Still, volatility was highest in the period from 1950-1974, meaning that the relative quiescence from 1975-2004 may itself be the anomaly. Also, the observed increase in volatility may be due to increased flexibility in energy infrastructure deployment, e.g. distributed energy resources and hydraulic fracturing wells, but the implications for energy decision-makers are largely the same regardless. Note that high concentrations of extreme errors begin *before* both the massive expansion of hydraulic fracturing *and* the Great Recession. For example, a concentration of extreme errors between 2005 and 2006 for natural gas production, consumption, and prices is visible in **Figure 4**. This suggests that the observed increase in volatility and unpredictability in this decade is due to a number of interlinked, unanticipated developments. While all the authors of this paper have worked on

topics related to energy systems, we do not claim to be experts on all the drivers of the observed results. We suggest that these developments may have contributed to the observed results:

In the late 1990s, natural gas was cheap and abundant, and projected to remain so for decades [35]. These expectations, coupled with newly restructured electricity markets in many US states, encouraged construction of natural gas electricity generation plants on an unprecedented scale, particularly between 1999 and 2005 [36]. However, large increases in offshore production and Canadian imports of natural gas predicted in the 1990s and early 2000s failed to materialize, driving up natural gas prices [30]. As a result, generation costs for new natural gas plants were higher than anticipated, and plant utilization was much lower [30]. After 2007, tight oil and natural gas production increased massively, driving down natural gas prices, and encouraging increased use of natural gas for electricity generation, displacing generation from coal [30]. Global oil prices rose substantially in the mid-2000s, peaking in 2008, largely as a result of increased demand in Asia, particularly China, and the Middle East [37]. High prices fostered the expansion of unconventional oil extraction in the US and internationally. The Financial Crisis of 2007-2009 and the ensuing Great Recession depressed demand for energy [30], placing a downward pressure on energy prices, including oil and natural gas. Vehicle transportation usage declined, with vehicle miles traveled peaking on a per-capita basis in 2005 and on a national total basis in 2007 [38]. Finally, industrial energy consumption fell due to deindustrialization and increased energy efficiency, a trend in many sectors of the economy. In combination, these factors led to an unexpected decline in total US energy consumption from its peak in 2007 [30].

The observed increase in the volatility and unpredictability of key energy-related quantities may suggest complex structural shifts in the US and world economies and energy systems. Any improvements in the world's most sophisticated energy system models would likely have been

overwhelmed by these changes. This turbulence may or may not continue. However, this analysis should serve as a stark reminder of the importance of considering the possibility of further surprises when planning for the future.

2.6. Data availability statement

Data that support the plots within this paper are available from the US Energy Information Administration (https://www.eia.gov/outlooks/aeo/archive.php), the Bureau of Economic Analysis (https://www.bea.gov/industry/xls/io-annual/GDPbyInd_VA_1947-2016.xlsx), and Oak Ridge National Laboratory (http://cdiac.ornl.gov/ftp/trends/emissions/usa.dat), and are described further in Supplementary Methods.

3. Chapter 3. Characterizing the relationship between low-income electric subsidies and the intra-day timing of electricity consumption

With Inês M. L. Azevedo, Dept. of Engineering and Public Policy, Carnegie Mellon University

3.1. Abstract

Electricity rate subsidies have long been used to provide low-income households with greater access to affordable energy services, which can bring enormous quality of life benefits. Such policies are likely to increase consumption and indeed this is arguably an intended consequence. Increases in electricity consumption necessitate additional electricity generation, with associated greenhouse gas (GHG) and criteria pollutant emissions, and may lead to increases in peak capacity requirements as well. We estimate the effect of the California Alternate Rates for Energy (CARE) subsidy on electricity consumption in northern California with a panel of interval electricity consumption data from more than 30,000 households using a difference-in-differences model. We find that CARE enrollment is associated with an increase in electricity consumption of 13% [10%, 15%], with a comparable increase in the hot Central Valley, a lower increase on the Coast, and a higher increase in the Inland Hills region in between. There is significant intra-day variation across all regions and seasons, with significant increases during current peak demand periods of 3-9pm as high as 6% [3%, 10%] in all cases except in the Central Valley. Using a time-invariant estimate of the price elasticity of demand, we estimate the climate and human health damages and peak capacity costs associated with this increase in consumption at $31M [$25M, $37M], $14M [$11M, $17M], $29M [$23M, $34M], respectively. Switching to intra-day estimates increases these costs by up to $4M, although this increase is not

statistically significant. In light of these externalities, it would be useful to compare the cost-effectiveness of energy subsidies with other methods of ensuring affordable access to energy services, including subsidies for energy efficiency, which may in some cases achieve the same result without increasing emissions or stressing the grid. Policy makers may want to consider such indirect costs as part of a holistic energy policy strategy, as they represent roughly 11% of annual expenditures on CARE disbursements. A time-invariant price elasticity of demand will likely be sufficient for most purposes.

3.2. Introduction

Throughout the United States and internationally, many electric utilities offer subsidized electricity rates to low-income households to ensure affordable access to important energy services [39,40]. As governments enact policies to reduce greenhouse gas (GHG) emissions from the electric power sector, subsidies will remain an option to help ensure that the associated costs are not borne disproportionately by low-income households.

Economic theory suggests that the implementation of subsidies will increase electricity demand in enrolled households through a price elasticity of demand [41,42]. Because electricity supply must always match instantaneous demand, the timing as well as the magnitude of this increase will affect the cost and reliability of the electric power system as a whole. In addition, the marginal emissions of GHG and criteria pollutants from the electric power sector change over time, meaning that consumption at different times can result in greater or smaller damages [43,44].

The importance of timing will likely become even more pronounced in an electric power system with high levels of non-dispatchable renewable electricity, such as wind and solar. California's high levels of solar photovoltaic electricity are already suppressing midday springtime wholesale prices to very low, or even negative values [45]. This means that increased consumption during

those hours places relatively little incremental burden on the system and, if this consumption prevents curtailment of renewable electricity generation, few additional emissions of GHG or criteria pollutants. However, as the sun sets and many people return home from work, a host of dispatchable fossil-fired electric power generators must come online to meet demand. Additional consumption during this period requires additional dispatchable generation capacity that must be available and likely adds GHG and criteria pollutant emissions from fossil fuel combustion as a result. Furthermore, incremental electricity demand during peak hours, e.g. additional air conditioning during the hottest hours of the year in California, increases requirements for peaking electricity generation capacity and transmission and distribution infrastructure. In the United States, all of these costs are generally borne by electric ratepayers.

Policymakers interested in ensuring affordable, low-carbon electricity for all can thus gain important insights into the broader impacts of low-income subsidies by considering not only the average effect of low-income subsidies on electricity demand, but also the intra-day timing of these effects across different regions and seasons. Such analysis can then inform the design of future time-of-use residential electricity rates that can help ensure affordable, clean electricity for all.

3.2.1. Background

Numerous studies estimate the price elasticity of residential electricity demand, the marginal change in electricity demand induced by a marginal change in electricity price, across different populations in different countries. Estimated values for the short-run, often defined as the response within the same month as a change in price, fall between 0 and -2.5 [41,42,46–53], with a median around -0.35 [41]. In the long-run, generally timescales of months to years, estimates range from 0 to -4.56 [41,42,46–53], with a median around -0.81 [41]. This wide range is partially attributable

to the time and geographical location at which each study was conducted, how long a time series was used, and inherent uncertainty in what are generally observational estimates of a causal parameter. However, it is also indicative of heterogeneity in local climates, and customer characteristics. Given this, it is reasonable to believe that the price elasticity of residential electricity demand may be different depending on the time of day. There is evidence that the price elasticity of national electricity demand tends to fall as gross domestic product increases in industrialized economies [54], but the relationship between income and price elasticity is less clear at the household level.

The widespread deployment of advanced metering infrastructure (AMI) by many electric utilities over the past decade and a half facilitates more detailed analysis of intra-day variation in the effect of various interventions on residential electricity consumption [55]. Jessoe et al. [56] use a randomized controlled trial to estimate the intra-day responsiveness of household electricity consumption to critical peak electricity pricing, in which electricity rates rise during anticipated peaks in system-wide electricity demand. Boomhower et al. [57] use quasi-experimental methods to estimate the timing of electricity savings from an air conditioner repair program. Qiu et al. [58] estimates hourly energy consumption effects of Energy Star and Leadership in Energy and Environmental Design (LEED) building certifications.

This analysis focuses on the California Alternate Rates for Energy (CARE) subsidy, which is available to low-income California households who receive electric or natural gas service from investor-owned utilities [39]. Eligible households must have income within 200% of the federal poverty level or must meet one of a number of categorical criteria, such as enrollment in other state or federal means-tested programs [39]. In 2012, just after the period for which we have data, the program provided an average subsidy of 33%, or $29 per month to 3.2 million California

households, with households in the Pacific Gas and Electric Co. (PG&E) service territory receiving an average discount of 42%, or $40 per month [39]. CARE is funded by a public purpose customer charge included in electricity rates, with $4 billion approved for the 2012-2014 budget cycle [39]. Thus, CARE constitutes a transfer payment from all non-CARE customers.

California has had low-income electricity subsidies since 1989, with substantial expansions in the aftermath of the 2001 Energy Crisis. In 2012, 32% of California households, and 30% of PG&E households, were eligible for CARE [39]. In 2011, 93% of eligible California households were enrolled in CARE, with 90% participation in the PG&E service territory [39], a 31% increase over 2008 enrollment [59]. This is consistent with estimates that approximately 10% of eligible households may be choose not to participate in CARE [60]. See the SI, Section 7.1 for further historical and demographic information.

3.2.2. Potential for intra-day variability in the price elasticity of demand

Existing estimates of the price elasticity of demand generally do not capture intra-day variability in consumption effects. CARE enrollment could also result in relatively time-invariant increases in electricity consumption, such as the purchase of an additional refrigerator. However, a low-income subsidy could also result in differential effects at different times in numerous ways. Lower prices may incentivize households to increase use of occupancy-dependent energy services, i.e. end uses that people only use when they are at home. This could include new capital purchases, e.g. a new television, or more intensive use of existing devices, e.g. watching more television.

Occupancy-dependent loads depend on occupancy patterns. Thus, one would expect such increases to occur primarily in the evenings after work or school and to a lesser extent in mornings as many occupants are preparing to go to work or school.

In warmer regions of California in particular, the presence of a subsidy may encourage adoption or more intensive use of air conditioning, a major determinant of peak demand. In the Central Valley, where roughly 95% of the population has some form of air conditioning [61], electricity demand approximately doubles during the summer in our sample (see the SI, Section 7.2.5 for further details). Conversely, electricity demand in the Coast and Inland Hills is fairly flat throughout the year, with a modest winter peak. The climate is substantially cooler on the Coast, where ~15% of households have air conditioning, and fairly moderate in the Inland Hills, where ~60% of households have air conditioning [61].

This suggests that CARE enrollment is unlikely to spur air conditioner adoption in the Central Valley but could encourage additional use of existing air conditioning. Households in warmer areas of the Inland Hills may adopt air conditioning. The cooler climate on the Coast suggests that CARE enrollment is unlikely to spur adoption of air conditioning and any additional air conditioners likely would not be used very often, although perhaps they would be used on hot summer peak demand days.

3.2.3. Aims

This study aims to estimate the relationship between enrollment in the CARE low-income electric subsidy in northern California and subsequent electricity consumption for households across different California regions and in different seasons. The primary goal is to establish whether there are significant intra-day differences in the estimated effect. That is, does the price elasticity of demand vary throughout the day? We then evaluate whether accounting for these intra-day differences substantively changes estimates of the indirect costs of this subsidy program.

3.3. Methods and data

3.3.1. Hourly panel of household electricity consumption

We use hourly electricity consumption data from a random sample of roughly 30,000 households in northern California, acquired from the Pacific Gas & Electric Co. (PG&E) service territory through the Wharton Customer Analytics Initiative. Data include enrollment and disenrollment dates for the CARE program and several other utility programs, as well as dates of participation in rebate programs for energy-efficient appliances and services. These data also include each household's census block, which allows matching both to neighborhood-level demographics from the 2010 Census and to temperature data from nearby weather stations.

Advanced metering infrastructure (AMI) was installed at dwellings during the study period of 2008-2011. For each dwelling, hourly data begin after the installation of AMI, which was staged across the three regions of the Central Valley, Inland Hills, and Coast, shown in the SI, Sections 7.2.1 and 7.2.2. As a result, the panel is unbalanced, with two or more years of data from most Central Valley dwellings and less than one year of data for most dwellings on the Coast. This stratified sample includes 8,597 dwellings on the Coast, 11,391 in the Inland Hills, and 10,217 in the Central Valley, totaling 30,205 dwellings.

Table 3 shows demographic statistics for the sample based on census block-level information. Households in the predominantly agricultural Central Valley tend to have lower incomes, lower home values, and less education than those in the Coast and Inland Hills, which together include most of the wealthy San Francisco Bay Area, including Silicon Valley. Rates of renting are higher on the more urban Coast than in the more suburban and rural Inland Hills and Central Valley.

Table 3. Summary statistics for 2010 census block neighborhoods of households in the sample*. The Central Valley has the lowest incomes and home values.

	Central Valley	Inland Hills	Coast	Full sample
Median Home Value*	282,000	586,000	597,000	479,000
Median Income*	51,800	78,500	63,400	65,600
Median % Renters	34	32	51	38
Median % w/ Bachelors (or higher)	17	38	40	32
Number of dwellings	8,597	11,391	10,217	30,426

* These values are medians from our sample of Census block neighborhood medians. The values are top-coded by the US Census at $1M and $250k, respectively. We report the values rounded to the nearest $1000 for median home value, and to the nearest $100 for median income values. Adapted from Meyer [62].

CARE enrollment surpasses 30% of the sample by the end of the study period. CARE households are distributed across all three regions, in rich and poor areas. In the Central Valley, households in the poorest tertile of census block income have CARE participation rates of 55% at the end of 2011. In the Inland Hills and Coast, CARE participation in poor census blocks is also above 40%. Even in census blocks with relatively high incomes, CARE participation is at least 10% in all three regions. See the SI, Section 7.2.3 for further details.

Households also had access to several other utility programs offered through PG&E. These included energy efficiency rebates, an air conditioner demand response program, a seasonal bill smoothing program, critical peak pricing, direct electricity purchases from an alternate supplier, and the option to purchase carbon emission offsets. These programs are described further in the SI, Section 7.2.4.

3.3.2. Regression specification

We estimate intra-day variation in the relationship between enrollment in the CARE low-income electric subsidy and electricity consumption using a difference-in-differences regression model. The main analysis uses the following regression specification, inspired in part by Boomhower et al. [57]:

$$\ln(y_{ith}) = \beta_h CARE_{it} 1_h + \alpha_c T_{ith} 1_{T_{ith}>65°F} + \alpha_w T_{ith} 1_{T_{ith}\leq 65°F} + \gamma_{ih} + \omega_{wh} + \epsilon_{ith} \quad (4)$$

Where y_{ith} is electricity consumption in kWh by dwelling *i* in day-of-sample *t* in hour-block *h*, where the day is divided into eight three-hour blocks, starting at 12am. The three-hour blocks from 3pm-6pm and 6pm-9pm roughly correspond with peak electricity pricing periods in PG&E's current residential time-of use rates, which are offered at 3pm-8pm or 4pm-9pm [63]. We consider $\ln(y_{ith})$ as the independent variable both because this analysis ultimately aims to estimate a price elasticity of demand, which is more easily calculated with logarithmic coefficients, and because the distribution of electricity consumption across households is approximately lognormal (see the SI, Section 7.1). $CARE_{it}$ is CARE enrollment status of dwelling *i* in day-of-sample *t*. This is an indicator variable, which takes the value of 1 when a household is enrolled in the program. 1_h is an indicator each hour-block of the day. T_{ith} is the temperature at dwelling *i* in day-of-sample *t* in hour-block *h*, using the average temperature from the three weather stations closest to the census block of dwelling *i* across the three hours in the hour-block, with data from the National Oceanic and Atmospheric Administration [64]. Electricity rates did not have substantial seasonal variation during the study period, but California has recently adopted seasonally-varying time-of-use rates to more accurately price

peak consumption [63]. In seasonal regressions, we divide the sample into two seasons based on time-of-use rates currently available through PG&E [63]. The four-month Summer is June-September, while Winter is the remaining eight months, October-May. $1_{T_{ith}>65°F}$ and $1_{T_{ith}\leq 65°F}$ are indicator functions for whether the average temperature in hour-block *h* is above or below 65°F. This piecewise linear representation of temperature response accounts for heating and cooling. γ_{ih} is a fixed effect term for dwelling *i* in hour-block-of-day *h*. ω_{wh} is a fixed effect term for hour-block *h* of week-of-sample *w*, with separate fixed effects for weekdays and weekends. $\epsilon_{ith}$is an error term corresponding to dwelling *i* in day-of-sample *t* in hour-block *h*. Intra-day coefficient estimates use the first hour-block, 12am-3am, as the baseline. Coefficients for other hour-blocks represent deviations from the baseline.

The key differences between this regression specification and that used in Boomhower et al. [57] are: 1) We estimate 8 β_h coefficients (one for each hour-block), while they estimate 288 β_{hm}coefficients (24 hours for each month); 2) For us *t* is the week-of-sample, separated into weekdays and weekends, while they treat *t* as the week-of-sample; 3) We control for temperature response. Boomhower et al. study air conditioner energy efficiency, meaning that the introduction of temperature controls would have potentially confounded estimation of efficiency gains for this highly temperature sensitive load [57].

We use Eq. 4 to estimate the hourly relationship between CARE enrollment and electricity consumption for the full sample and for regions within the sample. In all cases, we use cluster-robust and heteroskedasticity-robust standard errors.

In the baseline cases, we use the full sample, including dwellings that never enrolled in CARE, those that enrolled in CARE during the study period, and those that are enrolled in CARE throughout the entire study period. As a result, some dwellings have no pre-treatment data.

Including these dwellings helps capture any differential trends in dwelling electricity consumption between CARE and non-CARE households, such as income shocks that disproportionately affect low-income households.

To quantify the benefits of estimating an hourly price elasticity of demand, we compare the above intra-day estimates of the effect of CARE with time-invariant estimates. Eq. 5 uses the same dataset as Eq. 4 but drops intra-hour differentiation from the treatment effect estimate, β_s, and from the dwelling-level and week- and weekend-of-sample fixed effects, γ_i and ω_w.

$$\ln(y_{ith}) = \beta CARE_{it} + \alpha_c T_{ith} 1_{T_{ith}>65°F} + \alpha_w T_{ith} 1_{T_{ith}\leq 65°F} + \gamma_i + \omega_w + \epsilon_{ith} \quad (5)$$

We also consider a battery of robustness checks, described in detail in the SI, Section 7.5.3.

3.3.3. Bounding analysis: Estimating the value of characterizing intra-day variability in price-responsiveness

We estimate the climate change and human health damages, and electric power system peak capacity costs associated with the estimated increase in electricity consumption associated with enrollment in CARE. We compare costs estimated using intra-day and flat estimates.

In the base case, we estimate the marginal damages associated with changes in hourly electricity consumption from the Western Electricity Coordinating Council (WECC) from 2011 using marginal emission factors from Azevedo et al., described in further detail below in Section 3.3.4 [65]. We estimate peak electricity capacity costs at $170/kW-yr based on the cost of new entry for a natural gas combustion turbine plus transmission and distribution capacity costs from PG&E [66]. See the Section 3.3.5 below for further detail. Capacity costs are estimated using estimates of the increase in electricity consumption associated with CARE enrollment from 6-9pm, capturing

much of PG&E's residential peak pricing period in both the 3-8pm and 4-9pm options available to customers [63].

We estimate the total energy and power increase due to CARE assuming the average household in the sample is representative of the average household in PG&E, with a flat demand profile for simplicity. This corresponds to annual household consumption of 7,643 MWh at 0.87 kW of constant power demand. CARE enrollment reached roughly 34% of PG&E's 4,550,000 customers enrolled at the end of 2011. Thus, we apply the estimated time-invariant or intra-day increase associated with CARE enrollment to the roughly 1,530,000 the estimated households enrolled in CARE in 2012 in the PG&E service territory and compare the climate change, human health, and capacity costs associated with both methods [39].

3.3.4. Climate and human health effects

We compute the climate and human health effects of marginal changes in electricity generation at different hours of the day using hourly emission factors from Azevedo et al. for WECC in 2011 [65]. Marginal emission estimates for greenhouse gases and criteria pollutants are based on regression analysis of historical electricity generation patterns, described in Siler-Evans et al. [43]. Marginal human health damage estimates use the AP2 integrated assessment model, which links emissions in a particular location to human health damages [67].

We place a cost value on greenhouse gas and criteria pollutant emissions using a social cost of carbon of \$40/t(CO_2) and a value of statistical life of \$6M for consistency with similar policy analysis, recognizing that policy makers may wish to use a substantially higher or lower social cost of carbon [65]. Climate and human health damages are reported in 2010 dollars.

Confidence intervals for each hour-block are computed by adding confidence intervals for the baseline coefficient to the corresponding hour-block. Damages are then computed for each hour-

block and the resulting mean estimates and confidence intervals are then summed to produce a single confidence interval. Note that this represents an upper bound on the true confidence interval, as standard errors add in quadrature. However, the primary source of uncertainty is the confidence interval of the baseline estimate, which represents a low bound on the true confidence interval.

3.3.5. Cost of peak electricity consumption

We use a cost of incremental peak electricity consumption based on the California Public Utilities Commission's 2018 Avoided Cost Calculator [66]. Marginal generation capacity is met with a simple-cycle natural gas turbine at $1,250/kW, annualized over a 20-year lifetime at a 7.4% weighted average cost of capital to $121.7/kW-yr, with $12/kW-yr in fixed operations and maintenance (O&M) expenditures [66]. This analysis ignores variable O&M and energy costs, which will be small for a peaking plant that will operate for a very small fraction of the year. Thus, the annualized cost of peak generation is $133.7/kW-yr.

Historical average transmission costs across PG&E's 18 service divisions are roughly $35/kW-yr, with secondary distribution costs at $2.5/kW-yr [66].

This brings the total effective cost of additional peak consumption to $171.1/kW-yr, which we round to $170/kW-yr.

3.3.6. Statistical significance

Note that in the intra-day specification in Eq. 4, it is possible for there to be a statistically significant increase in electricity consumption in one hour-block of the day that falls within the confidence interval of the baseline estimate for the hours 12am-3pm. This adds complexity to the question of whether an increase in damages is statistically significant. In this analysis we say that a mean difference is statistically significant only if it falls outside the confidence interval

computed as the simple summation of the mean estimates and confidence intervals for each hour of the day, recognizing that this is, strictly speaking, an upper bound on the true confidence interval.

3.3.7. Limitations

This analysis is an observational study of a mature utility program attempting to characterize what is fundamentally a causal relationship between electric subsidies and electricity consumption. As a result, there is substantial potential for selection bias and, in this case, we were not able to identify a clear natural experiment or other technique that would enable a convincing improvement over the rich fixed effects difference-in-differences model in Eq. 4. The SI, Section 7.4, details some of the key limitations of the dataset, such as confounding of CARE enrollment with unobserved changes in income and employments status and lack of access to transparent electricity rate information. The section describes methods traditionally used to address such concerns and explains why they likely are not well-suited for this particular case. These methods include instrumental variables, propensity score matching, and regression discontinuity.

3.4. Analysis and results

Participation in the CARE program is associated, as expected, with an overall increase in electricity consumption and significant intra-day variation in all cases.

Figure 7 shows the estimated mean and intra-day increases in electricity consumption associated with enrollment in CARE. The black dashed lines represent the baseline estimate in the hours of 12am-3am, with black dotted lines representing the 95% confidence interval. In the full sample, electricity consumption increases by 11.6% with a 95% confidence interval of [9.2%, 14.2%], with a minimum of across all six scenarios of 7.0% [0.9%, 13.5%] on the Coast and a maximum

of 13.5% [8.9%, 18.3%] in the Inland Hills. Note that logarithmic coefficients and their confidence intervals are converted to percentages using the formula $e^x - 1$.

Intra-day coefficients are statistically significant at the 5% level if their 95% confidence interval does not contain the baseline value, the black dashed line in **Figure 7**. All cases except the Central Valley show significant intra-day increases in electricity consumption during 3-6pm (hours 15-17) or 6-9pm (hours 18-20), which fall within PG&E's current residential summer peak pricing periods of 3-9pm [63]. In addition, the Full Sample sees a significant increase within the system's historical afternoon peak period of roughly noon-6pm, apparently driven by the Coast and Summer [68].

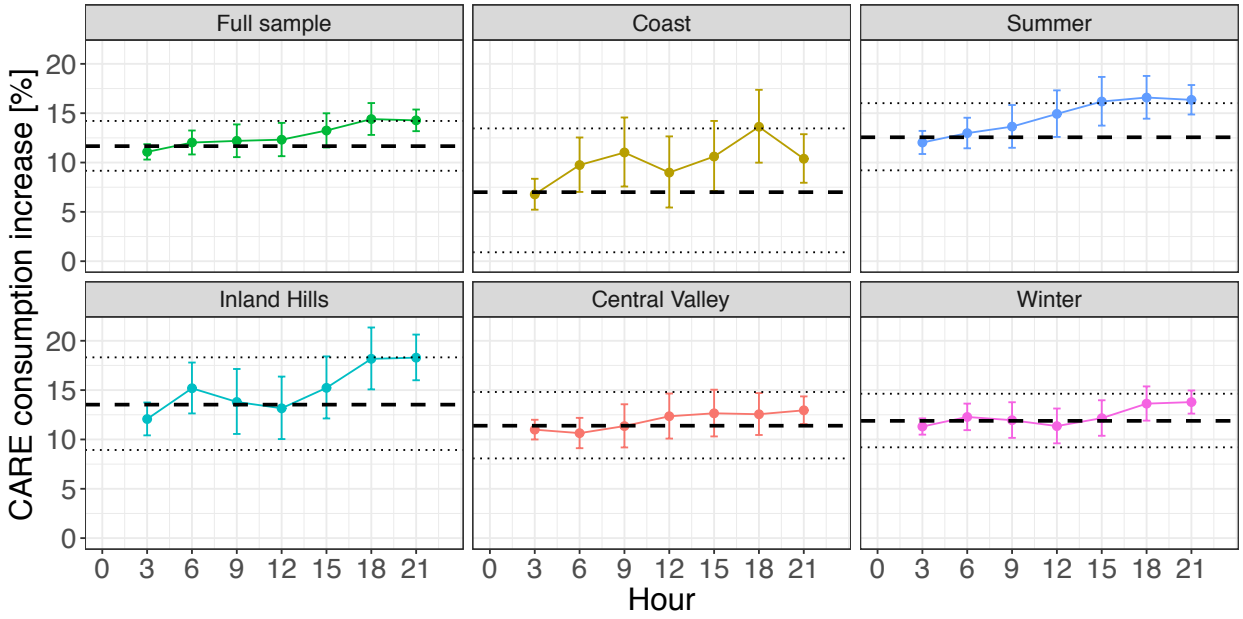

Figure 7. The estimated percent increase in electricity consumption due to CARE. Intra-day changes, with 95% confidence intervals, are relative to the estimate for the first three baseline hours of the day, 12am-3am, the black dashed line, whose 95% confidence interval is represented with black dotted lines. An increase is significant if its 95% confidence interval does not contain the black dashed line. All cases see baseline increases in electricity consumption between 7% and 14%, with significant increases during PG&E's current peak hours of 3-9pm, hours 15-20, in all cases. The absolute increase in consumption is largest in the Inland Hills, Central Valley, and Summer, but the intra-day effects are most pronounced on the Coast.

Note that although the percent increase associated with CARE participation is relatively flat in the Central Valley, there is more seasonal and intra-day variation in that warm region's electricity consumption due primarily to summer air conditioning demand (see the SI, Section 7.2.5). As a result, a constant percentage increase in electricity demand in the Central Valley may correspond to a greater absolute increase in on-peak electricity consumption than a time-varying increase in the Coast or Inland Hills.

The temperature coefficients in Eq. 4 are roughly symmetric in most cases. This may imply a roughly equivalent response to heating and cooling demand, but the fixed effects likely capture predictable seasonal or intra-day temperature response, potentially leading temperature coefficients to underestimate actual responsiveness to temperature. In the Full Sample, each degree Fahrenheit above or below 65°F is associated with an electricity consumption increase of 2.06% [2.01%, 2.12%] for heating and 2.51% [2.45%, 2.56%] for cooling. The largest coefficients are in the Summer, with 4.00% [3.91%, 4.02%] and 3.75% [3.70%, 3.80%] for heating and cooling, respectively. The relatively large summer heating coefficient may be due to the fairly rare and less predictable nature of heating in the summer, which would make it less likely that weekday/weekend-of-sample fixed effects would capture this temperature-dependent variation. The smallest temperature coefficients are on the Coast, with -0.51% [-0.61%, -0.41%] and -0.40% [-0.49%, -0.31%], respectively. This unintuitive apparent temperature response and the relatively small magnitude of temperature coefficients in all of these regressions, is likely because the fixed effects absorb much of the temperature-dependent variation in electricity consumption. See the SI, Section 7.5 for all temperature coefficients from Eq. 4.

Table 4 The time-invariant effects of CARE estimated in Eq. 5, shown in **Table 4**, are similar in magnitude to the average of the intra-day estimates. Temperature response coefficients are slightly larger but generally have the same order of magnitude, further suggesting that the fixed effects used in Eq. 4 likely capture some of the temperature response that has a predictable intra-day pattern. Weekday-of-sample and weekend-of-sample fixed effects still likely capture some seasonally-dependent temperature response.

Table 4. Estimated time-invariant effect of CARE on household electricity consumption by region, using Eq. 5. CARE coefficients are similar to the average value of the hourly coefficient in the corresponding scenario. Note that households are more responsive to changes in temperature during the summer and appear to be equally responsive to heating and cooling. Coefficients are roughly interpretable as a percent increase associated with enrollment in CARE or, for temperature coefficients, per °F of temperature deviation from 65°F. Weekday- and weekend-of-sample fixed effects likely capture some of the seasonally-dependent temperature response.

[%]	Full sample	Summer	Winter	Coast	Hills	Valley
CARE	12.7% [10.2, 15.1]	14.4% [11.0, 17.9]	12.3% [9.7, 14.9]	9.7% [3.8, 16.0]	14.9% [10.7, 19.3]	11.9% [8.6, 15.3]
Temp ≤ 65°F	2.14% [2.08, 2.20]	4.47% [4.39, 4.55]	-0.32% [-0.36, -0.28]	-0.32% [-0.42, -0.23]	0.03% [-0.05, 0.11]	1.21% [1.15, 1.27]
Temp > 65°F	2.53% [2.48, 2.59]	4.08% [4.01, 4.15]	0.01% [-0.03, 0.05]	-0.23% [-0.32, -0.14]	0.29% [0.21, 0.36]	1.71% [1.64, 1.77]

3.4.1. Elasticity estimation

We use the above estimates to compute the implicit price elasticity of electricity demand. Lacking detailed rate information, we assume all CARE households receive PG&E's average discount of 42% [39]. Thus, the Full Sample time-invariant effect of 12.7% [10.2%, 15.1%] translates to a price elasticity of electricity demand of -0.30 [-0.24, -0.36]. This is within the range of estimates for the short-run and long-run price elasticity of demand. It is slightly below, but not statistically distinguishable from the median short-run estimate from the literature of -0.35[41], and well below the median estimate of the long-run price elasticity of demand of -0.81[41].

3.4.2. Climate, human health, and electric power system externalities

CARE enrollment is associated with substantial incremental electricity consumption, resulting in associated climate, human health damages and capacity costs. However, note that the benefits of affordable access to energy services, such as improved physical and mental health and reduced reliance on high-interest short-term loans, are not quantified in this study and may well outweigh these costs [69,70]. The time-invariant full sample increase in electricity consumption of 12.6% [9.2%, 16.0%] for 34% of PG&E's customer base corresponds to an annual increase in demand of 1.48 TWh [1.20 TWh, 1.77 TWh] and 169 MW [137 MW, 203 MW] of incremental peak capacity [39].

This increase in electricity consumption is associated with climate and human health damages in the tens of millions of dollars per year. **Figure 8**A estimates climate damages associated with a time-invariant elasticity at $31.3M [$25.2M, $37.4M], with CO_2 emissions priced at $40/t($CO_2$). **Figure 8**B shows human health damages in the Full Sample at $14.0M [$11.3M, $16.8M] from human health effects of criteria air pollutants assuming a value of statistical life of $6M. However, the magnitude of these damages does not vary significantly if intra-day elasticities of demand are used instead of time-invariant elasticities.

In all cases, the increases in climate and human health damages are statistically significant in the sense that the confidence intervals do not contain zero. However, using intra-day or time-invariant elasticities results in essentially identical estimates of mean damages, within 2.5%, with confidence intervals spanning 39-126% of the mean estimate for time-invariant estimates and 43-228% for intra-day estimates. The magnitude of damages in **Figure 8**B is weighted by the fraction of the population, assumed to be evenly distributed through the three regions, and the

fraction of the year simulated, 171 days for summer (June-September) and 248 days for winter (October-May).

Capacity costs shown in **Figure 8**C are comparable to climate and human health damages, with a maximum value of $37.7M [$24.7M, $49.0M] based on the estimated effect of CARE in the Summer from 6pm-9pm. Using the Full Sample and a time-invariant estimate, costs are $28.8M [$23.2M, $34.4M]. In no case is there a statistically significant difference between capacity costs using a time-invariant or intra-day elasticity. The mean estimates differ by as much as $3.8M in the Summer case and by as little as $0.4M in the Central Valley case, representing a change of 12% and 5% compared to the time-invariant case, respectively. These changes are still within the confidence intervals of the intra-day and time-invariant estimates, which span 39%-126% and 43%-179% of the mean estimate, respectively.

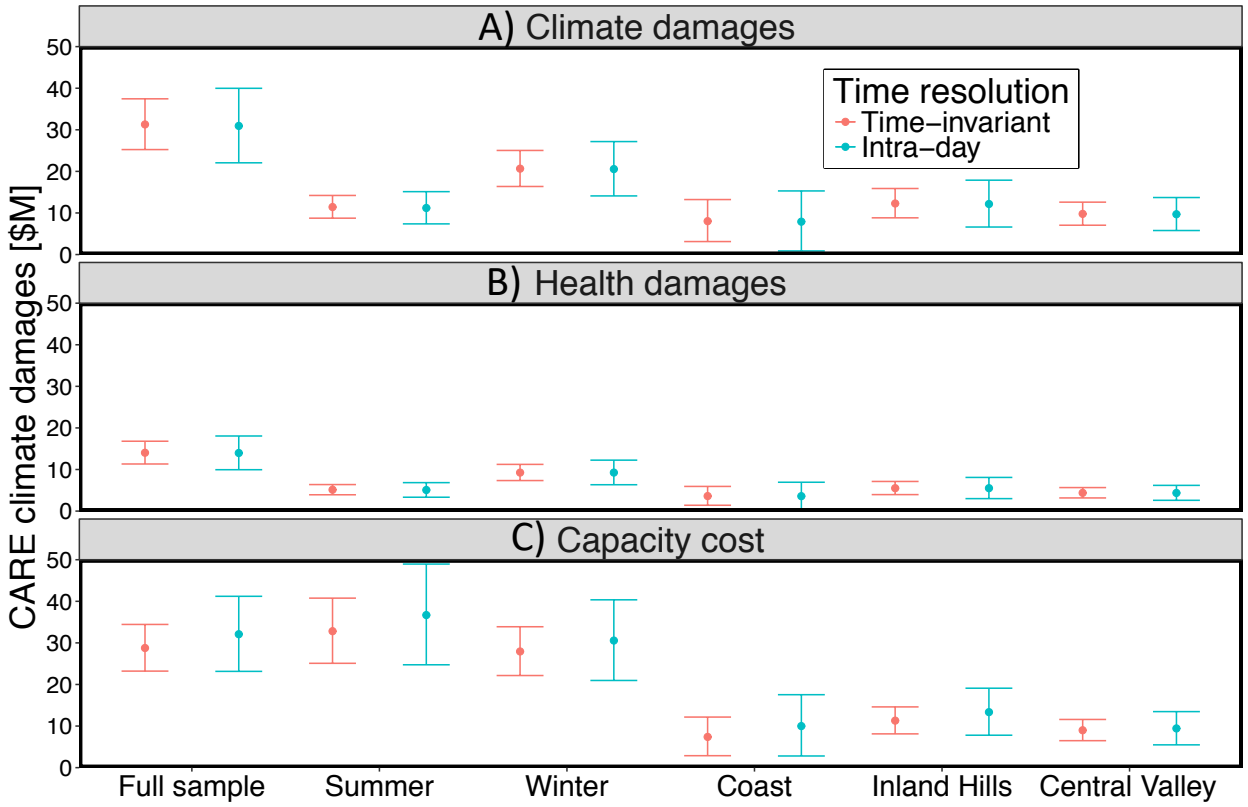

Figure 8. Comparison of societal costs associated with A) marginal climate damages from greenhouse gas emissions, B) human health damages from criteria pollutant emissions, and C) increases in peak electricity consumption attributable to the estimated increase in electricity consumption due to CARE during 2011, assuming a social cost of carbon of \$40/t(CO_2), a \$6 million value of statistical life, and a capacity cost of \$170/kW-yr [66], described in Section 3.3.5. Regional estimates consider consumption increases only in the regional population. Capacity costs assume the increase from 6pm-9pm translates to an increase in system-wide peak demand. Costs are on the order of millions to tens of millions of dollars in each scenario. Red and blue dots represent the mean and 95% confidence interval using intra-day or time-invariant estimates of the consumption increase due to CARE enrollment. In all cases, the mean estimates are close and the confidence intervals strongly overlap. Climate and human health damage estimates using intra-day elasticity estimates are slightly below those using time-invariant estimates, within \$400,000 and \$75,000, respectively, within 2.5% of the time-invariant mean estimates. Capacity costs using intra-day elasticity are \$0.4-3.8M, respectively, representing an increase of 5-12%.

3.4.3. Robustness checks

We estimate the robustness of the regression results using several subsamples of the data and alternative regression specifications.

We apply Eq. 4 to the following subsamples of the data: an event study that includes only households that enroll in CARE with at least 90 days of pre- and post-enrollment data; Only households that either enroll during the study period or never enroll in CARE, excluding those enrolled in CARE throughout the full study; Only weekdays; Only weekends; Only the years 2010 and 2011, aiming to exclude the effects of the 2008 Financial Crisis and the Great Recession; Only households above the 5^{th} percentile or below the 95^{th} percentile of average household electricity consumption, aiming to test the extent to which particularly high or low users drive the results; Households with median census block incomes in the bottom 40% or top 60% of the sample; Excluding perpetual CARE enrollees; and including only households that are always or never enrolled in CARE. These specifications are described in further detail in the SI, Section 7.5.3.

We also apply three alternative regression specifications: Eq. 4 using untransformed electricity consumption as the dependent variable, replacing ln(kWh) with kWh; Eq. 4 including intra-day controls for enrollment in other utility programs; A monthly specification using an aggregated version of the dataset to simulate traditional billing analysis.

We find that the main results are robust across all of the above specifications, with a baseline increase on the order of 10% and significant intra-day increases for at least one three-hour period between 3pm and 9am except for cases with only higher-income neighborhoods or including only households always or never enrolled in CARE.

Estimates of the baseline increase in electricity consumption vary from 10%, for households from poorer neighborhoods, to 14% for the monthly panel, with standard errors on the order of 2% in all cases.

Almost all variations on Eq. 4 reproduce increases in electricity consumption of roughly 2-3% for the period from 6pm-9pm. The major exceptions are the event study, which includes only households that enroll in CARE only once and have 90 days of pre-enrollment and post-enrollment data, and the cases that exclude households perpetually enrolled in CARE or include only households perpetually enrolled or never enrolled in CARE. In the event study, the baseline value is similar to the main results, 11.2% [7.5%, 14.9%], but intra-day increases are larger, with a significant increase of 3.6% [1.2%, 6.0%] from 3pm-6pm, 6.0% [3.9%, 8.1%] from 6pm-9pm, and 4.4% [3.0%, 5.8%] from 9pm-12am. This event study formulation considers only households enrolled in CARE and thus may miss trends controlled for in specifications that include all households. The case that excludes perpetual enrollees has a larger baseline increase of 18.2% as well as larger incremental evening increases of 4.1% [2.4%, 5.9%] from 6-9pm and 3.5% [2.3%, 4.7%] from 9pm-12am. The case that includes only households that are perpetually enrolled or never enrolled captures any differential secular trends between CARE and non-CARE households. This case shows a smaller baseline increase of 7.7% [0.02%, 15.9%] with no significant intra-day variation. This suggests that there may be secular differential trends between CARE and non-CARE households, but the observed intra-day effects likely correspond to behaviors associated with new CARE enrollment. Thus, the inclusion of perpetual enrollees in the main regression specification may have a downward bias on the estimated magnitude of intra-day effects.

Note that controlling for hourly effects of enrollment in other programs does not substantially change the results compared to the main analysis.

The linear formulation, which uses un-transformed electricity consumption as the dependent variable, also shows a baseline increase in electricity consumption of 0.75 [0.55, 0.95] kWh per three-hour period, or 0.25 [0.18, 0.23] kW, with significant increases from 3pm-12am, with a maximum increase from 6pm-9pm of 0.13 [0.08, 0.18] kW. For the average household in our sample, this constitutes a 2.3% [1.7%, 3.0%] increase in electricity consumption for the hours of 6pm-9pm. Note that the distribution of household electricity consumption is approximately lognormal, meaning that this linear specification is likely highly sensitive to changes in consumption for high-consumption households. See the SI, Section 7.3 for further discussion of the distribution of electricity consumption.

See the SI, Section 7.5 for further full regression results.

3.5. Conclusions and policy implications

Policy makers across the world are currently weighing options for deep decarbonization of the electric power system to prevent the worst effects of climate change. Such a transition likely requires substantial costs, at least in the short-run, raising equity concerns surrounding issues of who should bear those costs. Low-income electricity subsidies are one option available to policy makers to address equity concerns in an energy transition.

This analysis suggests that a major low-income electric subsidy program is, unsurprisingly, associated with an increase in electricity consumption in line with existing estimates of the price elasticity of demand.

We estimate annual climate and human health damages at roughly $31.3M [$25.2M, $37.4M] and $14.0M [$11.3M, $16.8M], respectively, with peak demand externality costs associated with

this increase in electricity consumption at roughly $28.8M [$23.2M, $34.4M]. This represents 2%, 4% and 5% of PG&E's 2012 expenditures on the CARE program, respectively, totaling 11% of expenditures [39].

Climate and human health damages from California's electric power sector have fallen rapidly in the past decade [71] and will likely decline further as renewable electricity continues to displace fossil fuels. However, increasing levels of variable renewable resources such as solar photovoltaic and wind electricity, are already placing strain on California's electric power system during shifting peak times [72].

Intra-day estimates find statistically significant variation in the response of households to subsidy enrollment. However, estimates of climate, human health, and peak demand externalities using intra-day or flat price responsiveness estimates both have similar mean values and strongly overlapping error bars.

Thus, policy makers and electric transmission and distribution system operators pursuing a low-carbon energy future can likely gain sufficient insight into the secondary effects of low-income electricity subsidies using time-invariant estimates of the price elasticity of demand, without accounting explicitly for intra-day variation in customer sensitivity to prices.

In addition, the human health effects of increased electricity consumption may be borne disproportionately by low-income households. These effects will likely be worse in areas with more pollution-intensive electric grids. This can be mitigated through policies to reduce air pollution, which California is already pursuing. This also highlights the potential role of energy efficiency policies, such as California's Energy Savings Assistance Program, in helping low-income households achieve a targeted level of energy services while decreasing or reducing growth in electricity consumption [39]. However, the efficacy of energy efficiency programs can be

difficult to measure because of uncertainty surrounding the counterfactual case in which there was no intervention. Econometric estimates of the US Weatherization Assistance Program, which provides heating efficiency measures to low-income households, suggests that engineering estimates substantially over-predicted realized savings [73]. Thus, any energy efficiency program intended to replace a subsidy must be subject to rigorous measurement and verification to ensure that households indeed gain increased access to the desired energy services.

The agreement of our estimates with the price elasticity of demand literature supports the validity of the estimation strategy employed in this study. However, given the innate limitations of this observational approach, selection bias and other confounding factors could bias these estimates upward or downward.

