Hydro power. Market might.

Grant R. McDermott *

This version: 2 January 2020.

Abstract

A central tenet of economic theory is that market power induces deadweight loss. This claim rests on an assumption that is difficult to verify empirically. Namely, that dominant firms produce less than the social optimum. I provide evidence of such restrictive behaviour using a rich dataset of Norwegian hydropower firms. The research design exploits exogenous variation in market power, arising from transmission bottlenecks and the formation of localized electricity markets. The unique production traits of hydropower production further helps to avoid empirical complications associated with marginal cost estimation and endogenous variation in the supply mix. This allows me to identify the causal impact of market power on firm behaviour without imposing strong structural assumptions on the data. I show that a firm's gaining pivotal status may cause it to withhold production by as much as a two to five percent. My results suggest that even nominally competitive markets are susceptible to strategic manipulation and welfare losses.

JEL Codes: Q25, Q41, L12, L13

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

*Department of Economics, University of Oregon. Email: grantmcd@uoregon.edu. This paper has benefited from comments and suggestions by Kyle Meng, Christopher Costello, Lucas Davis, Frederik Noack, Øivind A. Nilsen, Gunnar Eskeland, Morten Sæthre, Linda Nøstbakken, Timo Goeschl, Corbett Grainger, Gerard Doorman, and various seminar participants at UC Santa Barbara, the AERE Summer Conference, UMass Amherst, the University of Delaware, the University of Hawaii, and the University of Oregon. I am deeply thankful to Svein Taksdal and other members of the Norwegian Water and Energy Directorate for assisting me with data collection. Funding from CenSES (Research Council of Norway project no. 209697) is gratefully acknowledged.

1 Introduction

The great recession and its aftermath have reignited concerns about market power. A growing body of research now suggests that market power is to blame for a variety of economic ills. These range from stagnant wages and falling labour productivity shares (e.g. Autor et al., 2017; Azar et al., 2017; Benmelech et al., 2018), to a slowdown in aggregate output (e.g. De Loecker et al., 2019). Moreover, high industry concentrations and record markups are being observed at both a national and global level (Azar et al., 2018; De Loecker and Eeckhout, 2018), as well as in both traditional and nascent markets (Benmelech et al., 2018; Dube et al., 2018).

The central mechanism by which market power induces deadweight loss to society is, of course, a reduction in output. Economic theory tells us that dominant firms can raise prices (and thus profits) by restricting output. Yet, as uncontroversial as this idea is among economists, it is surprisingly hard to verify empirically. Firms do not enter, exit or come to dominate industries by chance. Nor do they acquire or merge with other firms at random. They do so based on their expectations about future profits and market conditions. Because these factors are unobservable to the empirical researcher, typically they must be modeled or assumed in a structural setting.[1] All of this is to say simply that firm behaviour is highly endogenous. We cannot measure how output or markups respond to, say, industry concentration and naively assign a causal interpretation.[2]

It is against this backdrop that the present paper describes a rare empirical setting in which the causal effect of market power on firm output can be gleaned *directly* from the data. In particular, I demonstrate how a large number of real-world firms respond to a series of on-going and plausibly exogenous shocks to the competitive structure of their industry. Moreover, that they do so in a manner that is remarkably consistent with the predictions of economic theory. My reduced-form estimates are easily interpretable and require minimal assumptions or imposed structure. While I will argue that my findings are generalisable to other sectors of the economy afflicted by market power,

[1]The notion that dominant firms will restrict output may be regarded as a *precept* of the typical structural model in industrial organization (I/O). Insofar as it arises directly from assumptions about profit maximization and equilibrium market behaviour (Reiss and Wolak, 2007).

[2]Even accurately measuring or drawing associative links between these variables may be problematic; as evidenced by a recent symposium of leading I/O scholars in the *Journal of Economic Perspectives* (Basu, 2019; Berry et al., 2019; Syverson, 2019).

I rely on a fortuitous confluence of factors that are manifest in a particular industry. Namely, the Norwegian hydropower industry.

Hydropower is the world's foremost source of non-fossil energy. It accounts for seven percent of global primary energy and 16 percent of global electricity generation (BP, 2016). This role is even more pronounced in countries like Norway — the setting for this paper — where it accounts for over 95 percent of national electricity generation.[3] It is hydropower's unique combination of production traits, however, that are of special interest to us here. The three traits that stand out are as follows. First, hydropower yields a homogeneous end-good in the form of electricity. This negates potential empirical complications related to product differentiation and branding effects. Second, unlike other modes of electricity production, hydropower is dispatchable and load-following. Which is to say that plants can immediately adjust output in response to market conditions.[4] Third, and perhaps most importantly, variable costs in hydropower production are negligible. Water arrives exogenously from nature, while employee wages and maintenance fees are best regarded as fixed costs independent of output. We are thus able to abstract from the complications associated with estimating firms' marginal cost curves. I shall expand on and clarify these claims at various points in the text. For the moment, though I hope to convey that hydropower presents specific empirical advantages, which have the potential to simplify important general economic questions about firm behaviour and output.

The other key component of this paper's contribution is research design. I exploit plausibly exogenous variation in local market conditions that arise due to transmission bottlenecks in the Norwegian electricity grid. These transmission bottlenecks lead to the formation of separate bidding areas, which manifest as a shock to the market power that firms command within their assigned areas. Together with the predominance of hydropower production in the region as a whole, it is this quasi-random allocation that allows me to cleanly identify the causal impact of market power on firm behaviour. Indeed, the fact that individual hydropower plants are randomly assigned to separate bidding areas during different periods — and that these bidding areas themselves are changing over time — means that my empirical setting constitutes a rare natural exper-

[3]Other prominent examples include Venezuela (70 percent of electricity generation), Brazil (65 percent), Canada (60 percent), and New Zealand (55 percent).

[4]C.f. Non-dispatchable renewables like wind or solar, which produce energy according to environmental conditions and cannot be controlled by system operators. On the other end of the scale, nuclear or coal plants are dispatchable but are typically run at or near maximum capacity regardless of immediate market conditions.

iment in which the actual *definition* of the market is constantly changing. Treatment is varying in both frequency (how often transmission constraints are binding) and distribution (how much market power is conferred to individual firms as a result).

These features allow me to make several contributions to the literature. First and foremost, rather than inferring noncompetitive outcomes through simulated counterfactuals or aggregate measures such as wholesale electricity prices, I am able to look at the behaviour of individual hydropower firms directly. The causal relationship between the producer's decision and market power is therefore determined at the firm (and, indeed, plant) level. This refinement not only establishes a close correspondence between the empirical set-up and underlying theory, but also permits the use of a conceptually straightforward regression framework in the form of panel fixed effects. My analysis imposes little-to-no structure on the data, nor does it hinge on strong assumptions about the market relationships in question. I also consider a much longer time period than many previous studies, with my sample running from 2000 until the end of 2013. In so doing, I hope to shed light on the way that dominant firms strategically utilise their resources, not just in the short-term, but in response to changing market conditions over the course of months, seasons and years. To the best of my knowledge, this paper is one of the first to provide direct empirical evidence — absent structural assumptions or simulated counterfactuals — of firm's strategically withholding production in a dynamic market setting.

To preview my results, I find that an increase in local market power leads to a distinctive pattern of resource shifting. Firms that gain pivotal status during the most inelastic demand periods of the year, for example, maintain demonstrably fuller reservoirs than their competitors. The effect is anywhere between two and five percent, depending on the econometric specification. The reverse is true during relatively elastic demand periods. Such strategic behaviour is apiece with the predictions of economic theory, which dictates that dominant hydropower firms will maximize profits by engaging in intertemporal price discrimination. Put simply, firms can increase profits by withholding supply in periods with relatively inelastic demand, driving up prices when consumers are least responsive to such changes.

The remainder of the paper is organised as follows. Section 2 provides a brief overview of related studies and further contrasts their approaches with my own. Section 3 describes the key institutional features of the Norwegian electricity market, while section 4 covers the basic theoretical framework. Section 5 introduces the dataset and discusses how various source materials were merged into a unified set. Section 6 discusses the econometric strategy for causal inference. Section 7 presents the empirical results and Section 8 concludes.

2 Related literature

Electricity markets feature prominently in the literature on empirical industrial organisation (e.g. Wolfram, 1999; Joskow and Tirole, 2000; Borenstein et al., 2000, 2002; Wolak, 2003; Müsgens, 2006; Sweeting, 2007; Puller, 2007; Mansur, 2007; Hortaçsu and Puller, 2008; Reguant, 2014; Davis and Hausman, 2016). In part this reflects electricity's fundamental importance to the overall economy, as well as the coincident wave of market liberalisation that swept the industry in the 1990s and 2000s (Joskow, 2008). However, it also speaks to the specific advantages that electricity offers as a subject for economic inquiry relative to other goods and sectors — product homogeneity, reduced complications in terms of branding, advertising, and so forth (Borenstein, 2016). Yet despite such advantages, much of this existing literature has still had to rely on a combination of aggregate data, simulation methods, model calibration and/or strong structural assumptions to support its findings. In contrast, my purely empirical approach leverages a mix of detailed plant-level data, hydropower's unique production characteristics and quasi-experimental variation in the economic parameters of interest, to identify evidence of noncompetitive behaviour in a dynamic market setting.

The studies closest to the present paper in terms of data richness and approach are those by Puller (2007), Hortaçsu and Puller (2008), Reguant (2014), and Davis and Hausman (2016). Both Puller and Davis and Hausman focus on the California electricity market, where they use plant-level data to document noncompetitive behaviour during the region's post-liberalisation energy crisis and following the shutdown of a major nuclear facility, respectively. Hortaçsu and Puller tread a similar path in evaluating the restructured Texas electricity market, where they find that larger firms are closer to (static) profit-maximising benchmarks than their smaller counterparts. In turn, Reguant uses Spanish data to show that startup costs in electricity generation can play an important role in rationalizing apparent deviations from optimal bidding behaviour.

Compared to these studies, the present work benefits from richer variation in market power over a longer period of time, as well as the specific empirical advantages afforded by my exclusive focus on hydropower generation. The firms in my dataset are directly and immediately comparable since they all use the same production technology and face the same (negligible) marginal costs. Similarly, while ignoring startup costs can lead to biased estimates of market power as per Reguant, the inherent dispatchability of hydropower renders such concerns moot in the present context. I would argue further that my identification strategy rests upon a more fundamental change in the parameters of interest. Rather than comparisons with simulated pricing data (Puller; Hortaçsu and Puller; Reguant) or a one-time shock to generation capacity (Davis and Hausman), I exploit recurring changes in the actual *definition* of the market being studied.[5] Despite these differences and the fact that we focus on different regions, the present analysis is still conducted in the same spirit as these four studies. Our collective ability to identify noncompetitive behaviour by firms is greatly enhanced by a combination of detailed data, quasi-experimental settings and transparent modeling approaches.

