


NHH



WATER. ENERGY. CLIMATE.

ESSAYS IN EMPIRICAL ENVIRONMENTAL ECONOMICS.

Grant R. McDermott

Dissertation submitted to the Department of Economics, Norwegian School of
Economics, in partial fulfillment of the requirements for the degree of PhD in
Economics.

To my family.

Contents

List of Figures	iii
List of Tables	v
Acknowledgements	vi
Introduction	viii
1 Electricity prices, river temperatures and cooling water scarcity	1
1.1 Introduction	2
1.2 Theoretical model and econometric approach	7
1.3 Data	12
1.4 Results	15
1.5 Concluding remarks	22
1.A Comparative statics	29
2 Sceptic priors and climate change mitigation	31
2.1 Introduction	32
2.2 Data	34
2.3 Econometric approach	36

2.3.1	Bayesian regression overview	36
2.3.2	Likelihood function	37
2.3.3	Priors	39
2.4	Results	41
2.4.1	Posterior regression results and TCR distributions	41
2.4.2	Future temperature predictions	43
2.4.3	Welfare implications and the social cost of carbon	46
2.5	Discussion	49
2.6	Concluding remarks	52
2.A	Bayesian multivariate linear regression	59
2.B	Supplementary tables and figures	61
3	Market might. Hydro power.	64
3.1	Introduction	65
3.2	Theoretical motivation	67
3.3	Market characteristics and data	69
3.3.1	The Norwegian electricity market	70
3.3.2	Hydropower reservoirs	71
3.3.3	Electricity flows and transmission constraints	74
3.4	Econometric approach	76
3.5	Results	80
3.6	Concluding remarks	85
3.A	Previous Elspot regimes	90
3.B	Market shares of top three producers	90
3.C	Reservoir volumes: Top three producers in NO2 bidding area	90

List of Figures

1	Global primary energy and electricity supply	x
2	Primary energy consumption, 1965–2014	xi
1.1	Final electricity demand: Germany vs OECD	5
1.2	Water intake and cooling	8
1.3	German mean air temperature	13
1.4	Effect of river temperature on electricity price	14
2.1	TCR densities	43
2.2	Recursive TCR estimates	44
2.3	Model fit and prediction: Noninformative priors	46
2.4	Predicted temperature anomaly by 2100: All priors	47
2.B.1	Posterior parameter densities: Noninformative priors	62
2.B.2	Social cost of carbon	63
3.1	Norwegian hydropower reservoirs	74
3.2	Binding transmission constraints: days	76
3.3	Binding transmission constraints: percent	77
3.4	Merit order of electricity generation	79
3.5	Effect of producer market share on reservoir volume	83
3.A.1	Elsport regime:, 1 Jan 2000 to 13 Dec 2003	91

3.A.2	Elsport regimes: 9 May 2004 to 31 Dec 2013	92
3.B.1	Market shares: 1 Jan 2000 to 13 Dec 2003	93
3.B.2	Market shares: 9 May 2004 to 5 Dec 2011	94
3.C.1	Mean reservoir volumes in NO2 bidding area	95

List of Tables

1.1	Regression results: Electricity price (I)	16
1.2	Regression results: Electricity price (II)	22
2.1	Data sources	35
2.2	Sceptic priors	40
2.3	Regression results: Global mean surface temperature	42
2.4	Covariate vectors for 2100 prediction	45
2.5	Social cost of carbon	48
2.B.1	Transient climate response (I)	61
2.B.2	Transient climate response (II)	61
3.1	Herfindahl index: Norwegian bidding areas	75
3.2	Regression results: Hydropower reservoirs	82

Acknowledgements

The year before starting graduate studies, I travelled from Cairo to Cape Town on a bicycle. The journey itself was a magnificent experience that renewed my sense of wonder in the world around me. But there were also times of hardship, when the road seemed too long and the elements too unyielding. During these moments, my travelling companions and I would succour ourselves with a piece of folk wisdom: “How do you eat an elephant?” (Answer: “One bite at a time.”)

I have often been reminded of that saying whilst working towards the completion of my PhD. Thinking about the whole can be overwhelming. The coursework, the grades, the promising research ideas that flounder, the endless data handling, the revisions and the rewriting. Yet, piece by piece and page by page, here I stand. Ready to take one last bite of this academic morsel. There are some very important people to thank.

I would not have started, let alone finished the PhD without the encouragement and constant support of my supervisor, Øivind Anti Nilsen. Thank you for taking a sincere interest in my research and my personal well-being as your student. Thank you too for being a strong advisor when I needed one, a good collaborator, and a firm friend.

I have benefited considerably from my interactions with other researchers and faculty members at the Department of Economics at NHH. I am hesitant to single out individuals, but Rögnvaldur Hannesson, Linda Nøstbakken, Erik Sørensen, and Gernot Doppelhofer have been especially influential in helping to refine my research ideas and approach to economic questions in general. Similarly, I must extend my deepest thanks to members of the Energy, Natural Resources and Environment research group at NHH, and in particular Gunnar Eskeland, for their friendships and our many stimulating conversations. And, lest I forget the wonderful administrative staff at NHH, Dagny Kristiansen, Irene Grønningsæter, Trude Gudmundset, and Norunn Økland are high on the list of people that helped to make

my years in Bergen such a positive experience.

My colleagues in the PhD programme have played an integral role in getting me to this point. To Sebastian, Kiki, Agnes, Elias, Tunç, Lukàš, Katrine, Morten, Magne, Patrick, Ole-Petter, Simen, and Øivind, among others: It has been a pleasure to work alongside all of you. Sometimes for more hours than is strictly healthy. Your friendships and advice have been a great source of motivation these last four years. Most of all, I am grateful to have been part of a graduate programme that fosters camaraderie above rivalry.

I was lucky to spend the final year of my PhD visiting the beautiful city of Santa Barbara in California. Here too, I was surrounded by genial colleagues who also happened to be first-rate researchers. I count myself even luckier to be continuing my career at UCSB now that the PhD is completed. I wish to express particular gratitude to my hosts — and new employers — Chris Costello and Steve Gaines at the Bren School of Environmental Science and Management.

The final acknowledgements must go to my family.

Mom and Dad, for teaching me the value of education, for encouraging me to ask questions, and for giving me the confidence to pursue my answers. Kathryn, for being a sibling, a friend and an intellectual compatriote.

To Donnacha. During these last few months, you have made a mockery of my sleep schedule, played havoc with my work productivity and shown little regard for my emotional state. Yet you remain the most precious thing in my life. Your mother and I are thankful beyond measure that you decided to join us. Here's hoping that I can be a cool "Dr." and surfer dad in the years to come.

And, lastly, to my beloved wife. Lorrainy, thank you for being a sounding board for all of my ideas, as well as my frustrations, and for reading all of my papers. Even the boring bits. Your unflagging support, kindness and love have truly been an inspiration. I cannot imagine having made it this far without you.

G.M.

Santa Barbara, September 2015

Introduction

The world's water, energy and climate systems are inextricably linked. Water plays a key role at every stage of the energy production process, including the generation of electricity, where it is used for producing hydropower and cooling thermal power plants (DOE, 2014; IEA, 2012). In return, substantial quantities of energy are needed to treat, transport and ultimately dispose of the freshwater resources that sustain our civilisation. At the same time, the composition of global energy technologies will determine the magnitude of future climate change in decades to come. And, of course, climate change threatens to cause major disruptions to freshwater supplies through changes to the earth's hydrological cycle (IPCC, 2014).

The present dissertation is an attempt to better understand the economic nature of these linkages. How do changes to one element propagate through to the others? What are the policy implications of such mutual dependencies and how is policy affected by the beliefs of economic agents? Moreover, how does market structure alter the utilisation of environmental goods and natural resources?

Providing answers to the above questions invites a number of complications. For example, a recurrent problem of valuing environmental goods and services is that these resources are not consistently priced, if at all. A related problem is that many environmental goods are characterised by indeterminate property rights regimes. We must therefore impute their economic value through revealed preference methods (i.e. travel costing and hedonic pricing), or elicit people's stated preferences through surveys (i.e. contingent valuation).¹ A third alternative is to look at the production process directly and see how a change in environmental inputs affects the value of some final good (e.g. McConnell and Bockstael, 2005). This latter approach best describes the methodological framework of Chapters 1 and 3. By ex-

¹There is a rich literature on these two general approaches and their respective strengths and weaknesses. Useful overviews are provided by Hanemann (1994), Pearce (2002), Phaneuf and Smith (2005), Palmquist (2005), and Carson and Hanemann (2005) amongst others. More recently, the *Journal of Economic Perspectives* issued a special symposium on contingent valuation, with contributions by Kling et al. (2012) and Carson (2012), as well as a more critical take by Hausman (2012).

plicitly characterising a firm's production decision in terms of market conditions and environmental inputs, we are able to draw a causal link between these inputs and the resulting economic output. This allows us to measure the economic value of environmental goods even in the absence of detailed market information about the goods themselves. An example from the first chapter is that insufficient access to cooling water can force thermal power plants to throttle production, thereby raising electricity prices in a quantifiable manner.

Unfortunately, it is not always possible to delineate a clear link between environmental amenities and economic outcomes. Sometimes the very status of an environmental good is unclear. This is particularly true from the perspective of individuals who do not get to experience its quality or effects directly. The prime exemplar is climate change. Climate change is by its very definition a global phenomenon that no individual or location experiences directly.² We must therefore infer its key attributes (e.g. global mean surface temperature) by compiling data from a network of satellites or observation stations located at different points around the world. Yet, this opens the door for policy polarization, as stakeholder beliefs are shaped by different priors and conflicting lines of personal experience. Such issues form the backdrop for Chapter 2, which distils support for climate policy as a matter of Bayesian learning. The goal is to determine under which conditions rational agents can reasonably disagree about something as important as climate policy — even when ostensibly faced with the same evidence.

The interweaving of water, energy and climate systems further threatens to impose paradoxical choices upon policy makers over the long term. For example, hydropower and nuclear power are the world's foremost sources of non-fossil fuel energy (excluding biomass). They are certainly the leading forms of low-carbon electric energy at present, and together account for over a quarter of global electricity generation (Figure 1). Yet, both are vulnerable to the effects of climate change due to their acute dependencies on water resources. Nuclear's low thermal efficiencies make it especially susceptible to cooling water scarcity, while retreating glaciers and altered rainfall patterns may serve to significantly reduce hydropower capacities. Countries that attempt to mitigate against climate change by investing heavily in these two energy sources are potentially undermining their future energy stability if other countries do not follow suit. At the very least, it suggests an additional layer of trade-offs for evaluating climate mitigation options versus adaptation pathways.

²Of course, the consequences of climate change may be felt directly at the local scale. Increased storm surge risk due to rising sea levels and so forth.

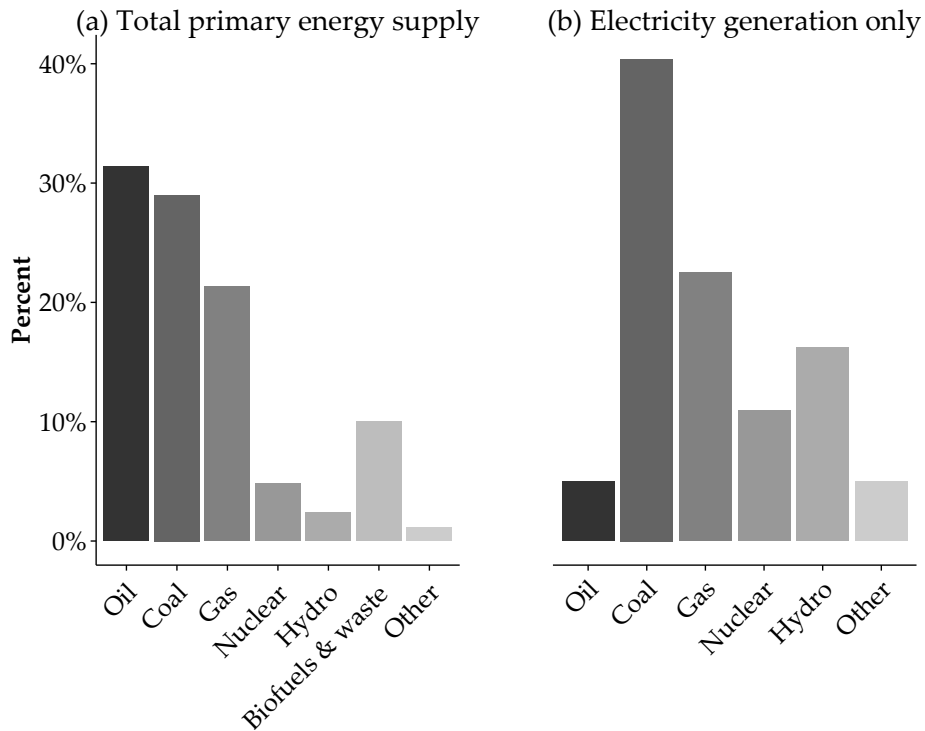


Figure 1: Contribution to global primary energy supply and electricity generation in 2012. Source: IEA (2014b)

Of course, climate change stands to affect the supply of energy in more ways than just freshwater availability. Disruptions to natural weather systems may bring adverse impacts to the wind power industry, for whom wind reliability and consistency are paramount. Conversely, the advent of ice-free summers in the Arctic would lengthen the drilling season there and potentially open up the global market to substantial new quantities of hydrocarbons. This promises to be a growing area of geopolitical intrigue, as Arctic nations both vie for resource supremacy and attempt to balance environmental concerns.³ An accessible overview of these and other vulnerabilities, as well as opportunities, is provided by Schaeffer et al. (2012).

It is furthermore important to contextualise such issues against the broader economic and political milieu. As I write this introduction, U.S. presidential hopefuls are contesting the Democratic and Republican primary nominations of their respective parties. Democratic Party frontrunner Hillary Clinton has just declared her opposition to the contentious Keystone XL oil pipeline, on the grounds that it is a distraction from the overarching challenges of climate change and the need to modernize energy infrastructure (Clinton, 2015). Yet, in a televised debate less than a week earlier, her Republican counterparts vowed to repeal the climate poli-

³C.f. Recent opinion articles in the popular press and elsewhere such as Borgerson (2013), Jacobsen (2015), Reiss (2015), and Myers (2015).




Figure 2: Primary energy consumption by fuel source (excluding biomass), 1965–2014. Source: BP (2015)

cies brought forward by President Obama’s administration, which they claimed would destroy the economy (Woolf, 2015). One candidate, Florida senator Marco Rubio, went on to characterise his position as one of strategic, national self-interest: “America is a lot of things[...] but America is not a planet.” A comment that, in its own way, succinctly highlights the free-rider problem inherent to all common-pool resources.

No doubt energy and the environment will continue to provide fodder for political rhetoric and theatre. What is equally true, however, is that climate policy has the potential to render specific energy technologies uneconomical; witness the ongoing debate over stranded fossil-fuel assets and the larger divestment movement (Van Renssen, 2014; IEA, 2014a; Carrington, 2015a,b). The coal industry has come under particular fire in recent years, with Norway’s sovereign wealth fund providing perhaps the most high profile case of strategic divestment.⁴ Yet, while much of the developed world looks to move beyond coal, no other energy source has made more substantive inroads at a global level since the turn of the millennium (Figure 2). Activism on its own is unlikely to have much of impact next to the growing energy demands of China, India and other emerging economies.

⁴The fund, formally the Norwegian Government Pension Fund Global, has since drawn criticism from environmental groups for offsetting its coal divestments with an increased holding in oil and gas companies (*ibid.*).

It should also be said that the war on coal, at least in North America, has been made politically and economically feasible thanks largely to the shale gas revolution (c.f. Grunwald, 2015). Shale gas — via its reliance on hydraulic fracturing, or “fracking” — comes with its own set of environmental concerns. Apart from methane leakages that may undermine its comparative climate credentials, these too are chiefly related to water. This dissertation will not address shale gas specifically. However, I shall note in passing that the improved thermal efficiencies of gas-fired power plants (relative to coal and nuclear plants) can help to ameliorate at least some of these water concerns. Indeed, numerous studies support the notion that a large-scale coal-to-gas reversion in the electricity sector, underpinned by increasing supplies from hydraulic fracturing, may actually lead to significant water *savings* over the energy production life-cycle (Jenner and Lamadrid, 2013; Laurenzi and Jersey, 2013; Meldrum et al., 2013; Chang et al., 2015).

Similar trade-offs and complementarities can be found all over the world. In Europe, for example, Germany has embarked upon a much-fêted programme of energy transition. This *Energiewende* envisages a sustained and dramatic increase in renewable energy sources over the coming decades. The country has simultaneously committed to phasing out its nuclear power fleet by 2022. As a result, Germany is now investing heavily in expanded grid capacity to its northerly neighbour, Norway, whose hydropower resources offer the most cost-effective way of accommodating the intermittency of large-scale wind and solar (Gullberg et al., 2014; Tønseth, 2014; Adomaitis, 2015). Once again, the broader point is that energy, water resources and climate policy are bound together in a common economic arc.

The individual chapters of this dissertation are best viewed as standalone papers, in the sense that they address distinct economic questions. However, they are unified by the common environmental themes described above. Similarly, while the chapters rely on a different econometric methods — instrumental variables, Bayesian regression, and fixed effects — they share an overarching empirical approach. In each case, a dataset has been constructed *de novo* from a variety of source materials so as to explore a specific aspect of the water-energy-climate nexus. The structure and institutional constraints of the data have then determined which methods are appropriate for answering the question at hand. My graduate studies have taught me that flexibility and resourcefulness are among the most vital components of the modern researcher’s toolkit. It is my hope that the mix of topics and methods in this dissertation will be interpreted in that light.

The following subsections describe the individual chapters in greater detail.

Chapter 1: Electricity prices, river temperatures and cooling water scarcity

Coauthored with Øivind A. Nilsen. Published in Land Economics 2014, 90(1): 131–148.

Chapter 1 establishes an empirical relationship between electricity prices and freshwater availability. While this relationship is self-evident in the case of hydropower production, our present context is the less obvious case of thermal-based power production.

We begin by defining a stylised theoretical model to capture the essence of the problem. The intuition is as follows. Thermal power plants produce electricity more efficiently in the presence of an external coolant, such as water taken from a nearby river. The extent of this efficiency boost depends on the temperature of the intake water and how much is being withdrawn. Plants can compensate for warmer intake water by withdrawing more of it. However, we assume that water becomes costly to withdraw as it gets scarcer. At the same time, a regulator sets a temperature cap on discharging used cooling water back into the environment in order to safeguard against thermal pollution. When this constraint is binding, producers must choose between reduced water intake, thus compromising plant efficiency, or incurring regulation costs. Either choice will result in an inward shift of the supply curve and a higher electricity price.

The primary challenge of moving from our theoretical framework to an empirical setting is that electricity prices and quantities are jointly determined by demand and supply forces. We use an instrumental variable approach to overcome this simultaneity problem. We are thus able to control for demand-side effects and isolate the impact of water-related shocks to electricity supply. Applying our econometric model on German data, we find that the electricity price increases in response to both lower river levels and higher river temperatures. The empirical results accord closely with our theoretical predictions and reflect the thermal power industry's vulnerability to water scarcity in an absolute and relative (i.e. regulatory) sense. One implication of our findings is that climate change will cause higher energy prices, not only through increased demand for cooling, but also through water-related supply shocks.

Chapter 2: Sceptic priors and climate change mitigation

Despite an overwhelming scientific consensus, many people are openly sceptical about anthropogenic climate change. The reasons for this persistent scepticism is a matter of some debate among researchers (Clark et al., 2013; Kahan et al., 2011, 2012). What seems beyond dispute, however, is that public scepticism has undermined attempts to enact meaningful climate policies on a global scale.

Chapter 2 begins with a simple question: How much evidence is needed to convince climate sceptics that they are wrong? I characterise the problem within a Bayesian framework, where a group of hypothetical sceptics are left to update their prior beliefs in conjunction with historical evidence for climate change. The primary finding is that instrumental climate data from the last century and a half overwhelms all but the most extreme priors. In Bayesian language, the posterior beliefs of the different sceptics gravitate strongly toward the mainstream consensus — i.e. climate change is real and caused by humans. I show further that the updated sceptic beliefs are consistent with a social cost of carbon that is significantly greater than zero. These results suggest that it should be possible to obtain broad support for a carbon price, provided that people simply have access to the available evidence.

The idealised structure of the Bayesian model in this chapter necessitates some caution in how we interpret its relevance for our current policy impasse. However, I argue that two points are particularly salient. First, the results can help to explain why climate change remains such a polarizing issue. As all intermediate positions are subsumed into the mainstream, only the most extreme sceptics remain wedded to their priors. No amount of new information is going to convince this group that it needs to reconsider its position on climate change. Second, the Bayesian framework is able to bridge the divide between competing theories of climate scepticism as a societal phenomenon. Using the insights of Jaynes (2003), I show that incorporating priors about source credibility allows for the same evidence to cause *diverging* beliefs among partisans. The policy implication is that additional evidence for climate change is unlikely to change beliefs unless sceptics experience the effects of climate change directly, or this evidence comes from sources that they regard as unambiguously trustworthy. Merely communicating additional evidence through existing channels and organisations — such as the Intergovernmental Panel on Climate Change — is unlikely to have much of an impact.

Chapter 3: Market might. Hydro power.

Chapter 3 concerns the most explicit link between water and energy: hydropower production. In particular, how do firm incentives for exploiting water resources vary along with changes in market concentration.

The distinctive characteristics of hydropower — seasonality, storability, negligible variable costs — has interesting implications for the way that water resources are utilised. Theory tells us that dominant hydropower firms will withhold production during periods with relatively inelastic demand (Førsund, 2015). This will allow them to recoup higher profits by driving up the price when consumers are the least responsive to such changes. Testing this in an empirical setting has proven difficult, however, because of data limitations. The contribution of this chapter is to bridge that gap by introducing a uniquely detailed dataset of Norwegian hydropower reservoirs, firms and electricity flows. Exogenous changes to bidding area divisions and transmission constraints allow me to cleanly identify variations in reservoir volumes arising from differences in local market share.

Consistent with theory, I find that market power causes firms to strategically reallocate their water resources across periods. An increase in regional market share is associated with higher reservoir volumes during summer months and lower reservoir volumes during winter months. Drawing on the insights of previous studies (Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008), I argue that this result is driven by the seasonal differences in demand elasticities. Electricity in the Norwegian winter is predominantly used for heating, which allows for numerous substitutes and, hence, a relatively higher elasticity of demand. In contrast, electricity consumption during summer is primarily given over to technical end-uses, which do not allow for easy substitution.

The size of this reallocation effect is modest next to other factors governing reservoir management, such as annual inflows from snow-melt. Yet, it may still cause the production profile of hydropower firms to diverge in meaningful ways if the differences in market share are large enough, and particularly when regional transmission constraints are binding.

References

- Adomaitis, Nerijus (2015). "Norwegian-German power cable project delayed until 2019," *Reuters*, 10 February 2015. Available: <http://www.reuters.com/article/2015/02/10/norway-germany-cable-idUKL5N0VK2PY20150210>.
- Borgerson, Scott (2013). "The Coming Arctic Boom," *Foreign Affairs*, 92(4): 76–89.
- BP (2015). *BP Statistical Review of World Energy*, British Petroleum, 64th edition.
- Bye, Torstein and Petter Vegard Hansen (2008). "How do Spot prices affect aggregate electricity demand?" Statistics Norway Discussion Papers No. 527. Available: <http://www.ssb.no/a/publikasjoner/pdf/DP/dp527.pdf>.
- Carrington, Damian (2015a). "Guardian Media Group to divest its £800m fund from fossil fuels," *The Guardian*, 1 April 2015. Available: <http://www.theguardian.com/environment/2015/apr/01/guardian-media-group-to-divest-its-800m-fund-from-fossil-fuels>.
- Carrington, Damian (2015b). "Norway's sovereign wealth fund drops over 50 coal companies," *The Guardian*, 16 March 2015. Available: <http://www.theguardian.com/environment/2015/mar/16/norways-sovereign-wealth-fund-drops-over-50-coal-companies>.
- Carson, Richard T. (2012). "Contingent valuation: A practical alternative when prices aren't available," *Journal of Economic Perspectives*, 26(4): 27–42, symposium on contingent valuation.
- Carson, Richard T. and W. Michael Hanemann (2005). "Contingent Valuation," in Karl-Göran Mäler and Jeffrey R. Vincent (eds.), "Valuing Environmental Changes," volume 2 of *Handbook of Environmental Economics*, chapter 17, Elsevier, pp. 821–936.
- Chang, Yuan, Runze Huang, Robert J. Ries, and Eric Masanet (2015). "Life-cycle comparison of greenhouse gas emissions and water consumption for coal and shale gas fired power generation in China," *Energy*, 86: 335–343.

- Clark, Dav, Michael Andrew Ranney, and Jacqueline Felipe (2013). "Knowledge Helps: Mechanistic Information and Numeric Evidence as Cognitive Levers to Overcome Stasis and Build Public Consensus on Climate Change," in N. Sebanz & I. Wachsmuth M. Knauff, M. Pauen (ed.), "Proceedings of 35th Annual Meeting of the Cognitive Science Society," pp. 2070– 2075.
- Clinton, Hillary (2015). "Why I oppose Keystone XL," *Medium*, 23 September 2015. Available: <https://medium.com/@HillaryClinton/why-i-oppose-keystone-xl-8698230df2a8>.
- DOE (2014). "The water-energy nexus: Challenges and opportunities," Report, U.S. Department of Energy, Washington D.C.
- Førsund, Finn R. (2015). *Hydropower Economics*, International Series in Operations Research and Management Science, Springer, second edition.
- Grunwald, Michael (2015). "Inside the war on coal," *POLITICO*, 27 May 2015. Available: <http://www.politico.com/agenda/story/2015/05/inside-war-on-coal-000002>.
- Gullberg, Anne Therese, Dörte Ohlhorst, and Miranda Schreurs (2014). "Towards a low carbon energy future — Renewable energy cooperation between Germany and Norway," *Renewable Energy*, 68: 216–222.
- Hanemann, W. Michael (1994). "Valuing the Environment Through Contingent Valuation," *Journal of Economic Perspectives*, 8(4): 19–43.
- Hansen, Petter Vegard (2004). "Regional electricity spot price responses in Norway," Statistics Norway Discussion Papers No. 2004/13. Available: http://www.ssb.no/a/english/publikasjoner/pdf/doc_200413_en/doc_200413_en.pdf.
- Hausman, Jerry (2012). "Contingent Valuation: From Dubious to Hopeless," *Journal of Economic Perspectives*, 26(4): 43–56, symposium on contingent valuation.
- IEA (2012). *World Energy Outlook 2012*, International Energy Agency, OECD Publishing.
- IEA (2014a). *World Energy Investment Outlook 2014*, International Energy Agency, OECD Publishing, special Report.
- IEA (2014b). *World Energy Outlook 2014*, International Energy Agency, OECD Publishing.