Our estimates are based on household response to a low-income electricity subsidy, which may also affect electricity consumption through social norms or other non-price mechanisms. A household receiving subsidized electricity with the explicit aim of ensuring access to energy services may increase consumption more than if they simply received a lower electricity price, e.g. due to a perceived expectation that they should increase electricity consumption. The opposite may also be true, e.g. if recipients wish to minimize their reliance on social assistance. Thus, these results may not be representative of the response of low-income households, or for that matter higher-income households, to decreases or increase in electricity price that are not mediated through a subsidy. These results also do not necessarily generalize to intra-day responses to changing price signals throughout the day, such as time-of-use pricing, as residential customers appear to respond to average electricity price, i.e. bills, rather than marginal prices [74].

States and utilities could likely benefit from empirically estimating price responsiveness for their own populations. A wide range of estimates of the price elasticity of demand exist in the literature. This is in part because of the uncertainty inherent in the available estimation strategies but. More fundamentally this is because people in different places in different contexts use electricity differently. It is unclear whether California policy makers and utilities explicitly model this effect and, if so, what price elasticity of demand they use.

Note that electric utilities are in a unique position to experimentally measure the price-responsiveness of their customers through randomized controlled trials, e.g. by randomly giving discounted electricity to some customers and comparing their consumption to that of similar customers.

The expanded deployment of inflexible resources such as wind and solar, places increasing value on understanding the flexibility of electricity demand. Advanced metering infrastructure offers policy makers and electric utilities a unique opportunity to better understand electricity demand and further integrate these insights into a rapidly transitioning electric power system. This in turn could ensure that a suite of policies succeeds in providing access to affordable, clean, low-carbon energy for all.

3.6. Data availability statement

Household electricity consumption and other household-level data were provided by Pacific Gas and Electric Company through the Wharton Customer Analytics Initiative via a non-disclosure agreement. Temperature data were provided by the National Oceanographic and Atmospheric Administration and are publicly available [64]. Census block demographic information are from the 2010 US Census and are publicly available [75].

4. Chapter 4. Electrofuel synthesis from variable renewable electricity: An optimization-based techno-economic analysis

4.1. Abstract

Policies consistent with a global warming limit of 2°C or lower will likely require substantial cuts in greenhouse gas (GHG) emissions from sectors such as aviation that currently have limited ability to substitute away from liquid fuels. This analysis characterizes the economic viability of *electrofuels*, synthesized from CO_2 from direct air capture (DAC) and hydrogen from electrolysis of water, powered primarily by solar or wind electricity. This optimization-based analysis compares the GHG mitigation cost-effectiveness of electrofuels with continued use of petroleum-based fuels offset by DAC with sequestration (DACS) under a range of techno-economic assumptions. Using today's technology, hydrocarbon electrofuels are likely to cost upwards of $12/gallon of gasoline equivalent (GGE) and are thus a substantially more expensive mitigation strategy than DACS. However, in a scenario in which cost targets for the next decade are met, electrofuels powered by the world's best solar resources could achieve mitigation cost parity with DAC if pre-tax petroleum fuel prices reach $3/GGE. Electrofuels could be the preferred strategy here if CO_2 sequestration costs are unexpectedly high. In a longer-term breakthrough scenario, the cost of electrofuels could fall below $3.00/GGE, making them a preferred mitigation strategy to DACS and potentially competitive with petroleum fuels. Electrofuel cost is most sensitive to the capital cost the DAC, electrolyzer, and renewable electricity, confirming their importance as priorities for research, development, and deployment (RD&D). Due to operational flexibility afforded by storage or supplementary natural gas or grid electricity interconnections wind brings cost reductions over solar only if it can produce bulk

electricity at a comparable cost with solar, despite its higher capacity factor. Without this operational flexibility, costs rise by more than 80%. This points to some less intuitive RD&D priorities, such as metallic phase change materials capable of storing heat above 900°C and low-cost, seasonal CO_2 storage.

4.2. Introduction and background

Avoiding 2°C or more of global warming likely requires the complete elimination of net carbon dioxide (CO_2) emissions from the global economy within four to seven decades [1]. This would require decarbonization of difficult-to-mitigate sectors, including aviation, long-distance road transportation, and ocean shipping, representing roughly 6% of global energy-related greenhouse gas (GHG) emissions [76].

These applications currently use energy-dense liquid fuels to operate capital-intensive fleets of airplanes, trucks, and ships with asset lifetimes of 20 years or more. Thus, even with research and development breakthroughs in electric or hydrogen propulsion systems, a rapid transition to these fuels would likely require premature retirement or retrofitting on a massive scale. As a result, a carbon-neutral hydrocarbon fuel with the ability to "drop in" to existing infrastructure could greatly reduce the cost of deep decarbonization in these sectors and buy time for a smooth transition to lower-cost technologies if they arise.

All hydrocarbons emit CO_2 when combusted, so achieving carbon neutrality requires that the carbon embedded in a hydrocarbon be sourced from the atmosphere. Biofuels achieve this by using carbon captured by plants, algae, or other living organisms. However, biofuel production generally results in substantial net life-cycle emissions [77], thus only partially mitigating emissions from displaced petroleum fuels.

Electrofuels are hydrocarbon or oxygenate fuels derived from CO_2 and H_2 from electrolysis of water through processes that rely primarily on electrical energy [78]. Potential products include methanol, dimethyl ether, methane, and Fischer-Tropsch liquid hydrocarbons such as gasoline, diesel, and jet fuel [78]. Techno-economic estimates of the cost of electrofuels range from $0.40-$135/GGE, with harmonized ranges of $7.70-10.80/GGE using today's technology and $6.20-8.10/GGE in 2030 [78]. Most analyses assume CO_2 is sourced at low cost from industrial waste, such as biofuel production or cement manufacturing [78].

Electrofuels only achieve substantial net CO_2 emissions reductions compared to petroleum-based fuels if they are powered by very low-carbon energy. The cost of variable renewable electricity, particularly solar photovoltaics (PV), has fallen precipitously in the past decade, making it one of the cheapest sources of electricity available [5,6]. Fasihi et al. [79–81] estimate the cost of electrofuel production from combined wind and solar installations using CO_2 from direct air capture (DAC) with point estimates of techno-economic parameters and an aggregated representation of variability in renewable electricity, finding costs on the order of $4/GGE.

This analysis uses an optimization-based techno-economic framework to characterize the range of possible costs of electrofuel production using CO_2 from DAC, powered primarily by variable renewable electricity. In addition to estimating the future cost of electrofuel production, this approach identifies system components and characteristics for which additional research, development, and deployment (RD&D) is likely to yield the greatest reductions in system cost. In addition, this work compares the cost of electrofuel production to the cost of simply continuing to use petroleum-based liquid fuels and using DAC with sequestration (DACS) to offset a corresponding amount of CO_2. Alternative mitigation strategies for deep decarbonization of hydrocarbon liquid-dependent applications, not modeled in this paper, include advanced

biofuels such as camelina-based jet fuel [82], electrification, hydrogen fuel cell designs, and carbon dioxide removal technologies such as biomass energy with carbon capture and sequestration, enhanced weathering, coastal blue carbon, afforestation and other forms of land management [83]. Biofuels capable of acting as jet fuel generally have substantial life-cycle greenhouse gas emissions, ranging from roughly 10% to 75% of petroleum jet fuel emissions. Staples et al. estimate that such fuels could avert at most about 70% of global aviation emissions, requiring that biofuel constitute over 85% of total jet fuel [84]. Thus, even with substantial cost reductions, a biofuel-based mitigation strategy for aviation would still require substantial carbon dioxide removal to achieve net-zero emissions. Hydrogen fuel cell airplanes are still in early stages of development [85]. Forms of carbon dioxide removal other than DAC may be available at lower cost, but their capacity is limited and may be needed for other applications [83].

4.3. Methods and data

4.3.1. System components

This analysis models the cost and operation of an electrofuel production system comprised of variable renewable electricity production, a DAC system heated either by an electric kiln or by natural gas, an electrolyzer, and Fischer-Tropsch hydrocarbon fuel synthesis infrastructure, transporting liquid fuel to market through a pipeline. To manage variability in electricity supply, the system can build storage of electricity, modeled as a lithium-ion battery system; heat, modeled as a molten salt, phase change material, or supercritical CO_2 system; hydrogen, either in a pressurized above-ground tank or in an underground formation; CO_2 in a pressurized tank; and fuel in a tank. The system can also build natural gas and grid electricity interconnections, paying a carbon price for the associated fossil CO_2 emissions. **Figure 9** shows a diagram of the electrofuel production system modeled here. Each component is introduced below.

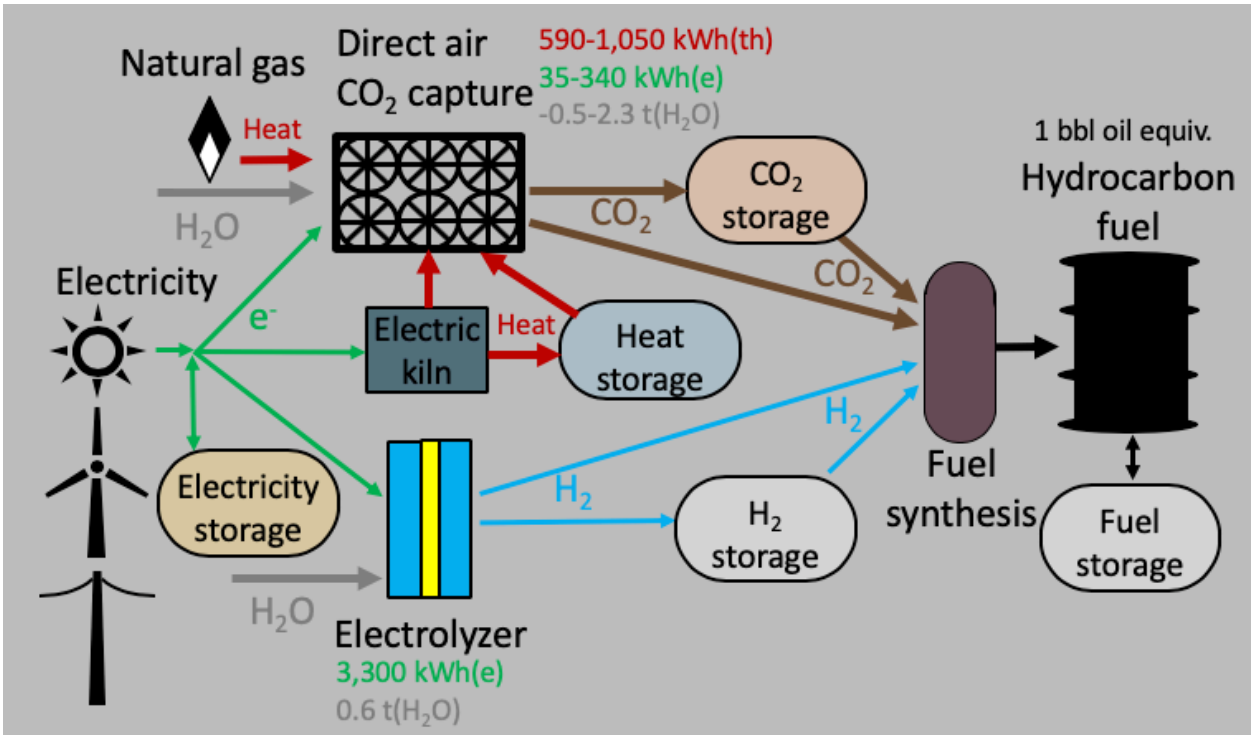

Figure 9. Process diagram for electrofuel production. The simulated electrofuel production system uses solar or wind electricity, supplemented with grid electricity and natural gas, to power an electrolyzer and a direct air CO_2 capture system, heated by a kiln powered by electricity, natural gas, or both. CO_2 and hydrogen are then converted into hydrocarbon fuels, such as diesel and jet fuel, through a fuel synthesis process, in this case the Fischer-Tropsch process. The fuel is transported to market through a pipeline. The system optimizes annual operation of all components, including storage of electricity, heat, carbon dioxide, hydrogen, and liquid fuel, in four-hour increments. Energy and material input requirements shown produce one barrel of oil equivalent of hydrocarbon fuel, such as jet fuel, diesel, or gasoline, assuming the electrolyzer and fuel synthesis are both 70% efficient.

Each technology is modeled using three sets of techno-economic parameters. The first case estimates the cost of building an electrofuel production system using the best available technology **Today**. The second is a **Next-decade** scenario using parameters that the literature suggests could be achieved within roughly a decade or less with significant RD&D efforts. The final case considers the potential for a **Breakthrough** in each of the component technologies in the next few decades. The Today case considers only a low-temperature DAC system, as high-

temperature systems have not yet been built commercially. The Next-decade and Breakthrough cases consider both low-temperature and high-temperature DAC systems. Parameters are based on the available literature, with values described in detail below and in Section 4.3.4.7.

4.3.1.1. Renewable electricity generation

The cost of wind and particularly solar PV electricity have fallen rapidly in the past decade [5,6], making them some of the lowest-cost sources of bulk electricity. Thus, this analysis considers both solar and wind electricity as a primary energy source for electrofuel production.

This analysis simulates renewable electricity production profiles from the System Advisor Model (SAM), produced by the National Renewable Energy Laboratory (NREL) [86]. Both solar PV and wind are simulated in locations within the United States with very high-quality resources: Tucson, Arizona and southern Wyoming, respectively, with capacity factors of 29.3% and 57.5%.

The production characteristics of the solar PV system follow default specifications from SAM [86] for a utility-scale PV installation with single-axis tracking. However, we assume the alternating current (AC) inverter is sized to the full rated capacity of the generation. Utility-scale solar PV installations often undersize the inverter to achieve a DC to AC ratio of 1.2 to reduce grid interconnection costs [86].

The production characteristics of the wind turbine also follow default specifications from SAM for a utility-scale installation [86]. The 1,500 kW GE 1.5sle turbines reach their maximum power plateau at wind speeds of 15-25 m/s, shutting off at higher wind speeds for safety [86].

For solar PV, the capital cost in the Today case is based on the installed cost of $900/kW(AC) ($600/kW(DC)), achieved in the best locations in the United States in 2017 [6]. This may be slightly generous due to the assumption of single-axis tracking. However, it may also be

conservative given the rate of cost declines in the past decade. Operations and maintenance (O&M) costs are set at 1% of capital cost per year in all cases [87]. With the Today case's 10% weighted average cost of capital (WACC), this corresponds to a levelized cost of electricity (LCOE) of $39/MWh(e). The Next-decade cases assume a cost structure similar to that in Mexico, where the average power purchasing agreement was $20.8/MWh(e) in Q4 of 2017, and where installers regularly sign unsubsidized power purchasing agreement contracts below $20/MWh(e) [88]. With the Next-decade case's 8% WACC, this amounts to a capital cost of $500/kW(AC), with an LCOE of $18.8/MWh(e). The Breakthrough case considers the possibility of further advances in solar PV technology, perhaps including high-efficiency DC-DC inverters to power the electrolyzer and electric kiln, both of which can use DC directly. This case assumes a capital cost of $400/kW with a WACC of 5%, which equates to an LCOE of $11.8/MWh(e).

For wind, the capital cost in the Today case is $1,500/kW(AC), in the lower range of installed cost in the US Midwest region in 2017 [87]. O&M costs in all cases are 2.5% of capital cost per year [87]. At a 10% WACC, this amounts to an LCOE of $40.2/MWh(e). In the Next-decade cases, capital cost falls to $1,250/kW(AC), or $29.4/MWh(e) at an 8% WACC, roughly matching the lowest-cost installations in the US in 2017 [5]. In the Breakthrough cases, capital cost falls to $1,000/kW(AC), or $19.0/MWh(e) at a 5% WACC, matching the lowest installed costs on record in 2006, before materials prices and other factors led to substantial cost increases on a kW(e) basis [5].

4.3.1.2. Electrolyzer

Electrolysis of water uses electricity to convert water into hydrogen gas and oxygen. Alkaline electrolysis has been industrially available for decades and has long been used to manufacture chlor alkali, an $8bn dollar industry in North America [89].

The prospect of using renewable hydrogen as a low-carbon fuel for transportation and other applications has spurred renewed interest in electrolysis. Emerging proton exchange membrane (PEM) and solid oxide electrolyzers, as well as next-generation alkaline electrolyzers, are projected to achieve substantial cost reductions and efficiency gains [78,90,91].

This analysis considers both alkaline and PEM electrolyzers. Solid oxide electrolyzers are likely to have higher capital costs [78,91], making it more difficult to justify intermittent operation.

Both alkaline and PEM electrolyzers typically operate at temperatures from 50-80°C [78,91], although alkaline electrolyzers are capable of operating above 100°C [91]. Solid oxide electrolyzers operate at much higher temperatures on the order of 850°C [92].

Produced oxygen can be sold or reused in the fuel synthesis infrastructure [93,94] or in an oxygen-fired natural gas kiln to produce process heat for the DAC system, producing a pure stream of CO_2 from combustion of natural gas that can either be sequestered or used as a feedstock for electrofuel synthesis.

For current alkaline and PEM electrolyzers, the stack must be replaced after 60,000-90,000 hours of operation, 7.6-11.4 years at a 90% capacity factor, incurring an additional 50-60% the capital cost of the electrolyzer at each replacement [91]. Stack replacement costs are incorporated explicitly into the Today case using the conservative assumption that the stack must be replaced every 9 years regardless of the level of utilization. At a 10% WACC, two stack replacements

over the system's 25-year lifespan, at 50% of upfront capital cost each, increase the effective net present value of the capital cost of the electrolyzer by 37.5%.

The Today case capital cost for alkaline electrolyzers is drawn from a manufacturer estimate for 2015, €$_{2012}$760/kW(e), $872.4/kW(e) in 2017 dollars [91,95,96]. Including the stack replacements above, this amounts to a present value of roughly $1,200/kW(e). The Next-decade and Breakthrough electrolyzer capital cost estimates apply the same escalation factor to a 2025 low-end cost estimate for alkaline electrolysis, €$_{2012}$370/kW(e), and a 2030 low-end cost estimate for PEM electrolysis, €$_{2012}$250/kW(e), resulting in effective capital costs of $584.1/kW(e) and $394.7/kW(e), which are rounded to $580/kW(e) and $390/kW(e). Using the same escalation factor does not fully capture the complexities of future technological advances, but the lower discount rate in the Next-decade and Breakthrough cases at least partially offsets the benefit of stack lifetime extensions.

4.3.1.3. Direct air CO_2 capture

Existing and proposed DAC technologies include both low-temperature amine systems, such as those used by Climeworks and Global Thermostat and high-temperature hydroxide-based systems, such as that used by Carbon Engineering [97]. This analysis does not model a moisture-swing adsorption system, such as that being developed by Infinitree [98].

All of the high- and low-temperature DAC systems considered use an air contactor, similar to a horizontal cooling tower, to expose atmospheric CO_2 to a sorbent, either an amine or a hydroxide. Once CO_2 is adsorbed by the sorbent, the sorbent solution undergoes an energy-intensive regeneration process through which it releases pure CO_2, which can then be used or sequestered. In low-temperature amine systems, this regeneration takes place through a temperature swing. Relatively low-temperature heat, approximately 100°C, is used to regenerate

the amine [98]. This temperature swing, from low to high temperatures, takes at least several hours [99]. Thus, the heat demand and CO_2 production of a low-temperature system operating at full capacity are not constant over time. The process also requires electricity, described below [98].

4.3.1.4. Carbon Engineering's high-temperature system

Carbon Engineering's air contactor uses an aqueous potassium hydroxide (KOH) sorbent to capture CO_2, forming K_2CO_3 and H_2O [97]. This K_2CO_3 then reacts with calcium hydroxide, $Ca(OH)_2$, in a pellet reactor to regenerate KOH, which is recycled, creating $CaCO_3$, which is dried into pellets. $CaCO_3$ is then heated to 900°C in a calciner, where it breaks apart into CaO and CO_2. CaO is then mixed with water in a slaker to reconstitute $Ca(OH)_2$, which is recycled into the pellet reactor to close the loop. The entire process requires 5.25 GJ of heat, including 4 GJ at 900°C for calcination, and 366 kWh(e) of electricity, much of which is used to compress CO_2 to 151 bar for pipeline transport [97]. A low-pressure configuration that produces CO_2 at 1 bar uses only 77 kWh(e) of electricity, implying that 289 kWh(e) are used for compression of CO_2. Carbon Engineering's proposed first-generation design uses oxygen-fired natural gas for all process heat requirements, co-capturing the resulting pure CO_2 stream with the CO_2 released during calcination [97]. As a result, the use of natural gas for process heat adds 0.3 t(CO_2) to each t(CO_2) captured from the atmosphere [97].

High-temperature electric kilns capable of reaching 900°C exist commercially but are largely for specialized manufacturing applications. Electric kilns on the hundred megawatt scale necessary for large-scale DAC do not appear to exist commercially. However, given the high efficiency of resistive heating and the existence of inexpensive industrial insulators, such products will likely materialize rapidly given market demand. Indeed, clever engineers can likely design kilns capable of using heat from both oxygen-fired natural gas, or even hydrogen, and electricity

simultaneously. For this reason, we consider a high-temperature electric resistive kiln capable of cofiring with oxygen-fired natural gas, using conservative cost and efficiency assumptions based on the unit cost of existing small-scale kilns described in Section 4.3.1.6.

Carbon Engineering's cost estimates for high-temperature DAC systems range from $94-232/t($CO_2$) [97], depending on the level of previous deployment, financing assumptions, system configuration, and other factors. This is substantially lower than previous cost estimates for high-temperature systems, due both to substantial differences in design choices and techno-economic assessment methods [100,101]. This corresponds to capital costs ranging from $694-1,046/(t($CO_2$)/yr) Keith et al. [97] estimate O&M costs for a system operating at a 90% capacity factor at roughly 4%/yr of capital cost.

The Next-decade case uses the baseline capital cost estimate for an early-build high-temperature DAC system, costing $1,146/t($CO_2$/yr). This system has heat and electricity demand of 5.25 GJ(th) and 366 kWh(e), respectively [97]. The system produces CO_2 at 151 bar, so CO_2 can be stored without additional compression energy. The Breakthrough case assumes a mature low-pressure system from Keith et al. [97], with capital cost of $694/t($CO_2$/yr). Heat demand remains 5.25 GJ(th), but electricity demand falls to 77 kWh(e) [97]. Thus, CO_2 storage requires an additional 289 kWh(e)/t(CO_2) for compression. O&M costs are modeled at 4% of capital cost per year, based on the O&M cost estimates from Keith et al. [97], which range from 3.8%-4% of capital cost per year. High-temperature systems are not modeled in the Today case, as they have not yet been built commercially.

4.3.1.5. Climeworks' low-temperature system

The cost of current low-temperature DAC systems is estimated at $500-600/t($CO_2$) [102]. Climeworks states that it has a clear engineering pathway to $200/t($CO_2$), with an aspirational

target of \$100/t(CO_2) [102]. Global Thermostat claims long-term costs could fall as low as \$15-50/t(CO_2), but have not yet described the details of how such costs could be achieved [103]. Climeworks' current technology, used in the Today case, is estimated to consume 2,200 kWh(th), 7.9 GJ(th), of heat and 700 kWh(e) of electricity [104]. Climeworks' long-run target, used in the Next-decade case, is 1,600 kWh(th), 5.8 GJ(th), of heat and 400 kWh(e) of electricity [104]. Global Thermostat's competing low-temperature technology, assumed in the Breakthrough case, claims only 4.4 GJ(th) of heat demand and 160 kWh(e) of electricity demand [98]. Because low-temperature systems can use low-grade waste heat, this analysis assumes the above cost estimates assume free heat. Climeworks claims a current lifetime of 20 years, used in the Today case, with a target of 30 years, used in the Next-decade and Breakthrough cases [104].

Estimates of the capital cost of low-temperature DAC systems are based on the above estimates of levelized cost and projections stated by manufacturers Climeworks and Global Thermostat to the media. Annualized capital and fixed O&M costs are derived from the above levelized cost estimates after energy costs are subtracted out, using Eq. 6. Climeworks' long-term levelized cost estimates likely assume low-temperature heat demand is met using waste heat or another low- or zero-cost heat source. As a result, this analysis assumes free heat in the capital cost calculations. These calculations assume energy demand for the Climeworks systems in the Today and Next-decade cases, with electricity cost at \$0.10/kWh(e), in line with industrial electricity rates in Europe, where Climeworks is based [105]. Electricity rates fall to \$0.05/kWh(e) in the Breakthrough case, closer to industrial rates in low-cost locations the United States, using energy demand projections from Global Thermostat [106]. This analysis assumes O&M costs of 4%/yr, based on O&M estimates for high-temperature systems, even though amine replacements may raise costs further. However, this does not affect the results of this analysis because the

optimization uses annualized capital and fixed O&M costs interchangeably and these capital costs are back-calculated to reproduce the above levelized cost statements.

$$Capex = (LevelizedCost - ElectricityPrice*(Electricity/t(CO_2)))/(CCF + O\&M\%) \quad (6)$$

Where CCF, the capital charge factor is computed using Eq. 7, with the lifetimes described above and the corresponding WACC from each scenario (10%, 8%, 5%).

$$CCF = WACC/(1 - (1 + WACC)^{-lifetime}) \quad (7)$$

The Today case assumes a levelized cost of \$600/t(CO_2), based on the cost of current commercial systems, falling to \$200/t(CO_2) in the Next-decade case, based on Climeworks' expected cost reductions for its current system, and \$100/t(CO_2) in the Breakthrough case, based on Climeworks' aspirational target. The resulting capital cost is \$3,029/(t(CO_2)/yr) in the Today case, \$1,118/t(CO_2) in the Next-decade case, and \$788/(t(CO_2)/yr) in the Breakthrough case. These numbers are rounded to \$3,030/(t(CO_2)/yr), \$1,120/(t(CO_2)/yr), and \$790/(t(CO_2)/yr), respectively.

Note that for low-temperature amine-based DAC systems such as those used by Climeworks and Global Thermostat, the process of temperature-swing adsorption of CO_2 implies a temporal profile to heat demand. Wurzbacher et al. suggest that the system takes roughly 120 to 180 minutes to heat from 20°C to 90°C, with the majority of the swing completed in 60 to 120 minutes [99]. This heating releases CO_2 bound to the amine. The system must then cool for some period of time. Thus, heat demand is not constant and would likely require heat storage or

dispatchable heat production, such as natural gas. These heat storage costs, not modeled here, would modestly increase the capital cost of a low-temperature system.

4.3.1.6. Electric kiln

Electric resistive heating is an established technology, used for space and water heating and a variety of other applications. However, electric kilns capable of efficiently generating heat at or above 900°C are primarily available for specialized manufacturing applications and do not appear to be commercially available on scales of 100 MW(e) or more. 900-1,100°C industrial high-temperature electric ceramic tunnel kilns are available for $10,000-120,000 at capacities of 85-150 kW [107]. Using the low numbers from both ranges for a 900°C kiln, this amounts to $125/kW(e). This unit cost would likely decline substantially at the much larger scales required for electrofuel production. This analysis assumes 95% kiln efficiency, due to the nearly 100% efficiency of resistive heating and the maturity of high-temperature industrial insulation in numerous applications [108]. Current specialized industrial electric kilns may operate at lower efficiencies, tolerating higher losses. Efficiency is held constant in all scenarios and varied in the sensitivity analysis.

Balancing the tradeoffs between cost reductions through economies of scale and cost increases due to greater use of insulation, the Today case for a high-temperature electric kiln assumes $125/kW(e), falling to $75/kW(e) and $25/kW(e) in the Next-decade and Breakthrough cases, respectively, assuming economies of scale outweigh additional insulation costs in the long-run. Requiring a high-temperature electric kiln is a conservative assumption for systems that use low-temperature DAC. However even at this level, electric kiln capital cost represents less than 1% of levelized system cost. Kiln lifetime is assumed at 10 years based on discussions with industry. O&M costs are assumed to be 1% of capital cost per year.

4.3.1.7. Fuel synthesis

The Fischer-Tropsch process, patented in the 1920s, was originally used to convert coal and other fossil fuels into liquid transportation fuels [109]. It was used extensively in Germany during the Second World War and by South Africa since the 1980s [109]. As a result, the technology is fairly mature. Fischer-Tropsch plants exist across the world [110].

A reverse water-gas shift (RWGS) reactor converts CO_2 and H_2 into synthesis gas, a mixture of CO and H_2. The process is mildly endothermic, but this heat can often be supplied using waste heat from another source, such as from the Fischer-Tropsch reaction itself [111].

A Fischer-Tropsch reactor takes in synthesis gas, and outputs hydrocarbons with a chain-length distribution that depends on the temperature, pressure, and catalysts used. These hydrocarbons must then be processed to produce liquid transportation fuels. Waxes are converted into shorter-chain hydrocarbons through a hydrocracking process. The product distribution for liquid transportation fuels is fairly flexible. This analysis focuses on jet fuel production. Jet fuel can be as much as 50% of the product or more, with the remainder split between diesel and naphtha, which can be refined into gasoline [81]. The remaining diesel and gasoline could likely be sold at a similarly competitive price for difficult-to-decarbonize applications such as long-distance road transport.

The full fuel synthesis process produces hydrocarbon fuels that embody 65-80% of the input energy from the H_2 [111]. The process is highly exothermic, producing much of the heat required to meet low-temperature DAC heat demand and other heat demand at or below its operating temperature of roughly 120-350°C [81,112].

Cost estimates for this full fuel synthesis apparatus range from $350-1,500/kW(fuel), depending on the scale, configuration, assumed system integration costs, and other factors [78]. O&M cost estimates range from 3-4 % per year of capital cost [81,113].

In the Today case, Fischer-Tropsch fuel synthesis has a capital cost of $800/kW(fuel) based on pessimistic assumptions for a large-scale, 200 MW(fuel) facility from Brynolf et al. [78]. The Next-decade case uses $450/kW(fuel), based on the reference case from Brynolf et al. [78]. The Breakthrough case uses $350/kW(fuel) based on the most optimistic case from Brynolf et al. [78]. O&M costs are set at 3% per year of capital cost [81].

The system has a 70% H_2 to fuel energy conversion efficiency in the Today and Next-decade cases, rising to 75% in the Breakthrough case [78,114]. Each MWh(fuel) requires 0.28 t(CO_2) of input, which is converted at 95% efficiency [78]. Unconverted CO_2 is assumed to be in a solid form akin to biochar, for which the cost of disposal is not considered. As a result, unconverted CO_2 is not released into the atmosphere and does not incur a CO_2 price.

The main analysis assumes that the fuel synthesis can ramp up and down at a rate of 10% of nominal capacity per hour and that each kW(fuel) of ramping incurs an energy penalty equivalent to the H_2 required to produce one kWh(fuel). The sensitivity analysis, described in Section 4.3.10, tests the sensitivity of the cost of electrofuel to these assumptions.

4.3.1.8. Electricity Storage

Fu et al. [7] estimate the current cost of a 60 MW(e), 240 MWh(e) lithium-ion (Li-ion) battery system for a utility-scale solar-plus-storage installation at $380/kWh(e). This estimate includes $22/kWh(e) in sales taxes and $18/kWh(e) in net profit, which are not included for other components in this analysis [7]. Thus, the Today case electricity storage capital cost is $350/kWh(e). Darling et al. estimate long-term Li-ion battery costs at as low as $225 [200,

250]/kWh(e) [115]. Kittner et al. propose a pathway to $100/kWh(e) for Li-ion batteries, which is used as the Breakthrough value [116]. The Next-decade case is roughly halfway between the Today and Breakthrough values at $250/kWh(e).

Efficiency estimates for of utility-scale Li-ion batteries range from 75-90% [117], with operational efficiencies often closer to 70-80% [118]. The Today case assumes electricity storage with 80% efficiency, rising to 85% in the Next-decade case and 90% in the Breakthrough case. O&M costs are estimated at 1% of capital cost per year, based on Zakeri et al. and Cole et al. [119,120].

This analysis assumes a 10-year battery lifetime [115–117]. Battery capital and O&M costs are annualized on a net present value basis without explicit consideration of the number of battery replacements that would be necessary over the full system's 25-year lifetime. If the electrofuel production system is retired before the end of a battery's functional life, the battery can likely be re-sold to recoup stranded costs. The same is true for the electric kiln.

Note that this analysis uses battery costs representive of scales over 100 MWh(e). Reducing scale to 30 MWh(e) could more than double the unit cost [7]. However, linear optimization does not allow cost to scale with capacity. Thus, storage costs represent a lower bound in cases where substantially less than 100 MWh(e) are built.

4.3.1.9. Hydrogen storage

Hydrogen (H_2) can be stored on an industrial scale in pressurized steel tanks or underground in salt caverns, depleted oil and gas wells, and other appropriate natural and man-made formations [121].

Cost estimates for steel tank H_2 storage range from $12.9-32.2/kWh($H_2$) [121], with a US Department of Energy (DOE) 2020 cost target of $500/kg($H_2$), or $15.0/kWh($H_2$) [122]. Geologic storage is less expensive, but its availability is location-specific. The cost of geologic storage

ranges from \$6.1-8.2/kWh(H_2), including compressor costs [121]. Pressurizing H_2 requires roughly 0.036 kWh(e) for tank storage, the Today and Next-decade cases, and 0.033 kWh(e) for geologic storage in the Breakthrough case [121]. Compressor capital cost is included in the capital cost of H_2 storage.

The Today case assumes steel tank storage at \$25/kWh(H_2) [121,122], with a Next-decade value assuming the DOE's target of \$15.0/kWh(H_2) is met [122], and a Breakthrough case using geologic storage at \$6/kWh(H_2) [121].

H_2 storage is assumed to be 99% efficient, not including compression energy. This is consistent with assessments that H_2 storage in structurally sound tanks and geologic formations, e.g. in salt caverns, does not have appreciable leakage concerns [123,124]. H_2 compression itself introduces leakage of roughly 0.05% of total mass flow [125], potentially rendering this a conservative estimate of H_2 storage efficiency.

4.3.1.10. Heat storage

Molten salts and phase change materials (PCM) are commonly used for industrial heat storage above 100°C [126,127]. Molten salt heat storage at up to 550°C has been applied in numerous utility-scale concentrated solar thermal power plants, with cost estimates ranging from \$20-40/kWh(th), with a DOE 2020 target of \$20/kWh(th) [86,128–130]. Such facilities have demonstrated thermal efficiencies as high as 93% [131]. Some concentrated solar thermal power plant designs are able to use molten salt directly as a working fluid for electricity generation, reducing heat exchanger losses. An electrofuel production system would likely require a heat exchanger.

Storing heat for high-temperature DAC requires higher-temperature storage. $Si_{49}Mg_{30}Ca_{21}$ phase change materials capable of storing heat at 865°C are available at an estimated $\$_{2013}52$/kWh(th), or $\$_{2017}53$/kWh(th) [127]. Metallic PCM capable of storing heat above the 900°C required for

calcination in the DAC process exist but are not widely commercially available [132]. In high-temperature DAC cases, we consider $Si_{49}Mg_{30}Ca_{21}$ phase change materials and assume that the heat is boosted to 900°C using a small amount of H_2. Hydrogen flows for this temperature boosting are currently not modeled, meaning that the system currently assumes that heat is stored at 900°C.

The Today case assumes molten salt heat storage at $25/kWh(th) [86,133]. For low-temperature systems, the Next-decade case assumes DOE's 2020 target of $20/kWh(th) for molten salt storage is met [129], using $15/kWh(th) supercritical CO_2 storage in the Breakthrough case [129]. High-temperature systems assume $Si_{49}Mg_{30}Ca_{21}$ phase change material storage at $55/kWh(th) in the Near-term case and reach parity with current molten salt storage at $25/kWh(th) in the Breakthrough case, assuming an alternative material [127]. Energy efficiency is assumed to be 90% in all cases, slightly below the best performance achieved for molten salt systems [131], and is varied in the sensitivity analysis.

4.3.1.11. CO_2 storage

Note that this analysis uses the term "storage" to refer to a system in which CO_2 can be deposited and easily withdrawn and "sequestration" to refer to a system in which CO_2 is deposited with the intention of permanently preventing it from returning to the atmosphere.

CO_2 becomes a dense liquid or supercritical fluid at high pressures [134]. Thus, it is generally cost-effective to store CO_2 in pressurized vessels. Refrigerated liquid CO_2 tanker ships have been proposed and analyzed for transportation of captured CO_2 to offshore sequestration locations [135]. Such systems have leakage rates on the order of 3% over the course of a voyage of a few days or weeks, and thus may not be suitable for cost-effective seasonal storage of CO_2 [135].

CO_2 could also be stored in suitable geologic features. Bulk CO_2 storage systems on the scale of tens of thousands of tons do not commercially exist but could likely achieve further cost reductions through economies of scale.

Appropriately pressurized steel tanks themselves cost roughly \$1000/t(CO_2) [136], while compressors cost roughly \$9000/t(CO_2) [137]. Energy required to compress CO_2 from 1 bar to 151 bar is modeled as 289 kWh(e), based on the difference in the energy requirement between high-pressure and low-pressure DAC systems in Keith et al. [97].

The Today case makes the conservative assumption that low-loss CO_2 storage is prohibitively expensive, setting its cost to \$50,000/t(CO_2). The Next-decade case assumes \$10,000/t(CO_2) for a tank and compressor system. The Breakthrough case considers \$1,000/t(CO_2) for an unspecified system. In open spaces with inexpensive land, a low-pressure system such as an ensemble of durable inflatable plastic cells could dramatically reduce compressor and material cost requirements. O&M costs are assumed at 5% of capital cost per year [136]. CO_2 storage is assumed to be lossless, with 100% efficiency. In practice, any losses could be captured through a secondary containment shell and used in fuel synthesis, which is essentially always in operation.

4.3.1.12. Fuel storage

Hydrocarbon fuel tank costs follow a roughly linear trend with a slope and intercept estimated at \$0.52/gallon and \$331,900/tank, respectively [138]. This fit has support for tanks of 500,000 gallons to 4 million gallons [138]. However, the modeled system does not build fuel tanks that large. This analysis extrapolates the linear fit to a 100,000-gallon tank, still larger than the amount of fuel storage that the system builds. Indeed, the availability of a fuel pipeline often obviates the need for a fuel storage tank. However, for the linear fit likely breaks down for smaller fuel tanks. This leads to unit costs of \$3.8/gallon for a 100,000-gallon tank, falling to

$0.6/gallon for the 500,000-gallon tank. For jet fuel, this translates to $0.1/kWh(fuel) and $0.016/kWh(fuel), respectively [139]. Thus, assuming a tank on the order of 100,000 gallons, the capital cost is set to $0.1/kWh(fuel) in all three cases. Fuel storage is assumed to be lossless, with 100% efficiency.

Because the system does not build fuel storage in most cases, its cost and other parameters are not considered in the sensitivity analysis.

4.3.1.13. Electric grid interconnection

The cost of an electric grid interconnection is modeled as the capital cost of constructing electric transmission infrastructure to the facility. The system assumes AC transmission with two substations, one on each end, to transform the voltage. A 230 kV, single-circuit, three-phase AC transmission line with Southwire 795kcmil ACSR conductors has an ampacity of 884-918 amps and a total power capacity of 350-366 MVA [140]. Pletka et al. estimate that such a transmission line costs $969,000/mi, with a multiplier of 1.05 for desert terrain, which describes the Arizona case [141]. For consistency, this analysis assumes the same cost for Wyoming. Costs on the plains are likely similar, but costs could escalate by a multiplier of as much as 1.75 for mountainous terrain [141]. The two 230 kV substations cost $1.7M each [141]. Assuming a power factor of 0.86, 350 MVA equates to 303 MW(e) [142].

In the Today case, the location is assumed to be 100 miles from the nearest interconnection point to external infrastructure, at 50 miles and 10 miles in the Next-decade and Breakthrough cases. This equates to $344/kW(e), $178/kW(e), and $45/kW(e), respectively, which are rounded to $340/kW(e), $180/kW(e), and $45/kW(e). The system builds interconnections of roughly 2MW(e)-500 MW(e). In a real system, constructing a grid interconnection at the low end of this range would likely not be economical, as the cost of cable, poles, etc. do not scale linearly with

power. In such cases, electricity storage or even on-site electricity generation from H_2 or liquid fuel would likely substitute for grid electricity to satisfy this very small electricity demand at negligible incremental cost.

Grid electricity is assumed to cost \$0.65/kWh(e), the average industrial electricity rate in Arizona [106]. Wyoming rates are \$0.69/kWh(e), but the Arizona value is used for consistency [106]. Grid electricity has the recent average carbon intensity of the Western Electricity Coordinating Council (WECC), the western US interconnection, 360 kg(CO_2)/MWh(e), or 794 lb(CO_2)/MWh(e) [143].

The model currently does not allow sale of excess electricity to the grid. In a decarbonizing electric power system with high levels of renewables, bulk electricity would likely be abundant and have little resale value at times when the system currently wastes excess electricity.

4.3.1.14. Natural gas pipeline

An 8" natural gas pipeline at a relatively low pressure of 200 psi [144] and a relatively slow flow rate of 33 ft/s [145] has 166 MW(th) of capacity, assuming natural gas at 37 MJ/m^3 at atmospheric pressure and ideal gas compression to 510 MJ/m^3 [146]. This analysis uses these values in a regression model from Parker et al. for 100-mile, 50-mile, and 10-mile pipelines in the Today, Next-decade and Breakthrough cases, for costs of \$225/kW(th), \$115/kW(th), and \$25/kW(th), respectively [147]. The system tends to build roughly 100-200 MW(th) of natural gas interconnection capacity, at less than 0.5% of levelized cost.

Natural gas price is \$3/MMBtu(th), \$0.01/kWh(th), representing recent low values for industrial customers [148]. Natural gas price is held constant through all scenarios and varied in the sensitivity analysis.

4.3.1.15. Fuel export pipeline

The unit cost of a hydrocarbon liquids pipeline depends on pipeline length, diameter and flow rate. Even a small, 8" jet fuel pipeline with a relatively slow 3 ft/s flow rate amounts to 1 GW(fuel) of capacity [139,149]. Pipeline costs are derived from a regression model in Rui et al [150]. The Today case assumes a 100-mile pipeline, with a cost of \$125/ft^3 this amounts to \$22.16/kW(fuel), rounded to \$20/kW(fuel) [150]. The Next-decade and Breakthrough cases assume 50-mile and 10-mile pipelines at the same unit cost, amounting to \$10/kW(fuel) and \$2/kW(fuel), respectively. At these costs, the system tends to build roughly 500 MW(fuel) of pipeline capacity, half the 1 GW(fuel) of the above 8" pipeline. Thus, these costs may be slightly optimistic, but fuel pipeline costs represent less than 0.1% of levelized cost in the final results. Note that a 100-mile 48" pipeline with a faster 10 ft/s flow rate has a capacity of 125 GW(fuel) and a unit cost of roughly \$1/kW(fuel) [150].

4.3.1.16. Water

Both Arizona and Wyoming are water-stressed regions. This analysis assumes a relatively high cost of \$1/t(H_2O), roughly the cost of desalinated water [151]. This could also represent the cost of building a pipeline to bring water from a lower-cost location.

4.3.1.17. CO_2 sequestration

Although it is not yet widely commercialized, permanent sequestration of CO_2 in sandstone and other appropriate geologic formations appears to be technically feasible in at least some locations. Industrial volumes of CO_2 from the oil and gas industry have been successfully sequestered in the Sleippner formation for over twenty years [152].

Most techno-economic estimates of the cost of transportation and sequestration of CO_2 fall between \$1-20/t(CO_2) [153].

However, the cost and efficacy of sequestration likely depends on formation-specific features. In locations close to seismically active faults, pressure accumulation from CO_2 injection in sandstone formations could induce earthquakes, potentially compromising the shale seal required to ensure CO_2 does not return to the surface [152]. This pressure can be relieved either by reducing the CO_2 injection rate at individual wells or by extracting a corresponding volume of brine. With these pressure management techniques, Anderson et al. find sequestration costs of \$55/t(CO_2) or more at high injection volumes [154].

Basalt formations are less ubiquitous than sandstone formations in the United States and their properties are not as well understood [83].

Due to this uncertainty in the long-term cost of large-scale sequestration, this analysis models only the cost of capturing CO_2. Whether electrofuels or DAC is a more cost-effective option for mitigating liquid hydrocarbon-dependent CO_2 sources then depends on the realized cost of sequestration, as well as the price of fossil liquid fuel.

4.3.2. Waste heat recycling

Waste heat from the fuel synthesis system and electrolyzer, generally produced at around 160-350°C and 50-80°C, respectively [91,112]. This waste heat can be recycled to meet heat requirements for low-temperature DAC, which uses heat at roughly 100°C. The high-temperature fuel synthesis heat can be mixed with lower-temperature electrolyzer heat to yield some amount of 100°C heat. High-temperature PEM electrolyzers operating between 100°C and 200°C and higher-temperature alkaline electrolyzers are under development [91]. By default, 100% of this waste heat is assumed to be reusable with 90% losses through a heat exchanger [81].