In addition to the studies named above, a number of authors have specifically examined the competitive structure of the Norwegian — and broader Nordic — electricity market. These 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. I too focus on local market power that arises from transmission bottlenecks.[6] However, the richness of my dataset allows me to go much deeper by tracking firm-level responses to changing bidding area configurations over time. These changes provide a key source of additional variation in local market power that enable me to cleanly identify its impact on hydro firm behaviour *vis-à-vis* the management of individual reservoirs. Moreover, all previous studies of the Nordic market rely on aggregate data. This places strong restrictions on the methods that can be used to 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 inherently dynamic nature of hydropower production. As Fridolfsson and Tangerås (2009,

[5] This claim will be motivated in greater detail in subsequent sections of the paper.

[6] Theoretical contributions to this literature and related empirical applications in other electricity markets include those by Joskow and Tirole (2000), Borenstein et al. (2000), Davis and Hausman (2016), and Bigerna et al. (2016).

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."[7]

Finally, the empirical analyses of the Nordic and other electricity markets to date have tended to focus on short-run deviations from competitive prices. Or, they have taken the form of single event studies that compare market conditions before and after an economic shock. Relatively little is known about the strategic behaviour of dominant firms over the medium- to long-term, much less how they respond to changing conditions in their markets over time. I attempt to shed light on these dynamic uncertainties in the present paper.

3 The Norwegian electricity market

With the enactment of the Energy Act of 1990, Norway become one of the first countries to deregulate and liberalise its electricity sector (Bye and Hope, 2005; Joskow, 2008). The other Nordic nations would steadily follow suit. By 2000, Norway, Sweden, Finland, and Denmark together formed the world's first — and still largest — multinational power exchange, *Nord Pool*. The primary market of the Nord Pool exchange is *Elspot*, a day-ahead market for the physical delivery of electricity. Elspot functions according to classic auction principles in which hourly supply and demand bids are first aggregated and then matched to determine a market clearing price, also known as the system price. In the absence of transmission constraints, all electricity is traded at the system price. However, when transmission constraints are binding, the Elspot market is split up into distinct bidding areas. During such periods, each bidding area effectively becomes its own market and is typically characterised by a distinct area price. This in turn confers potential local market power to dominant firms within those areas.

As per Figure 1, Norway is currently composed of five Elspot bidding areas: NO1 (east), NO2 (south), NO3 (mid), NO4 (north), and NO5 (west). Importantly, this configura-

[7]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[...]"

tion of bidding areas has been changed multiple times over the last decade and a half (c.f. Appendix B). New bidding areas have been added or removed and the boundaries between existing areas have been redrawn, as the Norwegian system operator, Statnett, has tried to optimise electricity access under the inherent physical and technical constraints of large-scale power grid. I will discuss this issue in more depth in the empirical section of the paper. For the moment, it need simply be said that such "regime" changes provide an additional layer of plausibly exogenous variation in local market power and are thus highly advantageous from an empirical perspective. The evolving configuration of bidding areas, in concert with binding transmissions constraints and data on individual reservoirs, will ultimately allow me to identify the causal effect of local market power on hydro firm behaviour by randomly assigning individual plants to distinct markets.

Turning to the broader characteristics of the Norwegian electricity market, hydropower dominates at both the national and regional level. Norwegian reservoirs generated around 130 TWh of electricity in 2013. This corresponds to over 95 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). For its part, Norwegian electricity demand tends to fluctuate significantly depending on the time of day, the day of week, and, indeed, the time of year. While large portions of this demand are not substitutable, the Norwegian electricity sector is designed to foster consumer responsiveness whenever conditions allow. Most buildings are equipped with in-house electricity meters that provide real-time information on consumption. Moreover, approximately 85–90% of end-user contracts are either tied directly to the spot price, or come in the form of variable price contracts that are highly-correlated with the spot price and may be terminated at short notice (Bye and Hansen, 2008). Norwegian consumers are thus directly exposed to the changes in the spot prices and can be expected to react accordingly.

The principal focus of this paper is strategic behaviour among hydropower firms over longer periods of time. It therefore makes sense to abstract from short-term fluctuations (e.g. day versus night, weekdays versus weekends) and concentrate on variations in demand over the course of months and seasons. With that in mind, winter electricity demand in Norway is approximately double that of summer, primarily because of indoor heating requirements. However, electrical heating is relatively easy to substitute with indoor fireplaces, fuel-burning heaters, warmer clothing, better insulation, etc. This permits greater flexibility among consumers, whereas summer electricity de-

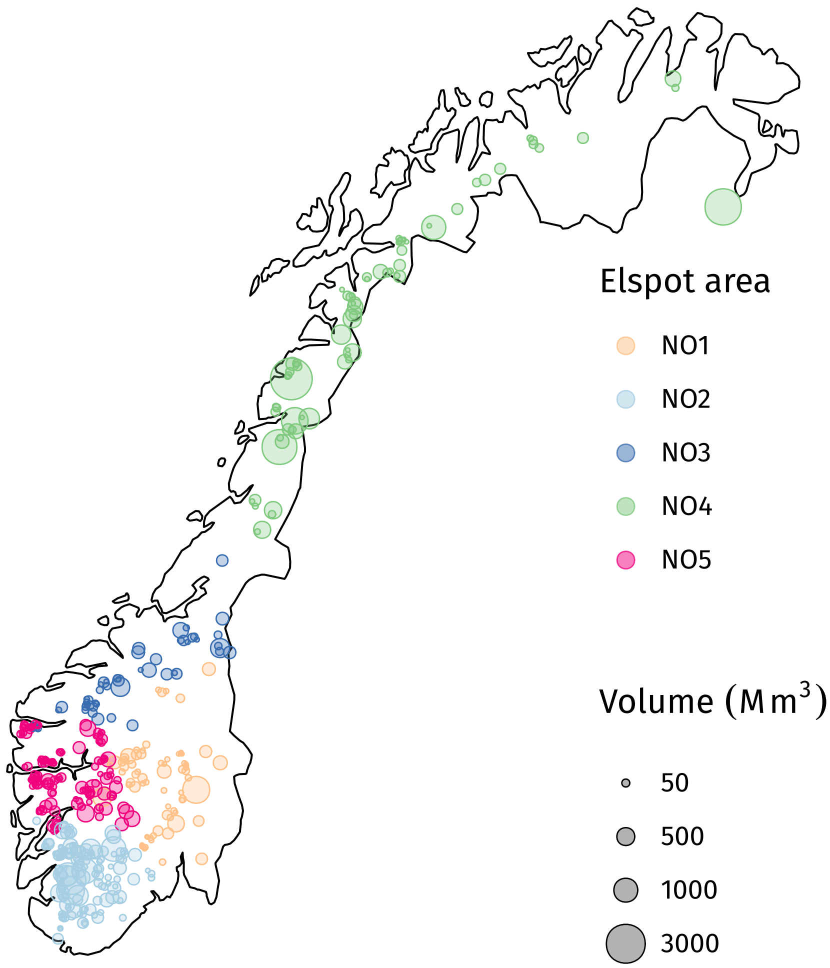

FIGURE 1: NORWEGIAN HYDROPOWER RESERVOIRS

Notes: Elspot bidding area configuration as of 2 December 2013 (i.e. regime Q).

mand is dominated by technical end-uses that do not allow for easy substitution. These seasonal differences in substitution and adjustment possibilities contribute to a perhaps surprising finding: Demand elasticities are significantly *higher* during the Norwegian winter than they are during summer (Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008).[8] I confirm this finding using detailed bid curve data in Section 5.2. However, it is first useful to cover the theoretical basis of dominant hydro firm behaviour, and why

[8]The phenomenon of a relatively more elastic winter electricity demand is quite common in high-latitude countries, where substitutable heating requirements account for a large portion of overall energy consumption. For example, Genc (2016) shows that this is true for Canada.

seasonal variations in demand elasticities are fundamental to their strategy.

4 Theoretical framework

As described in the introduction, variable costs are negligible in a hydropower system. Instead, production is largely determined by the opportunity cost of water — the so-called *water value*. Water inflows in a country like Norway are highly seasonal. Most reservoirs fill up rapidly from the end April as the result of glacial melt and spring rains. To a first approximation, however, water is a durable good that can be stored for long periods of time. The producer's decision thus collapses into one of how much water to use today versus save for tomorrow? Faced with this kind of dynamic allocation problem, let us consider how dominant firms might act strategically.

Førsund (2015) outlines the various ways in which market power affects producer behaviour in a hydropower system. 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,

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

where $p_t\left(q_t\right)$ 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, its water equivalent), and W is the known water endowment for the monopolist's reservoir. Let us further assume that period 1 is summer and period 2 is winter. This assumption has no bearing on the theoretical results, but will become useful as a reference for the empirical setup later in the paper.

The necessary first order conditions for profit maximization are

$$\frac{\partial L}{\partial q_t} = {p_t}'\left(q_t\right) \cdot q_t + p_t\left(q_t\right) - \lambda \leq 0 \qquad (2)$$
$$\left(= 0 \quad \text{for } q_t > 0\right)$$

and

$$\lambda \geq 0$$
$$\left(= 0 \quad \text{for } \sum_{t=1}^{2} q_t < W\right). \qquad (3)$$

The parameter λ denotes the shadow price on stored water, i.e. positive if the resource constraint in equation (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.[9] The first order conditions may then be written as

$$p_1\left(q_1\right)\left(1 + \frac{1}{\epsilon_1}\right) = p_2\left(q_2\right)\left(1 + \frac{1}{\epsilon_2}\right) = \lambda, \qquad (4)$$

where $\epsilon_t = \frac{p_t}{q_t}\frac{\partial q_t}{\partial p_t} < 0$ is the price elasticity of demand. Rearranging equation (4), it is easy to see that prices depend on the relative demand elasticities in each period. For example, we would have

$$p_1\left(q_1\right) > p_2\left(q_2\right) \quad \text{if} \quad \left|\epsilon_1\left(q_1\right)\right| < \left|\epsilon_2\left(q_2\right)\right|. \qquad (5)$$

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

$$q_1 < q_2 \quad \text{if} \quad \left|\epsilon_1\left(q_1\right)\right| < \left|\epsilon_2\left(q_2\right)\right|. \qquad (6)$$

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

[9] If 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.

under perfect competition, where production and prices are equalised across periods.[10] We can further generalise the difference between monopoly and perfect competition in this simple setup as

$$q_t^M < q_t^C \text{ and } q_{\hat{t}}^M > q_{\hat{t}}^C \quad \text{if} \quad |\epsilon_t(q_t)| < |\epsilon_{\hat{t}}(q_{\hat{t}})|, \tag{7}$$

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. Market power not only causes electricity 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 more elastic periods.

The generalized version of the above model with more than two periods follows the same pattern. Producers maximize profits by reallocating production away from relatively inelastic demand periods. Moreover, 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.[11] Market power leads to a strategy of shifting water use away from relatively inelastic demand periods to relatively elastic ones. Moreover, as described in Section 3, electricity demand in Norway is relatively more inelastic during the summer months. Our testable hypothesis is therefore that dominant firms will tend to maintain fuller reservoirs in summer and lower reservoirs in winter.