- IPCC (2014). *Climate Change 2014: Impacts, Adaptation, and Vulnerability. Contribution of Working Group II to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, [Field, C.B., V.R. Barros, D.J. Dokken, K.J. Mach, M.D. Mastrandrea, T. E. Bilir, M. Chatterjee, K.L. Ebi, Y.O. Estrada, R.C. Genova, B. Girma, E.S. Kissel, A.N. Levy, S. MacCracken, P.R. Mastrandrea, and L.L. White (eds.)], Cambridge University Press, 1st edition.
- Jacobsen, Stine (2015). "Norwegian government forced to reassess Arctic oil drilling boundaries," *Reuters*, 11 June 2015. Available: <http://www.reuters.com/article/2015/06/11/us-norway-oil-arctic-idUKKBN00R1GJ20150611>.
- Jaynes, Edwin T. (2003). *Probability Theory: The Logic of Science*, Cambridge University Press.
- Jenner, Steffen and Alberto J. Lamadrid (2013). "Shale gas vs. coal: Policy implications from environmental impact comparisons of shale gas, conventional gas, and coal on air, water, and land in the united states," *Energy Policy*, 53: 442–453.
- Johnsen, Tor Arnt (2001). "Demand, generation and price in the Norwegian market for electric power," *Energy Economics*, 23(3): 227–251.
- Kahan, Dan M., Hank Jenkins-Smith, and Donald Braman (2011). "Cultural cognition of scientific consensus," *Journal of Risk Research*, 14(2): 147–174.
- Kahan, Dan M., Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel (2012). "The polarizing impact of science literacy and numeracy on perceived climate change risks," *Nature Climate Change*, 2: 732–735.
- Kling, Catherine L., Daniel J. Phaneuf, and Jinhua Zhao (2012). "From Exxon to BP: Has Some Number Become Better Than No Number?" *Journal of Economic Perspectives*, 26(4): 3–26, symposium on contingent valuation.
- Laurenzi, Ian J. and Gilbert R. Jersey (2013). "Life cycle greenhouse gas emissions and freshwater consumption of Marcellus shale gas," *Environmental science & technology*, 47(9): 4896–4903.
- McConnell, Kenneth E. and Nancy E. Bockstael (2005). "Valuing the Environment as a Factor of Production," in Karl-Göran Mäler and Jeffrey R. Vincent (eds.), "Valuing Environmental Changes," volume 2 of *Handbook of Environmental Economics*, chapter 14, Elsevier, pp. 621–669.

- Meldrum, James, Syndi Nettles-Anderson, Garvin Heath, and Jordan Macknick (2013). "Life cycle water use for electricity generation: a review and harmonization of literature estimates," *Environmental Research Letters*, 8(1): 1–18.
- Myers, Stephen Lee (2015). "U.S. Is Playing Catch-Up With Russia in Scramble for the Arctic," *The New York Times*, 29 August 2015. Available: <http://www.nytimes.com/2015/08/30/world/united-states-russia-arctic-exploration.html>.
- Palmquist, Raymond B. (2005). "Property Value Models," in Karl-Göran Mäler and Jeffrey R. Vincent (eds.), "Valuing Environmental Changes," volume 2 of *Handbook of Environmental Economics*, chapter 16, Elsevier, pp. 763–819.
- Pearce, David (2002). "An Intellectual History of Environmental Economics," *Annual Review of Energy and the Environment*, 27(1): 57–81.
- Phaneuf, Daniel J. and V. Kerry Smith (2005). "Recreation Demand Models," in Karl-Göran Mäler and Jeffrey R. Vincent (eds.), "Valuing Environmental Changes," volume 2 of *Handbook of Environmental Economics*, chapter 15, Elsevier, pp. 671–761.
- Reiss, Bob (2015). "In the Race to Control the Arctic, the U.S. Lags Behind," *Newsweek*, 6 July 2015. Available: <http://www.newsweek.com/2015/07/17/cold-cold-war-349973.html>.
- Schaeffer, Roberto, Alexandre Salem Szklo, André Frossard Pereira de Lucena, Bruno Soares Moreira Cesar Borba, Larissa Pinheiro Pupo Nogueira, Fernanda Pereira Fleming, Alberto Troccoli, Mike Harrison, and Mohammed Sadeck Boulahya (2012). "Energy sector vulnerability to climate change: a review," *Energy*, 38(1): 1–12.
- Tønseth, Svein (2014). "Norway is Europe's cheapest "battery"," *Science Nordic*, 22 December 2014. Available: <http://sciencenordic.com/norway-europe%E2%80%99s-cheapest-%E2%80%9Cbattery%E2%80%9D>.
- Van Rensen, Sonja (2014). "Investors take charge of climate policy," *Nature Climate Change*, 4(4): 241–242.
- Woolf, Nicky (2015). "'America is not a planet': Republicans resist climate change action at debate," *The Guardian*, 16 September 2015. Available: <http://www.theguardian.com/us-news/2015/sep/17/marco-rubio-chris-christie-climate-change-republican-debate>.

Chapter 1

Electricity prices, river temperatures and cooling water scarcity

Abstract

Thermal-based power stations rely on water for cooling purposes. These water sources may be subject to incidents of scarcity, environmental regulations and competing economic concerns. This paper analyses the impact of water scarcity and increased river temperatures on German electricity prices from 2002 to 2009. Having controlled for demand-side effects, we find that electricity prices are significantly affected by both a change in river temperatures and the relative abundance of river water. An implication is that future climate change will affect electricity prices not only through changes in demand, but also through increased water temperatures and scarcity.

JEL Codes: Q25, Q41, Q5, C3

Keywords: Thermal-based power, water scarcity, water-energy nexus.

1.1 Introduction

Thermal-based power facilities, such as nuclear and coal-fired plants, are critically dependent on water for cooling purposes. This enables them to maintain high production efficiencies, but also means that they require access to vast quantities of water. To give an indication of scale, the thermal power industry accounts for roughly 40 percent of all freshwater withdrawals in the United States – a figure that places it alongside the agricultural sector (DOE, 2006). Unlike agriculture, the majority of these withdrawals are returned to their natural source. However, discharging used cooling water back into the environment presents problems of its own. Excess thermal energy absorbed during the heat exchange will naturally cause the water to warm up prior to being released back into the river or lake from which it was taken. This raises the ambient temperature of the water source itself and can ultimately bring detrimental effects to the aquatic ecosystem. Water temperatures at or above the mid-20s degree Celsius ($^{\circ}\text{C}$) mark are considered particularly dangerous to aqueous plants and certain fish species, since this leads to reduced oxygen levels and raised concentrations of ammonia (Langford, 1990). As a result, many countries have enacted environmental regulations on the maximum allowable temperature of discharge water from power stations, otherwise known as “thermal pollution”.¹

Within this context, an emerging literature has developed that seeks to analyse how thermal-based power production might be constrained by access to water resources. Some studies largely abstract from wider climate phenomena and focus primarily on what growing energy demand means for water consumption (for instance, see Feeley et al., 2008; NETL, 2009a,b,c). Others have specifically tried to incorporate climate change into their analysis and even suggest adaptive strategies available to the thermal power industry in coping with a warming world.² The Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC, 2007) synthesised several such studies in suggesting that future energy generation will be vulnerable to higher temperatures and a reduced availability of cooling water

¹The vulnerability to water scarcity, as well as problems related to thermal pollution, varies according to fuel type and cooling technology. For example, the low thermal efficiencies of nuclear plants make them particularly susceptible to water-related issues — see DOE (2006).

²The links between thermal-based power production, water scarcity and climate have also received growing attention in popular media formats. This includes news stories of European power plants shutting down during heat waves of the last decade (Gentleman, 2003; Godoy, 2006; Pagnamenta, 2009) and similar problems in the US (AP, 2008; Sohn, 2011; Eaton, 2012), as well as implications of the nuclear power sector’s dependency on water and recurring incidents of drought (Kanter, 2007; Dell’Amore, 2010). The linkage between power supply and water needs has also received increased attention in the wake of recent events at the Fukushima Daiichi nuclear plant in Japan (Chellaney, 2011).

for thermal power stations. Citing effects from the 2003 European heat wave as a precautionary example, the report declared: “[E]lectricity production was undermined by the facts that the temperature of rivers rose, reducing the cooling efficiency of thermal power plants (conventional and nuclear) and that flows of rivers were diminished” (*ibid.*, p. 367). Though typically local in focus, the list of individual studies covering equivalent issues includes Hurd and Harrod (2001); Arnell et al. (2005); Maulbetsch and DiFilippo (2006); Kirshen et al. (2008); Sovacool (2009); Sovacool and Sovacool (2009).

In terms of predictive capability, Koch and Vögele (2009) offer an adaptable framework that lends itself to scenario analysis. They construct an integrated water management model that can be used to simulate the interconnected effects of changing energy needs and water availability on thermal-based power production. This model is then applied to individual plants under various hypothetical climate and economic scenarios. However, whilst Koch and Vögele go on to comment on some broader socio-economic outcomes, they acknowledge that their simulations do not account for the fact that “water shortages affecting large regions[...] could have an impact on energy prices”. Förster and Lilliestam (2010) adopt a somewhat narrower approach by simulating the effects of climate change on a single large (hypothetical) nuclear plant in Central Europe that is reliant on once-through cooling technology. Their results indicate that annual load losses could be as high as 11.8 percent, with annual plant losses upwards of €100 million for the worst case scenarios. In turn, Van Vliet et al. (2012) provide a more general overview of how the power sector will be affected by climate change. Their simulations show that the average summer capacity of power plants in Europe and the United States will be reduced by 6.3 percent to 19 percent, or 4.4 percent to 16 percent, depending on cooling system type and climate scenario for 2031-2060. They too, however, make no explicit attempts to model for an effect on electricity prices.

In contrast to the simulation-based studies above, Linnerud et al. (2011) tread an empirical line. Using European data to analyse the impact that climate change may have on the nuclear industry, they find that an average temperature rise of 1°C reduces the supply of nuclear power by roughly half a percent. Finally, Kopytko and Perkins (2011) provide a discursive overview of the inherent vulnerabilities that nuclear power will be exposed to as a result of climate change. Among other things, they specifically highlight cooling water scarcity as an impediment to future investments in inland nuclear plants.

The purpose of this paper is to determine how electricity prices are affected by

access to cooling water. Despite increasing concern over the large volumes of water use in thermal-based power production, we are unaware of any studies that have analysed this question in an empirical setting.³ Our aim is to quantify what these effects are and, in so doing, provide a fresh look at how inadequate freshwater supplies can yield direct economic costs.

We use German data in this study as this has several advantages. First, at a total consumption level of around 550 TWh in 2010, the German power market is the largest in Europe and is characterised by a diverse mix of input fuels (Kristiansen, 2011). The market is also supplied by a large number of power plants scattered around the country. That said, only four companies are responsible for approximately 80 percent of total production capacity (Müsgens, 2006; Möst and Genoese, 2009). During the 2002–2009 review period of this paper, the country derived around 60 percent of its electricity needs from fossil fuels (mostly coal), 23 percent from nuclear, and the remainder from a combination of renewables.⁴ According to the International Energy Agency (IEA, 2011a) and as shown in Figure 1.1, these numbers very closely parallel those of the OECD region as a whole.⁵ Germany is therefore a very reasonable “representative agent”, from which one can make wider inferences about the impact that water scarcity has on thermal power production and, consequently, electricity prices.

The second factor in support of Germany as a case study is the availability of a wide series of relevant data – including wholesale electricity data and hydrological records — which makes it an amenable choice for conducting empirical analysis. The German power market was fully liberalised in 1998 and by 2001 the major electricity trading platforms had merged to form a single entity, the Leipzig-based European Energy Exchange (EEX). With regard to institutional settings, market participants on the EEX are able to trade a variety of products corresponding to different time horizons and derivative positions. A range of standardised products are also traded in the form of bilateral over-the-counter (OTC) agreements between direct counterparties, often concluded via brokerage firms. Our main focus in this pa-

³Boogert and Dupont (2005) suggest that water temperatures resulted in increased Dutch electricity prices (via a supply-side shock) during the 2003 heat wave. However, their concise paper does not offer any empirics beyond some descriptive statistics.

⁴The role of nuclear power in Germany has been highly contested over the last several decades, with renewed public interest in the wake of the stricken Fukushima Daiichi plant in Japan. Following a period of political flip-flopping on the issue, in 2011 the German government committed to phasing out nuclear power by 2022. We do not explicitly consider these developments, since they took place after our period of study.

⁵The relative role that thermal-based power plays is comparable even at a global level – where the contributions of fossil fuels and nuclear to total power generation in 2009 were 67 percent and 13 percent, respectively (IEA, 2011b).

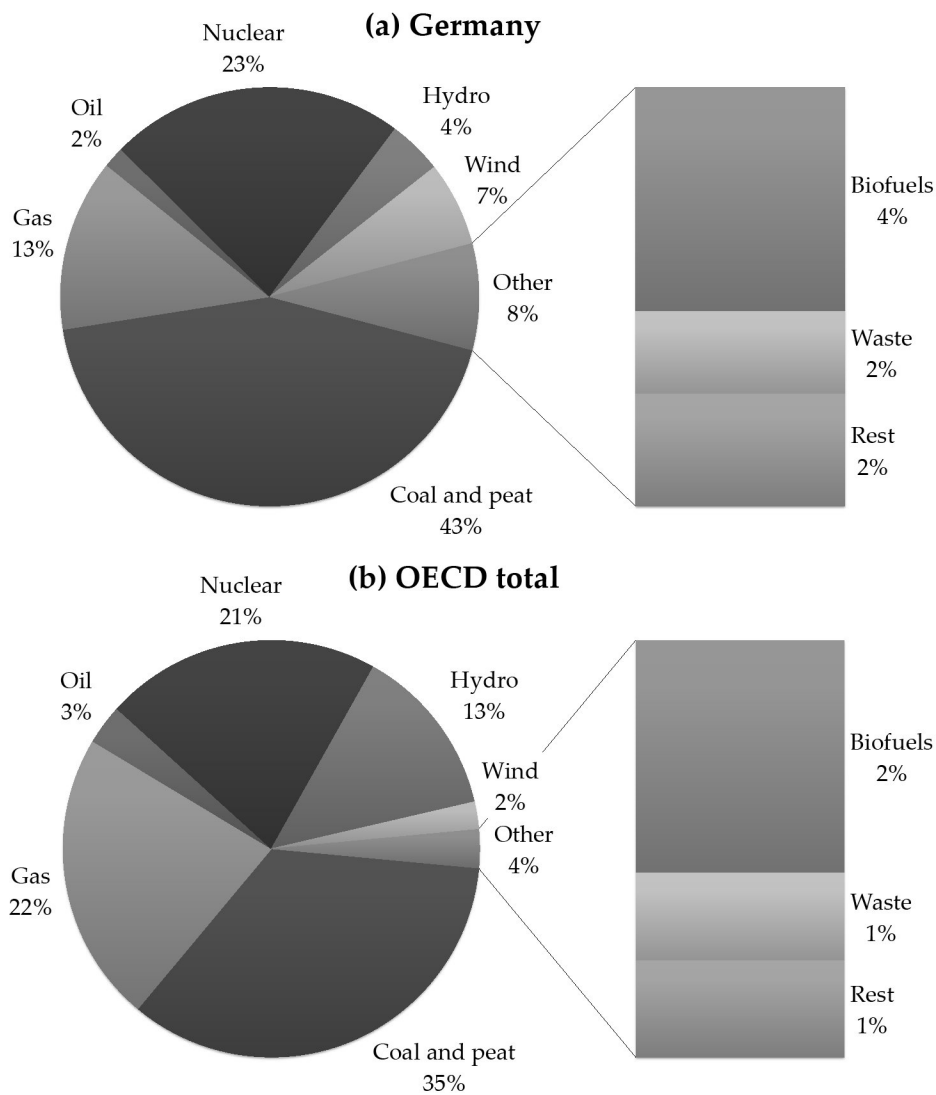


Figure 1.1: Contribution to final electricity demand in 2009. Source: IEA (2011a)

per is the day-ahead EEX spot auction, during which hourly electricity contracts and block contracts can be traded until midnight of the previous day.⁶ The EEX spot market accounts for approximately 30 percent of German electricity demand, and we would expect certainly marginal changes in river levels or temperatures to be reflected in these prices. Furthermore, the spot price also acts as a reference point for all upstream market participants, regardless of where and what they are trading. If not, there would be costless arbitrage opportunities (Viehmann, 2011),

⁶Instead of using average daily prices, one might be tempted to argue that higher-frequency data (e.g. hourly) would allow for a higher degree of analytical precision. There are several reasons why we do not follow this approach. First, the supply of produced electricity is not particularly flexible for the baseload plants that we focus on in this study. Second, environmental authorities are highly unlikely to respond to changes on an hourly basis, but rather issue daily orders to power plant operators. Third, daily data is consistent with our other data (e.g. hydrological).

Kristiansen (2011). The day-ahead spot market should thus not only reflect the underlying long-term demand and supply conditions of the power sector, but also respond to short-term shocks. This includes power plants being entered into constrained production due restrictions on their intake of cooling water.

The third and perhaps most important reason for using German data, is that the country's electricity sector has proven vulnerable to incidences of water scarcity and compromised water quality in the past. This has been especially true during very hot periods such as the European heat waves in 2003 and 2006, when Germany joined the likes of France and Spain in suffering from significantly reduced production capacities. A proximate cause of this outcome was the fact that river temperatures began to exceed the regulatory threshold imposed on thermal water pollution. Faced with a power system already straining under the pressure of unusually high demand for air-conditioning, German federal authorities initially provided emergency dispensation for power stations to flout environmental laws. However, they were eventually forced to uphold the standard restrictions on discharging hot water into the environment in order to protect river fauna and flora. At least 15 thermal plants had to be shut down or entered into constrained production because of water-related issues during the summer months of 2003 (Müller et al., 2007, 2008). Similarly, at least 12 plants were throttled during the 2006 heat wave so as to limit the discharge of thermal waste water into rivers.

It is against this backdrop that we can preview our key findings. Having successfully controlled for various demand-side factors, our empirical results indicate that the electricity price is significantly affected by both a change in river temperatures and the relative abundance of river water. Falling river levels are generally associated with a higher electricity price, while prices will also be driven higher once average river temperatures breach regulatory thresholds. For instance, we estimate that an increase in average river temperatures from 25 °C to 26 °C will bring about a near four percent increase in electricity prices over the course of a week.

The remainder of the paper is structured as follows. Section 2 presents the theoretical framework and discusses our econometric strategy. Section 3 describes the data. Section 4 presents the empirical findings together with a discussion on alternate specifications and aggregation issues. Concluding remarks are provided in Section 5.

1.2 Theoretical model and econometric approach

Our model aims to incorporate the thermal production process – albeit in a highly stylised manner – to account for the effect that increased river temperatures have on electricity output. In brief, thermal energy can be converted into electrical energy more efficiently in the presence of an external coolant, such as water. This result is famously rendered by Carnot’s theorem, which holds that the efficiency of a thermal-based engine is directly proportional to the temperature differential between its high and low temperature reservoirs (e.g. Langford, 1990). This allows excess heat from the production cycle to be transferred to the coolant and subsequently disposed of. The cooling technology that thermal-based power plants use may be divided into two broad categories: once-through or closed-circuit systems. The former requires that far greater volumes of water be withdrawn from natural sources, while the latter “consumes” more in the form of evaporation. That said, we abstract from such differences and instead focus on the core principle that cooling is essential to maintaining efficiency levels in any thermal-based power plant – irrespective of whether it is fuelled by coal, gas or nuclear energy. We therefore assume a simple production technology of

$$Q = A(T_{EW} - T) \cdot W. \quad (1.1)$$

In other words, the production of electricity Q is contingent on the difference in temperature of the discharge water at the outlet point, T_{EW} , and the cooling water at the intake point, T . Production will increase as this temperature difference increases, i.e. $A' > 0$.⁷ This formulation effectively takes the inner workings of the thermal engine as exogenous and instead focuses on the fact that surplus heat energy is transferred to the cooling water via a heat exchange. A higher temperature difference between the discharge water and its original source therefore implies a higher thermal efficiency (i.e. conversion of thermal energy into electrical energy). Importantly, the model also captures the possibility that thermal-based plants can use more cooling water, W , to compensate for a low temperature differential. Figure 1.2 provides a stylised depiction of the production cycle used in our model.

The production of electricity by thermal-based power stations is subject to the fol-

⁷An underlying assumption is that $T_{EW} \geq T$. In other words, there is a cooling effect due to the heat exchange that takes place in the plant condenser. The specification that we have used here is thus also indicative of the fact that the cooling effect becomes increasingly negligible as the temperature difference falls.




Figure 1.2: Water intake and cooling for a power plant

lowing constraint,

$$\frac{W}{S} \cdot T_{EW} + \frac{S - W}{S} \cdot T \leq \bar{T}. \quad (1.2)$$

This reflects the fact that environmental authorities set a cap, \bar{T} , on the temperature of the downstream river volume, S , which occurs as a result of the mixing between discharged cooling water, W , and the river water not used for cooling, $(S - W)$. Thus, $\frac{W}{S}$ is the share of total river water used for cooling. The constraint implies that rather than undergoing a complete shutdown, the plant has the option of reducing the flow of discharge relative to the volume of downstream mixing water when the temperature of each unit of discharged water, T_{EW} , is relatively hot. However, as the temperature of the river water itself approaches the regulatory limit (e.g. during very hot summer months), the plant has little room for manoeuvring and will likely have to decrease output.⁸

The strategic decision variable available to power plants in our theoretical frame-

⁸Of course, environmental authorities will also typically impose limits on the temperature of the discharged water itself – let us say T_{EW} – and/or on the temperature differential between river water at the intake point and the discharge (c.f. Mimler et al., 2009). In the interests of parsimony, however, we ignore these additional limits in our model. Indeed, one could argue that including a constraint, T , on the temperature of the downstream river volume, S , already serves to capture these effects indirectly.

work is quantity. It should be noted that electricity is a homogeneous product that cannot be stored, and demand must be perfectly balanced by supply at all times. Given the institutional settings of the German power market, our model allows for potential market power, but is also generalizable to a situation where the representative power plant behaves as a price-taker. We therefore model the profit, π , for thermal-based plants as

$$\pi = p(Q + F) \cdot Q - c(Q) - p_W(RL) \cdot W, \quad (1.3)$$

where $p(\cdot)$ denotes the inverse demand function and total electricity demand is the sum of power produced by the analysed plants, Q , together with electricity imports and the other sources that aren't dependent on cooling water (e.g. wind power), F . The cost function, $c(Q)$, captures the marginal costs associated with the production of additional quantities of electricity. In addition to these standard production costs, the latter part of the expression, $p_W(RL) \cdot W$, reflects the fact that there are costs associated with drawing cooling water, W , from its source. These are assumed to be a function of the river level, RL , such that $p_W' < 0$.

By substituting in the technology function for the water parameter, W , we have the following profit maximisation problem:

$$\max_Q p(Q + F) \cdot Q - c(Q) - p_W(RL) \cdot \frac{Q}{A(T_{EW} - T)}, \quad (1.4)$$

subject to

$$\frac{Q}{S \cdot A(T_{EW} - T)} \cdot T_{EW} + \left(1 - \frac{Q}{S \cdot A(T_{EW} - T)}\right) \cdot T \leq \bar{T}.$$

Taking the first-order condition with respect to Q yields

$$p^* \cdot \left(1 - \frac{1}{\epsilon}\right) = \frac{\partial c(Q^*)}{\partial Q} + \frac{p_w(RL)}{A(T_{EW} - T)} + \lambda \cdot \left[\frac{T_{EW} - T}{S \cdot A(T_{EW} - T)}\right], \quad (1.5)$$

where Q^* is the optimal level of produced energy (with corresponding optimal price, p^*), λ is the shadow price of the constraint, and ϵ denotes the price elasticity where we have integrated out the demand effects for electricity provided by

the other sources.⁹

Given that $p_W' < 0$, a lower river level will effectively have the same impact as an increase in production costs. Falling river levels will consequently reduce the supply of electricity and ultimately bring about an increase in price, i.e. $\frac{\partial p^*}{\partial RL} > 0$. The effect of river temperatures is slightly more complex as it will impact price through various channels. First, an increase in T will negatively impact the thermal efficiency of a plant. This effect could be mitigated by withdrawing more cooling water, although this will bring with it its own costs, since profit is a function of $p_w(RL) \cdot W$. Moreover, when the constraint is binding ($\lambda > 0$), the only way that a power plant can respond to an increased river temperature will be to reduce W and therefore lower the production. In either case, an increase in T will reduce the quantity of electricity by shifting the supply curve to the left. This in turn will lead to an increase in price, i.e. $\frac{\partial p^*}{\partial T} > 0$.¹⁰

It is well known that electricity prices and quantities are jointly determined in the market-clearing process. This simultaneity needs to be accounted for in the empirical estimation of our theoretical model. We therefore begin by defining the following supply equation:

$$\ln P_t = \beta_0 + \beta_1 \ln Q_t + \beta_2 \ln RL_t + \beta_3 \ln RT_t + \beta_4 \ln F_t + \beta_T' \mathbf{T}_t + v_t, \quad (1.6)$$

where P is the daily clearing price for electricity, Q is the daily electricity consumption, RL is the average river level, RT is the average river temperature, F is fuel (input) costs, and \mathbf{T} is a set of seasonal and trend variables.

The regressors of greatest interest to this study are river levels, RL , and river temperatures, RT . These two coefficients should reflect how electricity supply is constrained by diminished cooling water availability, due to either relative scarcity (i.e. falling river levels) or regulatory concerns (i.e. river temperatures breaching environmentally sensitive thresholds).

The aforementioned simultaneity of supply and demand means that simply regressing electricity prices on quantities using ordinary least squares (OLS) will generate inconsistent parameter estimates. We resolve this issue by adopting an

⁹This means that the model encompasses settings where the representative plant is a price-taker (i.e. $\epsilon \rightarrow \infty$), or where it exercises market power. Of course, a plant's ability to react to changes in demand or marginal costs will also depend on what fuel type they are. For instance, nuclear power plants are built for providing a constant baseload, while gas-fired plants are more flexible.

¹⁰Please see Appendix 1.A for a more detailed derivation of the comparative statics.

instrumental variable (IV) approach within a two-stage least squares (2SLS) framework. Instrumenting for Q should allow us to properly identify the causal effect that changing volumes have on electricity prices. Our set of instruments begins with a concept widely used in energy modelling, namely degree days (see, for instance, Halvorsen, 1975; Quayle and Diaz, 1980). Heating degree days (HDD) and cooling degree days (CDD) are complementary terms that capture the nonlinear effect that changing temperatures have on electricity demand. They do this by measuring the extent to which air temperatures fall outside a given comfort zone, which we define here as 18 °C to 22 °C.¹¹ The HDD variable measures how far the temperature drops below 18 °C on any given “cold” day (thus requiring heating), while CDD measures how far the temperature exceeds 22 °C on any given “hot” day (thus requiring cooling).¹² Our set of instruments is completed by a dummy variable that corresponds to non-working days, NWD . This variable is included to reflect the fact that electricity demand typically falls on weekends and public holidays due to reduced industrial activity. The reduced form regression equation is thus

$$\ln P_t = \pi_0 + \pi_1 \ln HDD_t + \pi_2 \ln CDD_t + \pi_3 \ln NWD_t + \pi_4 \ln RL_t + \pi_5 RT_t + \pi_T T_t + u_t. \quad (1.7)$$

The critical assumption for our chosen instrumental variables – CDD, HDD and NWD – is that they pass the exclusion restrictions requirement. That is, they affect prices only indirectly through changes in demand. The temperature discomfort associated with CDD and HDD is thus assumed to cause an increase in demand but have no bearing on direct supply. Given that we control for changing river levels and river temperatures separately, this seems to be a valid assumption. Similarly, it is extremely unlikely that the supply of electricity will be materially constrained by the fact that it is a weekday or public holiday. The main factors of production are not affected by the day of the week, for instance, and power companies will be able to operate at normal capacity irrespective of such considerations. To be sure of the exogeneity of the instruments, however, we conduct a Sargan-Hansen overidentification test to confirm our economic reasoning. Furthermore, the standard Durbin-Wu-Hausman specification test is used to test for endogeneity. The

¹¹This is a fairly standard range in the literature. Some studies (e.g. Bessec and Fouquau, 2008) contend that the turning point for temperate European countries occurs at slightly low intervals, from roughly 16 °C. Having tested this formally, however, there is no significant difference in using 16 °C or 18 °C as the threshold for HDDs for our data set.

¹²To illustrate, an aggregate daily temperature of 17 °C would correspond to one HDD, while a temperature of 15 °C would equate to three HDDs. Similarly, a temperature of 27 °C would correspond to five CDDs, and so forth.