This waste heat is likely not hot enough to meaningfully contribute to the 4 GJ of 900°C heat demand for calcination in high-temperature DAC systems, but could displace at least some of the

remaining 1.25 GJ of lower-temperature heat demand [97]. However, this analysis does not currently assume any heat recycling in high-temperature DAC scenarios.

4.3.3. Optimization-based techno-economic analysis

The operation of all system components is optimized to produce a fixed annual amount of liquid fuel at the lowest attainable cost using renewable electricity with a fixed generation profile, supplemented by natural gas for process heat and more expensive grid electricity.

The linear program optimization can build capacity for each of the thirteen system components: Production and storage of each of the five products, electricity, heat, hydrogen, carbon dioxide, and liquid fuel, as well as an electric grid interconnection and pipelines for importing natural gas and exporting liquid fuel. System operation is then optimized based on the year's renewable electricity production profile using a perfect foresight model. The perfect foresight assumption is likely justifiable for solar in the desert, but may be less defensible for wind, particularly if the system makes seasonal storage decisions based on knowledge of wind production months in advance.

Operation of the true system would likely also have integer variables, such as minimum capacity or operation levels. The potential effects of this relaxation are discussed in the SI, Section 7.7.

The linear optimization is outlined below, with the full formulation described in the Section 4.3.4.5:

Minimize	Annualized system cost
Subject to	Fixed total annual electrofuel production
	Fixed renewable electricity production profile throughout the year
	Production ≤ capacity for each component
	Storage level ≤ capacity
	Storage input ≤ production
	Storage output ≤ storage level
	Conservation of energy and matter
	Conservation of storage levels
	Ramping constraints

Because this optimization-based model can produce electrofuel, CO_2, H_2, or other products using electricity, it is named the "Power-to-X Optimization Tool", PtXOpt.

4.3.4. Optimization formulation

4.3.4.1. Indexes

Electrofuel production consists of flows of five primary products, electricity (*e*), hydrogen (*h*), heat (*q*), CO_2 (*c*), and fuel (*f*). Water (*w*) flows are also considered, but are derived from other decisions and are not modeled explicitly.

System components consist of production and storage of each product as well as external interconnections to the electric grid (*ex*), natural gas system (*gx*), and fuel distribution pipeline system (*fx*). Thus, there are thirteen system components.

4.3.4.2. Decisions

Capacity, x_i, for each component, *i*, is a time-invariant decision. Thus, there are thirteen total capacity decisions.

In each time period, each of the eight non-storage components has a production level ($p_{i,t}$) and ramping ($r_{k,t}$) and upward ramping decisions ($u_{k,t}$). Each of the five storage components, *l*, has three associated time-dependent decisions: Input ($\psi_{l,t}$), storage level ($s_{l,t}$), and output ($\xi_{l,t}$). Each of the five products can be disposed of as waste ($w_{j,t}$) as well. With four-hour resolution, a year-long simulation with 2,190 time periods thus requires thirteen capacity decisions and 72,270 operational decisions.

4.3.4.3. Objective function

This linear program optimization minimizes the annualized present value of constructing and operating an electrofuel production system to meet a fixed annual level of fuel production. The system models operation for a single representative year of renewable electricity production.

The objective function, Equation 8, includes the annualized capital and fixed O&M costs for each system component, using a weighted average cost of capital (WACC) and the component's lifetime to compute the annualized cost, c_i, of component i. These costs for each scenario are computed and shown in Section 4.3.4.7, **Table 8**.

$$\min_{x_i, p_{ikt}, \psi_{l,t}, s_{l,t}, \xi_{l,t}, r_{k,t}, u_{k,t}, w_{j,t}} \sum_{i,j,t} c_i x_i + v_k p_{k,t} + v_l(\psi_{l,t} + s_{l,t} + \xi_{l,t}) + \alpha_k u_{k,t} + \beta_j w_{j,t} \quad (8)$$

Time-varying operation, v_i, upward ramping, α_k, and disposal, β_j, costs are also included in the objective function. These costs are assumed to be zero for most system components (zero costs are set to $\$10^{-9}$/unit to discourage component vacuous use, with a negligible effect on the levelized cost of electrofuel). By default, use of fossil carbon and disposal of CO_2 incurs a carbon price (see Section 4.3.4.7).

Ideally, the system would minimize the cost of greenhouse gas mitigation, rather than the cost of electrofuel production. However, this would require dividing the objective function by the fossil carbon avoided through electrofuel use, making the problem nonlinear. This analysis approximates the same objective through the use of a carbon price based on the assumed cost of offsite DAC. By paying the carbon price, the electrofuel production facility is essentially paying another facility to offset any fossil carbon it opts to use. The carbon price is described further in Section 4.3.7.

Production, $p_{k,t}$, for each non-storage component, k, must be nonnegative and less than the installed level of production capacity. Production incurs variable cost, v_k. Variable operational costs for the electric grid and natural gas interconnections represent grid electricity and natural gas prices, subject to a carbon tax. These values are given for each scenario in Section 4.3.4.7,

Table 9. For all other quantities, production variable cost is set to $\$10^{-9}$/kWh or $\$10^{-9}$/t($CO_2$) to discourage vacuous production.

Production generally occurs with some associated efficiency or mass or energy requirement for conversion of inputs to outputs. $\eta_{i,j}$, represents the ratio of quantity produced by component *i* to the quantity of product *j* required to produce it. This is interpretable as an energy efficiency for components with a single energy input and a single energy output, such as the electrolyzer, whose efficiency, $\eta_{h,e}$, is the ratio of output hydrogen energy to input electrical energy. Values of $\eta_{i,j}$ are given in Section 4.3.4.7. The electrolyzer converts electricity and water into H_2, $\eta_{h,e}$ and $\eta_{h,w}$, and potentially reusable waste heat, $\eta_{h,q}$. Water flows are not modeled explicitly, but are represented assuming an all-in levelized cost of $\$1/m^3$ for desalinated water [151,155]. The electric kiln converts electricity into heat, $\eta_{q,e}$. The DAC system converts electricity and heat to CO_2, $\eta_{c,e}$ and $\eta_{c,q}$. The fuel synthesis system converts H_2 and CO_2 into liquid fuel, $\eta_{f,h}$ and $\eta_{f,c}$, producing a potentially reusable waste heat and water, $\eta_{f,q}$ and $\eta_{f,w}$. Natural gas is combusted to produce heat, $\eta_{gx,q}$. The fuel synthesis also incurs an upward ramping energy penalty, requiring hydrogen equivalent to production of $\eta_{fr,h}\eta_{f,h}$ kWh of electrofuel for each kW(fuel) of upward ramping. The DAC system also incurs ramping energy penalties of $\eta_{cr,q}\eta_{c,q}$ and $\eta_{cr,e}\eta_{c,e}$, requiring heat and electricity equivalent to 1 t(CO_2) of production for each t(CO_2)/hr of upward ramping.

Ramping, $r_{k,t}$, is the positive or negative change in a component's production relative to the previous time period. Each production component is subject to symmetric ramping constraints, requiring that $|r_{k,t}|$ not exceed a fixed percentage of installed capacity, ζ_k. ζ_k is 100% for all components except fuel synthesis and DAC, where it defaults to 10%.

Upward ramping, $u_{k,t}$, is defined as the nonnegative increase in a component's production from the previous period. Upward ramping incurs cost, α_k, to penalize increases in component operation level, which are often associated with increased energy requirements and other costs.

Storage input, $\psi_{l,t}$, must be less than or equal to total production of a given quantity, including reusable waste production potential, in each time period, including all potential sources of each product. For CO_2, H_2, and liquid fuel, there is only one source of production. Heat can come from the electric kiln, natural gas, or waste heat from the fuel synthesis and the electrolyzer. The model does not currently differentiate between low-temperature and high-temperature heat. For high-temperature DAC systems in particular, waste heat from the fuel synthesis and the electrolyzer is of limited utility, as it cannot replace calciner heat demand unless it is mixed with substantial amounts of much higher-temperature heat. For this reason, waste heat recycling is disallowed by default in high-temperature DAC scenarios. H_2 and CO_2 are stored in pressurized vessels with compression energy requirements, $\eta_{js,e,}$ described Sections 4.3.1.9 and 4.3.1.11. Storage input must be nonnegative and cannot exceed a fixed percentage of installed storage capacity, γ_ψ, 100% by default.

Storage level, $s_{l,t}$, tracks the cumulative amount of product in storage and is equal to the level in the previous time period plus any input and minus any output in the previous period. Storage level begins at zero by default in the first time period, must be nonnegative, and cannot exceed installed storage capacity.

Storage output, $\xi_{l,t}$, must be less than or equal to the storage level in the previous period. Storage operates with some efficiency, with losses incurred during output. Storage efficiencies, η_{js}, are described in Section 4.3.1 and are given in **Table 12**. Storage output must be nonnegative and cannot exceed a fixed percentage of installed storage capacity, γ_ξ, 100% by default.

Waste, $w_{j,t}$, for each of the five products allows for disposal of resources that would be uneconomic to store. By default, disposal costs, β_j, are numerically zero, set at $\$10^{-9}$/kWh to avoid vacuous waste for all quantities except CO_2, for which the disposal cost is the system-wide CO_2 price. See Section 4.3.7 for further discussion of the CO_2 price.

4.3.4.4. Conservation of product

Each of the five products is subject to a conservation constraint, which ensures that mass and energy are conserved. System inflows must match outflows in each period.

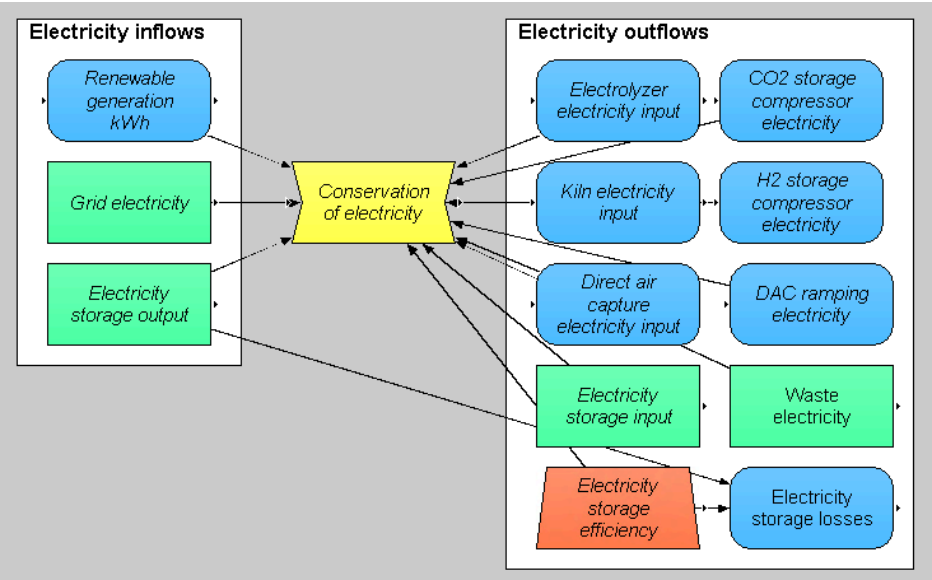

Figure 10. Conservation of Electricity. Inflows are renewable electricity generation, grid electricity, and electricity storage output. Outflows are consumption from the electrolyzer, electric kiln, and the DAC system, electricity storage input, DAC ramping electricity consumption, H_2 and CO_2 compressor energy consumption, and waste electricity. Storage losses are incurred upon discharge.

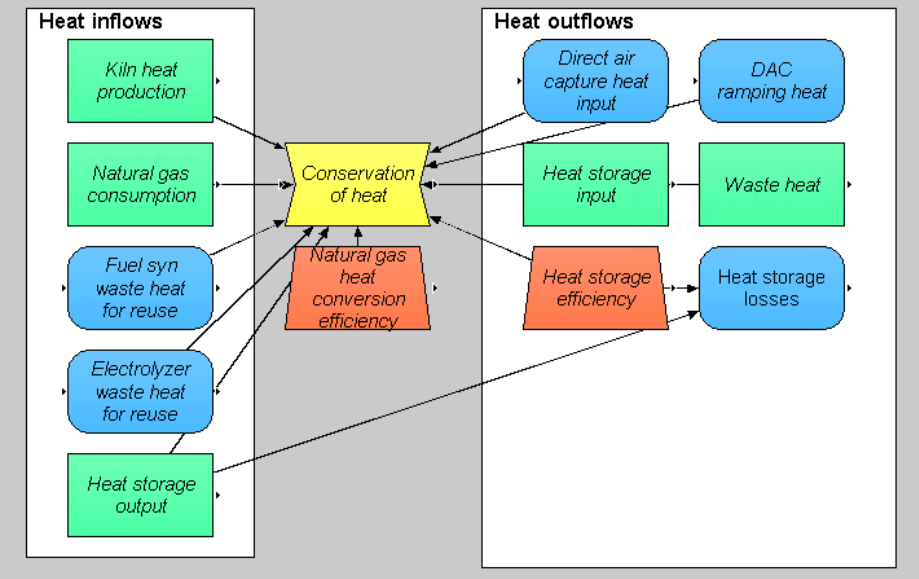

Figure 11. Conservation of Heat. Inflows are electric kiln and heat storage output, natural gas combustion, and recycled waste heat from the fuel synthesis system and electrolyzer. Outflows are consumption from the DAC system, DAC ramping heat, heat storage input, and waste heat. Storage losses are incurred upon discharge.

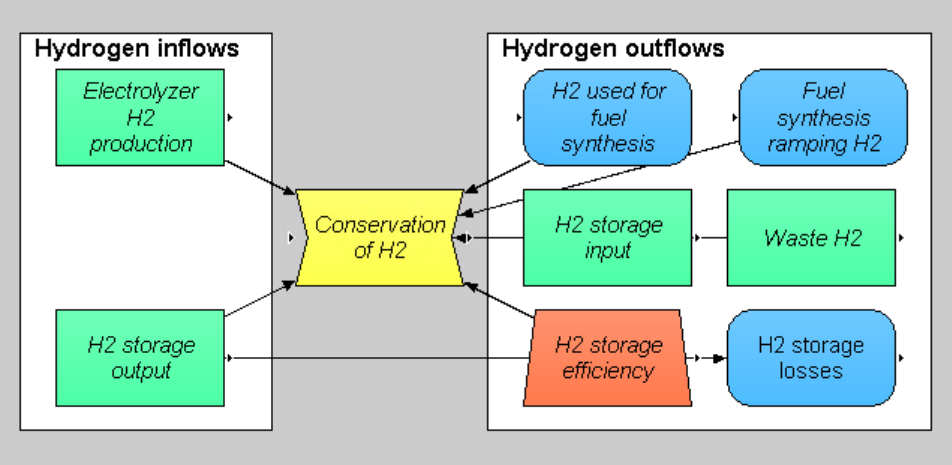

Figure 12. Conservation of Hydrogen. Inflows are electrolyzer and H_2 storage output. Outflows are consumption by the fuel synthesis system, fuel synthesis ramping, H_2 storage input, and waste H_2. Storage losses are incurred upon discharge.

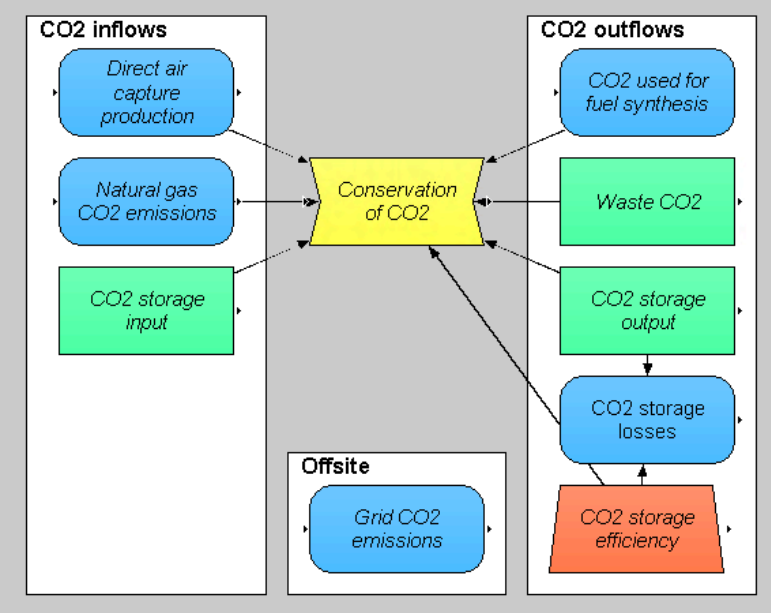

Figure 13. Conservation of carbon dioxide. Inflows are DAC production and CO_2 storage output and emissions from combustion of natural gas, which are assumed to be co-captured through oxygen-firing and can thus be converted into electrofuel. The system does not explicitly represent oxygen flows. However, H2 production through electrolysis produces more than enough oxygen for oxygen-firing of natural gas to meet DAC heat demand. Outflows are fuel synthesis and CO_2 storage input and waste CO_2. Storage losses are incurred upon discharge. Natural gas CO_2 emissions are captured through oxygen-firing and are thus available for reuse. Grid electricity CO_2 emissions occur offsite, and thus incur a carbon price but are not available for reuse within the facility.

Fuel inflows are fuel synthesis and fuel storage output. Outflows are fuel sent into the export pipeline and waste fuel.

4.3.4.5. Full formulation

Below is the optimization that underpins the PtXOpt model used in this analysis, Equation 8:

$$\min_{x_i,p_{ikt},\psi_{l,t},s_{l,t},\xi_{l,t},r_{k,t},u_{k,t},w_{j,t}} \sum_{i,j,k,l,t} c_i x_i + v_k p_{k,t} + v_l(\psi_{l,t} + s_{l,t} + \xi_{l,t}) + \alpha_k u_{k,t} + \beta_j w_{j,t} \tag{8.1}$$

(Levelized capital and fixed O&M + variable O&M + upward ramping + disposal)

$$\text{Such that } \sum_t p_{f,t} = P_{tot} \text{ (Fix total annual production)} \tag{8.2}$$

$$p_{k,t} \leq x_k \text{ (Production is below capacity)} \tag{8.3}$$

$$s_{l,t} \leq x_l \text{ (Storage level below capacity)} \tag{8.4}$$

$$\psi_{l,t} \leq \gamma_{l\psi} x_l \text{ (Storage input limit)} \tag{8.5}$$

$$\xi_{l,t} \leq \gamma_{l\xi} x_l \text{ (Storage output limit)} \tag{8.6}$$

$$s_{l,t} = s_{l,t-1} + \psi_{l,t-1} - \xi_{l,t-1} \text{ if } t \neq 0 \text{ (Storage level conservation)} \tag{8.7}$$

$$s_{l,0} = 0 \text{ for } t = 1 \text{ (Storage starts empty)} \tag{8.8}$$

$$r_{k,t} = p_{k,t} - p_{k,t-1} \text{ (Ramping definition)} \tag{8.9}$$

$$-\zeta_k x_k \leq r_{k,t} \leq \zeta_k x_k \text{ (Ramping limits)} \tag{8.10}$$

$$r_{k,t} \leq u_{k,t} \text{ (Upward ramping} \geq \text{ramping)} \tag{8.11}$$

$$0 \leq u_{k,t} \text{ (Upward ramping positive)} \tag{8.12}$$

$$p_{gx,t} \leq \eta_{gx,c}\eta_{cq} p_{c,t} \text{ (Natural gas carbon limit)} \tag{8.13}$$

(Conservation of electricity)

$$p_{e,t} + p_{ex,t} + \eta_{es}\xi_{e,t} = \frac{p_{h,t}}{\eta_{he}} + \frac{p_{q,t}}{\eta_{qe}} + \frac{p_{c,t}}{\eta_{ce}} + \frac{\psi_{h,t}}{\eta_{hse}} + \frac{\psi_{c,t}}{\eta_{cse}} + \frac{u_{c,t}}{\eta_{cre}} + \psi_{e,t} + w_{e,t} \tag{8.14}$$

(Conservation of hydrogen)

$$p_{h,t} + \eta_{hs}\xi_{h,t} = \frac{p_{f,t}}{\eta_{fh}} + \psi_{h,t} + \frac{u_{f,t}}{\eta_{frh}} + w_{h,t} \tag{8.15}$$

(Conservation of heat) (8.16)

$$p_{q,t} + \frac{p_{gx,t}}{\eta_{gxq}} + (1-\eta_{fh})\frac{p_{f,t}}{\eta_{fq}} + (1-\eta_{he})\frac{p_{h,t}}{\eta_{hq}} + \eta_{qs}\xi_{q,t} = \frac{p_{c,t}}{\eta_{cq}} + \psi_{q,t} + \frac{u_{c,t}}{\eta_{crq}} + w_{q,t}$$

(Conservation of carbon dioxide)

$$p_{c,t} + \frac{p_{gx,t}}{\eta_{gxc}} + \eta_{cs}\xi_{c,t} = \frac{p_{f,t}}{\eta_{fc}} + \psi_{c,t} + w_{c,t} \tag{8.17}$$

(Conservation of fuel)

$$p_{f,t} + \eta_{fs}\xi_{f,t} = p_{fx,t} + \psi_{f,t} + w_{f,t} \tag{8.18}$$

(Conservation of water) (8.19)

$$p_{w,t} = \frac{p_{h,t}}{\eta_{hw}} + \frac{p_{c,t}}{\eta_{cw}} + \frac{p_{w,t}}{\eta_{fw}} \tag{8.20}$$

$$\text{Where } i \in \{\Theta_p, \Theta_s, \Theta_x\}, \text{ (System components)}$$

$$j \in \Theta, \text{ (Products)}$$

$$k \in \{\Theta_p, \Theta_x\}, \text{ (Production or interconnection component)}$$

$$l \in \Theta_s, \text{ (Storage component)}$$

$$t \in 1, 2, 3, ..., T$$

$$x_i, p_{k,t}, \psi_{l,t}, s_{l,t}, \xi_{l,t}, u_{k,t}, w_{j,t} \in \mathbb{R}_+, \text{ (Nonnegativity)}$$

$$r_{k,t} \in \mathbb{R}, \text{ (Ramping can be negative)}$$

4.3.4.6. Indexes

Products and materials

$\Theta = \{$e = Electricity,
h = Hydrogen,
q = Heat,
c = Carbon dioxide,
f = Fuel,
g = Natural gas,
w = Water$\}$

Production components

$\Theta_p = \{ep = \text{Electricity production, (Renewable generation, exogenous)}$
$hp = \text{Hydrogen production, (Electrolyzer)}$
$qp = \text{Heat production, (Electric kiln)}$
$cp = \text{Carbon dioxide production, (Direct air } CO_2 \text{ capture)}$
$fp = \text{Fuel production, (Fischer-Tropsch with reverse water-gas shift)}\}$

Storage components

$\Theta_s = \{es = \text{Electricity storage},$
$hs = \text{Hydrogen storage},$
$qs = \text{Heat storage},$
$cs = \text{Carbon dioxide storage},$
$fs = \text{Fuel storage}\}$

External interconnections

$\Theta_x = \{ex = \text{Electric grid interconnection},$
$fx = \text{Fuel pipeline},$
$gx = \text{Natural gas pipeline}\}$

$t \in \{1, 2, 3, \ldots, T\}$ is the time period. With four-hour resolution T=2,190.

4.3.4.7. Optimization parameter values

The system produces P_{tot} = 4.5 TWh(fuel)/yr, equivalent to an average production of 7,570 bbl/d. The scale determined by the size of the DAC system modeled in Keith et al. [97], which produces 1 Mt(CO_2)/yr.

Table 5. Capital costs. Capital cost for the water interconnection is not modeled, as the water price is assumed to contain the interconnection cost. DAC and fuel synthesis costs are given per unit output capacity, while all other production costs are modeled in terms of input power requirements. Some costs have been converted from different units or dollar-years, as described in Section 4.3.1. All costs are in 2017 dollars.

Component	Units	Today	Next-decade	Breakthrough
Solar PV	$/kW(e)	900 [6]	500 [88]	400
Wind	$/kW(e)	1,500 [5]	1,250 [5]	1,000 [5]
Electrolyzer	$/kW(e)	1,200 [91]	580 [91]	390 [91]
Electric kiln	$/kW(e)	125 [107]	75	25
Direct air capture [high-temperature]	$/(t($CO_2$)/yr)	N/A	1,046 [97]	694 [97]
Direct air capture [low-temperature]	$/(t($CO_2$)/yr)	3,030 [102]	1,120 [102]	790 [102]
Fuel synthesis	$/kW(fuel)	800 [78]	450 [78]	350 [78]
Grid interconnection	$/kW(e)	340 [140,141]	180 [140,141]	45 [140,141]
Fuel pipeline	$/kW(fuel)	20 [156]	10 [156]	2 [156]
Natural gas pipeline	$/kW(th)	225 [147]	115 [147]	25 [147]
Electricity storage	$/kW(e)	350 [7]	250	100 [116]
H_2 storage	$/kW($H_2$)	25 [121,122]	15 [122]	6 [121]
Heat storage [Low-temperature]	$/kW(th)	25 [86,133]	20 [129]	15 [129]
Heat storage [High-temperature]	$/kW(th)	N/A	55 [127]	25
CO_2 storage	$/t($CO_2$)	50,000	10,000 [136,137]	1,000
Fuel storage	$/kW(fuel)	0.1 [138,139]	0.1 [138,139]	0.1 [138,139]
WACC	%	10%	8%	5%

Table 6. Operations and maintenance costs. Non-energy costs of operating and maintaining each component. Units are % of capital cost per year. Includes both fixed and variable O&M. Reasons for these choices are described further in Section 4.3.1.

Component	%/year
Solar PV [87]	1%
Wind [87]	2.5%
Electrolyzer [78,91]	3%
Electric kiln	1%
Direct air capture [97]	4%
Fuel synthesis [113]	4%
Grid interconnection	1%
Fuel pipeline	1%
Natural gas pipeline	1%
Electricity storage [119,120]	1%
H_2 storage	1%
Heat storage	1%
CO_2 storage [136]	5%
Fuel storage	1%

Table 7. System lifetime. Held constant in all cases. The 25-year lifetime of the DAC system dictates the full system lifetime and thus maximum component lifetime [97]. Thus, citations below may have a longer lifetime. Reasons for these choices are described further in Section 4.3.1.

Component	Years
Solar PV [87]	25
Wind [87]	25
Electrolyzer [78,91]	25
Electric kiln	10
Direct air capture [97,104]	25
Fuel synthesis [78]	25
Grid interconnection	25
Fuel pipeline	25
Natural gas pipeline	25
Electricity storage [115–117]	10
H_2 storage	25
Heat storage	25
CO_2 storage	25
Fuel storage	25

The model does not include a separate estimate of the cost of integrating components into a functioning system. Each component consumes and produces one or a few products, each of which can be transported at low cost through pipelines or wires. This suggests that system integration costs would be low. Still, this assumption biases results in favor of electrofuel production, which requires additional system integration costs over standalone DACS. Annualized capital and non-energy O&M costs are combined using Eq. 9 to form the c_i costs used in the optimization, shown in **Table 8**. Capital costs are annualized according to their lifetime and the weighted average cost of capital (WACC) using the capital charge factor described in Equation 7.

$$c_i = capex_i \, (CCF_i + O\&M\%_i) \tag{9}$$

Table 8. Annualized capital and non-energy O&M cost. Derived from **Table 5**, **Table 6**, and **Table 7** using Eq. 9.

Component	Units	Today	Next-decade	Breakthrough
Solar PV (c_{ep})	$/kW(e)-yr	108.2	51.8	32.4
Wind (c_{ep})	$/kW(e)-yr	202.8	148.3	96.0
Electrolyzer (c_{hp})	$/kW(e)-yr	168.2	71.7	39.4
Electric kiln (c_{qp})	$/kW(e)-yr	15.0	7.8	2.0
Direct air capture [high-temperature] (c_{cp})	$/(t(CO2)/yr)-yr	N/A	139.8	77
Direct air capture [low-temperature] (c_{cp})	$/(t(CO2)/yr)-yr	455.0	149.7	87.7
Fuel synthesis (c_{fp})	$/kW(fuel)-yr	120.1	60.2	38.8
Grid interconnection (c_{ee})	$/kW(e)-yr	40.9	18.7	3.6
Fuel pipeline (c_{fe})	$/kW(fuel)-yr	2.4	1.0	0.2
Natural gas pipeline (c_{ge})	$/kW(th)-yr	27.0	11.9	2.0
Electricity storage (c_{es})	$/kWh(e)-yr	60.5	39.8	14.0
H_2 storage (c_{hs})	$/kWh($H_2$)-yr	3.0	1.6	0.5
Heat storage (c_{qs})	$/kWh(th)-yr	3.0	5.7	2.0
CO_2 storage (c_{cs})	$/t($CO_2$)-yr	8008	1437	121
Fuel storage (c_{fs})	$/kWh(fuel)-yr	0.01	0.01	0.01

Table 9. Variable cost, v_i, is assumed to be zero in all cases except for external interconnections to grid electricity, natural gas, and water and carbon payments, which are added as well.

Component	Unit	Today
Grid electricity price (v_{ee}) [106]	$/kWh(e)	0.65
Natural gas price (v_{ge}) [148]	$/MMBtu(th)	3
Water (v_{we}) [151,155]	$/t($H_2O$)	1

Table 10. Fossil CO_2 emissions. For natural gas and grid electricity, carbon payments are represented as variable costs within v_k, as the product of the carbon price and the CO_2 intensity of the fuel.

Component	Units	Today	Next-decade	Breakthrough
Grid CO_2 intensity	lb(CO_2)/MWh(e)	794 [143]	397	0
Natural gas emissions	lb(CO_2)/MMBtu(th)	117 [157]	117 [157]	117 [157]

Table 11. Carbon price in each scenario. See Section 4.3.7 for further details.

CO_2 price	$/t($CO_2$)
Today, low-temperature [102]	600
Next-decade, high-temperature [97]	232
Next-decade, low-temperature [102]	200
Breakthrough, high-temperature [97]	126
Breakthrough, low-temperature [102]	100

Table 12. Efficiency and energy and material requirements. Solar and wind efficiencies are built into capital costs and the production profile. External interconnections are assumed to be lossless.

Component	Units	Today	Next-decade	Breakthrough
Electrolyzer electricity (η_{he})	kWh(e)/kWh(H_2)	65% [78,91]	70% [78,91]	75% [78,91]
Electrolyzer water (η_{hw})	kg(H_2O)/kWh(H_2)	0.27 [78]	0.27 [78]	0.27 [78]
Electric kiln (η_{qe})	kWh(e)/kWh(th)	95%	95%	95%
Direct air capture heat [high-temperature] (η_{cq})	GJ(th)/t(CO_2)	N/A	5.25 [97]	5.25 [97]
Direct air capture electricity [high-temperature] (η_{ce})	kWh(e)/t(CO_2)	N /A	366 [97]	77 [97]
Direct air capture water [high-temperature] (η_{cw})	t(H_2O)/t(CO_2)	4.7 [97]	4.7 [97]	4.7 [97]
Direct air capture heat [low-temperature] (η_{cq})	GJ(th)/t(CO_2)	7.9 [104]	5.8 [104]	4.4 [98,158]
Direct air capture electricity [low-temperature] (η_{ce})	kWh(e)/t(CO_2)	700 [104]	400 [104]	160 [98,158]
Direct air capture heat [low-temperature] (η_{cw})	t(H_2O)/t(CO_2)	-1 [104]	-1 [104]	-1 [104]
Fuel synthesis CO_2 (η_{fc})	t(CO_2)/MWh(fuel)	0.28 [78]	0.28 [78]	0.28 [78]
Fuel synthesis H_2 (η_{fh})	kWh(H_2)/kWh(fuel)	70 [114,159]	70 [114,159]	75 [114,159]
Electricity storage (η_{es})	%	80 [117,118]	85 [117]	90 [117]
H_2 storage (η_{hs})	%	99	99	99
Heat storage [Low-temperature] (η_{qs})	%	90 [131]	90 [131]	90 [131]
Heat storage [High-temperature] (η_{qs})	%	90 [131]	90 [131]	90 [131]
CO_2 storage (η_{cs})	%	100	100	100
Fuel storage (η_{fs})	%	100	100	100
Natural gas combustion (η_{gxq})	kWh(th)/kWh(th)	100	100	100

Table 13. Storage compression energy. Incurred upon storage input.

Component	Units	Today	Next-decade	Breakthrough
H_2 storage (η_{hse})	kWh(e)/kWh(H_2)	0.036 [121]	0.036 [121]	0.033 [121]
CO_2 storage (η_{cse})	kWh(e)/t(CO_2)	289 [97]	289 [97]	289 [97]

Table 14. Ramping energy penalty incurred for each hourly unit of upward ramping as a percentage of input energy requirements for one t(CO_2) for DAC and one kWh(fuel) for fuel synthesis. 100% means that for every kW(fuel) of ramping, the fuel synthesis infrastructure consumes H_2 equal to what it would use to produce 1 kWh(fuel). DAC ramping incurs analogous heat and electricity penalties based on input requirements per t(CO_2).

Component	Units	
Direct air capture (η_{cre}, η_{crh})	%	100%
Fuel synthesis (η_{frh})	%	100%

Table 15. Waste heat reusability rate. Includes heat exchanger losses. Waste heat recycling is only available for low-temperature DAC systems.

Component	Units	Today	Next-decade	Breakthrough
Electrolyzer (η_{hq})	%	0%	0%	90% [81]
Fuel synthesis (η_{fq})	%	0%	90% [81]	90% [81]

Ramping

All components except the DAC and fuel synthesis can ramp to full capacity within one hour (ζ_k = 100%). Renewable electricity is treated as exogenous. Existing alkaline and PEM electrolyzers are capable of cold start to minimum load times of 20 minutes or less, with the ability to ramp up and down to and from full capacity within seconds to minutes once they have begun operating [91]. Electric resistive heating simply requires a that current be applied to a heating element, which could be kept hot in an insulated containment vessel. Lithium-ion electricity storage can comfortably ramp on timescales of seconds or less. Gas and liquid storage systems are limited

primarily by compression and decompression, which likely would take minutes for foreseeable input and output demands.

Due to their operational complexity and potential for high operating temperatures, the DAC and fuel synthesis are assumed to take 10 hours to reach full capacity, thus $\zeta_{cp} = \zeta_{fp} = 10\%$. The sensitivity analysis varies this from 1% to 100%.

If time is aggregated below hourly resolution, ζ_k is multiplied by the number of hours per time period, with a maximum value of 100%. Thus, for four-hour resolution, the effective value of $\zeta_{cp} = \zeta_{fp} = 40\%$.

Table 16. Ramping limit. Maximum absolute hourly ramping as a fraction of installed production capacity.

Component	%
Solar PV (ζ_{ep})	100%
Wind (ζ_{ep})	100%
Electrolyzer (ζ_{hp})	100%
Electric kiln (ζ_{qp})	100%
Direct air capture (ζ_{cp})	10%
Fuel synthesis (ζ_{fp})	10%
Grid interconnection (ζ_{ex})	100%
Fuel pipeline (ζ_{fx})	100%
Natural gas pipeline (ζ_{gx})	100%

Storage power capacity

The flow rate in and out of storage components is limited only by storage capacity. In practice, modeled storage inflows and outflows are generally much smaller than rated capacity.

Natural gas carbon limit

CO_2 from combustion of natural gas is assumed to be co-captured in an oxygen-fired kiln used to heat the DAC process. This is based on the high-temperature system design in Keith et al. [97], which could be implemented to provide heat to a low-temperature system as well. In an electrofuel production system, oxygen from electrolysis of water would be abundantly available.

Oxygen flows are not modeled explicitly in this analysis and doing so would likely have only a marginal effect on the system operation and overall cost.

Co-captured CO_2 can be used for electrofuel production, displacing captured CO_2. Unconstrained, the model tends to prefer low-cost fossil CO_2 to captured CO_2. The model guards against excessive reliance on fossil CO_2 in two ways. First, the carbon tax gives fossil CO_2 an effective cost on par with captured CO_2. Second, the model includes a constraint that prevents instantaneous natural gas CO_2 emissions from exceeding instantaneous DAC CO_2 demand.

4.3.5. No storage scenarios

"No storage" scenarios assume storage and external grid electricity and natural gas interconnections are not available, constraining all system components to operate at the same capacity factor as the renewable electricity consumption. These cases remove ramping constraints and costs and do not allow any product to be wasted.

4.3.6. Levelized cost of fuel

The levelized cost of electrofuel is the optimal annualized electrofuel production system cost, the optimal objective function value, divided by the total quantity of electrofuel produced. This is reported in $/gallon of gasoline equivalent (GGE).

4.3.7. Carbon price

The model includes a carbon price to penalize the use of fossil carbon from natural gas and grid electricity, as well as waste of captured CO_2. By default, the CO_2 price is the levelized cost of DAC claimed by manufacturers under the technology assumptions in the corresponding scenario. The carbon price is based on cost claims by manufacturers, which may involve different assumptions than those modeled here, particularly the potential use of low-cost waste heat or on-site natural gas electricity generation. For low-temperature systems, this corresponds to

$600/t($CO_2$) for Today [102], based on Climeworks' estimated current costs, $200/t($CO_2$) in the Next-decade, and $100/t($CO_2$) in the Breakthrough case, based on Climeworks' statements to the media [102]. High-temperature systems use Carbon Engineering's estimated costs from Keith et al., with early build costs of $232/t($CO_2$) for the next-decade assuming a 12.5% capital return factor and Nth-of-a-kind costs of $126/t($CO_2$), assuming a 7.5% capital return factor, for a system producing pipeline-ready CO_2 at 151 bar [97]. These carbon prices assume free sequestration, in line with literature estimates that place sequestration costs as low as $2/t($CO_2$) [153]. This is also consistent with estimates of the cost of standalone DAC in this analysis, in which the highly uncertain cost of sequestration is treated as an unknown parameter whose value determines the cost differential that makes electrofuel production a more or less cost-effective mitigation strategy than DACS. See the SI, Section 7.9.1 for a simulation with no carbon price.

4.3.8. Electrofuel greenhouse gas mitigation cost-effectiveness

For electrofuels, the cost of mitigating fossil GHG emissions depends on the life-cycle carbon content of equivalent petroleum-based fuels and the relative price of the two fuels. This analysis uses petroleum jet fuel for comparison because there is less potential for fuel substitution in aviation than in ground or water transportation and thus likely a greater willingness to pay for a carbon-neutral hydrocarbon fuel.

The main analysis computes electrofuel mitigation cost, $C_{mitigation}$, using Eq. 10, assuming that the carbon price paid during electrofuel production is used to offset emissions from the use of fossil carbon.

$$C_{mitigation} = (C_{electrofuel} - C_{petro\text{-}fuel})/GHG_{petro\text{-}fuel} \quad (10)$$

Where $C_{electrofuel}$ is the production cost of electrofuel, $C_{petro\text{-}fuel}$ and $GHG_{petro\text{-}fuel}$ are the assumed unit cost and life-cycle GHG emissions intensity of equivalent petroleum-based fuels.

Combustion of a gallon of gasoline equivalent (GGE) of jet fuel produces 8.15 kg(CO_2)/GGE [160]. Oil refining is assumed to be 87% efficient, based on efficiency for low-sulfur diesel, which is chemically similar to jet fuel, resulting in an additional 1.22 kg(CO_2)/GGE [161]. Emissions from extraction and transportation of oil vary widely, with a median estimate of 1.04 kg(CO_2)/GGE [162]. Thus, the life-cycle carbon content of petroleum-based fuels, $GHG_{petro\text{-}fuel}$, is estimated at 10.4 kg(CO_2)/GGE(fuel).

Since 2008, unsubsidized jet fuel prices for domestic US flights have ranged from $1.50-3.20/GGE, averaging around $2/GGE [163]. This analysis treats $2/GGE as the baseline cost of jet fuel, considering a $3/GGE case as well.

Note that this is not a perfect comparison, as it does not account for life-cycle emissions associated with electrofuel production apart from direct emissions due to the use of fossil energy. This analysis also does not account for profit margins, which are included in oil prices. Taxes are not included here in electrofuel or DAC production or oil prices.

See the SI, Section 7.9.1 for a case without a carbon price.

4.3.9. DAC-only analysis

The electrofuel GHG mitigation cost is compared to the cost of a standalone DAC system under the same scenario assumptions. These cases use the same optimization formulation described in Section 4.3.4, but the system captures 1 Mt(CO_2) from the atmosphere per year instead of producing 4.5 TWh of electrofuel. Thus, DAC-only simulations do not include hydrogen or hydrocarbon fuel-related infrastructure. As with electrofuels, the levelized cost of DAC is the

objective function. The annualized cost of production is annualized system cost divided by the total CO_2 produced in the simulated year. This is reported in \$/t(CO_2).

With access to natural gas, the standalone DAC cases operate at roughly a 100% capacity factor. The levelized cost comparison with electrofuel production inflates the levelized cost of DAC by a factor of 1/0.9 to simulate a 90% capacity factor, which is the level targeted in Keith et al. [97]. In this analysis, electrofuel production scenarios tend to operate the DAC at capacity factors between 80% and 90%. The system could likely be taken down for maintenance during lulls in renewable electricity availability with a modest impact on annualized system cost.

4.3.10. Sensitivity analysis

The sensitivity analysis individually varies each parameter from its baseline value in the Next-decade high-temperature solar electrofuel case, with the following exceptions:

1. DAC and fuel synthesis ramping energy consumption and ramping rate limits, as well as compression energy requirements for H_2 storage, all of which are constant in the main analysis, are set to 0 and ten times their baseline value.
2. CO_2 compression energy remains at 0 kWh(e)/t(CO_2) in the baseline Next-decade case because the DAC system modeled produces pressurized CO_2. The high value remains at the level in the Today case. The low value remains at 0 kWh(e)/t(CO_2).
3. Kiln efficiency, held constant at 95% in the main analysis, varies from 75% to 99%.
4. Heat storage efficiency, held constant at 90% in the main analysis, varies from 70% to 95%.
5. A high heat storage cost of \$150/kWh(th) is used instead of the Today case cost of \$25/kWh(th), which represents the cost of molten salt heat storage [86,133]. This represents a case in which high-temperature thermal storage is roughly as expensive as the assumed

long-term cost of electricity storage. A low cost of $15/kWh(th) is used to represent the US DOE 2020 target for supercritical CO_2 thermal storage for concentrated solar electricity generation facilities [129]. Such storage would not reach 900°C but represents an option for low-temperature systems and an aspirational goal for high-temperature heat storage.

6. Lifetimes differ across components. The low and high cases consider lifetime changes of ±20% of the baseline values.

For computational reasons, the sensitivity analysis considers only a subset of the year, still with 4-hour resolution. Hours 1,501-2,500 reproduce a levelized electrofuel cost of $5.02/GGE, 6% below the baseline value simulated over the full year. As a result, this simulation will miss seasonal effects, allowing components to be undersized to match the simulated electricity production profile and potentially undervaluing seasonal CO_2 storage as well. However, most storage decisions take place over timescales of one or a few days, suggesting that this shorter simulation will yield insight into the system's sensitivity to key parameters in a full annual simulation.

Switching to 8-hour resolution allows the system to dispense with electricity storage simply due to aggregation effects, preventing exploration of important uncertainties. For this reason, sensitivity is not simulated at 8-hour resolution, even though this would allow a computationally tractable exploration of seasonal sensitivities.

4.4. Analysis and results

The levelized cost of electrofuel, shown in **Figure 14**, varies substantially across the six scenarios. A system running with Today's technology at a fixed capacity factor, 29.3% for solar and 57.5% for wind produces fuel at $25.60/GGE and $14.94/GGE, respectively. Adding storage

and electric grid and natural gas interconnections to the system allows capital-intensive assets such as the DAC system to run at a capacity factor over 89% in all solar cases and over 80% in all wind cases (see the SI, Section 7.8.1 for the capacity factor of each component). Under the same assumptions, adding storage and external interconnections reduces the levelized cost to $13.85/GGE and $12.75/GGE for solar and wind, respectively.

In the Next-decade cases, costs fall to $5.25-5.35/GGE or $5.89-5.68/GGE for solar- and wind-powered systems, respectively. Low-temperature DAC's higher capital and fixed O&M costs are offset by its ability to reuse waste heat from fuel synthesis. In the Breakthrough cases, costs fall further to $2.69-2.78/GGE for solar and $2.97-3.07/GGE for wind, with nearly all heat demand in the low-temperature case met by waste heat from the fuel synthesis and the electrolyzer.