[10]By definition, the price elasticity of demand facing a competitive firm is always perfectly elastic, i.e. $\epsilon \to \infty$. 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).

[11]For 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.

Table 1: Summary of data sources

Data	Description	Period	Frequency	Source
Reservoirs	Volumes and levels for the 500 largest reservoirs in Norway.	2000–2013	Daily	NVE[a]
Plants & firms	Regulatory limits, GIS data, etc.	N/A	N/A	NVE[b]
Elspot areas	Bidding area divisions and changes.	N/A	N/A	NVE[c], Nord Pool[d]
Electricity (I)	Prices, flows and transmission capacities.	2000-2013	Hourly	Nord Pool[e]
Electricity (II)	Bid curve data.	2014–2017	Hourly	Nord Pool[f]
Weather	Meteorological station data.	2000–2013	Hourly	NCDC[g]

[a] Available upon request. See: http://api.nve.no/doc/hydrologiske-data (Norwegian).
[b] http://gis3.nve.no/link/?link=vannkraft
[c] http://gis3.nve.no/link/?link=nettanlegg
[d] http://nordpoolspot.com/globalassets/download-center/elspot/elspot-area-change-log.pdf
[e] Proprietary. See: https://www.nordpoolgroup.com/services/Power-data-services/Product-details.
[f] https://www.nordpoolgroup.com/elspot-price-curves
[g] ftp://ftp.ncdc.noaa.gov/pub/data/noaa/

5 Data

This paper uses a novel dataset of Norwegian hydropower firms, reservoirs and electricity data. The data have been constructed from a variety of sources — both public and proprietary — and are summarised in Table 1. Below I describe the original data sources in greater detail, as well as the methods that have been used to merge these disparate parts into a unified dataset.

5.1 Hydropower reservoirs, plants and firms

Time series data for a large number of Norwegian hydropower reservoirs — representing approximately 90 percent of total system capacity — were obtained from the Norwegian Water and Energy Directorate (NVE). Most reservoirs are observed at a daily resolution from January 2000 to December 2013, and contain water readings in terms of both volumes (million m^3) and levels (m). The reservoirs in my dataset are subject to regulation regarding their maximum and minimum holding capacities. Operating firms are legally required to keep their reservoirs within these limits so as to guard against flooding and environmental degradation. To ensure comparability between the different reservoir sizes in my dataset, the reservoir data are therefore normalised as percentages of their respective maximum regulated capacities. After various data cleaning steps described in Appendix A, I am left with a panel dataset of 500 Norwegian hydropower reservoirs that comprises nearly 2 million daily observations. Each reservoir in my dataset is furthermore linked to a set of relevant covariates, including details about the specific hydropower plant that it supplies, the plant's average annual output, the operating firm in question, and various pieces of geographic information. These data were also obtained from the NVE.[12]

Tracing the evolution of local market structure over time requires that each power plant and its associated reservoir(s) are mapped to the correct Elspot bidding areas at every point in time. The information needed to do perform this mapping does not readily exist. However, using the current Elspot allocation as a fixed starting point — together with coordinate information contained in my dataset, the Elspot change log document hosted by Nord Pool[13], a map of the Norwegian electricity grid components[14], 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 B and denoted by the capitalised letters A to Q for convenience. The most recent Elspot regime, which went into effect on 2 December 2013, has already been highlighted in Figure 1.

At this juncture, it is worth pointing out something that is evident from the the maps in Appendix B. Namely, bidding areas are not consistently defined under the different regimes, even if they have been assigned the same area ID. For example, bidding area NO1 under regime A differs substantially in size and network coverage to bidding area NO1 under regime Q. In contrast, bidding area NO3 under regime B is exactly the same as bidding area NO4 under regimes M through Q. To ensure consistency, I henceforth adopt the term *zone* to describe a unique (and potentially reoccurring) Elspot footprint. All told, there are 23 such zones. These will serve as key fixed effects units in the estimation procedure and are summarised in Table C.1 in the appendices.

5.2 Electricity prices, flows and transmission constraints

All electricity data are obtained from Nord Pool. Electricity flows from one Elspot area to others via defined corridors and is limited by the capacity constraints of the trans-

[12]For example: http://gis3.nve.no/link/?link=vannkraft
[13]http://nordpoolspot.com/globalassets/download-center/elspot/elspot-area-change-log.pdf
[14]http://gis3.nve.no/link/?link=nettanlegg

mission lines that make up those corridors. Similarly, a bidding area can have several corridors attached to it, depending on the number of neighbouring areas that it shares borders with. I use the phrase "in-use hours" as a shorthand for the total number of hours that electricity was flowing to or from a bidding area, across all of its corridors, in a single day. Consequently, it is possible for an area to experience more than 24 in-use hours during a day if it is connected via several corridors.

Electricity flow data for the individual Norwegian bidding areas are available at an hourly level, by corridor, from November 2000 onwards.[15] I match these flows to the maximum transmission capacities for the relevant corridors at each hour. I then define the transmission constraint on a corridor as binding whenever the hourly flows — whether importing or exporting — are equal to the maximum transmission 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 the transmission constraint is moreover binding. If the spot price between two regions is equal during this period, then I cannot infer much beyond the fact that the constraint is not binding. I cannot reliably say which area is exporting to which.

All told, binding transmission constraints are a common occurrence. Figure 2 presents a normalized measure of constraint severity: The fraction of daily in-use hours that are constrained for individual bidding areas, expressed as a cumulative distribution. To summarize, just over a quarter of daily exporting periods (and a third of daily importing periods) in the sample experience no binding constraint. However, the severity rate rises to 50 percent — one constrained hour for every two in-use hours — for at least half the sample. Moreover, a full quarter of daily exporting periods (and 15 percent of daily importing periods) are completely constrained. That is to say, the transmission capacities available to bidding areas in this category were always binding, so that electricity flows along all corridors were constrained during every single hour of the day that they were in use.

[15] Real-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

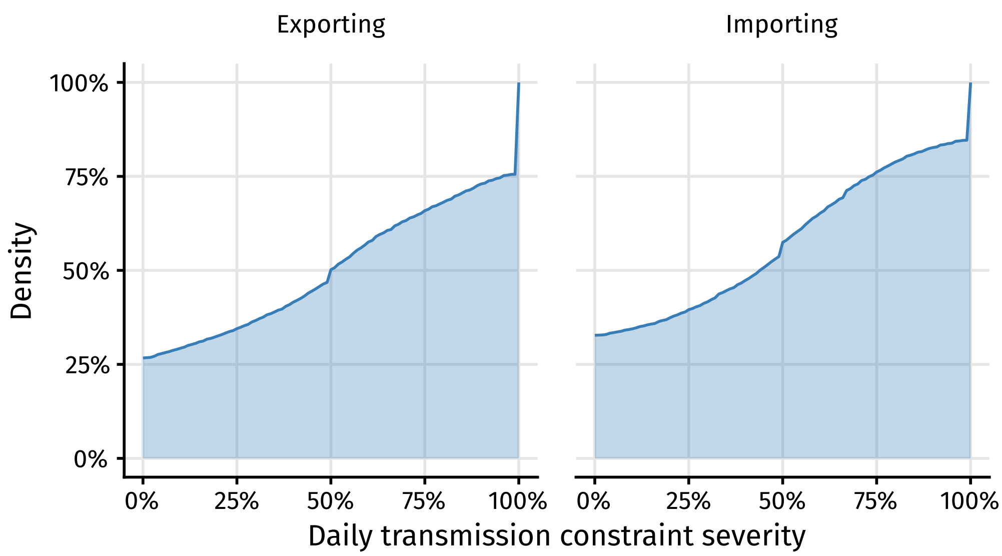

FIGURE 2: CUMULATIVE DISTRIBUTION OF TRANSMISSION CONSTRAINT SEVERITY

Notes: Severity reflects how often electricity flows to or from a bidding area were constrained by maximum transmission capacities.

5.3 Residual Supply Index (RSI)

Having established the institutional features underlying the reservoir and electricity data of this paper, we are now in a position to define a relevant measure of market power. Market power in a liberalised electricity sector is typically measured by the Residual Supply Index (RSI). Originally developed by the California Independent System Operator in the early 2000s, RSI is intended to overcome the shortcomings of traditional market power indicators — such as the Herfindahl Index and other concentration measures — as they apply to electricity markets.[16] It is formally defined as

$$RSI_f = \frac{\text{Total capacity} - \text{Firm } f\text{'s capacity}}{\text{Total Demand}}, \tag{8}$$

where total production capacity is the total regional supply capacity less net imports, and the other variables are self-explanatory. RSI thus captures the extent to which supply can meet demand, absent the supply of any one firm. Note that it is measured inversely to market power, such that a smaller RSI indicates greater market power. Similarly, an RSI less than one implies that a firm is *pivotal*. If $RSI_f < 1$, then the market

[16] See Sheffrin (2002), Twomey et al. (2005), and Newbery (2009).

cannot be satiated without f's production.

I calculate RSIs for every firm by estimating the daily output for each plant and then matching this to the observed electricity demand, supply, and net import values of the relevant market areas. While I do not observe (non-anonymous) daily bids of individual firms, I do know the average annual production of each plant. This allows me to obtain a daily equivalent by downscaling the annual numbers. Because this step involves some subjectivity, I construct four separate RSI measures that — taken together — should provide a robust insight into firm market power. The distinction between these four measures is a function of (i) how the relative competitive area is determined, and (ii) how average annual production is downscaled to the daily level.

1. **Constraint-weighted RSI:** Weights local RSI (i.e. in a firm's immediate bid area only) and system-wide RSI according to the relative frequency with which transmission constraints were binding for that day. For example, if the transmission corridors attached to a particular bidding area were constrained for 50 percent of the time, then a firm's constraint-weighted RSI is calculated as the simple mean of its local RSI and that which it commands within the Norwegian system as a whole. Similarly, if the electricity flows that day were entirely unconstrained, then a firm's weighted RSI is exactly equal to its system-wide RSI. As for downscaling, the daily GWh of each plant i is assumed to be the simple mean of its the annual production (i.e. $GWh_{i,d} = GWh_{i,y}/365$).

2. **Constraint-weighted (seasonally-adjusted) RSI:** Follows the same approach as the above, but adjusts the downscaling step to account for observed seasonal fluctuations in output. In other words, it suggests that daily production for all plants is higher during winter than it is in summer.

3. **Price-based RSI:** This measure treads a similar path to the first constraint-weighted approach above. However, rather than a weighted average of a firm's RSI in its local bidding area and the system as a whole, it is based on the actual formation of common price markets. These may comprise one or several bidding areas, or even the system as a whole when there are no transmission constraints. Such an approach has the advantage of accounting for constrained periods that may not be obvious when comparing electricity flows and capacities at a bidding area's immediate borders. For example, a bidding area may be part of a larger segmented

market arising from a bottleneck elsewhere in the system, even though electricity is freely flowing across its own corridors.