Kleibergen-Paap test, a robust variant of the Stock and Yogo (2005) test that allows for non-iid errors, is used to check the validity of our instruments (see Baum et al., 2007). This is complemented by a simple F -test of the instruments in the first-stage regression (Staiger and Stock, 1997).

1.3 Data

The data for this paper are collected from several different sources, with each series consisting of daily observations over a seven-year period from 2002 to 2009. German spot electricity prices and volumes are obtained from the EEX. These electricity data are available for both base (24-hour continuous) and peak (12 hours from 8am to 8pm) periods. However, we focus exclusively on the base series in this paper. Our primary motivation is that the power plants most vulnerable to water-related factors – such as nuclear and coal-fired plants – are all baseload electricity operators. Consequently, one would expect that the impact of any supply constraints to these plants will already be visible within the base price. Moreover, having a data point that runs over an entire day helps to ensure consistency with the other variables, which also cover daily time steps. It should also be noted that electricity prices in Germany are geographically uniform with no zonal differentiation.¹³ Both electricity prices and volumes are log-transformed for the regression analysis.

Air temperature data are obtained from *Deutscher Wetterdienst* (the German Meteorological Service). To compute aggregate temperature data, daily values are first collected for each capital city of the 16 German federal states. In the minority of cases where data limitations mean that a state cannot be represented by its capital, a significant counterpart city is used instead.¹⁴ The mean temperature recording in all of these cities (computed from 24 hourly observations) is then computed into a single daily mean temperature series for the entire country. This aggregating step is taken to ensure consistency with the uniform electricity prices across the German states. Next, we create a set of degree day interaction dummies, HDD and CDD. These variables are adjusted so as to reflect logged values, i.e. $D^{temp > 22^\circ C} \cdot \log(temp - 22^\circ C)$ and $D^{temp < 18^\circ C} \cdot \log(18^\circ C - temp)$. The daily mean air temperature series is shown in Figure 1.3 together with the designated

¹³The regulatory framework of the EEX does allow for the market to be broken up into different price zones when grid capacities are unable to fully execute the spot auction schedules, but this was not necessary during the review period of this study (Ockenfels et al., 2008).

¹⁴For instance, data for Wiesbaden, the capital of Hesse, were not available so these were substituted with data from the much bigger Frankfurt.

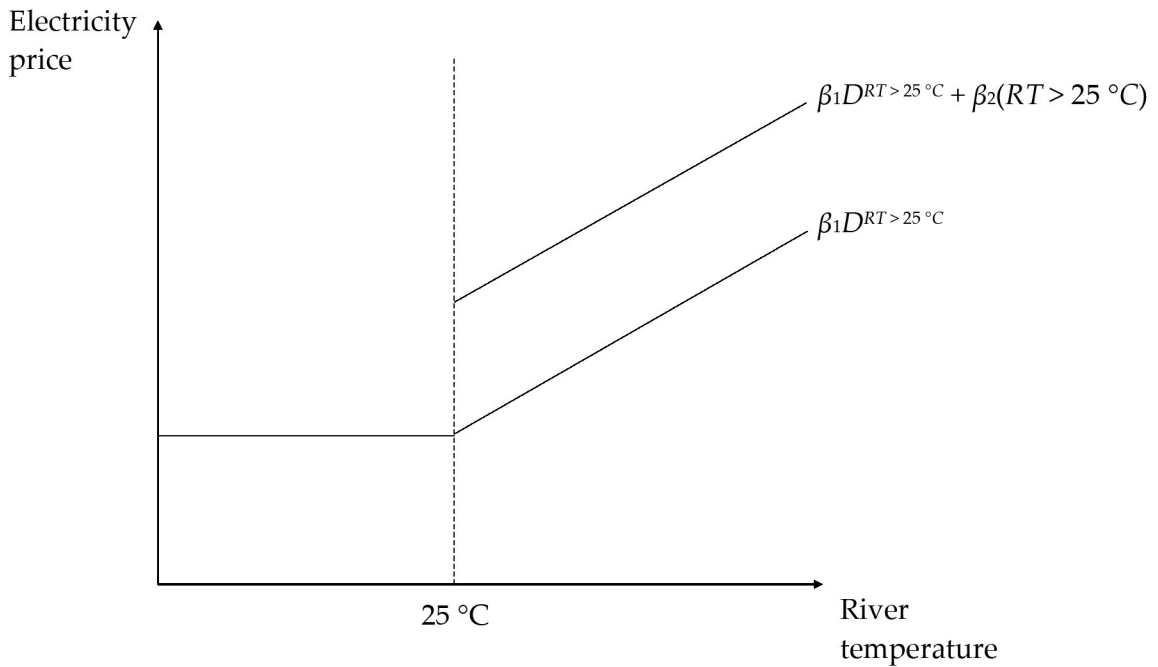


Figure 1.4: Effect of river temperature on electricity price

the regression model. However, we make two adjustments to the RT series to better capture how regulation of thermal pollution impacts electricity prices. The first is to generate a standard dummy variable that tests for a difference in price intercept when river temperatures exceed a defined regulatory limit of 25 °C. The second is to specifically measure the continued rise in temperature above this 25 °C threshold, i.e. $D^{RT > 25^\circ\text{C}} \cdot \log(RT - 25^\circ\text{C})$. This formulation is aimed at ensuring some flexibility and allows for a non-linear temperature effect around the regulatory threshold. It should be noted that our specification here is consistent with the theoretical model described in Section 1.2; a shadow price comes into play when the river temperature is greater than some regulatory limit, with a positive marginal effect on electricity prices as temperatures increase above that threshold. Figure 1.4 depicts this relationship in a stylised manner.

We use the log-transformed, 90-day moving average (MA) of Brent crude oil to account for the effect that changing fuel (i.e. input) costs have on power production. These are obtained from Bloomberg and adjusted for changes to the USD-EUR exchange rate. While oil-fired plants do not play a substantial role in the German electricity market, oil is widely used as a proxy for natural gas and it is even used within the power industry to forecast the general price movements of coal. The fact that daily spot prices are available for oil also makes it more amenable to our empirical analysis.

Finally, we include a number of parameters in the model to control for trend and seasonality. Month and year dummies are created to pick up the standard seasonal characteristics found in electricity data, as well as unaccounted trends in demand. (For example, those stemming from changes in consumers' aggregate income level over the review period.) Furthermore, since electricity consumption is expected to be highly correlated with economic activity, a dummy variable for non-working days (i.e. weekends and public holidays) is also included in the regression analysis.

1.4 Results

Our primary estimation results are presented in Table 1.1. Model (1) is characterised by a static setting that utilises only contemporaneous variables. Models (2) and (3) are dynamic in the sense that they include lagged electricity price and volume observations as additional regressors. All results have been calculated using heteroskedasticity- and autocorrelation-consistent (HAC) estimators (Newey and West, 1987).

Considering model (1) first, the rationale underpinning this "static" specification is that – given its role as an optimising market – the spot power exchange should effectively constitute a new market each day. The coefficient on volume (5.976) suggests that a one percent increase in the volume of consumed electricity will induce a six percent increase in the base electricity price. This implies that the daily power supply in Germany is highly inelastic, which we would expect given the very short-term nature of the data used in this study (i.e. daily observations).

Water scarcity, as measured by changes to the average river level, returns a negative coefficient in the static model (-1.025); indicating that the electricity price is expected to fall by around one percent for every one percent rise in river levels. This is consistent with our earlier hypothesis that electricity prices move in the opposite direction to the availability of cooling water, even after controlling for potential demand effects.

Model (1) also shows that there is a positive, statistically significant relationship between the electricity price and an aggregate river temperature over 25°C . Once this threshold is breached, the price rises by 0.277 percent with every additional percentage increase in river temperatures. To put this in perspective, a rise in average river temperatures from 25°C to 26°C would yield an approximate 1.2 percent increase in the price of electricity. The fact that the $D^{RT > 25^{\circ}\text{C}}$ dummy returns a

Table 1.1: Primary Models

Dependent Variable: Price (€/MWh)

	Model 1	Model 2	Model 3
<i>Coefficients</i>			
Volume (1,000 MWh)	5.976** (0.503)	8.267** (1.021)	8.233** (0.998)
<i>Predetermined Variables</i>			
L1.Price		0.624** (0.0541)	0.625** (0.0537)
L7.Price		0.138** (0.0533)	0.141** (0.0526)
L1.Volume		-3.729** (0.466)	-3.697** (0.451)
L7.Volume		-2.280** (0.372)	-2.272** (0.361)
<i>River Levels (cm)</i>			
Single series	-1.025** (0.234)	-0.389** (0.135)	
“Low”(0-33%)			-0.206 (0.232)
“Medium”(33-67%)			-0.148 (0.303)
“High”(67-100%)			-0.580** (0.223)
<i>River Temperature (°C)</i>			
D_{Riv25} (1 = $RT > 25$ °C)	0.0944 (0.156)	-0.0529 (0.0751)	-0.0273 (0.0745)
$RT - 25$ °C, if > 25 °C	0.277** (0.0550)	0.210** (0.0286)	0.218** (0.0288)
Brent (90-day MA, €/bbl)	0.528 (0.442)	0.163 (0.200)	0.149 (0.185)
<i>Tests</i>			
1SLS instrument joint significance test	123.65**	36.56**	37.13**
Instrument relevance test ^b	65.01**	23.64**	25.70**
Overidentifying restrictions test ^c	3.20	5.43	5.40
Autocorrelation test ^d	1604.23**	2.69	2.69
<i>Joint significance tests χ^2</i>			
Month Dummies	64.76**	35.33**	35.91**
Year Dummies	147.73**	42.45**	43.92**
<i>N</i>	2,922	2,915	2,915

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

Notes: All variables are entered in logarithmic form. Standard errors for the coefficients are reported in parentheses. For the tests, p-values are reported rather than test statistics. A constant term, year and month dummies are also included as regressors in the price equation. However, the estimated coefficients attached to these variables are not reported in the table. Heating degree days (*HDD*), cooling degree days (*CDD*), and a Non-working day (*NWD*) dummy are used as instruments in the first-stage regression.

^a Durbin-Wu-Hausman *F*-test includes the saved residuals from the first-stage regression in the second stage of the 2SLS estimation. H0: System is exogenous.

^b Kleibergen-Paap Wald *rk F* statistic. H0: System is underidentified and the instruments are not relevant (i.e. weak).

^c The Hansen *J* statistic for overidentifying restrictions is computed using HAC estimators. H0: Instruments are exogenous.

^d The autocorrelation test statistics is the *F* statistics of the coefficient from a regression where the residual from the main regression is regressed on its own lagged value.

statistically insignificant coefficient implies that there is no discontinuity around this 25 °C threshold. Again, this is consistent with our theoretical model in which, rather than simply shutting down, power plants have the option of reducing power to stay within the regulatory limits set by authorities.

While fuel costs are not of special interest to this paper, the coefficient on the 90-day MA for Brent crude is positive, albeit statistically insignificant. We would expect a positive sign given that fossil fuels serve as an important factor of production in generating electricity. The set of month and year dummies are not reported on an individual basis, but they are all jointly significant.¹⁶

Our use of an IV/2SLS approach has been motivated by the fact that prices and quantities are determined simultaneously. This is confirmed by the Durbin-Wu-Hausman test, which shows that endogeneity is a problem and that OLS should be discarded in favour of 2SLS. The Kleibergen Paap test indicates that our instruments are highly relevant (i.e. no weak instrument problem). We complement this with the recommendation of Staiger and Stock (1997) by using an *F*-test to test the relevance of the instruments in the first stage of the 2SLS. This *F*-statistic is significantly greater than their “rule of thumb” value of 10 and so we again reject the null of weak instruments. In addition, Hansen’s *J* test of overidentifying restrictions shows that they are also valid.¹⁷

Applying an augmented Dickey-Fuller test (ADF) on the residuals for model (1) shows that persistency in the data does not appear to be a problem.¹⁸ However, further testing does reveal the presence of positive autocorrelation. A probable explanation for this is misspecified dynamics in particular, our reliance thus far on a completely static model specification. Yet it could be argued that today’s elec-

¹⁶There is an increasingly negative coefficient on the year dummy coefficients until 2009, demonstrating that electricity prices have been increasing slowly relative to volumes over the years. Furthermore, the coefficients on the month dummies indicate that German electricity prices are typically higher in the summer months.

¹⁷ It could be argued that our *CDD* instrument has a potentially direct effect on electricity supply since high air temperatures – the basis for *CDD* – and river temperatures are correlated, albeit with a lag of several days. The increase in river temperatures implies a higher likelihood of regulatory enforcement. If so, there could be a direct link between *CDD* and supply of electricity that consequently violates the exogeneity assumption. In the spirit of an overidentification test, we have therefore run an auxiliary 2SLS that only uses our other two instruments; *HDD* and *NWD*. These two instruments are exogenous by assumption and we have no cause to think that they should be correlated with high river temperatures. We find that the predicted residuals of this auxiliary regression are not correlated with the debatable instrument *CDD* (*p*-value = 0.180). This suggests that *CDD* is a valid instrument in our setting.

¹⁸Although not reported, it is also tested whether persistency (i.e. non-stationarity) is a problem for the log-transformed electricity prices- and volumes using an ADF test. Having accounted for trends in the form of year and monthly dummies we are able to reject the null hypothesis of non-stationarity for these series.

tricity price is correlated with the previous day's price, or even that of the week before. This idea is given credence by the fact that electricity supply is comprised of quasi-fixed proportions of baseload and variable power. Baseload facilities such as nuclear and coal-fired power plants are typically constrained in their ability to change output levels.¹⁹ One might therefore argue that there is some "memory" in the power market system and that our modelling efforts could be improved by incorporating dynamic aspects.

The key results from two such dynamic models, which include one- and seven-day lags for both electricity price and volumes, are presented in the second and third columns of Table 1.1. These lags are chosen to account for the inertia from the previous day and weekday effects. Model (2) is a straightforward extension of our static model, while in model (3) we want to open up for the possibility that changes in water availability matter at different stages of relative abundance. Thus, the continuous river level series has been replaced by spline partitions. Since the majority of the coefficients of these two models are qualitatively indistinguishable, we consider them together.

It can immediately be seen that the coefficients on the lagged endogenous variables in both dynamic models are all statistically significant. This goes some way towards vindicating our suspicions that the German electricity spot exchange does not simply constitute a "new" market every day. Illustrating by way of model (2), the coefficient on the contemporaneous volume of electricity (i.e. 8.267) denotes the short-run, instantaneous impact of a change in quantity on price. The corresponding long-run multiplier is found by incorporating the lagged endogenous variables of our model and can be calculated as $[(8.267 - 3.729 - 2.28) / (1 - 0.624 - 0.138) =] 9.487$. Testing this figure reveals it to be statistically significant at the one percent level. The dynamic model specification therefore predicts that a one percent increase in electricity volumes will lead to a 9.5 percent increase in price over the course of a full week. Again, this describes a very inelastic supply curve, but it is representative of the inertia present within the power system.

Looking next at the effect of river temperatures, both dynamic models show that there is a positive impact on electricity prices once the 25 °C threshold is breached. A one percent increase in river temperatures above this mark will yield an increase

¹⁹The load-following capacity of baseload power is an important concept here. In particular, nuclear and coal-fired plants are normally run continuously at more-or-less constant level of output. This is both a matter of economic efficiency (since they have low variable costs in comparison with the high fixed costs that must be recouped), and technical efficiency (since these plants cannot readily alter power output in the same way that gas or hydro plants can). See, for instance, WNA (2011).

in contemporaneous prices slightly greater than 0.2 percent. The equivalent long-run effect is closer to 0.9 percent. Thus, a temperature rise from 25 °C to 26 °C would bring about an immediate price increase of approximately 0.9 percent, or equivalently, an increase of 3.8 percent over the next seven days. Once more, these effects are all statistically significant at the one percent level.

The key distinction between our two dynamic models lies in the way that they measure the impact of changing river levels. Model (2) is a straightforward extension of model (1) in that it uses a single, continuous series. Like its static counterpart, model (2) suggests a negative relationship with the electricity price: A one percent drop in river levels is associated with a 0.4 percent rise in the concurrent electricity price, while the relevant long-run multiplier suggests an approximate 1.6 percent rise over the course of a week.

For model (3), we split the river level series into three splines of equal size based on percentile distribution: i) “low” (0%–33%); ii) “medium” (33%–67%); and iii) “high” (67%–100%).²⁰ Only changes within the “high” river level category are shown to be statistically significant: A one percent drop in river levels within this range will lead to a 0.6 percent rise in contemporaneous prices, or a 1.8 percent rise in the long run. A potential explanation for the insignificance of the “medium” and “low” river level splines could be that those plants most reliant on water consumption – i.e. those most sensitive to water scarcity – have already been forced to power down by the time that rivers reach their lowest levels. Regardless, formal testing reveals that the coefficient on the “mid” spline is statistically indistinguishable from that of the “high” spline (p -value of 0.31). Conducting a similar test on the “low” spline coefficient reveals that it too is statistically identical to the “high” spline (p -value of 0.18). As a consequence of these tests, it makes sense to do away with the separate splines and simply include river levels as a single continuous series as in our preferred specification, model (2).

Running through the same set of statistical tests described previously, we are able to confirm the validity of our instruments (as well as the presence of endogeneity that necessitates the use of an IV approach in the first place). A more pertinent question concerning the extension towards a dynamic specification, however, is whether it removes the autocorrelation that was present in the static models. It is well known that in addition to efficiency concerns, inclusion of the lagged dependent variable will lead to biased coefficient estimates in the presence of serially correlated resid-

²⁰Admittedly, these splines are chosen somewhat arbitrarily. However, having experimented with different cut-off points, our conclusion is that the main results are robust to such changes.

uals. That said, testing reveals that autocorrelation is not present in any of the dynamic models. This adds further credibility to the notion that the dynamic specification of our model is preferable to its static counterpart.²¹

In addition to the primary models presented in Table 1.1, we have run a number of alternate specifications and supplementary regressions to confirm the robustness of our findings.²² First, instead of using month dummies, we have also tried to control for seasonal effects by incorporating a trigonometric wave function in the models.²³ Doing so does not change our results in any material way. Second, we have replaced the splines of model (3) with a set of dummies that similarly divides the river level series into equal thirds (“low”, “middle” and “high”) based on the overall distribution. Rather than measuring the potential difference in slope coefficients, the aim here is to assess whether there are any statistically significant differences in the intercepts of the different river level groups. Again, we find that electricity prices will be driven higher as river levels fall.²⁴

By focusing on high river temperatures (i.e. 25 °C and above), our goal has been to capture the impact that regulatory constraints have on power supply. Such regulatory interventions are a direct result of the increased river temperatures which are measured directly in our analysis. In addition, by holding our data together with shutdown dates collected by Müller et al. (2007), we find a jump in prices during the period of regulatory enforcement compared with those immediately preceding and/or following them.²⁵ There is also an uncanny overlap between these days and times where the average river temperature exceeds the 25 °C mark. We view this as supplementary evidence in favour of our analysis of the effects of regulatory actions and increased river temperature, as well as our defined regulatory threshold.²⁶

²¹While of lesser importance to this study, we again note that the coefficient on fuel costs remains positive yet insignificant under the dynamic specification.

²²While none of the results from the alternate specifications are reported here, they are available from the authors upon request.

²³The formulas used are $F(t) = \sin(2\pi t/365)$ and $G(t) = \cos(2\pi t/365)$, respectively, where t denotes time in days. This reflects the fact that a full seasonal cycle would complete once a year.

²⁴We have also included a log-transformed, 90-day MA of CO₂ future contracts as a proxy for the input costs of thermal-based electricity production. While we only have available data for the period 2006–2009, the basic results of our regression analysis are not altered by the inclusion of this CO₂ permit variable.

²⁵Müller et al. (2007) include information about shutdown dates over the period 19–31 July 2006 based on secondary sources, such as newspaper and other media articles. Such data are likely to be incomplete and imprecise and should therefore be used with great care.

²⁶The aggregate 25 °C threshold that we have defined does gloss over some river- and site-specific issues. The permitted mixing temperature from the thermal discharge in Germany varies between 23 °C and 28 °C Müller et al. (2007). Our choice of 25 °C is based on a careful reading of the literature, as well as some initial testing of different thresholds.

So far, we have followed an aggregated approach with regards to both river level and river temperature. This decision has primarily been motivated by the fact that Germany has a single electricity price common to all regions. Moreover, the purpose of this paper is to essentially test for *systematic* risks – particularly with regards to river temperatures and the regulation of thermal pollution. However, it could still be asked whether the individual rivers in our dataset are similar enough to warrant this type of aggregation. We have therefore subjected our data to several robustness checks. First, we construct a simple correlation matrix of the (detrended) individual river temperature series.²⁷ These correlation coefficients fall within the interval [0.80, 0.94] and are highly significant. We have also looked at the share of individual temperature observations that coincide with the days that our average river temperature series breaches the 25 °C threshold. The likelihood of an individual river exceeding 25 °C given that the average series does, is very high (= 0.993). These exercises illustrate the close correspondence between the individual temperature trends and the aggregated measure of critical river temperature that we have constructed.

The correlation coefficients for the individual river levels are less pronounced, but are still highly significant. To provide a more formal test of the disaggregated river level effects though, we incorporate data from each of the individual rivers separately into the regression analysis. More precisely, we run four new regressions based on the specifications of model (2) – each time replacing the average river level series with data from a single river. These results are presented in Table 1.2. As can be seen from the table, the individual river level coefficients are all negative and thus indicative of a higher electricity price when river levels fall. That is not to say that they all have the same marginal impact, although this is perhaps not surprising given that the importance of these rivers in terms of providing cooling water to Germany’s thermal industry can vary quite substantially.²⁸

In sum, we believe that these results serve to emphasize the validity of our predominantly aggregate approach. Again, the purpose of this paper is to test for systematic vulnerabilities and we would argue that focusing too much on individual trends and measurements actually distracts from the wider climate and its associated risks. The real danger implicit in climate change, for example, is that mean values are pushed closer to their regulatory thresholds, such that widespread capacity reductions become more commonplace. It therefore seems most appropriate to focus on the “average” effect, since this captures the systemic risk that comes

²⁷These series are detrended using the set of month and year dummies.

²⁸As expected, the other coefficients are extremely similar across the four different models.

Table 1.2: Individual River Levels

	Dependent Variable: Price (€/MWh)			
	Elbe	Main	Neckar	Rhine
<i>Coefficients</i>				
Volume (1,000 MWh)	8.099** (1.004)	8.256** (1.020)	8.081** (0.982)	8.277** (1.029)
<i>Predetermined Variables</i>				
L1.Price	0.644** (0.0530)	0.633** (0.0532)	0.623** (0.0527)	0.620** (0.0553)
L7.Price	0.158** (0.0540)	0.153** (0.0530)	0.168** (0.0531)	0.140** (0.0537)
L1.Volume	-3.738** (0.464)	-3.757** (0.460)	-3.672** (0.447)	-3.749** (0.472)
L7.Volume	-2.230** (0.374)	-2.330** (0.368)	-2.271** (0.360)	-2.284** (0.379)
<i>River Levels (cm)</i>				
Single series	-0.130 (0.0741)	-0.372* (0.171)	-0.511** (0.142)	-0.217** (0.0740)
<i>River Temperature (°C)</i>				
D_{Riv25} ($1 = RT > 25$ °C)	-0.0246 (0.0732)	0.00289 (0.0739)	0.00373 (0.0739)	-0.0683 (0.0752)
$RT - 25$ °C, if > 25 °C	0.208** (0.0301)	0.222** (0.0282)	0.227** (0.0284)	0.210** (0.0280)
Brent (90-day MA, €/bbl)	0.131 (0.194)	0.113 (0.189)	0.0755 (0.191)	0.137 (0.205)
<i>Tests</i>				
Endogeneity test ^a	478.40**	431.99**	442.59**	448.29**
1SLS instrument joint significance test	36.78**	36.67**	37.14**	36.63**
Instrument relevance test ^b	22.74**	24.55**	23.26**	29.37**
Overidentifying restrictions test ^c	5.27	5.78	5.58	5.23
Autocorrelation test ^d	4.27*	3.38	4.24*	2.75
<i>Joint significance tests (χ^2)</i>				
Month Dummies	35.25**	36.04**	37.95**	36.89**
Year Dummies	39.31**	38.95**	43.02**	42.16**
<i>N</i>	2,915	2,915	2,915	2,915

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

Notes: Based on Model (2) in Table 1.1. For each regression, the average river level series from Model (2) have been replaced with level data from an individual river. All other coefficients are excluded.

^{a b c d} See Table 1.1 for interpretation.

when rivers all across the country are breaching their regulatory thresholds at the same time.

1.5 Concluding remarks

This paper has sought to quantify how electricity prices are affected by the availability of cooling water. Our analysis is primarily motivated by the fact that water plays a critical role in the thermal production cycle, where tremendous volumes of freshwater are drawn every day to serve the cooling needs of thermal-based power plants around the world. At the same time, these water sources are subject to environmental regulations, competing economic concerns and periods of relative scarcity. We have argued that Germany serves as a good case study to investigate these issues

and have based our analysis on daily data taken over a period of seven years.

Having successfully controlled for various demand effects within a 2SLS regression framework, our results indicate that electricity prices are significantly affected by both falling river levels and higher river temperatures. The magnitude of these relationships varies according to the exact specifications of the regression model at hand and we have explored several contemporaneous and dynamic settings. Qualitatively, however, they all tell a very similar story: electricity prices are driven higher by falling river levels and high river temperatures. Under a fully contemporaneous setting, the electricity price is expected to rise by around one percent for every one percent that river levels fall. The dynamic specification, on the other hand, suggests that the price will rise at about half that rate in the short-run, before increasing to approximately one and a half percent in the long-run. With regards to river temperatures, the models imply that the price of electricity will increase by roughly one percent for every degree that temperatures rise above a 25 °C threshold. Incorporating the longer-run effects implied by a dynamic model shows that prices will rise by nearly four percent over the course of a week. In addition to this slope effect, we test for a price discontinuity on either side of this 25 °C threshold. However, we do not find evidence of a marked price jump once the threshold is breached. An explanation, which is consistent with our theoretical model and the surveyed literature, is that power plants reduce their output in stages rather than simply shutting down. This allows them some additional scope for managing thermal pollution, although a decrease in output – and hence increase in price – cannot be fully avoided.

One implication of our findings is that future climate change will impact electricity prices not only through changes in demand, but also as a result of increased cooling water scarcity. We believe that this type of analysis would lend itself to applications in a number of regions and countries – all of which are marked by a marked dependency on thermal-based power, at the same time as being prone to drought and periodic heat waves.