In all cases, the levelized cost of electrofuel is dominated by the capital and O&M costs of the DAC system, the electrolyzer, and the renewable electricity generation. The electrolyzer continues to operate intermittently, but at a capacity factor of 38-40% for solar and 57-62% for wind, above well above the 29.3% capacity factor for solar, but comparable to wind's 57.5% capacity factor. Note that for cases with storage, fuel synthesis capital represents only 4-7% of levelized cost. This is due in part to a capacity factor above 80% and a capacity requirement roughly 25% that of the electrolyzer, roughly 1/3 on an energy input basis assuming 70% efficiency. See the SI, Section 7.8.1 for installed capacities and capacity factors for all production components.

Note that for low-temperature cases, water produced from the DAC system is sufficient to supply the electrolyzer. However, water costs in high-temperature cases represent at most $0.05/GGE, or 1.6% of total cost.

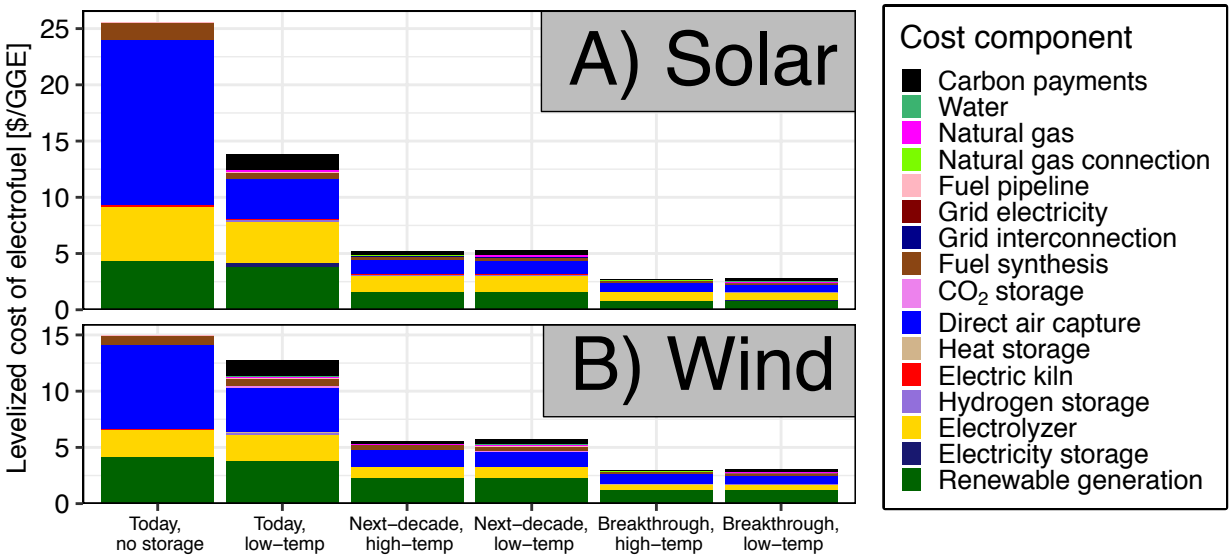

Figure 14. The levelized cost of electrofuel per gallon of gasoline equivalent (GGE) of liquid hydrocarbon fuel produced using solar (A) and wind (B) electricity. In all six cases, the three largest cost components are capital and non-energy operating costs for the direct air carbon capture (DAC) system, the electrolyzer, and the renewable electricity generation. Fuel synthesis and storage capital and operating costs are far smaller in all cases. However, storage of electricity, hydrogen, heat, and carbon dioxide, as well as grid electricity or natural gas interconnections, allow the direct air capture, fuel synthesis infrastructure, and other components to operate at a higher capacity factor than the renewable electricity, resulting in an $11.75/GGE savings for a solar-powered system built with today's technology, including a low-temperature DAC system, but only result in $2.19/GGE in savings for a wind-powered system due in part to an electrical capacity factor nearly double that of the solar. Most cases opt to use some natural gas or grid electricity despite carbon payments commensurate with the cost of direct air capture in the assumed scenario. Under Next-decade assumptions in which near-term cost and performance targets are achieved, the cost of production falls below $6/GGE for both the low-temperature system and high-temperature system, falling to $5.23/GGE in the solar low-temperature case. Longer-term technological Breakthroughs in each of the constituent technologies could lead to production costs of roughly $3/GGE or lower in all cases. Despite wind's higher capacity factor, the presence of storage and natural gas and electric grid interconnections make solar's lower levelized cost of electricity the decisive factor in achieving lower fuel cost than wind in the Next-decade and Breakthrough cases.

The system opts to use some natural gas for process heat in all scenarios in which it is available despite carbon prices of \$126-600/t(CO_2). Natural gas interconnection capacity is 99.5-291.7 MW(th). For cases without waste heat recycling, all but the low-temperature Next-decade and Breakthrough cases, this interconnection is used at a capacity factor of 61-83%. Cases with waste heat recycling have capacity factors of 21-49%. Note that many cases opt to use more natural gas than electricity for process heat. In cases with storage and without heat recycling, electric kiln capacity is 215.5-577.5 MW(e), used at capacity factors of 1-22%. The Next-decade low-temperature cases, which only reuse waste heat from the fuel synthesis, build 38.2-82.6 MW(e) of capacity at a 7% capacity factor. The Breakthrough low-temperature cases do not build electric kiln capacity.

All cases with storage build a grid interconnection of 9-55 MW(e). The solar Today case uses 10 MW(e) at a 17% capacity factor. The solar Next-decade cases use 48-53 MW(e) at a capacity factor of 28-31%. The solar Breakthrough cases use 9-14 MW(e) at a capacity factor of 3-6%. Wind cases build 31-55 MW(e) at a capacity factor of 1-2%

As a result of fossil energy use, primarily from natural gas, carbon credit payments represent as much as \$1.42/GGE or 10.2% of total cost in the Today solar scenario and \$0.44/GGE or 8% of total cost in the solar Next-decade high-temperature scenario. However, in no case does the system rely entirely on DAC process heat from natural gas.

Despite their enormous importance in lowering system cost, storage and external interconnection capital costs represent only 2-5% of total cost in all cases. For solar electrofuels, storage of electricity, hydrogen, and heat is generally diurnal, retaining enough electricity, heat, and hydrogen to run the DAC and fuel synthesis systems through the night. The electric grid and natural gas interconnections are direct substitutes for electricity and heat storage. In the Next-

decade cases, natural gas entirely replaces heat storage (see the SI, Section 7.8.1 for installed storage capacities). The Breakthrough cases use some heat storage, at less than 0.25% of total cost. If natural gas and grid electricity are not available in the Next-decade high-temperature solar case, the system builds 4.2 GW(th) of heat storage, representing $0.18/GGE of a total levelized cost of $5.68/GGE, resulting in a cost increase of $0.33/GGE over the corresponding case that allows use of natural gas and grid electricity. See the SI, Section 7.9.2 for further details of this case.

Removing the carbon price encourages further use of fossil energy. In the Next-decade high-temperature solar case, removing the carbon price reduces levelized cost from $5.35/GGE to $4.89/GGE, for a cost reduction of $0.46/GGE. In the corresponding case with a carbon price, carbon payments represent $0.44/GGE. System operation changes only slightly, with natural gas rising from $0.12/GGE to $0.13/GGE and grid electricity costs rising from $0.05/GGE to $0.12/GGE. This results in a 20% increase in the net fossil carbon intensity of electrofuel, from 1.9 kg(CO_2)/GGE to 2.28 kg(CO_2)/GGE, compared to a life-cycle carbon intensity of petroleum jet fuel estimated at 10.5 kg(CO_2)/GGE in Section 4.3.8.

Recall that the carbon payments roughly represent the cost of deploying direct air capture with sequestration to offset fossil emissions. Thus, without carbon payments, this increase in the fossil intensity of electrofuel results reduces net fossil CO_2 emission savings by approximately 20%, increasing the implicit mitigation cost of electrofuel accordingly.

CO_2 storage serves as a medium for seasonal energy storage in the Breakthrough cases, building 12-44 kt(CO_2) of capacity, representing roughly 3-12 days of DAC production at full capacity. In the Breakthrough high-temperature solar case, CO_2 is stored for nearly 200 days (see the SI, Section 7.9.3 for further details). The Next-decade cases build 1.5-9 kt(CO_2) storage with

charge-discharge periods as long as two months. Even with very expensive CO_2 storage in the Today case, the system builds 0.5-2.7 kt(CO_2) storage, highlighting the importance of CO_2 storage in maintaining a high capacity factor for the DAC system.

All cases opt to waste 4-8% of the electricity produced in the solar cases and 3-11% in the wind cases. This allows the system to reduce capacity requirements for capital-intensive components, thus lowering system cost. Low-temperature systems also waste a substantial amount of heat due to the high availability of costless recyclable waste heat from the fuel synthesis and, in the Breakthrough case, the electrolyzer. No scenario wastes H_2, CO_2, or fuel, as producing each incurs substantial efficiency losses without a corresponding reduction in storage costs compared to alternative methods of energy or material storage. See the SI, Section 7.8.1 for simulated waste levels.

Figure 15 compares the implicit GHG mitigation cost of electrofuel to the cost of standalone DAC under the same techno-economic assumptions, assuming \$2/GGE (A) and \$3/GGE (B) for equivalent petroleum-based fuels. The method for computing this implicit mitigation cost is described in detail in Section 4.3.5. The cost of standalone DAC does not include CO_2 sequestration costs. As a result, DACS is only a more cost-effective mitigation strategy if it is cheaper than mitigation through electrofuels by at least the cost of sequestration. Estimates of sequestration costs generally range from \$1-20/t(CO_2) with some at \$50/t(CO_2) or more once the cost of reservoir pressure management is included [153,154]. Higher petroleum fuel prices also reduce the mitigation cost of electrofuels, which is based on the cost differential between the two fuels.

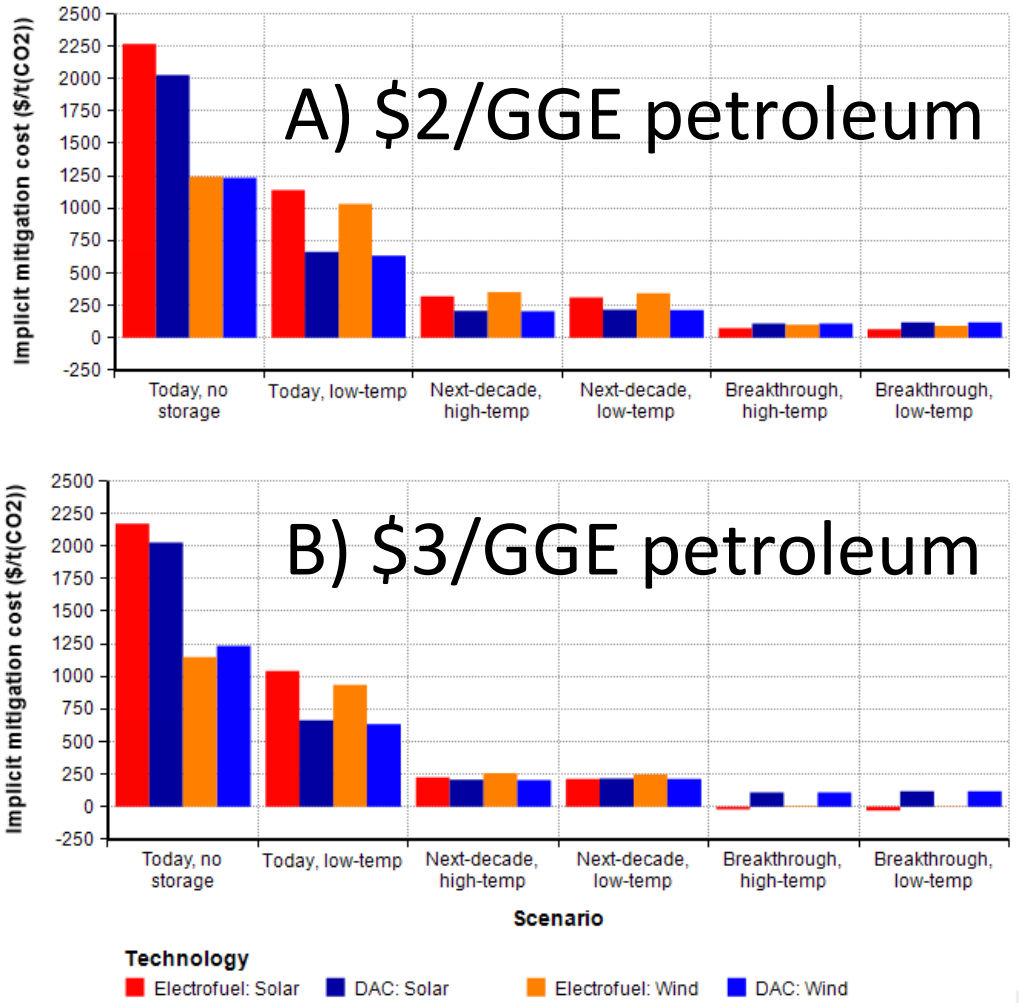

Figure 15. The cost of electrofuel production from solar or wind electricity as a greenhouse gas mitigation strategy compared with standalone direct air capture (DAC). Electrofuel mitigation costs are based on the cost increase relative to petroleum fuels at $2/GGE (A) and $3/GGE (B) and on the resulting carbon emissions avoided. Mitigation costs in the Today cases range from $630-$2,170/t(CO_2), falling to $210-1,150/t($CO_2$) in the Next-decade cases, and $-30-120/t($CO_2$) in the Breakthrough cases. DAC cost estimates do not include the cost of sequestration, estimates of which generally range from $1-20/t($CO_2$) with some at $50/t($CO_2$) or more [153,154], meaning that DAC with sequestration (DACS) is a lower-cost mitigation strategy than electrofuel production as long as sequestration costs and petroleum fuel prices do not negate the savings above. DAC is likely the more cost-effective option using Today's technology. Under Next-decade assumptions with petroleum fuels at $2/GGE, DAC is still $93-150/t($CO_2$) less expensive, requiring correspondingly high sequestration costs to achieve parity. If petroleum fuels rise to $3/GGE, wind electrofuels fall within $35-54/t($CO_2$) of DAC, within the high-end range of current sequestration

cost estimates [154]. In the Next-decade cases at $3/GGE, solar electrofuels approach or achieve cost parity with DAC, leading to cost differentials of $-3-17 t($CO_2$). In the Breakthrough cases, electrofuels are always a lower-cost mitigation strategy than DAC, with solar electrofuels achieving cost parity petroleum fuels at $3/GGE.

In the Today case with storage, electrofuels cost $304-475/t($CO_2$) more than DAC. This means that the cost of sequestration would have to reach this level or the cost of petroleum fuels would need to roughly double or triple for electrofuels to be a more cost-effective mitigation strategy. Note that without storage, wind-powered electrofuels are actually preferred to wind-powered DAC for petroleum fuels at $3/GGE.

In the Next-decade cases, electrofuels have mitigation costs of $214-322/t($CO_2$) with conventional fuels at $2/GGE, making them are $93-150/t($CO_2$) more expensive than DAC. With petroleum fuels at $3/GGE, this cost differential falls to $35-53/t($CO_2$) for wind cases, within the range of some estimates of the cost of sequestration with active pressure management [154]. For solar cases, the cost differential falls to $-3-17/t($CO_2$), meaning that in the low-temperature case, solar electrofuels are a more cost-effective climate mitigation strategy than DACS even if sequestration were costless.

Figure 16 shows the sensitivity of the levelized cost of solar electrofuel in the Next-decade high-temperature scenario to variation in individual parameters to high and low values. In most cases, the high values are from the Today low-temperature case and the low values are from Breakthrough high-temperature case. Sensitivity assumptions are described in detail in Section 4.3.10. The three most sensitive parameters are the capital cost of DAC, the electrolyzer, and renewable electricity generation, which introduce cost swings of $2.41/GGE, $1.91/GGE, and $1.37/GGE, respectively. The next most sensitive parameter is the weighted average cost of capital (WACC), which determines the present value of capital costs, with a cost swing of $1.45/GGE.

Electrolyzer and fuel synthesis efficiency are sensitive parameters, but the range of potential variation is small, resulting in cost swings of $0.39/GGE and $0.19/GGE, respectively. Lifetime changes of ±20% introduce a cost swing of $0.48/GGE.

Carbon price and natural gas price are two of the next most sensitive parameters, with cost swings of $0.46/GGE and $0.22/GGE, respectively. These parameters are largely substitutable, although the carbon price also affects the effective cost of grid electricity.

DAC heat demand varies by 4.4-7.9 GJ(th), 1,225-2,200 kWh(th), while DAC electricity demand varies from 77-700 kWh(e). DAC electricity demand introduces a cost swing of $0.28/GGE, while DAC heat demand introduces only $0.13/GGE. Thus, the availability of relatively low-cost natural gas and the higher cost of electricity storage and grid electricity render the system more sensitive to changes in electricity consumption than heat consumption.

Fuel synthesis capital cost, DAC electricity demand and H_2 storage capital cost are the last components with appreciable sensitivity, with cost swings of $0.24/GGE, $0.15/GGE and $0.20/GGE, respectively.

CO_2 storage introduces a cost swing of $0.13/GGE, demonstrating the modest but significant potential role of low-cost seasonal CO_2 storage to reduce system costs. This effect is likely underestimated as the sensitivity analysis simulates only 1000 hours of operation and thus cannot fully capture the benefits of seasonal storage.

Note that although the system is not very sensitive to the cost of individual forms of storage, this is largely because there is a high potential for substitution between different forms of storage as well as natural gas and electricity interconnections. **Figure 14** shows that the Today case without any storage or external interconnections sees a substantial increase in cost.

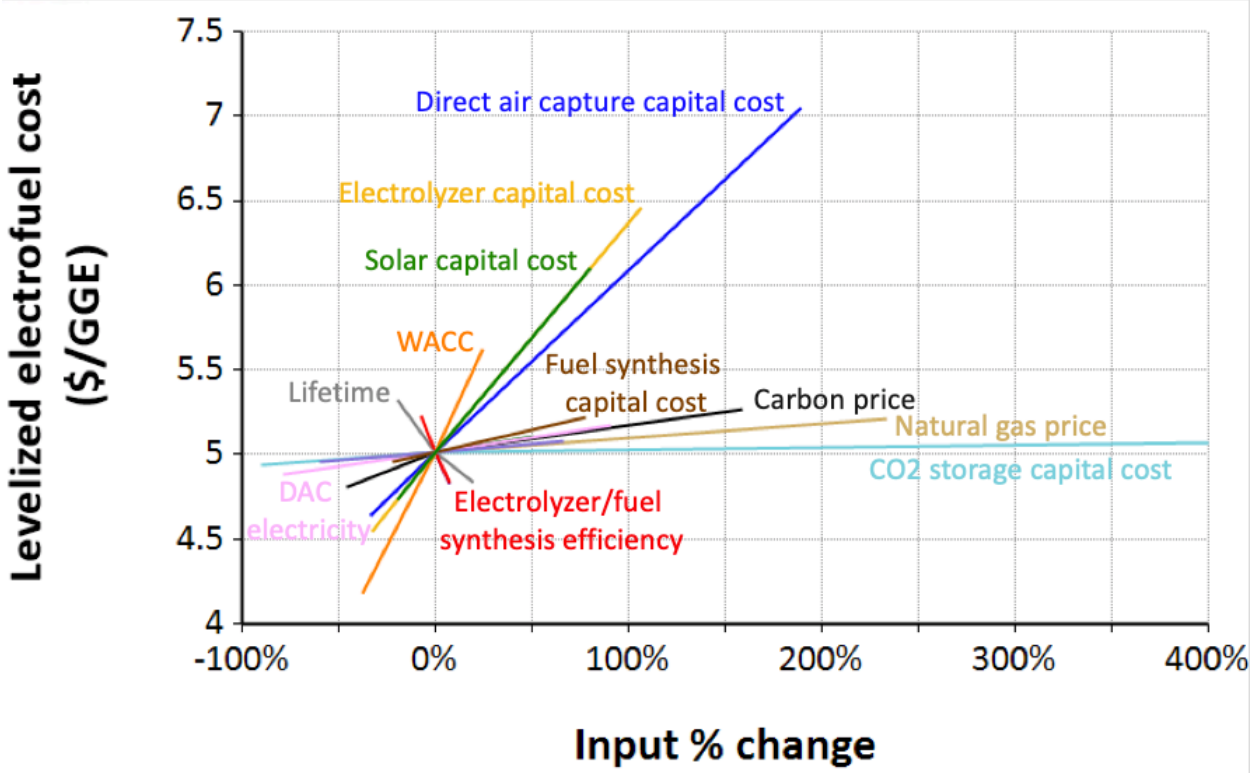

Figure 16. Sensitivity of the Next-decade high-temperature electrofuel production scenario, individually varying parameters to their Today and Breakthrough values. The four most sensitive parameters, which introduce production cost swings of over $1/GGE, are the capital and non-energy operating costs of direct air carbon capture system, the electrolyzer, the renewable electricity generation. The weighted average cost of capital and system lifetime, which substantially influence the levelized cost of these capital components, also have a major impact. Production cost is highly sensitive to electrolyzer and fuel synthesis efficiency, but the range of likely values is smaller, leading to a cost swing of less than $0.50/GGE. Fuel synthesis capital cost, natural gas price, and the carbon tax lead to cost swings of $0.22-0.46/GGE. The system is largely insensitive to changes in the cost of storage and interconnection to grid electricity and natural gas, largely because they can serve as substitutes for one another. Parameters with sensitivity less than $0.10/GGE are not shown.

4.5. Conclusions and policy implications

Hydrocarbon electrofuels have many characteristics that make them an attractive option for decarbonizing hard-to-reach sectors such as aviation, long-distance road travel, and heavy

shipping. They have a high energy density, which is particularly critical for aviation. They are chemically very similar to conventional jet fuel, diesel, and gasoline and can be used at high levels in existing engines with few if any engine modifications [112,164].

Using today's technology, the cost of electrofuel production is an expensive greenhouse gas mitigation option. However, if anticipated near-term reductions in the capital cost of DAC systems and electrolyzers materialize, an electrofuel production system co-located with the world's best solar resources, could fall below $6/GGE, implying greenhouse gas (GHG) mitigation costs that could compete with DAC at roughly $200/t($CO_2$) if the price of conventional fuel rises to $3/GGE.

Electrofuels could be competitive at lower conventional fuel prices if sequestration costs are substantially higher than anticipated. Much of the literature suggests that sequestration costs will be well below $50/t($CO_2$) [83]. However, many of these estimates do not account for reservoir pressure management techniques such as brine removal that may be necessary to avoid induced seismicity, which could increase sequestration costs to $50/t($CO_2$) or higher [152,154].

In addition, public acceptance may place limitations on sequestration. Research based on focus groups, interviews, and surveys suggests that public acceptance concerns are likely to be similar to those surrounding oil and gas extraction, but with the sequestration industry still in its infancy, public acceptance remains highly uncertain [165].

Within several decades, foreseeable breakthroughs in the cost of DAC, electrolyzers, and other technologies could bring electrofuel costs below $3/GGE, potentially reaching cost parity with petroleum fuels. The following specialized technological advances would facilitate these cost reductions:

- Electric kilns capable of supplying 900°C process heat for calcination at scales of hundreds of megawatts, prioritizing ensuring high efficiency over capital cost reductions.
- Heat storage systems capable of operating at or above 900°C. Metallic phase-change materials are capable of achieving these temperatures but are not widely commercially available. As an alternative, $Si_{49}Mg_{30}Ca_{21}$ phase-change materials capable of storing heat at 865°C are more mature [127]. Such systems could use a specialized calciner to combust small amount of oxygen-fired H_2 at 3,200°C [166], using O_2 from electrolysis, to boost this stored heat to 900°C or higher. This would be particularly useful if natural gas is unavailable for process heat.
- Low-cost direct-current (DC) to DC inverters to supply solar PV electricity with lower losses to the electrolyzer and perhaps the electric kiln.
- Alkaline or PEM electrolyzers designed for intermittent operation, with long stack lifetimes that depend primarily on intensity of use, rather than requiring replacement after a fixed number of years. Large capital cost reductions may justify substantial efficiency penalties.
- Low-cost, high-efficiency CO_2 storage. Current industrial CO_2 storage tends to operate in space-constrained environments, generally opting for high-pressure liquid CO_2 tank storage. As a result, compressor capital costs represent a large fraction of total capital costs for a CO_2 storage system. Electrofuel production systems would likely be built in remote areas, such as deserts, with fewer space constraints. Low-pressure CO_2 storage may allow substantial reductions in compressor and material costs as well as compression energy requirements. At \$1000/t(CO_2), roughly the cost of current high-pressure steel tanks without compressors, CO_2 storage can function as bulk seasonal energy storage.

The above recommendations demonstrate the role of optimization-based techno-economic analysis in identifying research, development, and deployment (RD&D) priorities for deep decarbonization of specific sectors. This model also demonstrates the ability of such models to quantify the value of storage and other forms of system flexibility for capital-intensive assets that rely on variable resources.

Achieving a warming limit of 2°C or less will likely require the complete elimination of net GHG emissions from the global economy within roughly half a century. In this context, decision-makers seeking to transform the global energy system must make substantial long-lived investments under enormous and largely irreducible technology uncertainty [167].

Such models can serve as screening tools to identify infrastructure pathways, dead-ends, and cross-sector interdependencies to inform the deployment of long-lived energy infrastructure and avoid technology lock-in. For instance, this analysis suggests that it may indeed be possible to achieve deep decarbonization in aviation, long-distance road travel, and heavy shipping with at most minor changes to the existing vehicle fleet (See the SI, Section 7.10). Thus, the development of electric and hydrogen alternatives certainly deserves serious investigation, but deep decarbonization does not depend solely on accelerated turnover of long-lived and capital-intensive airplanes, heavy trucks, and ocean freighters.

4.5.1. High- or low-temperature DAC?

Under near-term and long-term cost targets, high-temperature and low-temperature DAC systems produce electrofuel at comparable cost. High-temperature DAC systems may have lower capital cost but require higher-temperature heat, with associated engineering challenges. Low-temperature DAC systems are more capital-intensive, but can reuse waste heat from other

processes, reducing energy costs. These cost targets are based on inherently uncertain engineering estimates and assume substantial cumulative deployment.

Low-temperature DAC systems are currently commercially available and can be deployed in smaller units of capacity than high-temperature systems.

Because high- and low-temperature DAC systems are largely substitutable, an RD&D strategy that pursues both technologies simultaneously is most likely to produce substantial reductions in the cost of electrofuel.

4.5.2. Land and water use

Replacing all US jet fuel with electrofuel from systems modeled here would require roughly 370-5,500 mi^2 of land using wind electricity, amounting to 0.01-0.08% of US land area and 0.2-2.9% of Wyoming, assuming very high-quality wind resources are available in sufficient quantities [75]. Solar electrofuels would require 12,000-13,000 mi^2, or 0.2% of US land area and 5.5-5.9% of Arizona [75].

Low-temperature systems have the added benefit of producing enough water to supply the electrolyzer, whereas high-temperature systems require substantial water input. However, even assuming relatively expensive desalinated water is used, water costs represent less than 2% of the levelized cost of electrofuel. Still, replacing all US domestic jet fuel consumption with electrofuel using high-temperature DAC would require roughly 1.4 Gt(H_2O)/yr, about 0.3% of US water consumption and 17% of Arizona's water current consumption [168]. To avoid placing undue strain on local water resources, this would likely require construction of pipeline infrastructure, perhaps to transport desalinated water from the ocean or another non-potable resource. This would likely increase electrofuel costs by at most a few percent. See the SI, Section 7.11 for further details.

4.5.3. Potential vulnerabilities of DACS

Another potential concern surrounding the use of DACS to offset emissions from petroleum-based fuels is bookkeeping uncertainty. With electrofuels, CO_2 is captured before fuel is combusted. Airlines and other industries would likely purchase credits for DACS, which may include the option of paying for credits before any CO_2 is actually captured and sequestered. Thus, a system that relies on DACS or other forms of CO_2 removal to offset difficult-to-abate emissions is susceptible to weaknesses in policy, legal, and financial institutions, e.g. a collapse in a carbon trading marketplace, in ways that electrofuels are not.

Many projected pathways to a 2°C or 1.5°C warming target require net negative GHG emissions. In such a scenario, CO_2 removal technologies dictate the marginal cost of GHG mitigation. The cost of safely sequestering CO_2 in widely available saline aquifers within sandstone formations may have a strong nonlinear dependency on the injection rate [154]. If injection rates are already high, the marginal cost of sequestration could rise to \$50/t(CO_2) or higher [154], imposing a diminishing value on the expanded deployment of DACS. This is of particular importance for natural gas-fired DACS, which requires sequestration of at least 0.3 t(CO_2) from natural gas per t(CO_2) captured from the atmosphere [97].

4.5.4. Electrofuels as a pathway to DACS

Finally, electrofuel production represents a way to begin deploying DAC systems commercially at scale. A decarbonization pathway that fosters the electrofuel industry in the next few decades and could create a market for DAC, engendering cost reductions through innovation and learning through deployment. This would allow policy-makers to make more informed evaluations of the likely cost of mature DACS systems and the viability of mass deployment of standalone DACS as a mitigation strategy.

If aggressive RD&D in a suite of component technologies succeeds in fostering anticipated breakthroughs, near-term investment in electrofuels could result not only in cost-effective DACS, but also in low-cost electrofuels as a direct competitor with petroleum fuels, assuming oil prices only modestly above their recent ten-year average. In addition, low-cost electrofuels, particularly methanol or dimethyl ether, could serve as a dispatchable, seasonally storable, easy-to-transport, carbon-neutral energy carrier capable of enabling integration of high levels of variable renewable electricity into the electric power system.

4.6. Data availability statement

Annual solar and wind electricity production data are derived from the National Renewable Energy Laboratory's publicly available System Advisor Model using settings described in Section 4.3.1.1 [86].

5. Chapter 5. Conclusions and policy implications

Deep decarbonization compatible with a global warming limit of 2°C or less is likely achievable with a concerted, globally coordinated effort using the technologies that exist today and those that develop through a major research, development, and deployment (RD&D) effort. There are many conceivable pathways available to us, surrounded by enormous uncertainties. Given the long lifetime of energy infrastructure, pursuing a pathway to a dead-end could be incur substantial additional cost and may in fact render a desired warming limit unattainable.

Chapter 2 characterizes the depth of important uncertainties in the future production, consumption, and price of key energy quantities. Despite the commendable modeling efforts of the Energy Information Administration (EIA), whose National Energy Modeling System represents many facets of the US energy system in great detail, our ability to predict these quantities more than a few years into the future has not improved. If anything, emerging volatility and a host of technological and economic surprises have rendered the energy system even more difficult to predict than in the past. As such, a deep decarbonization strategy that relies heavily on precise predictions a decade or more into the future will likely be very vulnerable to uncertainty at a time when the recent history of the US energy system has been defined by surprises such as shale oil and gas, the financial crisis, the boom and bust cycle in natural gas, and the rise of low-cost wind and solar electricity.

Chapter 3 uses a wealth of observational data to estimate the demand side effects of a low-income electric subsidy in California, characterizing the associated effect on electric power system emissions and peak capacity requirements. Policy makers considering equity measures such as energy subsidies can use such estimates to foresee these indirect effects and account for them when crafting emissions reduction targets and resource adequacy plans.

Chapter 4 provides practical intuition into the future viability of electrofuels as a climate mitigation strategy. The analysis suggests that if direct air capture (DAC) is able to achieve targeted cost reductions, electrofuels could be a cost-effective alternative to simply using DAC with geologic sequestration of the captured CO_2 to offset emissions from aviation and other difficult-to-decarbonize applications that require liquid fuels. This is particularly true if the cost of electrolysis of water and renewable electricity continues to decline.

Perhaps more importantly, this analysis demonstrates the potential role of optimization-based techno-economic models in guiding RD&D investment in new technologies. The process of building the PtXOpt model forced me to think through every step of the engineered system I was designing. Can we use electricity to efficiently produce 900°C heat? How can we store that high-temperature heat? What are our options for storing massive amounts of CO_2 for months at a time? How much would all of this cost?

In some cases, it turns out that the cost of electrofuel production is largely insensitive to the likely range of costs of an emerging technology. Small-scale electric tunnel kilns currently cost roughly \$125/kW[107], and costs would likely go down at larger scales. Even at this cost, electric kiln capital cost represents less than 1% of total cost for an electrofuel system heated entirely with electricity.

In other cases, commercialization of new enabling technologies, such as metallic phase-change materials for high-temperature heat storage or low-cost, particularly if natural gas is unavailable for process heat. There is currently little to no market pull for such technologies at scale. If we would like to use them in the future, the time for RD&D is now.

Skillfully navigating toward deep decarbonization in the face of massive uncertainties requires three things. First, we must acknowledge the magnitude of what we do not know and strive to

characterize the large uncertainties we face. Second, we must use the wealth of data available to us to learn as much as we can about the fundamental dynamics of different components of the energy system, including human behavioral responses to changing incentives, so we can incorporate these factors into energy policy and decision-making. Third, we must use thoughtfully-constructed, interpretable quantitative models to gain insight into the likely effects of a wide range of energy technology and policy futures.

Taken together, these elements form a blueprint not for a golden road to deep decarbonization but for an adaptive decision-making framework to guide each step we take, using the results of each step to inform the next one. Such an approach can help balance tradeoffs between focusing limited resources on accelerating the deployment of the solutions that are most effective at the moment and investing in a wide array of technology and policy options to keep many pathways available.

5.1. Future work

The field of deep decarbonization research is in many ways reaching a new level of maturity since the first modern integrated assessment models were built thirty years ago. Exponential growth in data availability and computing power have laid the foundation for sophisticated analyses that only the most forward-thinking researchers could have conceived of at that time. This allows us to look back with fresh eyes on past energy policies and learn important lessons. At the same time, we now have first-hand experience with the potential for surprisingly rapid development of low-carbon technologies such as wind, solar photovoltaics, and lithium-ion batteries, as well as the potential for stagnation in much-anticipated technologies such as advanced nuclear and fossil electricity with carbon capture with sequestration.

Each of the three studies in this dissertation raises more questions than it answers. I list what I think are some of the most interesting research questions. I plan to work on at least some of them in the next stage of my career.

A decision-making framework capable of thoughtfully incorporating the uncertainties inherent in the future of the energy system must assimilate the particularities of the decisions at hand. There is little in the academic literature that documents how long-term energy projections and forecasts are actually used to make energy policy and investment decisions related to long-lived energy infrastructure. Before proposing changes to these decision-making practices, we need to understand how they currently work. To this end, I propose a series of interviews with analysts, managers, and forecasting departments at major energy companies and government regulatory agencies, such as public utilities commissions, characterizing how in-house and government long-term energy projections are actually used to inform energy investment and policy. Such research would draw on the management science and organizational behavior literature, informed by detailed knowledge of the engineering characteristics of the relevant components of the energy system for each company or agency interviewed.

Energy utilities and policy makers have only begun to harness the potential insights available through energy consumption data from advanced metering infrastructure. Chapter 3 demonstrates both the potential value and the limitations of simply making observational data available to academic researchers. While much can be learned from observational data, the most pertinent questions for those planning an energy transition are often causal relationships, which are difficult to estimate accurately without an experiment. I hope to work with energy utilities to implement experimental randomized controlled trials to accurately measure the responsiveness of residential energy demand to changes in prices and other interventions. This work would draw

on the experimental microeconomics literature, potentially employing machine learning techniques such as those developed by Athey et al. [169] to more accurately measure heterogeneous responses to different interventions.

The degree to which an energy decarbonization pathway succeeds will depend in large part on RD&D decisions for current and emerging technologies. Chapter 4 is a case study in the potential value of optimization-based techno-economic models for guiding RD&D decisions related to electrofuel production systems and component subsystems such as DAC. I hope to continue developing and using this and other models to further characterize the tradeoffs associated with potential pathways toward deep decarbonization.

In particular, I am interested in assessing the value of maintaining existing fossil pipeline distribution infrastructure as an option in case of major cost reductions in electrofuel production. Liquid or gaseous hydrocarbon or oxygenate electrofuels have numerous properties that make them attractive for certain difficult-to-decarbonize applications. Liquid fuel pipelines transport energy at one to two orders of magnitude lower cost than electric transmission lines and do not require phase or voltage synchronization. Liquids can also be stored seasonally at low cost. These attributes make low-cost liquid electrofuels an ideal energy carrier to provide flexible electricity in a low-carbon electric power system dominated by inflexible solar, wind, or nuclear electricity.

Building space and water heating is another such application. Electric heating options such as air-source heat pumps must currently switch to inefficient resistive heating at temperatures substantially below freezing. Power-to-gas, electrofuel methane, may thus be a more economical home heating option in very cold climates such as the US upper Midwest or much of Canada. However, it may be a few decades before power-to-gas achieves the necessary cost reductions.

As a result, a decarbonization strategy that prioritizes the expedited elimination of fossil fuels could result in the retirement of aging urban natural gas distribution infrastructure, essentially removing power-to-gas as a home heating option even if it achieves the requisite cost reductions. In addition, many of the component technologies studied in Chapter 4 have not yet been studied in detail in the context of electrofuel production. In many cases, detailed engineering analysis and expert elicitations would likely produce innovative component designs, new ideas for RD&D spending and a better-informed range of possible future costs.

In conclusion, addressing these and similar questions is likely too large a task to fit into a single researcher's career (even with a small army of graduate students). I look forward to pursuing some of these questions and working with the growing community of deep decarbonization decision researchers as we continue to figure out the nuts and bolts of solutions to one of the great problems of our age.

6. Supporting information for Chapter 2

6.1. Supplementary Note 1

The longest-running AEO side cases account for low and high oil price and low and high economic growth. Supplementary Figure 1 shows high and low constant-dollar oil price scenarios from AEO 1999 through AEO 2014 (years for which these data are digitally available). Supplementary Figure 2 shows the ratio of high to low oil price in these two scenarios. Note that the range between low and high scenarios begins to grow in AEO 2006 and rises substantially thereafter. This is no doubt influenced by the rapid rise in oil prices that began in the mid-2000s.

6.2. Supplementary Figures

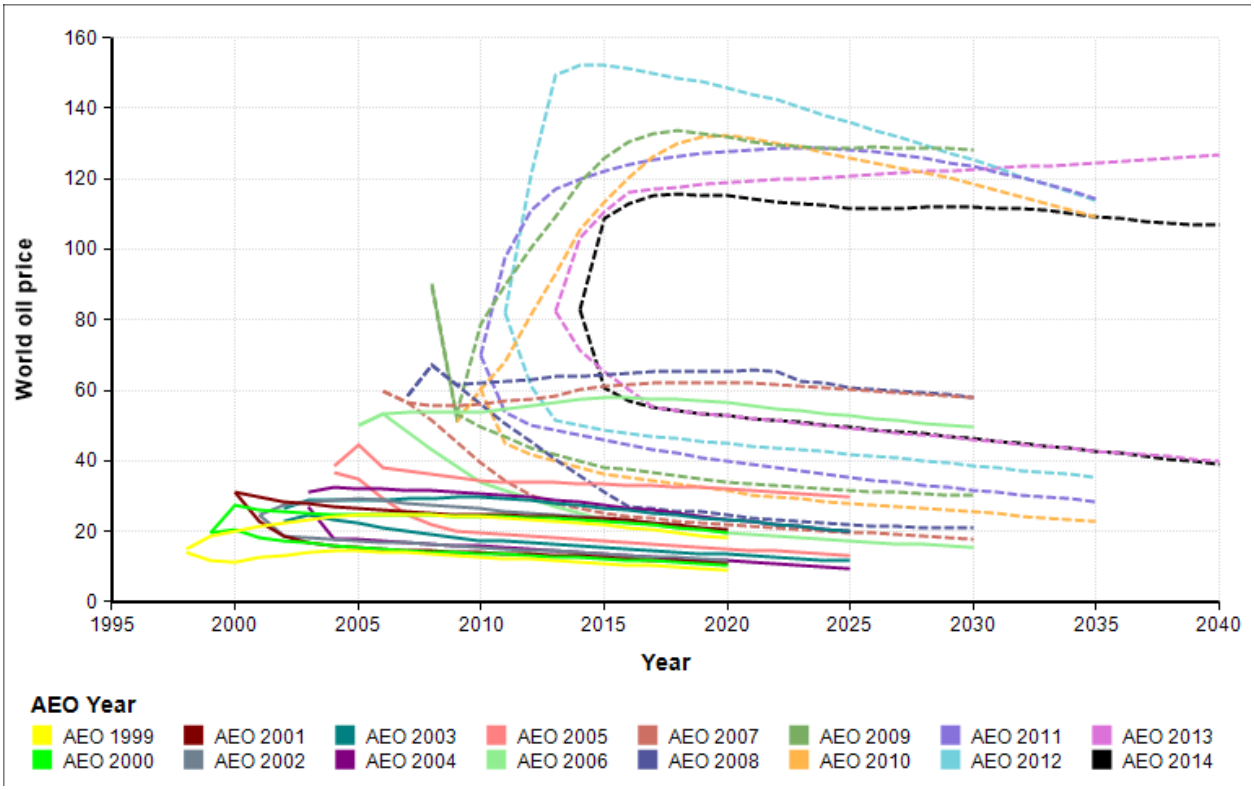

Figure 17. High and low oil price scenarios from AEO 1999-2014. a) displays high oil price scenarios, while b) displays low oil price scenarios. Note that the scenarios become increasingly divergent from AEO 2006 onward. Price is reported in constant 2005 dollars.

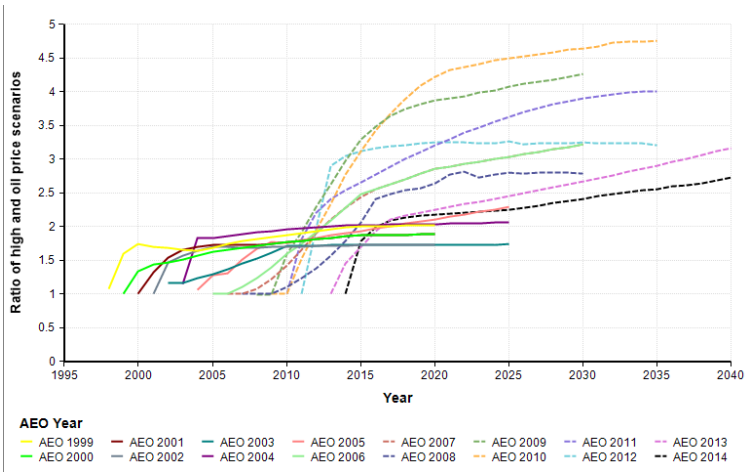

Figure 18. The ratio of high to low oil price scenarios. Note this ratio begins increasing sharply with AEO 2006, and grows to a maximum in AEO 2010.