4. **Price-based (seasonally-adjusted) RSI:** As per the previous price-based measure, but GWh adjusted to reflect seasonal fluctuations.

The distribution for each of these four RSI measures, over all regimes, is plotted in Figure 3. We see that there is generally strong agreement between all of them. However, the seasonally-adjusted and price-based measures tend to be more conservative, with fewer observations crossing the pivotal ($\text{RSI} < 1$) threshold.

5.4 Elasticity of electricity demand and implied reservoir volumes

A great advantage of working with market-based electricity data in recent years is the availability of detailed bid curve histories. Recovering demand and supply elasticities directly from the data is then a simple matter of estimating the relevant arc elasticities (e.g. Wolak, 2003; Bigerna et al., 2016). Which is to say, we need merely trace deviations from a clearing price along the demand (supply) curve and calculate the resulting slope changes as per any introductory economics textbook. The arc elasticity approach not only benefits from being relatively easy to execute, but should also yield more precise estimates than alternative approaches.[17]

Nord Pool has made anonomized bid curve data available since mid-2014 and I have obtained the hourly bid sheets until mid-2017.[18] While this three-year period falls after my primary study period (i.e. 2000-2013), the fact that I am interested in persistent differences in seasonal elasticities allows me to proceed apace. I begin by defining an arbitrary arc segment of 5 EUR on either side of the hourly clearing prices.[19] I then calculate the corresponding arc elasticity of demand (ϵ_D) for each hour and aggregate these up into daily means. Finally, I regress these daily elasticities on month dummies

[17]To the best of my knowledge, existing estimates of the elasticity of electricity demand in Norway — and, indeed, the wider Nordics — rely on indirect approaches such as instrumental variables, model simulation, and parameterization (e.g. Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008). Likely, this is because the availability of detailed bid curve data is a relatively new advent.

[18]https://www.nordpoolgroup.com/elspot-price-curves. See Figure D.1 in the appendices for a detailed comparison of hourly bid curve data across two representative summer and winter dates.

[19]Experimenting with different segment lengths yields very similar answers to the ones presented here. As a reference, the average Nord Pool clearing price over 2014–2017 was 25 EUR.

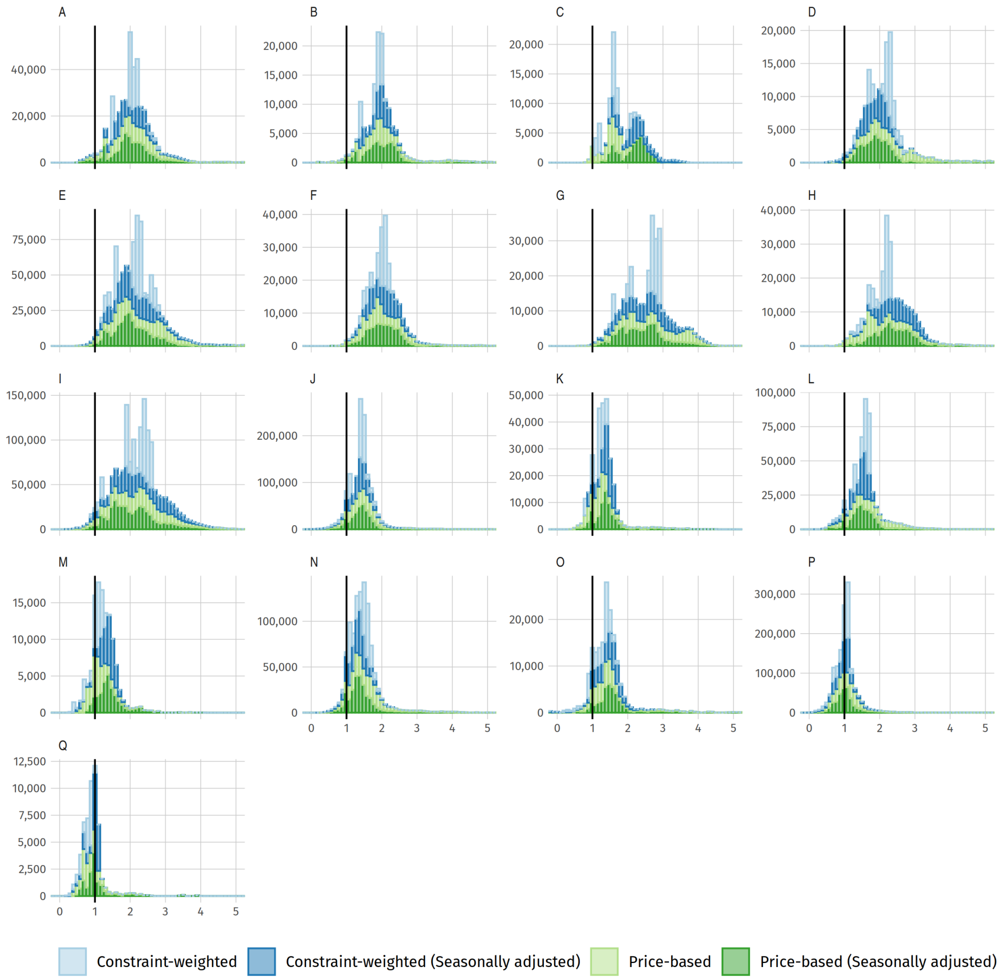

FIGURE 3: DISTRIBUTION OF RSI MEASURES

Notes: The black vertical line denotes an RSI value of 1. Observations to the left of this (i.e. < 1) are regarded as pivotal. The x-axis is truncated at an upper value of 5 and each facet represents a different Elspot regime.

and various other calendar effects. The key finding from this regression is summarized in Figure 4. The predicted demand elasticities exhibit a clear seasonal trend, with ϵ_D at its smallest absolute value in June and largest absolute value in November. These results correspond closely to the general finding of previous studies, which rely on alternate methods for estimating Norwegian electricity demand elasticities (e.g. Johnsen, 2001; Hansen, 2004; Bye and Hansen, 2008). We can therefore say with some confidence that electricity demand in Norway is at its most inelastic in the early summer months and

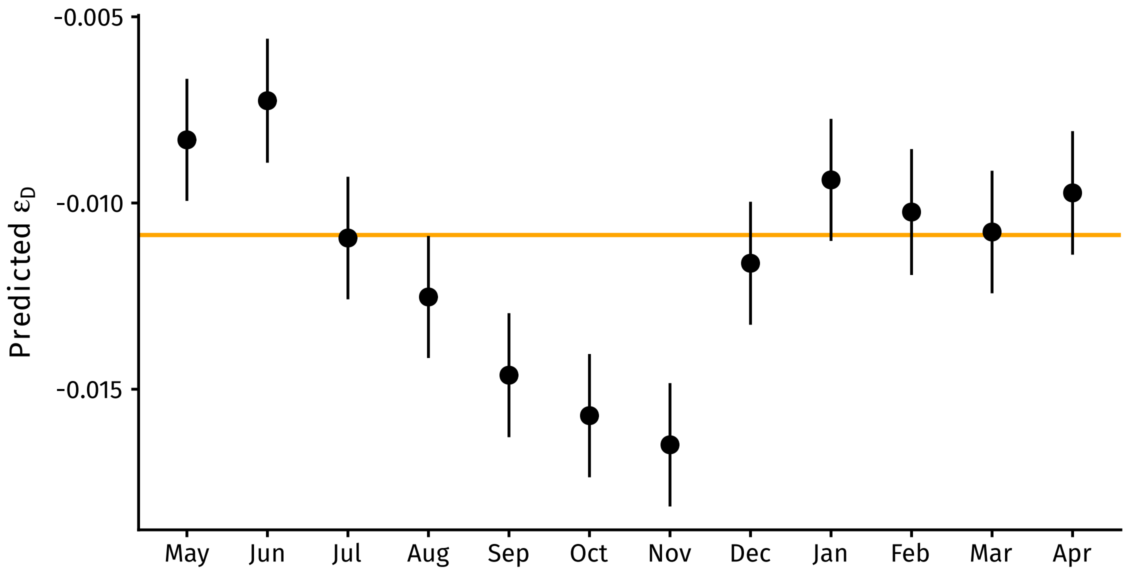

FIGURE 4: DEMAND ELASTICITY ESTIMATES

Notes: Estimates are obtained by regressing daily arc elasticity values, obtained from bid curve data over 2014–2017, on monthly dummies and other control variables. Dots denote point estimates and error bars denote 95 percent confidence intervals. The orange horizontal line highlights the median estimate and, following equation (6), serves as a reference case for periods when dominant firms would withold theoretically production (above the line) or oversupply (below the line) relative to the competitive outcome.

at its most elastic in the early winter months (or, late autumn months).

Importantly, these elasticity estimates provide the empirical link to a theory-motivated prediction of when (and how) variations in market power will cause deviations in reservoir management. Following equation (6), dominant hydropower firms will reallocate their water resources so as to withhold production during the relatively inelastic periods of the year. The opposite is true during the relatively elastic periods of the year. All other things being equal, the cumulative effect of this (proportional) production reallocation should lead dominant Norwegian firms to maintain relatively higher reservoir volumes in the summer months and relatively lower reservoir volumes in the winter months. Figure 5 maps out these implied deviations in production and reservoir volumes in stylised fashion, taking the theory of Section 4 and my empirical ϵ_D estimates here as given. The figure encapsulates the testable hypothesis that I shall subject my data to. If the theory about market power in a hydropower setting is correct, then my main regression analysis should yield coefficients that look roughly similar to the sinu-