References

- AP (2008). "Drought could shut down nuclear power plants," *Associated Press*, 23 January 2008. Available: <http://www.msnbc.msn.com/id/22804065/ns/weather/t/drought-could-shut-down-nuclear-power-plants/>.
- Arnell, Nigel, Emma L Tompkins, Neil Adger, and Kate Delaney (2005). "Vulnerability To Abrupt Climate Change In Europe," Tyndall Centre Technical Report 34, Tyndall Centre for Climate Change Research.
- Baum, Christopher F, Mark E Schaffer, and Steven Stillman (2007). "Enhanced routines for instrumental variables/GMM estimation and testing," *Stata Journal*, 7(4): 465–506.
- Bessec, Marie and Julien Fouquau (2008). "The non-linear link between electricity consumption and temperature in Europe: a threshold panel approach," *Energy Economics*, 30(5): 2705–2721.
- Boogert, Alexander and Dominique Dupont (2005). "The nature of supply side effects on electricity prices: The impact of water temperature," *Economics Letters*, 88(1): 121–125.
- Chellaney, Brahma (2011). "Japans Nuclear Morality Tale," *Project Syndicate*, 14 March 2011. Available: <http://www.project-syndicate.org/commentary/japan-s-nuclear-morality-tale>.
- Dell'Amore, Christine (2010). "Nuclear Reactors, Dams at Risk Due to Global Warming," *National Geographic News*, 26 February 2010. Available: <http://news.nationalgeographic.com/news/2010/02/100226-water-energy-climate-change-dams-nuclear/>.
- Dixit, Avinash K. (1990). *Optimization In Economic Theory*, Oxford University Press.
- DOE (2006). "Energy demands on water resources: Report to congress on the interdependency of energy and water," Technical report, U.S. Department of Energy,

- Washington D.C., report to Congress on the interdependency of energy and water.
- Eaton, Jon (2012). "Record Heat, Drought Pose Problems for U.S. Electric Power," *National Geographic News*, 17 August 2012. Available: <http://news.nationalgeographic.com/news/energy/2012/08/120817-record-heat-drought-pose-problems-for-electric-power-grid/>.
- Feeley, Thomas J., Timothy J. Skone, Gary J. Stiegel Jr, Andrea McNemar, Michael Nemeth, Brian Schimmoller, James T. Murphy, and Lynn Manfredo (2008). "Water: A critical resource in the thermoelectric power industry," *Energy*, 33(1): 1–11.
- Förster, Hannah and Johan Lilliestam (2010). "Modeling thermoelectric power generation in view of climate change," *Regional Environmental Change*, 10(4): 327–338.
- Gentleman, Amelia (2003). "France faces nuclear power crisis," *The Guardian*, 8 April 2010. Available: <http://www.theguardian.com/news/2003/aug/13/france.internationalnews>.
- Godoy, Julio (2006). "European Heat Wave Shows Limits of Nuclear Energy," *One World*, 26 July 2006. Available: <http://us.oneworld.net/article/european-heat-wave-shows-limits-nuclear-energy>.
- Halvorsen, Robert (1975). "Residential demand for electric energy," *The Review of Economics and Statistics*, 57(1): 12–18.
- Hurd, Brian and Megan Harrod (2001). "Water resources: Economic Analysis," in Robert Mendelsohn (ed.), "Global Warming and the American Economy: A Regional Assessment of Climate Change Impacts," chapter 5, Edward Elgar Publishing Ltd, pp. 106–131.
- IEA (2011a). *Energy Balances of OECD Countries 2011*, International Energy Agency, OECD Publishing.
- IEA (2011b). *World Energy Outlook 2011*, International Energy Agency, OECD Publishing.
- IPCC (2007). *Climate change 2007: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change*, [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)], Cambridge: Cambridge University Press, 1st edition.

- Kanter, James (2007). "Climate change puts nuclear energy into hot water," *The New York Times*, 20 May 2007. Available: <http://www.nytimes.com/2007/05/20/health/20iht-nuke.1.5788480.html>.
- Kirshen, Paul, Matthias Ruth, and William Anderson (2008). "Interdependencies of urban climate change impacts and adaptation strategies: a case study of metropolitan boston usa," *Climatic Change*, 86(1-2): 105–122.
- Koch, Hagen and Stefan Vögele (2009). "Dynamic modelling of water demand, water availability and adaptation strategies for power plants to global change," *Ecological Economics*, 68(7): 2031–2039.
- Kopytko, Natalie and John Perkins (2011). "Climate change, nuclear power, and the adaptation–mitigation dilemma," *Energy Policy*, 39(1): 318–333.
- Kristiansen, Tarjei (2011). "Power trading analytics and forecasting in germany," *The Electricity Journal*, 24(8): 41–55.
- Langford, Terry E. (1990). *Ecological effects of thermal discharges*, London: Elsevier Applied Science.
- Linnerud, Kristin, Torben K Mideksa, and Gunnar S Eskeland (2011). "The impact of climate change on nuclear power supply," *Energy Journal*, 32(1): 149–168.
- Maulbetsch, John S. and Michael N. DiFilippo (2006). "Cost and value of water use at combined cycle power plants," Technical Report CEC-500-2006-034, California Energy Commission, PIER Energy-Related Environmental Research.
- Mimler, S., U. Müller, S. Greis, and B. Rothstein (2009). "Impacts of climate change on electricity generation and consumption," in Walter Leal Filho and Franziska Mannke (eds.), "Interdisciplinary Aspects of Climate Change," volume 29, Frankfurt: Peter Lang Scientific Publishers, pp. 11–37.
- Möst, Dominik and Massimo Genoese (2009). "Market power in the German wholesale electricity market," *The Journal of Energy Markets*, 2(2): 47–74.
- Müller, U., S. Greis, and B. Rothstein (2007). "Impacts on water temperatures of selected german rivers and on electricity production of thermal power plants due to climate change," in P. Heneka, B. Zum Kley, G. Tetzlaff, and F. Wenzel (eds.), "8 Forum DKKV/CEDIM: Disaster Reduction in a Changing Climate," Karlsruhe.
- Müller, U, S Greis, and B. Rothstein (2008). "Möglicher Einfluss des Klimawandels auf Flusswassertemperaturen und Elektrizitätserzeugung thermischer Kraftwerke," Poster

- presented at *Tag der Hydrologie 2008*, 27-28 March 2008. Available: <http://www.iww.uni-hannover.de/tdh2008/Poster/Mueller.pdf>.
- Müsgens, Felix (2006). "Quantifying market power in the German wholesale electricity market using a dynamic multi-regional dispatch model," *The Journal of Industrial Economics*, 54(4): 471–498.
- NETL (2009a). "Estimating freshwater needs to meet future thermoelectric generation requirements," Report Number: DOE/NETL-400/2009/1339, U.S. National Energy Technology Laboratory.
- NETL (2009b). "Impact of drought on u.s. steam electric power plant cooling water intakes and related water resource management issues," Report Number: DOE/NETL-2009/1364, U.S. National Energy Technology Laboratory.
- NETL (2009c). "Water Requirements for Existing and Emerging Thermoelectric Plant Technologies," Report Number: DOE/NETL-402/080108, U.S. National Energy Technology Laboratory.
- Newey, Whitney K. and Kenneth D. West (1987). "A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix," *Econometrica*, 55: 703–708.
- Ockenfels, A., V. Grimm, and G. Zoetl (2008). "Electricity Market Design: The Pricing Mechanism of the Day Ahead Electricity Spot Market Auction on the EEX," Technical report, European Energy Exchange AG (EEX).
- Pagnamenta, Robin (2009). "France imports UK electricity as plants shut," *The Times*, 3 July 2009. Available: <http://www.thetimes.co.uk/tto/business/industries/utilities/article2198065.ece>.
- Quayle, Robert G. and Henry F. Diaz (1980). "Heating degree day data applied to residential heating energy consumption," *Journal of Applied Meteorology*, 19(3): 241–246.
- Sohn, Pam (2011). "River temperature forces nuclear plant to 50 percent power," *Times Free Press*, 4 August 2011. Available: <http://www.timesfreepress.com/news/2011/aug/04/river-temperature-forces-plant-to-50-percent/>.
- Sovacool, Benjamin K. (2009). "Running on empty: The electricity-water nexus and the US electric utility sector," *Energy Law Journal*, 30(1): 11–51.

- Sovacool, Benjamin K. and Kelly E. Sovacool (2009). "Preventing national electricity-water crisis areas in the United States," *Columbia Journal of Environmental Law*, 34(2): 333–393.
- Staiger, Douglas and James H. Stock (1997). "Instrumental variables regression with weak instruments," *Econometrica*, 65(3): 557–586.
- Van Vliet, Michelle T.H., John R. Yearsley, Fulco Ludwig, Stefan Vögele, Dennis P. Lettenmaier, and Pavel Kabat (2012). "Vulnerability of US and European electricity supply to climate change," *Nature Climate Change*, 2(9): 676–681.
- Viehmann, Johannes (2011). "Risk premiums in the German day-ahead Electricity Market," *Energy policy*, 39(1): 386–394.
- WNA (2011). "Nuclear power reactors," *World Nuclear Association*, March 2009. Available: <http://www.world-nuclear.org/info/inf32.html>.

Appendix

1.A Comparative statics: Effect of a change in river temperature on optimal quantity and price

We start from our first-order condition (eq. 1.5),

$$p^* \cdot \left(1 + \frac{1}{\epsilon}\right) = \frac{\partial c(Q^*)}{\partial Q} + \frac{p_w(RL)}{A(T_{EW} - T)} + \lambda \cdot \left[\frac{T_{EW} - T}{S \cdot A(T_{EW} - T)} \right]$$

and want to find the effect of a river-level temperature increase on the prices. Using standard comparative statics we define the problem as follows (e.g. Dixit, 1990, chap. 8):

$$\frac{dQ}{dT} = - \frac{\partial G(Q, T) / \partial T}{\partial G(Q, T) / \partial Q},$$

where

$$\begin{aligned} G(Q, T) &= \frac{W}{S} \cdot T_{EW} + \frac{S - W}{S} \cdot T \\ &= \frac{W}{S} \cdot (T_{EW} - T) + 1 \\ &= \frac{Q}{S \cdot A(\cdot)} \cdot (T_{EW} - T) + 1, \end{aligned}$$

which is the LHS of the constraint $G(Q, T) \leq T$. We have that

$$\frac{\partial G(\cdot)}{\partial Q} = \frac{T_{EW} - T}{S \cdot A(\cdot)} > 0$$

and

$$\begin{aligned}
 \frac{\partial G(\cdot)}{\partial T} &= \frac{Q}{S} \cdot \left[\frac{A'(\cdot) \cdot (T_{EW} - T) - A(\cdot)}{A(\cdot)^2} \right] + 1 \\
 &= \frac{A(\cdot)W}{S} \cdot \left[\frac{A'(\cdot) \cdot (T_{EW} - T) - A(\cdot)}{A(\cdot)^2} \right] + 1 \\
 &= \frac{1}{A(\cdot)} \cdot \left[\frac{W}{S} \cdot [A'(\cdot)(T_{EW} - T) - A(\cdot)] + A(\cdot) \right] \\
 &= \frac{1}{A(\cdot)} \cdot \left[\frac{S - W}{S} \cdot A(\cdot) + A'(\cdot) \cdot (T_{EW} - T) \right] \\
 &> 0,
 \end{aligned}$$

as long as $A'(\cdot) > 0$ (and $T_{EW} > T$). That means that $\frac{dQ}{dT} < 0$. This is an inward shift of the supply curve. With a given demand, and assuming that the supply curve is not completely flat, then a rise in river temperature will lead to an increased electricity price.

Chapter 2

Sceptic priors and climate change mitigation

Abstract

How much evidence would it take to convince sceptics that they are wrong about global warming? I explore this question within a Bayesian framework. I consider a group of stylised climate sceptics and examine how these individuals update their beliefs in light of the available evidence. I find that all but the most extreme priors are overwhelmed by the historical data. The resulting posterior distributions of climate sensitivity correspond closely to existing estimates from the literature. I show further that the updated sceptic beliefs are consistent with a social cost of carbon that is substantially greater than zero. I conclude by discussing the general conditions for consensus formation under Bayesian learning, its relevance to our current policy impasse, and offer some remarks about finding common ground in the future.

JEL Codes: Q54, Q58, C11

Keywords: climate change, climate scepticism, social cost of carbon, Bayesian econometrics

2.1 Introduction

Climate change has come to represent a defining policy issue of our age. Yet despite an overwhelming scientific consensus, many ordinary citizens and policy makers are sceptical about anthropogenic global warming.¹ What are we to make of this scepticism? And just how much evidence would it take to convince a climate sceptic that they are wrong? The present paper addresses these questions within an idealised Bayesian framework. My approach is to combine sceptic priors with available climate data, and thereby obtain posterior probabilities about climate change that are logically consistent with the beliefs of sceptics. I consider whether real people behave like rational Bayesian agents and discuss how apparent deviations from this ideal can be accommodated within the same conceptual framework. In so doing, I hope to shed light on the current policy impasse and offer some remarks about the possibility for finding common ground in the near future.

Numerous studies have explored the cultural factors and psychological motivations that underlie climate scepticism (e.g. Kahan et al., 2011, 2012; McCright and Dunlap, 2011a,b; Corner et al., 2012; Ranney et al., 2012; Clark et al., 2013, and references therein). My present concern is less with the origins of scepticism than what it represents — a set of beliefs about the likely causes of global warming, which will in turn affect how new information about those causes is interpreted. A convenient way to model such beliefs is by defining scepticism in terms of climate sensitivity, i.e. the temperature response to a doubling of CO₂. More precisely, we can map sceptic beliefs directly to subjective estimates of climate sensitivity, since they both describe the likely causes and probability distribution of future warming.

Climate sensitivity can be defined in several ways and it is important to distinguish between these variants for accurate analysis. I focus on the policy-relevant measure known as transient climate response (TCR). Formally, TCR describes the warming at the time of CO₂ doubling — i.e. after 70 years — in a 1% per year increasing CO₂ experiment (Allen and Frame, 2007; Otto et al., 2013; IPCC, 2013, Box 12.2). For the purposes of this paper, however, it will simply be thought of as the contemporaneous change in global temperature that results from a steady doubling of atmospheric CO₂. According to the the Intergovernmental Panel on Climate Change (IPCC, 2013), TCR is “likely” to be somewhere in the range of 1.0–2.5 °C. This corresponds roughly to a 66–100% probability interval in IPCC terminology. The IPCC further emphasises the inherently Bayesian nature of climate sensitivity estimates,

¹For review and further discussion, see: Oreskes (2004); Anderegg et al. (2010); Doran and Zimmerman (2011); Cook et al. (2013); Verheggen et al. (2014); Saad (2014); Tol (2014).

going so far as to state:

“[T]he probabilistic estimates available in the literature for climate system parameters, such as ECS [i.e. equilibrium climate sensitivity] and TCR have all been based, implicitly or explicitly, on adopting a Bayesian approach and therefore, even if it is not explicitly stated, involve using some kind of prior information.” (IPCC, 2013, p. 922)

To understand why classical (frequentist) methods are ill-suited for the task of producing credible estimates of climate sensitivity, recall that frequentism interprets probability as the limiting frequency in a large number of repeated draws. Such a narrow definition holds little relevance to the question of climate sensitivity, for which there exists but one unique value. There is no sample of “sensitivities” to draw from. The present paper also adopts a Bayesian approach to the question of climate sensitivity and its concomitant policy implications. However, it differs from the previous literature primarily on two accounts.

First, I deliberately focus on the beliefs of sceptics as a means for evaluating the case to mitigate against climate change. Priors for determining climate sensitivity are usually based on the judgements of scientific experts, or obtained more objectively through noninformative priors. Such approaches may have obvious scientific merit when it comes to establishing a best estimate of climate sensitivity. However, they are of limited relevance for understanding people’s motivations and voting behaviour when it comes to actual climate policy. My approach is to take sceptics at their word and work through to the conclusions of their stated priors. Although contrarian beliefs have generally been ignored in the policy literature to date, a handful of studies do try to consider policy options from the sceptic’s perspective. Van Wijnbergen and Willems (2015) show that climate sceptics actually have an incentive to reduce emissions, as it will facilitate learning about the true causes of climate change. Kagan (2014) introduces various levels of scepticism into a numerical model and solves for optimal climate policy under the threat of catastrophe. He shows that even complete sceptics will not follow an unbounded emissions path. The possibility of (exogenous) catastrophe causes agents to draw down on their capital stock in pre-emptive fashion, leaving investment insufficient for continued economic growth beyond some point. From a methodological perspective, the current paper is mostly empirical in nature. This contrasts with the game-theoretic and numerical approaches of Van Wijnbergen and Willems (2015) and Kagan (2014), who are attempting to pin down the strategic emissions and learning paths for climate

sceptics under future uncertainty. My purpose is less to provide a prescriptive emissions path than it is to establish the ground rules for thinking about climate policy today, given the information that is already available to us.

A second distinguishing feature of this paper is that the results are derived via conceptually straightforward time-series regression analysis. While climate scientists have typically relied on complex computer models to simulate TCR, a growing body of research is aimed at understanding the link between human activities and climate change through the purview of time-series econometrics. Much of this literature has concerned itself with the apparent non-stationarity of climate data over time. Suffice it to say that the present paper takes as its foundation newer studies by Gay-Garcia et al. (2009) and Estrada et al. (2013a,b), who argue convincingly that global surface temperatures and anthropogenic forcings can best be described as trend-stationary processes, incorporating at least one structural change.² Such matters notwithstanding, virtually all econometric studies of climate change attribution to date have been carried out in the frequentist paradigm. They do not consider the influence of priors, nor are they able to yield the probabilistic estimates that are characteristic of Bayesian analysis. An early exception is that of Tol and De Vos (1998), who are motivated to adopt a Bayesian approach because of multicollinearity in their anthropogenic emissions data. Multicollinearity does not plague more recently available datasets, which unitise and aggregate various radiative forcings into a single series (see Section 2.2). This allows us to pin down the influence of specific forcings with greater confidence, and so obtain a more precise estimate of climate sensitivity. Furthermore, Tol and De Vos do not consider the influence of overtly contrarian priors as a basis for affecting policy.

2.2 Data

The various data sources for this paper are summarised in Table 2.1. Global mean surface temperature data (1850–2013) are taken from the HadCRUT4 dataset, jointly compiled by the UK Met Office and the Climatic Research Unit at the University of East Anglia. Two alternate global temperature reconstructions — one provided by

²The upshot is to permit the use of level terms in an ordinary least squares (OLS) regression framework. Another group of researchers has argued that the instrumental temperature record contains a stochastic trend that is imparted by (and therefore cointegrates with) the time-series data of radiative forcings, such as the atmospheric concentration of greenhouse gases and solar irradiance. Proponents of this view include Stern and Kaufmann (2000); Kaufmann and Stern (2002); Kaufmann et al. (2006); Mills (2009); Kaufmann et al. (2010, 2013). The reader is referred to Estrada and Perron (2013) for a useful overview.

Table 2.1: Data sources

Variable	Product	Description	Period
GMST	HadCRUT4 ^a	Global mean surface temperature. Primary series. Compiled by the UK Met Office and the Climatic Research Unit at the University of East Anglia.	1850–2013
	CW2014 ^b	Secondary series. Compiled by Cowtan and Way (2014). Corrects for coverage bias in HadCRUT4.	1850–2013
	GISTEMP ^c	Secondary series. Compiled by the NASA Goddard Institute for Space Studies.	1880–2013
RF	RCP ^d	Total radiative forcing due to anthropogenic and natural factors (excluding volcanic aerosols). Compiled by Meinshausen et al. (2011). Historical data until 2005, simulated scenarios thereafter.	1750–2300
AER	RCP ^d	Radiative forcing due to volcanic stratospheric aerosols. Compiled by Meinshausen et al. (2011).	1750–2005
AMO	NOAA ^e	Atlantic Multidecadal Oscillation.	1856–2005
SOI	NCAR ^f	Southern Oscillation Index.	1866–2005

^a <http://www.metoffice.gov.uk/hadobs/hadcrut4/data/current/download.html>

^b <http://www-users.york.ac.uk/~kdc3/papers/coverage2013/series.html>

^c <http://data.giss.nasa.gov/gistemp/>

^d <http://www.iiasa.ac.at/web-apps/tnt/RcpDb>

^e <http://www.esrl.noaa.gov/psd/data/timeseries/AMO/>

^f <http://www.cgd.ucar.edu/cas/catalog/climind/soi.html>

Cowtan and Way (2014) and the other by the NASA Goddard Institute for Space Studies (GISTEMP) — are used as a check against coverage issues and other uncertainties. Radiative forcing data (1765–2005) are taken from the Representative Concentration Pathway (RCP) database, hosted by the Potsdam Institute for Climate Impact Research. These data include anthropogenic sources of radiative forcing like industrial greenhouse gas emissions, as well as natural sources like solar irradiance and volcanic eruptions. It is important to note that the forcings are defined in terms of a common unit, Watts per square metre (Wm^{-2}). This allows for aggregation into a composite series of total radiative forcing, which circumvents the attributional problems that we would otherwise encounter due to severe multicollinearities among the different series of greenhouse gases. Data for two major oceanic-atmospheric phenomena, the Atlantic Multidecadal Oscillation (AMO, 1856–2013) and the Southern Oscillation Index (SOI, 1866–2013), are taken from the U.S. National Oceanic and Atmospheric Administration (NOAA) and National Center for Atmospheric Research (NCAR). Summarising the common historic dataset, we have 140 annual observations running over 1866–2005.

2.3 Econometric approach

2.3.1 Bayesian regression overview

The Bayesian regression framework is less familiar to many researchers than the frequentist paradigm that is commonly taught in universities. For this reason, I provide a brief overview of the key principles of the Bayesian method and highlight some important distinctions versus the frequentist approach. A more in-depth discussion may be found in Koop (2006), Gill (2007), Albert (2009), Jackman (2009), and Kruschke (2014) among others.

A Bayesian regression model uses the logical structure of Bayes' theorem to estimate probable values of a set of parameters θ , given data X :

$$p(\theta|X) = \frac{p(X|\theta)p(\theta)}{p(X)}. \quad (2.1)$$

Here, $p(\theta|X)$ is known as the *posterior density* and serves as the fundamental criterion of interest in the Bayesian framework. The posterior asks, "What are the probable values of our parameters, given the observed data?" This stands in direct contrast to the first term in the numerator, $p(X|\theta)$, which is the familiar *likelihood function* that we recognise from frequentist statistics. The likelihood essentially reverses the question posed by the posterior; it instead evaluates how likely we are to observe some data for a given set of parameters (e.g. under certain distributional assumptions). The second term in the numerator, $p(\theta)$, is the *prior density* that one assigns to values of θ . The choice of prior can prove to be a contentious issue in Bayesian analysis and is often given primacy as a result. For the moment, it will suffice to say that the prior should encapsulate our knowledge about the parameters before we have analysed the current dataset. Insofar as we are interested in learning about θ , it is common practice to ignore the term in the denominator, $p(X)$. This is simply the marginal probability of the data and can be thought of as a normalisation constant, which helps to ensure that the posterior is a proper probability distribution (i.e. integrates to one) and can be calculated *ad hoc* if needed. For this reason, equation (2.1) is typically re-written as

$$p(\theta|X) \propto p(X|\theta)p(\theta). \quad (2.2)$$

Equation (2.2) embodies the mantra of Bayesian statistics: "The posterior is pro-

portional to the likelihood times the prior.” Solving for the posterior typically involves the combination of various integrals, which cannot be calculated analytically.³ Fortunately, the advent of increased computing power and software allows us to simulate the posterior density with relative ease using Markov Chain Monte Carlo (MCMC) routines. This can be done for virtually any manner of prior and data/likelihood combination. The upshot is that obtaining a valid posterior is simply a matter of (i) choosing a prior distribution for our regression parameters, i.e. regression coefficients and variances, and (ii) specifying a likelihood function to fit the data. For ease of exposition — how we map parameter values to beliefs about TCR will be determined by the specification of the regression model — I begin with the likelihood function.

2.3.2 Likelihood function

Consider the regression equation

$$GMST_t = \alpha_0 + \beta_1 RF_t + \gamma_2 AER_t + \delta_3 SOI_t + \eta_4 AMO_t + \epsilon_t, \quad (2.3)$$

where $GMST$ is the global mean surface temperature anomaly relative to the pre-industrial period (here defined as the 1851–1880 average), RF is total radiative forcing due to both anthropogenic and natural factors (excluding stratospheric aerosols), AER is the radiative forcing due to the release of stratospheric aerosols during volcanic eruptions, and SOI and AMO are scaled indices of these respective climatic phenomena. The subscript t denotes time. The simple OLS specification above believes the fact that it is a statistically valid regression model for estimating global temperature (Estrada and Perron, 2012; Estrada et al., 2013a). That being said, lumping natural and anthropogenic forcings together within the same composite RF term may seem a puzzling choice in the context of present paper. After all, our goal is to disentangle scepticism about the human role in climate change. And yet, this is a necessary step to ensure that the model does not become unphysical. The underlying forcings in the dataset are expressed in terms of a common unit (Wm^{-2}) and the model must therefore constrain them to have the same effect on temperature.⁴ Ignoring such physical constraints — e.g. by entering solar irradiance as a separate

³A well-known exception occurs in the case of so-called conjugate priors, but this places severe restrictions on type of questions that can be asked of the data.

⁴The forcing imparted by stratospheric aerosols, on the other hand, is of a more transitory nature. This explains why AER may be included as a separate component to RF in the regression equation (Estrada et al., 2013a).

regressor and further imposing subjective priors on the anthropogenic components of the model — nonetheless leads to very similar conclusions as the more correct formulation. Similarly, while it is possible to incorporate dynamic elements into the model, doing so would create interpretative complications. For example, running a dynamic specification where the lagged dependent variable is included as a regressor yields virtually identical results to the present set-up, once long-run effects are calculated. Such an approach would however make it difficult (if not impossible) to establish consistent beliefs about climate sensitivity.⁵ A static formulation allows us to avoid these inconsistencies at no real cost to model fit.

Equation (2.3) implies a likelihood function that is multivariate normal,

$$p(GMST|\boldsymbol{\beta}, \sigma^2) = \frac{1}{(2\pi\sigma^2)^{T/2}} \exp \left[-\frac{(GMST - \mathbf{X}\boldsymbol{\beta})'(GMST - \mathbf{X}\boldsymbol{\beta})}{2\sigma^2} \right], \quad (2.4)$$

where \mathbf{X} is the design matrix of explanatory variables; $\boldsymbol{\beta}$ is the coefficient vector; $\sigma^2 = \text{Var}(\epsilon)$ is the variance of the error term; and $T = 140$ is the observed number of time periods. Equation (2.4) could also be written as $GMST|\boldsymbol{\beta}, \sigma^2 \sim \mathcal{N}_T(\mathbf{X}\boldsymbol{\beta}, \sigma^2\mathbf{I})$. Further discussion of the multivariate normal likelihood function, and how it might combine with specific priors to yield a valid posterior, is provided in Appendix 2.A.

An important feature of equations (2.3) and (2.4) is that they define how we should map probabilities about the regression parameters to beliefs about climate sensitivity. Recall that TCR describes the contemporaneous change in temperature that will accompany a steady doubling of atmospheric CO₂ concentrations. It follows trivially that

$$\text{TCR} = \beta_1 * F_{2\times}, \quad (2.5)$$

where β_1 is the regression coefficient describing how responsive global temperatures are to a change in total radiative forcing; and $F_{2\times}$ is the change in forcing that results from a doubling of CO₂. For the latter, I use the IPCC's best estimate of

⁵Recall that the long-run multiplier of any variable in a ADL(p, q) model is determined by the coefficients on its own lags, as well as those on the dependent variable. Assuming we used an ADL(1,0) model, for instance, the long-run effect of a change in total radiative forcing would be $\beta_{LR} = \frac{\beta_1}{1-\theta}$, where θ is the coefficient on $GMST_{t-1}$. However, defining strong priors over β_1 can affect the posterior estimate for other coefficients (including θ) as a simple consequence of best-fit mechanics. Since both parameters are necessary to calculate TCR in such a dynamic setting — yet both are also varying — it will in general not be possible to make meaningful or consistent comparisons between prior and posterior beliefs.

$F_{2\times} = 3.71 \text{ Wm}^{-2}$ and further assume an additional $\pm 10\%$ variation to account for uncertainties over spatial heterogeneity and cloud formation (Schmidt, 2007; IPCC, 2001, Chapter 6).⁶ The key point is that assigning a distribution over the parameter β_1 will necessarily imply a distribution for TCR, and vice versa. We therefore have a direct means of linking prior and posterior probabilities of the regression parameters to beliefs about TCR. It also means that the primary goal of the regression analysis will be to determine probable values of β_1 . The rest of the parameters will take a backseat in the analysis that follows, acting largely as controls.