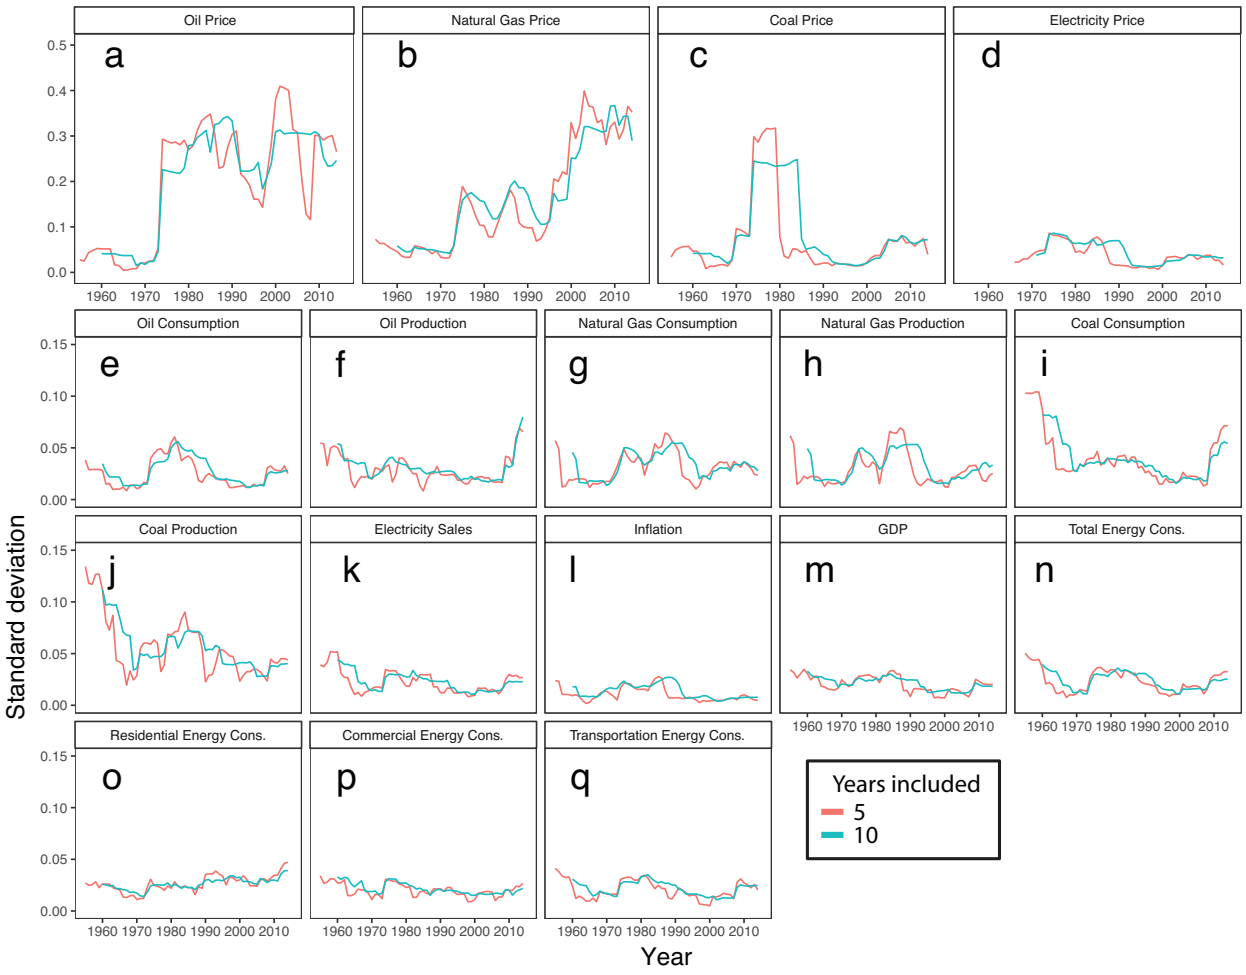

Figure 19. Five- and ten- year rolling standard deviation for all seventeen quantities. Note that there is a local minimum between 1995 and 2004 followed by a subsequent increase for all but three quantities (inflation, l, oil price, a, and natural gas price, b).

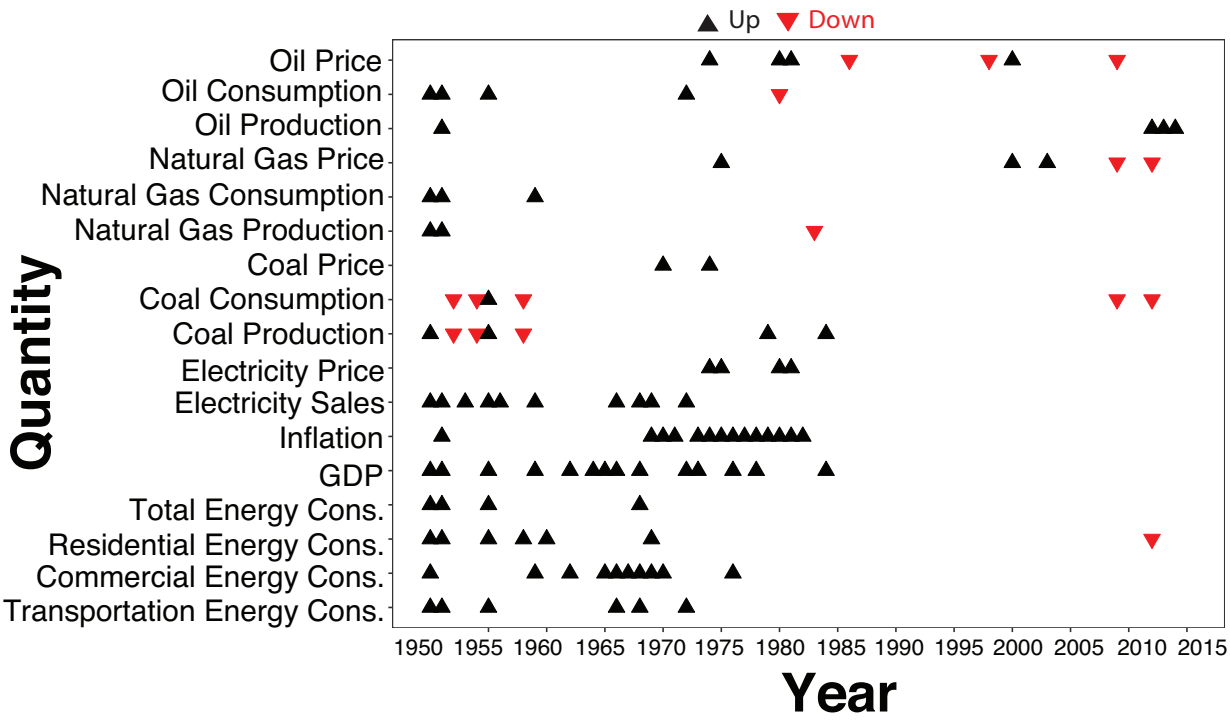

Figure 20. Extreme changes defined as a year with a year-over-year change of greater than 2σ or less than $1 - (1+2\sigma)^{-1}$ for each quantity. 1985-2005 has the smallest number of extreme changes of any period. 2005-2014 has the largest concentration of negative extreme changes of any decade.

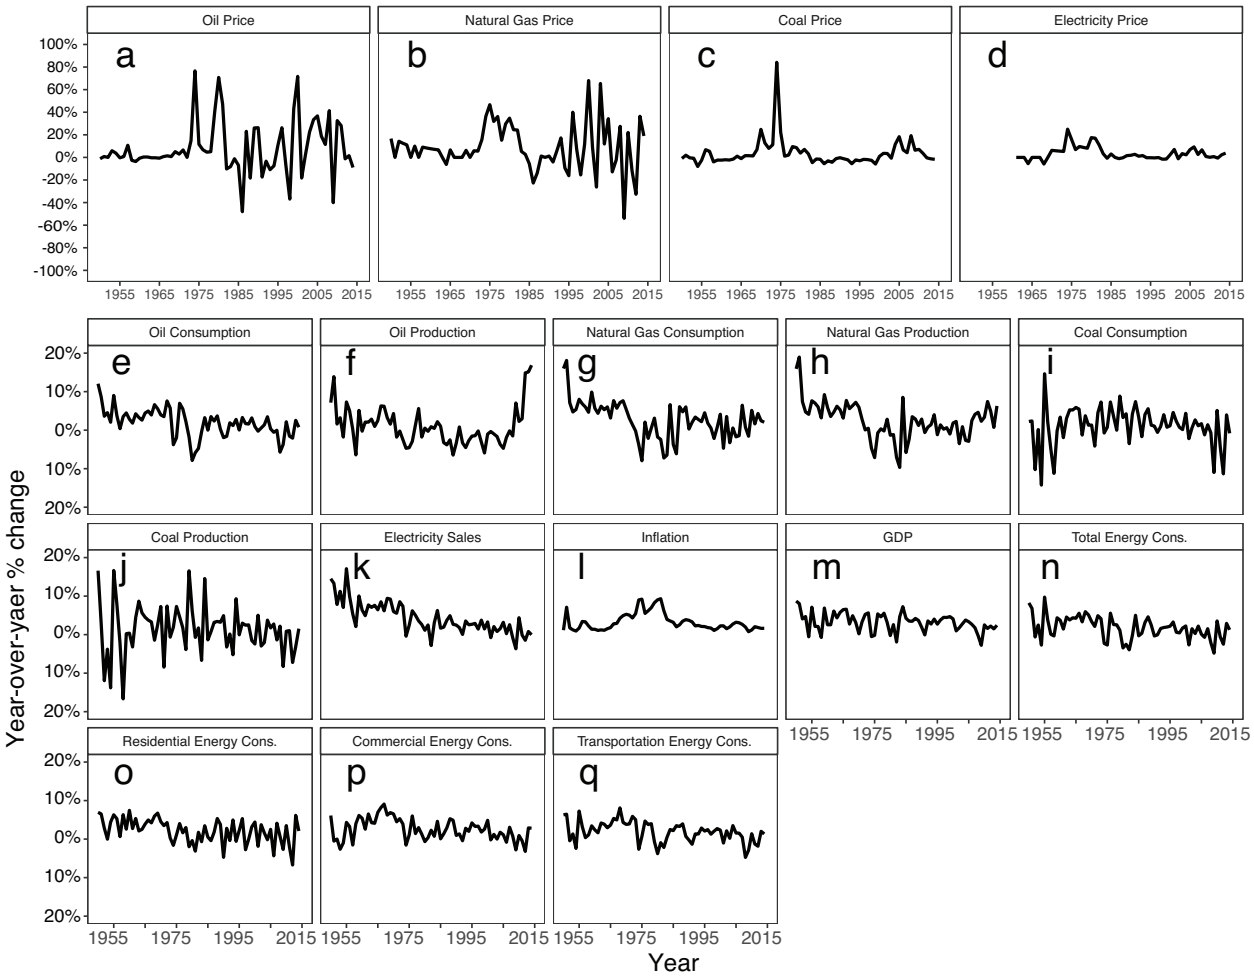

Figure 21. Year-over-year changes for the seventeen quantities examined. Data are from 1949-2014, except for electricity price, which begins in 1960. Note that price quantities (a-d) have substantially higher variance than non-price quantities (e-q). Note that eight quantities, natural gas price, electricity sales, inflation, GDP, and total, residential, transportation, and commercial energy consumption, have the most negative year-over-year change between 2005 and 2014.

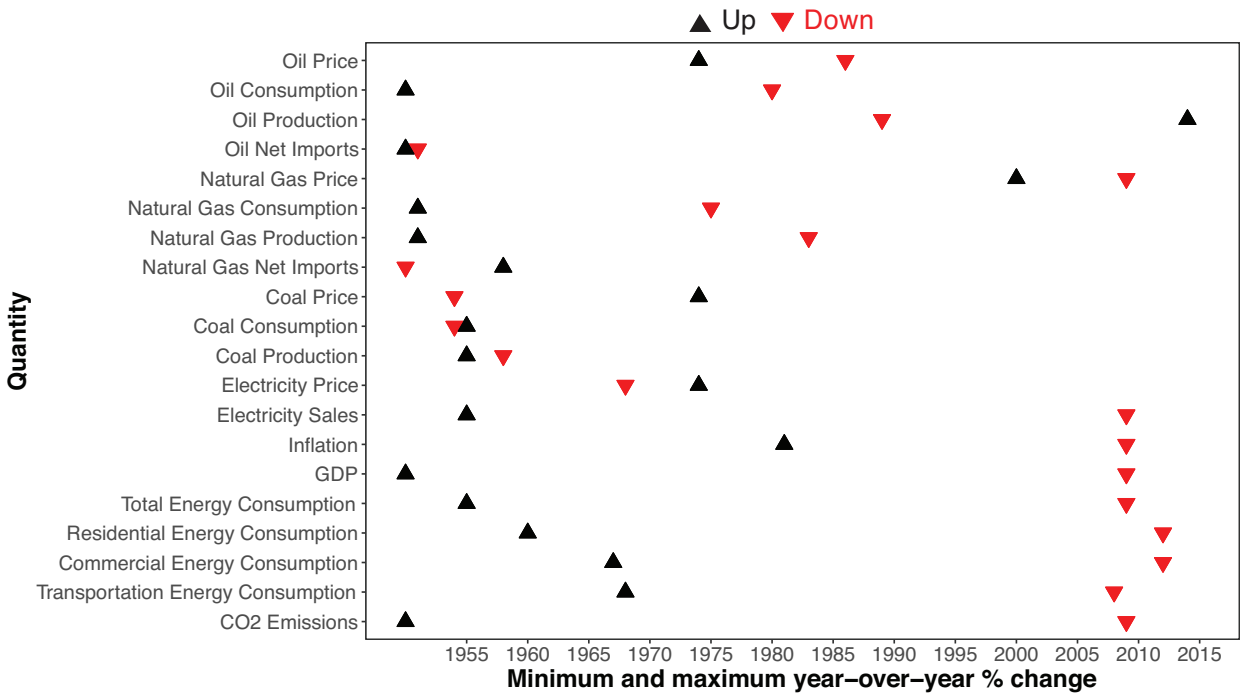

Figure 22. Extreme changes for twenty energy quantities, including three derivative quantities: Oil and natural gas imports, and CO2 emissions, from 1949 to 2014. The black and red triangles indicate the largest year-over-year increases and decreases in the quantities. Ten of forty extreme changes occur in 2005-2014.

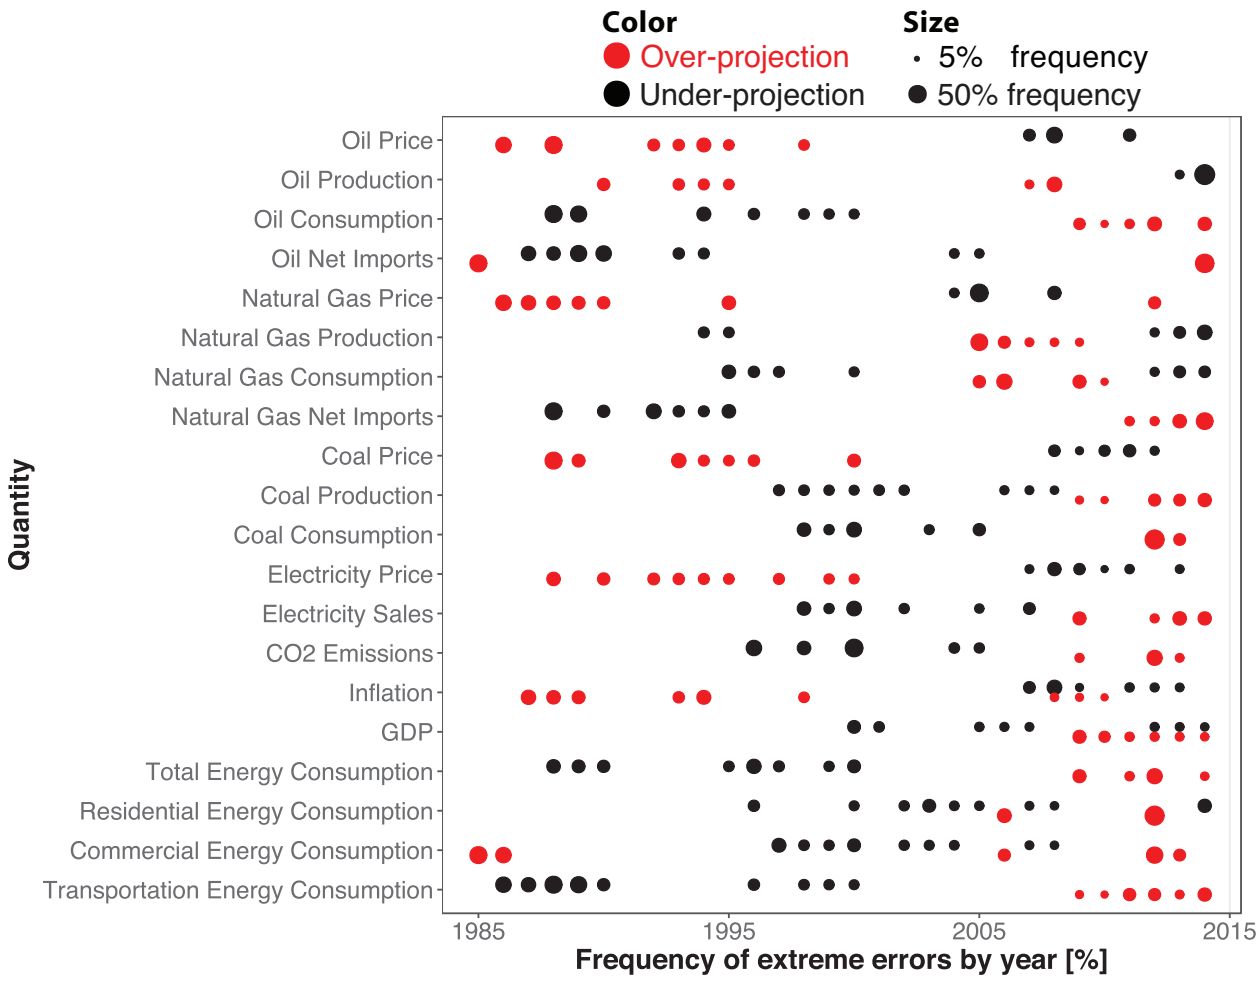

Figure 23. Annual frequency of extreme errors for twenty quantities, including three derivative quantities: Oil and natural gas imports, and CO2 emissions. Red and black circles correspond to over-projected and under-projected extreme errors, respectively. The size of each circle corresponds to the frequency of extreme errors in that year. Note the high concentration of extreme errors in 2005-2014.

Figure 24. The percentage point change in the frequency of extreme errors in 2005-2014 relative to 1985-1994 and 1995-2004 for different definitions of "extreme error" and subsets of the data, described in Supplementary Note 6. A positive number is an increase in the frequency of extreme errors in 2005-2014. Note that most quantities show increases in most scenarios. Note that some scenarios have a higher base frequency of extreme errors than others, so comparing the relative magnitude of an increase or decrease across scenarios may not be meaningful.

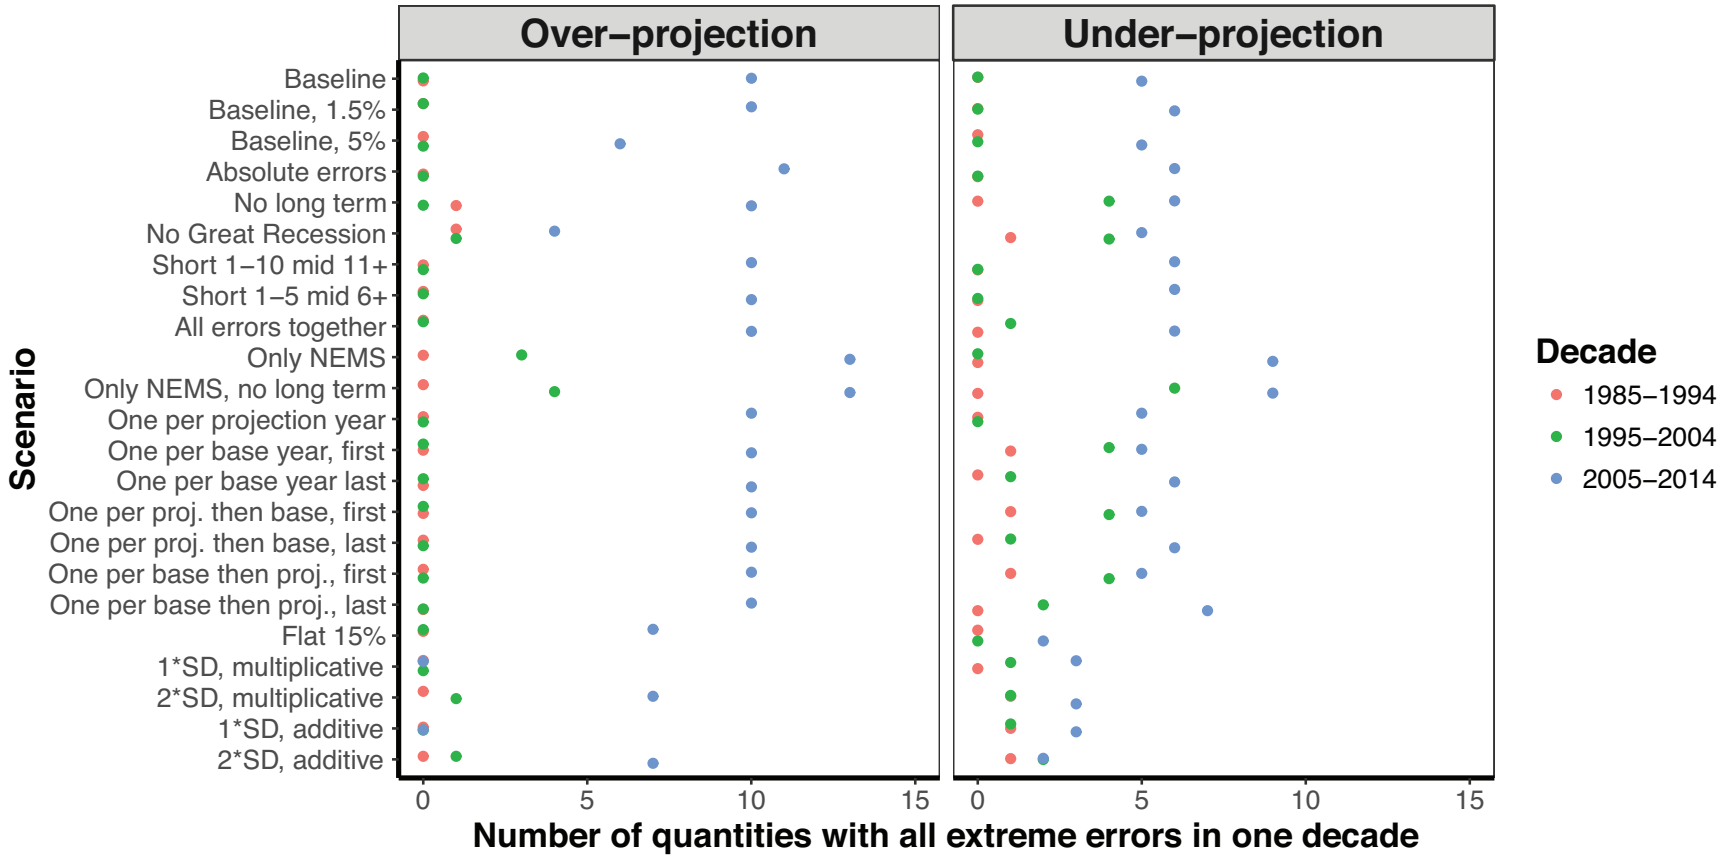

Figure 25. The total number of quantities, of the seventeen in the main manuscript, with all over-projected or under-projected extreme errors in a single decade, for each scenario. Note that the number in 2005-2014 is greater than in the previous two decades for both over-projected and under-projected extreme errors in all scenarios.

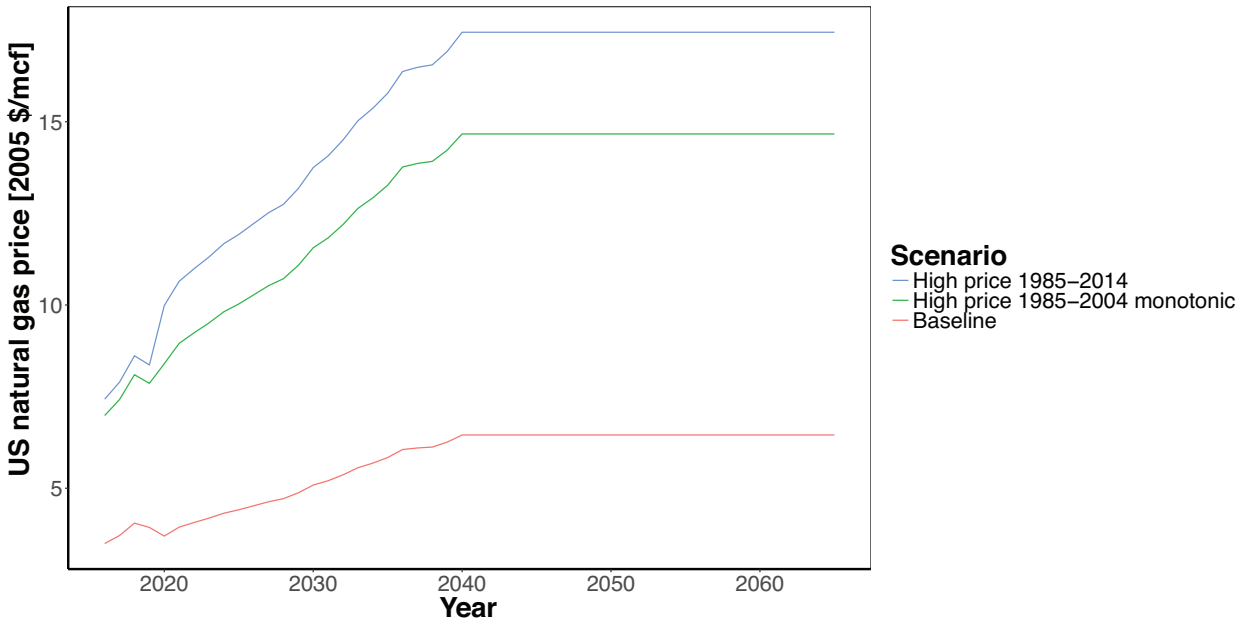

Figure 26. Domestic US natural gas price scenarios. Baseline is the AEO 2015 reference case, with constant interpolation after 2040. Two high price multiplier scenarios are considered, one based on the 2.5th percentile of historical errors from the AEO from 1985-2014, and one based only on 1985-2004, with an enforced monotonicity constraint to ensure the price multiplier increases or remains constant as projections go further into the future.

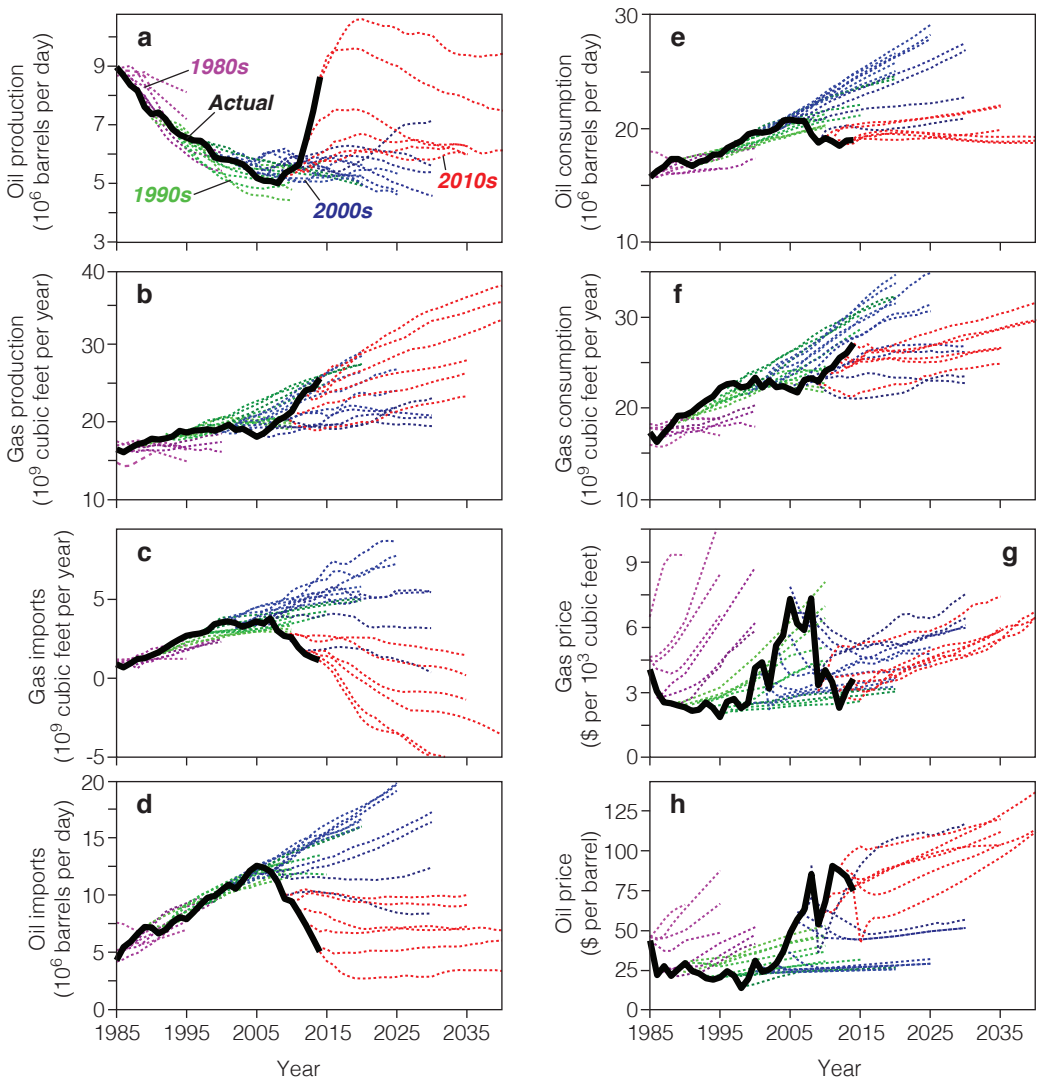

Figure 27. Annual Energy Outlook projections v. observed historical values for oil and natural gas quantities. Shows observed historical values (black solid line) for all quantities, and AEO projections out to 2040 (dashed lines). Most quantities tend to have distinct periods of over-projection and under-projection. [bbl = barrel; Mbbl = million barrels; Mcf = thousand cubic feet; tcf = trillion cubic feet]. Prices are in constant 2005 dollars.

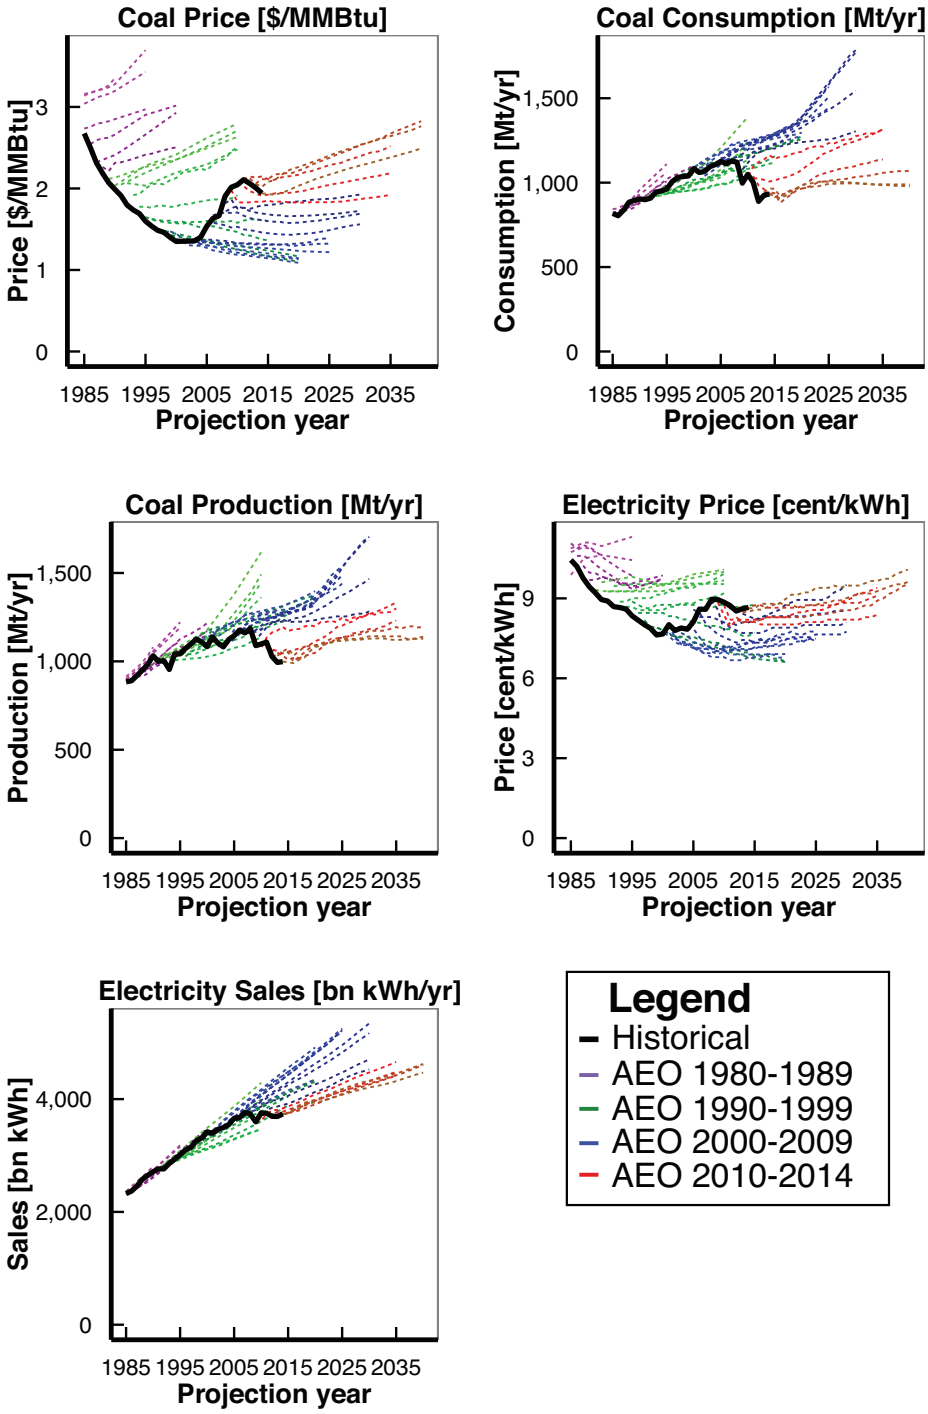

Figure 28. Annual Energy Outlook projections v. observed historical values for coal and electricity quantities. Shows observed historical values (black solid line) for all quantities, and AEO projections out to 2040 (dashed lines). Most quantities tend to have distinct periods of over-projection and under-projection. [MMBtu = million British thermal units (Btu); Mt = million short tons; kWh = kilowatt-hour]. Prices are in constant 2005 dollars.

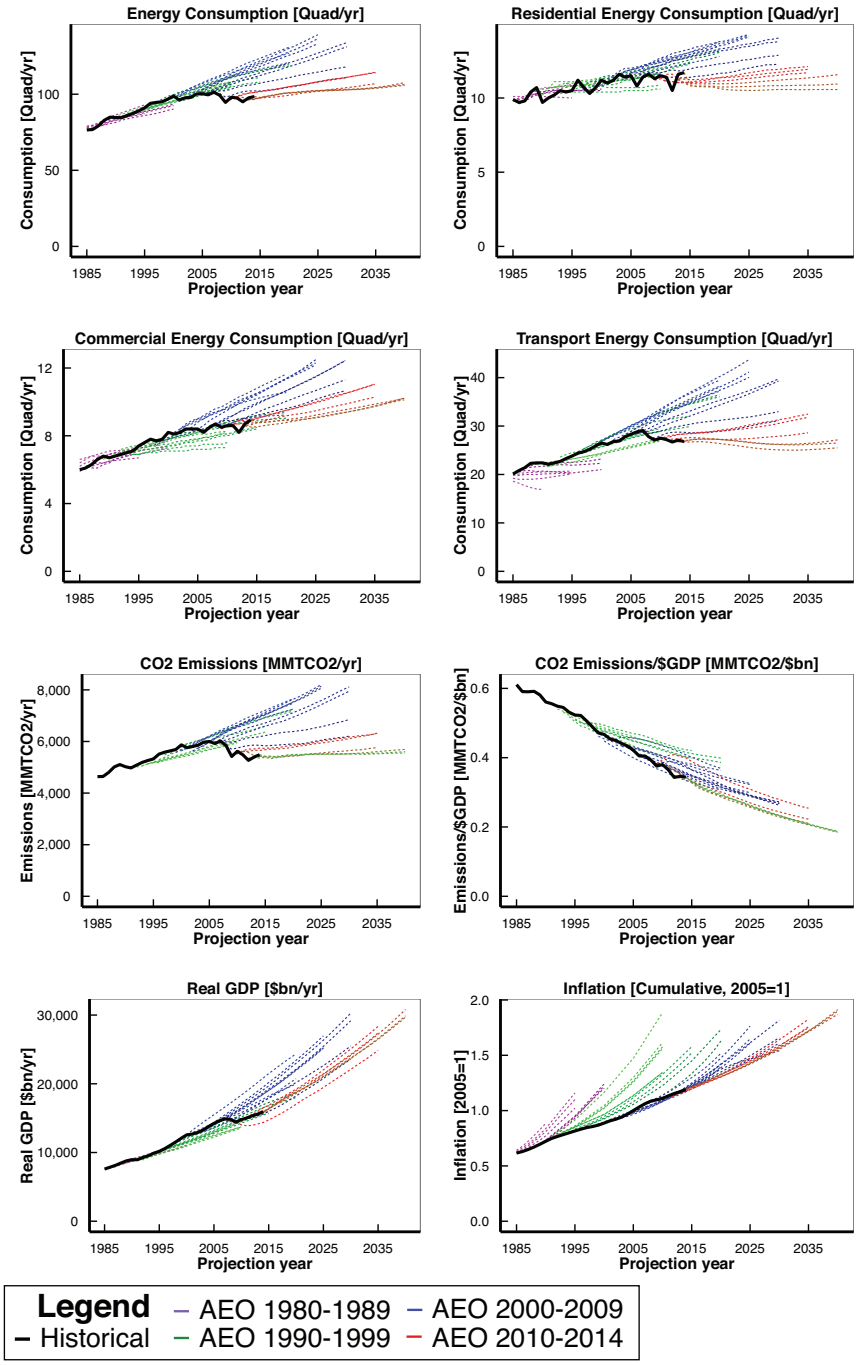

Figure 29. Annual Energy Outlook projections v. observed historical values for energy consumption by sector, CO_2, and macroeconomic quantities. Shows observed historical values (black solid line) for all quantities, and AEO projections out to 2040 (dashed lines). Most quantities tend to have distinct periods of over-projection and under-projection. [Quad = quadrillion btu; $MMTCO_2$ = Million metric tons carbon dioxide equivalent].

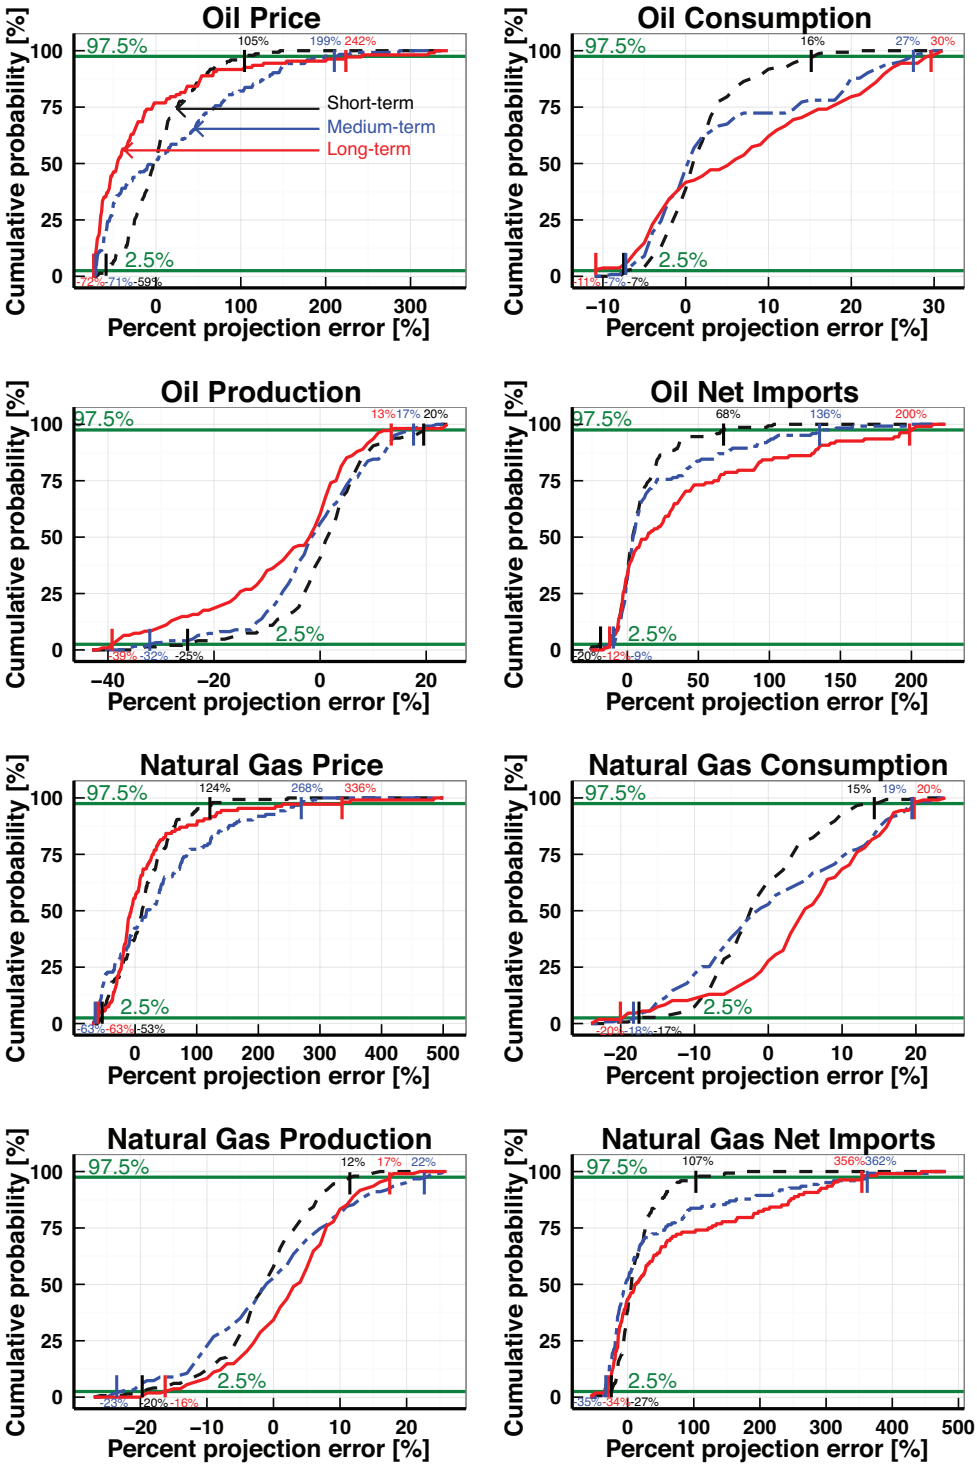

Figure 30. Error cumulative distribution functions and extreme error thresholds for oil and natural gas quantities.