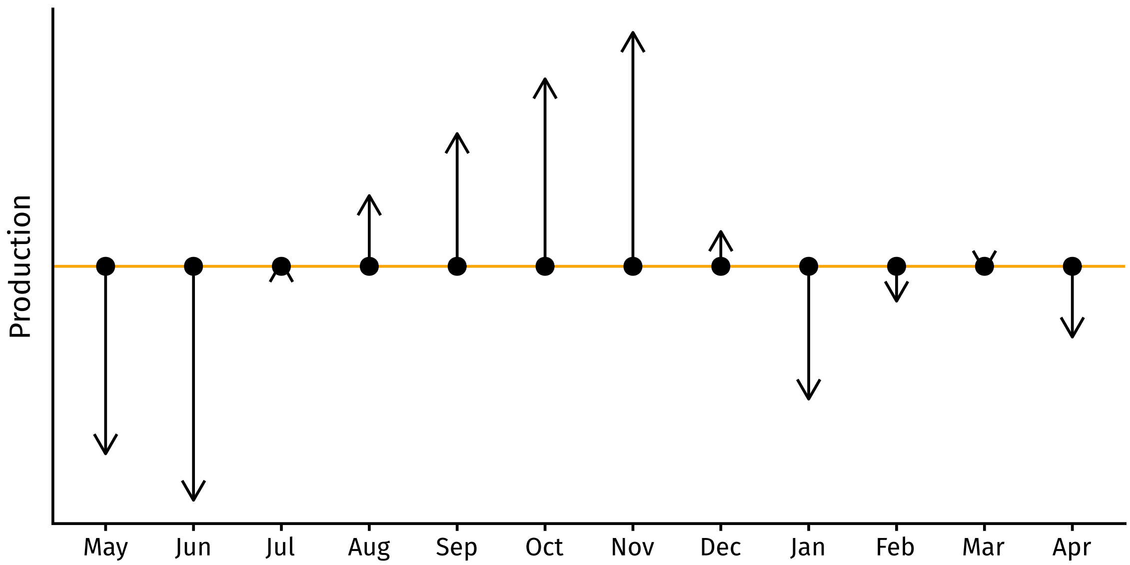

(A) Implied deviation in production

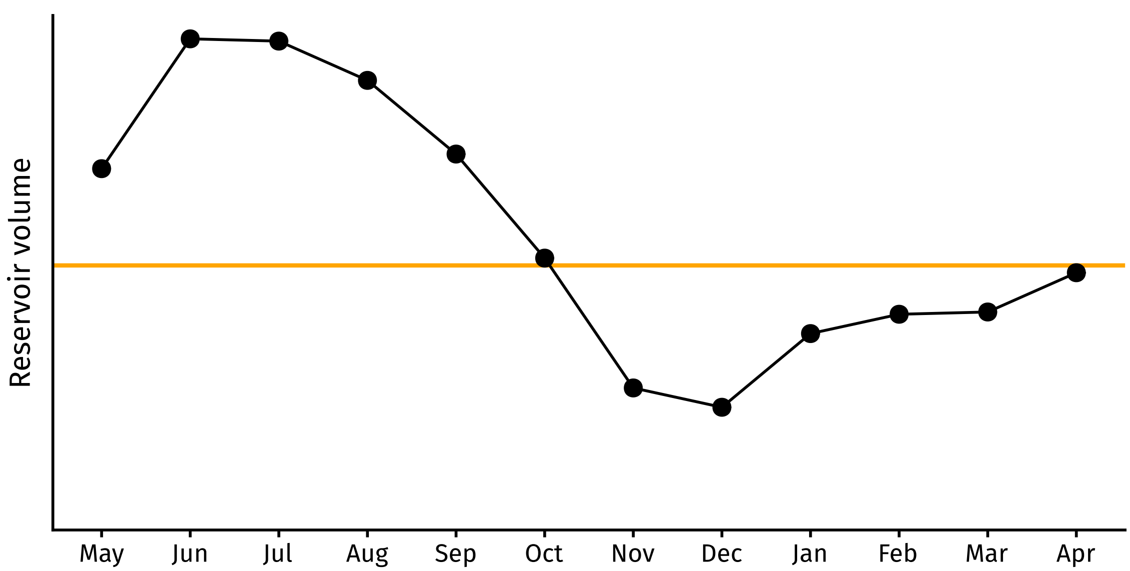

(B) Implied deviation in reservoir volume

Figure 5: Stylised implications of market power

Notes: Implied deviations are relative to a competitive outcome (orange lines) and based upon the empirical demand elasticity estimates shown in Figure 4. Panel (A) assumes that production follows a dominant hydropower firm's optimal allocation rule as per equation (6). Panel (B) takes this production allocation as given and maps it to reservoir volumes, assuming an approximate May-April hydro year.

soidal shape of Figure 5(B).

5.5 Weather data

Weather data are obtained from the National Climate Data Center (NCDC).[20] In order to derive a vector of daily weather data that is unique to each reservoir in my dataset, I follow a simple inverse-distance rule. Each reservoir is matched to its three nearest meteorological stations and weather conditions are assumed to be a weighted average of the conditions at these three stations. While this "nearest neighbours" approach is conceptually straightforward, there are two additional points worth noting. First, while snow depth and glacial melt are potentially valuable data for understanding Norwegian reservoir dynamics, these series are only available for a small subset (both temporally and spatially) of the meteorological stations that I have access to. I have thus excluded them from the analysis and focus predominantly on temperature and rainfall data. Second, the meteorological station dataset is unbalanced, with new stations coming into operation during the study period and others falling out of the NCDC records. Appendix A describes this issue in more depth, as well as the approach that I favour in dealing with the unbalanced station data. The upshot is that I determine the three nearest neighbours for each station on a per-year basis, according to which stations were operational during that year. This yields a dataset of 201 stations from which I derive the reservoir-specific weather data. A map showing the relative geographic distribution of these meteorological stations, as well as their years of availability relative to the full study period, is shown in Figure E.1.

6 Empirical strategy

6.1 Identification

A key assumption of my empirical framework is that changes to Norway's Elspot bidding areas over time provide a plausibly exogenous source of variation in local 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

[20] ftp://ftp.ncdc.noaa.gov/pub/data/noaa/

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 is in of itself testament to the underlying physical constraints, as Statnett (the system operator) seeks to best manage internal congestion problems, outages and scheduled maintenance periods, litigation procedures, and the laying out of new cables.[21]

The second component of my identification strategy — i.e. changes to the Elspot bidding areas only affect producer behaviour through changes in market share — is potentially complicated by the fact that the demand will also change with the redrawing of bidding areas. Yet, the demand complication is ultimately dealt with fairly easily through the inclusion of zone fixed effects.[22] This allows us to control for demand by grouping reservoirs at the zone level. Any residual differences between otherwise similar reservoirs can thus be interpreted as a causal result of market power, since producers within the same zone will always face the *same* demand, irrespective of whether transmission constraints are binding or not.

It is worth emphasizing that this latter claim would not necessarily 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. Depending on the supply characteristics of the redrawn bidding areas, the relative positions that two plants occupy along the merit curve would likely change if they did not have the same production technology. In other words, the probability of a producer serving a particular consumer — or, at least, the relative probability — would change and we would not be able to control for demand effectively. Fortunately, my exclusive focus on hydropower plants means that I am able to neatly sidestep this problem.[23] Not only am I

[21] The abuse of market power is not listed by Statnett as a reason for the redrawing of bidding areas. Instead, their communications to Nord Pool participants makes it clear that their decisions are based on exogenous factors like climate (e.g. drought) and the gradual completion of long-term grid connectivity objectives (e.g. subsea cables to new regions).

[22] Recall that a zone denotes a bidding area under a particular Elspot regime.

[23] At the very least, other studies that do not benefit from this unique institutional setting must control for varying plant types and bidding curves explicitly. For example, Puller (2007), Mansur (2007), Hortaçsu and Puller (2008), and Davis and Hausman (2016) rely on marginal cost estimates for different generation technologies to infer bidding behaviour.

always comparing like with like, but recall that the marginal costs of hydropower production are negligible. These plants occupy the lower rungs of the merit curve irrespective of other production technologies. Such factors further help to simplify the econometric analysis and again speak to the unique empirical advantages of hydropower in general and this research setting in particular.

6.2 Regression specification

Consider a fixed effects model that estimates water volumes in reservoir i $(i, ..., N)$, belonging to hydropower firm f $(f = 1, ..., F)$, at time t $(t = 1, ..., T)$. 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

$$V_{it} = \sum_{m=1}^{12} \beta_m M_{mt} + \sum_{m=1}^{12} \gamma_m M_{mt} \cdot \left(-RSI_{it,F(it)}\right) + \sum_{z=1}^{23} \eta_z Z_{it} + \boldsymbol{X\beta_X} + a_i + v_{it}, \tag{9}$$

where V is reservoir volume as a percentage of maximum regulated capacity, M_m is a set of month dummies, and RSI is a measure of the market power wielded by each firm. (The negative representation is used here to reflect an increase in market power.) The matrix $\boldsymbol{X}$ represents a set of additional controls such as (price-)zone dummies, as well as temperature and precipitation data. The model is completed by a composite error term, $a_i + v_{it}$, with the former component denoting 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 power terms. By interacting RSI with M_m, we are allowing for the fact that market power can have a differential effect on reservoir volumes, depending on how the price elasticity of demand varies by period.[24] Given the

[24]Following Balli and Sørensen (2013), I demean the reservoir-specific RSI terms prior to actually running any regressions on my computer. That is, the estimated model becomes $V_{it} = \sum_{m=1}^{12} \beta_m M_{mt} + \sum_{m=1}^{12} \gamma_m M_{mt} \cdot (RSI_{it,F(it)} - \overline{RSI}_i) + \boldsymbol{X\beta_X} + a_i + \epsilon_{it}$. This helps to safeguard against possibly spurious

observed variations in Norwegian demand elasticities (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 in the face of relatively more inelastic demand. Similarly, we expect these coefficients to then turn negative during the winter months as demand elasticities increase.

Finally, note that all of a firm's reservoirs in a realized (price-)zone share a common treatment; namely the same shock to local market power. I will therefore be double clustering my standard errors at the firm and (price-)zone level.

7 Results

My primary regression results are presented in Table 2. A visual equivalent in the form of faceted coefficient plots is presented in Figure 6. 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 (9) and show the marginal effect of increasing market power on reservoir volumes, contingent on month. Economic theory, together with my estimates of Norwegian electricity demand elasticities in Section 5.4, 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, these coefficients should be read as elasticities.

All four models exhibit a clear sinusoidal pattern. Increased market power generally leads to higher reservoir volumes — i.e. lower production — during the summer months and the opposite during the winter months. Taking Model (3) as a benchmark, a one unit decrease in RSI (i.e. increase on market power) corresponds to an approximate two percent increase in reservoir volumes during the peak summer months. Conversely, the same change in RSI yields a two percent decrease in volumes during the peak winter month of February. Again, this is all in accordance with our theoretical predictions from earlier; e.g. Figure 5b.

I run various alternate specifications to confirm the robustness of my results.[25] One

results due to unit-varying slopes. However, given that the underlying logic of the model is unchanged, I shall continue using the notation of equation (9) for simplicity.