2.3.3 Priors

Climate scepticism is a matter of degree. I account for this by defining a simple matrix of stylised sceptic priors as per Table 2.2. Rows in panel (A) differentiate climate sceptics according to their best estimate of TCR. A *lukewarmer's* prior for TCR is centred around a $1 \text{ }^\circ\text{C}$ mean, while a *denier's* is distributed around a $0 \text{ }^\circ\text{C}$ mean. A lukewarmer believes that humans are having an effect on the climate, but to a lesser extent than the mainstream consensus. A denier, on the other hand, believes that there is probably no relationship between human activity and global temperatures. Columns in panel (A) describe varying levels of prior certainty in the form of variance. A “moderate” level of certainty corresponds to a belief that the true value of TCR lies within a $1 \text{ }^\circ\text{C}$ range of one’s best estimate with 95% probability. In contrast, a person with “strong” certainty about their prior feels that a probability range of just $0.25 \text{ }^\circ\text{C}$ is appropriate. Following equation (2.5), obtaining priors over β_1 is simply a matter of dividing the respective TCR distributions by $F_{2\times} = 3.71 \text{ Wm}^{-2}$. These are the parameters that actually enter the Bayesian regression model and are shown in panel (B) of Table 2.2.

In addition to the subjective priors of our stylised sceptics, a noninformative prior provides a useful reference case for the analysis. Loosely speaking, noninformative priors are vague and should not privilege any parameter values before one sees the data. Candidates include uniform priors, a normally distributed prior with large variance, Jaynes’ maximum entropy prior, or the Jeffreys class of invariant prior. However, uniform priors have been fiercely criticized in the context of climate change attribution (Annan and Hargreaves, 2011; Pueyo, 2012; Lewis, 2013).

⁶It is worth noting that a number of studies which rely on time-series methods to derive an estimate of climate sensitivity — e.g. Kaufmann et al. (2006); Mills (2009); Estrada and Perron (2012) — do so under the assumption that $F_{2\times} = 4.37 \text{ Wm}^{-2}$. This outdated figure appears to be based on early calculations by Hansen et al. (1988). The climate sensitivity estimates of these studies may consequently be regarded as inflated.

Table 2.2: Sceptic priors

(A) TCR (°C)		
<i>Type</i>	Moderate	Strong
Lukewarmer	$\mathcal{N}(1, 0.250^2)$	$\mathcal{N}(1, 0.065^2)$
Denier	$\mathcal{N}(0, 0.250^2)$	$\mathcal{N}(0, 0.065^2)$

(B) Implied β_1		
<i>Type</i>	Moderate	Strong
Lukewarmer	$\mathcal{N}(0.270, 0.0674^2)$	$\mathcal{N}(0.270, 0.0175^2)$
Denier	$\mathcal{N}(0, 0.0674^2)$	$\mathcal{N}(0, 0.0175^2)$

Notes: Priors are based on an assumed normal distribution over the transient climate response (TCR). Rows denote an individual’s mean, or best, estimate of TCR. *Lukewarmers* are assumed to lie at the very low end of the IPCC’s likely range (1–2.5 °C), while *Deniers* expect TCR to be somewhere around zero. Columns denote uncertainty in the form of variance. A person with “moderate” convictions believes that the true value of TCR lies within a 1 °C interval of their prior mean (95% probability), while that interval falls to just 0.25 °C for someone with “strong” convictions. The implied priors for the regression parameter β_1 in panel (B) are obtained using the simple formula described in equation (2.5), i.e. $\beta_1 = \text{TCR}/3.71$.

I therefore favour a standard noninformative prior for the regression parameters, such that $g(\theta, \sigma^2) \propto \frac{1}{\sigma^2}$. Nonetheless, experimenting with other options shows that this choice does not materially affect the conclusions that we may draw from the analysis.

Two points merit further discussion before continuing on to the posterior results. The first is that our group of sceptics only hold subjective priors about TCR and, thus, β_1 . Noninformative priors are always assumed for the remaining parameters in the regression equation. The second point is to acknowledge that these sceptics are little more than stylised caricatures. Their priors are simply taken as given. We are not concerned with where these priors come from and why they are of a particular strength. However, such abstractions allow us to explore the way in which climate sceptics might interpret evidence for climate change. Moreover, it gives a sense of just how strong someone’s prior beliefs need to be, so as to preclude their acceptance of any policy interventions.

2.4 Results

2.4.1 Posterior regression results and TCR distributions

The posterior regression results for the various prior types are presented in Table 2.3. Beginning with the noninformative case in column (1), it is reassuring to note that all of the regression coefficients are credibly different from zero and of the anticipated sign. For example, global mean surface temperature is negatively correlated with *SOI*. This is to be expected since the El Niño phenomenon is defined by *SOI* moving into its negative phase. The posterior density of our main parameter of interest, the coefficient on *RF*, shows that global temperature will rise by an average of $0.4\text{ }^{\circ}\text{C}$ for every Wm^{-2} increase in total radiative forcing. Of greater interest, however, is the fact that the posterior estimates yielded by the group of sceptical priors in columns (2)–(5) are very similar to this noninformative case. With the exception of the Strong Denier, there is a clear tendency to congregate towards the noninformative parameter values.

Of course, the exact values of the regression parameters are themselves of somewhat limited interest. Rather, their primary usefulness is to enable the recovery of posterior beliefs about TCR. These are plotted in Figure 2.1. We see that the posterior TCR distributions are generally clustered around a best estimate of $1.5\text{ }^{\circ}\text{C}$, with a 95% credible interval somewhere in the region of $1.3\text{--}1.7\text{ }^{\circ}\text{C}$ depending on the prior. (See Table 2.B.1 in the Appendices for more information.) Figure 2.1 also makes clear that, excepting the Strong Denier, the posterior beliefs of the various sceptics fall comfortably within the IPCC “likely” range. However, the derived probability intervals are decidedly narrower and TCR values at the upper end of the spectrum are discounted accordingly. That being said, the HadCRUT4 record is known to suffer from slight biases due to incomplete coverage of *in situ* thermometer readings. I therefore re-run the Bayesian regression model using two alternate reconstructions of global temperature data as a check against this issue. Cowtan and Way (2014), hereafter CW2014, correct for the HadCRUT4 biases using an interpolation algorithm based on the “kriging” method. Similarly, the NASA Goddard Institute for Space Studies uses an extrapolation method to overcome coverage bias in GISTEMP, its own reconstruction of global surface temperatures. Re-running the model with these alternate series yields moderately higher TCR values. Under a noninformative prior, the posterior TCR means and 95% credible intervals are $1.6\text{ }^{\circ}\text{C}$ ($1.4\text{--}1.7\text{ }^{\circ}\text{C}$) for CW2014, and $1.7\text{ }^{\circ}\text{C}$ ($1.5\text{--}1.9\text{ }^{\circ}\text{C}$) for GISTEMP; see Table 2.B.2 in the appendices. While I omit them for brevity, the posterior results for the group

Table 2.3: Posterior regression results

Dependent variable: Global mean surface temperature anomaly (°C)

	Noninformative (1)	Moderate Lukewarmer (2)	Strong Lukewarmer (3)	Moderate Denier (4)	Strong Denier (5)
Total radiative forcing	0.415 [0.389, 0.441]	0.409 [0.384, 0.434]	0.359 [0.335, 0.381]	0.399 [0.374, 0.425]	0.109 [0.070, 0.148]
Stratospheric aerosols	0.047 [0.004, 0.091]	0.047 [0.004, 0.091]	0.047 [0.001, 0.093]	0.047 [0.004, 0.090]	0.042 [-0.055, 0.138]
SOI	-0.028 [-0.039, -0.016]	-0.028 [-0.040, -0.016]	-0.032 [-0.044, -0.019]	-0.029 [-0.041, -0.017]	-0.049 [-0.077, -0.023]
AMO	0.481 [0.407, 0.556]	0.481 [0.407, 0.555]	0.479 [0.400, 0.557]	0.481 [0.406, 0.556]	0.469 [0.299, 0.637]
Constant	-0.110 [-0.131, -0.088]	-0.106 [-0.127, -0.084]	-0.072 [-0.092, -0.050]	-0.099 [-0.120, -0.077]	0.097 [0.059, 0.138]

Notes: Mean values are given, with square brackets denoting 95% credible intervals. Columns are distinguished by different sets of prior covariates. Column (1) specifies noninformative priors over all regression parameters. Columns (2)–(5) each take a unique prior over the parameter β_1 , which is the coefficient pertaining to total radiative forcing — compare with Table 2.2. For the remaining parameters, noninformative priors are again assumed. The dataset consists of 140 annual observations over the period 1866–2005. Total radiative forcing (comprising anthropogenic forces and solar irradiance) and stratospheric aerosols are measured in Wm^{-2} . The Southern Oscillation Index (SOI) and Atlantic Multidecadal Oscillation (AMO) are measured as scaled indices.

of climate sceptics are similarly nudged higher towards the new noninformative distributions. Given that my explicit goal in this paper is to evaluate policy options from the perspective of climate sceptics, I continue to use the results from the Had-CRUT4 series as a default. Yet, it should be noted that this is a conservative choice that is likely to (at least marginally) understate the true level of warming.

Further insight into the updating behaviour of our stylised sceptics is provided by the recursive TCR estimates shown in Figure 2.2. It is apparent that stronger convictions about one’s prior beliefs have a greater dampening effect on posterior outcomes than the prior mean. The Moderate Denier converges quicker to the non-informative distribution than the Strong Lukewarmer, for example. That being said, certain sceptics will only converge to the noninformative distribution after “seeing” data from a number of decades. This does not alter the conclusions that we are able to draw from the Bayesian analysis.⁷ However, it does highlight the importance of using all of the available instrumental climate data for building any kind of policy

⁷As long as we have fully specified a prior that encapsulates someone’s initial beliefs, then we should in principle treat the full historical dataset as new information for updating those beliefs. As a corollary, concerns over the use of the full historical dataset would only hold sway in cases where priors already incorporate information that has been obtained from applying the model on a sub-sample of the dataset. In that case, we would need to exclude the sub-sample from the analysis to derive a valid posterior that avoids double counting.

Table 2.4: Covariate vectors for prediction in the year 2100

	RCP 2.6 420 ppmv CO ₂	RCP 4.5 540 ppmv CO ₂	RCP 6.0 670 ppmv CO ₂	RCP 8.5 940 ppmv CO ₂
RF_{2100}	2.626	4.281	5.522	8.340
<i>Due to CO₂</i>	85%	83%	86%	78%
<i>Due to solar</i>	7%	4%	3%	2%
\overline{AER}	0.017	0.017	0.017	0.017
\overline{SOI}	-0.079	-0.079	-0.079	-0.079
\overline{AMO}	-0.002	-0.002	-0.002	-0.002

Notes: Covariates are used to predict the global mean surface temperature anomaly in the year 2100. The Representative Concentration Pathways (RCPs) are a family of forcing scenarios developed for the IPCC (Van Vuuren et al., 2011). Each RCP has a core component of atmospheric CO₂ concentrations, measured in parts per million volume (ppmv). With regard to the covariates in the regression model, total radiative forcing (RF) and stratospheric aerosols (AER) are measured in Wm^{-2} . The Southern Oscillation Index (SOI) and Atlantic Multidecadal Oscillation (AMO) are measured as scaled indices. Future values for RF are taken from the RCP database. For the rest, historical mean values are used.

The point of using the RCPs is not to interpret them as precise forecasts far into the future. Rather, they serve as allegories to aid our understanding about the relative risks associated with different states of world. It is worth emphasising that each RCP presumes significantly different policy foundations. RCPs 2.6 and 4.5 correspond to lower emissions scenarios that would only be possible in the presence of an appropriately levied carbon price, global treaties that limit carbon leakage between regions, etc. In other words, these two scenarios take climate policy as antecedent to their outcomes. RCPs 6.0 and 8.5 do not incorporate such assumptions and instead correspond to what we might call “business as usual” paths — insofar as they inexorably relax assumed constraints on fossil fuel consumption. These scenarios are free of the regulations needed to drive emissions mitigation on a meaningful scale.

Figure 2.3 shows the temperature evolution for each RCP under the noninformative case, which we again take as the benchmark. (The figure also reaffirms the excellent model fit to historic observations.) The principal message is that CO₂ concentrations must be constrained to RCP 4.5 or lower, if we are to avoid a 2 °C rise in global temperature. Given the prominence of this particular threshold in international climate treaties and the popular narrative, the result is a reinforcement of commonly cited emissions targets such as 450 and 540 ppmv. On the other hand, we can expect to even breach 3 °C by the year 2100 if we continue along a truly unconstrained emissions path à la RCP 8.5.

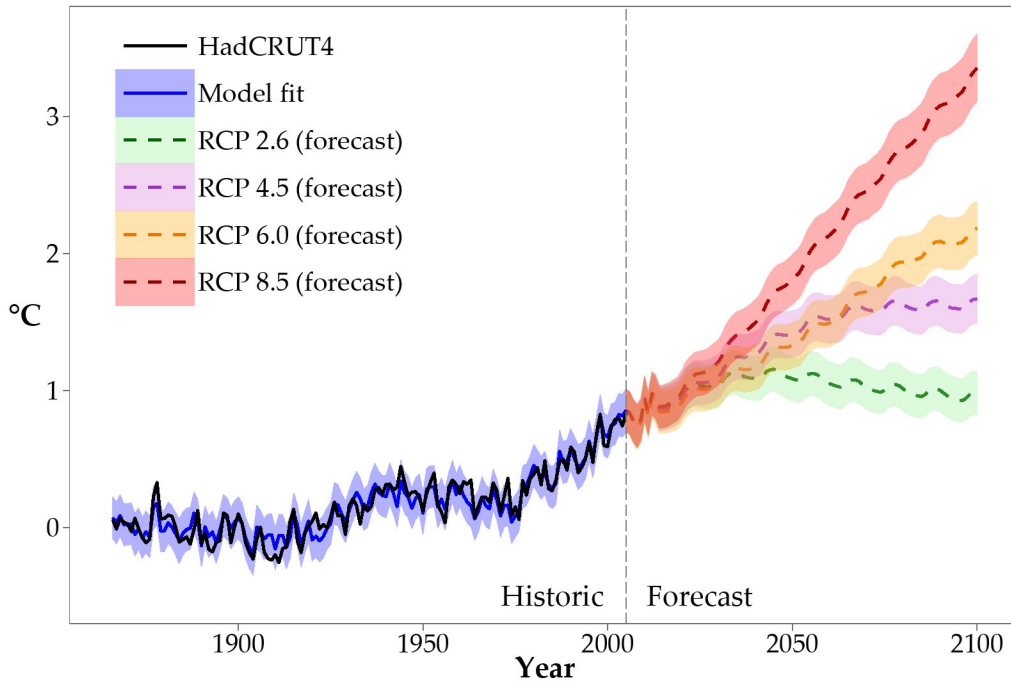


Figure 2.3: Model fit and prediction: Noninformative priors

Notes: Temperature anomaly relative to the 1851–1880 average. Shaded areas denote 95% credible intervals.

What of the predictions yielded by our group of climate sceptics? While it is straightforward to redraw Figure 2.3 for each prior type, a more intuitive comparison can be made by looking at the total warming that each person expects by the end of the century. Figure 2.4 plots the predictive temperature density in the year 2100 for all prior types by RCP scenario. Again, the data have a clear tendency to overwhelm even reasonably staunch forms of climate scepticism. Nearly all of the stylised sceptics would expect to breach the 2 °C threshold by 2100 under RCP 6.0, while a temperature rise of more than 3 °C is likely under under RCP 8.5. An exception can only be found in the form of the Strong Denier, whose extreme prior dominates the posterior in a way that obviates nearly all concern about large temperature increases.

2.4.3 Welfare implications and the social cost of carbon

Provided they consider enough data, even climate sceptics would seemingly agree that a 2 °C target requires limiting CO₂ concentrations to 540 ppmv. Of course, whether individuals actually subscribe to policy measures aimed at achieving this goal is dependent on many things; their choice of discount rate, beliefs about the

Table 2.5: Social cost of carbon (US\$2005 per tonne)

	Mean	Median	95% Probability interval
Noninformative	55.76	39.25	[11.32, 158.22]
Moderate Lukewarmer	71.14	37.61	[10.98, 154.54]
Strong Lukewarmer	39.64	25.70	[7.10, 103.12]
Moderate Denier	55.64	34.95	[10.10, 144.55]
Strong Denier	1.85	1.54	[-0.17, 5.54]

Results for each agent type are obtained from 10,000 simulation runs of PAGE09 (Hope, 2011a). Posterior TCR distributions serve as key inputs to the model, while the remaining parameters are set to the PAGE09 model defaults.

parameters (Hope, 2011a,b).

Table 2.5, together with Figure 2.B.2 in the appendices, present probability distributions for the SCC across all prior groups in 2005 US dollars. In each case, the model has been simulated 10,000 times to produce accurate estimates. The resulting distributions are highly skewed and characterised by extremely long upper tails. This is partly due to the fact that PAGE09, like most IAMs, assumes that economic damages are convex in temperature. In addition, PAGE09 allows for the possibility of large-scale discontinuities (e.g. melting of the Greenland ice sheet) at temperatures above 3 °C . Such low probability, high impact events would cause vast economic losses and yield some extreme SCC values as a consequence (Hope, 2011a,b). For this reason, I provide both the mean and median SCC values alongside the 95% probability interval.

Excepting the Strong Denier, the SCC for all prior types is comfortably larger than zero. The mean value ranges from \$40 to \$71 per tonne (2005 prices), while the 95% probability interval extends from around \$10 to upwards of \$100 per tonne.⁹ These results are consistent with the SCC estimates found within the literature. For example, an influential synthesis report by the United States Government’s Interagency Working Group on the Social Cost of Carbon (IWG, 2010, 2013) finds a mean SCC value of \$11–\$52 per tonne (2007 prices), depending on the preferred discount rate. The encouraging point from a policy perspective is that such congruency exists in spite of our deliberate consideration of climate scepticism. Another way to frame the SCC estimates presented here is to imagine that each prior type represents an equal segment of a voting population. We would then expect to see broad support for a carbon tax of at least \$20–\$25, if not double that. While such a thought experiment clearly abstracts from the many complications that would arise from

⁹The Moderate Lukewarmer obtains the highest mean value (higher than even the noninformative case), because of the larger variance attached to his posterior TCR estimate.

free-riding and so forth, again we see that nominal climate scepticism does not correspond to a mechanical dismissal of climate policy.

2.5 Discussion

We have seen that a non-trivial carbon price is consistent with a range of contrarian priors once we allow for updating of beliefs and, crucially, consider enough of the available data. An optimist might interpret these findings as a sign that common ground on climate policy is closer than many people think. On the other hand, they may also help to explain why the policy debate is so polarised in the first place. As all intermediate positions are absorbed into the mainstream, only the most hardcore sceptics will remain wedded to their priors. This group is unlikely to brook any proposals for reduced carbon emissions and no amount of new information will convince them otherwise. In that case, and given the persistent scepticism that one sees in opinion polls about climate change (e.g. Saad, 2014), it becomes reasonable to ask whether real-life sceptics (i) actually hold such extreme views, and (ii) are numerous enough to prevent political action. Such considerations are reinforced by the idealized nature of the analysis until now. Irrespective of the scientific merit of working through such a set-up, clearly normal people do not update their priors in lockstep with a Bayesian regression model, supported by large dataset of time-series observations.

A natural starting point for thinking about these issues is to take a closer look at the mechanisms underlying posterior agreement formation. The notion that partisans should converge toward consensus with increasing information has long been taken as a logical consequence of Bayes' theorem. Indeed, empirical evidence to the contrary has been cited as a weakness of the Bayesian paradigm and its relevance to real-life problems (Kahneman and Tversky, 1972). This is a misconception. Nothing in the Bayesian paradigm precludes the possibility of diverging opinions in the face of shared information. It may even be the case that the same information has a polarising effect on individuals, in that it pushes them towards opposite conclusions. Bullock (2009) discusses some of the conditions that will prevent partisans from reaching consensus on political issues, even when these agents are assumed to update their beliefs in a fully Bayesian manner.¹⁰ A more general framework is provided by Jaynes (2003), who shows that the perceived *trustworthiness* of new

¹⁰Such conditions include whether new evidence falls in between or outside the set of assigned prior probabilities, and whether partisans are learning about a variable that may itself be changing over time (e.g. a third party's support for some political proposal).

information is the key determinant for whether Bayesian agents converge toward consensus. Jaynes evinces the need to broaden our understanding of “prior” so that it incorporates not only our existing beliefs about some phenomenon S , but also our incredulity about any new data source that claims to tell us something about S . He provides several amusing examples, from extrasensory perception to drug safety scares, to drive home his point. Here I adapt one such example to our present problem:

Consider three individuals who hold different prior beliefs about climate change. Al is a “warmist”, Bob is a “lukewarmer” and Christie is a “denier”. Let us say that these labels are encapsulated by the probabilities that each person assigns to climate sensitivity S . For simplicity, we assume that only two states of the world are possible such that climate sensitivity is either high or low: $S \in S_L, S_H$. Denote by I an individual’s existing information about the world. Then, indexing by the first letter of a person’s name, the prior probability that each person assigns to a high climate sensitivity is $P(S_H|I_A) = 0.90$, $P(S_H|I_B) = 0.40$, and $P(S_H|I_C) = 0.10$.

Now imagine that the IPCC issues a statement in which it claims that climate sensitivity is high. How do the three individuals respond to this new data, $D = D_H$? The answer depends crucially on the regard that each person holds for the IPCC itself. Suppose we have the following situation. All three agree that the IPCC would present data supporting a high climate sensitivity if that was the true state of the world, i.e. $P(D_H|S_H, I_A) = P(D_H|S_H, I_B) = P(D_H|S_H, I_C) = 1.00$. However, they disagree on whether the IPCC can be trusted to disavow the high sensitivity hypothesis if the science actually supported a low climate sensitivity. In this case, we have $P(D_H|S_L, I_A) = 0.05$, $P(D_H|S_L, I_B) = 0.89$, and $P(D_H|S_L, I_C) = 0.05$. Putting it into words, Al and Christie hold wildly different priors about climate sensitivity. Yet, they both regard the IPCC as an upstanding institution which can be trusted to accurately represent the science of climate change. In contrast, Bob is not only somewhat sanguine about expected future warming; he is dubious about the very motives of the IPCC. He is, in effect, a cynic who believes that the IPCC is willing to lie in advancement of its agenda.

Recovering posterior beliefs about climate sensitivity is then a simple matter of modifying Bayes’ theorem to account for each person’s (dis)trust in the IPCC. For Al, we have

$$\begin{aligned}
P(S_H|D_H, I_A) &= \frac{P(D_H|S_H, I_A)P(S_H|I_A)}{P(D_H|S_H, I_A)P(S_H|I_A) + P(D_H|S_L, I_A)P(S_L|I_A)} \\
&= \frac{1.0 \times 0.9}{1.0 \times 0.9 + 0.05 \times 0.1} \\
&\approx 0.98.
\end{aligned}$$

Similarly, we get $P(S_H|D_H, I_B) \approx 0.43$ for Bob and $P(S_H|D_H, I_C) \approx 0.69$ for Christie. Comparing the prior and posterior beliefs of these three individuals, Al is even more of a believer in the high sensitivity hypothesis than before, having raised his subjective probability for S_H from 90% to 98%. Christie has experienced a still greater effect and has updated her subjective probability for S_H from 10% to 69%. In other words, she now attaches a larger probability to the high sensitivity hypothesis than the low sensitivity alternative. But what of cynical Bob? Well, it appears that he has not been swayed by the IPCC report in the slightest; both his prior and posterior probabilities suggest that S_H only has an approximately 40% chance of being true. Bob's extreme mistrust has effectively led him to discount the IPCC's high sensitivity claim in its entirety.

Extending the above framework to account for increasing granularity is conceptually straightforward. The principal insight remains the same. Trust is as much a determinant of whether beliefs are amenable to data, and whether individuals converge towards consensus, as the precision of the data itself. The relevance to the climate change debate is clear, particularly given evidence of scientific distrust among certain segments of the population (Malka et al., 2009; Krosnick and MacInnis, 2010; Gauchat, 2012; Leiserowitz et al., 2013; Fiske and Dupree, 2014; Hmielowski et al., 2014). Equally importantly, the modified Bayesian framework offers a bridge between competing explanations of climate scepticism as a phenomenon. For instance, the "deficit model" posits a lack of scientific knowledge and understanding as key drivers of scepticism, whereas advocates of the "cultural cognition" theory argue that group identity and value systems are more relevant (Clark et al., 2013; Kahan et al., 2011, 2012). A Bayesian model that incorporates perceptions of source credibility is able to accommodate both camps. Exposure to new scientific evidence can ameliorate a person's scepticism, but only if their priors allow for it. This includes whether factors such as cultural identity cause them to discount some sources of information more than others. Disentangling the root causes of such "information immunity" — i.e. whether sceptics do not update their beliefs because they are extremely sure of their priors, distrustful of climate experts,

or some combination thereof — remains an important area for future research.

2.6 Concluding remarks

The goal of this paper has been to explore the way in which prior beliefs affect our willingness to mitigate against climate change. The Bayesian paradigm provides a natural analytical framework and I have proposed a group of stylised sceptics to embody the degrees of climate scepticism that one encounters in the real world. The primary finding is that subjective priors are generally overwhelmed by the empirical evidence for climate change. Once they have updated their beliefs in accordance with the available data, most sceptics demonstrate a clear tendency to congregate towards the noninformative case that serves as an objective reference point for this study. Depending on the preferred reconstruction of global temperatures, the 95% posterior probability range for TCR under a noninformative prior is somewhere between 1.4 °C and 1.9 °C. This implies a tighter bound on climate sensitivity than suggested by the IPCC (2013), whose “likely” range for TCR is 1.0–2.5 °C. While such a finding would offer comfort against the most alarming rates of future warming, it should not be taken as evidence against the need for climate mitigation. Indeed, the updated beliefs of our various sceptics are shown to be consistent with a carbon price that is substantially greater than zero. I obtain a mean SCC range of approximately \$40–\$70 per tonne using the PAGE09 model of Hope (2011a). Only those with extreme *a priori* sceptic beliefs will find themselves in disagreement, or feel any confidence in the notion that unfettered emissions growth will *not* lead to sizeable future warming. Taken together, the results suggest that a rational climate sceptic — even one that holds relatively strong prior beliefs — could embrace policy measures to constrain CO₂ emissions once they have seen all of the data. These findings may ultimately be seen as a reflection of the compelling case for man-made climate change provided by the available evidence.

References

- Albert, J. (2009). *Bayesian Computation with R*, Springer, 2nd edition.
- Allen, Myles R. and David J. Frame (2007). "Call Off The Quest," *Science*, 318(5850): 582–583.
- Anderegg, William R. L., James W. Prall, Jacob Harold, and Stephen H. Schneider (2010). "Expert credibility in climate change," *Proceedings of the National Academy of Sciences of the United States of America (PNAS)*, 27: 12107–12109.
- Annan, J.D. and J.C. Hargreaves (2011). "On the generation and interpretation of probabilistic estimates of climate sensitivity," *Climatic Change*, 104(3-4): 423–436.
- Bullock, John G. (2009). "Partisan bias and the Bayesian ideal in the study of public opinion," *The Journal of Politics*, 71(03): 1109–1124.
- Clark, Dav, Michael Andrew Ranney, and Jacqueline Felipe (2013). "Knowledge Helps: Mechanistic Information and Numeric Evidence as Cognitive Levers to Overcome Stasis and Build Public Consensus on Climate Change," in N. Sebanz & I. Wachsmuth M. Knauff, M. Pauen (ed.), "Proceedings of 35th Annual Meeting of the Cognitive Science Society," pp. 2070– 2075.
- Cook, John, Dana Nuccitelli, Sarah A. Green, Mark Richardson, Bärbel Winkler, Rob Painting, Robert Way, Peter Jacobs, and Andrew Skuce (2013). "Quantifying the consensus on anthropogenic global warming in the scientific literature," *Environmental Research Letters*, 8(2): 024024.
- Corner, Adam, Lorraine Whitmarsh, and Dimitrios Xenias (2012). "Uncertainty, scepticism and attitudes towards climate change: biased assimilation and attitude polarisation," *Climatic Change*, 114(3–4): 463–478.
- Cowtan, Kevin and Robert G. Way (2014). "Coverage bias in the HadCRUT4 temperature series and its impact on recent temperature trends," *Quarterly Journal of the Royal Meteorological Society*, 140(683): 1935–1944.