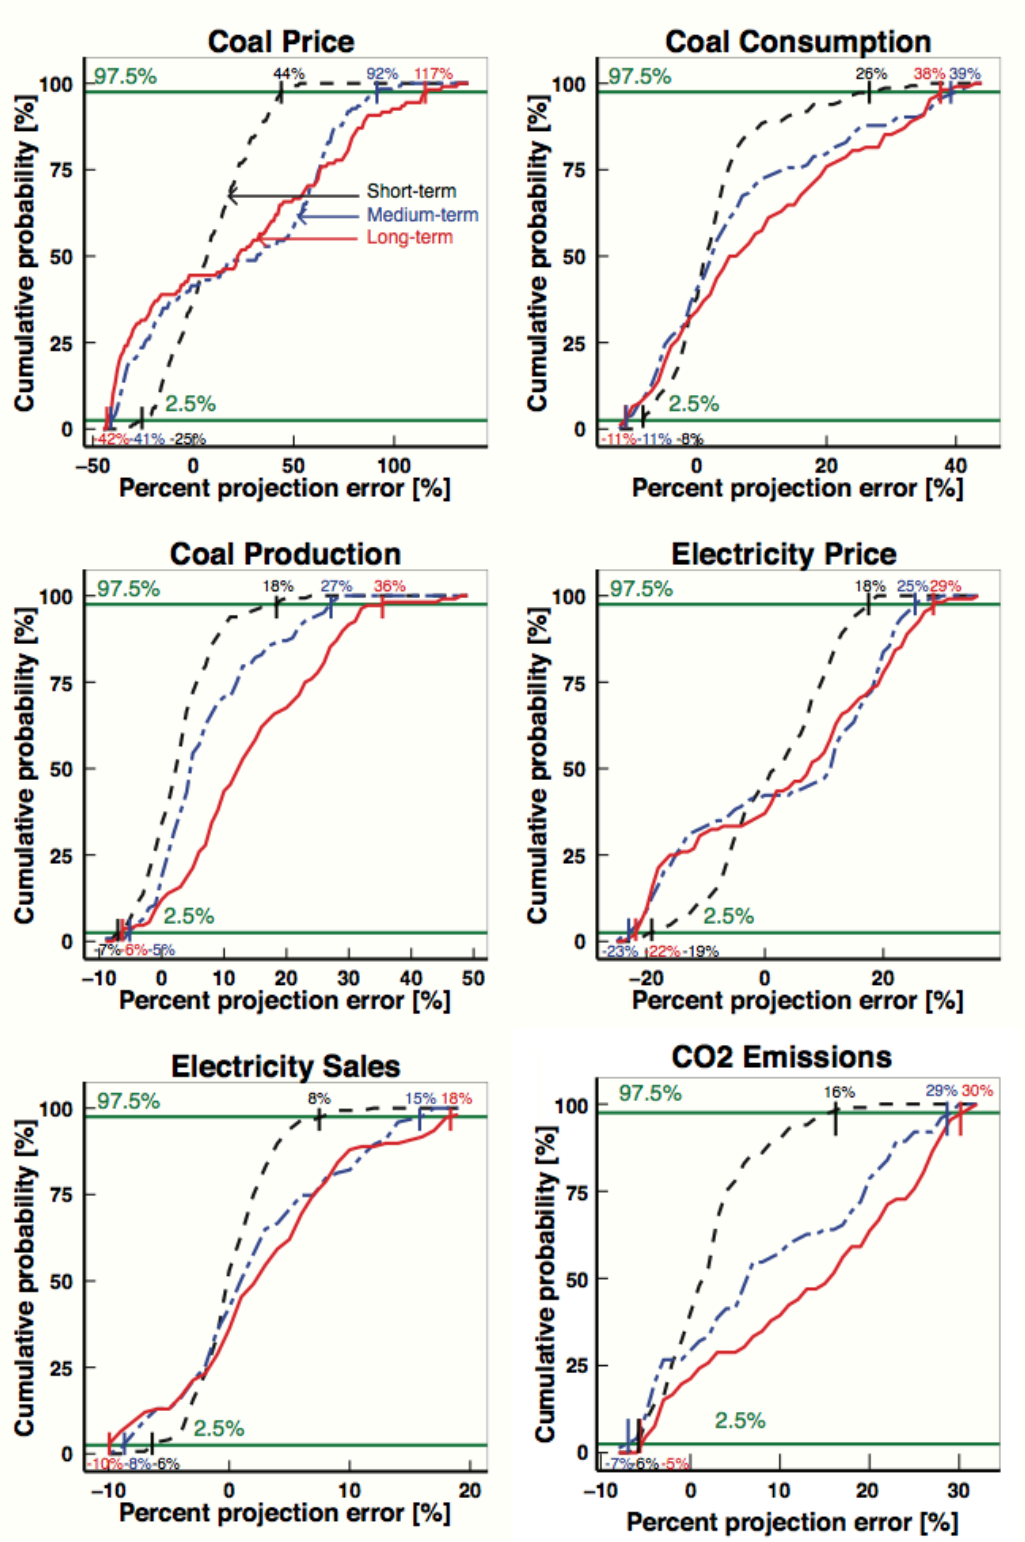

Figure 31. Error cumulative distribution functions and extreme error thresholds for coal and electricity quantities and CO_2 emissions.

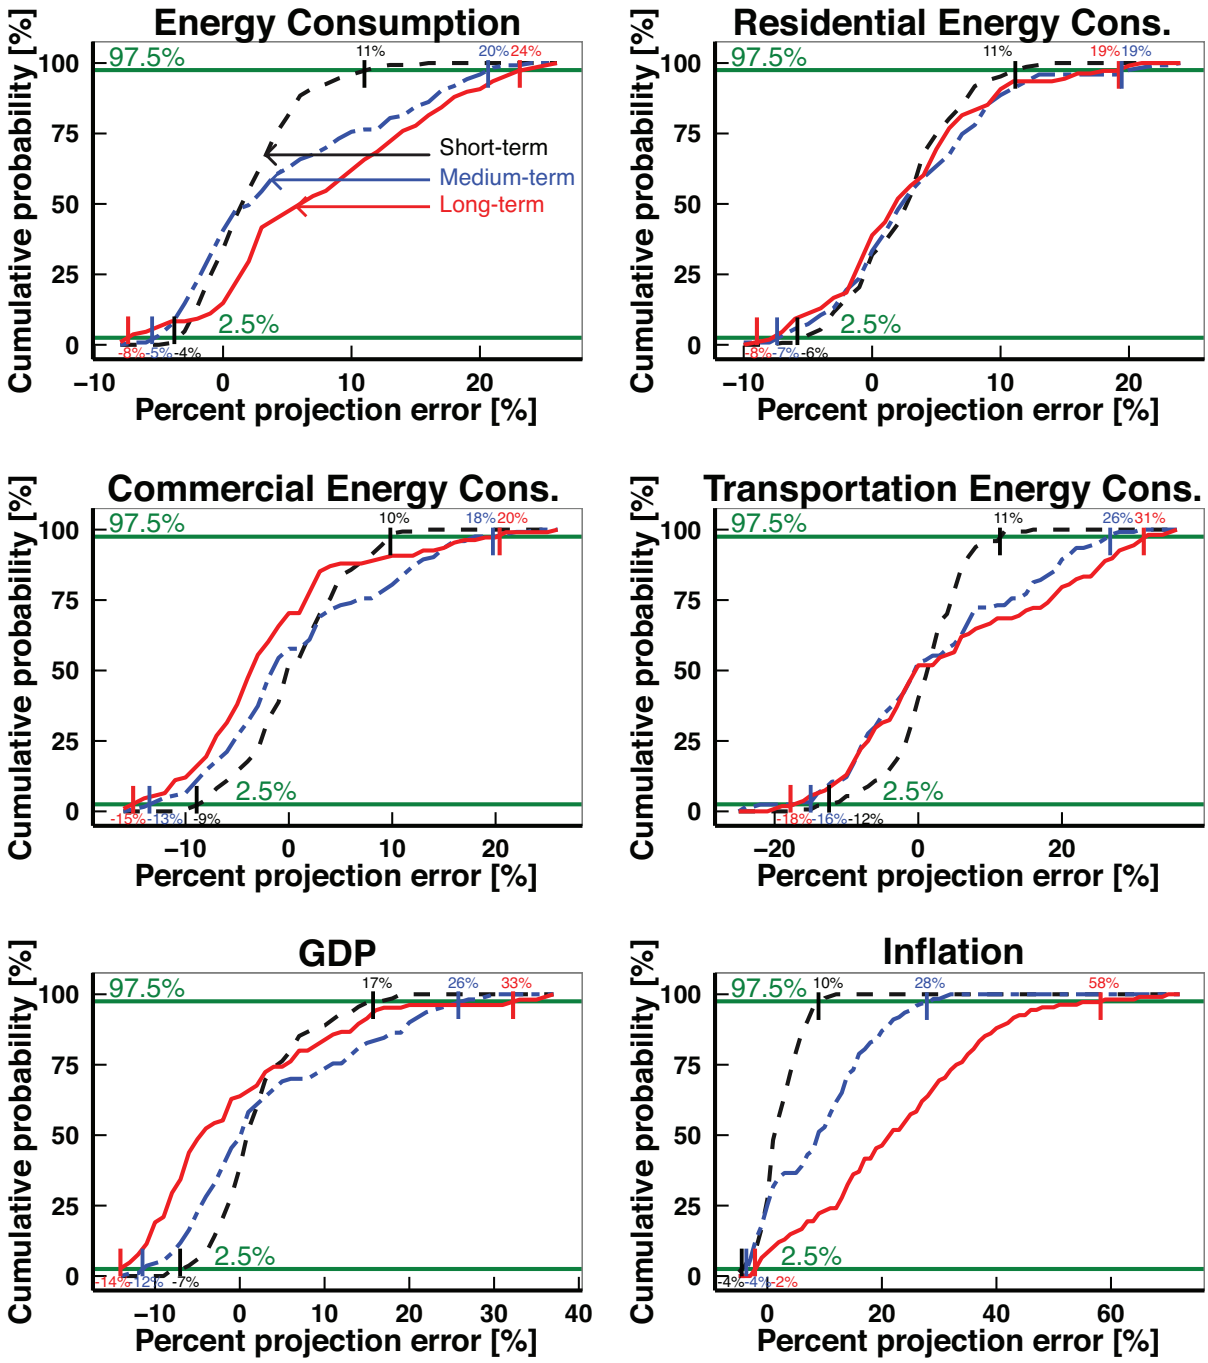

Figure 32. Error cumulative distribution functions and extreme error thresholds for energy consumption by sector, and macroeconomic quantities

6.3. Supplementary Note 2

The process by which the AEO reports are generated is not stationary over time in the statistical sense (see Supplementary Note 4). However, the EIA, which produces the AEO, has remained relatively stable as an organization, producing the AEO over the years with a consistent mission and approach. As a result, we claim that the AEO projections over the past three decades are the result of a relatively stationary organizational process, and thus it is appropriate to consider them as a single dataset. The legal mandate behind the AEO has remained constant over time. Section 57.a.2 of the Federal Energy Administration Amendments (FEAA) of 1976 charges the Director with producing:

> "[A]n annual report which includes … short-, medium-, and long-term energy consumption and supply trends and forecasts under various assumptions; and, to the maximum extent practicable, a summary or schedule of the amounts of mineral fuel resources, nonmineral energy resources, and mineral fuels that can be brought to market at various prices and technologies and their relationship to forecasted demands[170]."

The Department of Energy Organization Act of 1977, which replaced the Federal Energy Administration with the newly established Energy Information Administration, transferred this responsibility to the EIA in Section 205.c[171]. Fulfilling this obligation, every AEO report includes a reference case projection of the future trajectory of the US energy system assuming no departure from current policies, along with alternative scenarios[35,172].

Since its inception, the EIA has used techno-economic models of the US energy sector, such as the National Energy Modeling System, described below, to produce its projections[173]. Sections

52.a and 54.b.4 of the FEAA of 1976 mandate that the EIA maintain such a model[171]. In addition, the organizational and funding structure of the EIA was designed to ensure a high degree of organizational independence from other government agencies[174]. As a result, although there is not stationarity in the models used to generate projections, the mission of the EIA, and its organizational structure have been relatively stable since its creation in 1979. On this basis we assert that it is reasonable to treat AEO projections from 1982-2014 as a unified dataset for statistical analysis.

6.4. Supplementary Note 3

AEO projections come from the National Energy Modeling System (NEMS), a large energy-economic model of the US energy system, produced and maintained by the EIA. NEMS consists of 13 separate modules, representing different energy supply, demand, and conversion components of the US energy system. NEMS uses a computable equilibrium approach to determine a set of energy production, consumption, and price quantities that equilibrate the US energy market[175]. Prior to the introduction of NEMS in AEO 1994, the EIA used the Intermediate Future Forecasting System to produce AEO projections. IFFS had a modular structure, similar to NEMS, and used a similar method to locate an equilibrium[176].

NEMS has four energy supply modules, representing Oil and Gas Supply, Natural Gas Transmission and Distribution, Coal Market, and Renewable Fuels. The four demand modules are Residential, Commercial, Transmission, and Industrial Demand. The two energy conversion modules represent the Electricity and Liquid Fuels markets. The Macroeconomic Activity

Module models economic activity at national, regional, and industrial levels, producing projections of economic variables that drive energy supply and demand quantities[177]. The International Energy module includes major trends in the global oil market, and the effects of those trends on the US oil market[178]. The integrating module controls NEMS' solution algorithm and manages data flows between modules. All NEMS modules are fully documented and publicly available, with the exception of the Macroeconomic Activity Module, which relies heavily on the proprietary IHS Global Insight Model of the US Economy[175]. NEMS employs the Gauss-Seidel method of solving simultaneous linear equations, equilibrating modules iteratively, replacing module inputs with outputs from previously solved modules. Although some equations in NEMS are nonlinear, NEMS is expected to converge to an equilibrium because all equations are monotonic, either increasing or decreasing[175]. All of the quantities we analyze are computed endogenously in NEMS. The general GDP growth path is specified exogenously, but projected GDP growth is computed endogenously in the Macroeconomic Activity Module[177].

6.5. Supplementary Note 4

We wish to evaluate the significance of our key finding that year-on-year volatility and unpredictability increased in 2005-2014. The data do not allow traditional tests of statistical significance, such as t-tests, which would require that historical and projected values of energy quantities be statistically independent and identically distributed, generated by a process that is stationary over time.

The US energy system is a highly non-stationary system, constantly changing in unpredictable ways. The massive increase in oil and natural gas production from hydraulic fracturing is a

testament to this. In addition, the process that generates AEO projections is not stationary. The computational methods used to create AEO projections have changed over time, in accordance with advances in computational technology and changing national energy policy priorities.

Projected and historical values of energy quantities are highly correlated across quantities, particularly prices, production, and consumption of the same quantity (see Supplementary Table 3). In addition, projection errors for a single quantity from a single AEO projection, or for a single year are correlated (see Supplementary Note 9).

As a result, we assess the significance of our findings in two ways. We conduct a sensitivity analysis, described in Supplementary Note 6, and a Monte Carlo analysis, described in Materials and Methods.

6.6. Supplementary Note 5

Our primary definition of year-on-year volatility, based on the largest year-over-year percent increase and decrease for each quantity, allows straightforward comparison of values in a single year or small set of years. The more standard definition of volatility is based on the standard deviation of percent changes over time. We include two definitions of volatility based on such a standard deviation metric. The first computes annual volatility as the standard deviation over the previous five or ten years (see Supplementary Figure 3). The second defines extreme changes based on the standard deviation of year-over-year changes in each quantity (see Supplementary Figure 4).

Supplementary Figure 3 shows the five-year and ten-year standard deviation of each of the seventeen quantities considered in the primary analysis. The highest standard deviation in all quantities tends to occur before 1995. In both cases, all but three quantities (inflation, oil price, and natural gas price) have a local minimum between 1995 and 2004, followed by a subsequent increase. This suggests that by this metric as well, there is a general increase in volatility in 2005-2014 relative to at least the previous decade.

Supplementary Figure 4 defines extreme changes based on the standard deviation of year-over-year percent changes for each quantity. An extreme change is defined as a year with a year-over-year percent change of greater 2σ or less than $1 - (1+2\sigma)^{-1}$. This definition is multiplicatively symmetric, meaning that if $\sigma=1$, an extreme change is defined as a change of greater than 200% or less than -50%.

Most non-price quantities have strongly positive trends between 1950 and 1985, leading to a large number of upward extreme changes. Due to these upward trends, ten quantities do not have any negative extreme changes: Oil production, natural gas consumption, coal price, electricity price, electricity sales, inflation, GDP, and total, commercial, and transportation energy consumption. 1985-2005 has five extreme changes, the smallest number of extreme changes in any period. These extreme changes occur only in oil and natural gas price. 2005-2014 has the largest concentration of negative extreme changes of any decade, with six negative extreme changes in oil and natural gas price, coal consumption, and residential energy consumption. 2005-2014 also has three positive extreme changes in oil production, due to an increase in hydraulic fracturing.

6.7. Supplementary Discussion

Our method of computing year-on-year volatility and unpredictability draws on methods from extreme value theory and quantitative finance. Extreme value theory studies the statistical characteristics of either the most extreme single values of a distribution, or of the most extreme values in a distribution beyond some threshold value, sometimes called peaks-over-threshold (PoT)[179]. The thresholds can be computed in a number of ways, and there is often a subjective element to threshold selection[179].

In our study of year-on-year volatility, we use the former definition due to small sample size considerations, and the latter in the study of unpredictability.

The equivalent to extreme value theory's peaks-over-threshold (PoT) analysis in the finance literature is called the value-at-risk (VaR) metric[180]. VaR is commonly used to measure a portfolio's exposure to losses[181]. A study from McKinsey & Company found that of 18 leading financial institutions, all considered VaR in their decision-making processes, and 75% used historical simulation (HS) to compute VaR[181]. HS takes the desired percentile of the distribution of historical values (returns on investment in most cases)[181], which is exactly analogous to our method of computing extreme error thresholds, except that we consider historical projection errors instead of historical returns.

In many contexts, volatility is synonymous with the standard deviation of a quantity over time[182–184]. However, in energy markets the concept of volatility is deeply intertwined with price shocks, a phenomenon studied closely in the economics literature, particularly in oil markets[185–193]. Several papers in the energy finance literature use VAR of oil and electricity prices, looking at the most extreme changes, albeit on shorter timescales, primarily to improve quantitative

treatment of risk in financial trading models[179,180,194,195]. Due to our small sample size, a result of analyzing annual changes in quantities, our analysis of year-on-year volatility simply considers only the most extreme upward and downward year-over-year changes instead of determining a percentile of the distribution of historical changes.

6.8. Supplementary Note 6

We evaluate the robustness of the results of our unpredictability analysis through a sensitivity analysis, using twenty-three sensitivity scenarios, each with alternative definitions of extreme error or subsets of the data, described at the end of this section.

This sensitivity analysis tests the robustness of two main results, the increase in unpredictability, the frequency of extreme errors, in 2005-2014 relative to prior decades for the majority of quantities, and the high number of quantities for which 2005-2014 contains all over-projected extreme errors. Because some sensitivity scenarios define extreme error differently than the baseline, many have more or fewer overall extreme errors than others. As a result, one cannot directly compare the change in the frequency of extreme errors across decades. One can, however, compare the number of scenarios that have an increase in the frequency of extreme errors relative to other decades. If an increase or decrease holds for all scenarios, we say the increase is *robust*. If the result holds for a majority of scenarios, this still suggests that it is not simply an artifact of a particular definition of extreme error.

Supplementary Figure 8 shows the percentage-point change in unpredictability, the frequency of extreme errors, in 2005-2014 compared with 1985-1994 (left), and 1995-2004 (right) for all quantities and all sensitivity scenarios. For example, for oil prices there is an increase in unpredictability relative to 1995-2004 in the baseline, and in eighteen other scenarios, with a

decrease in four scenarios. For oil production, however, there is an increase relative to 1995-2004 for all 23 scenarios, thus this increase is robust. Note that Supplementary Figure 8 includes sensitivity results for the three derived quantities, oil and natural gas imports and CO_2 emissions. Reported results below are for the only the seventeen quantities listed in the main manuscript.

Relative to 1995-2004, unpredictability increases in the baseline scenario for fifteen of seventeen quantities considered in the manuscript. For sixteen quantities, there is an increase in the majority of scenarios, with robust increases in three quantities: Oil production, and natural gas production and consumption. Relative to 1985-1994, unpredictability increases in the baseline scenario for ten of seventeen quantities. For twelve quantities, unpredictability increases in the majority of scenarios, with robust increases in seven quantities: Gross domestic product, coal and natural gas production and consumption, electricity sales, and residential energy consumption. Note that quantities with baseline decreases in unpredictability relative to 1985-1994 tend to be prices or inflation. Thus, there is a convincing increase in the frequency of extreme errors for more quantities relative to 1995-2004, but there is a robust increase for more quantities relative to 1985-1994.

Supplementary Figure 9 shows the number of quantities, of the seventeen in the main manuscript, for which all over-projected and under-projected extreme errors fall in 2005-2014 for all twenty-three sensitivity scenarios. The total number of quantities, out of seventeen, for which this is true varies from scenario to scenario, largely due to a higher number of extreme errors in some scenarios than others. However, in all scenarios the number of quantities with all over-projected or under-projected extreme errors in 2005-2014 is always greater than the number in the previous two decades. This strongly supports our finding that many quantities saw unprecedented errors in at least one direction during 2005-2014.

Description of scenarios: Our sensitivity analysis uses 23 scenarios, each of which uses either an alternative definition of extreme projection error, or an alternative subset of the data.

11 scenarios use the baseline definition of extreme error, in some cases slightly perturbed. "**Baseline**", "**Baseline, 1.5%**", and "**Baseline, 5%**" use the baseline definition, with extreme error thresholds at the 2.5th and 97.5th, 1.5th and 98.5th, and 5th and 95th percentiles of the historical percent error distributions respectively (with thresholds computed separately for short-term, medium-term, and long-term errors). We include these scenarios to determine whether a more or less restrictive definition of extreme error, within the baseline framework, produces a similar increase in the frequency of extreme errors.

We include a series of scenarios with the baseline definition, but changes to the data. The "**Absolute errors**" scenario performs the baseline method using absolute, rather than percent errors, i.e. the simple difference between the projected and historical values. "**No long term**" removes from the data all projections beyond 10 years out into the future. This addresses the concern that the measured increase in the frequency of extreme errors could be primarily attributable to long-term projections.

The Financial Crisis and subsequent Great Recession likely had a substantial and unpredictable effect on many aspects of the US energy system. To address the possibility that these events entirely account for the measured increase in the frequency of extreme errors, we include a "**No Great Recession**" scenario, in which we remove projection years 2007-2014 from the analysis.

To address the potential effects of the baseline definition of projection intervals, we include two scenarios with alternative definitions. "**Short 1-10 mid 11+**" defines short-term as 1-10 years, and medium-term as 11-21 years, with no separate long-term interval. "**Short 1-5 mid 6+**"

defines short-term as 1-5 years, and medium-term as 6-21 years, with no separate long-term interval. "**All errors together**" groups all projection errors together into a single interval.

The "**Only NEMS**" scenario uses only projections that came from the NEMS model, i.e. AEO 1994 onward. "**Only NEMS, no long-term**" also removes projections with length more than 10 years because almost all of these lie in 2005-2014 (91 out of 94 for each quantity except CO2 emissions, for which 63 of 66 fall in 2005-2014).

Seven scenarios aim to control for the possibility that one or a few particularly high-error years or projections caused the increase in extreme errors. These scenarios apply the baseline definition of extreme error but allow only one extreme error per base year or projection year, or both. We describe these scenarios below:

"One per base year, first": $FreqExtremeErrors_q = N_{e,i,t} / N_{p,t}$, where $N_{e,i,t}$ is the number of base years (AEOs), i, for which the first (earliest projection year) extreme error is within time-period t. $N_{p,t}$ is the total number of projection values in time-period t. We divide by the total number of projection values, rather than the total number of base years within time-period t, because within a given time-period t, some base years will have more projection values than others.

"One per base year, last": Analogous to "One per base year, first", but using the last (latest projection year) extreme error from each base year (AEO).

"One per projection year": $FreqExtremeErrors_q = N_{e,t} / N_{p,t}$, where $N_{e,t}$ is the number of projection years that contain extreme errors within time-period t, and $N_{p,t}$ is the total number of projection values in time-period t.

"One per proj. then base, first": $FreqExtremeErrors_q = N_{e,p,t,i,f} / N_{p,t}$, where $N_{e,p,t,i,f}$ is the number of base years, i, whose first extreme error is also the first (earliest AEO) extreme error of a projection year within time-period t.

"One per proj. then base, last": $FreqExtremeErrors_q = N_{e,p,t,i,l} / N_{p,t}$, where $N_{e,p,t,i,l}$ is the number of base years, *i*, whose last extreme error is also the last (latest AEO) extreme error of a projection year within time-period *t*.

"One per base then proj., first": $FreqExtremeErrors_q = N_{e,p,i,t,f} / N_{p,t.}$, where $N_{e,p,i,t,f}$ is the number of projection years within time-period *t* whose first extreme error is also the first extreme error of base year *i*.

"One per base then proj., last": $FreqExtremeErrors_q = N_{e,p,i,t,l} / N_{p,t}$, where $N_{e,p,i,t,l}$ is the number of projection years within time-period *t* whose last extreme error is also the last extreme error of a base year *i*.

Alternative definitions of "extreme error": Five scenarios employ different definitions of extreme error. "**Flat 15%**" designates all percent projection errors above 15% or below -15% as extreme errors. We use 15% because all quantities examined have errors of at least 15%, whereas some energy production and consumption quantities do not have errors above 20%.

The standard deviation scenarios define extreme errors based on the standard deviation of the historical errors themselves. Each of these scenarios first computes the standard deviation of the historical percent errors, separated into the aforementioned short-term, medium-term, and long-term intervals, under the implicit (and incorrect) assumption that these errors are not correlated over projection year or base year. In the "**1*SD, additive**" scenario, we set extreme error thresholds at ± the computed standard deviation for each projection interval. The "**2*SD, additive**" scenario uses twice the computed standard deviation. The "**1*SD, multiplicative**" scenario places thresholds at $-1 + (1 + \sigma)^{\pm 1}$, where σ is the computed standard deviation. This definition is multiplicatively symmetric, meaning that if σ=100%, the extreme error thresholds are +100% and -50%. "**2*SD, multiplicative**" uses $-1 + (1 + 2\sigma)^{\pm 1}$.

6.9. Supplementary Note 7

If quantities are normalized by population, derived from the decadal US Census and linearly interpolated in the intervening years, the year-on-year volatility results change for three quantities.

The downward extreme change for oil production moves from 1989 to 1958. The upward extreme change for coal production moves from 1955 to 1979. The downward extreme change for commercial energy consumption moves from 2012 to 1953.

As a result, there are eight instead of nine of thirty-four extreme changes in 2005-2014. There are also six instead of seven of thirty-four extreme changes in 1975-2004. Thus the qualitative result of a relative increase in extreme changes in 2005-2014 relative to the preceding two decades still holds.

The question of population growth matters for projections as well. If the Annual Energy Outlook's projections of population growth are off, this will certainly increase error in other directions. Population projections, which we do not show in this paper, are historically based on the US Census, and tend to have relatively small error compared with other quantities, with a maximum error of less than 5%.

Given that AEO projections of population growth are fairly accurate, we expect that errors in population growth should have only a small effect on errors in energy production and consumption quantities.

6.10. Supplementary Note 8

We reproduce versions of Figures 3 and 4 from the main text including oil and natural gas imports, as well as CO_2 emissions from the energy system. These three quantities can be derived from the seventeen primary quantities, so they were excluded from the main analysis. Still results for these quantities can be informative in their own right.

Supplementary Figure 6 plots extreme changes for all twenty quantities, including the three derived quantities. Overall, ten of forty, as opposed to nine of thirty-four extreme changes occur in 2005-2014. One of the two extreme changes for CO_2 emissions occurs in 2005-2014, with the remaining five extreme changes for the three derived quantities occurring before 1960. Supplementary Figure 7 plots the frequency of extreme errors, including the three derived quantities, oil and natural gas imports, and CO_2 emissions. Including these quantities, all over-projected extreme errors occur in 2005-2014 for twelve of twenty quantities, including natural gas imports and CO_2 emissions.

6.11. Supplementary Note 9

AEO projections errors are correlated across different quantities, and also have serial correlations within a given quantity across projection years within a given AEO base year, and across AEO base years within a given projection year. We include tables of the following correlations in Supplementary Data 2: Pearson and Spearman cross-quantity error correlations by AEO, Pearson and Spearman error serial correlations by AEO and projection year, and Pearson and Spearman cross-quantity correlations for projections and historical values of each quantity. We include summary statistics across all seventeen quantities for each set of correlations in Supplementary Table 3.

For AEO projection errors, cross-quantity Pearson and Spearman correlations have a median and mean close to zero, meaning that there are a comparable number of positive and negative correlations. Spearman correlations for AEO errors have a median of about 0.75 in both directions, meaning that errors tend to increase or decrease in the same direction across different quantities. Pearson and Spearman cross-quantity correlations for AEO projection values each have a median value of about 0.95, meaning that many of the projected quantities have a roughly

linear relationship with one another. Spearman and Pearson cross-quantity correlations of the historical values of quantities have a median value around 0.60, implying substantial correlation between quantities, but not as much as the AEO projects.

For each form of energy, price and consumption quantities are substantially anti-correlated across AEO reports, but historical values are not necessarily positively correlated. From 1985-2014, the correlation between historical oil price and consumption is 0.34, while the median value across all AEO reports is -0.65. Similarly, historical correlation between natural gas price and consumption is 0.21, while the median across AEO reports is -0.38. This discrepancy suggests that the AEO may not be capturing feedbacks beyond simple supply and demand. By contrast, historical correlation between coal and electricity price and consumption is -0.77 and -0.60 respectively, while in the AEO reports it is -0.71 and -0.38 respectively. This suggests that coal and electricity market dynamics are better captured by traditional supply and demand than oil and natural gas markets.

We analyze the relationship between annual metrics of unpredictability and year-on-year volatility with historical quantities, including the energy intensity of GDP, which we define as total energy consumption divided by GDP (see Supplementary Methods). We define annual unpredictability as the total number of extreme errors divided by the total number of projections in a given year over all seventeen quantities from 1985-2014. We define year-on-year volatility as the fraction of all extreme changes over all seventeen quantities that occur in a single year from 1950-2014.

We find that oil price and natural gas production are the most highly correlated with annual unpredictability, at 0.40 and 0.41 respectively. Annual unpredictability is most anti-correlated with the energy intensity of GDP, at -0.25. The oil price result is consistent with the pivotal

importance of oil in the US economy. As unpredictability increases most in 2005-2014, the increase in natural gas production due to hydraulic fracturing and slowdown in energy consumption growth following the Financial Crisis. Year-on-year volatility is anti-correlated with most quantities, and is most strongly anti-correlated with oil consumption, -0.17, and total energy consumption, -0.15, with positive correlations for three quantities: Oil price, 0.08, electricity price, 0.11, and the energy intensity of GDP, 0.07. Correlations here are smaller than for unpredictability, suggesting that there may not be a clear relationship between our metric of year-on-year volatility and any of these quantities.

Further work is needed to formally decompose the drivers of this year-on-year volatility, unpredictability, and the above-mentioned anomalies in historical correlations. Such analysis could perhaps follow the work of[196].

6.12. Supplementary Note 10

There are two forms of serial correlation of concern: (1) projections from different AEO reports (different base years) for the same year (projection year) may be correlated, and (2) projections from a single AEO base year may be correlated across different projection years.

We bound the effect of serial correlation as follows:

First, we compute both forms of serial correlation for each quantity. Second, we estimate the probability of our results occurring by chance in a Monte Carlo analysis that only accounts for cross-quantity correlations using the projection errors from each AEO report. Third, we compare these Monte Carlo results to an analogous Monte Carlo that treats all correlations between all distinct pairs of quantities as a single parametric constant, e.g. 50%. Fourth, we compare the estimated probability of our results occurring by chance in the second and third steps and find the level of parametric cross-quantity correlation from the third step that approximately matches the

probability of results occurring by chance from the second step. Fifth, we estimate the likely effect of adding the serial correlations from the first step to the parametric cross-quantity correlations from the fourth step.

First, we compute both forms of serial correlation across all quantities (see the Supplementary Data 2). In Supplementary Table 3, we show the summary statistics of the distribution of ranked (Spearman) serial correlations from projection errors computed in both directions, across base years and projection years. In both directions, the median serial correlation is 0.74 and 0.77 respectively, with means of 0.62 and 0.67 respectively (see Supplementary Table 3). We have also computed the Pearson correlations and also report these in Supplementary Data 2. We find that the Pearson correlations are similar to Spearman correlations, with median correlations of 0.77 and 0.84 for correlations across base years and projection years respectively (see Supplementary Table 3).

Second, we estimate the probability of our results occurring by chance in a Monte Carlo analysis that only accounts for cross-quantity correlations using the projection errors from each AEO report (see Materials and Methods, Figures 5 and 6). We find that the probability of an increase in the frequency of extreme errors for at least 15 of 17 quantities between 1995-2004 and 2005-2014 is less than 5%, while the probability of all over-projected extreme errors for 10 of 17 quantities occurring between 2005 and 2014 is less than 0.1%.

Third, we compute an analogous Monte Carlo that treats all correlations between all distinct pairs of quantities as a single parametric constant, e.g. 50% (see Materials and Methods, Table 1). We compute the probability of our key results occurring by chance under parametric correlation.

Fourth, we compare the estimated probability of our results occurring by chance in the second and third steps and find the level of parametric cross-quantity correlation from the third step that approximately matches the probability of results occurring by chance from the second step. Using parametric cross-quantity correlations, we find that the probability of an increase in the frequency of extreme errors for at least 15 of 17 quantities between 1995-2004 and 2005-2014 is 1.6% at a parametric correlation level of 50%, and 5.7% at a correlation level of 75%, roughly bounding the results observed using cross-quantity correlations from the AEO. The probability of all over-projected extreme errors for 10 of 17 quantities occurring between 2005 and 2014 is less than 0.1% at as high as 90% parametric cross-quantity correlation. Thus, by this metric the empirical cross-quantity correlations from the AEO are roughly equivalent to 75% or less for our first key result, and 90% for our second key result.

Fifth, we estimate the likely effect of adding the serial correlations from the first step to the parametric cross-quantity correlations from the fourth step. It is likely that adding two forms of serial correlations, with average magnitude between 60% and 85%, would increase the probability of the first result, an increase in the frequency of extreme errors for at least 15 of 17 quantities between 1995-2004 and 2005-2014, above 5%. We expect that adding the same serial correlations would not move the probability of the second result, all over-projected extreme errors for 10 of 17 quantities occurring between 2005 and 2014, above 5%. Even with 99% parametric cross-quantity correlation, this probability is only 0.4%.

Thus, we believe that adding both forms of serial correlation may increase the probability of one, but not both of our main results occurring by chance to above 5%.

6.13. Supplementary Note 11

Here we illustrate a potential application of our methods, and the importance of including the years 2005-2014 in any analysis based on historical errors from the Annual Energy Outlook (AEO).

Suppose a firm in the year 2015 is considering investing in a liquefied natural gas (LNG) terminal in the United States to export LNG to markets in Europe and Asia. The firm considers a number of scenarios, varying important exogenous parameters to evaluate the soundness of the investment. We focus on the firm's evaluation of a "high export price" scenario, in which the going rate of exported LNG is $15/mcf, in the high range of the preceding ten years[197].

The firm uses domestic natural gas price projections from the AEO for its baseline scenario. Suppose the CEO would like to use historical errors from the AEO to create a "High domestic price" scenario for natural gas but believes that the volatile years 2005-2014 were an anomaly, and asks company analysts to include a scenario that excludes these years from their retrospective analysis.

The firm then uses a modified version of the baseline projection combined with information about historical projection errors to create a "High domestic price" scenario, based on the 2.5th percentile, γ, of historical errors in the AEO for natural gas prices for short-term (1-5 year), medium-term (6-10 year), and long-term (11+ year) projections. In other words, the firm computes the extreme error thresholds using historical AEO errors for all prior years, 1985-2014, and for only 1985-2004, excluding 2005-2014. The "High domestic price" scenarios assume the AEO baseline projections will be correspondingly low, using the formula:

$$price' = price/(1 - \gamma)$$

For both scenarios, the resulting values of γ are:

1985-2014: 0.53, 0.63, 0.63

1985-2004: 0.50, 0.56, 0.36

Note that γ is not monotonic in the 1985-2004 scenario, decreasing from medium-term to long-term. The analysts at this company enforce weak monotonicity over time.

1985-2004 monotonic: 0.50, 0.56, 0.56

The resulting price multipliers are:

1985-2014: 2.13, 2.70, 2.70

1985-2004 monotonic: 2.00, 2.27, 2.27

Other parameters used in the analysis are:

Terminal capacity: 1M mcf/d

(A mid-sized terminal among those approved by the Federal Energy Regulatory Commission as of May 1, 2017[198].)

Capital cost: \$5bn/(mcf/d)

(Based on the cost of the Sabine Pass LNG terminal, with a capacity of 1.067 bcf/d, and a capital cost of \$5 billion[199].)

Natural gas export price: \$15/mcf[197]

Capacity factor: 90%

LNG export cost: $3.37/mcf

(This is a low range of costs, with $2.25/MMBTU fixed costs, including liquefaction, $1/MMBTU transportation in addition to domestic natural gas price[200], with a conversion rate of 1.037 Mcf/MMBTU[201].)

Discount rate: 0.07

Asset lifetime: 50 years

(AEO prices and high domestic natural gas price scenarios use flat interpolation of prices beyond 2040.)

We find that under these circumstances, the AEO reference case natural gas price scenario has a net present value (NPV) of $26 billion. Under a high domestic natural gas price scenario that ignores the years 2005-2014 and enforces monotonicity in prices, the terminal still turns a substantial profit of $2.9 billion in NPV. However, even with this highly favorable export price, and low export costs, an LNG terminal under the "1985-2014" scenario loses $220 million.

In other words, a high domestic natural gas price scenario that excludes the years 2005-2014 produces a different, and perhaps misleading picture of the risks facing an LNG export terminal in the United States. *Even with high export prices, constructing an LNG export terminal is not a no-regrets scenario.*

We recognize that this is far from a comprehensive analysis of the decision to build an LNG export terminal. Many of the above assumptions would require sensitivity analysis, particularly natural gas export price. This analysis does, however, illustrate how excluding or including the years 2005-2014 from a retrospective analysis of projection errors from the Annual Energy Outlook could influence one major energy infrastructure investment decision.

6.14. Supplementary Methods

6.14.1. Data collection and processing.

Extreme error analysis: All projection data and observed historical values used in the extreme error analysis come from either the Annual Energy Outlook (AEO) retrospective reports, or from the individual AEO reports themselves. The single exception to this is Gross Domestic Product (GDP), which is derived from a combination of AEO projections of GDP growth and US Bureau of Economic Analysis (BEA) values of historical US GDP, described in further detail below.

All values for AEO 1993 and earlier come from the 2010 AEO retrospective report[172], except values for "projection year" 2010 (i.e. projections for the year 2010), which come from the individual AEO reports[30]. All values from AEO 1994 onward come from the 2014 AEO retrospective report[35], with exceptions listed below.

Projected values and observed historical values for projection year 2014 are not available in the 2014 AEO retrospective report. We draw these values from individual AEO reports.

We derive any missing projection values through linear interpolation between values from projection years before and after the missing value.

We use the term "base year" to describe the year in the AEO report name for AEO 1982-1987, and one year prior to that for AEO 1989-2015 as a result of a change in the AEO naming convention relative to release years. There is no AEO 1988 because EIA changed the naming convention in that year.

We analyze all projected price quantities in constant 2005 dollars to disentangle the effects of inflation projections embedded within nominal dollar projections.

6.14.2. Effects of interpolation on our results

Many of the earlier AEO reports produced by the EIA, listed below, only reported longer-term projections in five-year increments. The one exception is AEO 1990, which reported only five-year increments from 1990 to 2005. We include a comprehensive list of AEO reports with missing values below.

Missing values occur in the following AEO reports:

AEO 1998 skips from 2010 to 2015 to 2020. AEO 1996 skips from 2010 to 2015. The AEO 2014 retrospective reports values for AEO 1996 and 1998 up to 2012, which we use in our analysis. AEO 1990 only reported every 5 years after 1990, through to 2005. AEO 1987 skips from 1995 to 2000. AEO 1983 and 1984 skip from 1990 to 1995

In almost all cases, this interpolation simply follows the previous linear trend. For this reason, we do not believe linear interpolation introduces significant bias into our results. Two notable exceptions occur in AEO 1990, in which projections for final ten years reverse a previous upward trend for natural gas production and consumption. For natural gas consumption, this projection closely approximates the actual historical trend, thus interpolation does not introduce a clear bias. For natural gas production, there is a slight over-projection, thus interpolation of this concave curve introduces a slight bias toward smaller projection errors.

6.14.3. Adjustments to individual quantities

Coal production: Reported coal production values from AEO 2006 and beforehand include waste coal. From AEO 2007 on, we manually add waste coal to regular coal production to match AEO retrospective reports, which apparently include waste coal in all years but AEO 2013.

Because of this apparent discrepancy in the retrospective reports, for this quantity we use these computed values, drawn directly from the AEO reports.

CO_2 emissions: Carbon dioxide emissions were reported in units of carbon mass through AEO 2003. We convert these to CO_2 mass using the molecular mass ratio of carbon to carbon dioxide, 12/44 kg C/kg CO_2.

Oil volume quantities: Some oil volume quantities were reported in Mbbl/yr (million barrels per year) instead of Mbbl/day (million barrels per day). These are converted to Mbbl/day by dividing by the number of days in the year in question (accounting for leap years). To match the AEO retrospective reports, we define oil imports as follows.

AEO 1990-1993: Net imports (including Strategic Petroleum Reserve)

AEO 1996-2007: Crude oil imports + All refined products including ethanol

AEO 2008-2015: Crude oil imports + refined products + ethanol + biodiesel

Natural gas price: The AEO retrospective reports for natural gas price use an average price per thousand cubic feet (Mcf) in the lower 48 US states from AEO 1982-2012, switching to Henry Hub prices per million Btu ($/MMBtu) in AEO 2013-2014 (without reporting a lower 48 average). The AEO 2013 retrospective report apparently fails to perform a corresponding unit conversion [202]. The AEO 2014 retrospective report instead reports natural gas prices to electric generating plants [35]. For this reason, we use natural gas price data from the AEO 2013 retrospective report when available, in lieu of the AEO 2014 retrospective report. For both nominal and constant dollar values for AEO 2013-2015, we convert from $/MMBtu to $/Mcf using the EIA's reported average annual heat content of dry natural gas produced in the US, assuming each AEO report used the constant conversion factor from its release year. These

values are 1.024, 1.027, and 1.032 MMBtu/Mcf for AEO 2013, 2014, and 2015 respectively[203–205].

Coal price: While the AEO retrospective reports present coal prices on an energy basis, $/MMBtu, the AEO 1983-1993 reports present coal price on a mass basis, $/short ton. We convert these prices to $/MMBtu using values of coal energy content from the AEO reports. AEO 1983 and 1984 each include projected values for 1985 and 1995. AEO 1990-1993 each provide a single projected value. All of these values are between 20.8 and 21.2 MMBtu/short ton. We determine projected coal energy content for AEO 1983-1993 by linearly interpolating these values over all projection years[206–216], assigning years not between two projected values the constant value of the nearest projected value. We then estimate this conversion factor for AEO reports that do not include a projected energy content value through the same linear interpolation process across all AEO reports (across the AEO base year).

Inflation: AEO retrospective reports do not include projections of inflation. We use values of cumulative inflation, measured through the GDP or GNP Price Index, gathered from each individual AEO report. We compare this inflation rate with AEO's implied inflation projections in years for which there are price projections of both nominal and constant dollars. We confirm that of all reported inflation metrics, GDP Price Index is the closest match. For AEO 1998-2014 we use a quantity titled "GDP Chain-Type Price Index", for AEO 1987-1997 we use a quantity titled "GDP implicit price deflator", for AEO 1983-1986 we use a quantity titled "NIPA GNP Price Deflator", and for AEO 1982 we use a quantity titled "GNP Deflator". Note that GDP price index values were not available prior to AEO 1987.

Because the inflation projections come from different quantities over time (e.g. GDP and GNP Price Index), there is no single time series of historical inflation values that would match up in all

cases. We derive "observed historical values" from the most recent AEO that includes a projected value for the projection year in question. Because all AEO reports include projections for historical years, these values are all from reports released at least two years after the inflation value, except for projection year 2014, which comes from AEO 2015, with only a one-year lag.

GDP: The 2010 and 2014 AEO retrospective reports present GDP growth values in 2005 and 2009 chain-weighted dollars respectively. Individual AEO reports use different chain-weighted dollar-years. Chain-weighting is an extremely complex process, and conversion between chain-weighted dollar years is a highly nontrivial task[217]. To produce an intercomparable set of projections of US GDP, we combine percent GDP growth projections from the AEO with observed historical values of US GDP in 2009 chain-weighted dollars from the US Bureau of Economic Analysis (BEA) [32]. We get percent real GDP growth in chain-weighted 2009 dollars from the 2010 and 2014 AEO retrospective reports. The 2014 retrospective report goes through projection year 2013. For missing values and values for projection year 2014, we compute real GDP growth from projections of chain-weighted real US GDP, which we draw from AEO 1982-2015[203–216,218–236]. Note that this is an approximate method, which does not fully account for the complexities of the chain-weighting process.

Quantities for year-on-year volatility analysis: For our analysis of extreme year-over-year changes (year-on-year volatility) in observed historical values of energy quantities since 1949, except electricity price, which begins in 1960, we draw data from the EIA's Monthly Energy Review, November 2015 when available[31]. The exceptions are US GDP and inflation, which we draw from the US BEA, and CO_2 emissions, which we draw from Oak Ridge National Laboratory (ORNL) [32,33,237]. ORNL only reports CO_2 emissions through 2011, for subsequent

years, we use year-over-year changes derived from AEO retrospective CO_2 emissions data[35]. All prices here are in nominal dollars.