[25] An earlier version of this paper used market share — the proportion of total production controlled

Table 2: Regression results (RSI)

Dependent variable: Reservoir volume (% of maximum capacity)

	Constraint-weighted		Price-based	
	(1)	(2)	(3)	(4)
Summer months				
$-$RSI × May	2.727 (4.020)	-4.754*** (0.583)	-0.526 (0.835)	-2.493** (1.167)
$-$RSI × Jun	8.295*** (2.836)	1.661** (0.744)	1.625** (0.781)	0.777 (0.807)
$-$RSI × Jul	9.241*** (3.222)	1.980* (1.185)	1.516*** (0.521)	0.524 (0.591)
$-$RSI × Aug	9.919** (3.872)	2.256*** (0.871)	1.790*** (0.589)	1.107* (0.591)
$-$RSI × Sep	8.633** (3.461)	1.454 (0.998)	1.979*** (0.474)	1.210*** (0.406)
$-$RSI × Oct	7.106*** (2.281)	-0.250 (1.300)	0.633 (0.420)	-0.183 (0.326)
Winter months				
$-$RSI × Nov	6.816*** (1.408)	-1.155 (2.300)	0.718 (0.636)	-0.532 (0.668)
$-$RSI × Dec	4.770*** (1.619)	-1.850 (1.386)	0.304 (0.948)	-0.643 (0.815)
$-$RSI × Jan	2.549 (2.470)	-1.236 (1.164)	-0.702 (0.787)	-0.686 (0.596)
$-$RSI × Feb	0.627 (3.500)	-2.730*** (0.834)	-1.928** (0.908)	-1.887** (0.898)
$-$RSI × Mar	0.616 (4.279)	-3.870*** (1.260)	-0.822 (0.764)	-1.616** (0.785)
$-$RSI × Apr	3.353 (4.884)	-3.175*** (0.955)	-1.039*** (0.395)	-1.415* (0.781)
Fixed-Effects:				
Reservoir	✓	✓	✓	✓
Month	✓	✓	✓	✓
Hydro year	✓	✓	✓	✓
Zone	✓	✓	×	×
Price Zone	×	×	✓	✓
Observations	1,755,247	1,755,247	1,790,699	1,790,699
R2	0.50818	0.5065	0.51098	0.51076
Within R2	0.32963	0.32735	0.26059	0.26026

Notes: The table shows the effect of increasing market power (decreasing RSI) on reservoir levels. In particular, the coefficients on the key $-$RSI × month interaction terms denote the marginal effect of a one percent decrease in producer RSI on reservoir levels, conditional on month of year. A positive coefficient indicates withholding of supply as a producer gains market power, and vice versa. In addition to various fixed effects, the models control for local weather conditions via a group of temperature and precipitation covariates. These are of the expected sign and jointly significant, but omitted for brevity. Models (1) and (2) weight local RSI and system-wide RSI according to the severity of transmission constraints, with the latter adjusting for observed seasonal variations in output. Models (3) and (4) measure RSI according to the formation of common price zones, with the latter again adjusting for observed variations in seasonal output. Standard errors in parentheses are two-way clustered at the producer & (price-)zone level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

variation of interest is a discrete measure of market power. Namely, how does a firm's "pivotal" status (RSI < 1) impact its production profile? The results from this set of regressions are presented in Table 4. As before, we see the same sinusoidal pattern

by a firm — instead of RSI, and the qualitiative results are very similar.

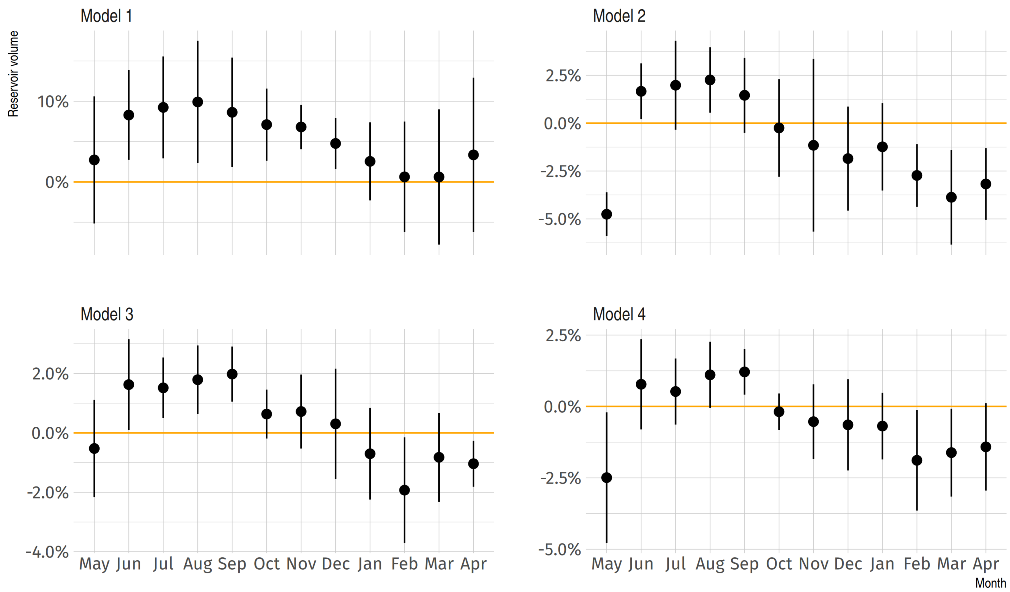

Figure 6: Change in reservoir volume due to increased market power

Notes: Marginal effect of a one percent decrease in producer RSI on reservoir volumes by month. Dots denote point estimates and the error bars show 95 percent confidence intervals. See Table 2 and the text for details.

reflective of strategic resource shifting. Firms that gain pivotal status clearly tend to keep higher reservoir volumes in summer and lower reservoir levels in winter.

While the remaining coefficients and controls have been omitted from the tables for brevity, these are all jointly significant and, where applicable, of the expected sign and magnitude. It should also be said that the comparative impact of market power 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.

Table 4: Regression results (pivotal producers)

Dependent variable: Reservoir volume (% of maximum capacity)				
	Constraint-weighted		Price-based	
	(5)	(6)	(7)	(8)
Summer months				
Pivotal × May	-5.881** (2.768)	-4.391** (2.204)	-0.996 (1.250)	-5.252*** (2.037)
Pivotal × Jun	1.169** (0.485)	2.608 (1.833)	1.854 (1.294)	0.007 (1.185)
Pivotal × Jul	2.666 (1.699)	4.499* (2.435)	2.683** (1.224)	1.764 (1.222)
Pivotal × Aug	3.128 (1.939)	5.570** (2.385)	3.635** (1.652)	2.750 (1.820)
Pivotal × Sep	2.787 (2.116)	4.883*** (1.884)	4.347** (1.788)	3.526** (1.738)
Pivotal × Oct	0.294 (1.847)	0.166 (1.466)	0.836 (1.485)	-0.781 (1.363)
Winter months				
Pivotal × Nov	-0.126 (1.473)	1.230 (1.633)	0.132 (1.297)	-1.419 (1.033)
Pivotal × Dec	-0.594 (0.568)	-0.158 (0.988)	0.189 (1.137)	-1.134 (0.939)
Pivotal × Jan	-0.356 (1.275)	0.744 (1.291)	-0.533 (0.773)	-1.786** (0.785)
Pivotal × Feb	0.041 (1.231)	0.998 (0.735)	0.471 (1.076)	-1.607 (1.670)
Pivotal × Mar	1.408 (1.796)	2.625*** (0.975)	2.479 (2.334)	0.264 (2.359)
Pivotal × Apr	1.338 (3.209)	3.091** (1.482)	2.909 (1.978)	1.118 (2.177)
Seasonal RSI	×	✓	×	✓
Weather controls	✓	✓	✓	✓
Fixed-Effects:				
Reservoir	✓	✓	✓	✓
Month	✓	✓	✓	✓
Hydro year	✓	✓	✓	✓
Zone	✓	✓	×	×
Price Zone	×	×	✓	✓
Observations	1,755,247	1,755,247	1,790,699	1,790,699
R2	0.50534	0.50587	0.51049	0.51066
Within R2	0.32576	0.32648	0.25984	0.2601

Notes: The table shows the effect of becoming a pivotal producer (RSI < 1) on reservoir levels. Models 5 through 8 thus represent discrete equivalents of Models 1 through 4, respectively, from Table 2. The same notes and general interpretations apply. Standard errors in parentheses are again two-way clustered at the producer & (price-)zone level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

8 Concluding remarks

How does market power affect firm behaviour? Microeconomic theory provides the answer familiar to every economist: Dominant firms will produce less than the social optimum (thereby raising prices and ensuring themselves higher profits). However, this fundamental tenet is surprisingly hard to verify with empirical data. A host of complications — from the unobserved factors underlying the decision to enter or exit a market, to uncertainties about marginal cost curves — prohibit us from causally identifying such behaviour among real-world firms. In the specific case of a hydro-based electricity system, theory tells us that dominant hydropower firms will 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.

I test this hypothesis using a novel data set of Norwegian hydropower reservoirs, power plants and electricity flows. Changes to bidding area divisions and binding transmission constraints provide the additional layers of exogenous variation that allow me to cleanly identify the causal impact of market power on firm behaviour. Consistent with the predictions of theory, I find that increased market power leads to a modest, yet definitive, intertemporal shifting of water resources. Taking Model (3) as a benchmark, a one percent decrease in producer RSI yields a 1–2 increase in reservoir volumes during summer and a similar percent decrease during winter. This market share effect is distinct from the regular seasonal patterns in reservoir volumes that arise from annual snow-melt inflows and so forth. The qualitative results are also robust to various specifications and model formulations.

Constructing a large dataset such as this from disparate sources 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 already been liberalised for nearly two decades by the start of my review period. The maturity of the market should give us at least some confidence regarding its stability and competitive structure. 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 competition authorities in effectively regulating electricity markets like the Nordic system, where hydropower comprises a major share of generation. Or, where transmission constraints create opportunities for producers to exercise local market power. To the best of my knowledge, this paper is one of the first studies to provide direct empirical evidence — absent structural assumptions or simulated counterfactuals — of noncompetitive behaviour in a dynamic market setting. While the magnitude of the effects that I identity here are modest, it is worth remembering that the Nordic system is generally regarded as a paragon among liberalised electricity markets. I therefore interpret my results as a lower bound on the extent to which market power is being exercised more generally in the economy. Any empirical advantages provided by the institutional features of this particular market should be tempered by these considerations. A larger lesson is that even our most advanced markets may be susceptible to the abuse of market power under relatively common conditions.

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.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen (2017). "Concentrating on the fall of the labor share," Working Paper 23108, National Bureau of Economic Research.

Azar, José A., Ioana Marinescu, and Marshall I. Steinbaum (2017). "Labor market concentration," Working Paper 24147, National Bureau of Economic Research.

Azar, José A., Ioana Marinescu, Marshall I. Steinbaum, and Bledi Taska (2018). "Concentration in US Labor Markets: Evidence From Online Vacancy Data," Working Paper 24395, National Bureau of Economic Research.

Balli, Hatice Ozer and Bent E. Sørensen (2013). "Interaction effects in econometrics," *Empirical Economics*, 45(1): 583–603.

Basu, Susanto (2019). "Are Price-Cost Markups Rising in the United States? A Discussion of the Evidence," *Journal of Economic Perspectives*, 33(3): 3–22.

Benmelech, Efraim, Nittai Bergman, and Hyunseob Kim (2018). "Strong employers and weak employees: How does employer concentration affect wages?" Working Paper 24307, National Bureau of Economic Research.

Berry, Steven, Martin Gaynor, and Fiona Scott Morton (2019). "Do increasing markups matter? lessons from empirical industrial organization," *Journal of Economic Perspectives*, 33(3): 44–68.

Bigerna, Simona, Carlo Andrea Bollino, and Paolo Polinori (2016). "Market Power and Transmission Congestion in the Italian Electricity Market," *The Energy Journal*, 37(2).

Borenstein, Severin (2016). "The Power and the Limits of Industrial Organization," *Review of Industrial Organization*, 48(3): 241–246.

Borenstein, Severin, James Bushnell, and Steven Stoft (2000). "The competitive effects of transmission capacity in a deregulated electricity industry," *RAND Journal of Economics*, 31(2): 294–325.