- Doran, Peter T. and Maggie K. Zimmerman (2011). "Examining the Scientific Consensus on Climate Change," *Eos, Transactions, American Geophysical Union (EOS)*, 90(3): 22–23.
- Estrada, Francisco, Pierre Perron, , and Benjamín Martínez-López (2013a). "Statistically derived contributions of diverse human influences to twentieth-century temperature changes," *Nature Geoscience*, Advance Online Publication: 1–6.
- Estrada, Francisco and Pierre Perron (2012). "Breaks, trends and the attribution of climate change: a time-series analysis." Working paper, Department of Economics, Boston University. Available: <http://people.bu.edu/perron/workingpapers.html>.
- Estrada, Francisco and Pierre Perron (2013). "Detection and attribution of climate change through econometric methods." Working paper, Department of Economics, Boston University. Available: <http://people.bu.edu/perron/workingpapers.html>.
- Estrada, Francisco, Pierre Perron, Carlos Gay-García, and Benjamín Martínez-López (2013b). "A Time-Series Analysis of the 20th Century Climate Simulations Produced for the IPCC's Fourth Assessment Report," *PLOS ONE*, 8(3): e60017.
- Fiske, Susan T. and Cydney Dupree (2014). "Gaining trust as well as respect in communicating to motivated audiences about science topics," *Proceedings of the National Academy of Sciences*, 111(4): 13593–13597.
- Gauchat, Gordon (2012). "Politicization of science in the public sphere a study of public trust in the United States, 1974 to 2010," *American Sociological Review*, 77(2): 167–187.
- Gay-Garcia, Carlos, Francisco Estrada, and Armando Sánchez (2009). "Global and hemispheric temperatures revisited," *Climatic Change*, 94(3-4): 333–349.
- Gill, Jeff (2007). *Bayesian Methods: A Social and Behavioral Sciences Approach*, Chapman & Hall / CRC, 2nd edition.
- Hansen, I. Fung, J., A. Lacis, D. Rind, S. Lebedeff, R. Ruedy, G. Russell, and P. Stone (1988). "Global climate changes as forecast by Goddard Institute for Space Studies three-dimensional model," *Journal of Geophysical Research*, 93: 9341–9364.
- Hmielowski, Jay D., Lauren Feldman, Teresa A. Myers, Anthony A. Leiserowitz, and Edward W. Maibach (2014). "An attack on science? Media use, trust in scientists, and perceptions of global warming," *Public Understanding of Science*, 23(7): 866–883.

- Hope, Chris (2011a). "The PAGE09 Integrated Assessment Model: A Technical Description," Cambridge Judge Business School, Working Paper Series (4/2011). Available: http://www.jbs.cam.ac.uk/fileadmin/user_upload/research/workingpapers/wp1104.pdf.
- Hope, Chris (2011b). "The Social Cost of CO2 from the PAGE09 model," Cambridge Judge Business School, Working Paper Series (5/2011). Available: http://www.jbs.cam.ac.uk/fileadmin/user_upload/research/workingpapers/wp1105.pdf.
- IPCC (2001). *Climate change 2001: The Scientific Basis. Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*, [Houghton, J.T., Y. Ding, D.J. Griggs, M. Noguer, P.J. van der Linden, X. Dai, K. Maskell, and C.A. Johnson (eds.)], Cambridge: Cambridge University Press, 1st edition.
- IPCC (2013). *Climate Change 2013: The Physical Science Basis. Contribution of Working Group I to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change*, [Stocker, T.F., D. Qin, G.-K. Plattner, M. Tignor, S.K. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex and P.M. Midgley (eds.)], Cambridge University Press, 1st edition.
- IWG (2010). "Social Cost Of Carbon For Regulatory Impact Analysis Under Executive Order 12866." Technical Support Document, Interagency Working Group on Social Cost of Carbon, United States Government, Washington, DC.
- IWG (2013). "Technical Update of the Social Cost Of Carbon For Regulatory Impact Analysis Under Executive Order 12866." Technical Support Document, Interagency Working Group on Social Cost of Carbon, United States Government, Washington, DC.
- Jackman, Simon (2009). *Bayesian Analysis for the Social Sciences*, John Wiley & Sons Ltd, 1st edition.
- Jaynes, Edwin T. (2003). *Probability Theory: The Logic of Science*, Cambridge University Press.
- Kagan, Mark (2014). "Climate Change Skepticism in the Face of Catastrophe," Tinbergen Institute Discussion Paper. Available: <http://papers.tinbergen.nl/12112.pdf>.
- Kahan, Dan M., Hank Jenkins-Smith, and Donald Braman (2011). "Cultural cognition of scientific consensus," *Journal of Risk Research*, 14(2): 147-174.

- Kahan, Dan M., Ellen Peters, Maggie Wittlin, Paul Slovic, Lisa Larrimore Ouellette, Donald Braman, and Gregory Mandel (2012). "The polarizing impact of science literacy and numeracy on perceived climate change risks," *Nature Climate Change*, 2: 732–735.
- Kahneman, Daniel and Amos Tversky (1972). "Subjective probability: A judgment of representativeness," *Cognitive Psychology*, 3(3): 430 – 454.
- Kaufmann, Robert K., Heikki Kauppi, Michael L. Mann, and James H. Stock (2013). "Does temperature contain a stochastic trend: linking statistical results to physical mechanisms," *Climatic Change*, 118(3–4): 729–743.
- Kaufmann, Robert K., Heikki Kauppi, and James H. Stock (2006). "Emissions, Concentrations, & Temperature: A Time Series Analysis," *Climatic Change*, 77(3–4): 249–278.
- Kaufmann, Robert K., Heikki Kauppi, and James H. Stock (2010). "Does temperature contain a stochastic trend? Evaluating conflicting statistical results," *Climatic Change*, 101(3–4): 395–405.
- Kaufmann, Robert K and David I Stern (2002). "Cointegration analysis of hemispheric temperature relations," *Journal of Geophysical Research: Atmospheres (1984–2012)*, 107(D2): ACL–8.
- Koop, Gary (2006). *Bayesian Econometrics*, John Wiley & Sons Ltd, 2nd edition.
- Krosnick, Jon A. and Bo MacInnis (2010). "Frequent Viewers of Fox News Are Less Likely to Accept Scientists Views of Global Warming," Report for The Woods Institute for the Environment, Stanford University. Available: <http://woods.stanford.edu/sites/default/files/files/Global-Warming-Fox-News.pdf>.
- Kruschke, John K. (2014). *Doing Bayesian Data Analysis: A Tutorial with R, JAGS, and Stan*, Academic Press, 2nd edition.
- Leiserowitz, Anthony A., Edward W. Maibach, Connie Roser-Renouf, Nicholas Smith, and Erica Dawson (2013). "Climategate, public opinion, and the loss of trust," *American behavioral scientist*, 57(6): 818–837.
- Lewis, Nicholas (2013). "An Objective Bayesian Improved Approach for Applying Optimal Fingerprint Techniques to Estimate Climate Sensitivity," *Journal of Climate*, 26(19): 7414–7429.

- Malka, Ariel, Jon A. Krosnick, and Gary Langer (2009). "The association of knowledge with concern about global warming: Trusted information sources shape public thinking," *Risk Analysis*, 29(5): 633–647.
- McCright, Aaron M. and Riley E. Dunlap (2011a). "Cool dudes: The denial of climate change among conservative white males in the United States," *Global Environmental Change*, 21(4): 1163–1172.
- McCright, Aaron M. and Riley E. Dunlap (2011b). "The Politicization of Climate Change and Polarization in the American Public's Views of Global Warming 2001–2010," *Sociological Quarterly*, 52(2): 155–194.
- Meinshausen, Malte, Steven J. Smith, K. Calvin, John S. Daniel, M.L.T. Kainuma, J.F. Lamarque, K. Matsumoto, S.A. Montzka, S.C.B. Raper, K. Riahi, et al. (2011). "The RCP greenhouse gas concentrations and their extensions from 1765 to 2300," *Climatic Change*, 109(1-2): 213–241.
- Mills, Terence C. (2009). "How robust is the long-run relationship between temperature and radiative forcing?" *Climatic Change*, 94(3-4): 351–361.
- Nordhaus, William D. (2014). "Estimates of the Social Cost of Carbon: Concepts and Results from the DICE-2013R Model and Alternative Approaches," *Journal of the Association of Environmental and Resource Economists*, 1(2): 273–312.
- Oreskes, Naomi (2004). "The Scientific Consensus on Climate Change," *Science*, 306: 1686.
- Otto, Alexander, Friederike E.L. Otto, Olivier Boucher, John Church, Gabi Hegerl, Piers M. Forster, Nathan P. Gillett, Jonathan Gregory, Gregory C. Johnson, Reto Knutti, et al. (2013). "Energy budget constraints on climate response," *Nature Geoscience*, 6: 415–416.
- Pueyo, Salvador (2012). "Solution to the paradox of climate sensitivity," *Climatic Change*, 113(2): 163–179.
- Raney, Michael Andrew, Dav Clark, Daniel Lee Reinholz, and Sarah Cohen (2012). "Changing global warming beliefs with scientific information: Knowledge, attitudes, and RTMD (Reinforced Theistic Manifest Destiny theory)," in & R.P. Cooper N. Miyake, D. Peebles (ed.), "Proceedings of the 34th Annual Conference of the Cognitive Science Society. Austin, TX: Cognitive Science Society," pp. 2228–2233.

- Saad, Lydia (2014). "A Steady 57% in U.S. Blame Humans for Global Warming," *Gallup*, 18 March 2014. Available: <http://www.gallup.com/poll/167972/steady-blame-humans-global-warming.aspx>.
- Schmidt, Gavin (2007). "The CO₂ problem in 6 easy steps," *RealClimate.org*, 6 August 2007. Available: <http://www.realclimate.org/index.php/archives/2007/08/the-co2-problem-in-6-easy-steps/>.
- Stern, David I. and Robert K. Kaufmann (2000). "Detecting a global warming signal in hemispheric temperature series: A structural time series analysis," *Climatic Change*, 47(4): 411–438.
- Tol, Richard S. J. (2014). "Quantifying the consensus on anthropogenic global warming in the literature: A re-analysis," *Energy Policy*, 73: 701–705.
- Tol, Richard S. J. and Aart F. De Vos (1998). "A Bayesian Statistical Analysis of the Enhanced Greenhouse Effect," *Climatic Change*, 38: 87–112.
- Van Vuuren, Detlef P., Jae Edmonds, Mikiko Kainuma, Keywan Riahi, Allison Thomson, Kathy Hibbard, George C. Hurtt, Tom Kram, Volker Krey, Jean-Francois Lamarque, et al. (2011). "The representative concentration pathways: an overview," *Climatic Change*, 109(1-2): 5–31.
- Van Wijnbergen, Sweder and Tim Willems (2015). "Optimal Learning on Climate Change: Why climate skeptics should reduce emissions," *Journal of Environmental Economics and Management*, 70: 17–33.
- Verheggen, Bart, Bart Strengers, John Cook, Rob van Dorland, Kees Vringer, Jeroen Peters, Hans Visser, and Leo Meyer (2014). "Scientists Views about Attribution of Global Warming," *Environmental Science & Technology*, 48 (16): 8963–8971.

Appendix

2.A Bayesian multivariate linear regression

Suppose we are interested in describing the mean variation of some dependent variable y in response to a vector of k explanatory variables x_1, \dots, x_k over time. If our sample runs over T periods, then we have

$$y_t = \beta_1 x_{t1} + \dots + \beta_k x_{tk} + \epsilon_t, \quad t = 1, \dots, T, \quad (2.A.1)$$

where β_1, \dots, β_k are unknown regression parameters. This can be written in expected conditional value form as

$$E(y_t | \beta, X) = \beta x_{t1} + \dots + \beta_k x_{tk}, \quad t = 1, \dots, T. \quad (2.A.2)$$

Letting $X_t = (x_{t1}, \dots, x_{tk})$ denote the row vector of k predictors for the t^{th} period and $\beta = (\beta_1, \dots, \beta_k)$ the column vector of regression coefficients, we can simplify to

$$E(y | \beta, X) = X\beta. \quad (2.A.3)$$

In a linear regression model, the y_t observations are assumed to be conditionally independent given values of the parameters and explanatory variables. We also assume equal variances across periods, $\text{Var}(y_t | \theta, X) = \sigma^2$, where $\theta = (\beta_1, \dots, \beta_k, \sigma^2)$ is the vector of unknown regression parameters. Finally, it is assumed that the errors, $\epsilon_t = y_t - E(y_t | \beta, X)$ are independent and normally distributed, such that

$\epsilon \sim \mathcal{N}(0, \sigma^2)$. Our model can now be written in matrix notation for all periods as

$$y|\beta, \sigma^2 \sim \mathcal{N}_T(\mathbf{X}\beta, \sigma^2\mathbf{I}), \quad (2.A.4)$$

where y is the vector of observations; \mathbf{X} is the $T \times k$ design matrix; \mathbf{I} is the identity matrix; and $\mathcal{N}_T(\mu, A)$ denotes a multivariate normal distribution of dimension T with mean vector μ and variance-covariance matrix A . Mapping this model to a valid posterior density will depend on the particular choice of prior. However, the general form is given by the joint distribution

$$p(\beta, \sigma^2|y) = p(\beta|y, \sigma^2)p(\sigma^2|y). \quad (2.A.5)$$

To further illustrate using an example described by Albert (2009), suppose we apply a standard noninformative prior of $g(\beta, \sigma^2) \propto \frac{1}{\sigma^2}$. Then the posterior distribution of the coefficient vector β , conditional on the error variance σ^2 , is multivariate normal

$$\beta|y, \sigma^2 \sim \mathcal{N}(\hat{\beta}, V_\beta\sigma^2), \quad (2.A.6)$$

where $\hat{\beta} = (X'X)^{-1}X'y$ is the OLS estimate of the true value of β ; and $V_\beta\sigma^2 = (X'X)^{-1}\sigma^2$ is the variance-covariance matrix. If one defines the inverse gamma density $\mathcal{IG}(a, b)$, proportional to $y^{-a-1}\exp(-b/y)$, then the marginal posterior distribution of σ^2 is itself inverse gamma

$$\sigma^2|y \sim \mathcal{IG}\left(\frac{n-k}{2}, \frac{(y - X\hat{\beta})'(y - X\hat{\beta})}{2}\right). \quad (2.A.7)$$

2.B Supplementary tables and figures

Table 2.B.1: Transient climate response (°C)

	Prior	Posterior
Noninformative	-	1.5 [1.4, 1.7]
Moderate Lukewarmer	1.0 [0.5, 1.5]	1.5 [1.3, 1.7]
Strong Lukewarmer	1.0 [0.9, 1.1]	1.3 [1.2, 1.5]
Moderate Denier	0.0 [-0.5, 0.5]	1.5 [1.3, 1.7]
Strong Denier	0.0 [-0.1, 0.1]	0.4 [0.3, 0.6]

Mean estimates are given, with square brackets denoting 95% credible intervals.

Table 2.B.2: Transient climate response (°C): Results from different temperature series

	Mean	95% Credible interval	Effective sample period
HadCRUT4 ^a	1.5	[1.4, 1.7]	1866–2005
CW2014 ^b	1.6	[1.4, 1.8]	1866–2005
GISTEMP ^c	1.7	[1.5, 1.9]	1880–2005

All estimates are computed using noninformative priors. The effective sample period is limited either by overlap with other regression covariates (^{a,b}), and/or because data begin at a later date (^c).

^a Compiled by the Hadley Centre of the UK Met Office and the Climatic Research Unit at the University of East Anglia. Serves as the default temperature series in this paper.

^b Cowtan and Way (2014). Corrects for biases in the HadCRUT4 series due to incomplete global coverage.

^c Compiled by the NASA Goddard Institute for Space Studies.

Chapter 3

Market might. Hydro power.

Abstract

I test for evidence of strategic behaviour in the Norwegian electricity sector using a uniquely detailed dataset of hydropower reservoirs. Consistent with theory, I find that market power leads to an intertemporal reallocation of water resources across periods. An increase in market power causes firms to withhold production when demand is at its most inelastic. This would permit dominant firms to recoup higher profits while consumers are least responsive to price changes. The effects are modest next to other factors governing reservoir management, such as annual snow-melt. Yet, they may still cause the production profile of hydropower firms to diverge in meaningful ways if the differences in market share are large enough, and particularly when regional transmission constraints are binding.

JEL Codes: Q25, Q41, L12, L13

Keywords: market power, electricity markets, hydropower, renewable energy

3.1 Introduction

Much of economics is concerned with deviations from the ideal of competitive markets. A standard result from microeconomic theory is that dominant firms will tend to produce less than the social optimum, thereby raising prices and ensuring higher profits (e.g. Mas-Colell et al., 1995). How this behaviour manifests itself in reality will depend on the specific institutional characteristics of the market in question. Moreover, testing theoretical predictions in an empirical setting can prove challenging due to the limitations of real-world data. The present paper contributes to our understanding of these issues within a particular empirical context. Namely, how does market power affect firm behaviour in a hydro-based electricity market?

Hydropower has several unusual features that make it interesting from an economic perspective. To begin with, the variable costs of operating a hydropower plant are negligible. Employee wages and maintenance fees are best regarded as fixed costs independent of output. At the same time, reservoir inflows are both stochastic and highly seasonal. Yet, water is a durable good that can be stored for long periods of time. The producer's decision collapses into one of how much water should be used today and how much should be saved for tomorrow. In other words, production in a hydropower system is a dynamic optimisation problem determined by the opportunity cost of water — the so-called *water value*. Together with variations in demand, this creates opportunities for exploiting market power via intertemporal price discrimination. A firm with market power can increase profits by withholding supply in periods with relatively inelastic demand; hence driving up price when consumers are least responsive to such changes. Dominant hydropower firms will reallocate their water resources across periods in accordance with this strategy.

The present paper provides empirical evidence of such strategic behaviour among Norwegian hydropower firms. An increase in market power leads to a modest, but distinctive, pattern of intertemporal resource shifting. Compared to previous studies, I use a much richer dataset that links individual hydropower reservoirs to specific firms, as well as broader market data such as electricity flows and transmission constraints. These data allow me to make several contributions to the literature. First and foremost, rather than inferring evidence of market power solely through aggregate measures such as wholesale electricity prices, I am able to look at the behaviour of individual hydropower plants *directly*. The relationship between water resource management and market power is therefore determined at the firm-reservoir level. This refinement not only establishes a close correspon-

dence between the empirical set-up and underlying theory, but also permits the use of a conceptually straightforward regression framework (i.e. panel fixed effects). Changes to bidding area divisions and binding transmission constraints provide the additional layers of exogenous variation that enable me to cleanly identify the effect of local market power on firm behaviour. In addition to the above, I consider a longer time period than previous studies. My sample runs over 14 years from the beginning of 2000 until the end of 2013. In so doing, I hope to shed light on the way that dominant firms will strategically utilise their water resources, not just in the short-term, but in response to changing market conditions over the course of months, seasons and years.

Related studies on anti-competitive behaviour in the Norwegian — and broader Nordic — market include Johnsen et al. (1999), Hjalmarsson (2000), Steen (2004), Kauppi and Liski (2008), and Mirza and Bergland (2012). A review is provided by Fridolfsson and Tangerås (2009). The general finding is one of healthy competition, but with some scope for exercising local market power due to constraints in transmission capacities. However, as noted, all of the above studies rely on aggregate data. This places strong restrictions on the methods that can be used to support causal inference. For example, several studies make use of the well-known Bresnahan (1982) and Lau (1982) framework for estimating market power in the absence of marginal cost data. Yet the Bresnahan-Lau model ultimately presumes that firms face a static production decision at each point in time. It is therefore of limited use for understanding the intertemporal aspects of hydropower production.¹ Moreover, the empirical analyses to date have tended to focus on short-run deviations from competitive prices. There is still some uncertainty regarding strategic behaviour among dominant firms over long periods of time. I attempt to resolve these uncertainties in the present paper.

The remainder of the paper is organised as follows. Section 3.2 provides the theoretical underpinnings for understanding firm behaviour in a hydropower system. Section 3.3 introduces the dataset and describes the key institutional settings of the Norwegian electricity market. Section 3.4 describes the econometric strategy for

¹Fridolfsson and Tangerås (2009, p. 3689) are led to remark in their review: “A lack of firm level data may explain why only the Bresnahan-Lau model or the even less demanding methodology proposed by Johnsen et al. [sc. 1999] has been applied to the Nordic power market. This suggests that more detailed data would be highly valuable for [determining] market power in the Nordic market for wholesale electricity.” Similarly, Kauppi and Liski (2008, p. 35): “Our approach to efficient allocations and those distorted by imperfect competition is aggregative. Analysis exploiting more detailed information on capacities, usage, and regional heterogeneity is therefore called for. If such data becomes available, one could potentially estimate hydro usage policies directly from the data[...].”

causal inference, before moving on to the ensuing results in Section 3.5. Section 3.6 concludes.

3.2 Theoretical motivation

In his influential text on hydropower economics, Førsund (2015) outlines the ways in which market power can affect firm behaviour in various hydropower systems. The general result is to incentivise a reallocation of water from periods with relatively inelastic demand to periods with relatively more elastic demand. The contribution of this paper is empirical and, as such, I will not describe these theoretical permutations in detail. However, it will be useful to recapitulate a version of the simplest case — i.e. monopoly with no uncertainty, outside trade or reservoir constraints — to get a sense of the underlying intuition.

Consider the profit maximization problem of a hydropower monopoly in a two period setting,

$$\begin{aligned} \max \sum_{t=1}^2 p_t(q_t) \cdot q_t \\ \text{s.t. } \sum_{t=1}^2 q_t \leq W, \end{aligned} \tag{3.1}$$

where $p_t(q_t)$ is an inverse demand function with standard properties (e.g. price decreasing in quantity), q_t is the quantity of electricity demanded by consumers (or, more precisely, the water equivalent thereof), and W is the known water endowment for the monopolist's reservoir. As is standard in this literature, we may further take period one to be summer and period two to be winter. This assumption has no bearing on the theoretical results, but will become useful as a reference for the empirical setup later.

The necessary first order conditions for profit maximization are

$$\begin{aligned} \frac{\partial L}{\partial q_t} = p_t'(q_t) \cdot q_t + p_t(q_t) - \lambda \leq 0 \\ (= 0 \quad \text{for } q_t > 0) \end{aligned} \tag{3.2}$$

and

$$\begin{aligned} \lambda &\geq 0 \\ & (= 0 \quad \text{for } \sum_{t=1}^2 q_t < W). \end{aligned} \quad (3.3)$$

The parameter λ denotes the shadow price on stored water, i.e. positive if the resource constraint in equation (3.1) is binding and zero otherwise. Without loss of generality, let us assume that the shadow price is positive and that the monopolist also produces in both periods.² The first order conditions may then be written as

$$p_1(q_1) \left(1 + \frac{1}{\epsilon_1}\right) = p_2(q_2) \left(1 + \frac{1}{\epsilon_2}\right) = \lambda, \quad (3.4)$$

where $\epsilon_t = \frac{p_t}{q_t} \frac{\partial q_t}{\partial p_t} < 0$ is the price elasticity of demand. Rearranging equation (3.4), it is easy to see that prices depend on the relative demand elasticities in each period. For example, we would have

$$p_1(q_1) > p_2(q_2) \quad \text{if} \quad |\epsilon_1(q_1)| < |\epsilon_2(q_2)|. \quad (3.5)$$

Since we have assumed a downward-sloping demand curve, the above corresponds to

$$q_1 < q_2 \quad \text{if} \quad |\epsilon_1(q_1)| < |\epsilon_2(q_2)|. \quad (3.6)$$

The monopolist solution thus involves a reallocation of production across periods, contingent on the elasticity of demand. This contrasts with the social solution that arises under perfect competition, where production and prices are equalised across periods.³ We can further generalise the difference between monopoly and perfect competition in this simple setup as

$$q_t^M < q_t^C \quad \text{and} \quad q_{\hat{t}}^M > q_{\hat{t}}^C \quad \text{if} \quad |\epsilon_t(q_t)| < |\epsilon_{\hat{t}}(q_{\hat{t}})|, \quad (3.7)$$

²If we instead assumed that the shadow price is zero (due to a non-binding resource constraint), then the ensuing results are largely unchanged but for some fraction of water that remains unused.

³By definition, the price elasticity of demand facing a competitive firm is always perfectly elastic, i.e. $\epsilon \rightarrow \infty$. That competition leads to equal prices and quantities across periods is easily shown by solving the above set-up as a social optimisation problem that maximizes total welfare. See Førsund (2015).

where the M and C superscripts denote the monopolist and competitive outcomes, respectively. The monopolist is able to recoup higher profits by withholding supply — hence driving up the electricity price — during the relatively inelastic period when consumers are least responsive to such changes. In other words, market power not only causes prices and quantities to diverge from their social optimums, but also implies an observable difference in the way that reservoirs are managed. Reservoirs belonging to dominant firms will tend to be relatively fuller during inelastic periods than they would otherwise have been under competition. The reverse is true during elastic periods.

The theoretical extensions that Førsund (2015) and others (e.g. Hansen, 2009; Mathiesen et al., 2013) explore beyond the simple case presented here may be regarded as variations on a theme. While each variation has the ability to ameliorate or exacerbate market distortions in its own way, the substantive result is largely unchanged.⁴ Market power leads to a strategy of shifting water use from relatively inelastic demand periods to relatively elastic ones. As we shall see, in the Norwegian context this amounts to comparatively fuller reservoirs during summer, with dominant firms undersupplying relative to the competitive outcome, and vice versa during winter.

3.3 Market characteristics and data

This paper introduces a rich and uniquely detailed dataset of Norwegian hydropower reservoirs, which has been constructed from a variety of sources. It is hence worth describing the original sources in some detail, as well as the methods that have been used to merge these disparate parts into a unified dataset.⁵ However, I first provide a brief overview of the Norwegian electricity market and its main institutional features.

⁴For example, trade with outside regions can moderate the intertemporal disparities yielded by the standard monopoly model. However, accounting for transmission constraints brings us back towards the original result.

⁵Data work was primarily executed in the *R* programming environment. While the subsequent subsections provide a descriptive overview of methods, full data cleaning and merging documentation is available upon request.

3.3.1 The Norwegian electricity market

Norway liberalised its electricity sector in 1991, placing it in the vanguard of a broader deregulation movement (Bye and Hope, 2005; Joskow, 2008). Together with its neighbours — Sweden, Finland and Denmark — Norway would go on to form the world’s first, and still largest, multinational power exchange, *Nord Pool*.⁶ The foundation of this exchange is the day-ahead *Elspot* auction for physical delivery of electricity. Hourly supply and demand bids are aggregated and then matched to determine a market clearing price, commonly referred to as the *system price*. A key institutional feature of the Nord Pool exchange is that countries are sub-divided into distinct Elspot *bidding areas*. Absent transmission constraints, area prices are equal to the common system price. However, when transmission constraints are binding, each area becomes its own separate market and prices can diverge as a result. This potentially confers local market power to dominant firms within those areas.