Low and high oil price scenarios are derived from the following sources:

AEO 1999: https://www.eia.gov/outlooks/archive/aeo99/results.html
AEO 2000: https://www.eia.gov/outlooks/archive/aeo00/results.html
AEO 2001: https://www.eia.gov/outlooks/archive/aeo01/results.html
AEO 2002: https://www.eia.gov/outlooks/archive/aeo02/results.html
AEO 2003: https://www.eia.gov/outlooks/archive/aeo03/results.html
AEO 2004: https://www.eia.gov/outlooks/archive/aeo04/results.html
AEO 2005: https://www.eia.gov/outlooks/archive/aeo05/results.html
AEO 2006: https://www.eia.gov/outlooks/archive/aeo06/aeolowprice.html
https://www.eia.gov/outlooks/archive/aeo06/aeohighprice.html
AEO 2007: https://www.eia.gov/outlooks/archive/aeo07/aeolowprice.html
https://www.eia.gov/outlooks/archive/aeo07/aeohighprice.html
AEO 2008: https://www.eia.gov/outlooks/archive/aeo08/aeolowprice.html
https://www.eia.gov/outlooks/archive/aeo08/aeohighprice.html
AEO 2009: https://www.eia.gov/outlooks/archive/aeo09/aeolowprice.html
https://www.eia.gov/outlooks/archive/aeo09/aeohighprice.html
AEO 2010: https://www.eia.gov/outlooks/archive/aeo10/aeolowprice.html
https://www.eia.gov/outlooks/archive/aeo10/aeohighprice.html
AEO 2011: https://www.eia.gov/outlooks/archive/aeo11/topic_prices.cfm
AEO 2012: https://www.eia.gov/outlooks/archive/aeo12/topic_prices.cfm
AEO 2013: https://www.eia.gov/outlooks/archive/aeo13/topic_prices.cfm
AEO 2014: https://www.eia.gov/outlooks/aeo/data/browser/#/?id=12-AEO2014&cases=ref2014~highprice~lowprice&sourcekey=0

Further notes: For AEO 2009 values, we use the updated AEO report produced with funds from the federal stimulus package at that time[233]. These values match those reported in AEO retrospective reports.

6.15. Supplementary Tables

Table 17. Mean values of short-, medium-, and long-term distributions of AEO errors for each of twenty quantities, including derived quantities. Note that fifty-two of sixty cases have positive mean drift.

	Short-term (1-5 yrs)	Medium-term (6-10 yrs)	Long-term (11+ yrs)
Oil price	4.3%	16.8%	-15.8%
Oil consumption	1.8%	4.6%	6.9%
Oil production	0.6%	-1.6%	-7.0%
Oil imports	9.1%	21.3%	39.6%
Natural gas price	14.6%	41.0%	19.1%
Natural gas consumption	-1.4%	0.3%	4.3%
Natural gas production	-1.7%	-0.6%	2.4%
Natural gas imports	15.0%	42.0%	67.5%
Coal price	8.0%	21.9%	21.0%
Coal production	2.8%	6.5%	9.4%
Coal consumption	2.7%	7.1%	13.9%
Electricity price	1.3%	3.3%	3.7%
Electricity sales	0.1%	2.1%	2.7%
GDP	2.1%	3.2%	-0.8%
Inflation	2.0%	8.9%	22.4%
Total energy cons.	2.0%	4.5%	7.4%
Residential energy cons.	2.4%	3.3%	2.5%
Commercial energy cons.	0.5%	0.7%	-2.2%
Transportation energy cons.	0.9%	2.4%	4.7%
CO2 emissions	1.9%	8.7%	13.3%

Table 18. Median values of short-, medium-, and long-term distributions of AEO errors for each of twenty quantities, including derived quantities. Note that forty-two of sixty cases have positive median drift.

	Short-term (1-5 yrs)	Medium-term (6-10 yrs)	Long-term (11+ yrs)
Oil price	-3.1%	-0.2%	-45.7%
Oil consumption	0.6%	0.2%	5.6%
Oil production	1.7%	-1.8%	-1.9%
Oil imports	3.5%	4%	13.8%
Natural gas price	11.6%	18.6%	-6.6%
Natural gas consumption	-2.5%	-1.3%	4.9%
Natural gas production	-1.3%	-1.3%	3.1%
Natural gas imports	5.7%	-3.1%	12.2%

Coal price	6.9%	31.5%	22.2%
Coal production	1.3%	2.3%	5.5%
Coal consumption	2.2%	4.7%	12.3%
Electricity price	1.8%	11.0%	7.7%
Electricity sales	-0.1%	1.0%	2.3%
GDP	0.9%	0.1%	-4.5%
Inflation	1.2%	8.7%	21.2%
Total energy cons.	1.5%	2.0%	6.0%
Residential energy cons.	2.9%	2.6%	1.8%
Commercial energy cons.	0.0%	-1.3%	-3.7%
Transportation energy cons.	1.2%	-0.4%	-0.4%
CO2 emissions	1.4%	6.1%	15.8%

Table 19. Summary statistics for cross-quantity and serial correlations for AEO projection errors, AEO projection values, and historical values of the seventeen selected energy quantities. See Supplementary Data 2 for all correlation values.

	Median	Mean	Standard deviation
Pearson error corr. cross-quantity	0.12	0.07	0.67
Spearman error corr. cross-quantity	0.10	0.06	0.66
Pearson error serial corr. by AEO	0.84	0.71	0.36
Pearson error serial corr. by proj. yr	0.77	0.67	0.36
Spearman error serial corr. by AEO	0.77	0.67	0.35
Spearman error serial corr. by proj. yr	0.74	0.62	0.37
Pearson proj. cross-quantity corr.	0.94	0.52	0.70
Spearman proj. cross-quantity corr.	0.97	0.54	0.70
Pearson cross-quantity historical corr.	0.62	0.34	0.62
Spearman cross-quantity historical corr.	0.59	0.35	0.59

6.16. Supplementary Note 12

Supplementary Figures 14-16 show cumulative distribution functions of projection errors for each quantity by projection interval. Extreme error thresholds, the 2.5^{th} and 97.5^{th} percentiles of each distribution, are also displayed. In Supplementary Tables 1 and 2, we also compute mean and median drift for each quantity.

Extreme error thresholds vary dramatically in magnitude from quantity to quantity, with a maximum over-projection threshold value of 362% for medium-term projections for natural gas imports, and a minimum magnitude under-projection threshold value of -2% for long-term inflation. Prices tend to have larger-magnitude thresholds than energy production and consumption quantities or macroeconomic indicators such as GDP.

Considering twenty quantities, including the three derived quantities, short-term thresholds tend to be smaller in magnitude than medium-term or long-term thresholds, for 19 of 20 over-projection thresholds and 15 of 20 under-projection thresholds. However, the number of quantities for which the magnitude of the long-term threshold is greater than the medium-term threshold falls to 17 of 20 for high thresholds and 16 of 20 for low thresholds. Thus, errors tend to increase with projection length.

All quantities have at least some mean and median drift (see Supplementary Tables 1 and 2). Median mean drift across all twenty quantities is 2.0% for short-term, 4.5% for medium-term, and 5.8% for long-term projections. Median drift tends to be slightly smaller, with a median median drift across all twenty quantities of 1.3% for short-term, 1.5% for medium-term, and 5.2% for long-term projections. Mean and median drift tend to be positive, or over-projections, with fifty-two of sixty means positive, and forty-two of sixty medians positive.

7. Supporting information for Chapter 3

7.1. Demographics of the CARE program

The majority of eligible California households (51%) reside in single-unit structures in 2011, with 32% in multi-family housing of five or more units, 11% in two-to-four-unit structures, and 6% in mobile homes [39]. 64% rent their homes, compared to 55% of all households. 93% of low-income households California reside in more urbanized counties, although 15% of low-income PG&E customers reside in largely rural counties [39]. Most low-income households, 54% spoke a language other than English at home, with 38% primarily speaking Spanish and 9% using an Asian language [39]. 20% are "linguistically isolated, with no household member aged 14 or older who can speak English fluently [39]. This is double the general population rate of 10% [39]. 42% reported Hispanic ethnicity, followed by 36% white, 10% Asian, and 9% African American [39].

7.2. Additional data description

7.2.1. Regions within PG&E

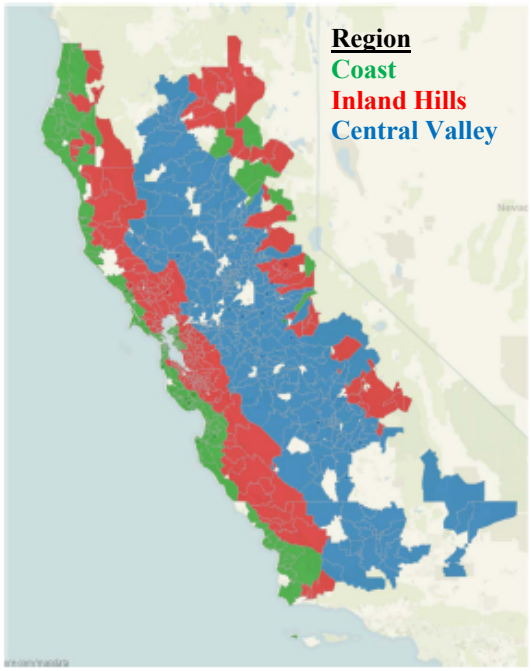

Figure 33. Regions in the Pacific Gas and Electric Company (PG&E) service territory. PG&E randomly selected approximately 10,000 dwellings from each of the region to construct the sample. Note that these regions were grouped by climate and thus are not fully geographically connected, particularly the Coast and Inland Hills. However, these disconnected eastern areas classified as Coast and Inland Hills are in sparsely-populated mountainous locations and thus likely have little impact on our results. Figure from the Wharton Customer Analytics Initiative.

7.2.2. Deployment of advanced metering infrastructure

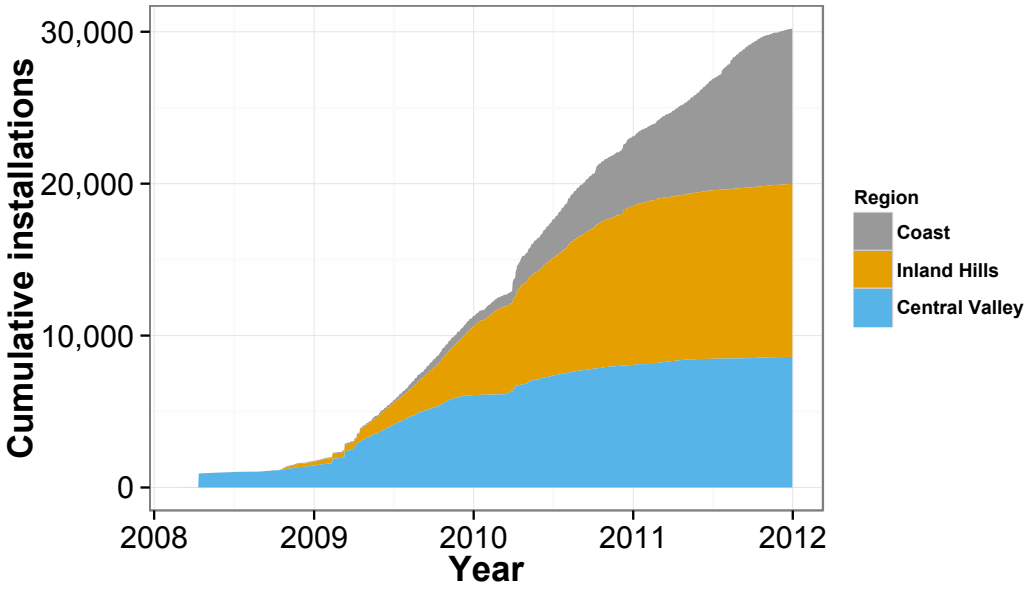

Figure 34. Advanced metering infrastructure deployment within the sample, March 1, 2008 to December 31, 2011 by region. Deployment began in the Central Valley, followed by the Inland Hills, followed by the Coast. The result is an unbalanced panel, with more observations in later time periods. Also shown in Sherwin et al. [238].

7.2.3. CARE enrollment by region and neighborhood income

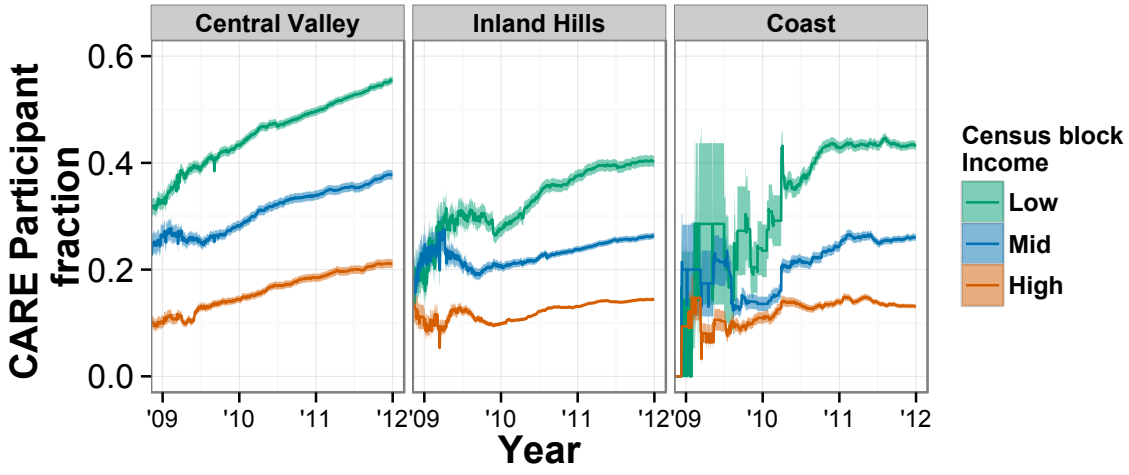

Figure 35. Enrollment rate in CARE as a fraction of households in the dataset over time, by region and median census block median income, with thresholds of $52,252.33 and $81,572.00, the 1/3 and 2/3 quantiles of households in our sample respectively. Shaded areas are 95% probability interval, considering sample error. A large fraction of CARE participants lives outside low-income census blocks. The increasing trend in CARE participation is likely in part due to AMI deployment decisions, not changes in population enrollment. Also shown in Sherwin et al. [238].

7.2.4. Other utility programs

Table 20. Description of the PG&E programs and the total and maximum number of participants observed in the dataset. Includes total and maximum participation. Adapted from Meyer [62].

Program	Total Part.	Max Part.	Description
Rebates	2,804	2,773	Energy efficiency rebates subsidize the purchase of efficient appliances, services, and household equipment through an after-purchase mail-in rebate.
BPP	2424	1,574	Balanced Payment Plan: Provides a bill smoothing service, in which PG&E calculates the household's average monthly utility bill and the customer pays a flat amount for each monthly billing cycle. This value is an average annualized value
CARE	10,193	9,337	California Alternate Rates for Energy: Subsidizes monthly energy bills based on income and occupant criteria such as enrollment in other means-tested programs.
Climate Smart	147	147	Households pay for carbon offsets through monthly utility bills
Direct Access (DirAccess)	668	631	Allows customers to purchase their electricity from alternative (non-PG&E) power providers, using PG&E as the distribution company (New customers have not been able to join the Direct Access program since the California energy crisis in 2001, though existing customers have been able to remain in the program)
Smart AC	1,069	856	Allows customers to opt in to a central air conditioning curtailment program during peak-load events during the summer cooling season.
Smart Rate	154	117	Lower average electricity tariff (3¢/kWh reduction) in exchange for accepting a higher rate (60¢/kWh) during peak hours in some days during the summer cooling months. These days are communicated to the consumer a day ahead via text, email, or phone

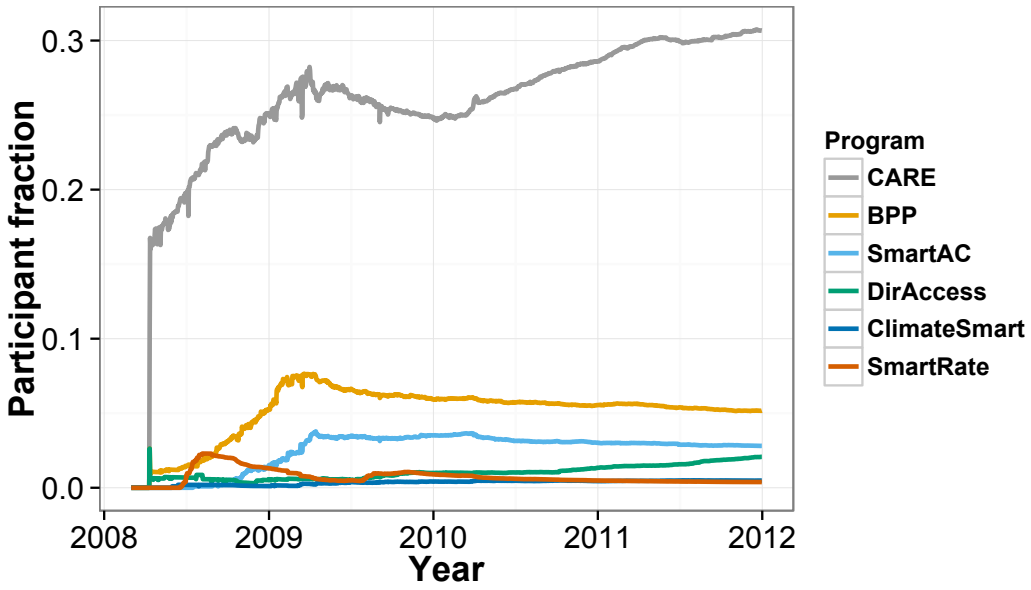

Figure 36. Enrollment rate in PG&E programs as a fraction of households in the dataset over time. The CARE low-income subsidy is by far the most prevalent. We exclude DirAccess, ClimateSmart, and SmartRate from detailed analysis due to low participation rates, and the fact that PG&E no longer allows new enrollment in DirAccess. Also shown in Sherwin et al. [238].

7.2.5. Variation in electricity consumption across seasons and regions

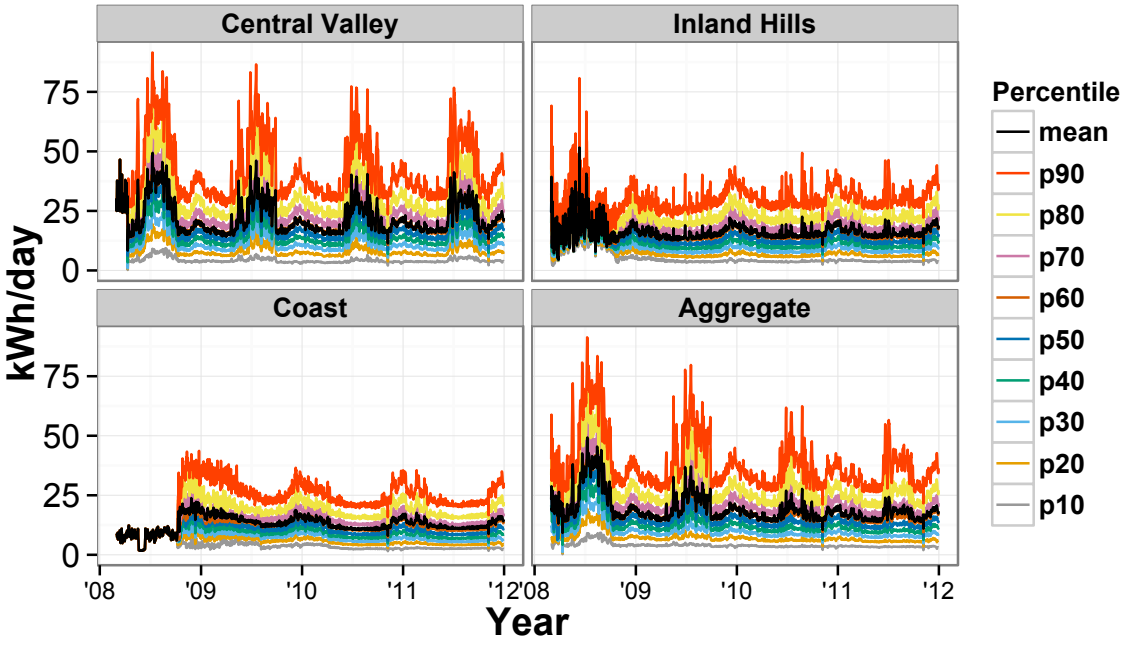

Figure 37. Deciles of daily household electricity consumption shaded by region and day, from the 10th percentile to the 90th percentile. Also shown in Sherwin et al. [238].

7.3. Approximate lognormality of residential electricity consumption

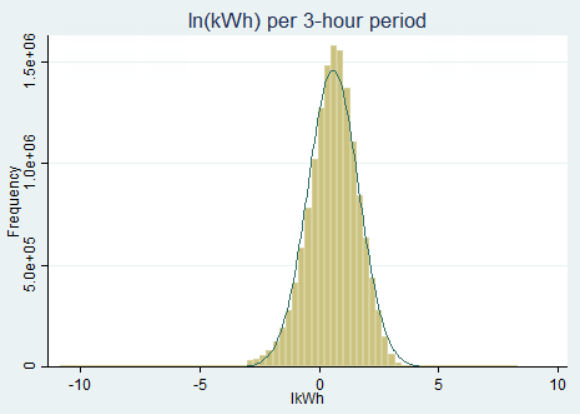

Figure 38. A histogram of the logarithm of electricity consumption in the full panel with 3-hour resolution. Note that the distribution of electricity consumption is approximately lognormal, further motivating the use of ln(kWh) as the dependent variable in the main analysis.

7.4. Limitations to causal identification

7.4.1.1. Confounding with income and employment

Households that newly enroll in CARE may do so because they are newly eligible for the program. This likely equates to a decline in household income, potentially due to a change in employment status, or a change in household size that would likely coincide with CARE enrollment. A reduction in income would likely trigger curtailment of electricity consumption through an income elasticity of demand. Depending on the size of the income shock, this effect could be comparable in magnitude to any change in electricity consumption due to a subsidized electricity rate. Increases in household size would also likely increase electricity demand. Intra-day variation in the timing of these effects is unknown, potentially confounding estimates of the intra-day effects of the CARE program.

Changes in employment status, such as a transition from full-time to part-time or shift work or to unemployment may introduce changes in household occupancy patterns that could substantially change both the magnitude and timing of electricity consumption. Although our data do not explicitly include unemployment, except at the census block level, one could conceivably extract unemployment information from interval electricity consumption data alone using machine learning or other statistical methods. However, such an approach would likely require ground truth measurements of employment status, which we do not have.

7.4.1.2. Lack of transparent electricity rate information

During the study period, PG&E's residential electricity rates were based on a multi-tier inclined block structure, in which prices increase after a customer surpasses one of several threshold values [74]. These thresholds differ across regions depending on local weather and other factors. In addition, the CARE subsidy is not a fixed percentage of rates, but varies across regions, with

discounts that depend on a customer's consumption tier [39]. As a result, the discount received by a household enrolled in CARE varies substantially across regions. In 2012, the average CARE discount in for PG&E customers was 42%, while for it was 24% for Southern California Edison customers [39].

The data do not include enough information about customer rates to reconstruct the magnitude of the CARE discount for individual customers or for the sample as a whole. We use the PG&E population average discount of 42% to estimate the price elasticity of electricity demand [39]. This lack of rate information is one reason we believe that the primary contribution of this analysis is estimation of the likely magnitude of intra-day variation in the price elasticity of demand, rather than estimation of the magnitude of the price elasticity of demand itself.

7.4.1.3. Instrumental variables

Observational studies often employ plausibly exogenous instrumental variables to mitigate the effect of selection bias on their causal estimates. For a means-tested program such as CARE, one could conceive of instrumental variables based on unemployment, perhaps using exogenous localized employment shocks. The Financial Crisis and subsequent Great Recession of 2007-2008 and the accompanying collapse of much of the California housing market likely had differential effects on unemployment across the PG&E service territory, potentially affecting urban, rural, and suburban household differently. However, even with time-varying household-level employment data, AMI deployment was in its early stages in 2008, with less than 10% of our sample, almost all in the Central Valley, covered by December 31, 2009.

7.4.1.4. Propensity score matching

In the absence of a clear natural experiment or instrumental variable, many econometricians rely on propensity score matching or synthetic controls to improve observational causal estimates.

Such an approach compares treated households with non-treated households that have similar observable characteristics. Roughly 90% of eligible households were enrolled in CARE during the study period [39]. As a result, even with individual-level demographic information, which we do not have, only a small fraction of non-enrolled households have incomes comparable to CARE households. In addition, it is likely that eligible non-enrolled households are different in important ways from CARE households [60], and thus may not be suitable matches.

As a rough approximation to matching, we compute separate estimates for subsamples based on median census block income represented in the sample, i.e. we compare households to peers in neighborhoods with similar median income. This approximates matching in that households in poorer neighborhoods tend to have higher CARE enrollment and may be more similar to each other in important unobserved ways than households in wealthier neighborhoods. See the SI, Sections 7.5.3.5 and 7.5.3.6 for these results.

7.4.1.5. Regression discontinuity

In many ways, means-tested programs are well-suited for a regression discontinuity approach, which can leverage the income threshold eligibility criterion to compare consumption in households with incomes above and below the threshold. We have only census block-level income information, which is likely too coarse and imprecise to add meaningful information beyond the difference-in-differences model. Even with such data, which electric utilities do not generally possess, there is large potential for confounding with other means-tested programs, many of which have similar eligibility criteria. In addition, there is a potential for households near the eligibility threshold to underreport income to gain access to such programs, which would further confound causal estimates.

7.5. Coefficient values from main figures and tables and robustness checks

The regression tables below correspond to the results in **Figure 7** and **Table 1**, applying Eq. 4 for intra-day estimates and Eq. 5 for time-invariant estimates.

Coefficients for cases in which the dependent variable is ln(kWh) are reported below. These are converted to percentages in the main text using the formula $e^x - 1$.

Regressions use a high-dimensional fixed effects (HDFE) regression specification.

7.5.1. Time-invariant regressions

	Full sample	Coast	Inland Hills	Central Valley	Summer	Winter
CARE	0.119*** (0.0112)	0.093** (0.0285)	0.139*** (0.019)	0.112*** (0.0153)	0.135*** (0.0154)	0.116*** (0.0117)
Temp ≥ 65°F	0.021*** (0.0003)	-0.003*** (0.0005)	0 (0.0004)	0.012*** (0.0003)	0.044*** (0.0004)	-0.003*** (0.0002)
Temp < 65°F	0.025*** (0.0003)	-0.002*** (0.0005)	0.003*** (0.0004)	0.017*** (0.0003)	0.040*** (0.0003)	0 (0.0002)
Constant	1.303*** (0.0053)	1.382*** (0.0104)	1.590*** (0.0066)	1.601*** (0.0077)	0.872*** (0.0084)	1.643*** (0.0042)
R-squared	0.51	0.56	0.51	0.47	0.58	0.51
N clusters	30,080	10,118	11,288	8,465	29,082	30,030
N observations	143,558,079	28,636,816	57,521,638	56,464,270	50,893,578	92,664,497

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7.5.2. Intraday regressions

	Full sample	Coast	Inland Hills	Central Valley	Summer	Winter
CARE baseline	0.110***	0.068*	0.127***	0.108***	0.118***	0.112***
	(0.0115)	(0.0299)	(0.0211)	(0.0154)	(0.0154)	(0.0124)
CARE3am	-0.005	-0.002	-0.013	-0.004	-0.005	-0.005
	(0.0036)	(0.0075)	(0.0076)	(0.0046)	(0.0053)	(0.0038)
CARE6am	0.003	0.025*	0.014	-0.007	0.004	0.004
	(0.0055)	(0.0128)	(0.0114)	(0.0071)	(0.007)	(0.0061)
CARE9am	0.005	0.037*	0.002	0	0.01	0.001
	(0.0075)	(0.0161)	(0.0148)	(0.01)	(0.0097)	(0.0082)
CARE12pm	0.006	0.018	-0.003	0.009	0.021*	-0.005
	(0.0077)	(0.0169)	(0.0143)	(0.0104)	(0.0105)	(0.0081)
CARE3pm	0.014	0.033*	0.015	0.011	0.032**	0.002
	(0.0078)	(0.0164)	(0.0139)	(0.0107)	(0.0108)	(0.0082)
CARE6pm	0.024***	0.060***	0.040**	0.01	0.035***	0.015*
	(0.0072)	(0.0166)	(0.0135)	(0.0096)	(0.0095)	(0.0078)
CARE9pm	0.023***	0.031**	0.041***	0.014*	0.033***	0.017**
	(0.0049)	(0.0114)	(0.01)	(0.0063)	(0.0066)	(0.0053)
Temp ≥ 65°F	0.020***	-0.005***	0	0.013***	0.039***	-0.004***
	(0.0003)	(0.0005)	(0.0004)	(0.0003)	(0.0003)	(0.0002)
Temp < 65°F	0.025***	-0.004***	0.002***	0.018***	0.037***	-0.0005*
	(0.0003)	(0.0005)	(0.0003)	(0.0003)	(0.0002)	(0.0002)
Constant	1.311***	1.407***	1.600***	1.592***	0.948***	1.649***
	(0.0052)	(0.0106)	(0.0064)	(0.0076)	(0.0068)	(0.0041)
R-squared	0.558	0.616	0.564	0.51	0.634	0.562
N clusters	30,072	10,115	11,287	8,461	29,051	30,023
N observations	143,557,960	28,636,757	57,521,610	56,464,239	50,893,226	92,664,394

Standard errors in parentheses.
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7.5.3. Robustness checks

7.5.3.1. Only weekdays
Includes only electricity consumption during weekdays.

7.5.3.2. Only weekends
Includes only electricity consumption during weekends.

7.5.3.3. Only 2010 and 2011
This specification uses Eq. 4 but excludes all years before 2010, which may capture short-run dynamics caused by the Financial Crisis of 2008 and the subsequent recession.

7.5.3.4. No consumption extremes
This specification uses Eq. 4 but excludes dwellings with average annual electricity consumption above the 5^{th} percentile or below the 95^{th} percentile of the sample.

7.5.3.5. Lower-income neighborhoods
Dwellings in the bottom 40% of census block income of our sample.

7.5.3.6. Higher-income neighborhoods
Dwellings in the top 60% of census block income of our sample.

7.5.3.7. Event study
This specification uses Eq. 4 with a restricted sample considers only dwellings that enroll in CARE a single time during the study period with at least 90 days of pre-and post-enrollment data. This reduces the sample to 1,524 dwellings. Households with multiple enrollments are an interesting case as well, but likely leave the program due to changes in income, household size, or other factors that would substantially influence electricity consumption.

7.5.3.8. No perpetual enrollees
This specification uses Eq. 4 with a restricted sample includes dwellings that never enroll in CARE and those that enroll during the study period but excludes the roughly 8,000 dwellings that are enrolled in CARE for the entire study period or enroll multiple times throughout their time with PG&E.

7.5.3.9. Always/never enrolled

This specification uses Eq. 4 with a restricted sample includes only dwellings that never enroll in CARE and those that enroll only once in CARE and are enrolled in CARE for the entire study period. This excludes the roughly 5,500 households that enroll in CARE during the study period or enroll multiple times throughout their time with PG&E.

7.5.3.10. Un-transformed kWh

This specification uses a modified version of Eq. 4, replacing ln(kWh) with kWh as the dependent variable. These coefficients are interpretable as a linear increase in kWh of electricity consumed per 3-hour period. Divide by three to convert to kW.

7.5.3.11. 10% data subsample

This specification uses Eq. 4 with a random 10% subset of electricity consumption readings from the full sample. For computational tractability, the same random subset is used for the robustness check that includes hourly coefficients for other utility programs. The qualitative results are similar to those from the full sample.

7.5.3.12. Interactions with other programs

This specification uses a modified version of Eq. 4, including intra-day coefficient estimates for all available utility programs, described in the SI, Section 7.2.4.

7.5.3.13. Monthly analysis

This specification uses an aggregated monthly version of the dataset and Eq. 11, below. This mimics traditional billing analysis of residential energy consumption.

$$\ln(y_{im}) = \beta CARE_{im} + \alpha_c CDD_{T>65°F} + \alpha_w HDD_{T\leq 65°F} + \gamma_i + \omega_m + \epsilon_{im} \quad (11)$$

CARE enrollment is determined based on whether a household was enrolled for at least half of a month. Hourly temperature is converted into heating and cooling degree-days using a 65°F set point. *m* is the month-of-sample. We use household-level and month-of-sample fixed effects.

7.5.3.14. Robustness checks using intra-day specification

	Full sample	Weekdays	Weekends	2010 and 2011 only	No consumption extremes	Lower-income neighborhood	Higher-income neighborhood	Event study	No perpetual enrollees	Always/ never enrolled
CARE baseline	0.110*** (0.0115)	0.110*** (0.0116)	0.111*** (0.0116)	0.116*** (0.0131)	0.105*** (0.0118)	0.096*** (0.0157)	0.122*** (0.0169)	0.106*** (0.0169)	0.167*** (0.0138)	0.074* (0.0375)
CARE3am	-0.005 (0.0036)	-0.004 (0.0039)	-0.008* (0.0036)	-0.008 (0.0042)	-0.005 (0.0037)	-0.001 (0.0051)	-0.011* (0.0049)	-0.007 (0.0054)	-0.009 (0.0044)	-0.011 (0.0092)
CARE6am	0.003 (0.0055)	0.005 (0.0059)	0 (0.0056)	0.003 (0.0063)	0.006 (0.0058)	0.003 (0.0075)	0.005 (0.0082)	0.017* (0.008)	0.002 (0.0069)	0.023 (0.0135)
CARE9am	0.005 (0.0075)	0.003 (0.008)	0.01 (0.0076)	0.004 (0.0085)	0.007 (0.0078)	0.009 (0.0104)	0 (0.0108)	0.023* (0.0113)	0.006 (0.0094)	0.022 (0.0198)
CARE12pm	0.006 (0.0077)	0.005 (0.0082)	0.007 (0.0076)	0 (0.0086)	0.007 (0.008)	0.012 (0.0108)	-0.003 (0.0108)	0.021 (0.0116)	0.007 (0.0095)	0.03 (0.0239)
CARE3pm	0.014 (0.0078)	0.016 (0.0083)	0.009 (0.0076)	0.012 (0.0087)	0.016 (0.0082)	0.019 (0.011)	0.007 (0.0108)	0.035** (0.0116)	0.023* (0.0095)	0.019 (0.0238)
CARE6pm	0.024*** (0.0072)	0.026*** (0.0076)	0.019** (0.007)	0.030*** (0.0082)	0.026*** (0.0076)	0.027** (0.0099)	0.02 (0.0104)	0.058*** (0.0104)	0.041*** (0.0086)	0.012 (0.0223)
CARE9pm	0.023*** (0.0049)	0.024*** (0.0052)	0.020*** (0.0048)	0.028*** (0.0056)	0.024*** (0.0051)	0.023*** (0.0065)	0.023** (0.0074)	0.043*** (0.0071)	0.034*** (0.0059)	-0.002 (0.0135)
Temp ≥ 65°F	0.020*** (0.0003)	0.019*** (0.0003)	0.026*** (0.0003)	0.019*** (0.0003)	0.021*** (0.0003)	0.021*** (0.0004)	0.018*** (0.0004)	0.024*** (0.001)	0.020*** (0.0003)	0.018*** (0.0003)
Temp < 65°F	0.025*** (0.0003)	0.023*** (0.0003)	0.030*** (0.0003)	0.024*** (0.0003)	0.025*** (0.0003)	0.026*** (0.0004)	0.022*** (0.0004)	0.030*** (0.001)	0.024*** (0.0003)	0.022*** (0.0003)
Constant	1.311*** (0.0052)	1.326*** (0.005)	1.250*** (0.0061)	1.308*** (0.0056)	1.308*** (0.0055)	1.172*** (0.0089)	1.436*** (0.0063)	1.274*** (0.0173)	1.357*** (0.0049)	1.378*** (0.0078)
R-squared	0.558	0.569	0.553	0.573	0.465	0.534	0.573	0.512	0.578	0.596
N clusters	30,072	30,068	30,023	30,059	26,667	12,021	18,051	1,524	22,118	24,666
N observations	143,557,960	102,690,602	40,867,134	124,571,172	130,214,929	57,097,674	86,460,276	10,192,045	105,394,612	112,628,194

Standard errors in parentheses.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

7.5.3.15. Intra-day specification with other programs

	Full sample	10% sample	All programs
CARE12am	0.110*** (0.0115)	0.108*** (0.012)	0.108*** (0.012)
CARE3am	-0.005 (0.0036)	-0.011 (0.006)	-0.011 (0.006)
CARE6am	0.003 (0.0055)	0.005 (0.0076)	0.005 (0.0076)
CARE9am	0.005 (0.0075)	0.01 (0.0092)	0.01 (0.0092)
CARE12pm	0.006 (0.0077)	0.004 (0.0097)	0.004 (0.0097)
CARE3pm	0.014 (0.0078)	0.020* (0.0097)	0.020* (0.0097)
CARE6pm	0.024*** (0.0072)	0.026** (0.0092)	0.026** (0.0092)
CARE9pm	0.023*** (0.0049)	0.027*** (0.0073)	0.027*** (0.0073)
Temp ≥ 65°F	0.020*** (0.0003)	0.020*** (0.0003)	0.020*** (0.0003)
Temp < 65°F	0.025*** (0.0003)	0.025*** (0.0003)	0.025*** (0.0003)
Rebate12am			0.022 (0.0136)
Rebate3am			0.001 (0.0079)
Rebate6am			0.033*** (0.01)
Rebate9am			0.013 (0.0121)
Rebate12pm			0.007 (0.0118)
Rebate3pm			0.023 (0.0121)
Rebate6pm			0.048*** (0.0122)
Rebate9pm			0.035*** (0.0095)
BPP12am			0.081*** (0.021)
BPP3am			-0.011 (0.0109)
BPP6am			-0.014 (0.0137)
BPP9am			-0.032 (0.0167)
BPP12pm			-0.01 (0.0168)
BPP3pm			0.017 (0.0161)
BPP6pm			0.015 (0.0166)

	Full sample	10% sample	All programs
BPP9pm			0.025 (0.0133)
ClimateSmart12am			-0.09 (0.3071)
ClimateSmart3am			0.122 (0.1007)
ClimateSmart6am			0.323 (0.1687)
ClimateSmart9am			0.074 (0.2471)
ClimateSmart12pm			-0.221 (0.3097)
ClimateSmart3pm			0.099 (0.1858)
ClimateSmart6pm			-0.36 (0.3484)
ClimateSmart9pm			-0.176 (0.3208)
DirAccess12am			0.076* (0.0317)
DirAccess3am			-0.016 (0.0178)
DirAccess6am			0.026 (0.0191)
DirAccess9am			-0.005 (0.0211)
DirAccess12pm			0.017 (0.0228)
DirAccess3pm			0.01 (0.0252)
DirAccess6pm			0.005 (0.021)
DirAccess9pm			0.017 (0.0191)
SmartAC12am			0.056* (0.0268)
SmartAC3am			0.013 (0.0144)
SmartAC6am			0.021 (0.0177)
SmartAC9am			-0.018 (0.0208)
SmartAC12pm			-0.022 (0.0187)
SmartAC3pm			-0.001 (0.019)
SmartAC6pm			0.009 (0.0199)
SmartAC9pm			-0.016 (0.0173)
SmartRate12am			0.058 (0.0371)
SmartRate3am			0.028 (0.0325)

	Full sample	10% sample	All programs
SmartRate6am			-0.004
			(0.0237)
SmartRate9am			-0.038
			(0.0302)
SmartRate12pm			-0.039
			(0.035)
SmartRate3pm			0.014
			(0.039)
SmartRate6pm			-0.011
			(0.0411)
SmartRate9pm			-0.013
			(0.031)
Constant	1.311***	1.313***	1.303***
	(0.0052)	(0.0054)	(0.0057)
R-squared	0.558	0.565	0.565
N clusters	30,072	30,023	30,023
N observations	143,557,960	14,354,313	14,354,313

7.5.3.16. Linear electricity consumption independent variable

	Linear kWh
CARE baseline	0.748*** (0.1036)
CARE3am	-0.087** (0.0308)
CARE6am	-0.106* (0.0524)
CARE9am	-0.056 (0.0815)
CARE12pm	0.062 (0.0903)
CARE3pm	0.231* (0.0983)
CARE6pm	0.393*** (0.0825)
CARE9pm	0.295*** (0.0488)
Temp ≥ 65°F	0.310*** (0.0098)
Temp < 65°F	0.353*** (0.0096)
Constant	3.608*** (0.1487)
R-squared	0.686
N clusters	30,112
N observations	146,162,009

7.5.3.17. Monthly specification, ln(kWh)

	Monthly
CARE	0.133*** (0.012)
Cooling degree-days	0.003*** (0.00003)
Heating degree-days	0.000006 (0.000004)
Constant	0.583*** (0.007)
R-squared	0.705
N clusters	30,125
N observations	621,241

Supporting information for Chapter 4

7.6. Effects of aggregating from 1-hour to 4-hour resolution

The main analyses aggregate hourly renewable electricity production profiles to 4-hour blocks, beginning with the first hour of the year. This can have the effect of reducing peaks and increasing local minimum electricity production, particularly for solar.

In the solar Next-decade case with a high-temperature DAC, the levelized cost of electrofuel is $5.43/GGE with 1-hour resolution, falling to $5.36/GGE in the 4-hour resolution case used in the main analysis, a difference of 1.2%. Aggregation effects are likely to be smaller for wind production, where the system does not have as strict a diurnal pattern and regularly attains maximum power for hours at a time (see the SI, Section 7.8.2).

Aggregation to 8-hour blocks can introduce a small amount of electricity consumption through the night, unphysically eliminating the need for electricity storage or grid electricity in some solar cases.

7.7. Potential effects of integer variables and constraints and nonlinearities

In practice, operation of an electrofuel production system would likely also have integer variables, such as minimum capacity or operation levels, particularly for high-temperature components such as the fuel synthesis or high-temperature DAC. In practice, however, the system tends to operate the DAC and fuel synthesis at high levels with minimal intra-day ramping, suggesting that adding such integer variables or constraints likely would not substantially change the results.

There are also many nonlinearities in true engineered systems, which system designers and operators would need to consider when building and operating systems. Accounting for these nonlinearities is outside the scope of this analysis.

7.8. Detailed results

7.8.1. System operation characteristics

Table 21. Installed capacity for each production component.

Solar PV	Today, no storage	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Renewables [MW(e)]	5339	4741	4110	4047	3470	3427
Electrolyzer [MW(e)]	3859	2876	2713	2741	2361	2419
Kiln [MW(e)]	1136	408	215	83	201	0
DAC [t(CO_2)/yr]	4307	1045	1153	1165	1064	1284
Fuel synthesis [MW(fuel)]	1756	603	615	626	619	624
Grid connection [MW(e)]	0	10	48	54	14	9
Fuel pipeline [MW(fuel)]	1756	603	615	626	619	624
Natural gas [MW(th)]	0	262	192	214	177	99

Wind	Today, no storage	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Renewables [MW(e)]	2715	2504	2079	2035	1768	1713
Electrolyzer [MW(e)]	1962	1819	1755	1812	1521	1600
Kiln [MW(e)]	577	322	223	38	222	0
DAC [t(CO_2)/yr]	2190	1164	1274	1307	1265	1332
Fuel synthesis [MW(fuel)]	893	715	807	816	768	788
Grid connection [MW(e)]	0	50	48	55	31	34
Fuel pipeline [MW(fuel)]	893	715	807	816	768	788
Natural gas [MW(th)]	0	292	212	240	211	186

Table 22. Capacity factor for each production component.

Solar PV	Today, no storage	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Renewables	0.29	0.29	0.29	0.29	0.29	0.58
Electrolyzer	0.39	0.39	0.38	0.39	0.38	0.62
Kiln	0.11	0.14	0.07	0.12	Null	0.15
DAC	0.91	0.89	0.94	0.96	0.95	0.82
Fuel synthesis	0.85	0.83	0.82	0.83	0.82	0.72
Grid connection	0.17	0.28	0.31	0.06	0.03	0.02
Fuel pipeline	0.75	0.75	0.49	0.83	0.26	0.67
Natural gas	0.85	0.83	0.82	0.83	0.82	0.72

Wind	Today, no storage	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Renewables	0.58	0.58	0.58	0.58	0.58	0.58
Electrolyzer	0.62	0.60	0.58	0.60	0.57	0.62
Kiln	0.15	0.14	0.07	0.22	Null	0.15
DAC	0.82	0.81	0.83	0.83	0.90	0.82
Fuel synthesis	0.72	0.64	0.63	0.67	0.65	0.72
Grid connection	0.02	0.01	0.02	0.01	0.02	0.02
Fuel pipeline	0.67	0.67	0.45	0.61	0.21	0.67
Natural gas	0.72	0.64	0.63	0.67	0.65	0.72

Table 23. Installed storage capacity.

Solar PV	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Electricity [MWh(e)]	734	0	0	83	208
Hydrogen [MWh(H_2)]	9840	9225	9581	8988	9309
Heat [MWh(th)]	974	0	0	57	649
CO_2 [t(CO_2)]	561	2228	1587	32208	11943
Fuel [kWh(fuel)]	0	0	0	0	0

Wind	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Electricity [MWh(e)]	0	0	0	0	0
Hydrogen [MWh(H_2)]	9548	4252	5374	8069	10488
Heat [MWh(th)]	169	0	0	0	283
CO_2 [t(CO_2)]	2731	9026	7559	33390	44035
Fuel [kWh(fuel)]	0	0	0	0	0

Table 24. Average waste electricity as a percentage of total renewable electricity production.