Borenstein, Severin, James B. Bushnell, and Frank A. Wolak (2002). "Measuring Market Inefficiencies in California's Restructured Wholesale Electricity Market," *American Economic Review*, 92(5): 1376–1405.

BP (2016). *BP Statistical Review of World Energy*, British Petroleum, 65$^{\text{th}}$ edition.

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.

Davis, Lucas and Catherine Hausman (2016). "Market Impacts of a Nuclear Power Plant Closure," *American Economic Journal: Applied Economics*, 8(2): 92–122.

De Loecker, Jan and Jan Eeckhout (2018). "Global market power," Working Paper 24768, National Bureau of Economic Research.

De Loecker, Jan, Jan Eeckhout, and Gabriel Unger (2019). "The Rise of Market Power and the Macroeconomic Implications," *Quarterly Journal of Economics*, Forthcoming.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri (2018). "Monopsony in online labor markets," Working Paper 24416, National Bureau of Economic Research.

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.

Genc, Talat S. (2016). "Measuring demand responses to wholesale electricity prices using market power indices," *Energy Economics*, 56: 247–260.

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.

Hortaçsu, Ali and Steven L. Puller (2008). "Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market," *The RAND Journal of Economics*, 39(1): 86–114.

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.

Mansur, Erin T. (2007). "Do Oligopolists Pollute Less? Evidence From A Restructured Electricity Market," *The Journal of Industrial Economics*, 55(4): 661–689.

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.

Müsgens, Felix (2006). "Quantifying Market Power in the German Wholesale Electricity Sector Using a Dynamic Multi-Regional Dispatch Model," *The Journal of Industrial Economics*, 54(4): 471–498.

Newbery, David (2009). "Predicting Market Power in Wholesale Electricity Markets," European University Institute Robert Schuman Centre for Advanced Studies. Working Paper RSCAS 2009/03. Available: https://cadmus.eui.eu/bitstream/handle/1814/10620/EUI_RSCAS_2009_03.pdf.

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.

Puller, Steven L. (2007). "Pricing and firm conduct in California's deregulated electricity market," *The Review of Economics and Statistics*, 89(1): 75–87.

Reguant, Mar (2014). "Complementary Bidding Mechanisms and Startup Costs in Electricity Markets," *The Review of Economic Studies*, 81(4): 1708–1742.

Reiss, Peter C. and Frank A. Wolak (2007). "Structural Econometric Modeling: Rationales and Examples from Industrial Organization," in James J. Heckman (ed.), "Handbook of Econometrics," volume 6, chapter 64, Elsevier, pp. 4277–4415.

Sheffrin, Anjali (2002). "Predicting Market Power Using the Residual Supply Index," in "FERC Market Monitoring Workshop," pp. 3–4.

SSB (2015). "Electricity, annual figures, 2013," *Statistics 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.

Sweeting, Andrew (2007). "Market Power In The England And Wales Wholesale Electricity Market 1995–2000," *The Economic Journal*, 117(520): 654–685.

Syverson, Chad (2019). "Macroeconomics and market power: Context, implications, and open questions," *Journal of Economic Perspectives*, 33(3): 23–43.

Twomey, Paul, Richard Green, Karsten Neuhoff, and David Newbery (2005). "A Review of the Monitoring of Market Power," Massachusetts Institute of Technology Center for Energy and Environmental Policy Research. Working Paper 05-002. Available: http://ideas.repec.org/p/cam/camdae/0504.html.

Wolak, Frank A. (2003). "Measuring Unilateral Market Power in Wholesale Electricity Markets: The California Market, 1998-2000," *American Economic Review*, 93(2): 425–430.

Wolfram, Catherine D. (1999). "Measuring Duopoly Power in the British Electricity Spot Market," *American Economic Review*, 89(4): 805–826.

Appendices

A Data cleaning and and merging

A.1 Hydropower reservoirs, plants and firms

As noted in the main text, reservoir data are normalised in terms of their 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 filtered 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.[26]

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

A.2 Weather data

The major issue with respect to the weather data used in this study is the fact that the NCDC panel dataset is unbalanced. While some Norwegian meteorological stations in the dataset are operational over the entire 2000–2013 study period, others come in

[26] While 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.

and out of existence. This presents a challenge insofar as I have to decide between two contrasting approaches for determining the (three) nearest neighbour stations to each reservoir, which are in turn used for calculating a reservoir's weather data vector via the inverse-distance weighted average rule. The first approach is to consider only meteorological stations that are available over the full 14-year study period. However, this excludes approximately 90 percent of the stations that would qualify as a "nearest neighbour" during the period when they were operational. I therefore favour a second approach, which is to proceed year-by-year and base the nearest neighbour rule on the stations that were operational for that particular year. This latter approach ensures that as much useful information is retained as possible and that the weather data for every reservoir in my dataset is as accurate as it can be. Experimentation shows that this choice has a negligible effect on the main findings.

B Previous Elspot regimes

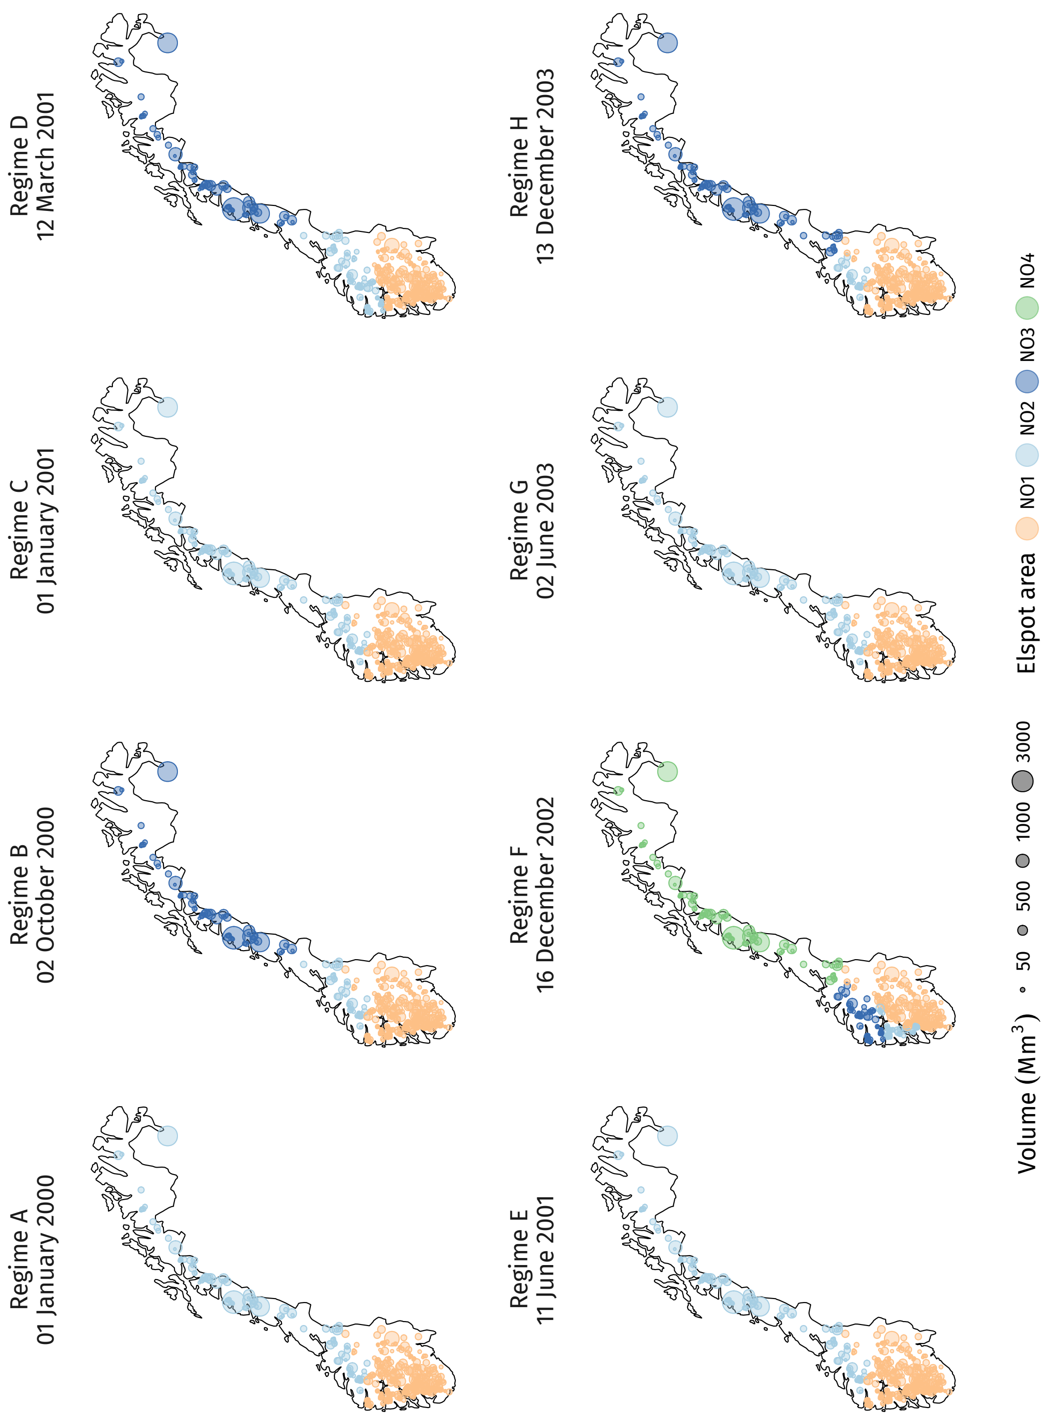

FIGURE B.1: ELSPOT REGIMES A – H (1 JANUARY 2000 – 13 DECEMBER 2003)