Norway is currently comprised of five Elspot bidding areas: NO1 (east), NO2 (south), NO3 (mid), NO4 (north), and NO5 (west). It should be noted that this configuration of bidding areas has changed multiple times over the last decade and a half. New bidding areas have been added (and sometimes removed) and boundaries between existing areas have been redrawn. Such “regime” changes are important from an empirical perspective, because they provide an exogenous source of variation in local market power. More specifically, the changing division of bidding areas, in combination with data on individual hydro reservoirs and binding transmissions constraints, will allow me to identify the causal effect of local market power on firm behaviour. I expand upon this point in the empirical section of the paper.

We have seen from theory that relative demand elasticities are central in determining how dominant hydropower firms would reallocate production across periods. Electricity demand in Norway fluctuates significantly over the course of hours, days of the week, and time of year. The principal focus of this paper is the behaviour of hydropower firms over the longer-term. It therefore makes sense to abstract from short-term fluctuations (e.g. day versus night, weekdays versus weekends) and concentrate on the demand variations over the course of months and seasons. Winter electricity consumption in Norway is approximately double that of summer, primarily because of indoor heating requirements. At the same time, elec-

⁶Hydropower predominates at both the national and regional level. Norwegian reservoirs alone generated around 130 TWh of electricity in 2013. This corresponds to 96 percent of national electricity generation and more than a third of the 380 TWh generated in the Nordic region as a whole (NordREG, 2014; SSB, 2015).

trical heating is substitutable (wood, fuels, etc.) and hence permits greater flexibility among consumers. In contrast, summer electricity demand is dominated by technical end-uses that do not allow for easy substitution. These seasonal differences in substitution and adjustment possibilities contribute to a somewhat surprising finding: Demand elasticities are significantly *higher* during the Norwegian winter than they are during summer (Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008). I take this observation as given for interpreting the empirical results later in the paper. Norwegian electricity demand is relatively more inelastic during summer, and relatively more elastic during winter.

3.3.2 Hydropower reservoirs

Time series data for the 500 most important hydropower reservoirs — representing approximately 97 percent of total Norwegian system capacity — were obtained from the Norwegian Water and Energy Directorate (NVE). For most reservoirs, these data are observed at a daily resolution from January 2000 to December 2013, and contain water readings in terms of both volumes (million m³) and levels (m). It should be noted that reservoirs in Norway are subject to regulation regarding their maximum (and minimum) holding capacities. Hydropower firms are held responsible for keeping their reservoirs within these boundaries in order to protect against dangerous flooding and environmental degradation.

To ensure comparability between the different reservoir sizes in my dataset, the reservoir data are normalised as percentages of their respective maximum regulated capacities. While this normalisation procedure is generally straightforward, a number of reservoirs in the dataset suffer from discrete jumps in measurement values, while others have conflicting regulatory limits ascribed to them. These anomalies may reflect adjustments to the base measurement value (e.g., metres above sea level versus a local reference point), regulatory changes, and, in some rare cases, building out of extra capacity. To account and correct for such anomalies, the data are run through an automated filter to detect large, discrete jumps in measurement values and other outliers. These are then corrected as best as possible by reconciling the data with the various regulatory limits, and by comparing the volume and levels series for consistency.⁷

A related problem is that some reservoirs exhibit distinctly unnatural trends. Most

⁷While the data analysis part of this study does not utilise the levels series, it nonetheless provides a very useful counterpoint to the volumes series for this reason.

obviously, long streaks of the same recurring value. This issue is effectively limited to small reservoirs in the dataset and almost certainly constitutes measurement error. Streaks extending over 10 or more consecutive observations are thus discarded from the analysis. As a final check, time-series plots of all 500 individual reservoirs are examined manually to check for abnormalities that the automated filters may have missed, leaving a handful of cases to be corrected as per the above. Any remaining data anomalies that could not be reconciled in a satisfactory manner, or rationally accounted for, have been dropped from the analysis. This leaves a clean dataset of 499 Norwegian hydropower reservoirs approaching 2 million daily observations.

Following the logic of several previous studies (e.g. Steen, 2004; Mirza and Bergland, 2012), this paper focuses on local market power in the Norwegian electricity sector that arises from internal transmission bottlenecks.⁸ However, the richness of my dataset allows me to go further by tracking changes to the way that Elspot bidding areas have been configured over time. These changes create an additional source of variation in local market power that allows me to cleanly identify its impact on hydro firm behaviour vis-à-vis the management of individual reservoirs. The next step in the data compilation process thus involves linking each hydropower reservoir to a set of relevant covariates, including firm information, current Elspot region, and geographic coordinates. These data are again obtained from the NVE and provide the means for determining the market power that each firm wields within a designated Elspot bidding area (as well as the Norwegian electricity sector as a whole) at a given point in time.

I define market power in this paper as the share of summed maximum reservoir capacities that each producer wields within a designated bidding area. More specifically, the market share S of hydropower firm f , in a bidding area comprising F firms in total, is given as

$$S_f = \frac{\sum_{i=1} \bar{R}_{if}}{\sum_{j=1}^F \sum_{i=1} \bar{R}_{ij}}, \quad (3.8)$$

where \bar{R}_i is the maximum regulated capacity of reservoir i . This “share of overall capacity” definition is perhaps best regarded as proxy for true market power; in the sense that it captures potential, rather than actual, output of a firm relative to its competitors. However, it thereby avoids the endogeneity problem associated with

⁸Earlier theoretical contributions and empirical applications to other markets include those by Joskow and Tirole (2000) and Borenstein et al. (2000).

definitions that *do* measure actual output. Namely that changes in output have a direct impact on contemporaneous reservoir water volumes, which serve as the dependent variable in my econometric model.⁹

Tracing the evolution of local market power (market share) over time now involves three steps: (i) mapping reservoirs to the correct bidding area under a particular Elspot regime, (ii) summing the maximum reservoir capacities within these areas by firm, and (iii) comparing the summed firm capacities to the overall reservoir capacity within an area. The information needed to complete these steps does not readily exist. However, using the current Elspot allocation as a fixed starting point — together with the series of reservoir coordinates already obtained, the Elspot change log document hosted by Nord Pool¹⁰, a map of the Norwegian electricity grid components¹¹, and several other sources — I am able to manually back out the divisions of 16 earlier regimes going back to the beginning of 2000. These are depicted in Appendix 3.A. The most recent Elspot regime, which went into effect on 2 December 2013, is shown on its own in Figure 3.1. I will at times refer to these various regimes by the capitalised letters A to Q for convenience. Furthermore, note that the bidding areas are not consistently defined under the different regimes, even if they have been assigned the same name. For example, NO1 under regime A is very different in size and coverage to NO1 under regime Q. To avoid potential confusion on this matter, I will use the term *zone* to denote specific regime-area combinations, e.g. A-NO1 versus Q-NO1.

Based on the above distinctions, there are several ways that one could measure (regional) market power within the Norwegian electricity system, and how this changes over time. Table 3.1 shows the Herfindahl index measures for the various bidding areas under different Elspot regimes. Approximately 36 percent of the 55 realised zones are characterised by an index score indicating high concentration ($H > 0.25$), 42 percent are moderately concentrated ($0.15 < H < 0.25$), while the remaining 22 percent are unconcentrated ($H < 0.15$). As an alternative to the Herfindahl index, we can also look at the market shares wielded by the top three producers in each area. These are shown in Appendix 3.B and provide further evidence of potential market power at the regional level. The leading hydropower firm in some

⁹Similarly, note that this definition of market power is expressed in terms of volume, rather than actual energy output. Although the physical volume of water held by a hydropower reservoir does not perfectly correlate with its true energy potential — other factors such as flow rate and fall height can influence the final outcome — it should again serve as a good proxy. Moreover, one could reasonably expect variations between reservoir volumes and energy potential to average out across producers.

¹⁰<http://nordpoolspot.com/globalassets/download-center/elspot/elspot-area-change-log.pdf>

¹¹<http://gis3.nve.no/link/?link=nettanlegg>




Figure 3.1: Norwegian hydropower reservoirs by Elspot bidding area (as of 2 December 2013)

areas may command as much as 60 percent of available water volumes, depending on the regime.

3.3.3 Electricity flows and transmission constraints

All electricity data are obtained from Nord Pool. The flow of excess electricity from one Elspot area to another takes place via predefined corridors. That is, up to the capacity constraints of the transmission lines making up that corridor. Note too that a bidding area can have several corridors attached to it, depending on the number of neighbouring areas that it shares borders with. Electricity flows for the individual Norwegian bidding areas are available at an hourly basis, by corridor, from November 2000 onwards.¹² These flows are then matched to the maximum transfer capacities for the relevant corridors at each hour. The transmission constraint on

¹²Real-time flows can be viewed on the Statnett website: <http://www.statnett.no/en/Market-and-operations/Data-from-the-power-system/Nordic-power-flow/>

Table 3.1: Herfindahl index of the Norwegian Elspot bidding areas over time

Regime	Date	NO1	NO2	NO3	NO4	NO5
A	01 Jan 2000	0.139	0.348	–	–	–
B	02 Oct 2000	0.139	0.192	0.443	–	–
C	01 Jan 2001	0.139	0.348	–	–	–
D	12 Mar 2001	0.143	0.147	0.443	–	–
E	11 Jun 2001	0.139	0.348	–	–	–
F	16 Dec 2002	0.159	0.243	0.215	0.373	–
G	02 Jun 2003	0.139	0.348	–	–	–
H	13 Dec 2003	0.137	0.294	0.373	–	–
I	29 May 2004	0.139	0.348	–	–	–
J	18 Nov 2006	0.139	0.192	0.443	–	–
K	17 Nov 2008	0.139	0.348	–	–	–
L	13 Apr 2009	0.139	0.192	0.443	–	–
M	11 Jan 2010	0.152	0.216	0.192	0.443	–
N	15 Mar 2010	0.205	0.257	0.185	0.443	0.171
O	05 Sep 2011	0.205	0.244	0.185	0.443	0.208
P	05 Dec 2011	0.205	0.244	0.185	0.443	0.208
Q	02 Dec 2013	0.325	0.237	0.185	0.443	0.200

a corridor is defined as binding — whether importing or exporting — whenever the hourly flows reach the maximum transfer capacity. It is furthermore possible to infer binding transmission constraints during the nine months prior to November 2000 by looking at differences in spot prices between neighbouring regions. For instance, if the price in NO2 is higher than NO1, then it is reasonable to infer that NO2 is importing electricity from NO1, and that the transfer constraint is also binding. If the spot price between two regions is equal, then we do not know much beyond the fact that the constraint is not binding. We cannot reliably say which area is exporting to the other, and vice versa.

All told, a binding constraint is a relatively common occurrence. Figure 3.2 shows that bidding areas can generally expect to reach the maximum transfer capacity along one of their corridors at least once per day on average. As per Figure 3.3, the picture is qualitatively similar even if we further normalise by the total in-use hours for each bidding area. For example, in 2000, bidding area NO1 ran up against the maximum transfer capacity 78 percent of time that it was exporting power to one of its neighbours, and 45 percent of the time that it was importing power.




Figure 3.2: Percentage of operating days with a binding transmission constraint of at least one hour along at least one corridor

3.4 Econometric approach

Consider a fixed effects model to estimate water volumes in reservoir i (i, \dots, N), belonging to hydropower firm f ($f = 1, \dots, F$), at time t ($t = 1, \dots, T$). The primary observation unit is the reservoir, while the secondary observation unit is the firm. Following the notation of Abowd et al. (2008), the reservoir–firm relationship may be conceptualised through a link function, $f = F(i, t)$, which indicates that firm f is managing reservoir i at time t . The regression model may thus be written as

$$V_{it} = \sum_{m=1}^{12} \beta_m M_{mt} + \sum_{m=1}^{12} \gamma_m M_{mt} \cdot S_{it,F(i,t)} + \mathbf{X}\boldsymbol{\beta}_X + a_i + v_{it}, \quad (3.9)$$

where V is reservoir volume as a percentage of maximum regulated capacity, M_m is set of month dummies, and S is the market share wielded by the operating firm as per equation (3.8). By interacting S with M_m , we are allowing for the fact that mar-




Figure 3.3: Binding transmission constraints as a percentage of total in-use hours along all corridors

ket share can have a differential effect on reservoir volumes, depending on how the price elasticity of demand varies by period.¹³ The regression model is completed by a group of additional controls, X , while the first component of the composite error term, $a_i + v_{it}$, denotes an unobservable reservoir-specific effect — e.g. idiosyncratic operating characteristics or hydrological conditions — that we eliminate from the model via within group transformation. The key parameters of interest in the above regression model are the γ_m coefficients pertaining to the interacted market share terms. Given the observed variations in Norwegian demand elasticities — i.e. relatively more inelastic when it is warmer (Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008) — we would expect a positive sign on these coefficients during the summer months as dominant producers withhold their production, before turning

¹³Following Balli and Sørensen (2013), the reservoir-specific means may also be subtracted from the continuous market share variable prior to running the regression, so as to safeguard against possibly spurious results in the interaction term. That is, the estimated model becomes $V_{it} = \sum_{m=1}^{12} \beta_m M_{mt} + \sum_{m=1}^{12} \gamma_m M_{mt} \cdot (S_{it,F(it)} - \bar{S}_i) + a_i + \epsilon_{it}$. Given that the underlying logic of the model is unchanged, I shall nonetheless continue using the notation of equation (3.9) for simplicity.

negative in the winter months.

While the naive specification of regression equation (3.9) captures the essence of the underlying theory, it suffers from two, related conceptual limitations: It does not explicitly account for the conditions that make local market power possible, and it does not address a potential identification problem related to changing demand. Fortunately, the nature of dataset allows me to correct for both limitations simply by adding a few additional controls to the model. To see why this is the case, it is worth taking a brief digression to discuss some of the specific empirical advantages of the dataset.

This paper aims to exploit the boundary changes that have been made to the Norwegian Elspot bidding areas over time. A key underlying assumption is that these changes yield a plausibly exogenous source of variation in regional market power. My identification strategy thus relies on the fact that (i) the Elspot area divisions are determined by outside factors, and (ii) they affect a hydropower firm's production decision only via changes in market share. It is relatively straightforward to argue for the former. The reason that separate bidding areas exist in the first place is that geographical and technical constraints limit the flow of electricity that is physically possible between two regions (Steen, 2004; Mirza and Bergland, 2012; ENTSOE, 2015, etc.). That these bidding areas have been redrawn over time as the system operator seeks to best manage internal congestion issues, scheduled maintenance, outages, and the laying out of new cables, is in of itself testament to the underlying physical constraints.

The second component of my identification strategy — i.e. changes to the Elspot bidding areas should affect producer behaviour only through changes in market share — is potentially complicated by the fact that the demand will also change with a redrawing of the bidding areas. Yet, the demand complication is dealt with fairly easily through the inclusion of zone dummies. This controls for demand by grouping reservoirs at the zone level. (Recall that a zone denotes a bidding area under a particular Elspot regime.) Any residual differences between otherwise similar reservoirs can then be interpreted as a causal result of market share, since producers within the same zone will always face the same demand; irrespective of whether transmission constraints are binding or not.

It is worth noting, however, that this latter claim would not hold true if we were to compare plants with different production technologies (e.g. hydro versus nuclear). Electricity in modern power systems is dispatched according to the merit order of production, with plants ranked in ascending order of their marginal costs; see Fig-

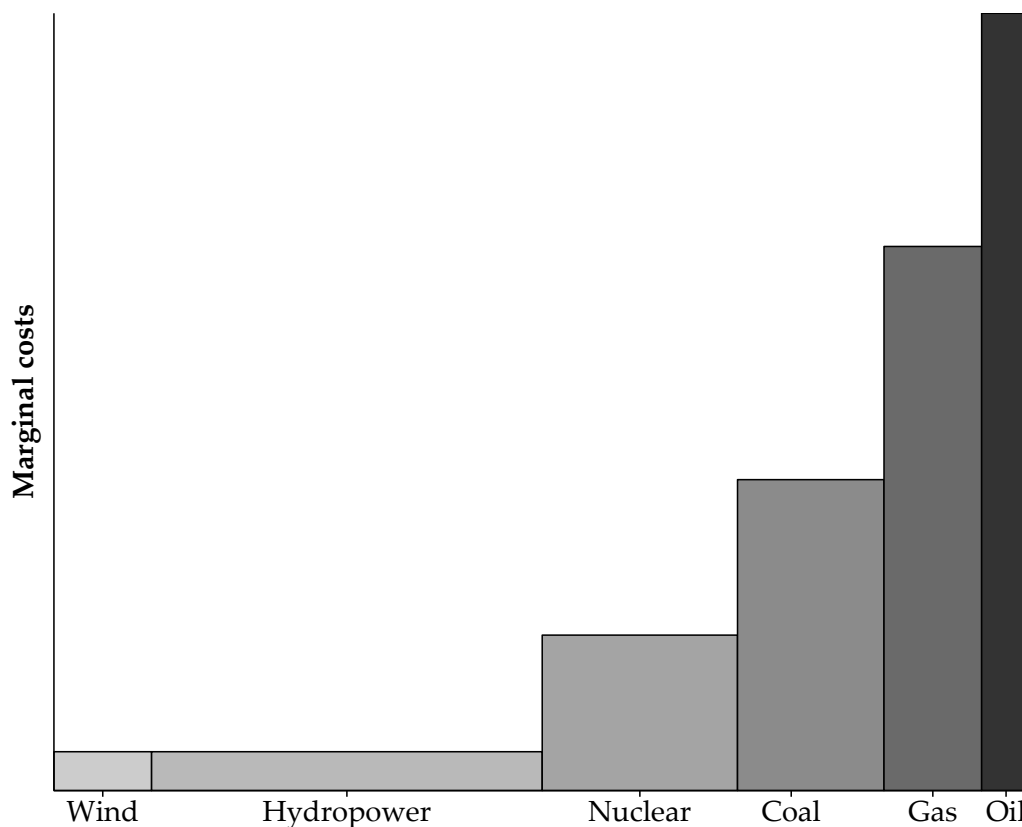


Figure 3.4: Merit order of electricity generation

Figure 3.4. Depending on the supply characteristics of the redrawn bidding areas, the relative positions that two plants occupy along the merit curve would therefore likely change if they did not have the same production technology. In other words, the probability of them serving the same consumer — or, at least, their relative probabilities — would diverge and we would no longer be able to control effectively for demand. The exclusive focus on hydropower plants in the present paper circumvents this problem. The fact that the marginal costs of hydropower production are negligible, thereby ensuring that these plants always occupy the lower rungs of the merit curve, further simplifies the econometric analysis (Førsund, 2015; Kauppi and Liski, 2008).

Having considered the issues of identification and regional market conditions, let us return to the regression model. Equation (3.10) expands on the earlier specification in equation (3.9) by explicitly introducing zone and regime dummies into model.¹⁴ Employing a reservoir–firm link function as per before, we now estimate

¹⁴Including regime dummies alongside the zone dummies may at first seem redundant. After all, zones are already denoting regime-area combinations. However, doing so allows for the possibility that the configuration of bidding areas under a particular regime brings about an aggregate effect. For example, some regimes may comprise a less efficient configuration of bidding areas, which raises

water volumes of hydropower reservoir i , belonging to firm f , in zone z , and at time t as

$$\begin{aligned}
 V_{it} = & \sum_{m=1}^{12} \beta_m M_{mt} + \sum_{m=1}^{12} \gamma_m M_{mt} \cdot S_{it,F(it)} \\
 & + \sum_{r=1}^{17} \delta_r R_{rt} + \sum_{z=1}^{55} \eta_z Z_{it} + \mathbf{X}\boldsymbol{\beta}_X + a_i + v_{it},
 \end{aligned} \tag{3.10}$$

where R_r is a set of regime dummies, Z_z is a set of zone (i.e. regime-area) dummies, and the remaining parameters are as before. Given that producers will be able to exploit their local market power most effectively when transmission constraint are binding, we may also want to limit the analysis to dates where this condition is true. An alternative approach would be to account for binding transmission constraints through additional dummy variables, or even a continuous variable measuring the extent to which the constraints were binding relative to the desired time period. However, both of these alternatives would significantly increase the parameter space and lead to three-way interaction terms, which do not lend themselves to clear interpretation.

As a final practical matter, recall that the compiled dataset comprises nearly 2 million daily observations over 500 reservoirs. Implementing the above regression on a panel dataset of this size is problematic because of software and technical limitations associated with in-memory performance (e.g. see Wickham, 2014). To ease the computational burden, as well as reduce the time series dimensionality of the problem, the data are collapsed into their weekly means or, where appropriate, modes.¹⁵ This leaves a more manageable dataset of approximately 335,000 observations.

3.5 Results

The main empirical results are presented in Table 3.2 and Figure 3.5. While each column in the table depicts a different model specification, the highlighted coefficients may all be interpreted in the same way. These correspond to the γ_m parameters described in regression equations (3.9) and (3.10), and show the marginal effect of increasing market share on reservoir volumes, contingent on month (season). Our

the possibility of transmission bottlenecks and congestion within the system as a whole. This would have potential knock-on effects beyond the individual constraints pertaining to any one zone.

¹⁵For example, a bidding area is defined as experiencing a binding constraint in electricity flow at the weekly scale if this is true for most days during that particular week.

theoretical framework — together with the empirical literature on Norwegian demand elasticities — suggests that the sign on these variables should be positive during the summer months and negative during the winter months. Moreover, since both reservoir volume and market share are measured in percentages, the coefficients should be read as elasticities. While the remaining coefficients and interaction dummies have been omitted from the table for brevity, these are all jointly significant and, in the case of the month dummies, of the expected sign and magnitude.

Model (1) is a direct application of the naive specification in equation (3.9). The results are reasonably encouraging with respect to theory. We see that higher market share generally leads to a fall in reservoir volumes (i.e. increased production) during winter. However, evidence for the reverse during summer is less clear. More importantly, we cannot be sure that these results are empirically valid, since there has been no attempt to control for potential changes in demand due to redrawn bidding areas, nor binding transmission constraints. Model (2) seeks to address these shortcomings by introducing zone and regime dummies — as per regression equation (3.10). The new specification yields an improved correspondence between the empirical results and theory. We now obtain coefficients that are of the expected sign and statistically significant across both seasons. Yet, we are still not explicitly controlling for transmission constraints. The last two model specifications are aimed at addressing this issue. Model (3) imposes a filter that limits the sample to observations that experienced a binding transmission constraint of at least one hour.¹⁶ It almost goes without saying that is a rather weak requirement — less than three percent of the observations fall out of the sample — and it does not come as a surprise that the coefficients are barely changed by its introduction. In contrast, Model (4) imposes a more meaningful filter on the sample. Here we exclude any observations where binding transmission constraints account for less than 50 percent of total in-use hours for that week. That is, if a reservoir belongs to a zone where power was flowing to or from its neighbours for 100 hours in a week, then we would filter out this reservoir if the relevant transmission corridors were constrained for less than 50 hours. Approximately a third of observations drop out of the analysis as a result of this more stringent filter.

Model (4) most accurately captures the conditions that make local market power possible. It hence represents the preferred regression specification and is deserving of a more detailed discussion. We see that increasing market power causes

¹⁶For the weekly data used in the analysis, this would apply to reservoirs that lie within a zone where the majority of days during that week experienced a constraint of at least one hour.

Table 3.2: Regression results

Dependent Variable: Reservoir volume

	(1)	(2)	(3)	(4)
<i>Market share effect</i>				
Summer months				
May	-0.197*** (0.023)	-0.073*** (0.023)	-0.025 (0.024)	-0.132*** (0.034)
June	0.005 (0.024)	0.127*** (0.024)	0.137*** (0.024)	0.142*** (0.031)
July	-0.017 (0.024)	0.108*** (0.024)	0.125*** (0.024)	0.242*** (0.030)
August	0.009 (0.024)	0.118*** (0.024)	0.128*** (0.024)	0.269*** (0.029)
September	0.061** (0.023)	0.036 (0.024)	0.061** (0.024)	0.206*** (0.028)
October	0.072*** (0.024)	0.036 (0.024)	0.037 (0.024)	-0.016 (0.032)
Winter months				
November	0.076*** (0.024)	0.036 (0.024)	0.041* (0.024)	-0.054 (0.037)
December	-0.135*** (0.022)	-0.075*** (0.022)	-0.075*** (0.023)	-0.151*** (0.032)
January	-0.249*** (0.023)	-0.125*** (0.023)	-0.143*** (0.024)	-0.119*** (0.031)
February	-0.321*** (0.025)	-0.203*** (0.025)	-0.258*** (0.026)	-0.380*** (0.036)
March	-0.213*** (0.023)	-0.098*** (0.023)	-0.152*** (0.024)	-0.323*** (0.034)
April	-0.031 (0.023)	0.088*** (0.024)	-0.020 (0.025)	-0.161*** (0.034)
Month dummies	Yes	Yes	Yes	Yes
Regime-zone dummies	No	Yes	Yes	Yes
Constrained transmission	No	No	Weak	Stringent
<i>N</i>	335,462	335,462	325,750	229,461
<i>R</i> ²	0.357	0.386	0.386	0.394
Adjusted <i>R</i> ²	0.356	0.386	0.386	0.393

p*<0.1; *p*<0.05; ****p*<0.01

Coefficients estimated using reservoir fixed effects. Standard errors for the coefficients are reported in parentheses. While certain controls and dummy variables have been omitted for brevity, these are jointly highly significant. The table shows the effect of a one percent increase in producer market share on reservoir volumes, expressed as a percentage of maximum capacity. Columns (1) through (4) denote increasing model controls and restrictions on the sample. See text for details.