	Today, low-temp	Next-decade, high-temp	Next-decade, low-temp	Breakthrough, high-temp	Breakthrough, low-temp
Solar PV	7.7%	6.3%	5.9%	4.4%	4.2%
Wind	11.4%	5.2%	3.7%	2.8%	2.8%

7.8.2. Component operation profiles

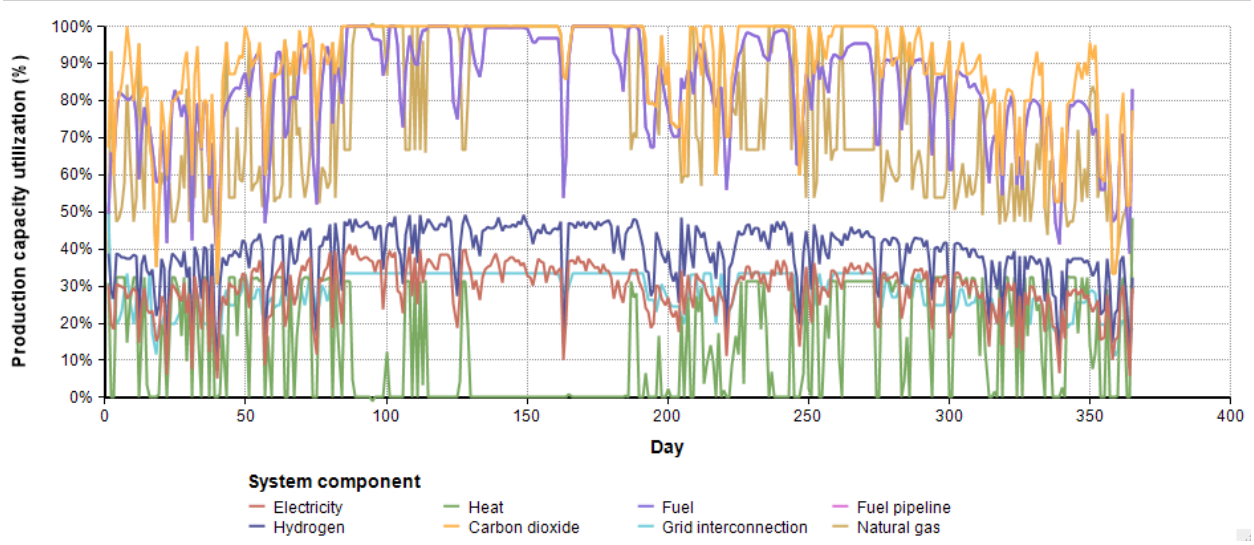

Figure 39. Annual average operation profile for production components for the solar Next-decade high-temperature case. The DAC operates almost constantly during the summer, made possible through use of natural gas, grid electricity, and heat and electricity storage.

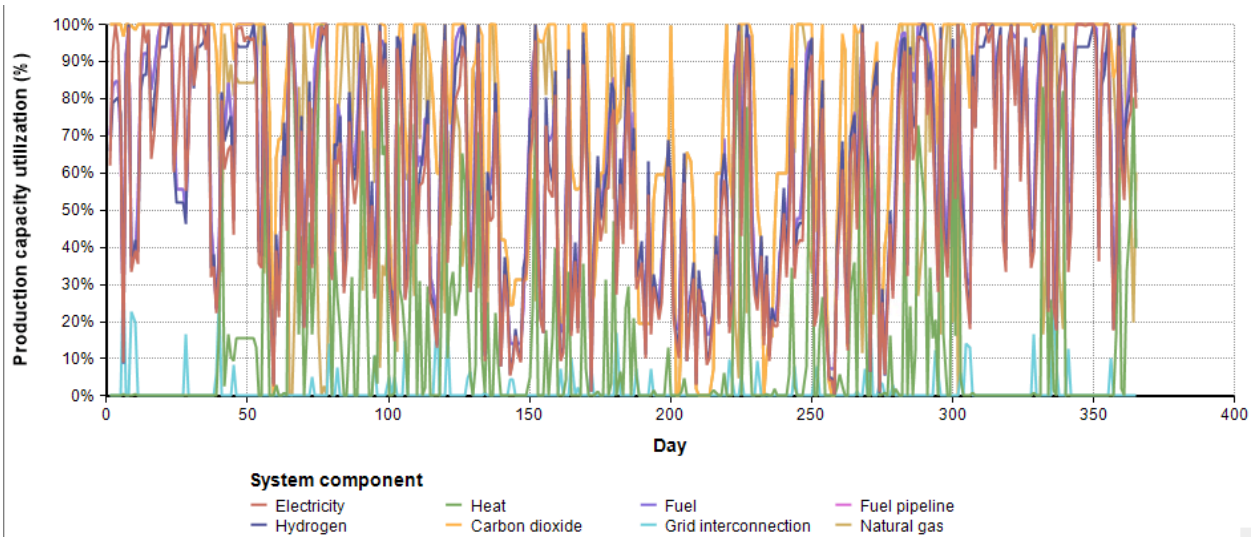

Figure 40. Annual average operation profile for production components for the wind Next-decade high-temperature case. Despite a capacity factor above 57%, wind's irregular and highly seasonal production profile require substantial curtailment of DAC production, particularly during the summer months.

7.9. Supplementary results

7.9.1. No carbon price

Removing the carbon price in the solar Next-decade high-temperature case results in electrofuel at $4.89/GGE, a cost reduction of 8.5% from the corresponding case with a carbon price. The bulk of this cost savings is simply due to the absence of $0.44/GGE in carbon payments. The net fossil carbon content of the resulting electrofuel rises by 21%, thus reducing the effective mitigation value of the electrofuel.

7.9.2. No fossil carbon

Removing natural gas and grid electricity requires the system to operate entirely using renewable electricity, eliminating the possibility of directly attributable fossil greenhouse gas emissions (excepting life-cycle emissions from materials and construction). In the solar Next-decade high-temperature case, this increases electrofuel cost to $5.68/GGE, an increase of 6%.

7.9.3. Seasonal storage of CO_2

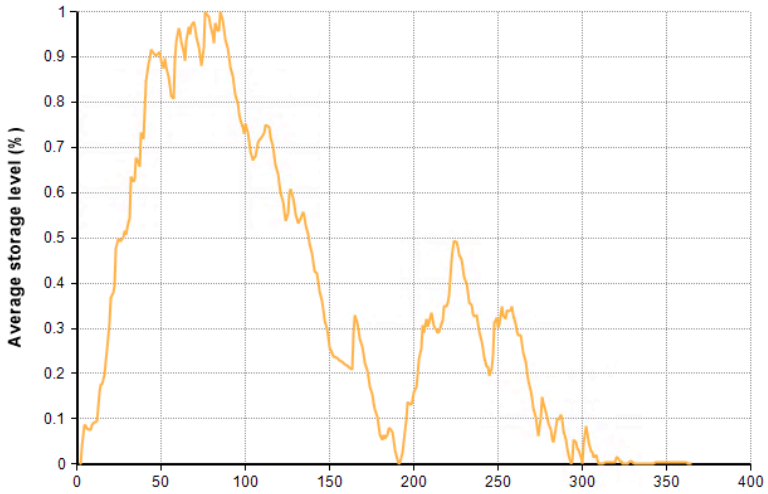

Figure 41. CO_2 storage levels as a fraction of capacity in the solar Next-decade high-temperature case. Note that CO_2 is stored for almost 200 days, from winter to summer, allowing the system to undersize the DAC while making use of abundant summertime solar irradiation for hydrogen production for fuel synthesis.

7.10. Barriers to mass adoption of electrofuels

Fischer-Tropsch liquids are currently used for a variety of commercial applications. Fischer-Tropsch-based synthetic paraffinic kerosene is approved by the ASTM for blending into jet fuel at up to 50% [164]. Higher levels of blending would require either blending of supplementary aromatics or engine modifications [164].

Fischer-Tropsch gasoline and diesel have been used commercially in South African vehicles for decades with few if any vehicle modification requirements [112].

7.11. Land and water use

Estimates of solar land use are based on average total direct area impacts of systems larger than 20 MW(e) for 1-axis tracking installations, 8.3-9.0 acre/MW(e) [239].

Estimates of wind land use of 0.1-1.5 MW(e)/ha are based on permanent direct area impact for utility-scale wind turbines [240].

8. References

1. Fuss, S. *et al.* Betting on negative emissions. *Nat. Clim. Change* **4**, 850–853 (2014).
2. *Climate change 2014: mitigation of climate change: Working Group III contribution to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*. (Cambridge University Press, 2014).
3. Smil, V. Perils of long-range energy forecasting: reflections on looking far ahead. *Technol. Forecast. Soc. Change* **65**, 251–264 (2000).
4. Koomey, J., Craig, P., Gadgil, A. & Lorenzetti, D. Improving long-range energy modeling: A plea for historical retrospectives. *Energy J.* **24**, 75–92 (2003).
5. Wiser, R. *et al. 2017 Wind Technologies Market Report*. (U.S. Department of Energy, Office of Energy Efficiency & Renewable Energy, 2017).
6. Bolinger, M. & Seel, J. *Utility-Scale Solar: Empirical Trends in Project Technology, Cost, Performance, and PPA Pricing in the United States – 2018 Edition*. (2018). doi:10.2172/1477381
7. Fu, R., Remo, T. & Margolis, R. 2018 U.S. Utility-Scale Photovoltaics-Plus-Energy Storage System Costs Benchmark. *Renew. Energy* 32 (2018).
8. Boden, T., Marland, G. & Andres, R. J. *National CO2 Emissions from Fossil-Fuel Burning, Cement Manufacture, and Gas Flaring: 1751-2014*. (Carbon Dioxide Information Analysis Center, Oak Ridge National Laboratory, U.S. Department of Energy, 2017).
9. *Value added by Industry*. (U.S. Bureau of Economic Analysis, 2017).
10. Huss, W. R. Can electric utilities improve their forecast accuracy? The historical perspective. *Public Util. Fortn.* 3–8 (1985).

11. Huss, William R. Comparative analysis of company forecasts and advanced time series techniques using annual electric utility energy sales data. *Int. J. Forecast.* **1**, 217–239 (1985).
12. Huss, W. R. What makes a good load forecast. *Public Util. Fortn.* 3–11 (1985).
13. Nelson, C. R. & Peck, S. C. The NERC fan: A retrospective analysis of the NERC summary forecasts. *J. Bus. Econ. Stat.* **3**, 179–187 (1985).
14. Landsberg, H. Energy in transition: A view from 1960. *Energy J.* **6**, 1–18 (1985).
15. Huntington, H. G. Oil price forecasting in the 1980s: What went wrong? *Energy J.* **15**, 1–22 (1994).
16. Sohn, I. Long-term energy projections: What lessons have we learned? *Energy Policy* **35**, 4574–4584 (2007).
17. *Issues in Midterm Analysis and Forecasting 1996*. (Energy Information Administration, 1996).
18. Shlyakhter, A. I., Kammen, D. M., Broido, C. L. & Wilson, R. Quantifying the credibility of energy projections from trends in past data: The US energy sector. *Energy Policy* **22**, 119–130 (1994).
19. Kaack, L. H., Apt, J., Morgan, M. G. & McSharry, P. Empirical prediction intervals improve energy forecasting. *Proc. Natl. Acad. Sci.* **114**, 8752–8757 (2017).
20. Craig, P. P., Gadgil, A. & Koomey, J. G. What can history teach us? A retrospective examination of long-term energy forecasts for the United States*. *Annu. Rev. Energy Environ.* **27**, 83–118 (2002).

21. Bezdek, R. H. & Wendling, R. M. A half century of long-range energy forecasts: errors made, lessons learned, and implications for forecasting. *J. Fusion Energy* **21**, 155–172 (2002).

22. Linderoth, H. Forecast errors in IEA-countries' energy consumption. *Energy Policy* **30**, 53–61 (2002).

23. Auffhammer, M. The rationality of EIA forecasts under symmetric and asymmetric loss. *Resour. Energy Econ.* **29**, 102–121 (2007).

24. O'Neill, B. C. & Desai, M. Accuracy of past projections of US energy consumption. *Energy Policy* **33**, 979–993 (2005).

25. Winebrake, J. J. & Sakva, D. An evaluation of errors in US energy forecasts: 1982–2003. *Energy Policy* **34**, 3475–3483 (2006).

26. Considine, T. J. & Clemente, F. A. Gas-market forecasts: Betting on bad numbers. *Public Util. Fortn.* 53–59 (2007).

27. Fischer, C., Herrnstadt, E. & Morgenstern, R. Understanding errors in EIA projections of energy demand. *Resour. Energy Econ.* **31**, 198–209 (2009).

28. Wara, M., Cullenward, D. & Teitelbaum, R. Peak electricity and the clean power plan. *Electr. J.* **28**, 18–27 (2015).

29. Gilbert, A. Q. & Sovacool, B. K. Looking the wrong way: Bias, renewable electricity, and energy modelling in the United States. *Energy* **94**, 533–541 (2016).

30. *Annual Energy Outlook reports, 1982-1987, and 1989-2015*. (Energy Information Administration).

31. *Monthly Energy Review, November 2015*. 85 (Energy Information Administration, 2015).

32. *Current-Dollar and 'Real' Gross Domestic Product*. (U.S. Department of Commerce Bureau of Economic Analysis, 2015).
33. *Gross Domestic Product: Implicit Price Deflator [A191RI1A225NBEA]*. (US. Bureau of Economic Analysis, 2016).
34. *Energy Policy Act Transportation Study: Interim Report on Natural Gas Flows and Rates*. (Energy Information Administration, 1995).
35. *Annual Energy Outlook Retrospective Review: Evaluation of 2014 and Prior Reference Case Projections*. (Energy Information Administration, 2015).
36. *Form EIA-860 detailed data, 2013*. (Energy Information Administration, 2014).
37. Kilian, L. & Hicks, B. Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *J. Forecast.* **32**, 385–394 (2013).
38. Sundquist, E. & McCahill, C. *For the first time in a decade, U.S. per capita highway travel ticks up*. (State Smart Transportation Initiative, 2015).
39. Evergreen Economics. *Needs Assessment for the Energy Savings Assistance and the California Alternate Rates for Energy Programs*. (Evergreen Economics, 2013).
40. Mahoney, M. & O'Boyle, M. *A National Survey of Electric and Gas Utility Rate Structures for Low-Income Customers*. (Arizona State University, 2013).
41. Espey, J. A. & Espey, M. Turning on the Lights: A Meta-Analysis of Residential Electricity Demand Elasticities. *J. Agric. Appl. Econ.* **36**, 65–81 (2004).
42. Azevedo, I. M. L., Morgan, M. G. & Lave, L. Residential and Regional Electricity Consumption in the U.S. and EU: How Much Will Higher Prices Reduce CO2 Emissions? *Electr. J.* **24**, 21–29 (2011).

43. Siler-Evans, K., Azevedo, I. L. & Morgan, M. G. Marginal Emissions Factors for the U.S. Electricity System. *Environ. Sci. Technol.* **46**, 4742–4748 (2012).

44. Siler-Evans, K., Azevedo, I. L., Morgan, M. G. & Apt, J. Regional variations in the health, environmental, and climate benefits of wind and solar generation. *Proc. Natl. Acad. Sci.* **110**, 11768–11773 (2013).

45. Brancucci Martinez-Anido, C. & Hodge, B. *Impact of Utility-Scale Distributed Wind on Transmission-Level System Operations*. (2014). doi:10.2172/1159355

46. Lijesen, M. G. The real-time price elasticity of electricity. *Energy Econ.* **29**, 249–258 (2007).

47. Labandeira, X., Labeaga, J. M. & López-Otero, X. A meta-analysis on the price elasticity of energy demand. *Energy Policy* **102**, 549–568 (2017).

48. Miller, M. & Alberini, A. Sensitivity of price elasticity of demand to aggregation, unobserved heterogeneity, price trends, and price endogeneity: Evidence from U.S. Data. *Energy Policy* **97**, 235–249 (2016).

49. Alberini, A. & Filippini, M. Response of residential electricity demand to price: The effect of measurement error. *Energy Econ.* **33**, 889–895 (2011).

50. Alberini, A., Khymych, O. & Ščasný, M. Response to Extreme Energy Price Changes: Evidence from Ukraine. *Energy J.* **40**, (2019).

51. Paul, A. C., Myers, E. C. & Palmer, K. L. A Partial Adjustment Model of U.S. Electricity Demand by Region, Season, and Sector. *SSRN Electron. J.* (2009). doi:10.2139/ssrn.1372228

52. Chang, Y., Choi, Y., Kim, C. S., Miller, J. I. & Park, J. Y. Disentangling temporal patterns in elasticities: A functional coefficient panel analysis of electricity demand. *Energy Econ.* **60**, 232–243 (2016).

53. Bernstein, M. A. & Griffin, J. M. *Regional differences in the price-elasticity of demand for energy*. (Citeseer, 2006).

54. Chang, Y., Choi, Y., Kim, C. S., Miller, J. I. & Park, J. Y. Disentangling temporal patterns in elasticities: A functional coefficient panel analysis of electricity demand. *Energy Econ.* **60**, 232–243 (2016).

55. Cooper, A. *Electric Company Smart Meter Deployments: Foundation for a Smart Grid*. (The Edison Foundation: Institute for Electric Innovation, 2017).

56. Jessoe, K. & Rapson, D. Knowledge is (Less) Power: Experimental Evidence from Residential Energy Use. *Am. Econ. Rev.* **104**, 1417–1438 (2014).

57. Boomhower, J. & Davis, L. W. *Do Energy Efficiency Investments Deliver at the Right Time?* (National Bureau of Economic Research, 2017).

58. Qiu, Y. & Kahn, M. E. Better sustainability assessment of green buildings with high-frequency data. *Nat. Sustain.* **1**, 642–649 (2018).

59. Norris, S. CARE NEUAC 2013. (2013).

60. *Final Report on Phase 2 Low Income Needs Assessment*. (KEMA, Inc. Prepared for: California Public Utilities Commission, 2007).

61. Palmgren, C., Stevens, N., Goldberg, M. & Rothkin, K. *2009 California Residential Appliance Saturation Study*. (California Energy Commission, 2010).

62. Meyer, R. M. Analysis of Selected Regulatory Interventions to Improve Energy Efficiency. (Carnegie Mellon University, 2014).

63. PG&E. Take control with Time-of-Use rate plans. (2019).

64. Menne, M. J. *et al.* Global Historical Climatology Network - Daily, Version 3 [California]. (2012).

65. Azevedo, I. L., Horner, N. C., Siler-Evans, K. & Vaishnav, P. T. Electricity Marginal Factor Estimates. (2017).

66. CPUC. *2018 Avoided Cost Calculator*. (California Public Utilities Commission, prepared by Energy + Environmental Economics, 2018).

67. Muller, N. Z. AP2 Model. (2017).

68. PG&E. Find out if Peak Day Pricing is right for your business. *Pacific Gas & Electric* (2017). Available at: https://www.pge.com/en_US/business/rate-plans/rate-plans/peak-day-pricing/peak-day-pricing.page.

69. Drehobl, A. & Ross, L. *How Energy Efficiency Can Improve Low Income and Underserved Communities*. 56 (American Council for an Energy-Efficient Economy, 2016).

70. Skumatz, L. A. *Non-Energy Benefits/Non-Energy Impacts (NEBs/NEIs) and Their Role & Values in Cost-Effectiveness Tests: State of Maryland*. 71 (Natural Resourced Defense Council, prepared by Skumatz Economic Research Associates Inc., March 312014).

71. Holland, S. P., Mansur, E. T., Muller, N. & Yates, A. J. Decompositions and Policy Consequences of an Extraordinary Decline in Air Pollution from Electricity Generation. *NBER Work. Pap.* (2018). doi:10.3386/w25339

72. Denholm, P., O'Connell, M., Brinkman, G. & Jorgenson, J. *Overgeneration from Solar Energy in California. A Field Guide to the Duck Chart*. (2015). doi:10.2172/1226167

73. Fowlie, M., Greenstone, M. & Wolfram, C. *Do energy efficiency investments deliver? Evidence from the weatherization assistance program*. (National Bureau of Economic Research, 2015).

74. Ito, K. Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing. *Am. Econ. Rev.* **104**, 537–563 (2014).

75. *United States Summary: 2010 Population and Housing Unit Counts*. (U.S. Census Bureau, 2012).

76. Davis, S. J. *et al.* Net-zero emissions energy systems. *Science* **360**, eaas9793 (2018).

77. IEA. *Technology Roadmap: Biofuels for Transport*. (International Energy Agency, 2011).

78. Brynolf, S., Taljegard, M., Grahn, M. & Hansson, J. Electrofuels for the transport sector: A review of production costs. *Renew. Sustain. Energy Rev.* **81**, 1887–1905 (2018).

79. Fasihi, M., Bogdanov, D. & Breyer, C. Economics of Global Gas-To-Liquids (GtL) Fuels Trading Based on Hybrid Pv-Wind Power Plants. in 1–20 (International Solar Energy Society, 2016). doi:10.18086/swc.2015.09.03

80. Fasihi, M. & Breyer, C. Synthetic Methanol and Dimethyl Ether Production based on Hybrid PV-Wind Power Plants. in 15

81. Fasihi, M., Bogdanov, D. & Breyer, C. Techno-Economic Assessment of Power-to-Liquids (PtL) Fuels Production and Global Trading Based on Hybrid PV-Wind Power Plants. in *Energy Procedia* **99**, 243–268 (2016).

82. Agusdinata, D. B., Zhao, F., Ileleji, K. & DeLaurentis, D. Life Cycle Assessment of Potential Biojet Fuel Production in the United States. *Environ. Sci. Technol.* **45**, 9133–9143 (2011).

83. Committee on Developing a Research Agenda for Carbon Dioxide Removal and Reliable Sequestration *et al. Negative Emissions Technologies and Reliable Sequestration: A Research Agenda*. (National Academies Press, 2018). doi:10.17226/25259

84. Staples, M. D., Malina, R., Suresh, P., Hileman, J. I. & Barrett, S. R. H. Aviation CO 2 emissions reductions from the use of alternative jet fuels. *Energy Policy* **114**, 342–354 (2018).

85. Baroutaji, A., Wilberforce, T., Ramadan, M. & Olabi, A. G. Comprehensive investigation on hydrogen and fuel cell technology in the aviation and aerospace sectors. *Renew. Sustain. Energy Rev.* **106**, 31–40 (2019).

86. Freeman, J. M. *et al. System Advisor Model (SAM) General Description (Version 2017.9.5)*. (2018). doi:10.2172/1440404

87. *Levelized Cost and Levelized Avoided Cost of New Generation Resources in the Annual Energy Outlook 2018*. 20 (U.S. Energy Information Administration, 2018).

88. Deign, J. Mexico's Record Solar Prices Fall Below the Average Cost of Energy From Gas and Coal. *Greentech Media* (2018).

89. *North America Chlor-alkali Market Analysis By Product (Caustic Soda, Chlorine, Soda Ash), Trends & Dynamics, Competitive Landscape, Value Chain Analysis, Price Trend Analysis, And Forecasts, 2018 - 2024*. (Grand View Research, 2017).

90. Schmidt, O. *et al.* Future cost and performance of water electrolysis: An expert elicitation study. *Int. J. Hydrog. Energy* **42**, 30470–30492 (2017).

91. Luca Bertuccioli *et al. Development of Water Electrolysis in the European Union*. 160 (E4tech Sàrl with Element Energy Ltd for the Fuel Cells and Hydrogen Joint Undertaking, 2014).

92. Graves, C., Ebbesen, S. D. & Mogensen, M. Co-electrolysis of CO2 and H2O in solid oxide cells: Performance and durability. *Solid State Ion.* **192**, 398–403 (2011).
93. van Vliet, O. P. R., Faaij, A. P. C. & Turkenburg, W. C. Fischer–Tropsch diesel production in a well-to-wheel perspective: A carbon, energy flow and cost analysis. *Energy Convers. Manag.* **50**, 855–876 (2009).
94. Rafati, M. *et al.* Techno-economic analysis of production of Fischer-Tropsch liquids via biomass gasification: The effects of Fischer-Tropsch catalysts and natural gas co-feeding. *Energy Convers. Manag.* **133**, 153–166 (2017).
95. CPI-Average Price Data. (2018).
96. Euro (EUR) to U.S. dollar (USD) annual average exchange rate from 1999 to 2017. (2018).
97. Keith, D. W., Holmes, G., St. Angelo, D. & Heidel, K. A process for capturing CO2 from the atmosphere. *Joule* (2018).
98. Sandalow, D., Friedmann, J. & McCormick, C. *Direct Air Capture of Carbon Dioxide: ICEF Roadmap 2018*. (2018).
99. Wurzbacher, J. A., Gebald, C., Brunner, S. & Steinfeld, A. Heat and mass transfer of temperature–vacuum swing desorption for CO2 capture from air. *Chem. Eng. J.* **283**, 1329–1338 (2016).
100. Socolow, R. *et al. Direct Air Capture of CO2 with Chemicals A Technology Assessment for the APS Panel on Public Affairs*. (American Physical Society, 2011).
101. Herzog, H. *Assessing the Feasibility of Capturing CO2 from the Air*. 16 (Massachussets Institute of Technology, Laboratory for Energy and the Environment, 2003).
102. Gertner, J. The Tiny Swiss Company That Thinks It Can Help Stop Climate Change. *New York Times* (2019).

103. Kintisch, E. Can Sucking CO2 Out of the Atmosphere Really Work? *MIT Technology Review* (2014).

104. Direct Air Capture & Storage (DACS) Factsheet for Researchers. (2018).

105. Industrial electricity prices in selected countries in Europe 2017, by amount consumed (in euro cents per kilowatt hour). (2019).

106. 2017 Average Monthly Bill- Industrial. (2017).

107. Yixin Sunny Furnace Co. Ltd. High Temperature Electric Ceramic Tunnel Kiln For Sale. (2019).

108. DOE. *Electric Resistance Heating*. (U.S. Department of Energy, 2019).

109. Dry, M. E. The Fischer–Tropsch process: 1950–2000. *Catal. Today* 15 (2002).

110. Leckel, D. Diesel Production from Fischer−Tropsch: The Past, the Present, and New Concepts. *Energy Fuels* **23**, 2342–2358 (2009).

111. Becker, W. L., Braun, R. J., Penev, M. & Melaina, M. Production of Fischer–Tropsch liquid fuels from high temperature solid oxide co-electrolysis units. *Energy* **47**, 99–115 (2012).

112. Dry, M. E. High quality diesel via the Fischer-Tropsch process - a review. *J. Chem. Technol. Biotechnol.* **77**, 43–50 (2002).

113. Kreutz, T. G., Larson, E. D., Liu, G. & Williams, R. H. Fischer-Tropsch Fuels from Coal and Biomass. 86 (2008).

114. Tremel, A. *Electricity-based Fuels*. (Springer International Publishing, 2018). doi:10.1007/978-3-319-72459-1

115. Darling, R. M., Gallagher, K. G., Kowalski, J. A., Ha, S. & Brushett, F. R. Pathways to low-cost electrochemical energy storage: a comparison of aqueous and nonaqueous flow batteries. *Energy Env. Sci* **7**, 3459–3477 (2014).

116. Kittner, N., Lill, F. & Kammen, D. M. Energy storage deployment and innovation for the clean energy transition. *Nat. Energy* **2**, (2017).

117. Luo, X., Wang, J., Dooner, M. & Clarke, J. Overview of current development in electrical energy storage technologies and the application potential in power system operation. *Appl. Energy* **137**, 511–536 (2015).

118. Schimpe, M. *et al.* Energy efficiency evaluation of a stationary lithium-ion battery container storage system via electro-thermal modeling and detailed component analysis. *Appl. Energy* **210**, 211–229 (2018).

119. Zakeri, B. & Syri, S. Electrical energy storage systems: A comparative life cycle cost analysis. *Renew. Sustain. Energy Rev.* **42**, 569–596 (2015).

120. Cole, W. J., Marcy, C., Krishnan, V. K. & Margolis, R. Utility-scale lithium-ion storage cost projections for use in capacity expansion models. in *2016 North American Power Symposium (NAPS)* 1–6 (IEEE, 2016). doi:10.1109/NAPS.2016.7747866

121. Ramsden, T., Kroposki, B. & Levene, J. *Opportunities for Hydrogen-Based Energy Storage for Electric Utilities*. 17 (United States Department of Energy, 2008).

122. *2015 Annual Progress Report: DOE Hydrogen and Fuel Cells Program*. (2015).

123. Reuß, M. *et al.* Seasonal storage and alternative carriers: A flexible hydrogen supply chain model. *Appl. Energy* **200**, 290–302 (2017).

124. Lord, A. S. *Overview of geologic storage of natural gas with an emphasis on assessing the feasibility of storing hydrogen.* SAND2009-5878, 975258 (2009). doi:10.2172/975258

125. Ripepi, G. *Hydrogen storage for variable renewable electricity integration: Techno-economic analysis of a Lined Rock Cavern system*. (Chalmers University of Technology, 2018).

126. Tian, Y. & Zhao, C. Y. A review of solar collectors and thermal energy storage in solar thermal applications. *Appl. Energy* **104**, 538–553 (2013).

127. Kotzé, J. P. Thermal energy storage in metallic phase change materials. 195

128. Herrmann, U., Geyer, M. & Kearney, D. Overview on Thermal Storage Systems. 23

129. *SunShot Vision Study*. (United States Department of Energy, 2012).

130. Islam, M. T., Huda, N., Abdullah, A. B. & Saidur, R. A comprehensive review of state-of-the-art concentrating solar power (CSP) technologies: Current status and research trends. *Renew. Sustain. Energy Rev.* **91**, 987–1018 (2018).

131. Reddy, R. G. *Novel Molten Salts Thermal Energy Storage for Concentrating Solar Power Generation*. (2013). doi:10.2172/1111584

132. Nomura, T. *et al.* Microencapsulation of eutectic and hyper-eutectic Al-Si alloy as phase change materials for high-temperature thermal energy storage. *Sol. Energy Mater. Sol. Cells* **187**, 255–262 (2018).

133. Stutz, B. *et al.* Storage of thermal solar energy. *Comptes Rendus Phys.* **18**, 401–414 (2017).

134. Seo, Y., Huh, C., Lee, S. & Chang, D. Comparison of CO 2 liquefaction pressures for ship-based carbon capture and storage (CCS) chain. *Int. J. Greenh. Gas Control* **52**, 1–12 (2016).

135. *IPCC special report on carbon dioxide capture and storage*. (Cambridge University Press, for the Intergovernmental Panel on Climate Change, 2005).

136. Seo, Y., Lee, S., Kim, J., Huh, C. & Chang, D. Determination of optimal volume of temporary storage tanks in a ship-based carbon capture and storage (CCS) chain using life cycle cost (LCC) including unavailability cost. *Int. J. Greenh. Gas Control* **64**, 11–22 (2017).
137. McCoy, S. T. The Economics of CO2 Transport by Pipeline and Storage in Saline Aquifers and Oil Reservoirs. 267
138. EPIC. *Phase 2 Report: Interregional Transmission Development and Analysis for Three Stakeholder Selected Scenarios And Gas-Electric System Interface Study. Volume 12, Appendices to Section 11*. (Eastern Interconnection Planning Collaborative, 2015).
139. Chevron. Aviation Turbine Fuel Performance. (2000).
140. Southwire. ACSR: Aluminum conductor. Steel reinforced. Bare. (203AD).
141. Pletka, R., Khangura, J., Rawlins, A., Waldren, E. & Wilson, D. Western Electricity Coordinating Council | CAPITAL COSTS FOR TRANSMISSION AND SUBSTATIONS. 35
142. McFayden, S. Three Phase Current - Simple Calculation. (2009).
143. Power Sector Carbon Index. (2018).
144. NaturalGas.org. The Transportation of Natural Gas. (2013).
145. CheCalc.com. Natural Gas Pipeline Sizing. (2019).
146. Brennand, T. Natural gas, a fuel of choice for China. *Energy Sustain. Dev.* **5**, (2001).
147. Parker, N. Using Natural Gas Transmission Pipeline Costs to Estimate Hydrogen Pipeline Costs. 86
148. EIA. United States Natural Gas Industrial Price. (2019).
149. Eng-Tips.com. Crude oil velocity in pipelines for various API grades. (2006).

150. Rui, Z., Metz, P. A., Reynolds, D. B., Chen, G. & Zhou, X. Historical pipeline construction cost analysis. *Int. J. Oil Gas Coal Technol.* **4**, 244 (2011).

151. Pearson, R. J. *et al.* Energy Storage via Carbon-Neutral Fuels Made From CO$_{2}$, Water, and Renewable Energy. *Proc. IEEE* **100**, 440–460 (2012).

152. Zoback, M. D. & Gorelick, S. M. Earthquake triggering and large-scale geologic storage of carbon dioxide. *Proc. Natl. Acad. Sci.* **109**, 10164–10168 (2012).

153. Rubin, E. S., Davison, J. E. & Herzog, H. J. The cost of CO2 capture and storage. *Int. J. Greenh. Gas Control* **40**, 378–400 (2015).

154. Anderson, S. T. & Jahediesfanjani, H. Net Costs of Brine Production to Expand the Pressure-Limited CO2 Storage Capacity of the Mount Simon Sandstone. in (2018).

155. WateReuse, W. *Seawater Desalination Costs*. 20 (WateReuse Association, 2012).

156. Rui, Z., Metz, P. A., Reynolds, D. B., Chen, G. & Zhou, X. Regression models estimate pipeline construction costs. 8

157. EIA. How much carbon dioxide is produced when different fuels are burned? (2018).

158. Ishimoto, Y. *et al.* Putting Costs of Direct Air Capture in Context. *SSRN Electron. J.* (2017). doi:10.2139/ssrn.2982422

159. Perner, J., Unteutsch, M. & Lövenich, A. *The Future Cost of Electricity-Based Synthetic Fuels*. 96 (Agora Energiewende, Agora Verkehrswende, 2018).

160. *Carbon Dioxide Emissions Coefficients*. (Energy Information Administration, 2016).

161. Xie, X., Wang, M. & Han, J. Assessment of Fuel-Cycle Energy Use and Greenhouse Gas Emissions for Fischer−Tropsch Diesel from Coal and Cellulosic Biomass. *Environ. Sci. Technol.* **45**, 3047–3053 (2011).

162. Masnadi, M. S. *et al.* Global carbon intensity of crude oil production. *Science* **361**, 851–853 (2018).

163. BTS. *Airline Fuel Cost and Consumption (U.S. Carriers - Scheduled)*. (United States Bureau of Transportation Statistics, 2018).

164. *Standard Specification for Aviation Turbine Fuel Containing Synthesized Hydrocarbons*. (ASTM International, 2018).

165. Wong-Parodi, G., Dowlatabadi, H., McDaniels, T. & Ray, I. Influencing Attitudes toward Carbon Capture and Sequestration: A Social Marketing Approach. *Environ. Sci. Technol.* **45**, 6743–6751 (2011).

166. *Flame Temperatures Gases*. (Engineeringtoolbox.com, 2003).

167. Sherwin, E. D., Henrion, M. & Azevedo, I. M. L. Estimation of the year-on-year volatility and the unpredictability of the United States energy system. *Nat. Energy* **3**, 341–346 (2018).

168. Maupin, M. A. Summary of Estimated Water Use in the United States in 2015. (2015).

169. Athey, S. & Imbens, G. Recursive partitioning for heterogeneous causal effects. *Proc. Natl. Acad. Sci.* **113**, 7353–7360 (2016).

170. Staggers, H. *Federal Energy Administration Act Amendments of 1976*. (1976).

171. Ribcoff, A. A. *Department of Energy Organization Act*. (1977).

172. *Retrospective Review Annual Energy Outlook 2010*. (Energy Information Administration, 2011).

173. Murphy, F. H., Conti, J. J., Shaw, S. H. & Sanders, R. Modeling and forecasting energy markets with the intermediate future forecasting system. *Oper. Res.* **36**, 406–420 (1988).

174. Vinsel, L. J. The crusade for credible energy information and analysis in the United States, 1973–1982. *Hist. Technol.* **28**, 149–176 (2012).

175. *Integrating Module of the National Energy Modeling System: Model Documentation 2014*. (Energy Information Administration, 2014).

176. Gabriel, S. A., Kydes, A. S. & Whitman, P. The National Energy Modeling System: a large-scale energy-economic equilibrium model. *Oper. Res.* **49**, 14–25 (2001).

177. *Model Documentation Report: Macroeconomic Activity Module (MAM) of the National Energy Modeling System*. (Energy Information Administration, 2014).

178. *International Energy Module of the National Energy Modeling System: Model Documentation 2014*. (Energy Information Administration, 2014).

179. Marimoutou, V., Raggad, B. & Trabelsi, A. Extreme Value Theory and Value at Risk: Application to oil market. *Energy Econ.* **31**, 519–530 (2009).

180. Fong Chan, K. & Gray, P. Using extreme value theory to measure value-at-risk for daily electricity spot prices. *Int. J. Forecast.* **22**, 283–300 (2006).

181. Mehta, A., Neukirchen, M., Pfetsch, S. & Poppensieker, T. *Managing market risk: Today and tomorrow*. (McKinsey & Company, 2012).

182. Zareipour, H., Bhattacharya, K. & Cañizares, C. A. Electricity market price volatility: The case of Ontario. *Energy Policy* **35**, 4739–4748 (2007).

183. Hadsell, L., Marathe, A. & Shawky, H. A. Estimating the volatility of wholesale electricity spot prices in the US. *Energy J.* **25**, 23 (2004).

184. *Handbook of high Frequency Trading*. (Academic Press, 2015).

185. Kilian, L. & Hicks, B. Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *J. Forecast.* **32**, 385–394 (2013).

186. Kilian, L. *Oil price volatility: Origins and effects*. (WTO Staff Working Paper, 2010).

187. Kilian, L. Oil Price Shocks: Causes and Consequences. *Annu. Rev. Resour. Econ.* **6**, 133–154 (2014).

188. Ebrahim, Z., Inderwildi, O. R. & King, D. A. Macroeconomic impacts of oil price volatility: mitigation and resilience. *Front. Energy* **8**, 9–24 (2014).

189. Bernanke, B. S., Gertler, M., Watson, M., Sims, C. A. & Friedman, B. M. Systematic monetary policy and the effects of oil price shocks. *Brook. Pap. Econ. Act.* **1997**, 91–157 (1997).

190. Kellogg, R. The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling †. *Am. Econ. Rev.* **104**, 1698–1734 (2014).

191. Covi, G. Puzzling out the first oil shock. History, politics and the macroeconomy in a forty-year retrospective. *Hist. Econ. THOUGHT POLICY* 57–91 (2015). doi:10.3280/SPE2015-002004

192. Kim, D. H. What is an oil shock? Panel data evidence. *Empir. Econ.* **43**, 121–143 (2012).

193. Hamilton, J. M. What is an Oil Shock? *J. Econom.* 363–398 (2003).

194. Herrera, R. & González, N. The modeling and forecasting of extreme events in electricity spot markets. *Int. J. Forecast.* **30**, 477–490 (2014).

195. Van Robays, I. Macroeconomic uncertainty and oil price volatility. *Oxf. Bull. Econ. Stat.* (2016).

196. Feng, K., Davis, S. J., Sun, L. & Hubacek, K. Drivers of the US CO2 emissions 1997–2013. *Nat. Commun.* **6**, 7714 (2015).

197. *Price of Liquefied Natural Gas Exports*. (Energy Information Administration, 2017).

198. North American LNG Import/Export Terminals: Approved. (2017).

199. A liquid market. *The Economist* (2012).

200. Fisher, J. Henry Prices Too High to Support New Long-Term LNG Contracts, BofA Says. *NGI's Daily Price Index* (2015).

201. EIA. What are Ccf, Mcf, Btu, and therms? How do I convert natural gas prices in dollars per Ccf or Mcf to dollars per Btu or therm? (2017).

202. *Annual Energy Outlook Retrospective Review: Evaluation of 2013 and Prior Reference Case Projections*. (Energy Information Administration, 2014).

203. *Annual Energy Outlook 2013 with Projections to 2035*. (Energy Information Administration, 2013).

204. *Annual Energy Outlook 2014 with projections to 2040*. (Energy Information Administration, 2014).

205. *Annual Energy Outlook 2015 with projections to 2040*. (Energy Information Administration, 2015).

206. *1982 Annual Energy Outlook: With Projections to 1990*. (Energy Information Administration, 1983).

207. *Annual Energy Outlook 1983: With Projections to 1995*. (Energy Information Administration, 1984).

208. *Annual Energy Outlook 1984: With Projections to 1995*. (Energy Information Administration, 1985).

209. *Annual Energy Outlook 1985: With Projections to 1995*. (Energy Information Administration, 1986).

210. *Annual Energy Outlook 1986: With Projections to 2000*. (Energy Information Administration, 1987).

211. *Annual Energy Outlook 1987: With Projections to 1987*. (Energy Information Administration, 1988).

212. *1989 Annual Energy Outlook: Long-Term Projections*. (Energy Information Administration, 1989).

213. *1990 Annual Energy Outlook: Long-Term Projections*. (Energy Information Administration, 1990).

214. *1991 Annual Energy Outlook: With Projections to 2010*. (Energy Information Administration, 1991).

215. *1992 Annual Energy Outlook: With Projections to 2010*. (Energy Information Administration, 1992).

216. *Annual Energy Outlook 1993: With Projections to 2010*. (Energy Information Administration, 1993).

217. *Annual Energy Outlook Evaluation 2005*. 23 (Energy Information Administration, 2006).

218. *Annual Energy Outlook 1994: With Projections to 2010*. (Energy Information Administration, 1994).

219. *Annual Energy Outlook 1995: With Projections to 2010*. (Energy Information Administration, 1995).

220. *Annual Energy Outlook 1996: With Projections to 2015*. (Energy Information Administration, 1996).

221. *Annual Energy Outlook 1997: With Projections to 2015*. (Energy Information Administration, 1996).

222. *Annual Energy Outlook 1998: With Projections to 2020*. (Energy Information Administration, 1997).

223. *Annual Energy Outlook 1999: With Projections to 2020*. (Energy Information Administration, 1998).

224. *Annual Energy Outlook 2000: With Projections to 2020*. (Energy Information Administration, 1999).

225. *Annual Energy Outlook 2001: With Projections to 2020*. (Energy Information Administration, 2000).

226. *Annual Energy Outlook 2002: With Projections to 2020*. (Energy Information Administration, 2001).

227. *Annual Energy Outlook 2003: With Projections to 2025*. (Energy Information Administration, 2003).

228. *Annual Energy Outlook 2004: With Projections to 2025*. (Energy Information Administration, 2004).

229. *Annual Energy Outlook 2005: With Projections to 2025*. (Energy Information Administration, 2005).

230. *Annual Energy Outlook 2006: With Projections to 2030*. (Energy Information Administration, 2007).

231. *Annual Energy Outlook 2007: With Projections to 2030*. (Energy Information Administration, 2007).

232. *Annual Energy Outlook 2008: With projections to 2030*. (Energy Information Administration, 2008).

233. *An Updated Annual Energy Outlook 2009 Reference Case Reflecting Provisions of the American Recovery and Reinvestment Act and Recent Changes in the Economic Outlook*. (Energy Information Administration, 2009).

234. *Annual Energy Outlook 2010: With projections to 2035*. (Energy Information Administration, 2010).

235. *Annual Energy Outlook 2011: With projections to 2035*. (Energy Information Administration, 2011).

236. *Annual Energy Outlook 2012: With projections to 2035*. (Energy Information Administration, 2012).

237. Boden, T. A., Marland, G. & Andres, R. J. *Global, Regional, and National Fossil-Fuel CO2 Emissions*. (Carbon Dioxide Information Analysis Center, Environmental Sciences Division, Oak Ridge National Laboratory, 2011).

238. Sherwin, E. D., Azevedo, I. M. L. & Meyer, R. M. Characterization of utility programs' enrollment by income and region. in *Consumption, Efficiency & Limits* (European Council for an Energy Efficient Economy, 2017).

239. Ong, S., Campbell, C., Denholm, P., Margolis, R. & Heath, G. *Land-Use Requirements for Solar Power Plants in the United States*. (2013). doi:10.2172/1086349

240. Denholm, P., Hand, M., Jackson, M. & Ong, S. *Land Use Requirements of Modern Wind Power Plants in the United States*. (2009). doi:10.2172/964608