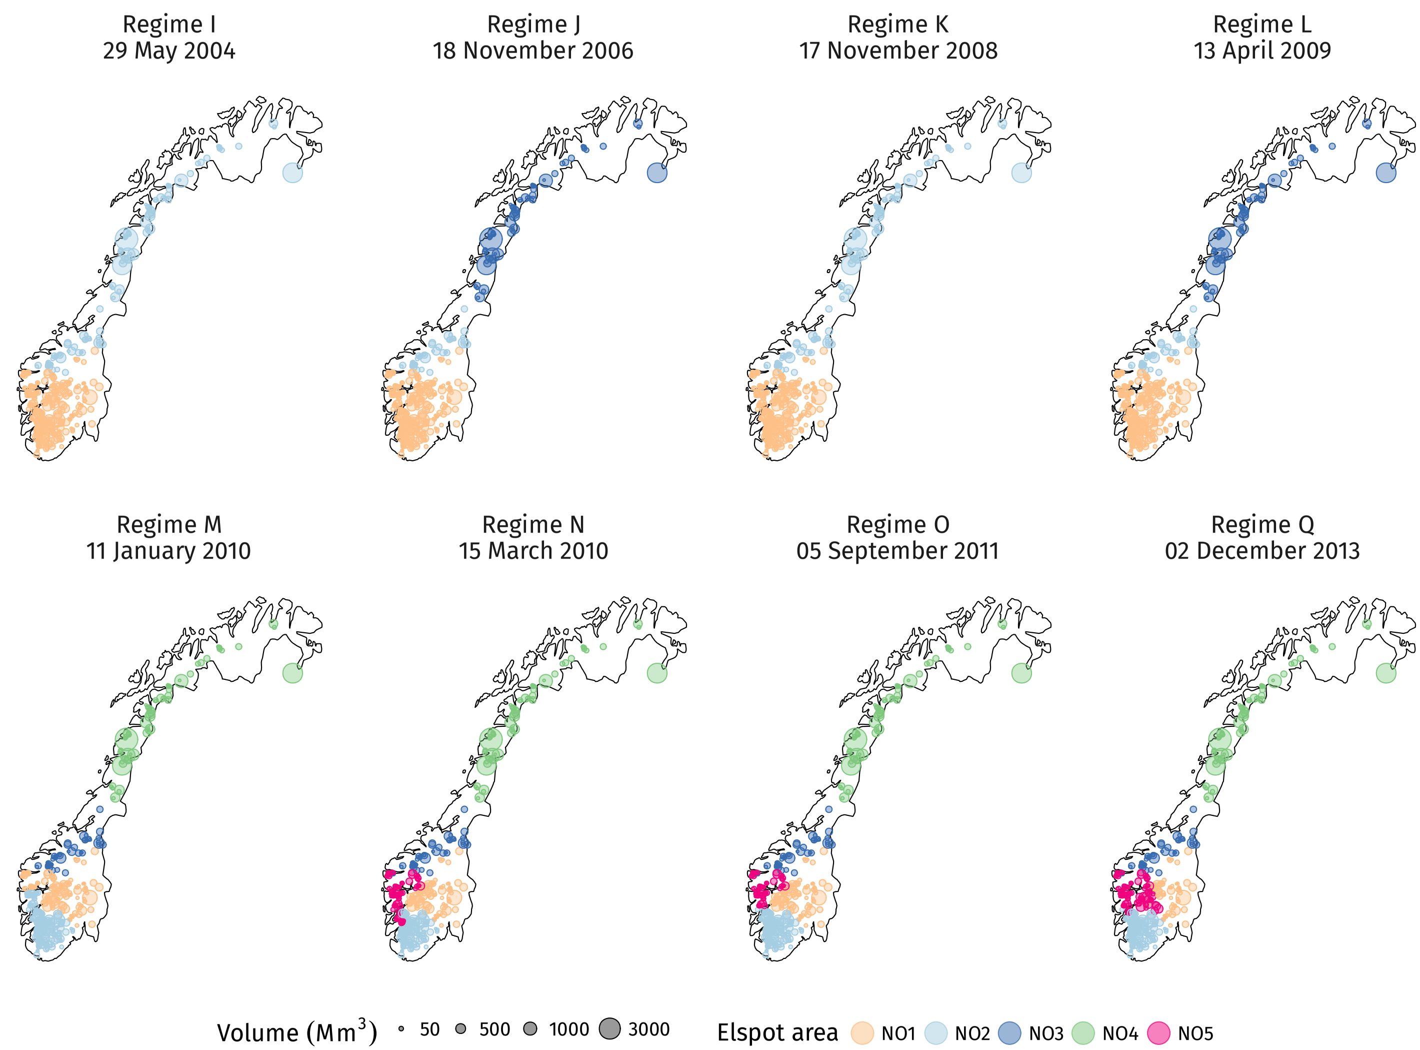

Figure B.2: Elspot regimes I – Q (29 May 2004 – 5 December 2011)

C Elpsot zones

TABLE C.1: ELSPOT ZONES

Zone	First instance	Repeat instances
1	A (NO1)	B (NO1); C (NO1); E (NO1); G (NO1); I (NO1); J (NO1); K (NO1); L (NO1)
2	A (NO2)	C (NO2); E (NO2); G (NO2); I (NO2); K (NO2)
3	J (NO2)	L (NO2); M (NO3)
4	B (NO3)	D (NO3); J (NO3); L (NO3); M (NO4); N (NO4); O (NO4); P (NO4); Q (NO4)
5	D (NO1)	–
6	D (NO2)	–
7	F (NO1)	–
8	F (NO3)	–
9	F (NO2)	H (NO3)
10	F (NO4)	–
11	H (NO1)	–
12	H (NO2)	–
13	M (NO1)	–
14	M (NO2)	–
15	N (NO1)	O (NO1); P (NO1)
16	N (NO3)	O (NO3); P (NO3); Q (NO3)
17	N (NO2)	–
18	N (NO5)	–
19	O (NO2)	P (NO2)
20	O (NO5)	P (NO5)
21	Q (NO1)	–
22	Q (NO5)	–
23	Q (NO2)	–

Notes: Zones describe particular geographic and network extents, which may be common to multiple Elspot bidding areas under different regimes A–Q. See main text for additional details.

D Bid curves

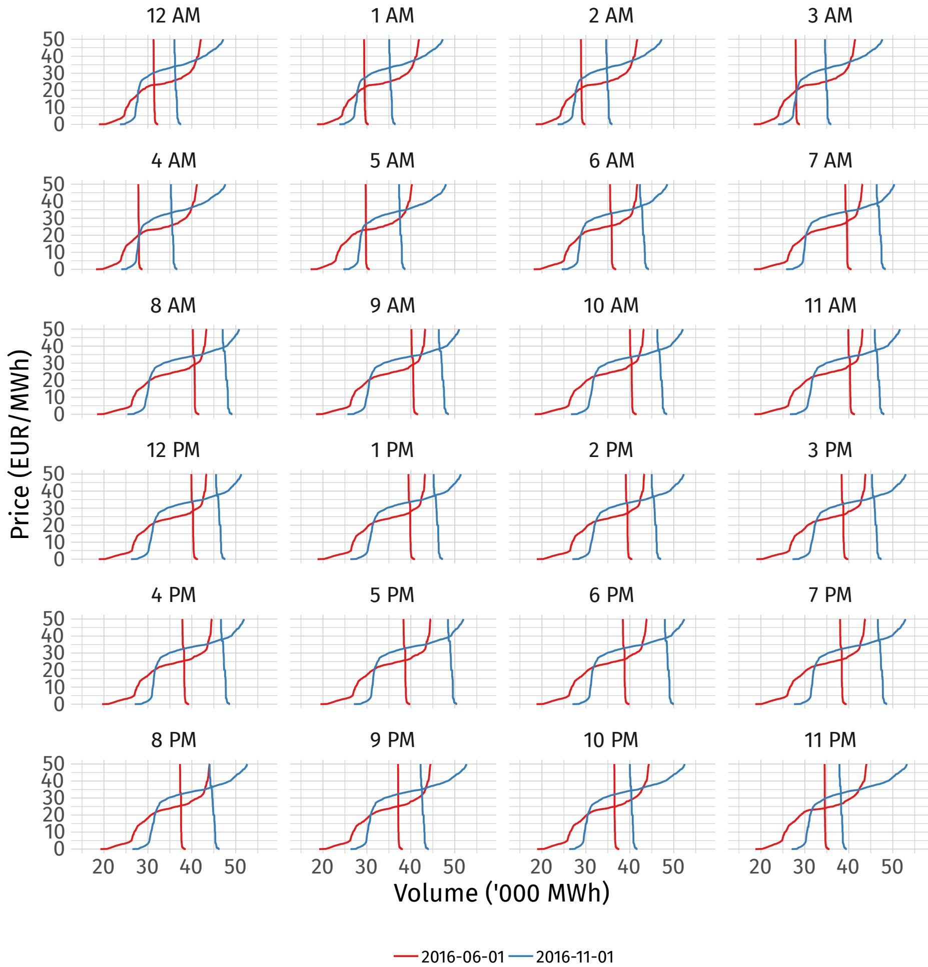

Figure D.1: Hourly Elspot bid curves

Notes: June 1st (red) and November 1st (blue) are selected as representative days from the months where electricity demand is respectively at its the most inelastic and elastic. The y-axes have been truncated at 0 and 50 EUR/MWh to aid visual inspection.

E Meteorological stations

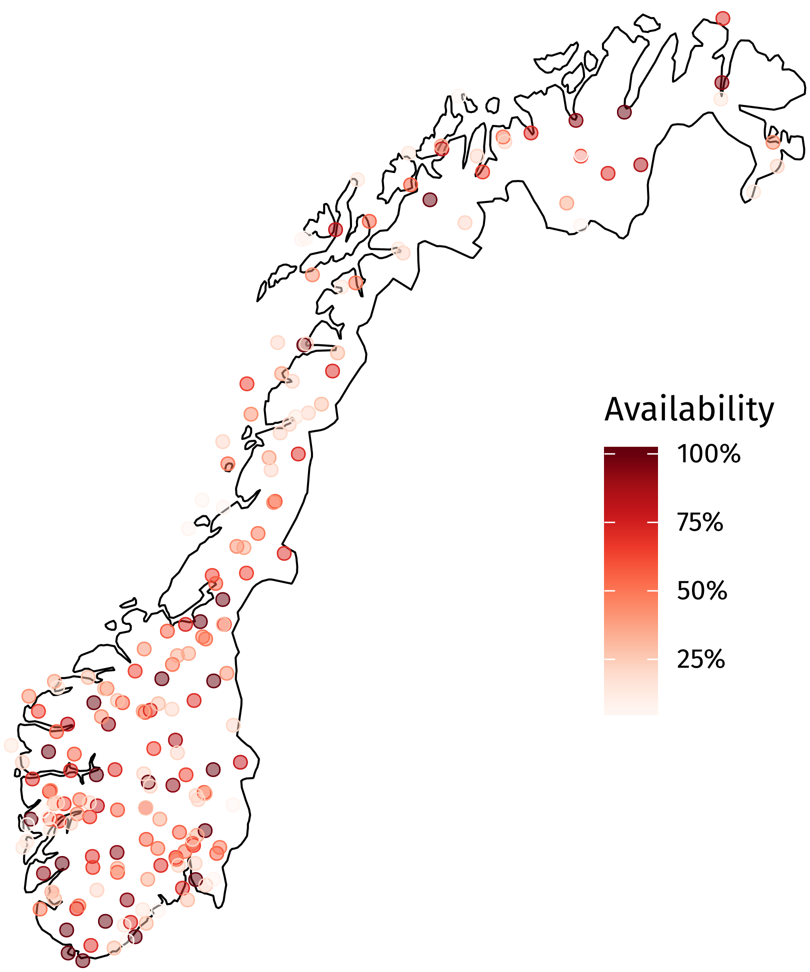

Figure E.1: Meteorological stations

Notes: Stations are colour-coded by years of availability relative to the full study period of 2000–2013. A station with 100% availability thus provides weather records for all 14 years, while a station with 50% availability only provides records for seven years.