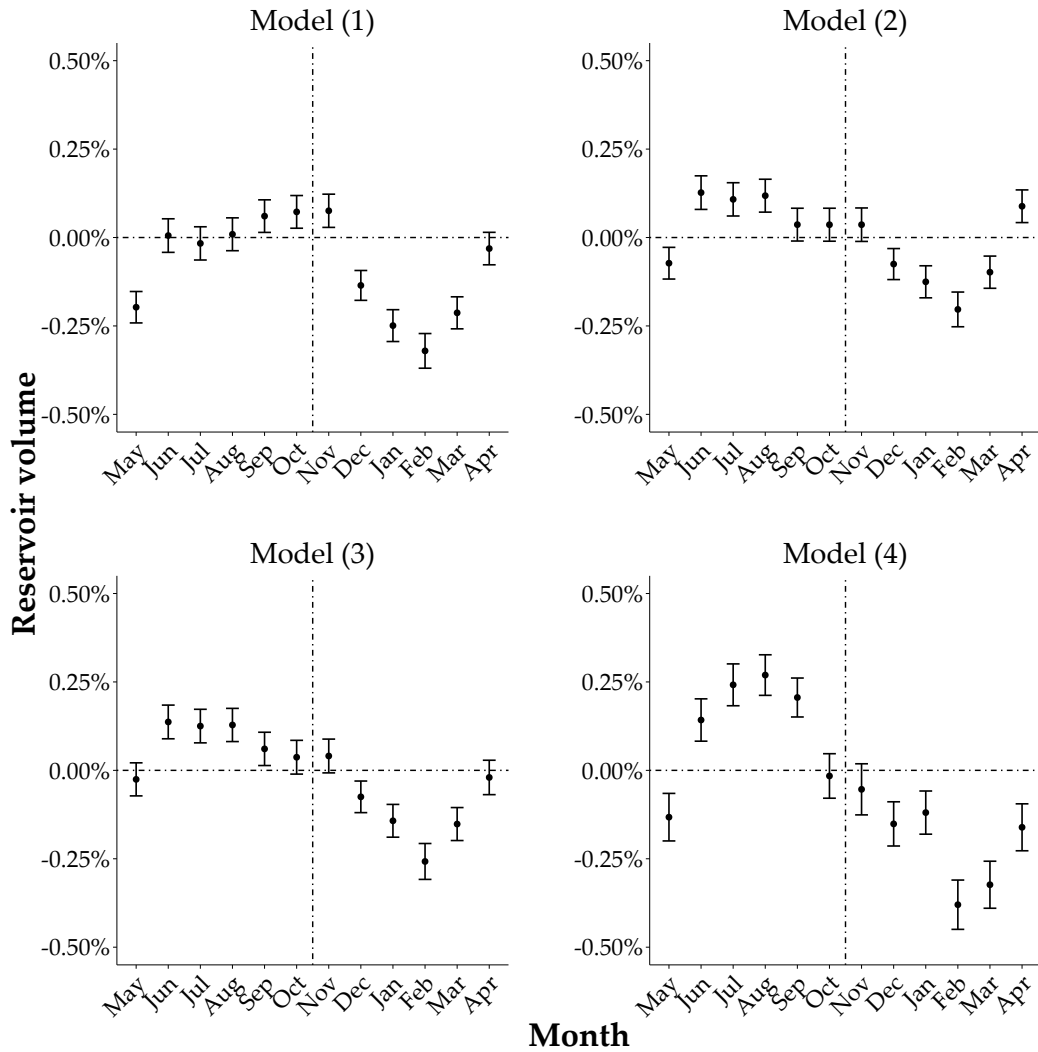


Figure 3.5: Marginal effect of a one percent increase in producer market share on reservoir volumes by month. Dots denote point estimates and the error bars show 95 percent confidence intervals. The stippled horizontal and vertical lines help to identify statistical significance, and summer and winter seasons, respectively. Models are distinguished by increasing controls and restrictions on the sample, with Model (4) being the preferred specification. See Table 3.2 and the text for more detail.

hydropower firms to reallocate their water resources in a distinct, intertemporal manner. Taking the peak summer month of August as an example, a one percent increase in producer market share corresponds to 0.3 percent increase reservoir volumes on average. Conversely, the same change in market share yields a 0.4 percent mean decrease in volumes during the peak winter month of February. All of this accords with the predictions of economic theory and our knowledge of demand elasticities in Norway. Dominant hydropower firms will restrict their production during the Norwegian summer when the demand for electricity is relatively more inelastic, and oversupply during the winter when the elasticity of demand is rela-

tively more elastic.¹⁷

It should be said that the comparative impact of this market share effect on reservoir volumes is modest next to the role of other factors like snow-melt runoff and changes in aggregate electricity demand. Reservoir volumes typically vary over a range of 70 to 85 percent of maximum regulated capacity as one moves from the pre-melt trough in early spring to the autumn peak. Yet, it still suggests that firms within the same bidding area will operate their reservoirs in meaningfully distinct ways when the differences in local market share are large enough. As an example, consider bidding area NO2 from regime M onwards (i.e. January 2010 – December 2013). The regional market shares for the top three firms in this area were stable over the period, with *Statkraft SF* reservoirs accounting for roughly 40 percent of regional volume capacity, *Agder Energi Produksjon* 21 percent, and *Sira Kvina Kraftselskap* 16 percent. Model (4) implies a seasonal differential between the reservoirs belonging to these firms of up to seven percent (allowing for model uncertainty). Indeed, the difference is even greater when we look at the actual reservoir volumes averaged across the three firms for this period — see Figure 3.C.1. Of course, we have not controlled for any confounding factors in this simple example and so it should not be taken as anything more than a piece of descriptive evidence in support of the formal econometric analysis. The fact that the qualitative pattern is consistent with the main empirical results is probably of greater significance.

As a final robustness check, I have also run a set of auxiliary regression models that use an algorithmic routine to compute multi-level fixed effects in an efficient iterative procedure (Gaure, 2013).¹⁸ This effectively allows us to control for higher-level fixed effects (e.g. firms, zones) in a way that does not require specifying factor dummies directly in the regression equation. As expected, the results closely approximate those provided in Table 3.2 and the main qualitative finding is thus unchanged. Higher market share is associated with fuller reservoirs in the summer and emptier reservoirs in the winter. These results are available from the author upon request.

¹⁷These results correspond reasonably well to the aggregate intra-yearly variations in Nordic reservoir levels identified by Kauppi and Liski (2008). Their simulations suggest that, relative to the competitive solution, market power resulted in reservoirs being underutilised during the summer months of 2003–2005. See Figure 5 of their paper.

¹⁸The routine is based on a series of Gauss-Seidel iterations also described by Guimarães and Portugal (2010).

3.6 Concluding remarks

This paper has been motivated by a simple question: How does market power affect firm behaviour in a hydro-based electricity system? The answer, according to economic theory at least, is clear. Dominant hydropower firms can be expected to reallocate their water resources away from periods with relatively inelastic demand for electricity, to periods with relatively elastic demand. This would allow them to recoup higher profits by restricting supply when consumers are least responsive to the resulting price increase. In the case of Norway, this translates to fuller hydro reservoirs in the summer months, followed by lower reservoirs in the winter months.

I test this hypothesis using a uniquely detailed dataset of approximately 500 Norwegian hydropower reservoirs. Exogenous changes to bidding area divisions and transmission constraints allow me to cleanly identify variations in reservoir volumes arising from differences in local market share. Consistent with the predictions of theory, the empirical findings reveal a modest, yet definitive, intertemporal pattern in the way that market power alters utilisation of water resources. Taking my preferred regression specification as a benchmark, a one percent increase in producer market share yields a 0.2–0.3 increase in reservoir volumes during the peak summer months, and a 0.3–0.4 percent decrease during the peak winter months. This market share effect is distinct from the regular seasonal patterns in reservoir volumes that arise from annual snow-melt inflows and so forth.

Constructing a large dataset such as this from disparate sources obviously entails numerous choices in how one compiles the data. For instance, an implicit simplifying assumption in my empirical analysis has been that firm ownership of reservoirs remains constant over the review period. This may not be an entirely benign assumption and could mask some important results if there was significant merger and acquisition activity during that time. On the other hand, the Norwegian electricity sector had been liberalised for nearly two decades by the beginning of the study period. The relative maturity of the market should at least give us some confidence regarding the stability of the competitive structure over this period. In a similar vein, the effects of partial- and cross-ownership have not been considered (c.f. Amundsen and Bergman, 2002). Fully accounting for these issues is a potentially fruitful topic for future research.

Such caveats notwithstanding, the results of this paper may be interpreted as empirical vindication of the underlying theory. More to the point, they can help guide

policy and competition authorities in effectively regulating electricity markets like the Norwegian system, where hydropower comprises a major share of generation, or transmission constraints create opportunities for producers to exercise local market power.

References

- Abowd, John M., Francis Kramarz, and Simon Woodcock (2008). "Econometric Analyses of Linked Employer–Employee Data," in László Mátyás and Patrick Sevestre (eds.), "The Econometrics of Panel Data: Fundamentals and Recent Developments in Theory and Practice," chapter 22, Springer Science & Business Media, pp. 727–760.
- Amundsen, Eirik S and Lars Bergman (2002). "Will Cross-Ownership Re-Establish Market Power in the Nordic Power Market?" *The Energy Journal*, 23: 73–95.
- Balli, Hatice Ozer and Bent E. Sørensen (2013). "Interaction effects in econometrics," *Empirical Economics*, 45(1): 583–603.
- Borenstein, Severin, James Bushnell, and Steven Stoff (2000). "The competitive effects of transmission capacity in a deregulated electricity industry," *RAND Journal of Economics*, 31(2): 294–325.
- Bresnahan, Timothy F. (1982). "The oligopoly solution concept is identified," *Economics Letters*, 10(1): 87–92.
- Bye, Torstein and Petter Vegard Hansen (2008). "How do Spot prices affect aggregate electricity demand?" Statistics Norway Discussion Papers No. 527. Available: <http://www.ssb.no/a/publikasjoner/pdf/DP/dp527.pdf>.
- Bye, Torstein and Einar Hope (2005). "Deregulation of electricity markets: the Norwegian experience," *Economic and Political Weekly*, 40(50): 5269–5278.
- ENTSOE (2015). *Principles for determining the transfer capacities in the Nordic power market*, European Network of Transmission System Operators for Electricity, Available: <https://www.nordpoolspot.com/globalassets/download-center/tso/principles-for-determining-the-transfer-capacities.pdf>.
- Førsund, Finn R. (2015). *Hydropower Economics*, International Series in Operations Research and Management Science, Springer, second edition.

- Fridolfsson, Sven-Olof and Thomas P. Tangerås (2009). "Market power in the Nordic electricity wholesale market: A survey of the empirical evidence," *Energy Policy*, 37(9): 3681–3692.
- Gaure, Simen (2013). "lfe: Linear group fixed effects," *The R Journal*, 5(2): 104–117.
- Guimarães, P. and P. Portugal (2010). "A simple feasible procedure to fit models with high-dimensional fixed effects," *Stata Journal*, 10(4): 628–649(22).
- Hansen, Petter Vegard (2004). "Regional electricity spot price responses in Norway," Statistics Norway Discussion Papers No. 2004/13. Available: http://www.ssb.no/a/english/publikasjoner/pdf/doc_200413_en/doc_200413_en.pdf.
- Hansen, Petter Vegard (2009). "Inflow uncertainty in hydropower markets," *The Scandinavian Journal of Economics*, 111(1): 189–207.
- Hjalmarsson, Erik (2000). "Nord pool: A power market without market power," Göteborg University. School of Business, Economics and Law. Working Papers in Economics, no. 28. Available: http://www.ssb.no/a/english/publikasjoner/pdf/doc_200413_en/doc_200413_en.pdf.
- Johnsen, Tor Arnt (2001). "Demand, generation and price in the Norwegian market for electric power," *Energy Economics*, 23(3): 227–251.
- Johnsen, Tor Arnt, Shashi Kant Verma, and Catherine D Wolfram (1999). "Zonal pricing and demand-side bidding in the Norwegian electricity market," Program on Workable Energy Regulation (POWER) working paper series, No. PWP-063. Available: <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.16.2190&rep=rep1&type=pdf>.
- Joskow, Paul L. (2008). "Lessons learned from electricity market liberalization," *The Energy Journal*, 29(2): 9–42, Special Issue. The Future of Electricity: Papers in Honor of David Newbery.
- Joskow, Paul L. and Jean Tirole (2000). "Transmission rights and market power on electric power networks," *RAND Journal of Economics*, 31(3): 450–487.
- Kauppi, Olli and Matti Liski (2008). "An Empirical Model of Imperfect Dynamic Competition and Application to Hydroelectricity Storage," MIT-CEEPR WP-2008-012. Available: <http://web.mit.edu/ceepr/www/publications/workingpapers/2008-011.pdf>.
- Lau, Lawrence J. (1982). "On identifying the degree of competitiveness from industry price and output data," *Economics Letters*, 10(1): 93–99.

- Mas-Colell, Andreu, Michael D. Whinston, and Jerry R. Green (1995). *Microeconomic Theory*, Oxford University Press.
- Mathiesen, Lars, Jostein Skaar, and Lars Sørgard (2013). "Electricity Production in a Hydro System with a Reservoir Constraint," *The Scandinavian Journal of Economics*, 115(2): 575–594.
- Mirza, Faisal Mehmood and Olvar Bergland (2012). "Transmission congestion and market power: the case of the Norwegian electricity market," *The Journal of Energy Markets*, 5(2): 59–88.
- NordREG (2014). "Nordic Market Report 2014: Development in the Nordic Electricity Market," Report 4/2014, Nordic Energy Regulators, Available: <http://www.nordicenergyregulators.org/wp-content/uploads/2014/06/Nordic-Market-Report-2014.pdf>.
- SSB (2015). "Electricity, annual figures, 2013," *Statisttics Norway*, 25 March 2015. Available: <http://www.ssb.no/en/energi-og-industri/statistikker/elektrisitetaar/aar/2015-03-25>.
- Steen, Frode (2004). "Do bottlenecks generate market power? An empirical study of the Norwegian electricity market," Norwegian School of Economics. Discussion Paper 26/03. Available: <http://www.nhh.no/Files/Filer/institutter/sam/Discussion%20papers/2003/26.pdf>.
- Wickham, Hadley (2014). "Advanced R," in John M. Chambers, Torsten Hothorn, Duncan Temple Lang, and Hadley Wickham (eds.), "The R Series," Chapman & Hall / CRC, 1st edition.

Appendix

3.A Previous Elspot regimes

3.B Market shares of top three producers

**3.C Reservoir volumes for top three NO₂ producers,
Jan 2010 – Dec 2013**

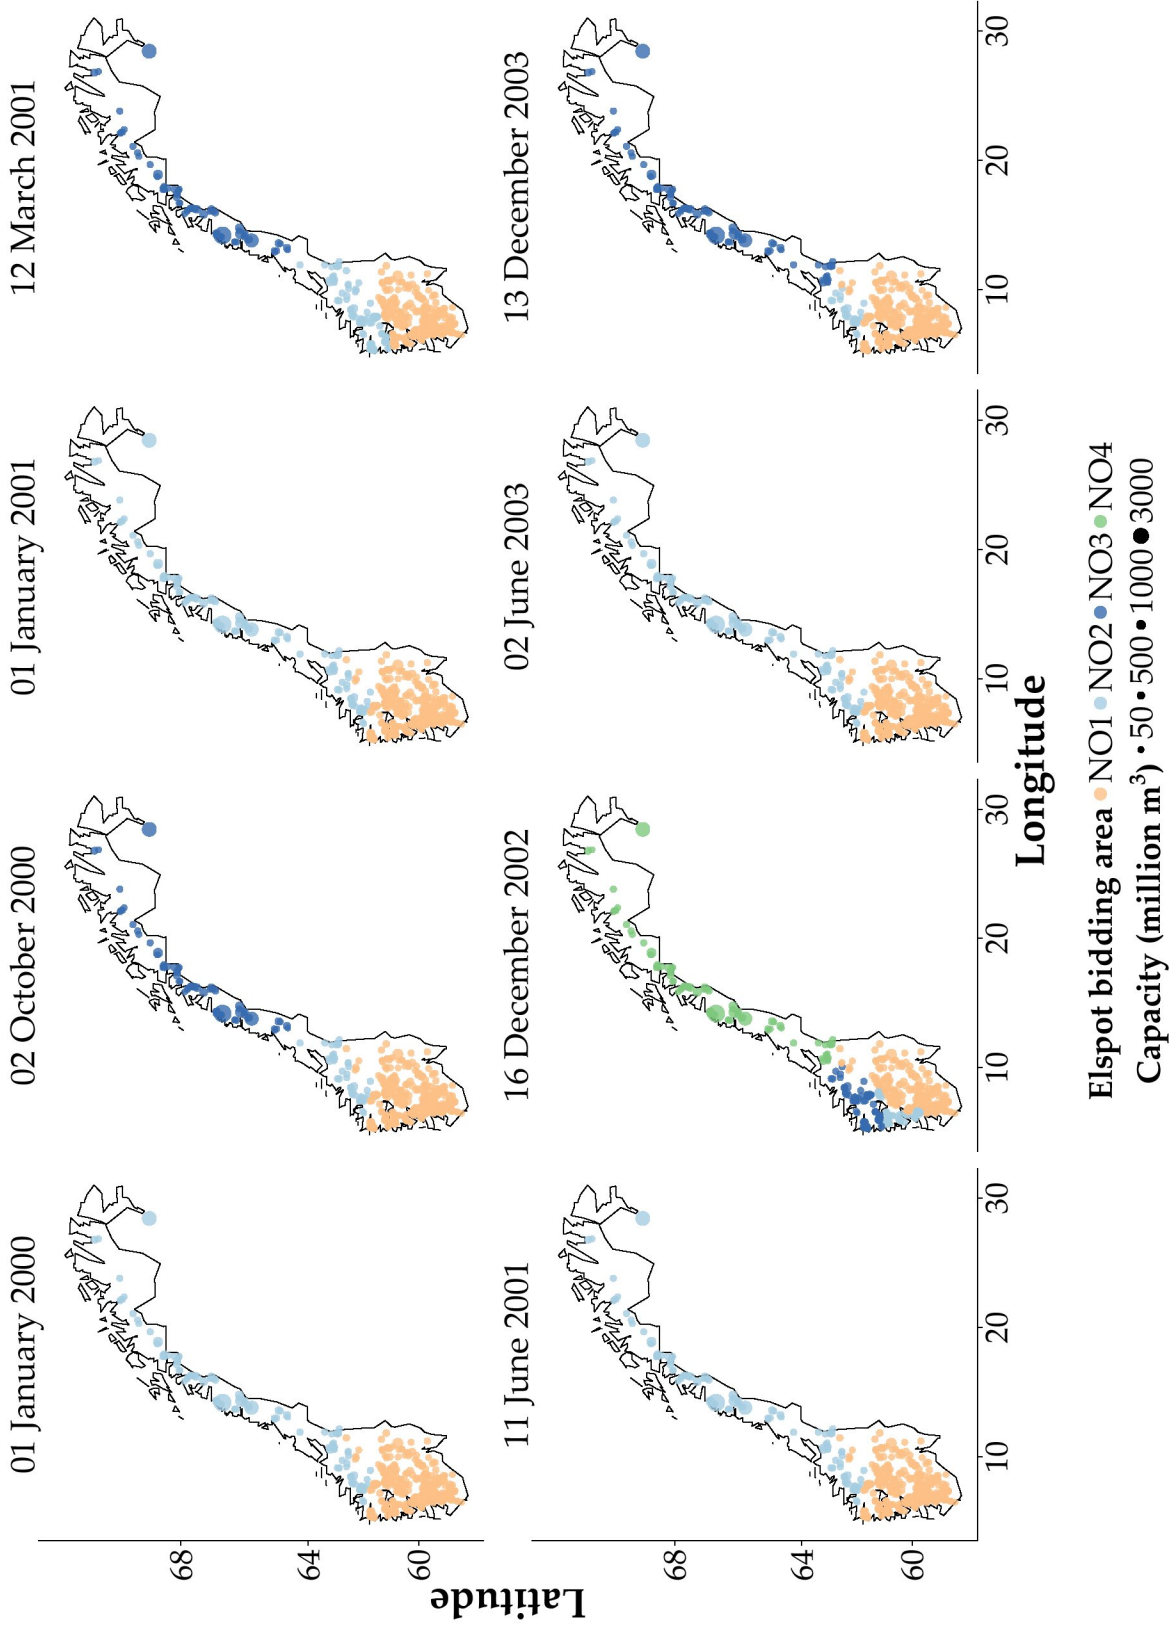


Figure 3.A.1: Elspot regimes A – H (1 January 2000 – 13 December 2003)

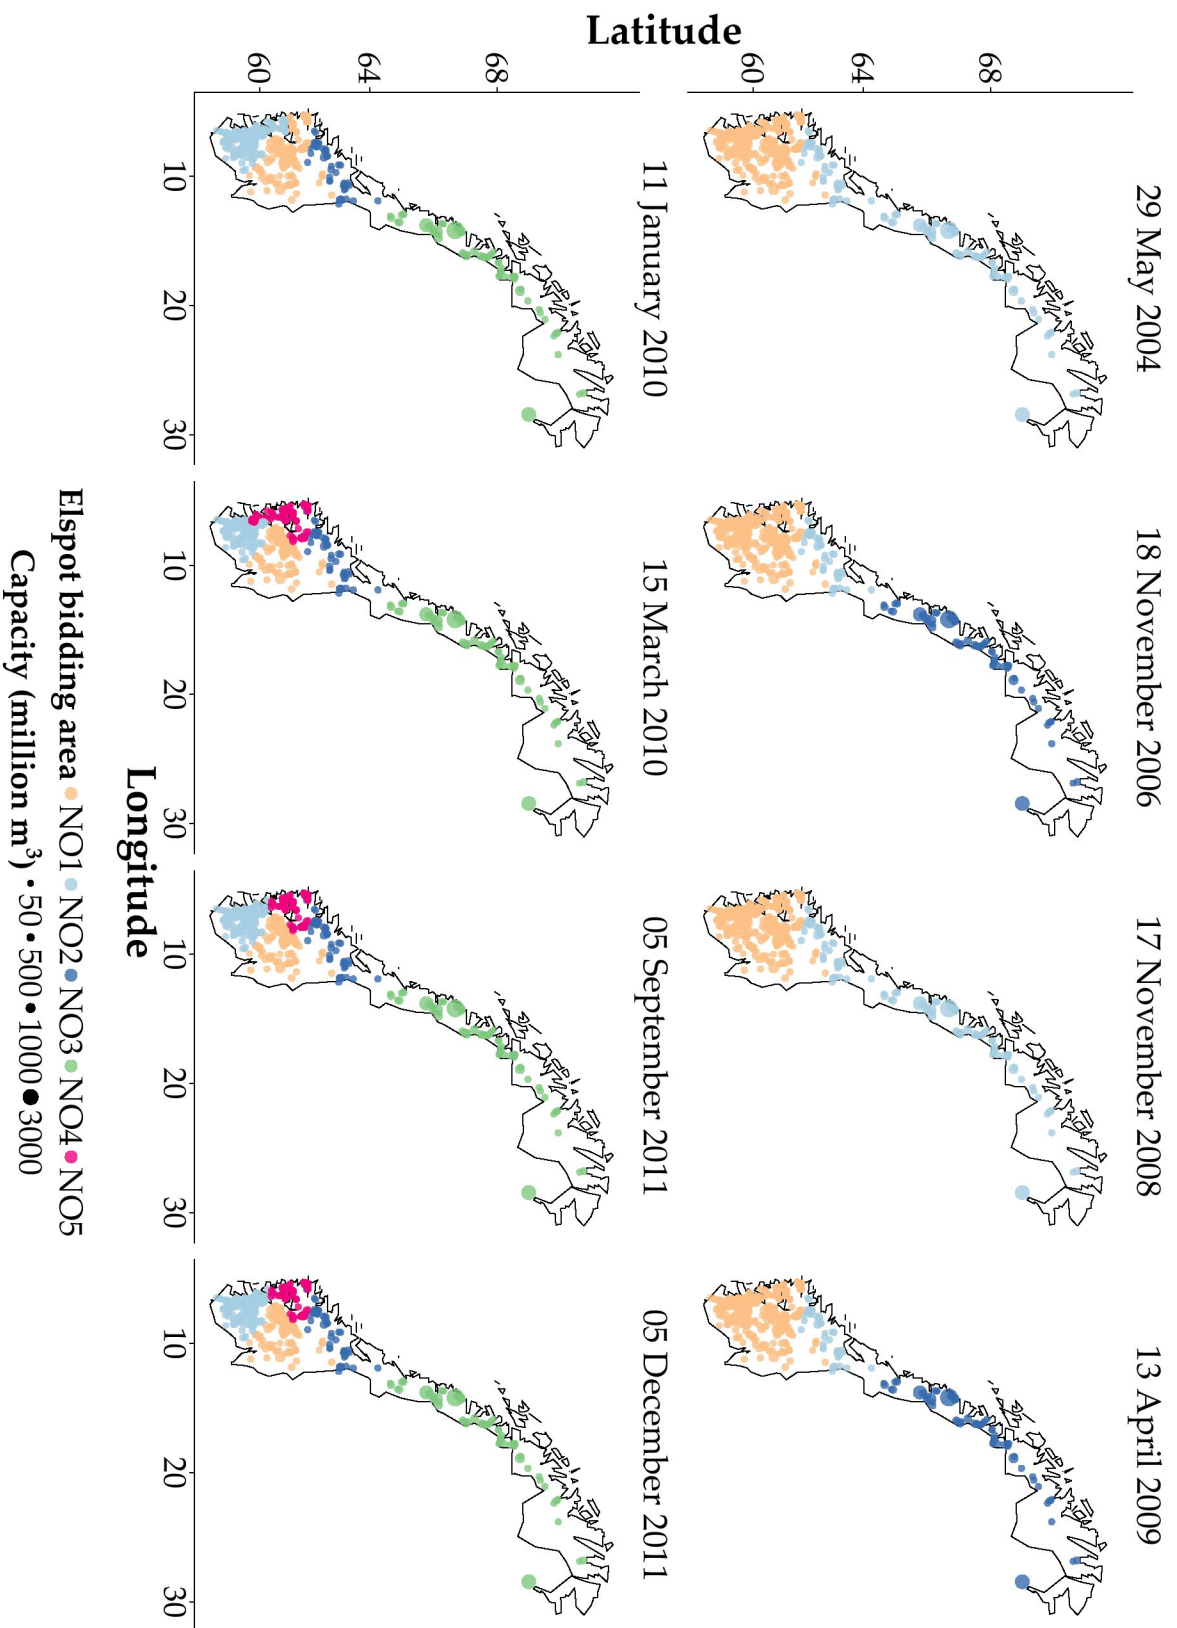


Figure 3.A.2: Elspot regimes I - P (29 May 2004 – 5 December 2011)




Figure 3.B.1: Market shares of top three producers in each bidding area (Elspot regimes A – H)




Figure 3.B.2: Market shares of top three producers in each bidding area (Elspot regimes I – Q). *Note: Regime P excluded for readability as it is identical to regime O.*

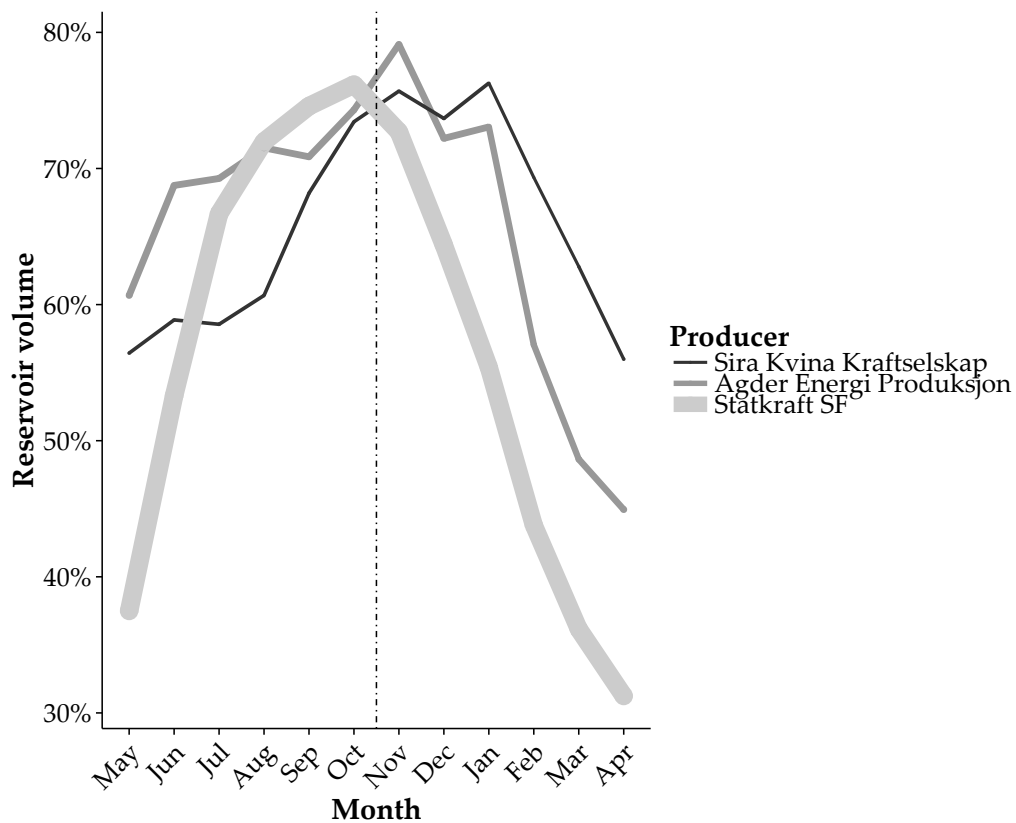


Figure 3.C.1: Mean monthly reservoir volumes for the top three producers in the NO2 bidding area from January 2000 to December 2013 (i.e. regimes M to Q). The width of the individual lines is representative of the relative regional market shares over this period: *Statkraft SF* (40 percent), *Agder Energi Produksjon* (21 percent), and *Sira Kvina Kraftselskap* (16 percent).